Trend-following strategies for tail-risk hedging and alpha generation

Artur Sepp

https://artursepp.com, artur.sepp@gmail.com

23 April 2018

Introduction

Because of the adaptive nature of position sizing, trend-following strategies can generate the positive skewness of their returns, when infrequent large gains compensate overall for frequent small losses. Further, trend-followers can produce the positive convexity of their returns with respect to stock market indices, when large gains are realized during either very bearish or very bullish markets. The positive convexity along with the overall positive performance make trend-following strategies viable diversifiers and alpha generators for both long-only portfolios and alternatives investments.

I provide a practical analysis of how the skewness and convexity profiles of trend-followers depend on the trend smoothing parameter differentiating between slow-paced and fast-paced trend-followers. I show how the returns measurement frequency affects the realized convexity of the trend-followers. Finally, I discuss an interesting connection between trend-following and stock momentum strategies and illustrate the benefits of allocation to trend-followers within alternatives portfolio.

Contents

Key takeaway

1. The skewness and the convexity of strategy returns with respect to the benchmark are the key metrics to assess the risk-profile of quant strategies. Strategies with the significant positive skewness and convexity are expected to generate large gains during market stress periods and, as a result, convex strategies can serve as robust diversifiers. Using benchmark indices on major hedge fund strategies, I show the following.
   - While long volatility hedge funds produce the positive skewness, they do not produce the positive convexity.
   - Tail risk hedge funds can generate significant skewness and convexity, however at the expense of strongly negative overall performance.
   - Trend-following CTAs can produce significant positive convexity similar to the tail risk funds and yet trend-followers can produce positive overall performance delivering alpha over long horizons.

   See *Risk-profile of quant strategies*
2. Trend-following strategies adapt to changing market condition with the speed of changes proportional to the trend smoothing parameter for the signal generation. As result, when we measure the realized performance of a trend-following strategy, the return measurement frequency should be low relative to the expected rebalancing period of the trend-following strategy. Using the data of SG Trend-following CTAs index, I show that trend-followers are expected to produce both the positive skewness and convexity for monthly, quarterly and annual returns. As a result, trend-following strategies should not be seen as diversifiers for short-term risks measured on the scales less than one month. Overall, I recommend applying quarterly returns for the evaluation of the risk-profile of a trend-following strategy.
   See *The risk-profile of Trend-Following CTAs as function of return measurement frequency*
3. By analyzing quarterly returns on the SG trend-following CTAs index conditional on the quantiles of quarterly returns on the S&P 500 index, I show that trend-following CTAs can

serve as diversifiers of the tail risk. On one hand, the trend-followers generate significant positive average returns with positive skewness conditional on negative returns on the S&P 500 index. On the other hand, the trend-followers generate large positive returns, but with insignificant skewness conditional on large positive returns on the S&P 500 index. Conditional on index returns in the middle of the distribution during either range-bound or slow up-drifting markets, the trend-followers generate negative returns yet with significant positive skewness.

See *Trend-following CTAs as hedge against the tail risk*

4. The nature of trend-followers is to benefit from markets where prices and returns are auto-correlated, which implies the persistence of trends over longer time horizons. I present the evidence that the recent underperformance of trend-followers since 2011 to 2018 could be explained because the lag-1 autocorrelation of monthly and quarterly returns on the S&P 500 index become significantly negative in this sample period. The negative autocorrelation indicates the presence of the mean-reverting regime, even though the overall drift is positive, in which trend-followers are not expected to outperform. I introduce an alternative measure of the autocorrelation that can be applied to test for the presence of autocorrelation in short sample periods. I show that my autocorrelation measure has a strong explanatory power for returns on SG trend-following CTAs index.

   See *Autocorrelation as explanatory factor for trend-followers returns*
5. To quantify the relationship between the trend smoothing parameter, which defines fast-paced and slow-paced trend-followers, and the risk profile of fast-paced and slow-paced trend-followers, I create a quantitative model for a trend-following system parametrized by a parameter of the trend smoothing and by the frequency of portfolio rebalancing. The back-tested performance from my model has a significant correlation with both BTOP50 and SG trend-following CTAs indices from 2000s using the half-life of 4 months for the trend smoothing.

   See *Construction of Trend-following strategy for the S&P 500 index*
6. Using the trend system parametrized by the half-life of the trend smoothing, I analyze at which frequency of returns measurement the trend-following strategy can generate the positive convexity. The key finding is that the trend-following system can generate the positive convexity when the return measurement period exceeds the half-life of the trend smoothing and the period of portfolio rebalancing. I recommend the following.
   - If a trend-following strategy is sought as a tail risk hedge with a short-time horizon of about a quarter, allocators should seek for trend-followers with relatively fast smoothing of signals with the average half-life less than a quarter.
   - If a trend-following strategy is sought as an alpha strategy with a longer-time horizon, allocators should seek for trend-followers with medium to low smoothing of signals with the average half-life between a quarter and a year.

   An alternative way to interpret the speed of the trend smoothing is to analyze the trend-following strategy beta to the underlying asset. For the slow-moving smoothing, the strategy maintains the long exposure to the up-trending asset with infrequent rebalancing. As a result, the higher is the half-life of the trend smoothing, the higher is the beta exposure to the index. Thus, while fast-paced trend-followers can provide better protection during sharp short-lived reversals, they suffer in periods of choppy markets. There is an interesting article on Bloomberg that some of fast-paced trend-following CTAs fared much better than slower-paced CTAs during the reversal in February 2018.

See *Risk profile of S&P 500 Trend-following strategy*

7. I examine the dependence between returns on the trend-following CTAs and on the market-neutral stock momentum. I show that the trend-followers have a stronger exposure to the autocorrelation factor and a smaller exposure of higher-order eigen portfolios. As a result, the trend-following CTAs produce the positive convexity while stock momentum strategies generate the negative convexity of their returns.
   See *Trend-following vs Stock Momentum*
8. The allocation to trend-following CTAs within a portfolio of alternatives can significantly improve the risk-profile of the portfolio. In the example using HFR Risk-parity funds and SG trend-following CTAs index, the 50/50 portfolio equally allocated to Risk-parity funds and trend-following CTAs produces the drawdown twice smaller than the portfolio fully allocated to Risk-parity funds. The 50% reduction in the tail risk is possible because the occurrence of the drawdowns of Risk-parity HFs and Trend-following CTAs are independent. While trend-followers tend to have lower Sharpe ratios than Risk-parity funds, trend-followers serve as robust diversifiers with 50/50 portfolio producing the same Sharpe ratio but with twice smaller drawdown risk.
   See *Benefits of Trend-following CTAs for Allocations in Alternatives*

Risk-profile of quant strategies

All quantitative strategies have specific risk profiles characterized by the skewness of returns, performance in tail events, and cyclicality risk, when strategies perform only in certain market cycles. I have discussed the cyclicality risk of quantitative strategies in detail here.

I consider the following key risk metrics to evaluate the risk profile and the performance attribution of quantitative investment strategies.

1. **Realized historical volatility of strategy returns**. The volatility serves as a normalizing factor to compare different strategies on the same scale. I emphasize that, while the volatility serves only as a static measure of the second-order risk, on one hand, the volatility targeting can be applied to align return profiles of different strategies and, on the other hand, the volatility targeting along with dynamic forecasting of volatility can serve as risk-control. It is typical for trend-following programs to apply the volatility targeting both as a scaler and risk-control of portfolio positions.
2. **Skewness of realized returns**. The skewness serves as a static measure of the third order risk-profile of investment strategies. Strategies with the negative skewness typically include carry and mean-reversion strategies, in which frequent small gains are followed by infrequent large losses. Strategies with positive skewness tend to include trend-following and long volatility strategies, in which case strategies tend to produce infrequent large gains followed by a series of frequent small losses.
3. **Convexity of realized returns with respect to the flagship index or benchmark.** I define the convexity as the beta coefficient of strategy returns to the square of returns on the benchmark. In this way, the convexity measures the dynamic risk of strategy performance in tails of the performance of the flagship index. On one hand, the strategies with negative convexity underperform considerably the index in stressed periods with large negative performance on the index and yet they tend to underperform the benchmark when it

produces large positive gains. On the other hand, the strategies with positive convexity tend to outperform the benchmark in both negative and positive periods with strong emphasis on delivering positive performance in stressed periods.

The quantitative way to analyze the convexity profile of a quantitative investment strategy is to estimate the quadratic (parabolic) regression of returns on the strategy predicted by returns on the flagship or benchmark index and index returns squared:

$$
\begin{aligned}
StrategyReturns = Alpha& \\
&+LinearBeta * IndexReturns \\
&+ConvexityBeta * IndexReturns\hat{}2
\end{aligned}
$$

The estimate of the linear beta indicates the direct first-order exposure to the performance of the benchmark index. Market-neutral strategies have insignificant linear betas.

The estimate of the convexity beta assesses the second-order exposure indicating how the strategy performs in markets with strong bias to either downside or to upside. The convex strategies benefit from extreme returns on the benchmark index while the concave strategies suffer in extreme markets.

The alpha is the estimate of the excess return that the strategy can generate. I caution that typically the explanatory power of this regression may not be strong because of limited sample size and noise. However, the estimate of the convexity beta helps to understand the behavior of strategies in tail events.

The skewness and convexity profiles of major quant strategies are illustrated below in Figures 1 and 2, respectively. One of drawbacks of the skewness and convexity statistics is that both measures depend on the measurement frequency of realized returns. It is typical that the skewness and convexity profiles are different for daily and monthly returns on quantitative strategies for the trend-following strategies. I will investigate this topic further.

Illustrations using hedge fund indices

To illustrate the risk profile of different quantitative strategies, I apply hedge fund indices that show the aggregated performance of niche quant hedge funds (HF):

1. **Short Vol HF index** is CBOE Eurekahedge Short Volatility Index measuring the equal-weighted performance of 13 hedge funds which have short exposures to the implied volatility.
2. **Relative Value Vol HF index** is Relative Value Volatility Hedge Fund Index of 35 hedge funds that run volatility strategies with long, short, or neutral exposure using the relative value approach.
3. **Long Vol HF index** is Long Volatility Index of 11 hedge funds that have net long exposures to the implied volatility.
4. **Tail Risk HF index** is Tail Risk Hedge Fund Index including 8 hedge funds that seek to generate significant upside during market stress periods.
5. **SG Trend-following CTAs** is SG Trend Index (NEIXCTAT Index) consisting of 10 systematic trend-following commodity trading advisors (CTAs).

Both the CBOE HF and SG CTA indices are equal-weighted and reconstituted annually accounting for the survivorship bias. CBOE HF indices are updated on monthly basis while the SG trend-following CTAs index is updated daily. The time series of monthly returns for the first three Eurekahedge HF indices are available from January 2005, the data for Tail Risk HF index and SG Trend-following CTAs are available from January 2008 and January 2000, respectively.

In the figure below, I show the realized skewness on the HF strategies. The short vol strategies are the most negatively skewed, which is a typical feature of this strategies. The positive performance over the stressed periods on short volatility strategies is seen as a compensation to bear the skewness risk of these strategies. I presented some ideas for filtering and hedging of short vol strategies to improve their risk-adjusted performance.

The relative value vol funds have insignificant skewness by combining short and long exposures to implied volatility, while the long volatility funds can generate positive skewness by maintaining the net long exposure to the implied volatility. The tail risk hedge funds have a strong positive skewness which come however at the expense of extended periods of losses during normal markets.

The trend-following CTAs do not appear to generate significant skewness for monthly returns, however the skewness of quarterly returns becomes significant. I also point out because we have a limited number of monthly returns – about 200 observations – the sample volatility of the skewness estimate is relatively high of about 0.17.

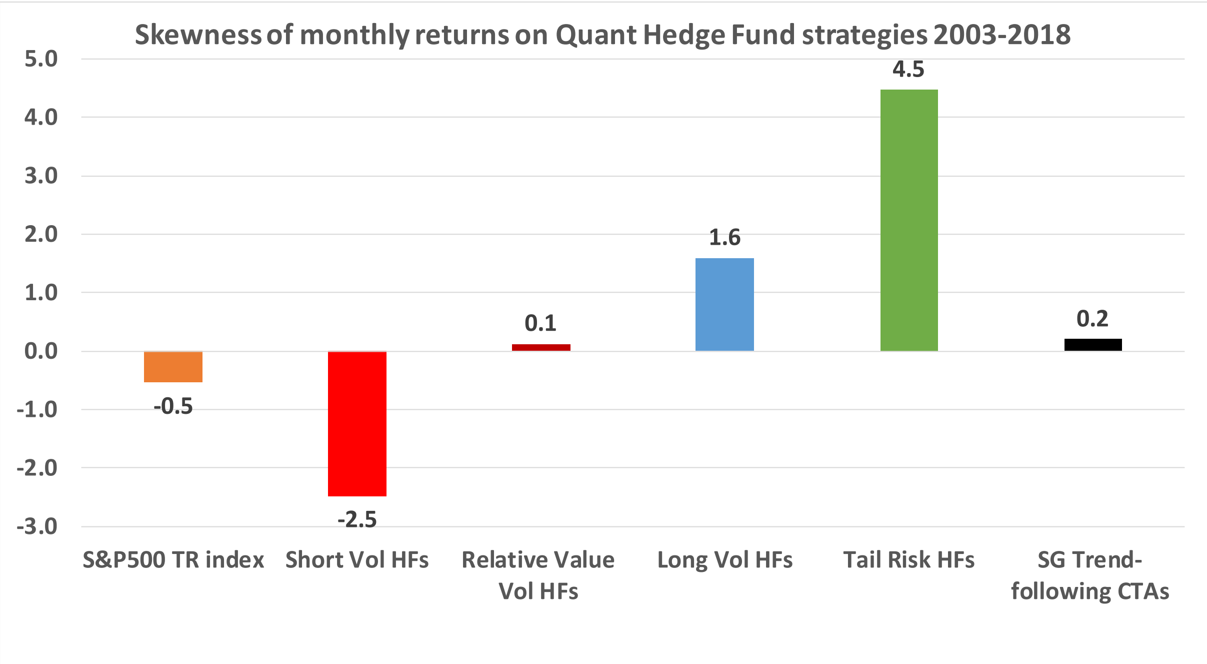

In the figure below, I show the realized convexity with respect to the S&P 500 index using monthly returns. The overall explanatory power of the regression is not strong, only about 10%, However, it provides the indication on the tail risk of quant strategies.

Short volatility HFs generate strong negative convexity losing during stress periods. Relative value and long vol HFs have an insignificant convexity profile. It is instructive that for long Vol HFs, the positive skewness does not produce positive convexity: while the long vol HFs are expected to produce infrequent large gains, the arrival of these gains is not expected to occur in the tail. In contrast tail risk HFs do produce both positive skewness and convexity, however at a cost of negative overall performance which can be seen from their negative alpha. Trend-following CTAs produce significant positive convexity even though they have insignificant positive skewness.

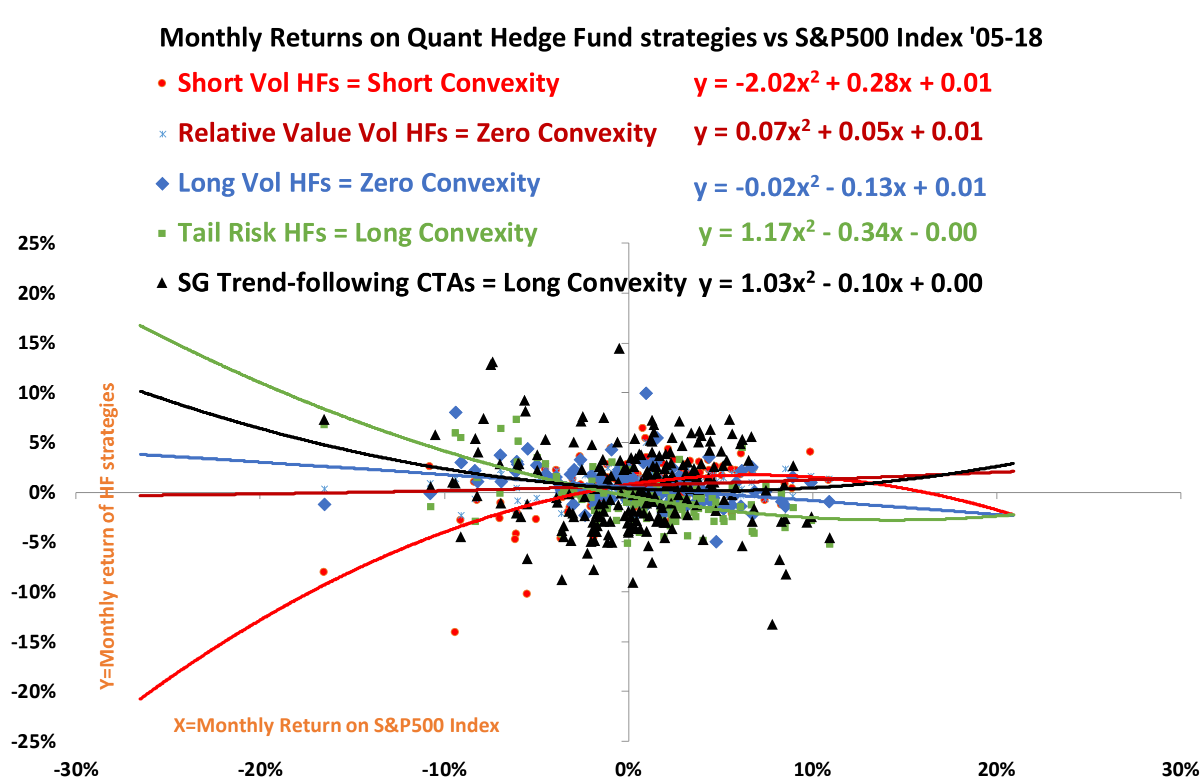

The risk-profile of Trend-Following CTAs as function of return measurement frequency

The return measurement frequency is defined by the non-overlapping periods for computing realized returns on the strategy. Common examples include daily, weekly, monthly, quarterly, and annual periods. The return measurement frequency has a strong impact on the realized risk-profile of a trend-following strategy. The reason is that the trend-following strategy need to adapt to changing market condition with the speed of changes proportional to the length of observation window for the signal generation. Fast-paced CTAs apply narrow window of recent return to determine the direction of the trend while slow-paced CTAs use longer windows. As a result, the sign of the exposure and the realized performance can be different for both slow- and fast-paced CTAs, in particular, when the measurement frequency is slow.

I start with the analysis of the skewness and convexity of realized returns on trend-following strategies using SG Trend-following CTAs index with time series data from January 2000 to March 2018. I will analyze the impact of window for signal generation later. I apply the S&P 500 index as the benchmark.

Realized Skewness

In the figure below, I show the skewness measured at different frequencies. The skewness of returns on the S&P 500 index is almost zero for daily returns while it becomes negative with magnitude of about -0.5 for returns sampled at lower frequencies. At the same time the skewness of Trend-following CTAs is negative for daily and weekly returns, slightly positive for monthly returns and strengthens further for quarterly and annual returns.

The indication is that the trend-following CTAs are more likely to have bigger gains than losses when their returns are sampled at slower frequencies. As a result, because of the dynamic nature of CTAs, the positive skewness manifests at slower frequencies.

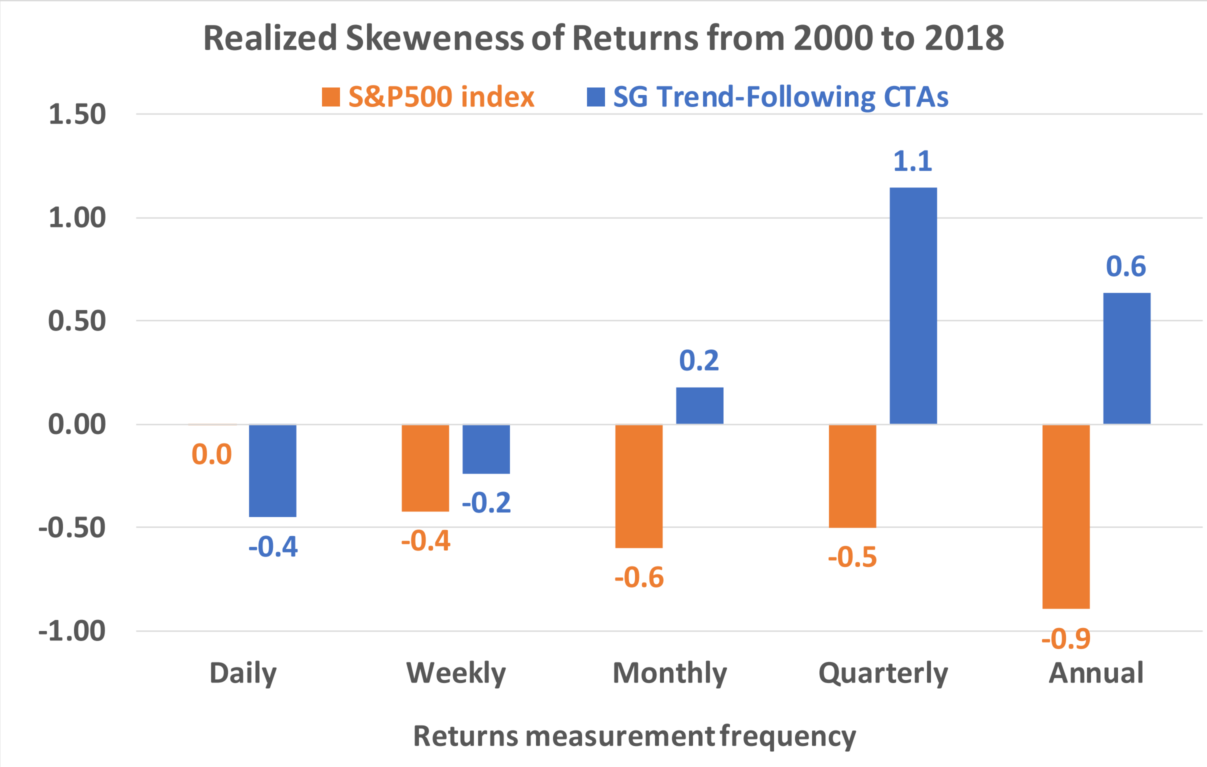

Realized Convexity

1. **For daily and weekly frequencies of return measurement**, we observe a negative convexity on trend-following CTAs. We also observe a similar pattern of negative convexity for weekly frequency of returns measurement. Trend-following CTAs should not be sought for short-term protection.

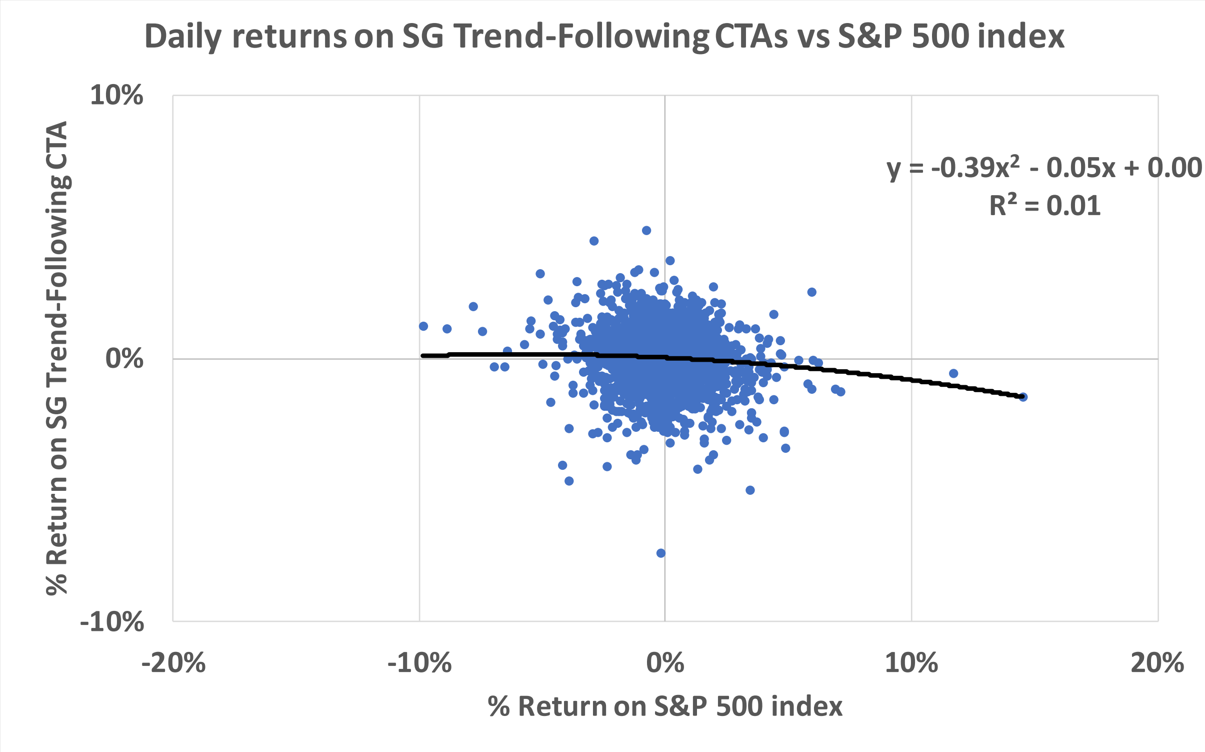

2. **For monthly and quarterly frequencies of return measurement**, we observe that the convexity of trend-following CTAs becomes positive and strong. Similarly, for quarterly frequency, we observe that the magnitude of the trend-following CTAs becomes smaller than for monthly returns, but the explanatory power of the regression increases.

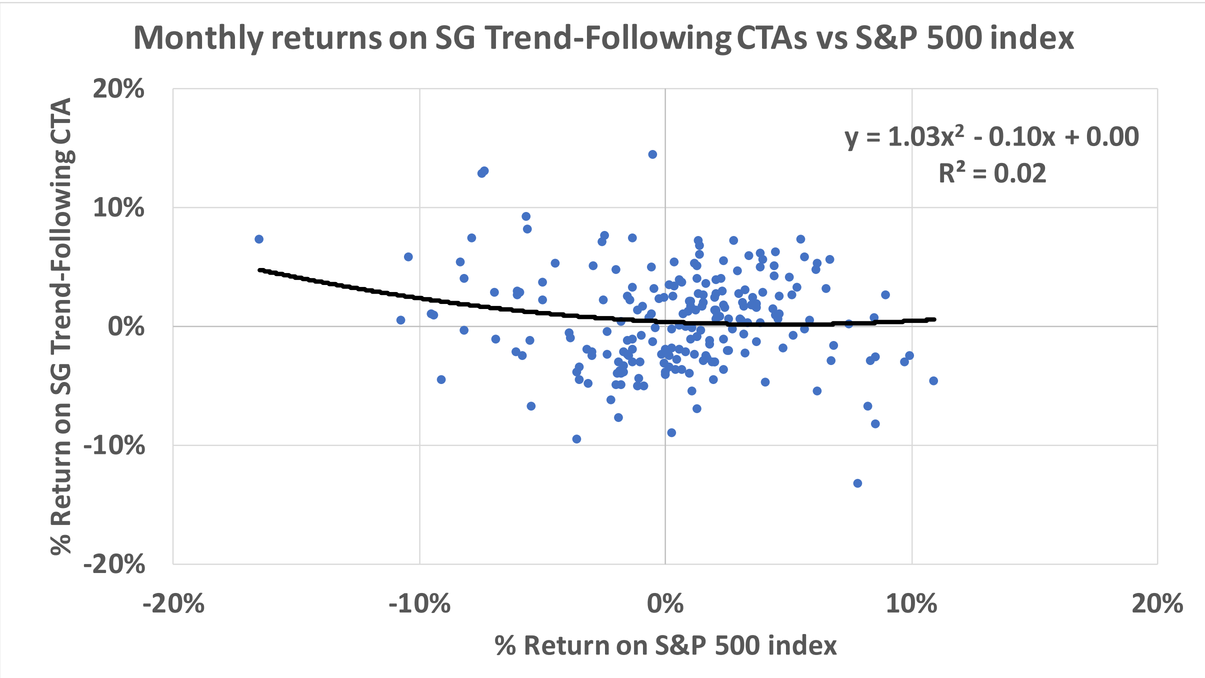

3. **For annual frequency of return measurement**, we observe that the magnitude of the TF CTAs increases, and the explanatory power of the regression strengthens.

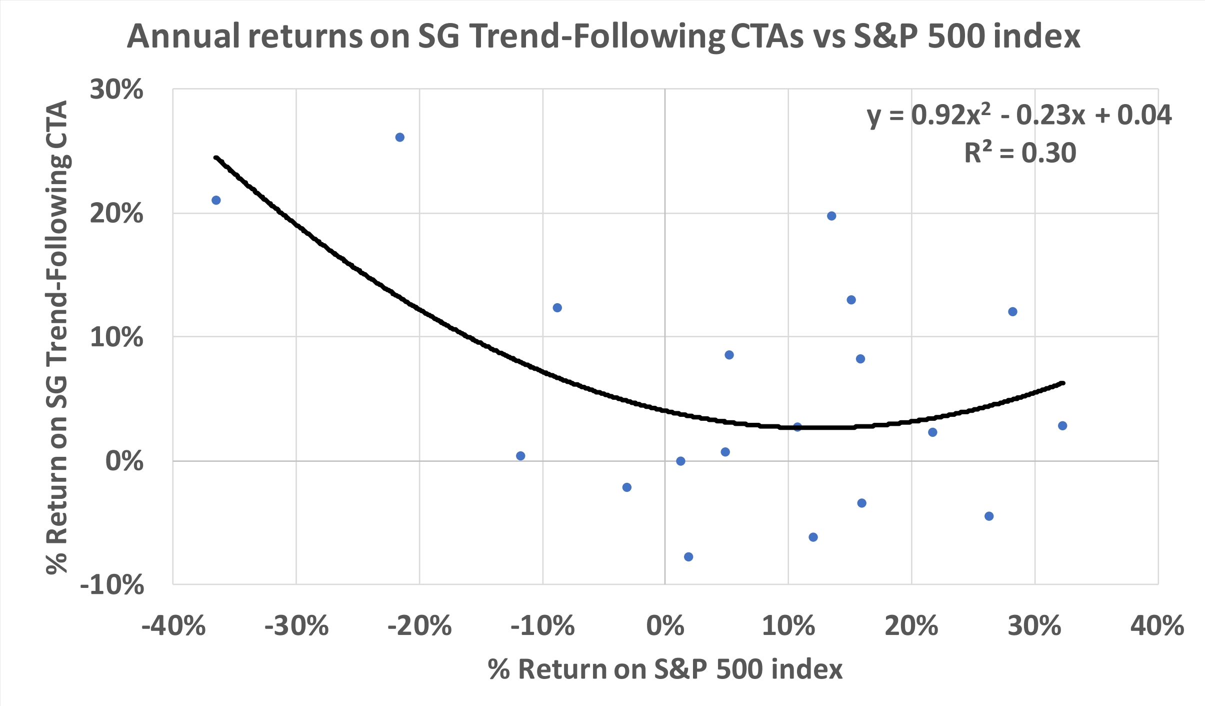

Summary

In the figure below, I illustrate the linear beta and the convexity beta for the trend-following CTAs as functions of return measurement frequency.

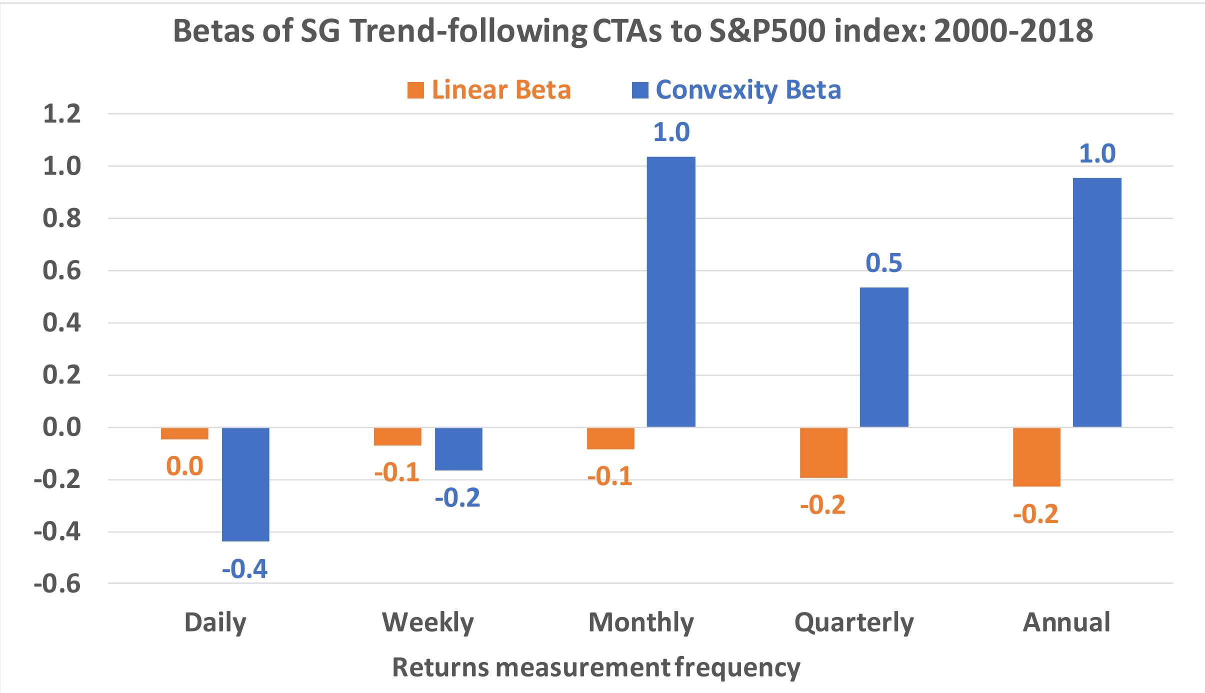

Trend-following CTAs as hedge against the tail risk

Now I analyze conditional quarterly returns on trend-followers using SG Trend-following CTAs index from January 2000 to March 2018. To compute the conditional return, I divide quarterly returns on the S&P 500 total return index into the four buckets corresponding to four quantiles measured with 25% intervals:

1. 0-25%Q includes quarterly returns between -35% to -1% and represents the stressed market.
2. 25-50%Q includes quarterly returns between -1% to 2% and represents the range-bound markets.
3. 50-75%Q includes quarterly returns between 2% to 6% and represents the normal up-trending market.
4. 75-100%Q includes quarterly returns between 6% to 35% and represents a strong bullish market.

In the figure below, I plot the conditional quarterly returns. We see that the trend-following CTAs deliver strong conditional performance when markets are either in the stress or in the bullish regime. The worst period of performance is the range-bound market. Trend-followers also produce negative performance in the period with normal up-trending market.

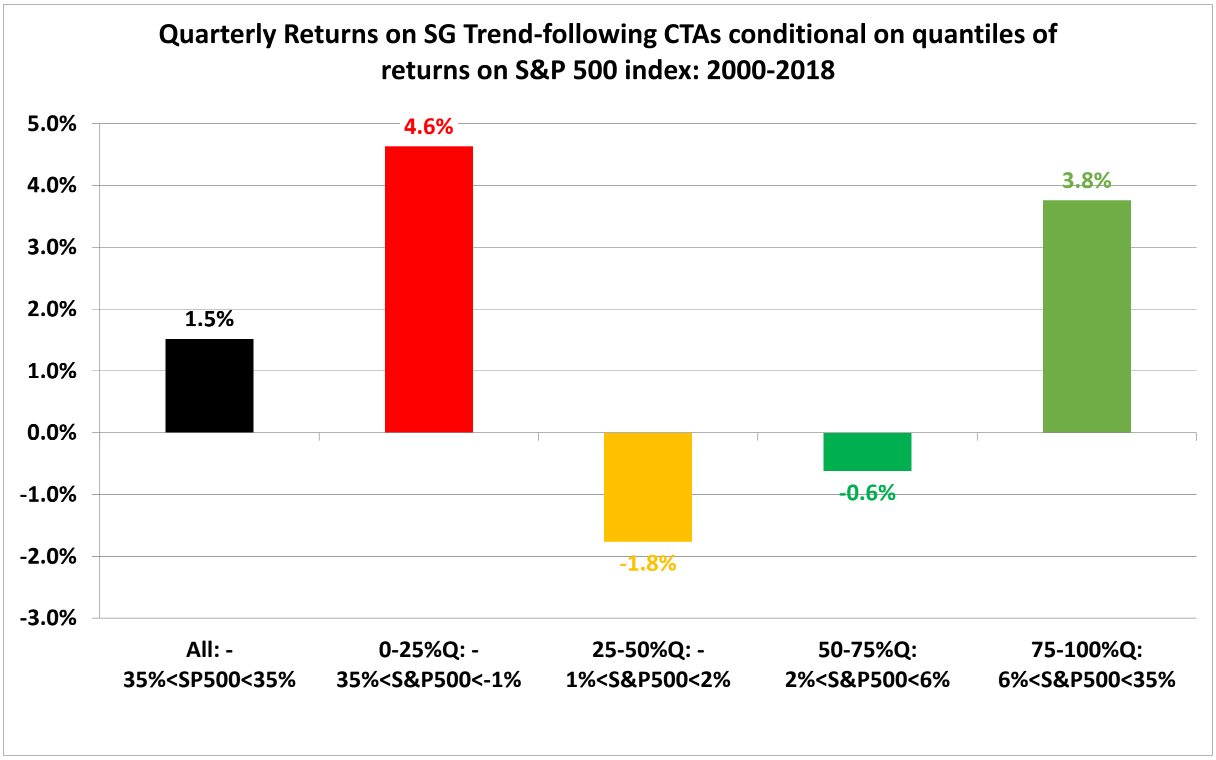

In the next figure I analyze the conditional performance of the trend-following CTAs by the risk reward frontier with the reward measured by the average of conditional quarterly returns and risk measured by the skewness of conditional quarterly returns. We see that overall the skewness on trend-following CTAs does not depend on market regimes for all but the bullish period. In the bullish period, while the trend-following CTAs still deliver the performance, the skewness of returns in this period is almost zero. As a result, the trend-following CTAs can experience both large gains and large losses in bullish markets.

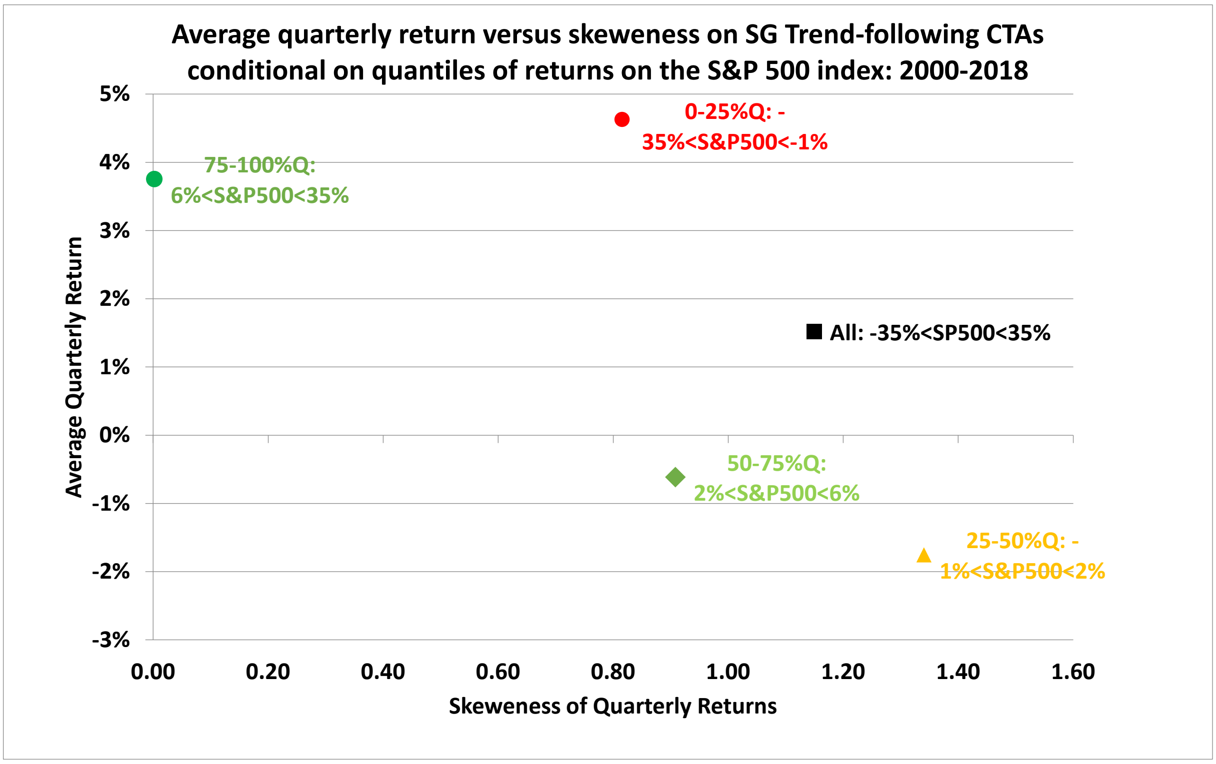

Autocorrelation as explanatory factor for trend-followers returns

Previously I showed the evidence that the trend-followers outperform when the S&P 500 index is either in a bear market or in a strong bull market. Now I turn my analysis into understanding the performance of trend-followers conditional on the autocorrelation of the S&P 500 returns. The nature of trend-followers is to benefit from markets where prices and returns are auto-correlated implying the persistence of trends.

The visual inspection of the past performance of the SG trend-following CTAs index indicates the two regimes:

1. Outperformance period from January 2000 to June 2011 with the realized Sharpe ratio of 0.49 vs Sharpe of 0.05 on the S&P 500 total return (TR) index.
2. Flat performance period from June 2011 to March 2018 with the realized Sharpe ratio of 0.06 vs Sharpe of 1.13 on the S&P 500 TR index.

One variable that has changed the sign in a significant way during these two periods is the lag-1 autocorrelation between quarterly returns on the S&P 500 index (the same conclusions are reached when using monthly returns):

1. During the outperformance period, the lag-1 autocorrelation is positive and insignificant.
2. During the flat performance period, the lag-1 autocorrelation is negative and significant.

In the table below, I summarize the realized quantities using quarterly returns.

Period			Sharpe		Lag-1 autocorrelation of quarterly returns	
	Start	End	Trend-followers	S&P 500 TR index	S&P 500 TR index	95% Significance Bounds
All Sample	Jan-2000	Mar-2018	0.35	0.35	0.06	± 0.23
Outperformance	Jan-2000	Jun-2011	0.49	0.06	0.13	± 0.29
Flat performance	Jun-2011	Mar-2018	0.06	1.13	**-0.43**	± 0.37

In Figure below I visualize the Sharpe ratios in the three samples as function of the lag-1 autocorrelation.

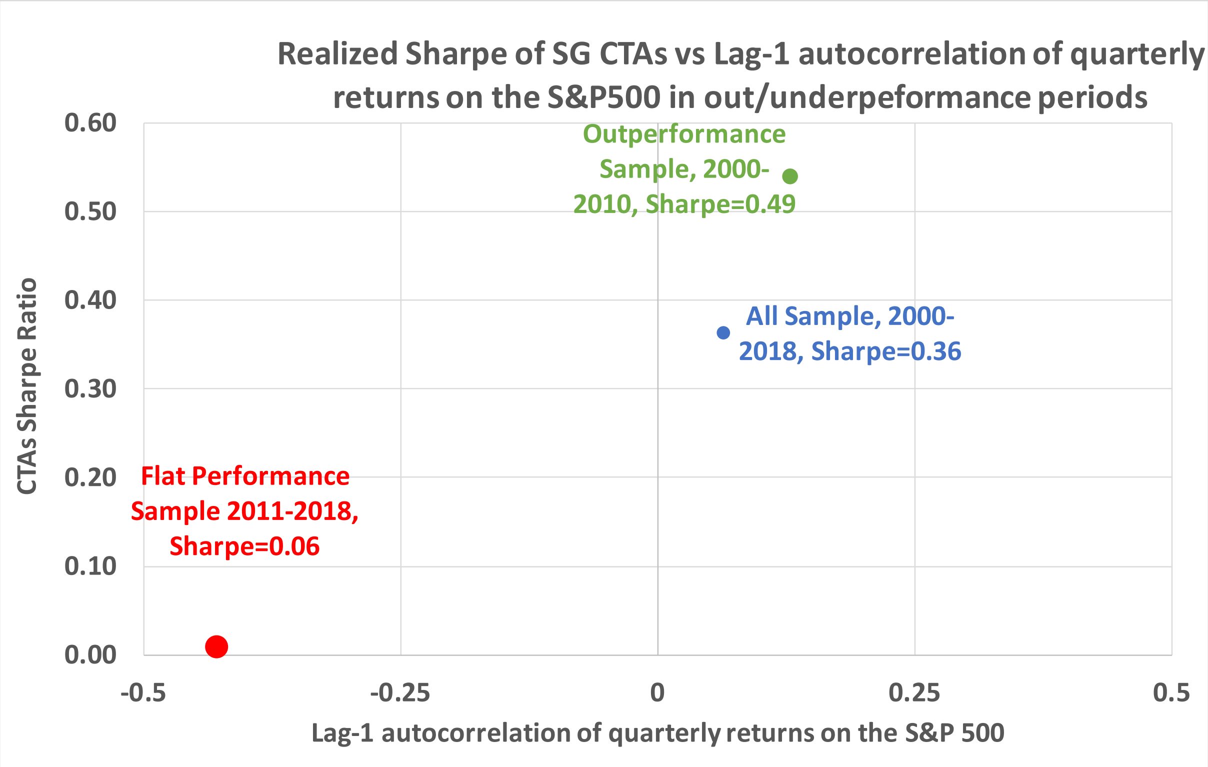

Measuring Autocorrelation

The traditional measure of the autocorrelation using lagged return is noisy for a short period of observations, say one month or one quarter.

An alternative way to measure the autocorrelation is to investigate the realized volatility of returns sampled at different frequencies:

1. When the autocorrelation is negative, returns revert to the mean over lower frequencies and the volatility of returns measured at lower frequencies, say, monthly or quarterly, will be lower than the volatility of returns measured at higher frequencies, say, daily.
2. When the autocorrelation is positive, returns tend to grow over lower frequencies and the volatility of returns at low frequency tend to increase than the volatility of returns at low frequency.
3. When the autocorrelation is zero, returns sampled at low and high frequencies tend to have the same volatility.

I introduce the following ratio of the volatility of returns measured at low frequency (weekly, monthly, etc) and the volatility of daily returns:

$$Autocorrelation\ Vol\ Ratio = \frac{Volatility\ of\ Returns\ At\ Low\ Frequency}{Volatility\ of\ Returns\ At\ Daily\ Frequency} - 1$$

The figure below illustrate the how the autocorrelation affects the autocorrelation vol ratio using a theoretical quantitative model. In the presence of strong positive autocorrelation, the volatility of quarterly returns (about 60 days in returns measurement), the ratio is expected to be around two indicating that quarterly returns are three time more volatile than daily returns. In opposite, in the presence of strong negative autocorrelation, the volatility of quarterly returns is reduced in half compared to daily returns.

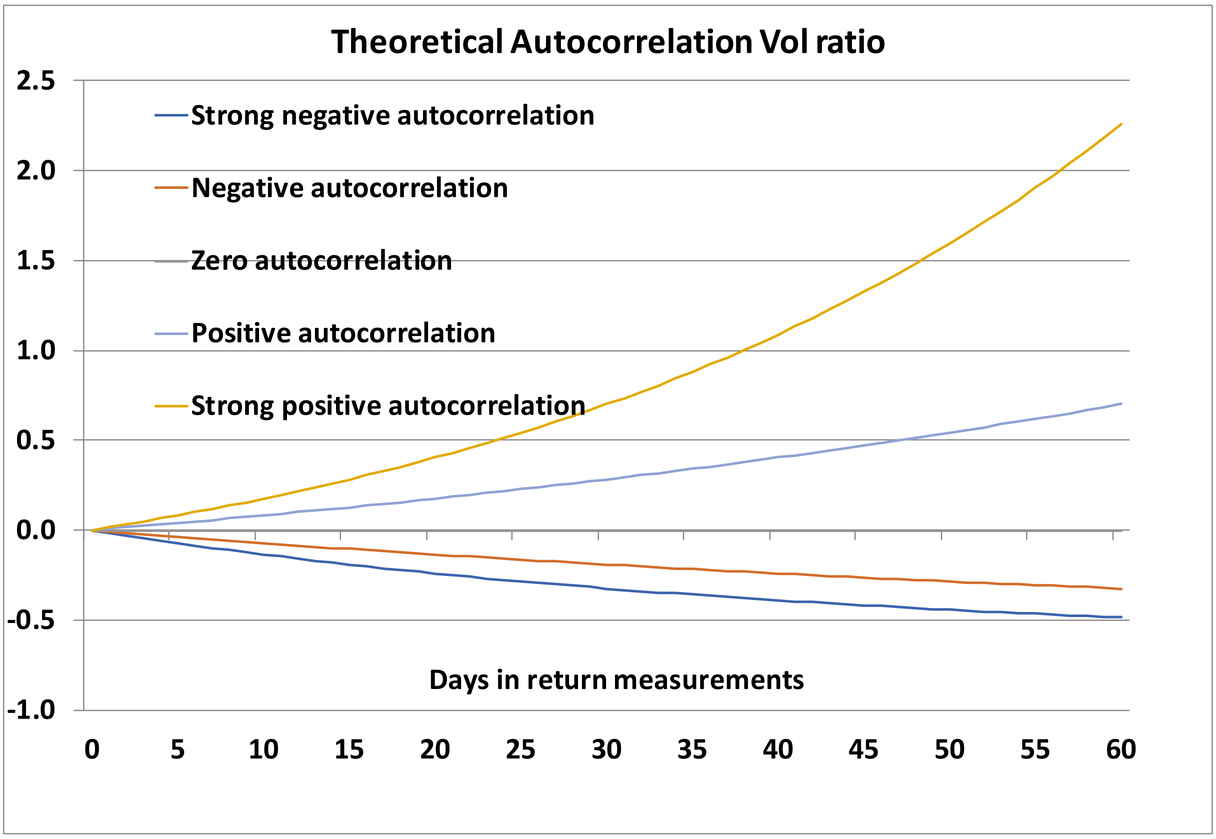

To measure the autocorrelation vol ratio from the empirical data, I introduce the following statistics defined on quarterly periods:

$$Autocorrelation\ Vol\ Ratio = c\frac{|Quarterly\ Return|}{Average\ |Demeaned\ Daily\ Returns|} - 1$$

where for each quarter I compute the average of absolute values of demeaned daily returns for denominator and apply the absolute value of the quarterly return for the numerator. Here c=sqrt(4/260) is the normalization constant so the expected value of the statistics is zero when returns are not serially correlated. I also tried more sophisticated estimators for daily volatility, but the outcomes have not been very different from each other, so I stick with the simplest measure.

The interpretation of the statistics for the measuring the autocorrelation vol ratio is the following:

1. When daily returns are large but cancel each other because of the strong mean-reversion, the realized value of the statistic will be close to -1.
2. When daily returns are smooth the statistics will be negative because of [the triangle inequality for sums](#).
3. When returns are strongly autocorrelated the statistics will be positive.

In the figure below, I show the cutoffs for the 12.5% empirical quantiles of the realized statistics using quarterly returns on the S&P500 index. We see that only about 25% of observations of the autocorrelation vol ratio is positive, indicating that autocorrelation is present only about 25% of the time with returns exhibiting mean-reversion most of the time.

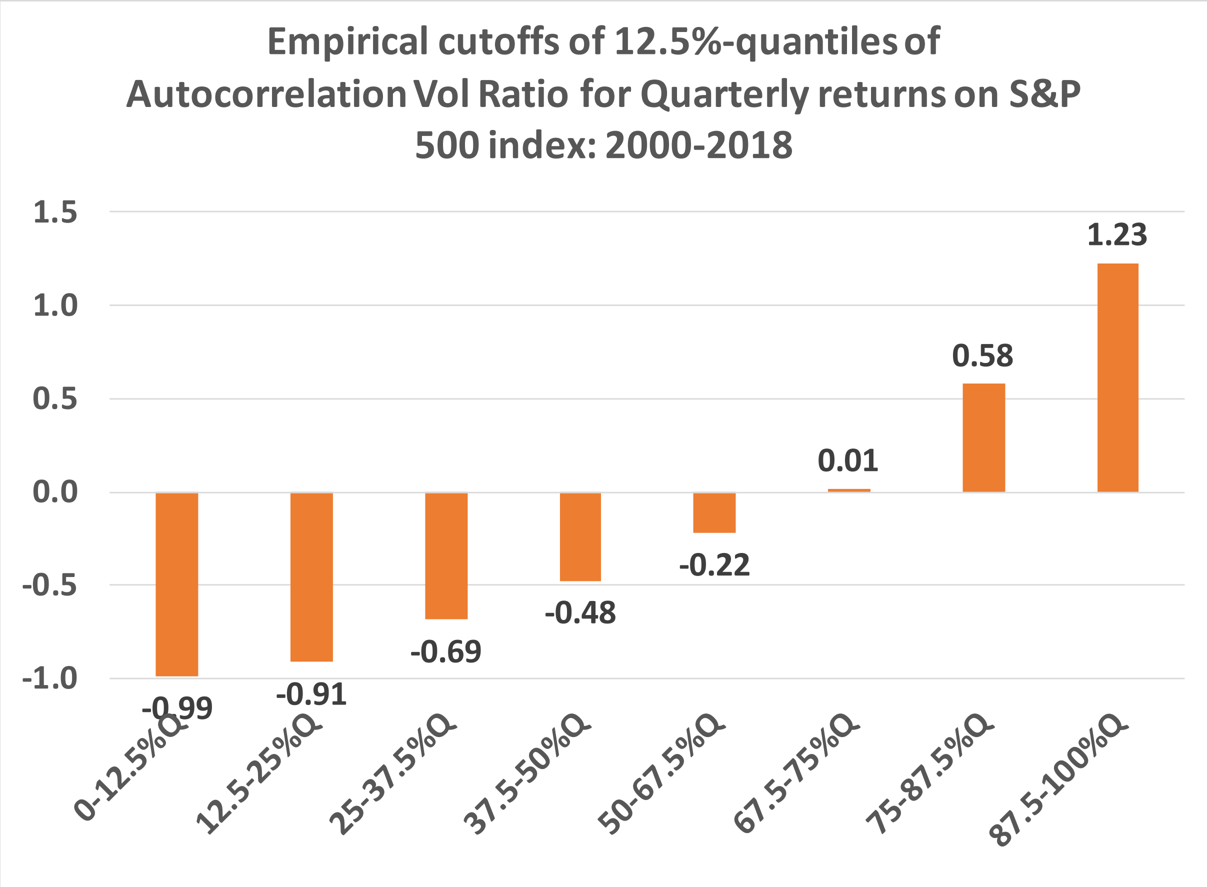

In the figure below I show the realized statistics against the quarterly returns on the S&P 500 index. We see that the statistics has a V-shaped curve. When quarterly returns are small, the value of the autocorrelation vol ratio is low. When the magnitude of return is high, the autocorrelation vol ratio increases but not in a linear way implied by the absolute value of the realized return.

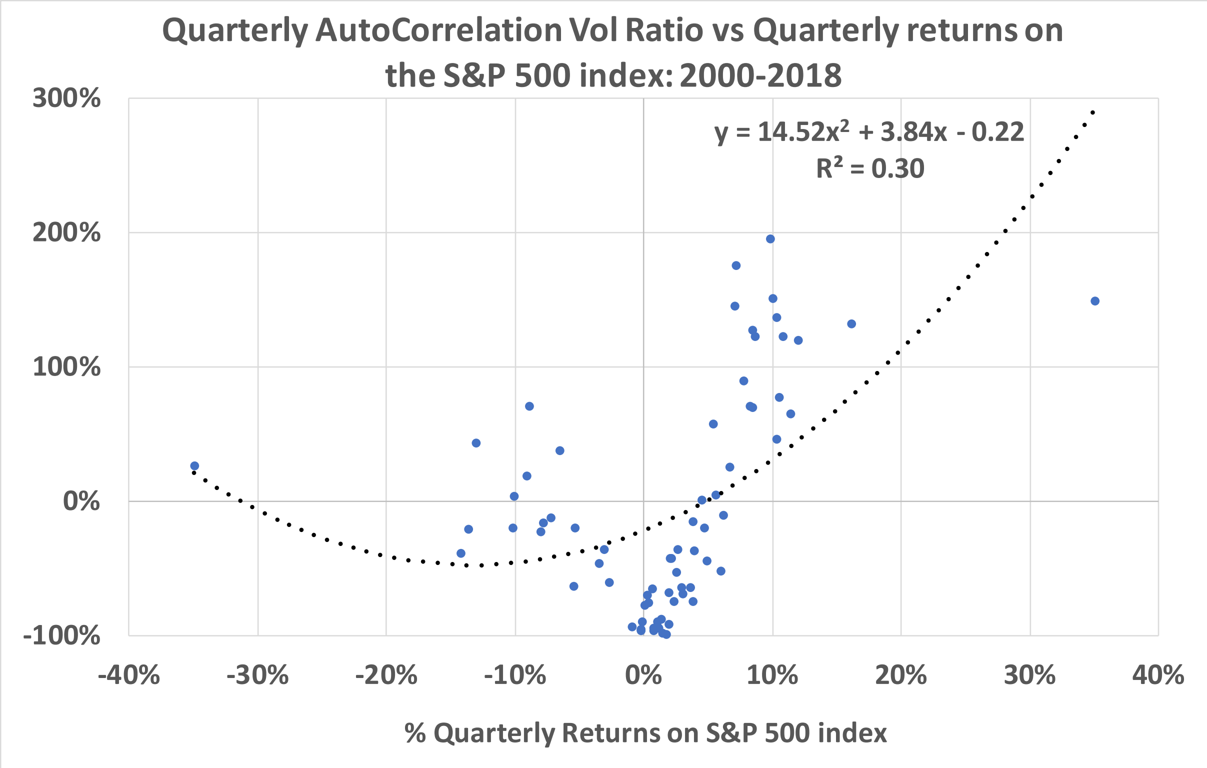

Returns on Trend-following CTAs conditional on autocorrelation

I divide the quarterly autocorrelation vol ratio on the S&P 500 index into the four buckets corresponding to four quantiles measured with 25% intervals:

1. 0-25%Q includes values between -100% to -68% interpreted as the regime with strong mean-reversion.
2. 25-50%Q includes values between -66% to -22% interpreted as the regime with moderate mean-reversion.
3. 50-75%Q includes values between -22% to 57% interpreted as the regime with moderate autocorrelated trend.
4. 75-100%Q includes monthly returns between 57% to 195% interpreted as the regime with strong autocorrelated trend

In the figure below, I illustrate the conditional quarterly returns on the SG trend-following CTAs benchmarked with returns on the S&P 500 index.

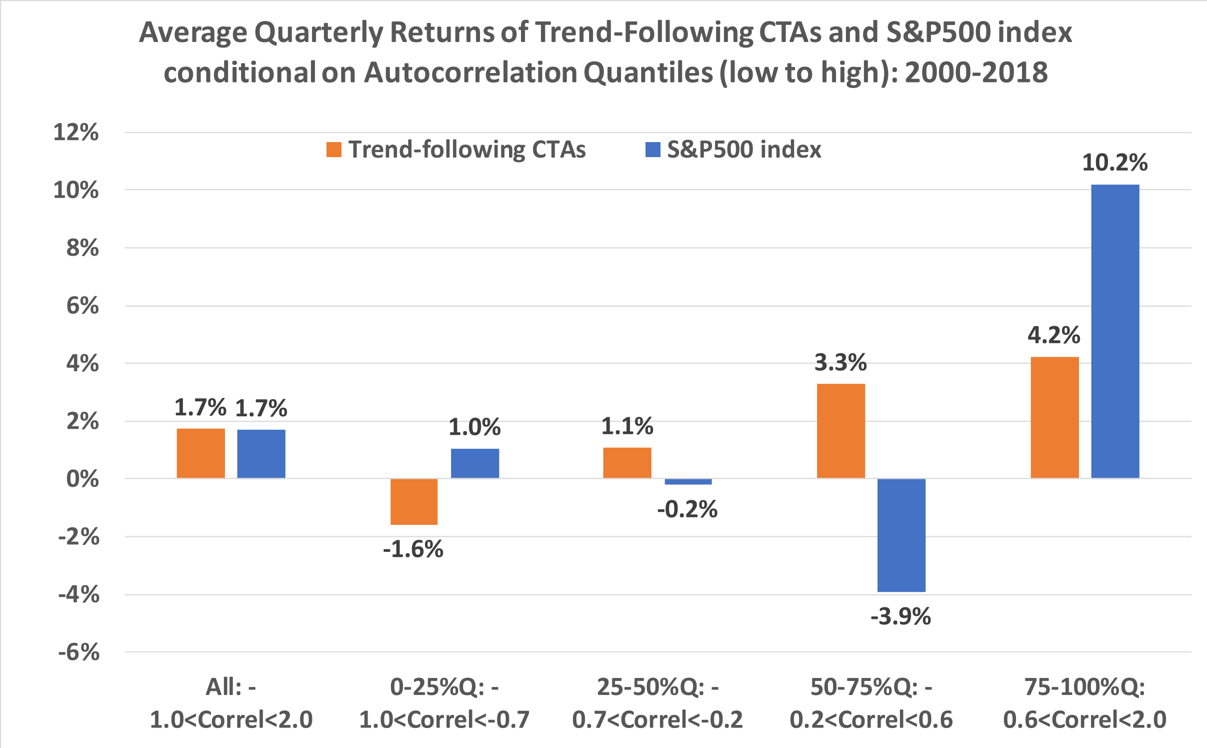

Construction of Trend-following strategy for the S&P 500 index

My aim is to understand in a quantitative way the relationship between the three variables:

1. The parameter of the trend smoothing that defines the slow- and fast-paced trend-followers,
2. The frequency of portfolio rebalancing of the trend-following system,
3. The return measurement frequency of the performance on the trend-following system.

For this purpose, I apply a typical construction of trend-following systems presented by researchers from Capital Fund Management (CFM) with some minor modifications. To keep my analysis simple, I parametrize the trend-following system by one parameter of the trend smoothing to obtain a range from fast-paced to low-paced trend-following strategies. Also, I apply the trend system only for the S&P 500 index, which I also use as the flagship benchmark, so that I can understand the convexity profile as a function of the trend smoothing parameter and frequency of portfolio rebalancing. My sole focus is on the risk profile of the trend-following strategy not on the optimization of its performance.

Trend smoothing

The trend signal is constructed using the volatility-normalized returns:

$$Volatility\ Normalized\ Return = \frac{Return}{Volatility}$$

where return is measured over uniform and preferably non-overlapping periods (daily, weekly, monthly, etc) and the volatility is the realized volatility of returns in this period, which is either measured at higher frequencies or extrapolated using autoregressive volatility models such as ARCH or GARCH. The model for volatility measurement and forecast is important to have smooth time series of volatility-normalized returns. I have addressed this topic in my research paper on volatility modelling and in my recent presentation on machine learning for volatility prediction.

The time series of the volatility-normalized return is defined on the dates of the signal update and portfolio rebalancing consisting of non-overlapping periods: weekly, monthly, quarterly, etc. At each re-balancing time t, I apply the exponential moving average (EMA) of past volatility-normalized returns:

$$EMA(R; HalfLife) = w_1 R_{t-1} + w_2 R_{t-2} + \cdots$$

with R_t being the volatility normalized return observed at period t and w_1,w_2 are the weights defined by the exponential smoothing as function of the signal half-life measured in months:

$$w_n(HalfLife) = 2^\wedge\left(-\frac{n}{HalfLife}\right)$$

The interpretation of the half-life of, say, 4 months is that the four of past months contribute 50% of the weight for the signal generation while all prior months contribute the remaining 50% of the weight. On one hand, the higher is the half-life, the less weight is put on front months with the trend-following system being backward-looking over the extended period of past observations. On the other hand, the smaller is the half-life, the more weight is put on the front months with the system becoming more reactive to recent outsized returns.

The magnitude of the half-life defines pace of the trend-following system:

- Low-paced trend-followers apply large values of the half-life parameter to track longer-dated trends with little impact from recent outsized returns.
- Fast-paced trend-followers use shorter windows with small values of the half-life parameter to track shorter-dated trends with more weight on recent outsized returns.

In my analysis, I will apply the five values of the half-life parameter that define the range of 5 trend-following systems as illustrated in the figure below.

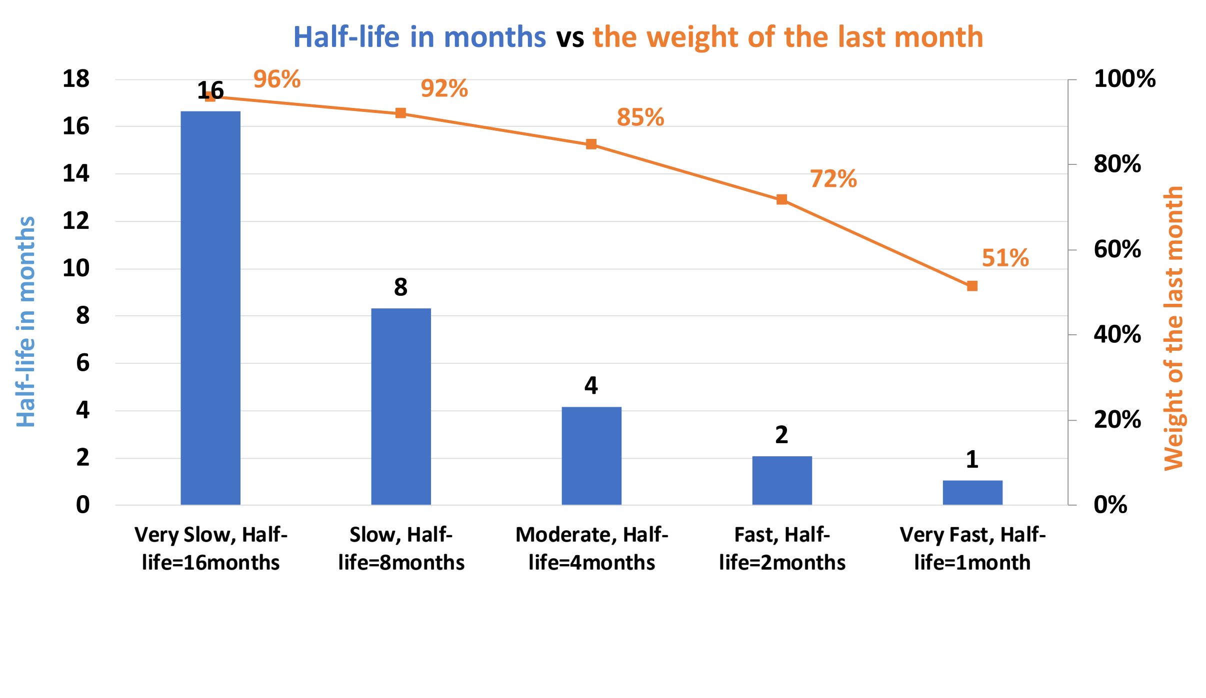

Position size generation

One of the advantages of using the volatility-normalized returns for signal generation is that the time series of the volatility-normalized returns has approximately unit standard deviation across different assets. One of the most important results from CFM paper is that, if the trend-following system uses the EMA of volatility normalized returns to size the position for the next period, the trend-following system is convex with respect to the underlying asset.

I follow the same approach as in CFM paper, but I apply some normalization of the signal to compute the notional size at each portfolio rebalancing date as follows:

$$NotionalSize = f(E(VolNormalizedReturn; HalfLife))$$

where f is a signal normalization function that also caps the notional between -100% and 100%. As a result, the system will be always allocated to the underlying asset. The allocation will increase when the EWA of volatility-normalized returns increases and vice versa. To avoid excess positions, the notional will be kept within +-100%.

The figure below shows the EWA signal and the position size for the fast-paced trend-following system with the fast signal updated monthly. We see that the position size is changed frequently.

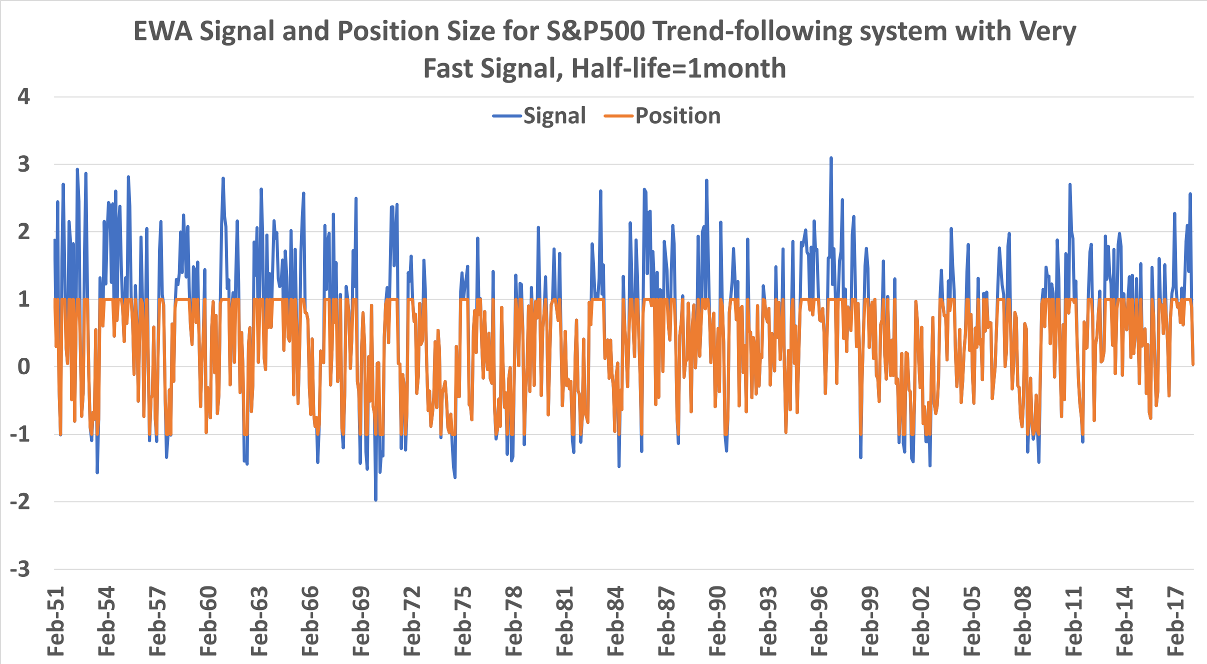

The figure below illustrates the medium-paced system in which the position size changes less frequently.

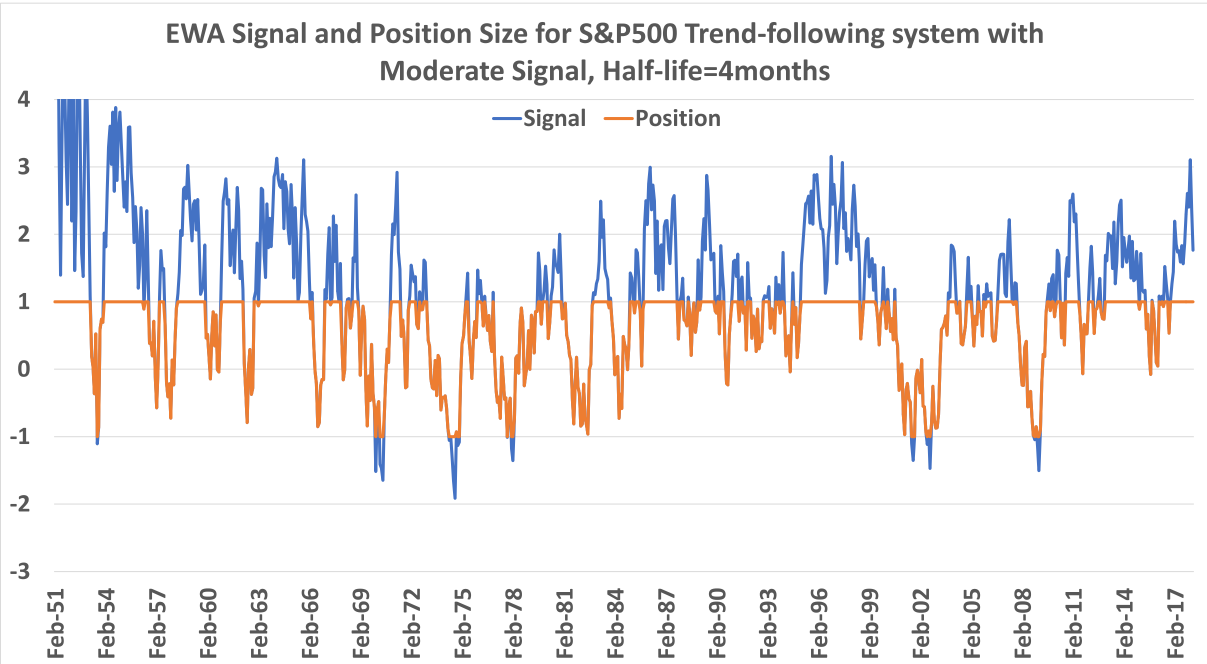

Finally, the figure below illustrates the slow-paced trend-follower with very infrequent rebalancing of the position keeping predominantly long exposure to the S&P 500 index because of the strong positive performance of the index.

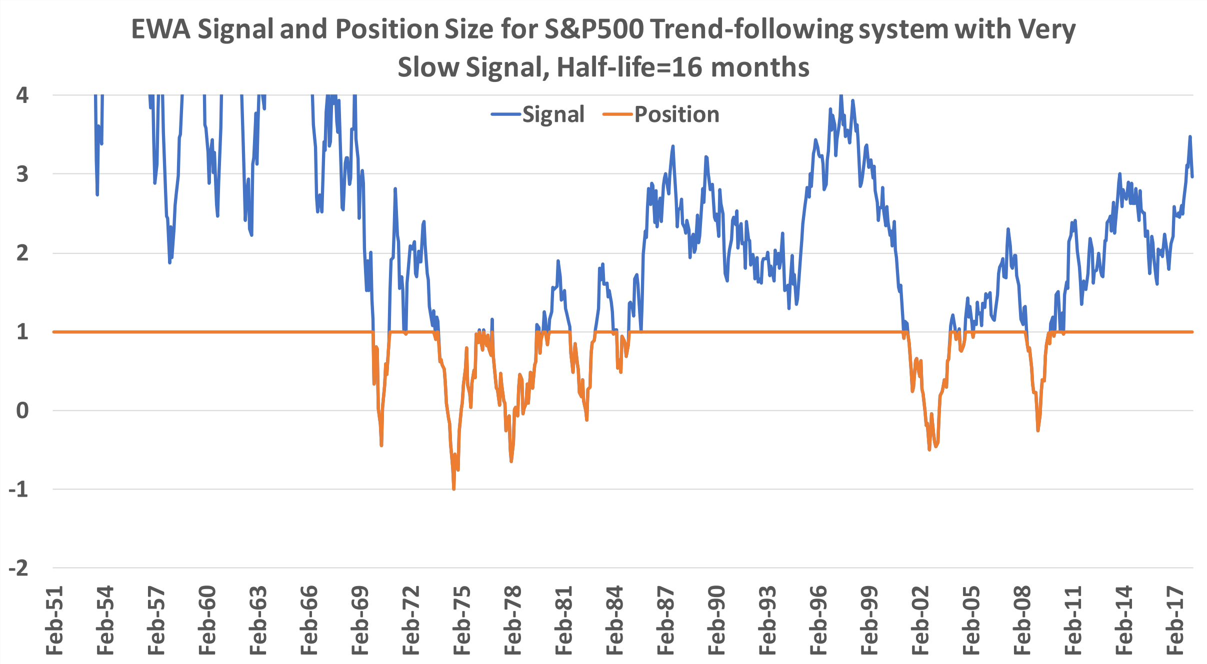

S&P 500 Trend-following strategy

To construct the time series of the S&P 500 trend-following strategy I apply the data of S&P 500 index from 1950. I use the S&P500 index as a proxy for the S&P 500 futures contract. When the dividend yield on the index is close to the short-term interest rate, there is very little difference in the performance between the index and the index future contract. Again, my focus is on the risk-profile of the trend-following strategy, so I omit the roll-costs which are not significant for index futures and do not alter the risk-profile of the trend-following strategy.

The time series for the S&P 500 index is available from 1928 but in the current form the index started in 1957, when the index expanded to the current number of 500 constituents. I apply the period starting from 1950, with the period from 1928 to 1950 having minor change in the overall performance.

The goal is to analyze the convexity of the strategy to the S&P 500 index which will serve as the benchmark. The strategy serves as close as possible as the tail protection for the long position in the S&P 500 index.

As the indicative benchmarks I use two trend-following CTA indices.

1. Barclay BTOP50 Index which is the benchmark for managed futures including trend-following CTAs with monthly time series available from 1987.
2. SG trend-following CTAs which I have already applied through my note with time series starting from 2000.

By varying the half-life of the signal, I obtain the highest correlation between monthly returns with the two benchmarks using the moderate smoothing with the half-life of 4 months and monthly rebalancing of the position.

In the table below, I report the correlation of the strategy with benchmarks. Remarkably, my strategy has a strong correlation with both the BTOP and SG trend-following CTAs from 2000s. Given that the typical strategic allocation to stock indices in CTA programs is between 20 to 40%, the obtained correlation of about 40% is indicative that my replication matches the stock index profiles of CTAs.

Correlation of monthly returns for S&P 500 trend strategy			
	1951-2018	1987-2018	2000-2018
S&P 500	0.40	0.48	0.03
BTOP 50		0.18	**0.36**
SG Trend-following CTAs			**0.42**

In the table below, I report the realized Sharpe ratios on the S&P 500 trend-following strategy and benchmarks. In terms of the performance, my S&P 500 trend-following strategy provides a good match to a typical CTA program.

Realized Sharpe ratios for S&P 500 trend strategy			
	1951-2018	1987-2018	2000-2018
S&P500	0.58	0.60	0.30
S&P 500 trend strategy	0.54	0.65	**0.54**
BTOP 50		0.78	**0.50**
SG Trend-following CTAs			**0.45**

In the figure below, I show the time series of the logarithm of the NAV for the strategy and benchmarks. The starting point for BTOP and SG trend-following CTAs is set to equal the corresponding NAV of the S&P 500 trend strategy.

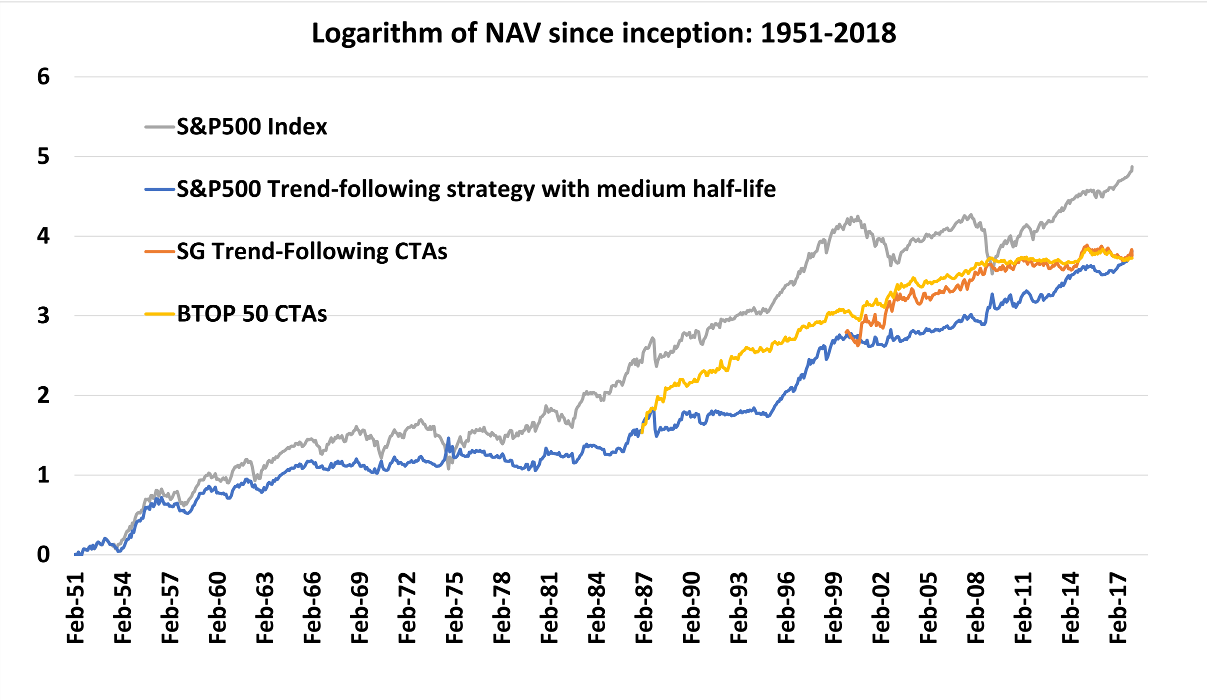

Risk profile of S&P 500 Trend-following strategy

The trend-following system developed in the previous section is parametrized by the two internal parameters:

1. The half-life of the trend smoothing from slow-paced to fast-paced trend-following systems.
2. The frequency of portfolio rebalancing for each system: weekly, monthly, etc.

Intuitively, slow-paced systems should be rebalanced infrequently, while fast-paced systems should be rebalanced frequently. I apply one external parameter of return measurement frequency to measure the risk-profile of each specification of the trend-following systems.

Convexity profile

First, I illustrate the convexity profile using the S&P 500 trend-following strategy with trend smoothing using medium half-life of 4 months. For the benchmark, I also use the S&P 500 index. In contrast to the previous analysis using SG trend-following CTAs, in the case of applying the trend-following system only on the S&P 500 index and analyzing the realized convexity with respect to the S&P 500 index reduces the amount of the noise. As a result, we can see a clearer relationship between the trend-following strategy on the flagship index and the index itself.

Monthly returns. We observe a clustered profile because the trend-following strategy is expected to be either short or long the S&P 500 index over whole month, because the half-life of 4 months is well above the return observation frequency of one month. As a result, the strategy has no clear convexity profile.

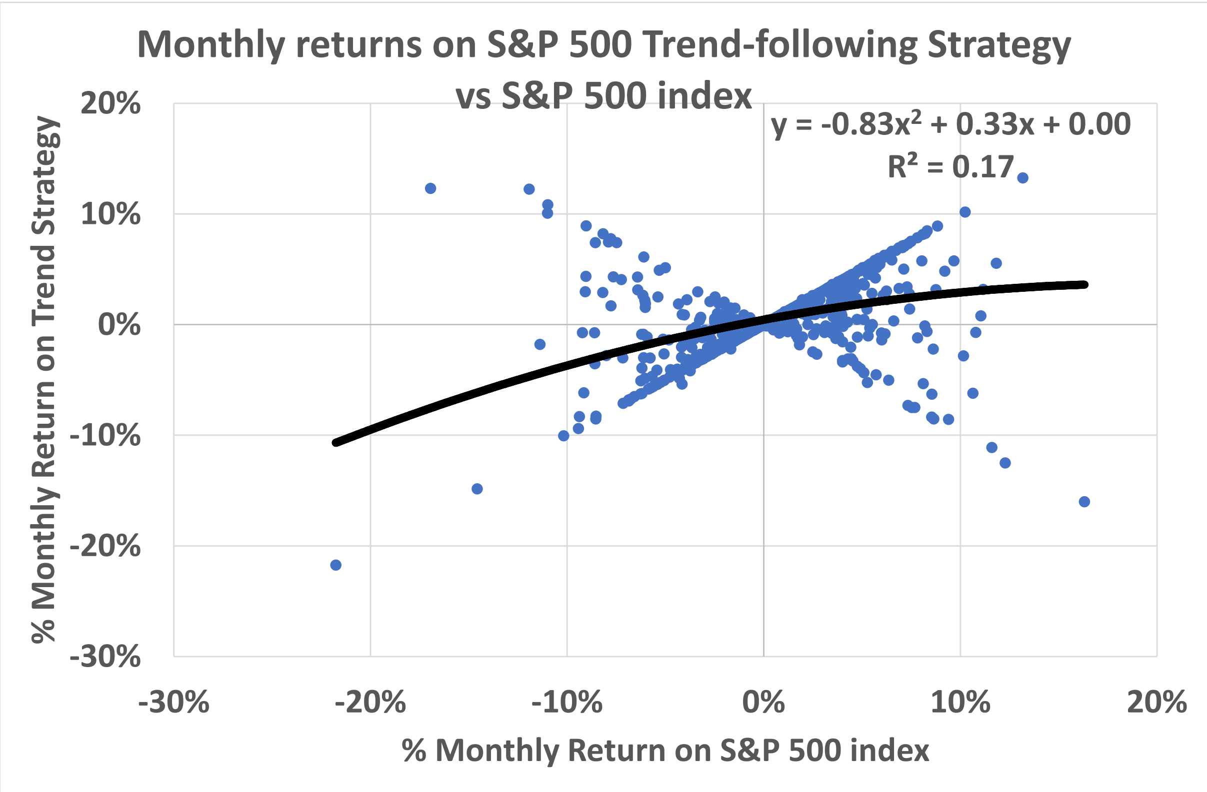

Quarterly returns: we observe a clearer convex profile because the quarterly return measurement is comparable with the half-life of four month for the medium-paced strategy. For certain quarters we see that the strategy can generate large gains for the large losses on the S&P 500 index. However, the strategy also produces a couple of outliers. On one hand, it generates a large loss of the same size as a large loss on the S&P 500 index; on the other hand, it produces a large loss when the S&P 500 index generates a large gain.

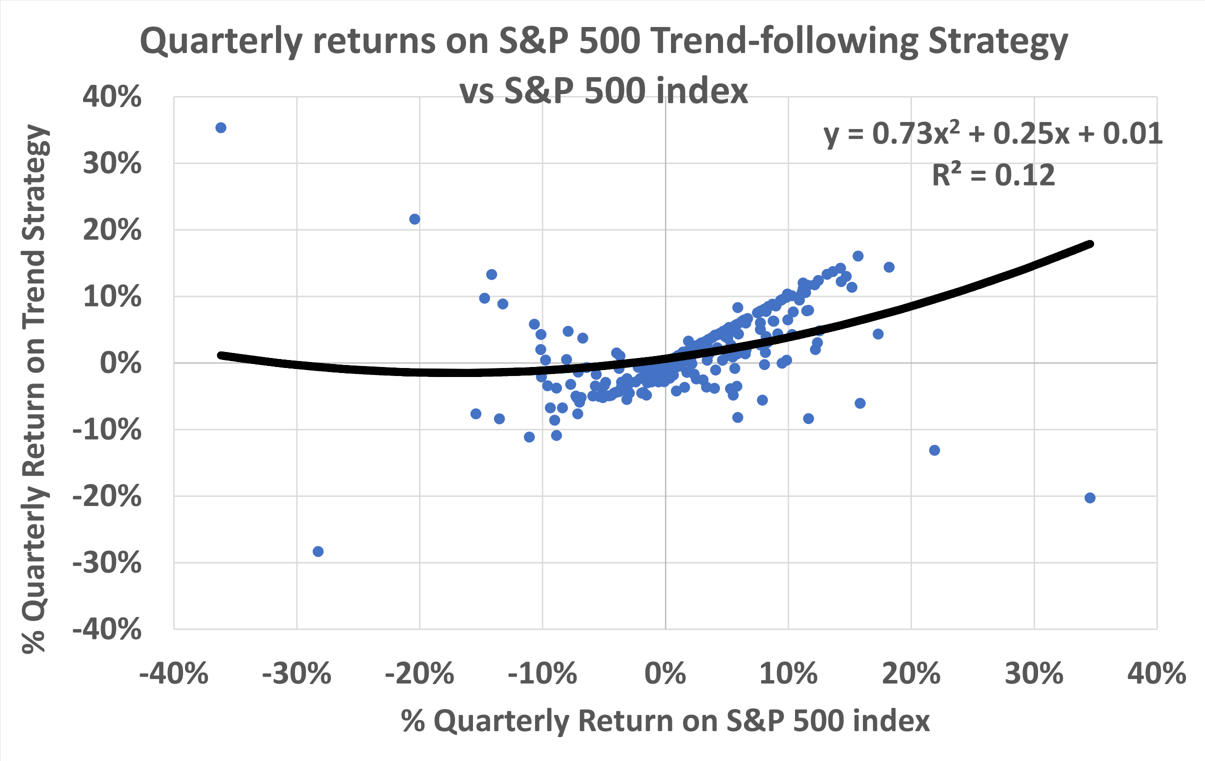

Annual returns. We observe a clear convex profile with the strong explanatory power. The idiosyncrasies observed at lower frequencies of returns sampling disappear for the annual sampling because the length of returns sampling exceeds the half-life of the trend signal.

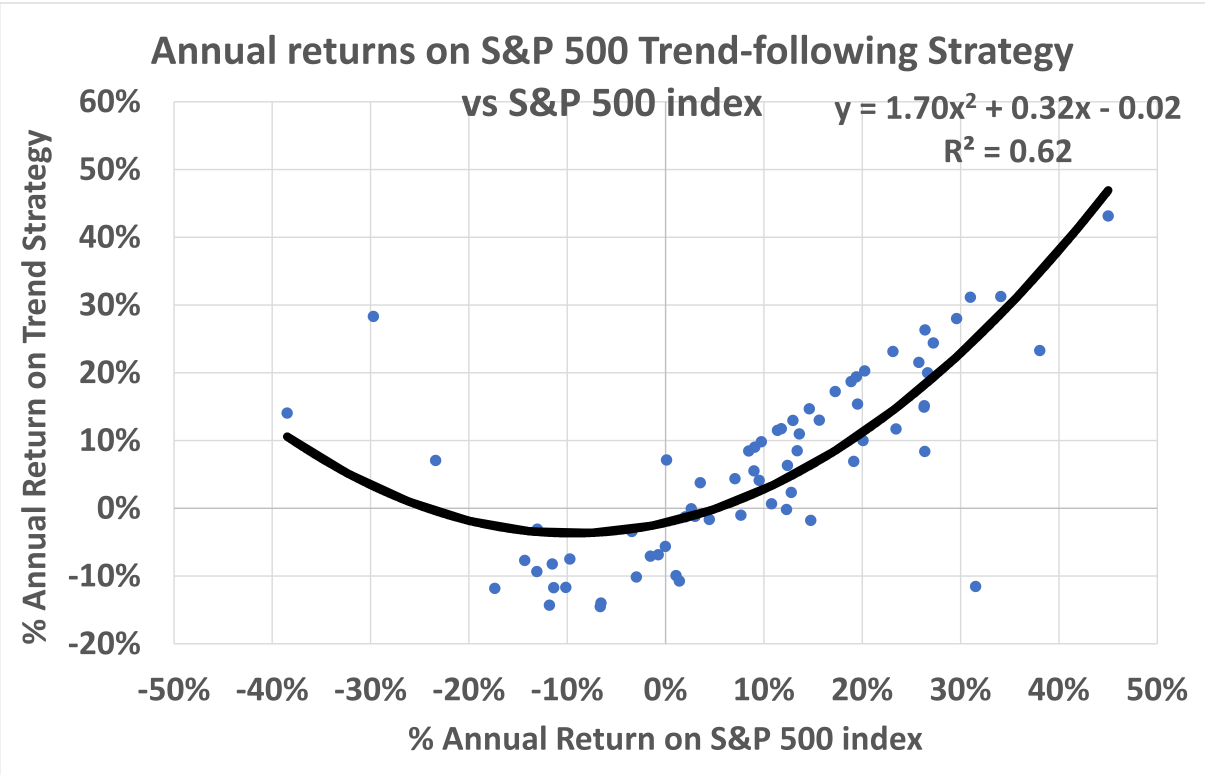

These three observations about the impact of return frequency measurement are similar to results obtained using the SG trend-following CTA index. The key conclusion is that the trend-following system can generate the positive convexity when the return measurement period exceeds the half-life of the trend-smoothing and the period of portfolio rebalancing.

Impact of the frequency of the position rebalancing

Now, I analyze how the strategy risk-profile is affected by the frequency of the position rebalancing. For the consistency, I always define the half-life of the signal in months and the strategy is rebalanced monthly (as before), quarterly, and annually. I have also checked strategies with the weekly rebalancing where the half-life is defined on the weekly scale. The strategy with weakly rebalancing has a similar risk-profile to the strategy with monthly rebalancing. As a result, I only analyze the trend-following strategies with different values of the half-life defined on the monthly scales and rebalanced monthly, quarterly, and annually.

Convexity Betas

First, I analyze the convexity betas as functions of the trend half-life and the frequency of position rebalancing.

Monthly returns: the portfolio rebalancing does not have any impact because the observation frequency of realized returns is too fast relative to the adaptation speed of the trend-following system. The convexity betas are negative for all systems except for the very fast system with the annual rebalancing (in this case, the strategy is rebalanced annually using past months returns with the half-life of one month).

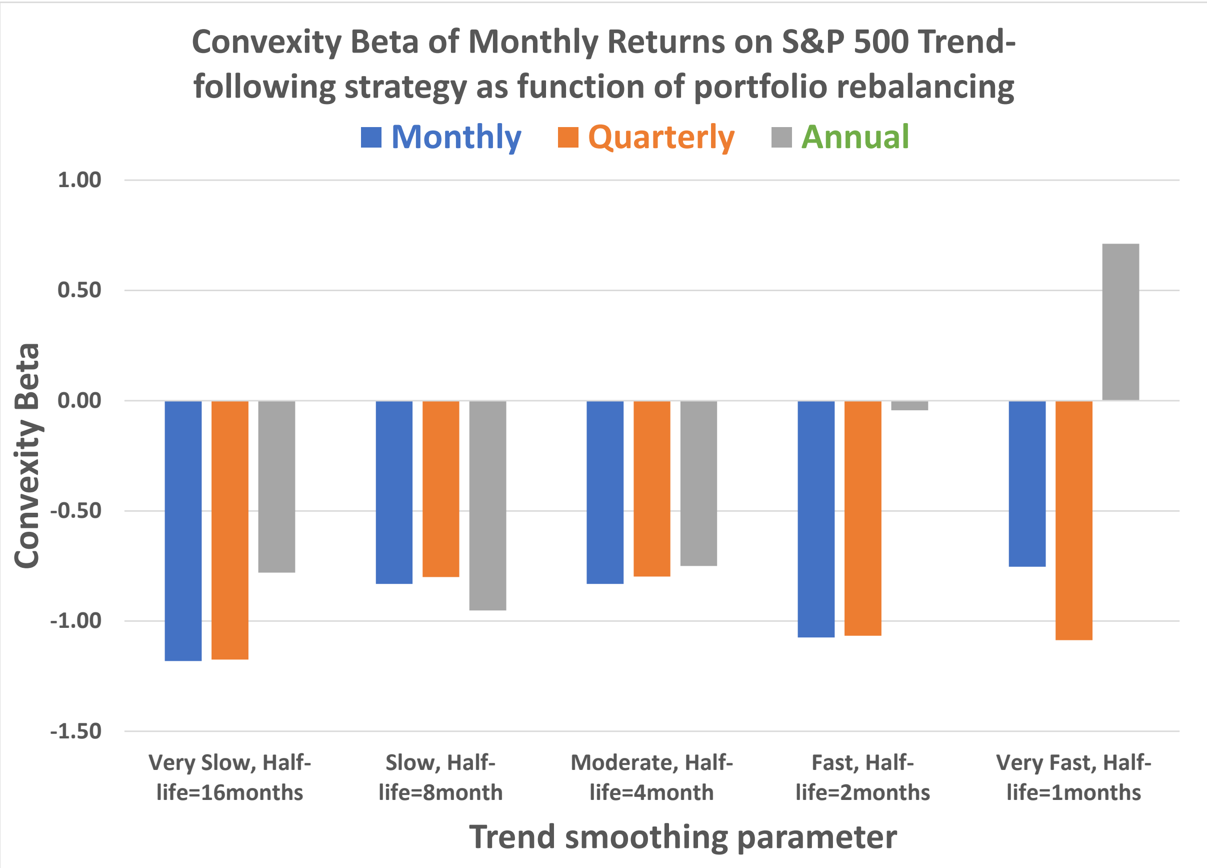

Quarterly returns: the trend-following strategy with monthly updates generates positive convexity starting with the moderate-paced system with the half-life of 4 months. As for fast and very fast systems, the half-life declines further, the convexity of the strategy strengths. As an indication, trend-following strategies with moderate to fast trend smoothing can adjust within one quarter and generate positive convexity. However, the strategies with quarterly updates are unable to generate the convexity as their rebalancing is the same as the returns measurement frequency.

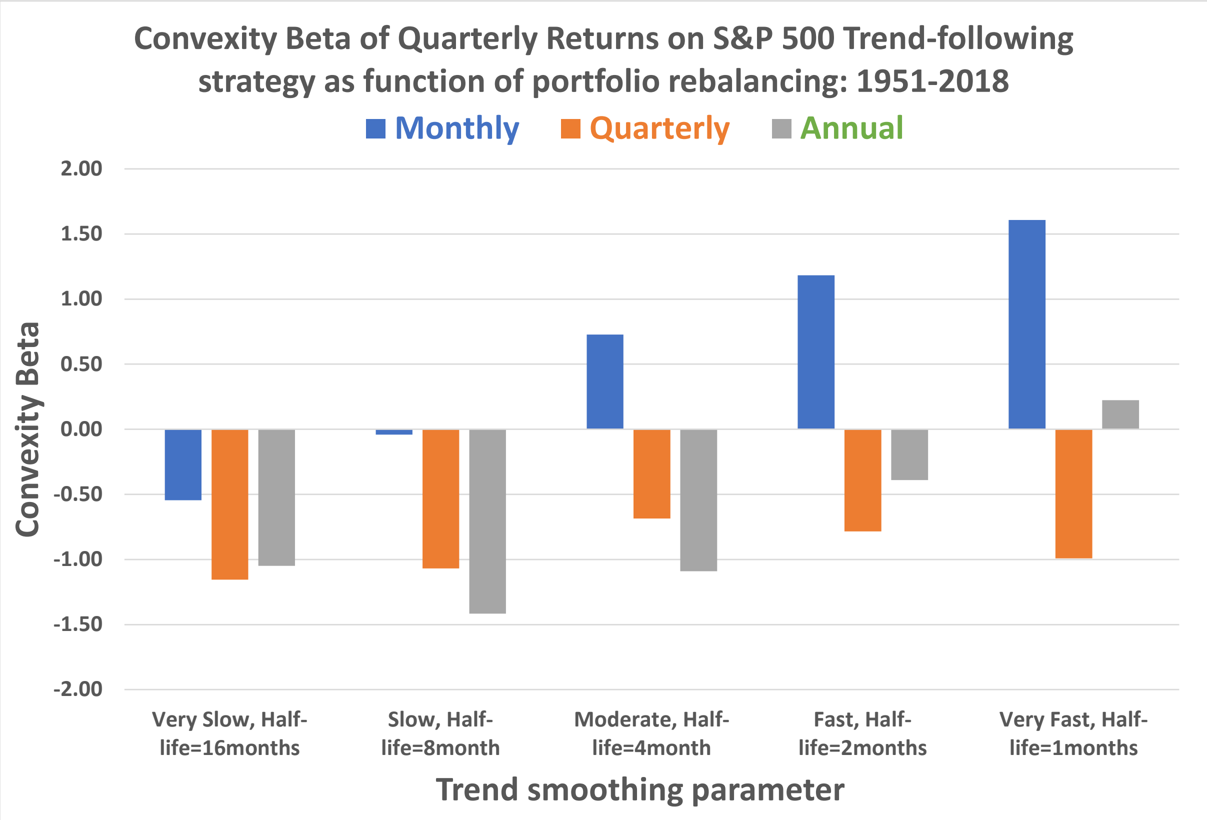

Annual returns: the trend strategies with monthly and quarterly updates generate positive convexity if the half-life is relatively small starting 8 months. The strategy with annual rebalancing is unable to generate the convexity as its adjustment period is too narrow relative to the return measurement frequency.

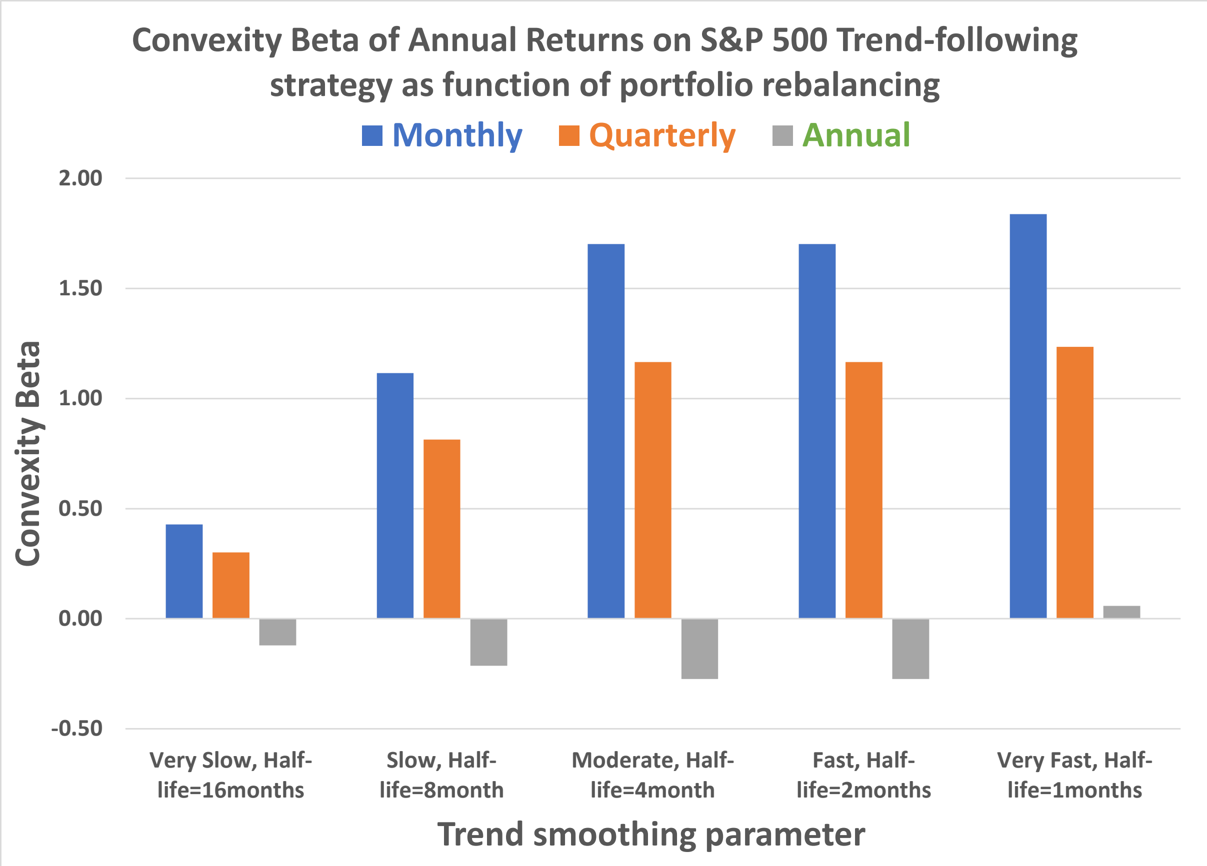

Conclusion: the target frequency of return sampling for investor portfolio and whether the trend-following strategy is sought as a hedge to generate positive convexity is important to decide over the strategy:

- For monthly returns, the convexity beta is not significant for all trend-followers no matter of the portfolio rebalancing frequency.
- For quarterly returns, the positive convexity is produced by trend-followers with monthly rebalancing and with fast trend-smoothing with the half-life more than about 4 months.
- For annual returns, the positive convexity is produces by trend-followers with monthly and quarterly rebalancing and with relatively fast trend-smoothing with the half-life less than 8 months.

Skewness

The profile of the skewness of strategy returns is similar to the convexity profile considered above. In the figure below I show the realized skewness for quarterly returns. Only the trend-following strategy with monthly rebalancing

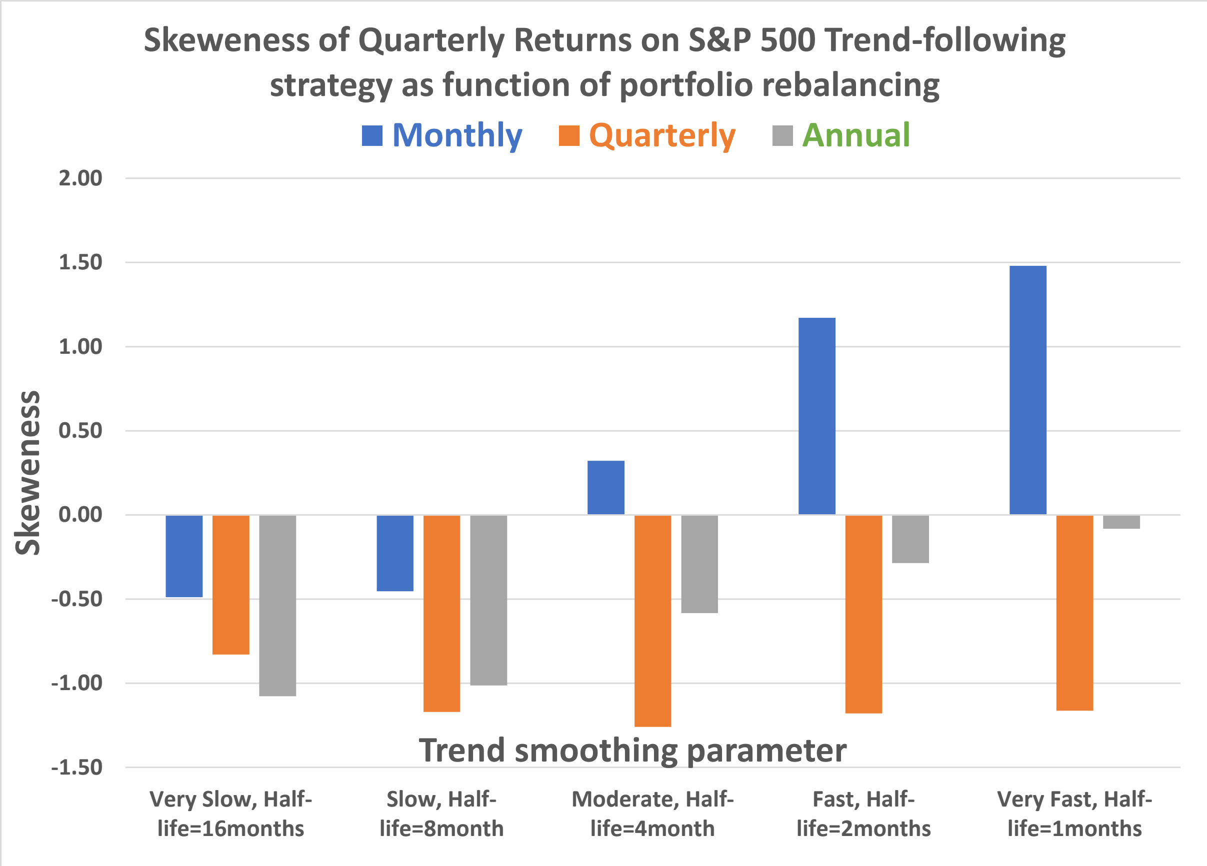

Linear beta

In the figure below, I plot the linear beta of the S&P 500 trend-following strategy for the quarterly returns. We see that, for the slow-moving trend smoothing, the strategy maintains the long exposure to the index with infrequent rebalancing. As a result, the higher is the half-life of trend smoothing, the higher is the strategy linear exposure to the index. The strategy rebalancing frequency does not significantly impact the linear beta.

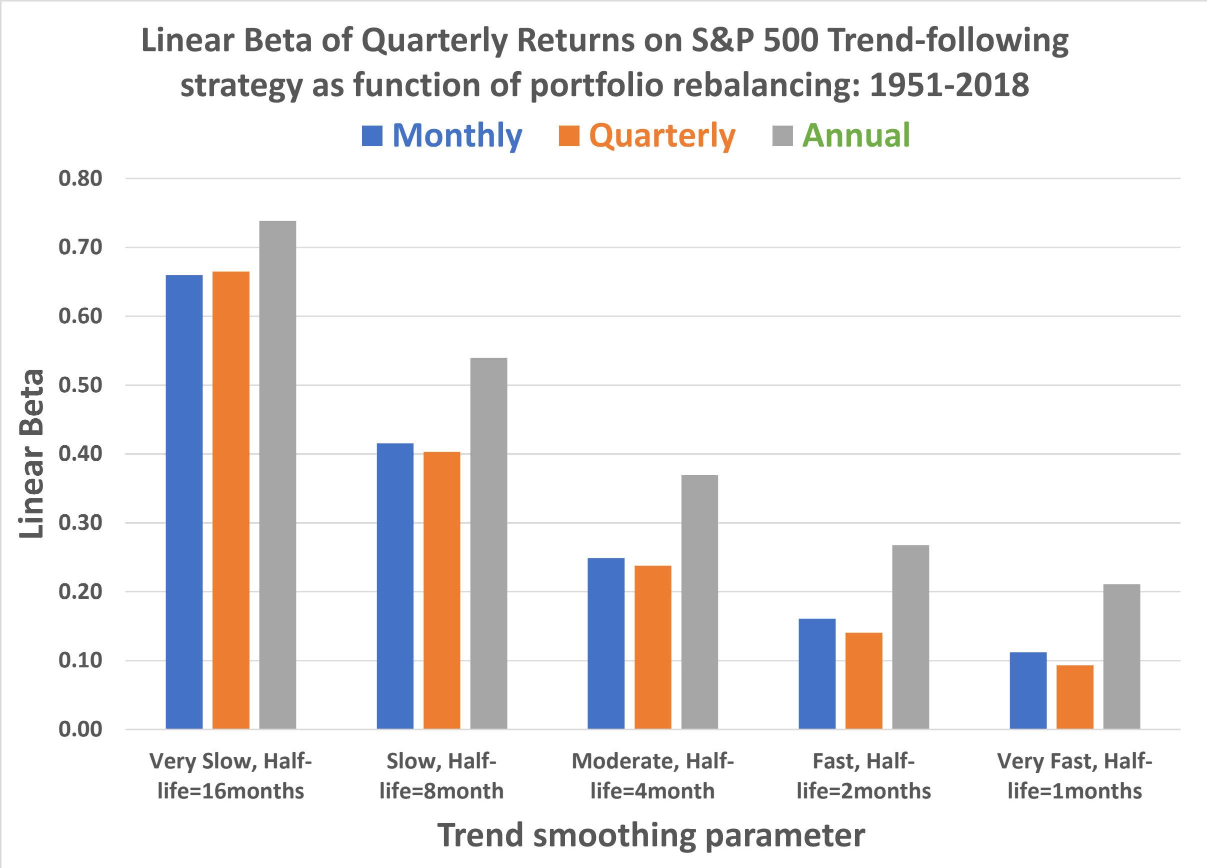

Performance

In the figure below, I show the realized Sharpe. As the half-life increases the strategy Sharpe declines. This may be a feature that the performance on S&P 500 index was positive and significant over the sample period (7.4% p.a. with realized Sharpe of 0.52).

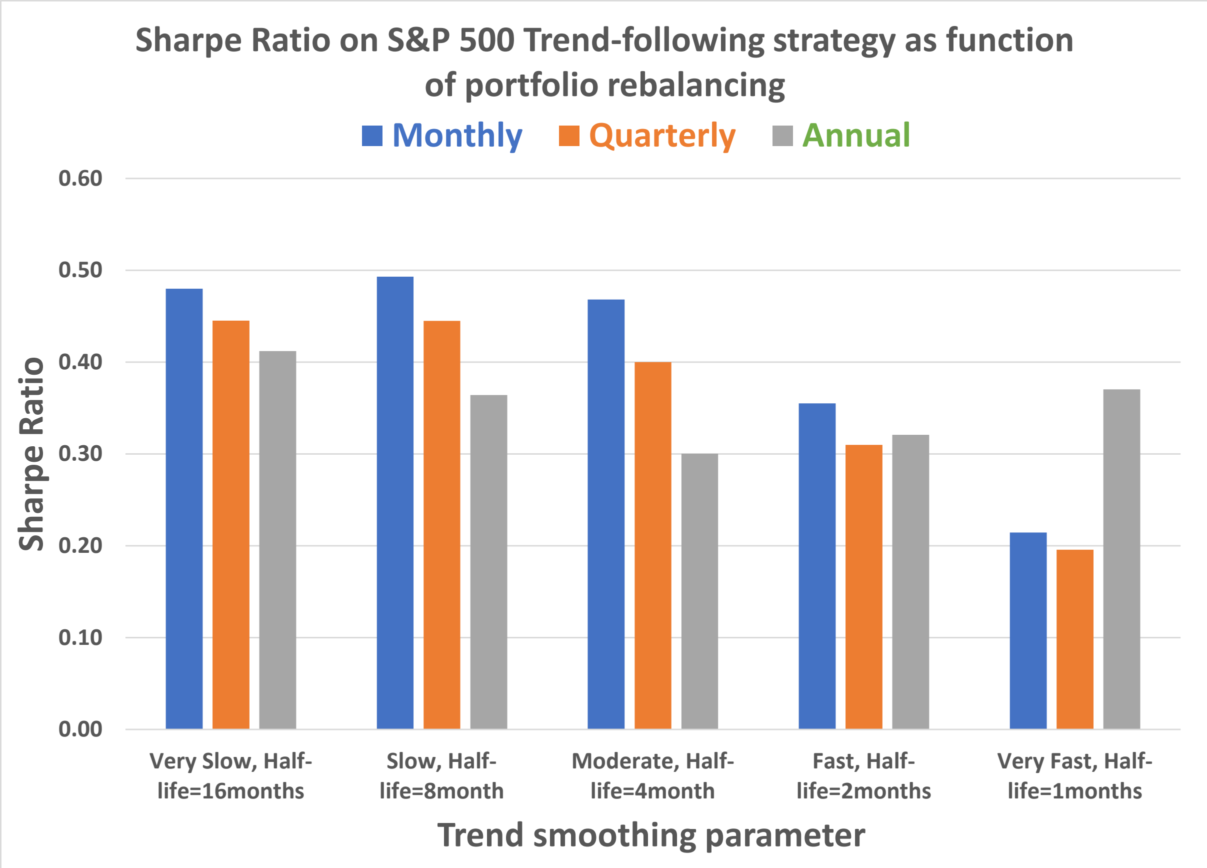

In figure below, I show the logarithm of the NAV of S&P 500 trend-following strategies as function of the half-life. The strategies with very fast signal and low half-life tend to underperform when the underlying asset is upward drifting over long horizons. While these strategies may have a good risk-profile when returns are measured infrequently, the too frequent rebalancing is detrimental for the performance, especially, in the presence of transaction costs.

The trend-following strategy with moderate trend smoothing with the half-life of about 4 months has the optimal risk-profile with reasonable performance with respect to the underlying index.

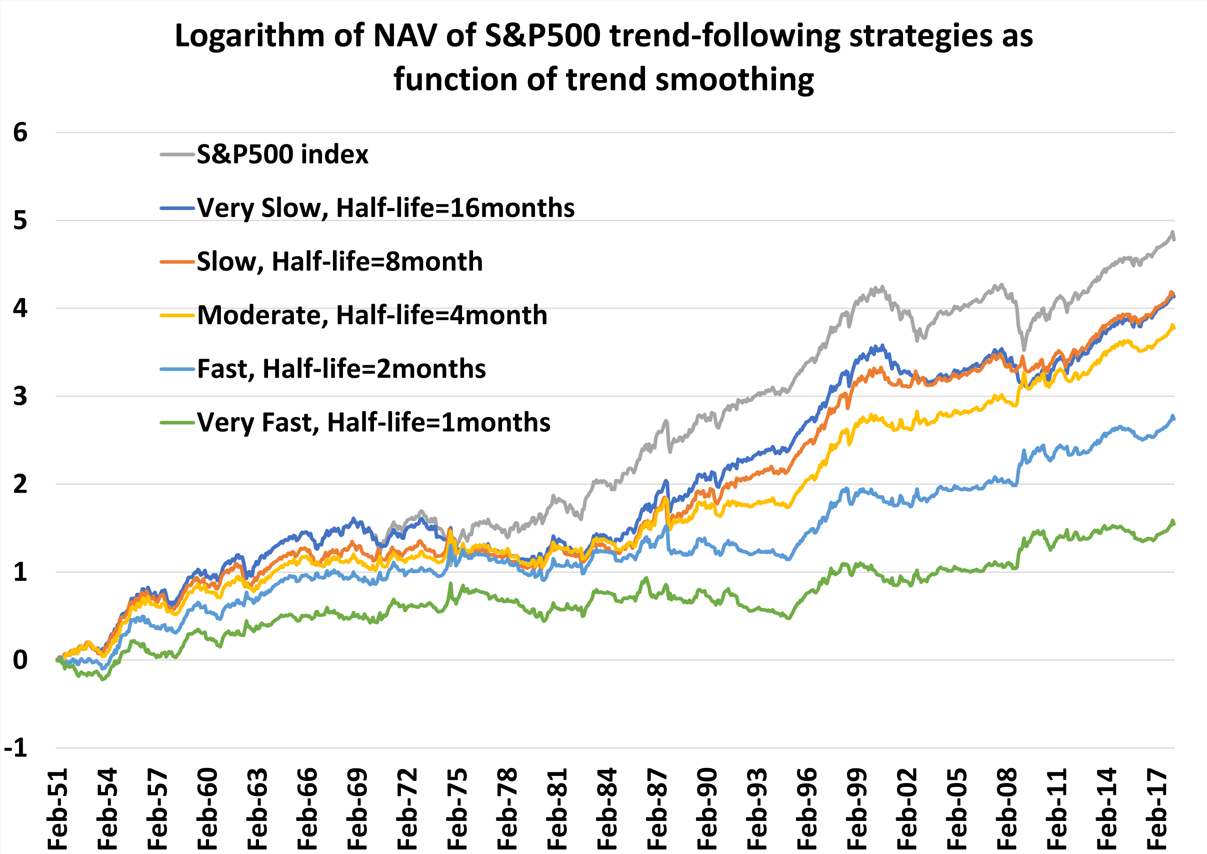

Trend-following vs Stock Momentum

One of the interesting questions is how the trend-following on futures is different from the stock momentum and if there is any overlap between the two.

The essence of the trend-following strategies is to apply smoothing and extrapolation of recent history for both signal generation and position sizing for trading in major futures markets. The ground of the equity momentum is to compute the return over the past 12 months minus the last month return for each stock in the defined universe and generate signals using the cross-sectional ranking.

In practice, the stock momentum strategies are implemented as follows.

1. **The long-short momentum**: the top performing stocks are selected for long leg and the bottom of worst performing stocks are selected for the short leg. Typically, one uses top 80-90% quantiles for the long leg and 10-20% for the short leg with all positions allocated equal weights. Practical problems are the construction of the short side where some stocks are hard to borrow and borrow costs can be high. For this reason, this approach is not common in asset management, in particular, for long-only managers. In literature the long-short momentum factor is associated with Carhart four-factor model.
2. **The long only format**: the portfolio is constructed by allocating to the top quantile of stocks with equal position sizes. All major index providers and ETF managers construct momentum indices in this format.
3. **The market-hedged momentum**: the long leg is constructed as for the long only format while the short leg is constructed using the flagship index for the defined stock universe. Typically,

the short leg is implemented using index futures with position size equal to the notional size of the long leg. The market neutral format is one of the recent developments in the alternative beta industry. The market-hedged momentum factor is popular among providers of the absolute return solutions.

I apply the following strategies in my analysis.

1. **Long-only Momentum** strategy obtained using MSCI momentum index for US stocks constructed in the long only format (Bloomberg ticker: M2US000$ Index).
2. **Market-hedged Momentum** strategy computed using MSCI index for the long leg and S&P 500 index for the short leg. I use notional matching for the long and short legs with annual rebalancing. I apply the leverage of 2 so that the in-sample volatility of the strategy is close to the volatility of the long-only momentum index.
3. **SG Trend-Following CTAs** is the flagship index of 10 trend-following CTAs.
4. **S&P 500 Trend-following strategy** is my own trend-following strategy using the medium-paced signal as described before.

As the benchmark I apply the S&P500 total return index. In the table below, I report the risk-return profile of this strategies from 2000 to April 2018.

Risk-Return Statistics using monthly returns: 2000-2018ytd					
	Return p.a.	Vol	Sharpe	Skewness	MaxDD
S&P 500 TR Index	4.9%	14%	0.34	-0.53	-55%
Long-only Momentum	7.1%	15%	0.47	-0.71	-56%
Market-hedged Momentum	4.7%	15%	0.31	-0.56	-38%
SG Trend-Following CTAs	5.1%	14%	0.36	0.20	-21%
S&P500 Trend Strategy	5.2%	11%	0.47	0.05	-24%

In the figure below I show the scatter plot of monthly returns

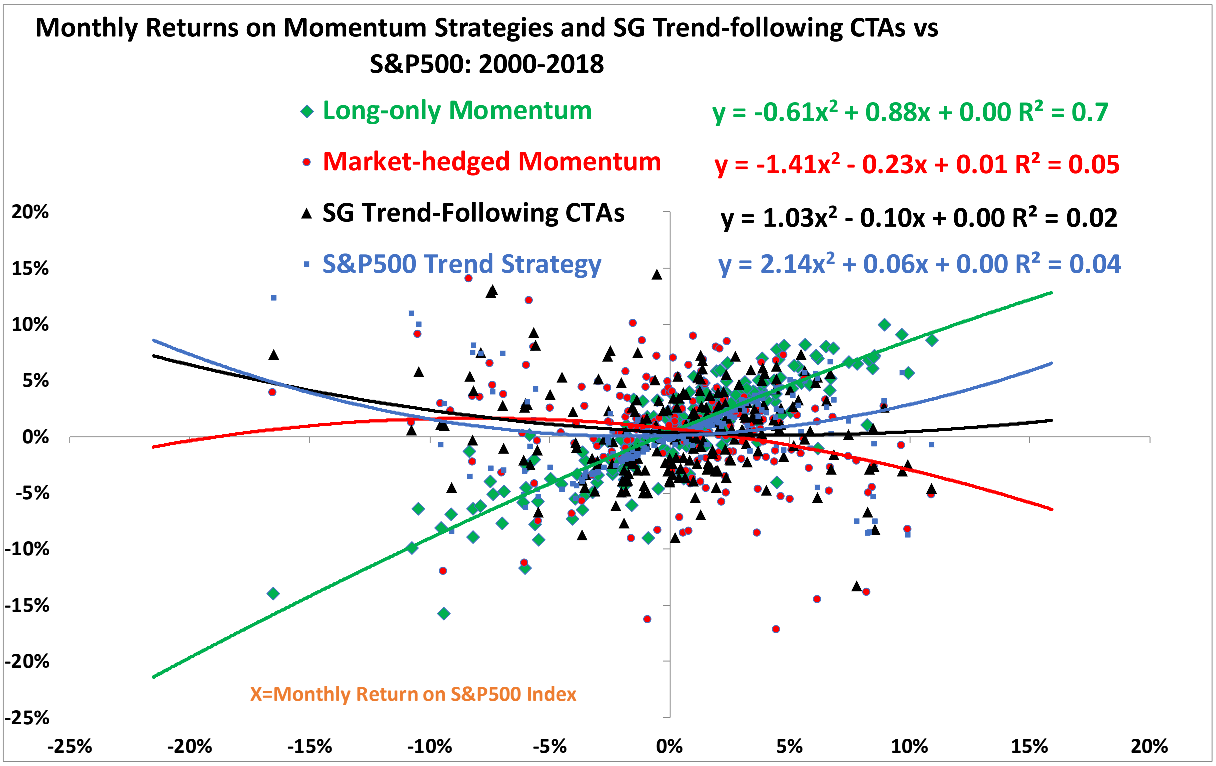

In table below I show the correlation matrix of monthly returns. Market-hedged momentum and Trend-following strategies are insignificantly correlated to the S&P 500 index. Market-hedged momentum is modestly correlated to the momentum index with 33% correlation and to the SG CTAs and S&P 500 trend-following strategy with correlations of 28% and 20% respectively.

Correlation matrix using monthly returns: 2000-2018ytd					
	S&P 500	Long-only M	Market-hedged M	SG Trend-Following	S&P500 Trend Stra
S&P 500 TR Index	1.00				
Long-only Momentum	0.86	1.00			
Market-hedged Momentum	-0.20	0.33	1.00		
SG Trend-Following CTAs	-0.11	0.04	0.28	1.00	
S&P500 Trend Strategy	0.05	0.15	0.20	0.42	1.00

It appears that the stock momentum is somewhat correlated to the trend-following strategies. To examine the dependency further I implement the principal component analysis (PCA) as follows.

1. I apply the PCA on correlation matrix of monthly returns.
2. I compute the returns on eigen portfolios which are orthogonal portfolios that obtained from the PCA decomposition.
3. I apply the Lasso regression to find the significant betas of strategy returns to the returns on eigen portfolios. The obtained regression explains about 90% of total variation for each strategy using their exposures to only one or two eigen portfolios.

I reach the following conclusions.

- The betas to the first eigen portfolio can be interpreted as the overall exposures to the stock market. Returns on long-only stock momentum are fully explained by the first eigen portfolio with the highest beta of 0.8, while the market-hedged momentum and trend-following strategies have a very small exposure to the first eigen portfolio and thus to the overall market direction.
- The second eigen portfolio appears to be specific to the joint component of both the trend and the momentum. The second eigen portfolio can be interpreted then as the exposure to the auto-correlation, which is the key source of returns on these strategies. The long-only momentum strategy itself has a zero exposure to the auto-correlation factor because the strategy returns are dominated by the first-order exposure to the market direction.
- The third and fourth eigen portfolios appear to be specific to the market hedged momentum and trend-following CTAs, respectively.

As a result, there is indeed some dependence between the market-hedged stock momentum and the trend-following CTAs with trend-following CTAs having stronger exposure to the auto-correlation factor. Market-hedged momentum has a stronger exposure to the third order eigen portfolio and thus it loses the convexity property of the trend-following CTAs.

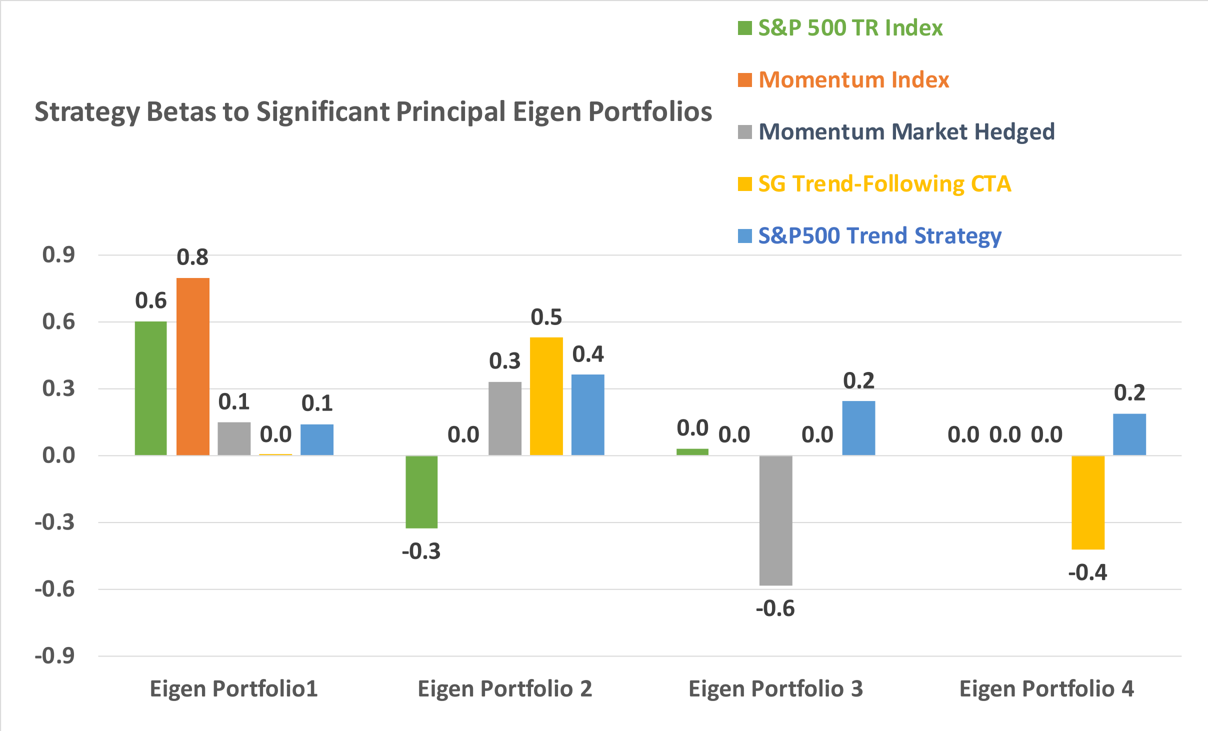

Benefits of Trend-following CTAs for Allocations in Alternatives

Because trend-following CTAs have positive convexity, they serve as robust solutions for diversification and reduction of tail risks.

For this analysis, I will apply the following strategies.

1. HFR Risk Parity Vol 10 Institutional Index (HFRPV10I Index) is the benchmark hedge fund index which includes funds with over $500MM AUM and which target the volatility of 10%. This category of funds includes relative value, carry, long-short, etc.
2. SG Trend-Following CTAs index is the flagship index of 10 trend-following CTAs.
3. **50/50 Portfolio** with 50% allocations to the HFR Risk Parity and SG Trend-Following CTAs with annual rebalancing.

In the figure below, I show the risk profile of these strategies using quarterly returns since March 2003 (the inception date of the HFR index) to April 2018. Clearly, funds with volatility targeting produce negative convexity because they typically need to de-leverage in downside markets. While HFR Risk Parity Index generates significant alpha, the index has a significant first-order exposure to the performance of the S&P 500 index. SG Trend-following CTA index produces a significant positive convexity using quarterly returns. As a result, trend-followers can generate positive performance in downside markets.

Importantly, when we use equal allocation to the Risk Parity strategies with Trend-following CTAs, the portfolio has insignificant convexity profile yet generating significant alpha.

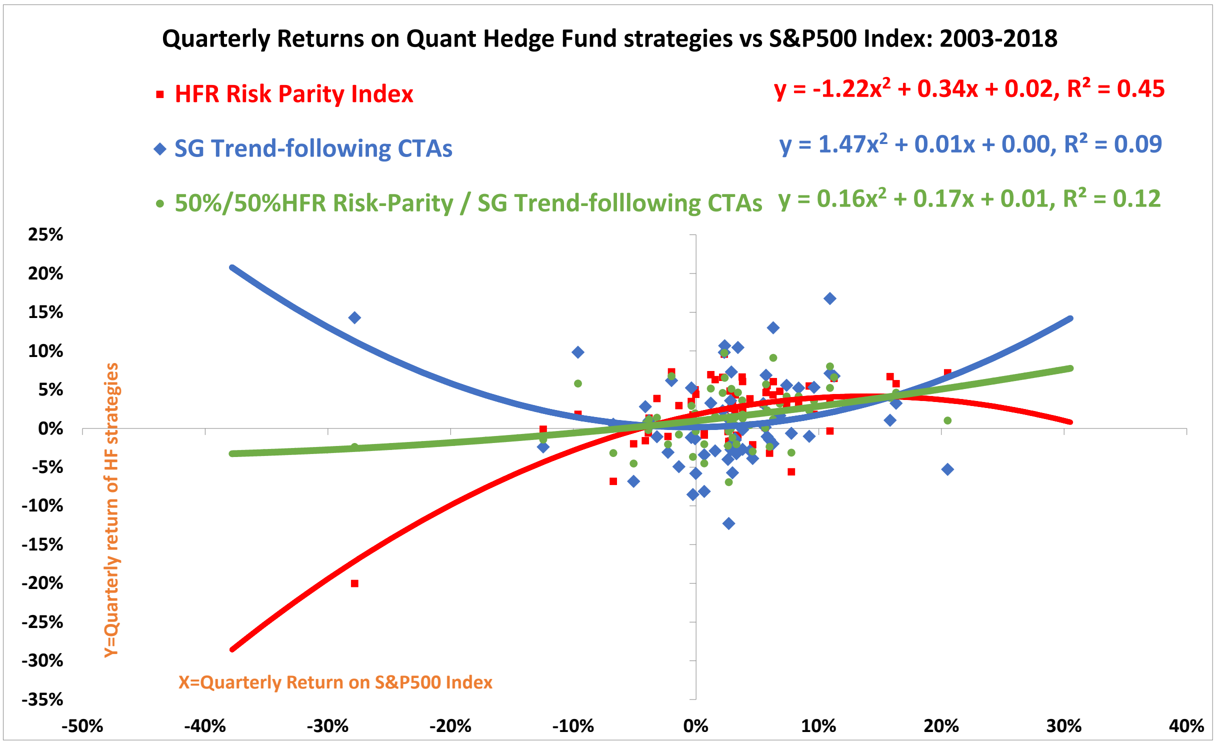

In the figure below, I illustrate the convexity profile of the strategies in comparison to the S&P 500 total return index with the risk measured by the maximum standard deviation and the reward measured by the realized Sharpe ratio. It is illuminating that while both the Risk parity and Trend-following CTAs have the maximum drawdown of about -22%, the 50/50 portfolio has the maximum drawdown of -11% which is the perfect reduction in half because the occurrence of the drawdowns of Risk-parity HFs and Trend-following CTAs are independent (in this sample). While CTAs trend-following have a lower Sharpe than Risk-parity funds, trend-followers serve as robust diversifiers with the 50/50 portfolio having nearly the same Sharpe but with twice smaller drawdown risk.

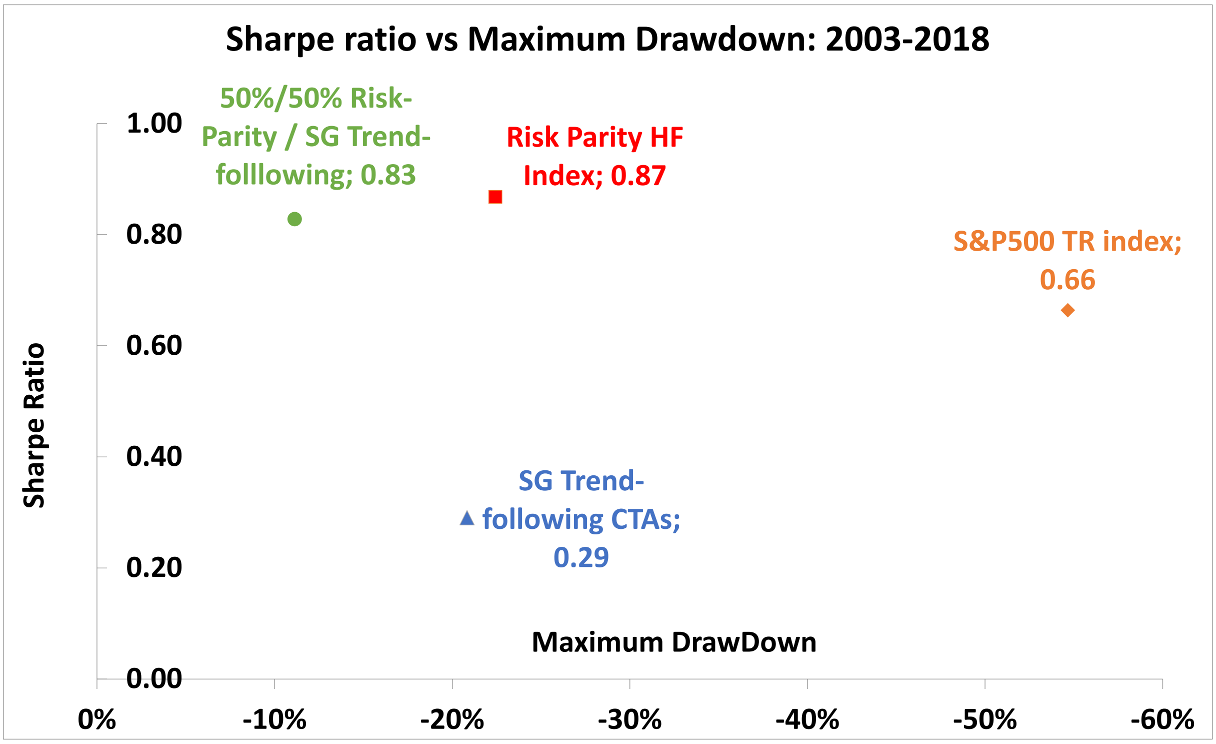

References

- Artur Sepp (2018) "Machine Learning for Volatility Trading", https://youtu.be/CwiSvzyEyMY
- Artur Sepp (2017) "Diversifying Cyclicality Risk of Quantitative Investment Strategies", https://artursepp.com/2017/12/01/diversifying-cyclicality-risk-of-quantitative-investment-strategies-presentation-slides-and-webinar-qa/
- Artur Sepp (2017) "Allocation to systematic volatility strategies using VIX futures, S&P 500 index puts, and delta-hedged long-short strategies", https://artursepp.com/2017/09/20/allocation-to-systematic-volatility-strategies-using-vix-futures-sp-500-index-puts-and-delta-hedged-long-short-strategies/
- Artur Sepp (2016) "Volatility Modelling and Trading" https://ssrn.com/abstract=2810768
- Tung-Lam Dao, Trung-Tu Nguyen, Cyril Deremble, Yves Lemperiere, Jean-Philippe Bouchaud and Marc Potters (2017) "Tail protection for long investors: trend convexity at work", Journal of Investment Strategies 7(1), pages 61-84, https://arxiv.org/pdf/1607.02410.pdf