Replicating CTA Positioning:
An Improved Method

Lars N. Kestner
July 2020

Summary

Analysis of systematic strategies is a current topic of focus, centering on the impact these strategies have on various financial markets. Risk parity, option overwriting, volatility targeted equity indices, and trend following strategies receive the majority of this attention. In this paper, we focus on the dynamic trading of trend following strategies and detail an improved method for estimating their actions across markets.

» A simple replication model employed on 16 futures markets explains over 75% of the variation in a trend following benchmark.

» This replication model is able to estimate trend follower positions without lag.

» Using estimates of total funds allocated to trend following managers, we can use our replication model to estimate positions by specific market and the expected trading flows when individual markets move.

STATISTICALLY APPLIED TRADING
www.satquant.com
New York

Introduction

Analysis of systematic strategies is a current topic of focus, centering on the impact these strategies have on various financial markets. Risk parity, option overwriting, volatility targeted equity indices, and trend following strategies gain the majority of this attention. In this paper, we focus on the dynamic trading of trend following strategies and detail an improved method for estimating their actions across markets.

Trend following strategies gained interest in recent decades as a non-correlated diversifier to more economically-sensitive asset classes. Performance of trend following CTAs during the financial crisis of 2008 was strong, resulting in growth of assets under management. As assets have grown, the trading sizes of these managers has become larger relative to the total market. This has led to market analysts attempting to forecast the aggregate size of positions and also the potential changes in their positions. We believe most of these estimates are inaccurate, largely because of incorrect estimation techniques used. Analysts frequently force fit CTA return data using regression techniques with a rolling window of returns for calculations. This paper will describe how to improve estimates and reduce lag for these estimated exposures using replication techniques.

Performance of trend followers was quite strong during 2007 and 2008.

Source: Bloomberg and Saτ calculations

Performance Attribution

Performance explanation and attribution is an important task as investors attempt to understand an asset or portfolio's behavior. For a transparent investment, such as an index fund linked to the S&P 500, that attribution is simple and can be calculated using the weights and returns of the underlying components. But for private investment vehicles such as hedge funds, the sparse reporting frequency and lack of detail around their investment process makes this attribution more difficult. Analysts usually resort to factor-based techniques where portfolio returns are regressed on known investment returns from broad asset benchmark returns and intra-asset class strategies, using statistical techniques to minimize fitting errors. This analysis is ripe for misattribution as portfolios can change instantaneously while the estimation techniques need additional data to determine these changes. As such, exposure estimates are likely to lag the true change in portfolio strategy. Without a better understanding of the investment process, the best an analyst can accomplish is to minimize this lag.

Certain parts of the non-traditional investment world have been studied in detail and their investment process is generally understood. With these types of investment managers, their risk exposure can be estimated more accurately using methods other than regression-based techniques. Trend following CTAs are one group of managers that fit this category. Trend followers (often given the designation of CTAs or managed futures traders due to the investment products they trade) generally buy on market strength and short on market weakness, with algorithms determining when to trade. Much has been learned about how these managers take risks since academics began studying their returns. Because we now generally understand the investment process of trend followers, we can estimate their positions more accurately in real-time and without lag.

Building the Model

We present a replication model that explains the majority of variability in a trend following CTA benchmark. We begin by calculating the average weekly return over a set lookback window. Trading positions will be larger when the momentum intensity is higher. Next, we normalize average returns by a measure of realized return volatility and multiply by the square root of the number of weeks in the lookback window. This result is the relative position size for a market on any given week. The volatility normalization creates a like-for-like ability to compare returns of low volatility markets such as government bonds with higher volatility markets such as crude oil, while the window scaling normalizes the results of varying the lookbacks. We also cap these exposures at defined thresholds to avoid situations where most of the portfolio's variance is created in just a few outlier market moves. Each market is given an equal weight once adjusted for volatility and no adjustment is made for correlations among markets (i.e. naive risk parity). We scale the leverage of our model to match the volatility of our benchmark using in-sample results of the replication portfolio. For our benchmark we use the *SG Trend Index*, a sub-index of the *SG CTA Index* that aggregates the performance of managers trading trend following strategies.

We limit our trading to 16 of the most liquid futures markets originating from four market classes. The model will trade equity, interest rate, foreign exchange, and commodity futures. The number of markets and sectors in our replication model is much smaller than a typical diversified trend follower would trade.

The replication model trades 16 futures markets split across four asset classes.

Equities	ES - S&P 500 Emini	Z - FTSE 100	TY - US 10 year	G - Long gilts	Rates
	VG - Euro Stoxx 50	NK - Nikkei 225	RX - Euro Bund	JB - Japan 10 year	
	EC - Euro	BP - British Piound	CL - Crude Oil	GC - Gold	
Currencies	JY - Japanese Yen	AD - Aussie Dollar	HG - Copper	S - Soybeans	Commodities

Source: Saт calculations

We investigate the robustness of our replication model by varying three parameters: the momentum lookback, the point where we cap the maximum signal (the normalized momentum cap), and the lookback for the realized volatility calculation. We measure the robustness of replication by calculating the coefficient of determination (R^2) between our replication strategy's returns and that of the benchmark. The goal is to maximize the replication to the benchmark, not the return performance of our model.

We use 5 possibilities for each of the three parameters, leading to 125 total tests. All calculations are performed weekly on the last active trading day. Trades are made using closing prices and do not take into account explicit or implicit trading costs.

We study 125 total parameters: 5 for the momentum lookback, 5 for the normalized momentum cap, and 5 for the volatility lookback.

Parameter list
Momentum lookback: 4, 8, 16, 32, and 52 weeks
Normalized Momentum Cap: 0.0, 0.5, 1.0, 1.5, and 2.0 normalized momentum scores
Volatility lookback: 20, 30,60, 90, and 180 days

Source: Saт calculations

Results are represented by a combination of parameter values. For example, "8/1/90" refers to the parameter selection of 8 weeks for the momentum lookback, capping the normalized momentum at 1.0, and using 90 days for the window to calculate realized volatility.

The chart below graphs all 125 parameter combinations sorted from lowest R^2 to highest R^2. The two "jumps" in the graph can be described generally as where the momentum lookback shifts from 4 weeks to 8 weeks and then from 8 weeks to 16 weeks.

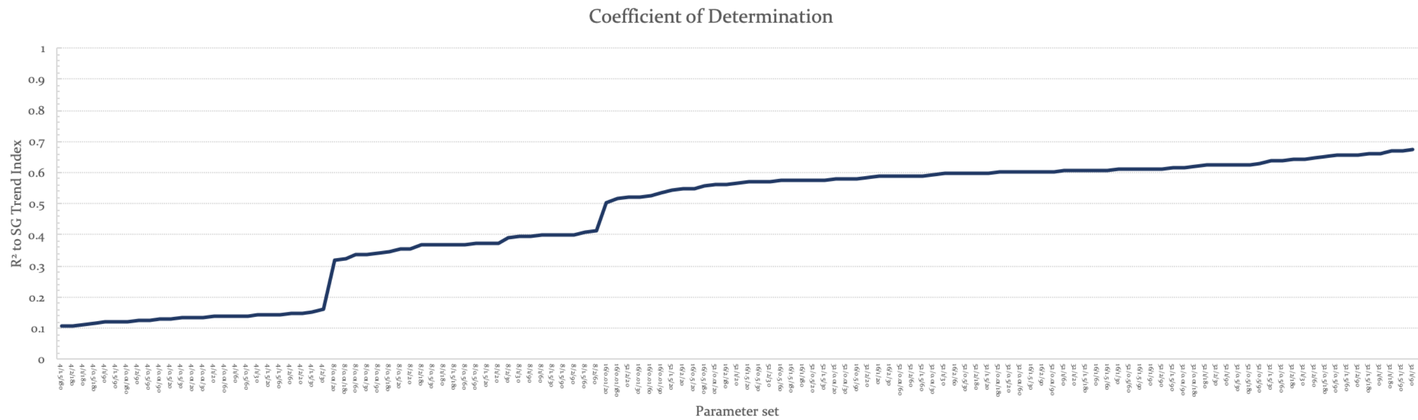

Source: Bloomberg and Saт calculations

In the next three graphs we investigate the replication stability by holding two parameters at constant values while varying the other parameter across values.

First, we isolate the effect of the momentum lookback. Below are the R^2 results for the 5 momentum lookbacks (4 weeks, 8 weeks, 16 weeks, 32 weeks, and 52 weeks) while holding the normalized momentum cap at 1.0 and the volatility lookback constant at 90 days. Note the increase in R^2 as the momentum lookback increases from a 4 week lookback to 8 weeks then to 16 weeks, plateauing with the addition of 32 weeks and 52 weeks.

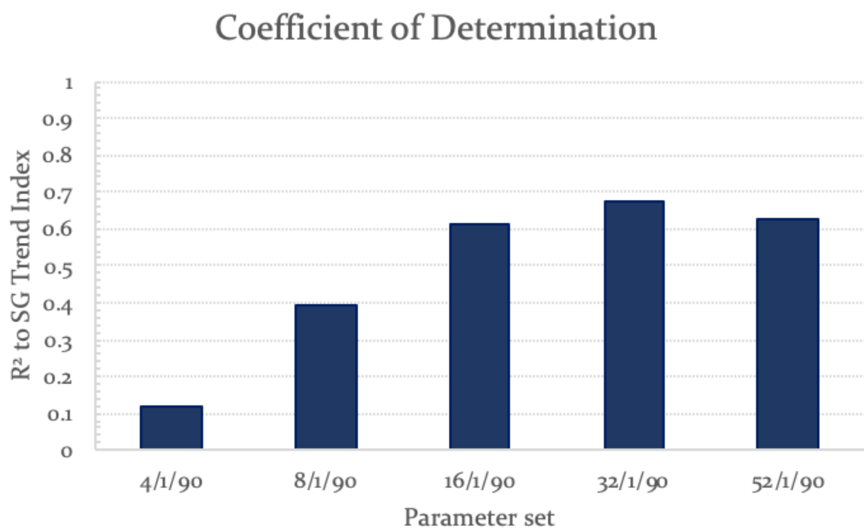

Source: Bloomberg and Saɪ calculations

We graph R^2 for varying momentum lookbacks while holding the normalized momentum cap at 1.0 and the volatility lookback at 90 days.

Next we vary the normalized momentum cap while holding the momentum lookback at 32 weeks and the volatility lookback at 90 days. Unlike the variability we saw in the momentum lookback, there is no such result when the normalized momentum cap moves through the parameter points. Replication does appear to peak near a cap value of 1.0.

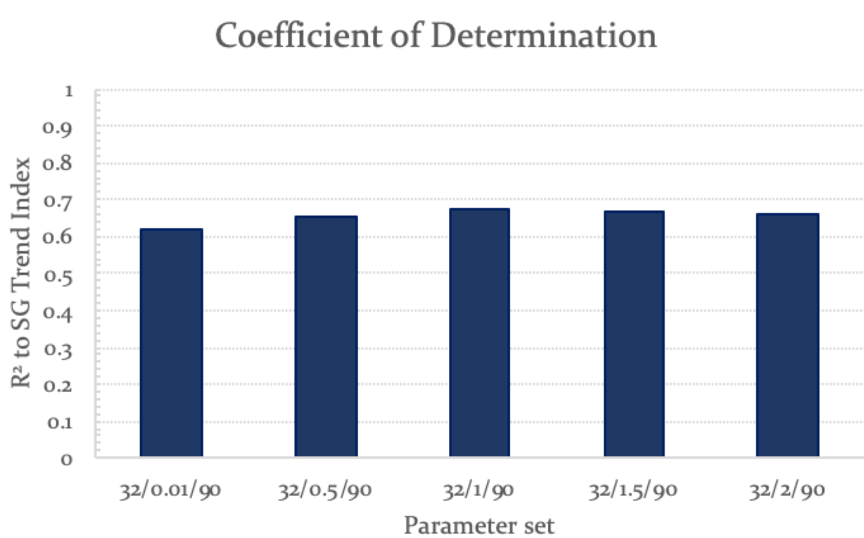

Source: Bloomberg and Saɪ calculations

We graph R^2 for varying normalized momentum caps, while holding the momentum lookback at 32 weeks and the volatility lookback at 90 days.

Lastly, we hold the momentum lookback at 32 weeks and the normalized momentum cap at 1.0 but vary the volatility lookback. We see little difference when varying the volatility lookback, but replication appears to peak near 90 days.

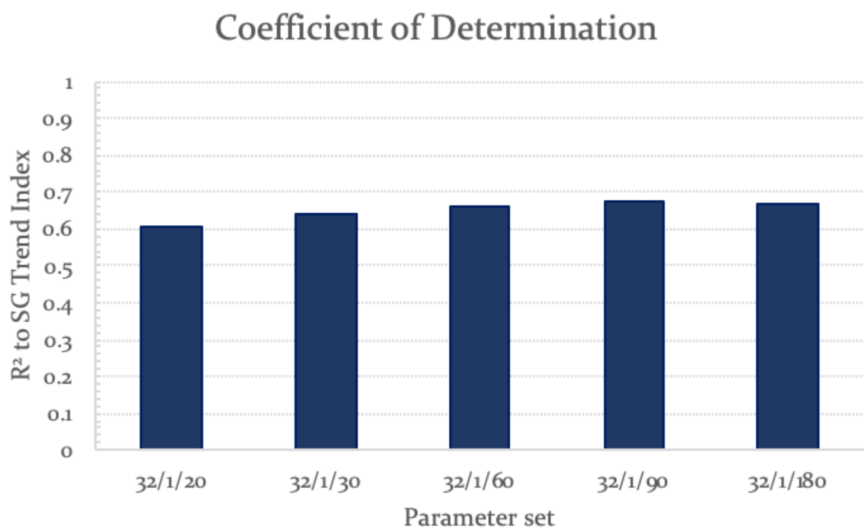

Source: Bloomberg and Sat calculations

We graph R^2 for varying volatility lookbacks while holding the momentum lookback at 32 weeks and the normalized momentum cap at 1.0.

We find that the benchmark is best approximated by momentum lookbacks around 32 weeks, normalized momentum caps around 1.0, and volatility lookbacks of 90 days. Changing momentum lookbacks has the largest sensitivity to replication success with the choice of normalized momentum caps and volatility lookback less sensitive.

Model Replication Results

Below is a time series graph of three specific parameter sets compared to the SG Trend Index from 2015-2019. Momentum lookbacks of 16 weeks, 32 weeks, and 52 weeks were selected, each with a normalized momentum cap of 1.0 and a volatility lookback of 90 days. The shape of our models generally follows the benchmark. This is not surprising as each model yielded R^2 results in the mid 60%s.

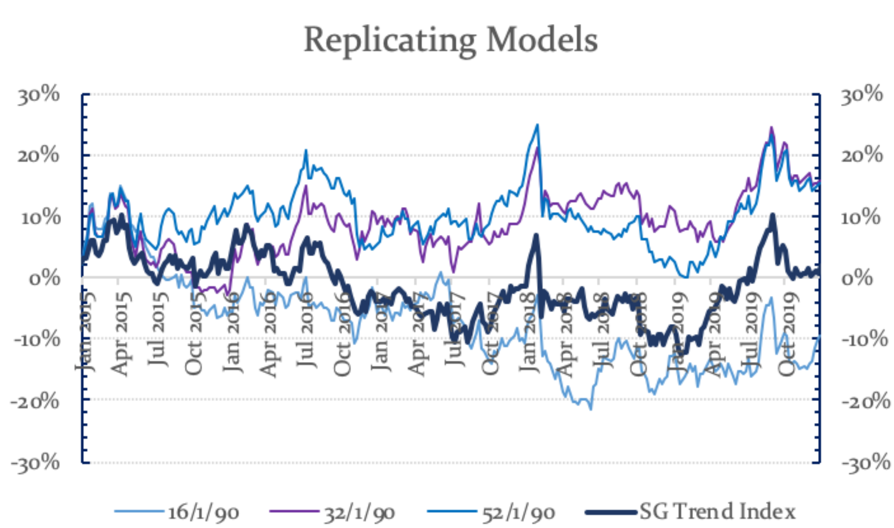

Source: Bloomberg and Sat calculations

Three potential replication models are compared with the SG Trend Index. Momentum lookbacks vary across 16 weeks, 32 weeks, and 52 weeks while holding the normalized momentum cap at 1.0 and the volatility lookback at 90 days.

Building on the strong correlation to the benchmark, we use an ensemble method to improve fit further. While the returns of trend following CTAs are correlated, the strategies these managers utilize will vary across trend speeds. Combining a range of parameters in our replication model improves our results by better matching how the actual managers trade – across trend speeds.

When we average the performance of our 3 models (momentum lookbacks of 16, 32, and 52 weeks; normalized momentum cap of 1.0; volatility lookback of 90 days) the R^2 to the benchmark climbs above 75%. Expressed graphically, we see the tight relationship between the replication model and the benchmark on both time series and return scatter graphs. This ensemble of three parameter sets will be our model for replicating returns of the benchmark.

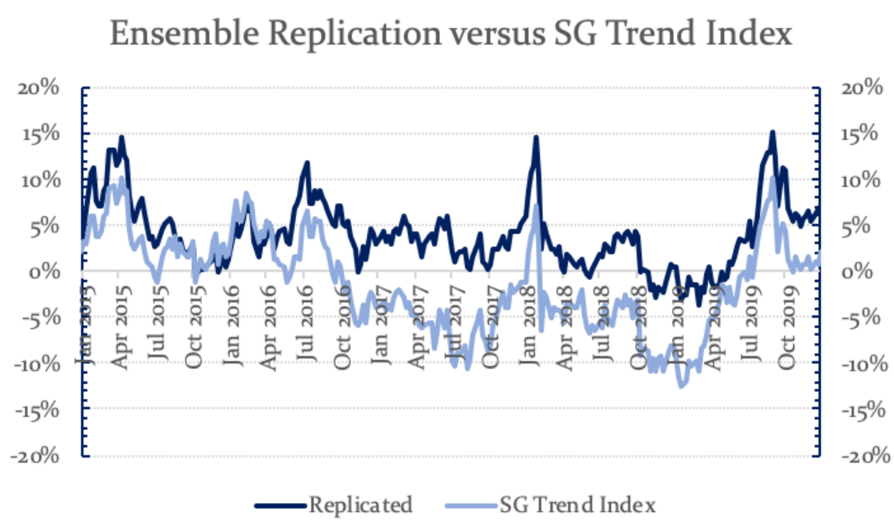

An ensemble approach of averaging results across 16 week, 32 week, and 52 week momentum lookbacks improves the replication. Both time series and scatter graphs display the successful replication.

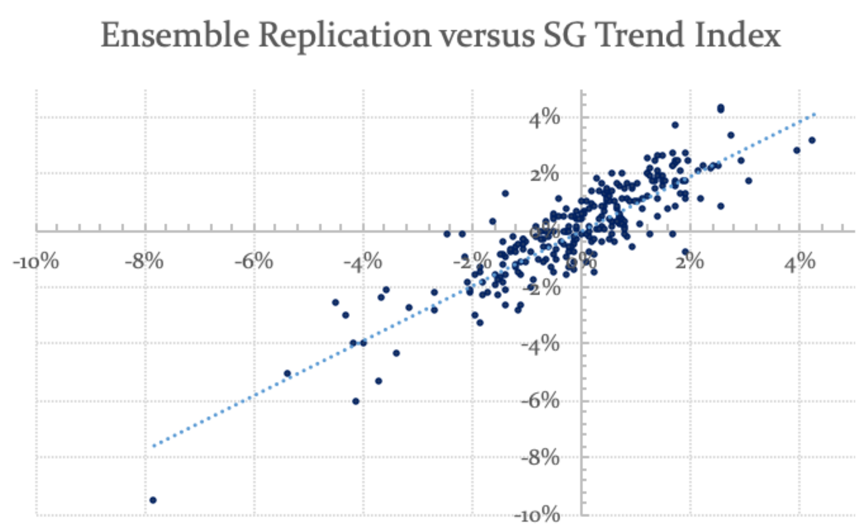

Source: Bloomberg and Saτ calculations

Trend Followers' Estimated Exposures

Now that we have identified a model to replicate trend following CTA returns, we can estimate the risk exposures that trend following managers have to each of our 16 markets.

Below is the percentage exposure per $1 notional allocated to the benchmark aggregated by sector. That is, a +100% exposure to equities would approximate a $1 long position per dollar invested in the benchmark (or managers comprising the benchmark).

We calculate the notional exposures over time by sector.

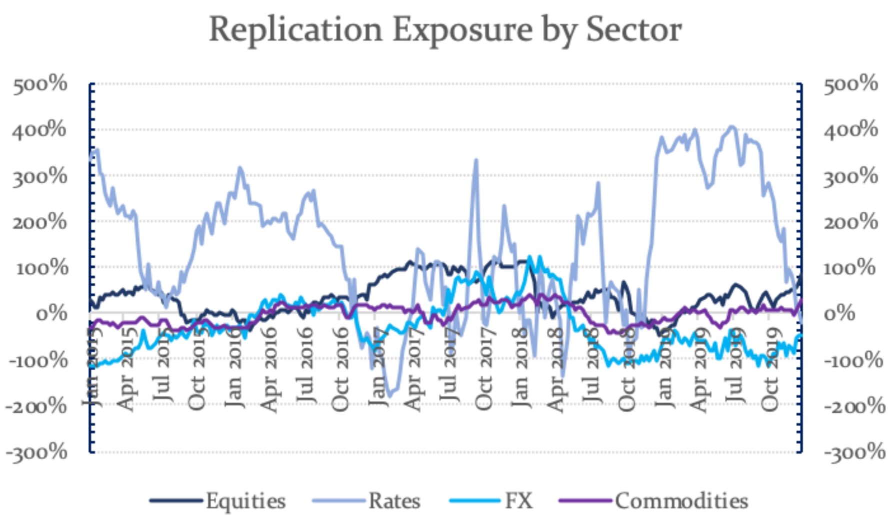

Source: Bloomberg and Sat calculations

The Wall Street Journal estimated that there is $300BB in managed money allocated to trend followers. If we assume this $300BB under management mimics the SG Trend Index, we can utilize the output of our replication model and estimate the position sizes by market traded. The analysis is performed in real-time, is not dependent on delayed published data, and will capture changes in manager positions immediately so long as our replication model continues to accurately mimic the benchmark.

Using the $300BB estimate and calculating positions from our replication model, the chart below estimates the total dollar size in longs or shorts to each of the four equity markets in our portfolio over the past 5 years. In a future note we will take an additional step to determine how much trend followers need to buy or sell based on tomorrow's move using tomorrow's price as an input to our replication model.

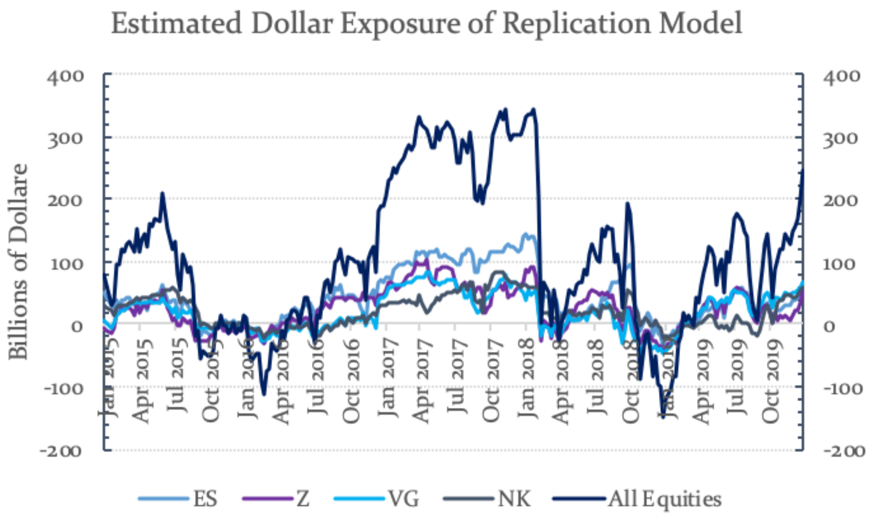

Source: Bloomberg and Saт calculations

We calculate the dollar notional of positions from our replication model across equities markets assuming a total $300BB allocation to trend following strategies.

We noted that replication would be superior compared with factor-based approaches to determine trend follower positioning. Next, we compare results of the two methods side-by-side. The graph below displays our replication model's results for percentage notional exposure to the S&P 500 and all equity markets combined. We also graph betas for a univariate regression where the SG Trend Index is our dependent variable and returns on the S&P 500 are our independent variable. We use a 20 week rolling window to estimate regression betas.

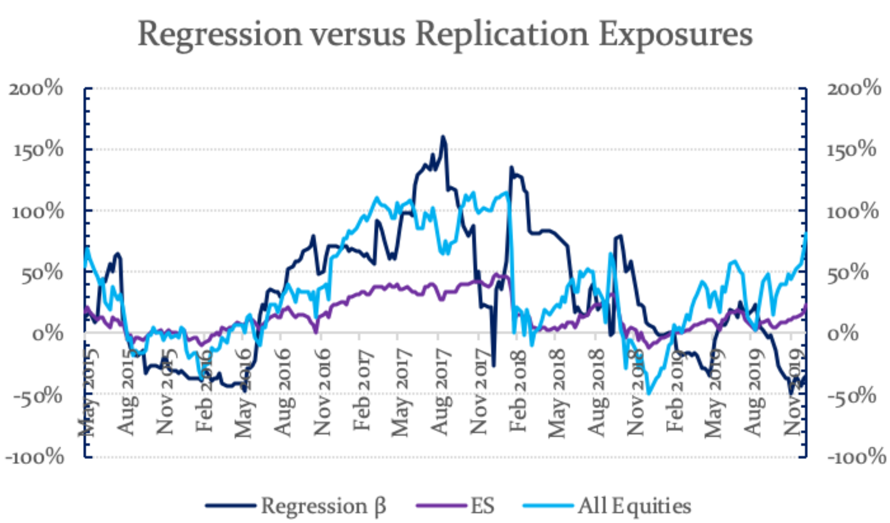

Source: Bloomberg and Saт calculations

This graph compares the estimated trend following positions in S&P 500 futures using two methods: 1.) a linear regression beta resulting from regressing returns of the SG Trend Index on returns of the S&P 500 and 2.) positions from our replication model.

Note the smoothness of the replication position estimates compared to the noisy output of the regression. At turning points, the replicating portfolio positions will adjust quickly to changes in market trends while the forced fits from the regression take time to catch up with significant lag.

Conclusion

We have created a simple replication model to mimic the performance of trend following managers. This model is able to explain more than 75% of the weekly variability in a well-known trend following benchmark, the SG Trend Index. The results of the replication model suggest it is a very effective technique to estimate the position sizes of trend following managers.

Using this replication model, market analysts can better estimate the amount of futures these trend followers have positioned and also the size that needs to be traded for dynamic rebalancing based on day to day or week to week changes in trend.