

Trend-Following: What's Luck Got To Do With It?

Nastja Bethke¹, Ed Tricker²

Abstract

Trend-following strategies experienced a wide range of performance outcomes during February and March 2020. In this paper, we illustrate how the precise sequence of events that occurred led to some strategies performing unexpectedly poorly, while others performed unexpectedly well. Our results suggest that should we see a similar situation replay again, where markets fall the same amount in the same period, but along a slightly different path, we would generally expect to see broadly positive performance and a narrower range of outcomes.

¹ Senior Quantitative Research Analyst

² CIO - Quantitative Strategies

1. Introduction

The global health crisis of February and March 2020 caused financial markets to experience drastic price movements, in many cases at speed unprecedented in history. Volatile markets can sometimes present a good opportunity for systematic strategies such as trend-following. However, if volatility is extreme, it can also present challenges. This note investigates the actual and simulated performance of two popular types of trend-following models during this crisis, to shed some light on what behavior can be expected, or what behavior is potentially quite rare.

2. A Motivating Example

Throughout February and March, we observed a much broader dispersion of performance across trend-following models than one might typically expect. Certain types of models or certain specific parameterizations of models appeared to yield better performance. But what can we infer from this result? Was it the product of skill, luck, or some combination of the two? Can one assume that the best performers this time would be the best performers if a similar event were to repeat?

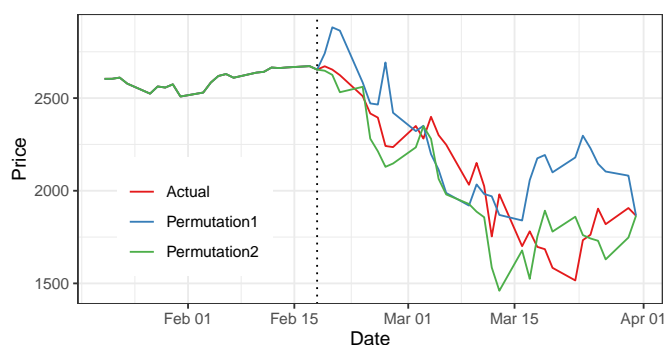
To investigate, we start with a simple experiment. We begin with the actual returns of the S&P 500 futures contract as the market declined from February into March. Using this data, we produce simulated price series by reordering the returns. In mathematics, a reordering of this type is referred to as a *permutation*. Each permutation yields a new price trajectory that experiences the same decline as the actual S&P 500 over the same period, but in a slightly different way since the returns are now in a different order. Two examples of such permutations are shown in Figure 1a¹.

Taking the simulated price trajectories, we apply a simple moving average crossover, or MaCo, trend-following strategy², the resulting P&L of which is shown in Figure 1b. It is striking that relatively similar price trajectories yield dramatically different P&L outcomes.

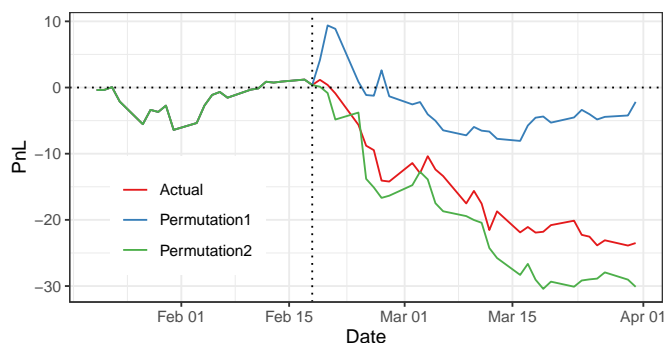
¹Note, one must be careful not to distort the autocorrelation when changing the order of returns. In this case, because the market move was so sustained, a permutation of returns does not materially alter the autocorrelation

²MaCo trend-followers compare fast and slow moving price averages to de-

This experiment, while simple, serves as motivation for a more in-depth investigation. It demonstrates that sometimes it is not just how much a market gains or loses, but also how it gets there. In this sense, trend-following is *path-dependent*. It suggests that small changes in market behavior can lead to large changes in performance - trend-following performance over this period was relatively *fragile*.



(a) Actual S&P 500 trajectory compared with two permutations



(b) Comparison of trend-following performance on different prices.

Figure 1. Examples of permuted prices and the results apply trend-following to them. Observe that similar price trajectories yield very different performance.

termine whether to go long or short a market and represent a well-understood generic form of trend-following. For more details, see the GCM paper *The Speed of Trend-Following (March 2018)*.

3. A More General Experiment

While the experiment in Figure 1 is helpful, it would not be appropriate to draw any conclusions from a handful of examples. To further investigate the performance of trend-following in February and March, we can estimate its statistical distribution.

To do this, we will make use of a series of Monte Carlo simulations that involve randomly sampling the data repeatedly to construct a large number of what-if scenarios for the S&P 500. In effect, we replicate our initial experiment 1000s of times, covering a range of price trajectories. Figure 2 shows a plot of these simulated price trajectories. Observe that while they all start and end in the same place, they take different routes.




Figure 2. Simulated price trajectories for the S&P 500. Note that the start and end points are the same, but the route taken differs.

As before, we can run MaCo trend-following on these simulated price trajectories and examine the performance. In Figure 3, we plot a histogram of the returns for all the simulations. As with our motivating example, even though in each case the market fell by the same amount in the same period, we observe a wide range of outcomes. The simple MaCo trend-following strategy generally struggles to produce positive returns in this market (not surprising given the speed and extreme nature of the reversal). Still, performance ranges from broadly flat to significant losses.

We can also use this distribution of performance to rank the actual observed results. That is, we can compare the results of trend-following on the precise price trajectory actually experienced by the S&P with performance on very similar simulated trajectories. Highlighted in Figure 3 by the red line, we see that actual performance was among the worst of all the 1000s of paths tested. Indeed, only 9% of our simulated trajectories yielded worse performance. By this metric, one can consider observed performance relatively ‘unlucky’ - since, in the overwhelming majority of cases, the same returns but in a slightly different order would have led to significantly better performance.

3.1 An Even More General Experiment

In the experiment above, we only considered one market. While a trend-following strategy can suffer unusually poor performance in one market, how about across a range of markets? To investigate, we repeated the experiment across dozens of assets, covering equity indices, commodities, FX, and fixed income.

Again, we plot a histogram summarizing performance across

the simulations (Figure 4). We again observe a range of outcomes, and again it is striking that the actual realized performance is still one of the worst of all possible outcomes. It is also worth noting that the mean return of the simulations is positive, while the actual realization was negative. Such a result is remarkable - one might reasonably expect that poor luck in some markets would be offset by good luck in others. The result demonstrates that the events of February and March happened in a specific way that led to a near worst-case performance for this type of trend-following strategy.




Figure 3. The distribution of outcomes from repeated trials of permuted data. Less than 9% of permutations lead to worse performance than actually observed.




Figure 4. The distribution of outcomes from repeated trials of permuted data applied to all markets. Less than 4% of permutations lead to worse performance than actually observed.

4. Good and Bad Luck

In the previous section, we detailed experiments that demonstrated that a MaCo trend-following strategy exhibited performance towards the extreme end of its expected range - in a sense it was unlucky. But what about other types of trend-following?

Another popular type of trend-following strategy is known as a breakout. As the name suggests, breakout models identify when the price breaks outside of a trading range, and then go long or short accordingly. Historically, breakout trend-following strategies have proven to sometimes be more nimble in volatile market conditions (although they are prone to whipsawing when

conditions are more benign). As a result, one might expect them to perform better in February and March.

To investigate, we repeated the previous analysis on this second type of trend-following strategy. As before, we generated 1000s of trajectories for all markets. The resulting range of returns is summarized in Figure 5.

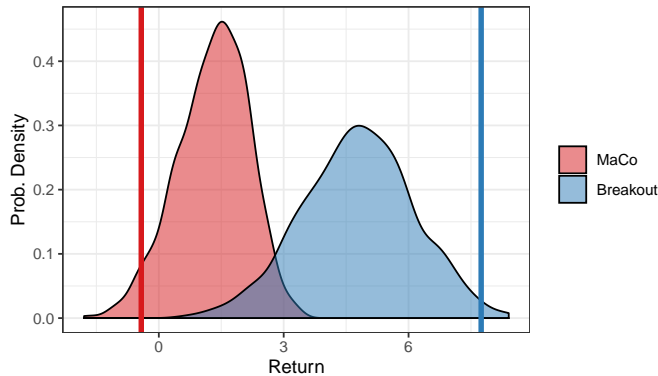


Figure 5. Comparison of simulated return distributions for MaCo and breakout trend-followers. Actual performance highlighted by vertical lines. We observe that realized performance was at opposite ends to the expected distribution.

Once again, we observe a range of outcomes. On average, as expected, the breakout models outperform the MaCo models that could not keep up with the speed of the market changes. However, the results also show that the actual realized performance of these models was almost as good as possible. Of all the permutations considered, the reality of February and March was almost a best-case scenario. Even small modifications to the price trajectories lead to significantly worse performance. To be clear, breakout models generally outperformed MaCo models - however, the magnitude of the realized performance difference is a result of MaCo performing towards the lower end of its range, while breakout models performed towards the upper end. An interesting example of both bad and good luck during identical market conditions.

We can also compare different variants of the same model. For example, in Figure 6 we compare the performance distribution of two parameterizations of the same trend-following strategy. We start with the same (30,90) MaCo model used previously, and contrast results with a faster (20,60) variant. This faster model uses a significantly shorter lookback.

The results of the simulations suggest that, *on average*, in the events of February-March, the faster model could be expected to perform slightly better. However, for the specific sequence of events that we observed, the (30,90) MaCo model suffers almost worst-case performance, while the faster model realized much closer to its average. Once again, we observe a large difference in actual performance, which is unlikely to be repeated if a similar market move occurs in the future.




Figure 6. Comparison of simulated return distributions for MaCo trend-following models of different speeds. Actual performance highlighted by vertical lines. We observe that while the faster MaCo20 model performed significantly better in the actual events of February-March than the slower MaCo30 model, on average we would expect performance of the two models to be much closer.

5. Conclusion

Using Monte Carlo simulations, we have shown that we can view the poor and good performance of representative MaCo and breakout trend-followers during the crisis of February and March 2020 as somewhat rare events. Moving average crossover strategies appear to have been unlucky, in the sense that most other sequences of market returns leading to the same overall price drop would have resulted in better performance. In contrast, breakout models have been lucky: few other parallel universes experiencing the same overall downturn would have seen such good returns. While this may be a frustrating conclusion for those who employ MaCo models versus breakouts, it suggests that a similar future market scenario will likely be kinder to the former. Pivoting a trend-following strategy to only trade breakout models, or significantly changing trend-following speed, may lead to disappointing results next time.

Legal Disclaimer

THIS DOCUMENT IS NOT A PRIVATE OFFERING MEMORANDUM AND DOES NOT CONSTITUTE AN OFFER TO SELL, NOR IS IT A SOLICITATION OF AN OFFER TO BUY, ANY SECURITY. THE VIEWS EXPRESSED HEREIN ARE EXCLUSIVELY THOSE OF THE AUTHORS AND DO NOT NECESSARILY REPRESENT THE VIEWS OF GRAHAM CAPITAL MANAGEMENT. THE INFORMATION CONTAINED HEREIN IS NOT INTENDED TO PROVIDE ACCOUNTING, LEGAL, OR TAX ADVICE AND SHOULD NOT BE RELIED ON FOR INVESTMENT DECISION MAKING.