

WHITE PAPER – DECEMBER 2023
FOR INVESTMENT PROFESSIONALS ONLY
NOT FOR PUBLIC DISTRIBUTION

Selecting a Trend-Following CTA

Executive Summary

We outlined in our previous paper, [The Case for a Strategic CTA Allocation](#), the advantages of trend-following strategies in the current environment.

In this paper, we continue the analysis with the aim of helping allocators at the manager selection step of the investment process. To do this, we examine the four most important strategy design considerations:

- 1) The universe of assets traded.
- 2) The trend signal.
- 3) The inclusion of diversifying signals.
- 4) Risk management.

Understanding the role and interaction of these four considerations is vital for identifying trend-following strategies that are a good fit for an investor's portfolio and objectives.

Introduction

Our last paper, [The Case for a Strategic CTA Allocation](#), explored the benefits to investors from allocating to a trend-following CTA strategy. Almost uniquely among liquid hedge fund strategies, trend following provides highly valuable diversification to a portfolio of equities and bonds, and thanks to its cash efficiency, the strategy is particularly suited to an environment with higher interest rates.

Our aim for this paper is to equip allocators with the information required to distinguish between trend-following CTAs and select a strategy aligned with their investment objectives.

What is Important for Trend Following?

Trend-following strategies, in the simplest terms, take long positions in assets that have increased in price and short positions in assets that have declined in price. While the idea is straightforward, the practical application of such a strategy involves four key elements:

- 1. Universe:** Which assets are traded? (page 4)
- 2. Trend signal:** The algorithm to determine the direction and strength of the trend in each asset. (page 8)
- 3. Diversifying signals:** Should complementary signals be included? (page 17)
- 4. Risk management:** The framework for understanding and managing the strategy's risk. (page 20)

We discuss these points in detail over the remainder of this paper.

Why Manager Selection Matters

Managers vary in how they approach the four key considerations listed above and, as a result, the performance they achieve varies widely over the long term as well as the short.

This variation is illustrated by Figure 1, where we have compared the performance of the SG Trend Index with that of its current and past constituents. The SG Trend Index is designed to track the returns of the largest trend-following CTAs, is reconstituted annually, and has data back to 2000.¹


As managers offer their strategies at different levels of volatility, we have normalised the returns to a 10% annualised volatility to put the strategies on the same footing. Despite this adjustment, the net performance benefits of selecting the right manager are considerable, with the normalised returns over the full period ranging from 3.1% to 7.7% per annum. At the bottom end of the spectrum, some members of the index produced disappointing returns and shut down their strategies altogether.²

¹ The index's methodology, along with details of current and past constituents, is available on Société Générale's [Prime Services Indices](#) website.

² High-profile examples include strategies operated by John W. Henry & Co (Wall Street Journal, [Henry to Exit Money Game](#), 13 Nov 2012), Tudor Investment Corp (Bloomberg, [Tudor to Return Money From Managed-Futures Fund Amid Loss](#), 25 Mar 2014) and GSA (Financial Times, [Hedge fund GSA Capital to become private trading firm](#), 19 Nov 2021).

Figure 1: Return dispersion between CTA strategies makes manager selection critical

SG Trend Index and underlying constituents since 2000, adjusted to 10% annualised volatility



Source: Winton, Société Générale, Hedge Fund Research, Morningstar, as at 30 November 2023. **Past performance is not indicative of future performance.** The analysis includes net returns reported for all current index constituents and all past constituents are included for as long they reported returns. Returns have been adjusted to a 10% annualised volatility. Strategies that were launched after January 2000 have been rebased to the value of the SG Trend Index at the time of the first reported return.

The Value of Experience

Perhaps because the basic idea of trend following is beguilingly simple, new managers enter the space every so often. Often the implementation is based on a simple analysis of past industry returns, without much attention paid to the four key considerations listed above.³ Consequently, these strategies have tended to underperform the wider industry over the long term, leading to a change of approach or to the manager shutting down their strategy altogether.

³ For example: Brian Hurst, Yao Hua Ooi and Lasse Heje Pedersen, *Demystifying Managed Futures*, *Journal of Investment Management*, 2013.

1. The Universe of Assets: Maximising Diversification

Background on Futures Markets

One of the remarkable aspects of futures markets is that no matter how varied the underlying asset (gold, bonds, oil, stock indices, etc.), trading a future on one asset is like trading a future on any other.

This feature means that a firm can manage systematic strategies across a diverse collection of futures, without needing to employ specialists for each market. The resulting diversification is one of the great virtues of CTA strategies.

Starting in the 1970s, CTAs focused initially on the established commodity futures trading on exchanges in Chicago, New York and London. Over time the universe broadened: first, through the rapid development of financial futures in the 1980s; and second, with the launch of regional futures contracts on exchanges around the world during the 1990s and early 2000s.

The main developments since the 2008 global financial crisis have been: 1) an increase in the variety of liquid instruments available to trend followers – including interest rate swaps, credit default swaps OTC-traded futures and cash equity baskets; and 2) the opening up of China’s vibrant futures markets to global investors, through internationalised futures, the qualified foreign investor programmes, and offshore swaps.

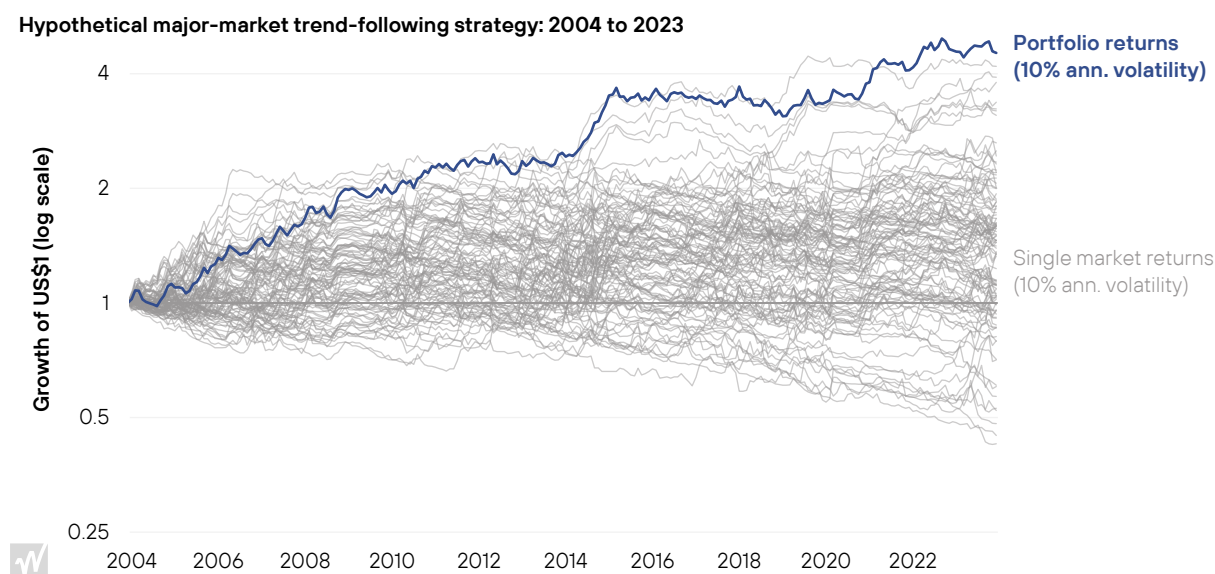
It is common now for CTAs and investors to distinguish “major markets” – the large and liquid futures markets and currencies traditionally traded by CTAs – from “alternative markets”, which span OTC markets, emerging market futures, and small or non-standard futures.

These developments are positive for trend following as market diversification is a critical part of the strategy. Consider, for example, a major-market trend-following strategy. Applied to each of its 100 underlying markets separately, the strategy only shows a faint edge: a Sharpe ratio of 0.2 on average over the past 20 years. However, due to the low mean correlation of 0.07 between the underlying returns, the overall result when combined in a portfolio is a 0.8 Sharpe ratio.

This diversification effect is visible in Figure 2, which compares the strategy returns in each market with the combined portfolio.⁴

⁴ The interaction between market-level Sharpe ratios, correlations and the number of markets, along with how they influence the range of portfolio-level Sharpe ratios and drawdowns can be explored using our interactive [The Future](#) tool.

Figure 2: The power of diversification in trend following



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees, net of transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

The Value of Universe Expansion: Number vs. Effective Number of Markets

Thinking about trend following in this way, while useful, has also led managers to the wrong conclusion about the value of adding new markets. If one assumes that each market in a trend-following portfolio is mostly independent of the others, then the marginal improvement in the Sharpe ratio from adding a new market is proportional to one over the square root of the number of markets. The marginal improvement diminishes significantly as the number of markets already traded becomes large.


This is true as far as it goes, but this reasoning is sometimes used to suggest that going beyond the major markets may not be worth the effort and additional operational complexity.

The main flaw with this conclusion is that the first 50 to 100 markets are not independent of one another – there are in fact clusters of markets that move together. The returns of developed market equity index futures, for example, show a mean inter-market correlation of 0.7 over the past 10 years, and there are similar clusters among the fixed income and commodity markets. This means that the number of uncorrelated exposures – sometimes called the “effective number” of markets – is much smaller than the count of the assets traded. As a result, the marginal gain from adding a new genuinely uncorrelated market is larger than the naïve argument suggests.

We have used principal component analysis in Figure 3 to find how many independent factors explain 80% of a hypothetical trend-following strategy’s variance. This provides us with an estimated 21 different exposures, or effective markets, that explained most of our strategy’s returns over the past decade. The same analysis for a highly diversified trend-following strategy applied to 280 different markets – allocating risk equally across major markets, alternative markets and Chinese futures – identified 42 effective markets.

The Sharpe ratio improvement of increasing a strategy’s 21 uncorrelated exposures to 42 is much more worthwhile than an increase from the 100 individual assets to 280.

Figure 3: Estimating the number of “effective” markets



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

Over the past 10 years, alternative markets and Chinese futures have trended strongly and been highly profitable for CTAs. Interestingly though, the argument we lay out above in favour of trading these instruments does not assume that this will necessarily be true in the future. Even if the Sharpe ratio of each alternative market individually is on par with a typical major market – a reasonable assumption given that the CTA footprint in these markets has increased – the diversification benefits would still enhance the overall portfolio performance.

At the other end of the spectrum to the managers that choose to focus only on the major markets, there are those that seek to maximise the number of markets traded, with little regard for whether the additions are diversifying or not. Examples of such markets include developed market interest rate swaps that duplicate exposures captured by government bond futures or equity sectors that are highly correlated with stock index futures. The main culprits for this type of universe proliferation are large alternative-market CTA strategies seeking to increase their capacity.

Capacity, in general, is an issue once moving beyond the 100 to 150 largest futures and OTC markets available to CTAs. The most scalable trend-following implementations must place more weight on the largest markets and allocators seeking out exposure to alternative markets should identify managers with a disciplined approach to capacity.

Maximising Diversification

Given the faint edge that trend following has in any individual market, the overarching goal when setting market and sector weights should be to maximise portfolio diversification. This enables the strategy to participate most effectively in trends, wherever they emerge.

While these weights are usually set for the long term, the positions the strategy takes in each market and sector over the short term are more dynamic, driven by the trend-following signal and market volatility. At any given moment the strategy takes the most risk where the trends are strongest.


Some managers try to go one step further than this. In addition to the dynamic behaviour of the signal, they also modify their long-term expected risk exposures based on their expectations for trend-following performance in a specific sector.

Upweighting in advance the markets or sectors where trends are going to emerge would improve returns considerably. But effective methods of doing this that improve on the strategy's inbuilt dynamic weighting are difficult to identify. Moreover, much caution is required when researching such methods due to the risk of overfitting. Overfitting occurs when a model has too many parameters relative to the amount of data available and is thus fitted to non-permanent features of the dataset. The result is disappointing returns when the model is used in live trading.⁵

Similarly, when performance is disappointing in certain sectors for periods of time – such as a poor overall run for trend following in commodities and currencies through the 2010s – then a manager might be tempted to underweight or remove those markets from the portfolio. Had they done this, they would have missed out on the strong trends in commodities and currencies since 2020.

Figure 4 illustrates this lack of relationship in sector performance between one period and the next over the past half century using a hypothetical trend-following strategy. Each green marker plots the sector return over a period on the y-axis against its return in the previous period on the x-axis. If periods of negative performance tended to follow positive performance (and vice versa), the markers would be oriented from the bottom right of the plots to the top left.

Figure 4: Lack of predictability in sector-level trend following performance



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees, net of transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

Investors should seek out CTAs that aim to maximise diversification across their strategies, both in the markets traded and the weights they assign to those markets. Strategies that have market weights, or even market exclusions, based on past or recent performance, should be treated with caution as they are likely to be overfitted.

⁵ This topic is discussed at length in our December 2013 paper *Blinded by Optimism*.


2. The Trend Signal

Trend-Following Indicators

The starting point for trend following is a rule, or “indicator” which takes in the price history of an asset and calculates a numerical measure of the direction and strength of the recent trend. There are a lot of different rules that can do this, and some common examples are listed below and shown in Figure 5:

- **Time-series momentum:** whether a market’s return has been positive or negative over a period.
- **Moving-average crossovers:** whether a market’s short-term moving average price is trading above or below its long-term moving average price.
- **Breakouts:** whether a market’s price is trading above or below a breakout channel, which is usually based on a measure of price volatility around a long-term moving average.


Figure 5: Examples of common trend-following indicators



Source: Winton. Indicative trend-following models using similar lookback windows.

All these indicators are designed to capture the same effect: the price trend. It is therefore unsurprising that the direction of the signal is the same most of the time when the lookback parameters are set to similar values. This is visible in Figure 6, which shows stylised versions of the three types of indicator applied to LME copper since 2014. Over the long term, these three indicators produce returns with a correlation of 90% or more when applied across a diverse portfolio of futures markets.

Figure 6: Different indicators based on the same periods of data produce similar positions



Source: Winton, as at 30 November 2023. Indicative trend-following indicators using similar lookback windows.

Simplicity vs. Complexity

Despite the high correlation between trend signal functions designed to capture the same effect, some managers make a virtue of combining large numbers of indicators in their strategies. Rather than improving returns, this adds complexity and creates other problems.

First, adding new parameters may improve the strategy's backtested performance modestly, but – as with attempting to time sector performance – it also increases the risk of overfitting.

Second, the combination of multiple, overlapping indicators makes it harder to interpret the strategy's results and truly understand the interaction between other, more important parameters, such as speed, response function and volatility. Is a new indicator adding value because of the indicator itself or because the indicator results in a slightly different speed?

The focus should instead be on an efficient and uncomplicated indicator that can produce an intuitive trend signal, ideally with only incremental changes each day to minimise transaction costs. Over 26 years of research and experience we have found moving-average-based signals to be as effective as any.

The Key Choice: Speed

All trend indicators have a parameter that controls how fast the signal changes and how quickly it absorbs new information about changes in market direction, usually called the "speed" of the signal.

The choice of this parameter is a key decision for the manager of a trend-following CTA.

Sometimes investors like to measure the impact of the choice of speed for a strategy by looking at the average holding period for a position – the typical length of time for which the strategy maintains a trade in one direction. But this mixes up a feature of the strategy (its responsiveness) with market behaviour (the typical duration of trends).

A purer way to think about speed is to measure the time required to close out a position when a strong trend starts to reverse. This focuses on the responsiveness of the signal and is also very important for helping understand how a strategy will react at moments where a trend is in the process of reversing.

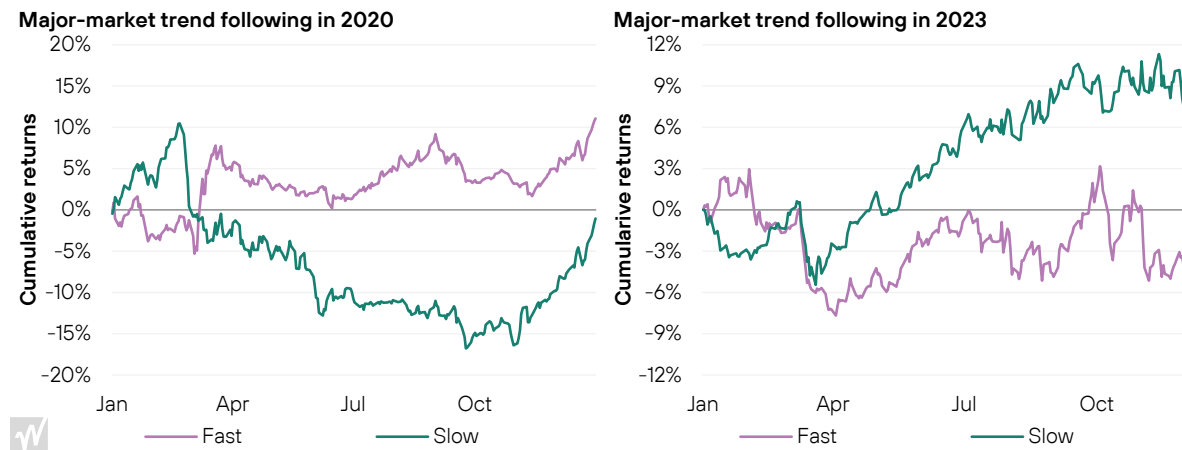
Faster is Not Necessarily Better

There is a common misconception with systematic investment strategies: the idea that faster is always better. The logic to this line of thought is that the more responsive a strategy is, the more likely it is to adapt its positions and make money when markets change direction.

The reality in trend following is more nuanced. While faster models are quicker to identify new trends and cut their positions more rapidly when markets turn, these models then suffer if the previous trend resumes. Slower models, on the other hand, perform well in long steady trends, but can underperform at turning points.

In Figure 7, we show two recent market episodes where the speed of a major-market trend-following strategy led to vastly different results. The "fast" strategy uses an 8- and 24-day moving-average crossover indicator, while the "slow" strategy uses 32- and 96-day averages. Both strategies are otherwise identical, allowing us to directly compare the effect of speed in each scenario.

Figure 7: Different speeds of trend-following strategy can lead to very different results



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

When market reversals play out over weeks, such as in equities after the global outbreak of Covid-19 in 2020, then faster strategies can turn around in time to generate profits on the downside. This type of situation is difficult for slower strategies, which reduce their positions more slowly. In this episode, by the time slower strategies had turned short, stock markets began to rally, leading to further losses.

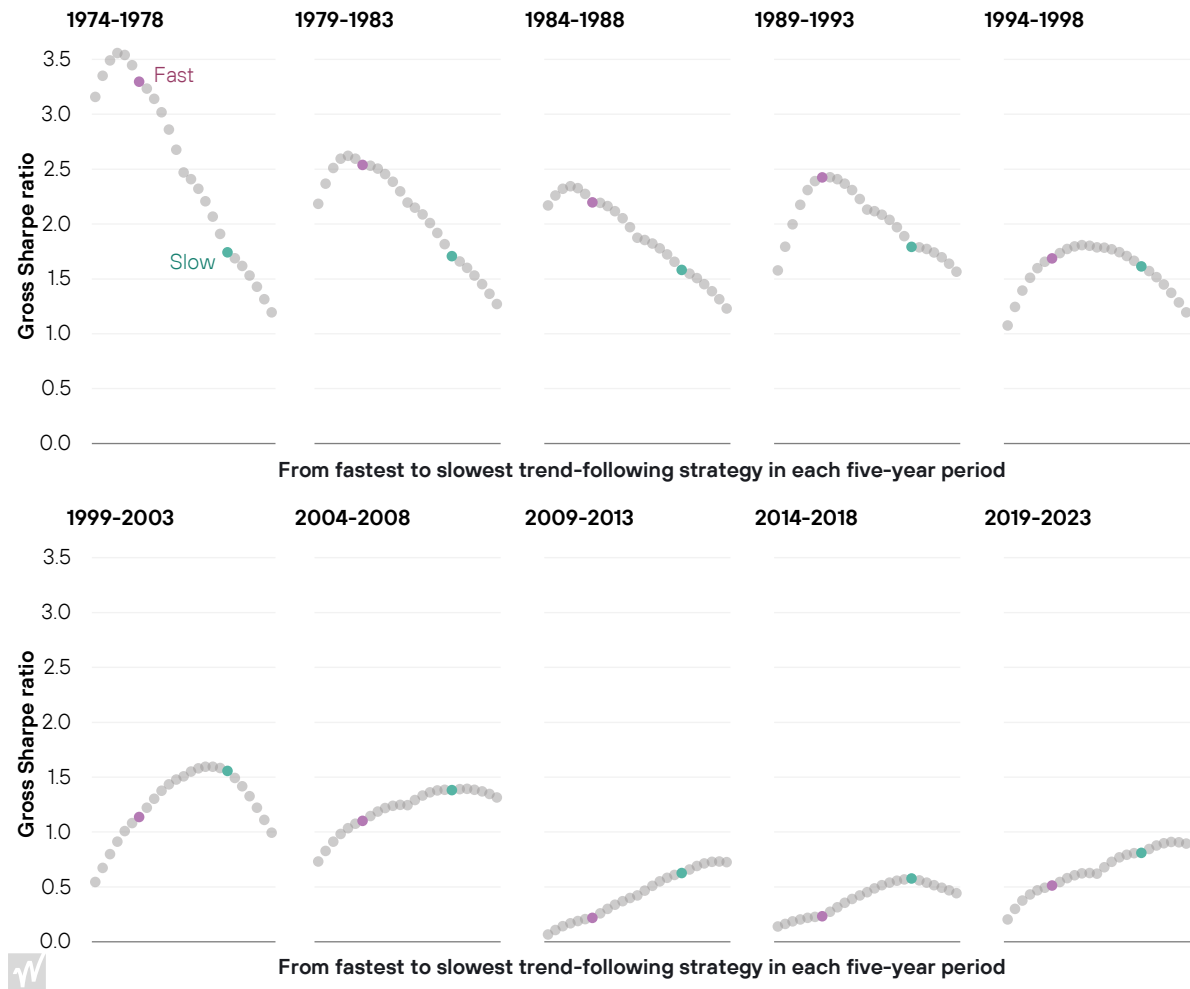
A counterexample to 2020 is 2023. Most speeds of trend following had similar positions heading into the collapse of Silicon Valley Bank in March. This meant that speed did not make a difference when the previous downtrend in government bonds reversed and most trend-following strategies declined together. Faster strategies, however, cut their bond positions too quickly and turned long, which led to losses when the previous trend resumed. This whipsawing market behaviour continued over the remainder of the year, while slower strategies benefited from gradual, longer-term trends.

The Long-Term Decline of Fast Trend Following

The performance of different speeds of trend-following strategy varies unpredictably over shorter time horizons. Deciding whether faster or slower strategies perform better over the multi-year holding periods that are relevant to most investors is therefore an empirical question that can be settled only by looking at real market data.

Over the longer term, there are clear trends in the data. In Figure 8, we show the Sharpe ratio of 25 different speeds of trend-following strategy over the past 50 years split into 10 discrete five-year periods. We have highlighted in purple and green the speeds shown in Figure 7.

Figure 8: Long-term performance evolution of 25 different speeds of trend-following strategy



Source: Winton, as at 30 November 2023. The purple and green dots correspond with the “fast” and “slow” speeds shown in Figure 7. The speeds in this analysis range from 4/12 to 64/192 day moving average crossovers. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

Starting with the earliest three five-year periods in the top left of Figure 8, we see that in the 1970s and 1980s the fastest trend signals were the most profitable. These backtests do not include transaction costs, so in that period a careful optimisation of costs against gross Sharpe ratio was necessary.

Moving along the plots, the analysis reveals how the performance of faster models degraded over the 40 years to 2013, while slower strategies proved more resilient. Since the 2008 financial crisis, the faster models have performed poorly both relative to slower models and in absolute terms.

Our approach to this data is empirical rather than ideological. When Winton first identified this phenomenon in the late 2000s⁶, we tilted our strategies towards slower models. However, if our research identified a change in the long-term performance trends of different speeds, then we could speed up our strategy. Interestingly, in the recovery of trend following performance more generally since 2019, slower models have continued to lead the way.

Why Has the Performance of Faster Models Declined?

The decrease in the Sharpe ratio of faster trend signals over time seems clear from Figure 8, but in 2010 not everyone interpreted the data in the same way.

As trend following became better understood and CTAs lost their “black box” image, some began to describe the strategy as a “beta” or a “risk premium”, borrowing the language of the Capital Asset Pricing Model. The implication was that the performance of trend-following strategies was an intrinsic and ultimately static feature of futures markets – a reward that should be expected to follow from taking on a particular type of risk.

Coming from that point of view, changes in performance over time are dismissed as a statistical variation, and one arrives at performance expectations for different speeds of trend following by doing the longest backtests possible. This naturally favours the faster signals, because their performance was so much better in the 1970s and the 1980s.

Our view in the late 2000s was that this argument relied on historical data that was no longer relevant. And the record since then has supported our decision to slow down the strategy. But if the decline in the performance of fast trend following was not just natural statistical variation, then what did cause it? We do not have firm proof, but the evidence available suggests that the increase in the assets under management of trend-following CTAs is the most likely culprit.

Sharpe Ratio vs. Convexity

The discussion above focused on maximising the Sharpe ratio of a trend-following strategy. But standalone performance is not the only consideration for a CTA allocation. The reason investors allocate to trend following is usually due to the strategy’s ability to generate returns that are diversifying for traditional equity and bond investments.

Over long periods of time, trading speed does not particularly matter in this regard: all 25 speeds of trend-following strategy shown in Figure 8 have had low to zero correlation with equities and bonds over periods of 10 years or more. However, some allocators are looking for a strategy that can provide returns more specifically during bad months or quarters for other parts of their portfolio.

⁶We published our findings in December 2013 once it had become apparent that many of our peers had also slowed down their strategies (see [Historical Performance of Trend Following](#)).

A diagram demonstrating this property of trend following that rose to prominence in the 2010s is a plot of the return for the strategy against the return for the stock market along with a curve of best fit.⁷ As Figure 9 shows, for trend-following CTAs, this line indicates a convex return profile; that is, strong performance at the same time as extreme returns for the stock market. The shape is sometimes known as the “trend following smile”.

Caution is required in the interpretation of these charts. First, there are few data points for the extreme equity returns, even when using a half century of data. The addition or removal of a single data point can change the line of best fit and you would be hard pressed to spot any relationship if it was not for this line. Second, presenting the strategy’s returns in this way ignores the path-dependent nature of trend following performance. For example, a sudden drop in the equity market after a long uptrend, as was the case in October 1987 and February 2018, led to losses for trend followers.

We show how the fast and slow strategy used in the previous analyses fare in terms of convexity in Figure 9. The charts on the left consider the full 1974 to 2023 period, while the charts on the right examine the post-2000 environment, when slower trend-following strategies have done better than faster strategies.

On this measure, the faster strategy looks more appealing with its stronger returns in difficult months and quarters for equities. However, since 2000, this extra convexity has come at a cost, with lower returns on average and near-zero average returns in more muted months and quarters for equities.

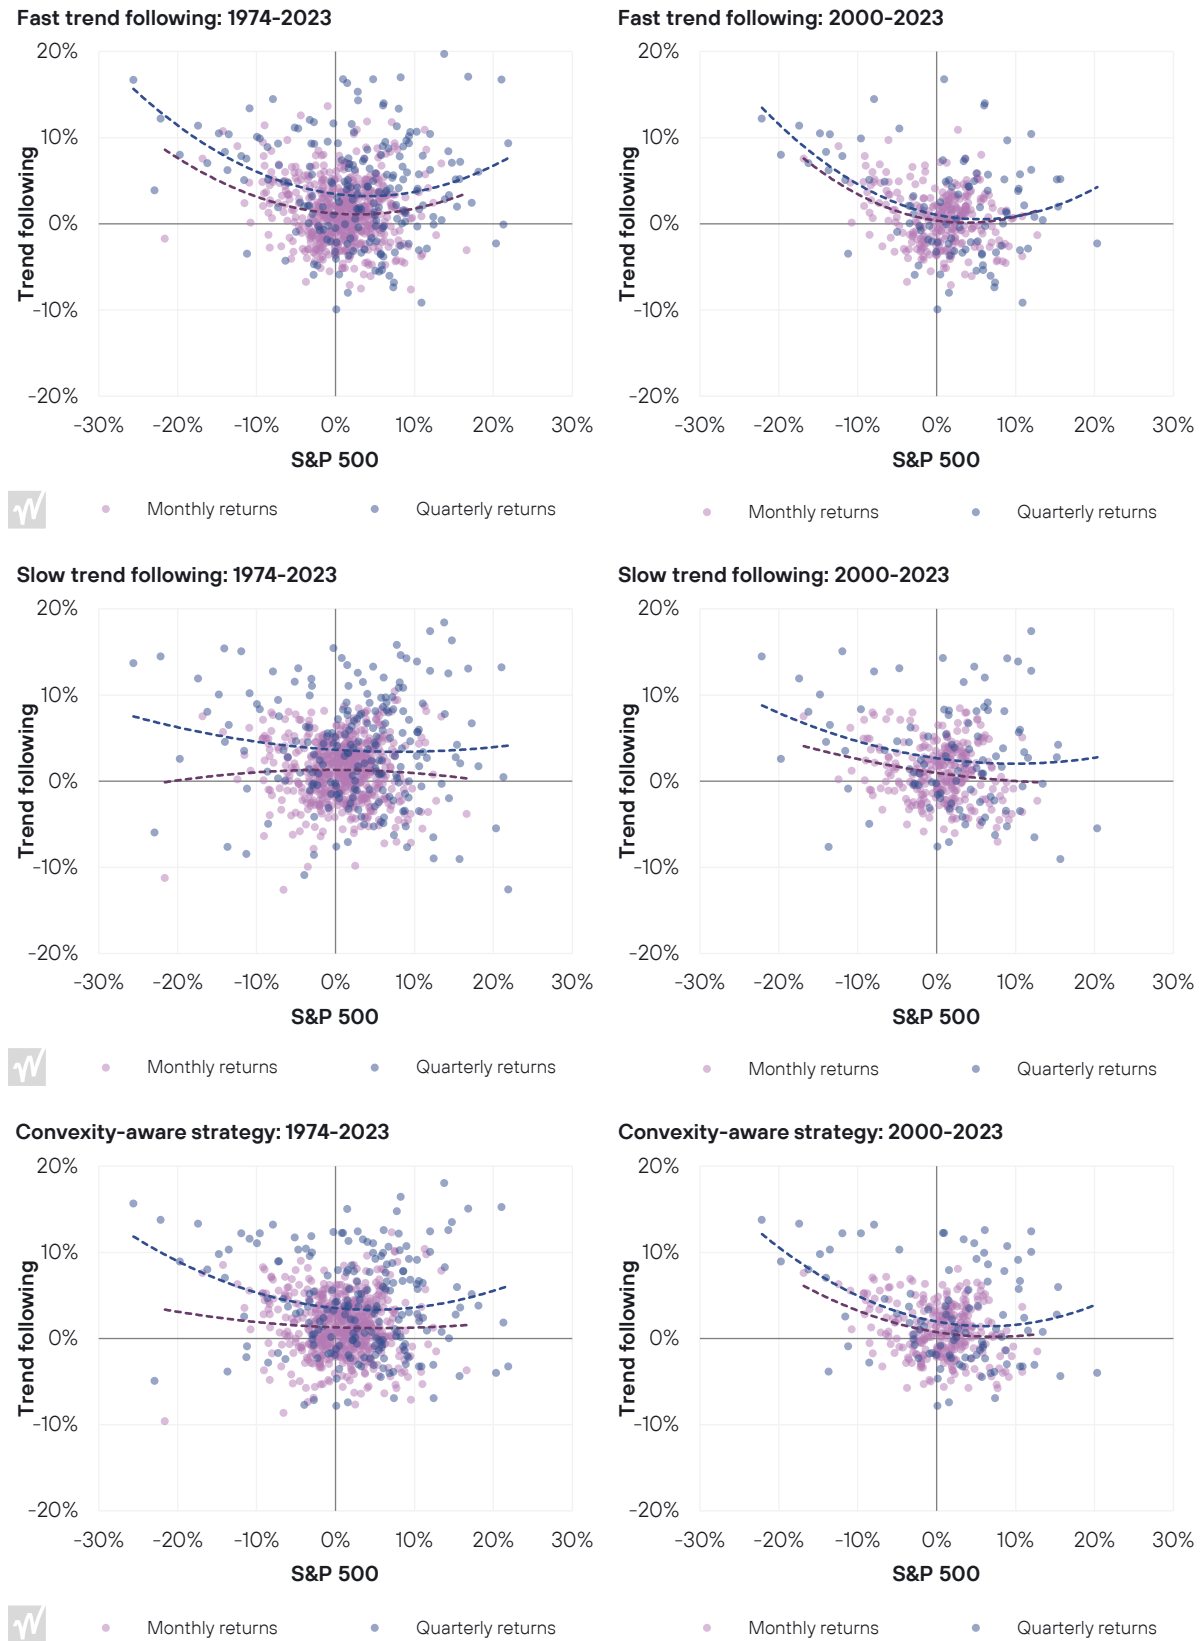
There are several ways trend followers can address the emergence of this trade-off between returns and convexity. One option is to add other diversifying signals, which we will examine in the next section of this paper.

Another method – which we use at Winton – is to tailor the blend of trading speeds across sectors. Rather than trading all sectors at the same speed, our trend-following CTA strategies are tilted to shorter-term models in stock indices, the sector where it is most effective to trade faster from a portfolio diversification perspective. The remaining commodities, currencies and fixed income sectors are then tilted to slower models to maximise returns. The results of this “convexity-aware” approach are shown in the final two plots of Figure 9.

Regardless of preferences around trading speed, Figure 8 and Figure 9 demonstrate how important this topic is for allocators seeking to identify strategies with trading speeds that are aligned with their investment objectives.

⁷ While appearing in the marketing literature of trend-following CTA strategies from the 2010s, the origins of this idea date back to at least a 2001 paper that found trend-following strategies produced returns like those of an options straddle strategy: William Fund, David A. Hsieh, *The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers*, 2001.

Figure 9: Faster strategies show more “convexity” than slower strategies



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

Counteracting Trend Exhaustion


Trend-following strategies do not assume that trends will last forever. When price movements suggest that a trend is coming to an end, then the model responds by reducing the position and eventually switching direction; the faster the strategy, the faster this process happens, albeit after the fact. Managers can also incorporate features into their models that seek to reduce the profits given back when trends reverse by taking action ahead of the reversal.

One feature that tends to do this is sizing positions inversely to volatility. Volatility-weighting rose to prominence in the 2000s and is now employed universally by the world's largest trend-following CTAs. If, all else being equal, market A has twice the return volatility of market B, then a strategy would take a position in market A that is half the size of a position in market B. This serves to normalise an investment universe with annualised volatilities ranging from low single digits (short-term interest rates) to 30% or more (energies), enabling trend-following strategies to be applied consistently across a diverse collection of markets.

The reason that volatility weighting can sometimes mitigate trend exhaustion is that when market volatility spikes, then the strategy reduces its position, even if the trend is strengthening. This played out in several commodity markets after Russia's invasion of Ukraine in February 2022, with trend-following strategies responding to the elevated market volatility by reducing positions before markets gave back their gains from April that year.

Another approach is to apply a "filter" or a "response function" to the raw trend indicator. Figure 10 compares a few choices. The linear function is equivalent to just using the indicator as it is, while the other examples build the position in the emerging trend early and dampen the position when the trend signal rises towards the top of its range. We illustrate examples of when this feature would have added value in Figure 11: gold during 2020 and US Treasuries in 2022.

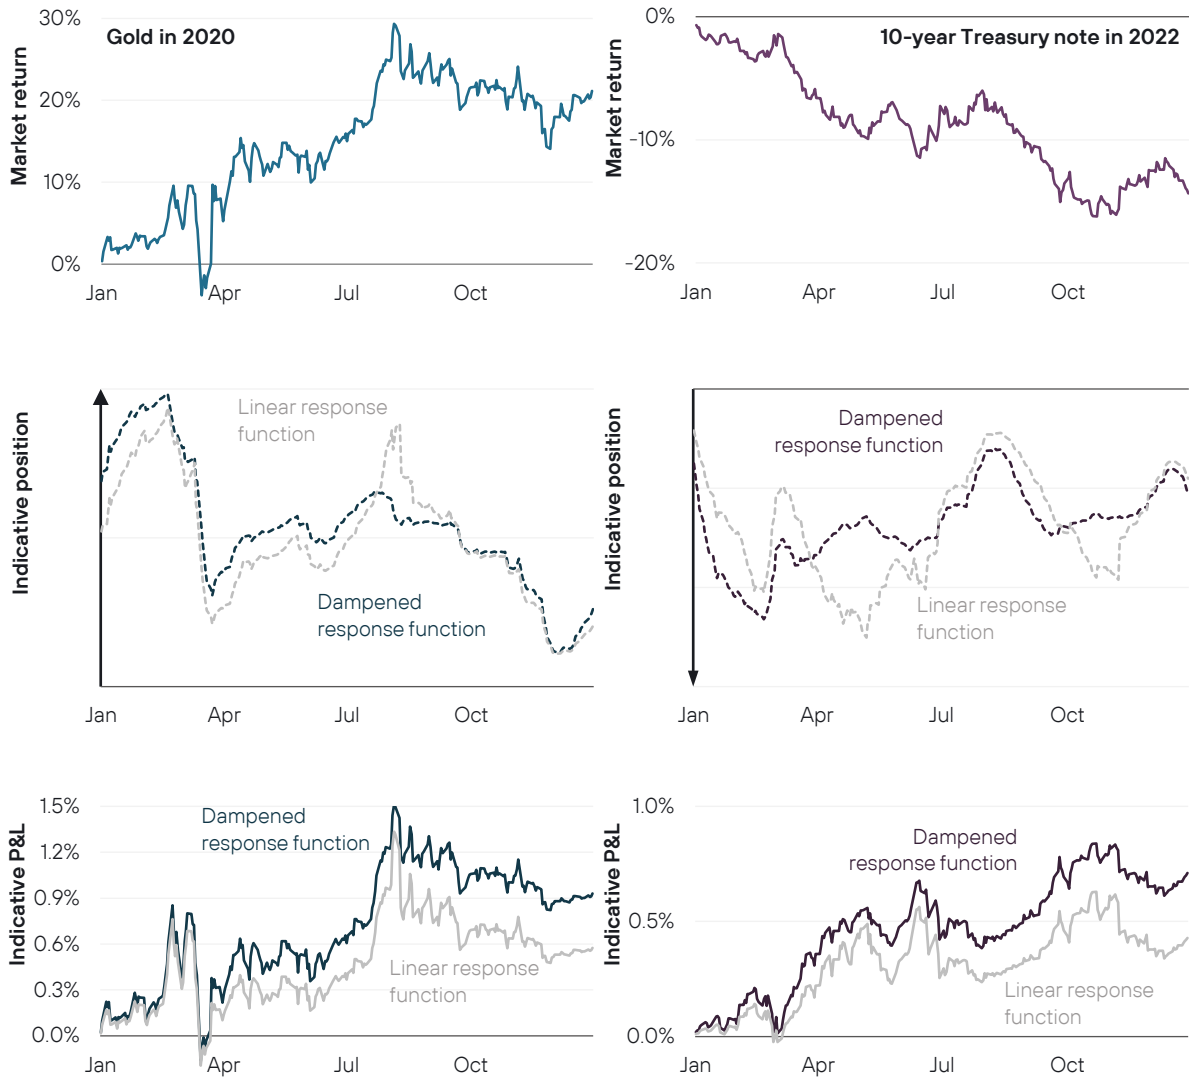
Figure 10: Examples of trend-following strategy response functions



Source: Winton.

The dampened function leads to a modest performance improvement over the long term on average, with much of the advantage coming at times when trend following struggles. The downside is that it can lead to counterintuitive position changes. For example, when the indicator is weakening from its maximum level, the response function can cause the position to increase. This issue can be mitigated by tapering the response function at extreme values, such as in the rightmost plot of Figure 10.

Figure 11: Dampened response function trade examples



Source: Winton. **Past performance is not indicative of future results.** The above charts show a hypothetical simulation of a signal applied to an instrument over a selected time period. The signal P&L is shown gross of fees and net of transaction costs. The instrument and time period have been selected for illustration purposes. This example is not used to demonstrate the adviser's track record and is shown only to demonstrate how a particular signal works. Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

3. Diversifying Signals

Why Go Beyond Trend Following?

Although the scope of systematic futures trading is very broad, in practice, CTAs predominantly use trend-following signals to determine their positions. The main reason for this is that trend following has proven highly successful at generating diversifying returns for investors over the past 50 years.


Complementary non-trend signals still have a role to play, however. In particular, they can help to smooth out the drawdowns that trend following suffers periodically, such as in the periods 2009 to 2013 and 2015 to 2018. This in turn makes it easier for investors to stick with the allocation and then benefit during periods like 2020 to 2022 when trend following is needed the most.

In order to diversify the trend-following exposure, a non-trend signal must have low correlation with trend following and positive performance, which together increases the overall strategy's Sharpe ratio. The low correlation also means that reallocating risk to these signals does not necessarily require a commensurate reduction in the exposure to trend following, which helps the strategy maintain its attractive diversifying properties.

We can illustrate this "free lunch" effect by constructing a hypothetical trend-following strategy that allocates 80% of risk to the convexity-aware trend-following strategy used earlier in Figure 9 and 20% of risk to diversifying non-trend signals. The diversifying component has a long-term correlation to trend following of 0.2 and the full strategy operates at an annualised volatility of 10%.

Figure 12 shows the annualised return volatility through time of this multi-signal strategy, along with the volatility of the two sub-components. For comparison, we have included a pure trend-following strategy operating at 10% annualised volatility. The annualised volatility for the full 20-year period for each line is shown in brackets in the legend. The trend component of the multi-signal strategy is capturing 92% of the trend-following exposure of the pure trend strategy over the long term, despite a 20% allocation to diversifying signals.

Figure 12: Diversifying signals can be added without much reduction in trend exposure



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

What are Diversifying Signals?

The non-trend risk allocation in Figure 12 is split equally between multi-asset carry signals and an ensemble of systematic macro signals, which are both applied to the same universe of major futures markets as the core trend-following strategy.


Carry is an established systematic signal that has been traded by trend-following CTAs in currencies since at least 1992 and the concept was expanded to other asset classes in the early 2010s. The signal profits from futures prices converging towards the underlying spot price. Like trend following, carry has high capacity and can be applied consistently across many markets and multiple asset classes. However, unlike trend following, the signal generally benefits from stable prices, which makes it highly complementary.

The rest of the diversifying component combines a range of fundamental, seasonal, and cross-sector signals. These signals are harder to generalise across managers as they are often proprietary, capture less well-known effects, and vary widely by implementation. This is evident in the dispersion of returns for standalone systematic macro strategies offered by CTA managers, which have a mean correlation with one another of 0.3. For the purposes of this analysis, we use Winton's own systematic macro signals, many of which have been traded for more than a decade.

To assess the effect of including such signals on the strategy's performance, we penalise the simulated returns of the two diversifying components to reduce their long-term Sharpe ratios to 0.7. This has been done partly to mitigate hindsight bias and partly to demonstrate how such signals can add value, even at a performance level below the 0.9 gross Sharpe ratio displayed by the standalone trend-following strategy over this period.

Figure 13 shows the results of this analysis. The 20% allocation to diversifying signals improves returns over the 20-year period by 2.2% per annum, increasing the gross Sharpe ratio of the strategy from 0.9 to 1.1. This return improvement is broadly consistent through time, with the right plot in Figure 13 showing how the strategy with the diversifying signals outperforms the pure trend-following strategy over all rolling five-year periods.

Figure 13: The effect of adding a 20% allocation to diversifying signals



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.


How Much Risk Should be Allocated to Diversifying Signals?

The inclusion of non-trend signals by CTA managers has occasionally been criticised on the grounds that it can dilute the attractive diversifying properties associated with trend following, particularly the convexity feature highlighted in Figure 9.⁸

As we illustrated with trend following itself, the reality is more nuanced and depends on the signals traded and their implementation. For example, currency carry rightly has a reputation for being a risk-on strategy, particularly in emerging markets. But the same is not true in fixed income and commodities. The key then is to understand clearly the characteristics of a manager’s diversifying signals, how they interact with its trend-following strategy, and ensure that they are weighted appropriately.

We can illustrate this point by examining the multi-signal strategy assessed in Figure 13, which is designed specifically to maintain the attractive properties associated with trend following. As Figure 14 shows, when the returns of both this strategy and the pure trend-following strategy are plotted versus the S&P 500 for the same 20-year period, the convexity profiles are indistinguishable.

Figure 14: The effect of adding a 20% allocation to diversifying signals on “convexity”



Source: Winton, as at 30 November 2023. **Past performance is not indicative of future results.** The above results are based on a hypothetical strategy shown for research purposes only. Returns are gross of fees and transaction costs and exclude interest earned on cash. **They do not reflect actual trading results and are not representative of a strategy or investment product.** Hypothetical performance has inherent limitations, some of which are disclosed in the Appendix.

⁸ For example; Risk.net, [Some trend followers are less than ‘pure’ – study](#), 26 Mar 2019; FT Alphaville, [Stay Friends with Your Trends](#), 23 Aug 2022

4. Risk Management

The paper so far has covered the choice of assets in a CTA strategy, along with different aspects of the trend and non-trend signals that are used to determine a strategy's exposures. We now turn to perhaps the most important aspect of the portfolio construction – understanding and managing the overall risk.

The Aim of Risk Management

Large and unexpected market moves, like those that occurred in 2008, the onset of the Covid-19 pandemic in 2020, or in March 2023, are a fact of life. There is nothing a CTA manager, or anyone else, can do to stop them occurring. And since they are unexpected, there is always a chance that they will act against current positions and cause losses.

In order to achieve a return above the risk-free rate, any strategy must take on some amount of risk. What matters then is ensuring that the strategy's risk levels are kept under control, so that adverse events do not destroy the rationale for the investment in the long run.

Historical risk events for systematic investors, like the 2007 Quant Quake or the collapse of Long-Term Capital Management in 1998, provide a clear lesson: when large market moves occur, the best outcome is achieved by continuing to trade the strategy as designed and without significant intervention.

Liquidity is often strained in chaotic, fast-moving market environments, and transaction costs can be orders of magnitude higher than usual. This is the worst time to be intervening in the strategy to cut risk. If anything, it is preferable to avoid trading at all in the most volatile moments.

This sounds easier than it is. The temptation for a manager to cut risk and "fix the problem" can become overwhelming when losses are large. The conclusion is that risk management must therefore be done in advance, and that CTA strategies need to be designed so that dramatic market events do not require an interruption to the systematic process.

Risk is Not the Same as Volatility!

There is an immediate problem with the ambition of coping with losses from dramatic market events: such a loss is very difficult to estimate. How does a manager even begin to guess the risk of an event like a war, a pandemic, or something worse that we have yet to experience?

What happens in practice in the CTA industry is a classic example of a behavioural heuristic, as described by Daniel Kahneman in *Thinking, Fast and Slow*.

"...when faced with a difficult question, we often answer an easier one instead, usually without noticing the substitution."

The difficult question: "what is my risk level?" – is replaced by the (much) easier question: "what is my portfolio volatility?" As Kahneman says, the substitution is often unconscious – when asked about target risk levels, many CTA managers will reply by quoting their target volatility level!

One can sympathise with this. The volatility of a strategy's returns is easy to measure and target, and the process of maintaining a fixed target volatility is easy to automate. But maintaining a constant level of volatility does not result in a consistent level of risk.

Moreover, a risk management approach based solely on volatility can lead to increased positions at exactly the time when the risk of sudden and extreme losses is at their highest. This was the case in March 2023, as we will discuss below.

March 2023 Case Study

To calculate the expected volatility of a portfolio we need three ingredients: the positions held in each asset, the volatilities of the assets, and the correlations between the assets. The positions are facts the manager knows with certainty, but future asset volatilities and correlations can only be estimated.


As we discussed in the first section of this paper, low correlations between the assets in a strategy's trading universe results in valuable diversification. But there is another way strategies can produce diversification when correlations are higher, which deserves more caution.

In early March 2023, the correlation between equities and bonds was positive, but CTAs were long equities and short bonds. The combination of positive correlation and positions pointing in the opposite direction meant that the equity and bond exposures were offsetting, reducing the overall volatility of the portfolio.

Managers who target a fixed level of volatility therefore leveraged up at this moment. However, despite their strategy's volatility holding constant, their risk level was increasing, because the correlation between equities and bonds is not stable and is prone to sudden reversals.

This reversal is exactly what occurred when Silicon Valley Bank collapsed and other regional US banks became fragile. The stock-bond correlation suddenly became negative, as concerns about the safety of bank deposits drove a dramatic flight into US Treasuries, while the equity market fell. Together, these moves led to negative returns many times larger than would have been expected based on portfolio volatility estimates alone.

Figure 15: Diversification assumed within portfolio volatility disappeared in March 2023



Source: Winton, Société Générale. Equity-bond correlation is based on S&P 500 and US 10-year-note futures returns.

Beyond Volatility Targeting

While the collapse of Silicon Valley Bank came as a surprise, the elevated risk of a correlation reversal was not so difficult to see. A risk of a very similar kind had appeared almost exactly 10 years prior, in the so-called “Taper Tantrum” of May 2013. In this instance CTAs were long both bonds and equities, and at the time, the correlation between them was negative, again reducing the portfolio volatility. When fears grew that the US Federal Reserve would end quantitative easing prematurely, bonds and equities fell together. This led to large losses for CTAs who had leveraged up on the assumption that the positions would always be offsetting.⁹

The lesson from these episodes is that to manage portfolio risk effectively, a CTA manager must incorporate many different metrics into their risk management, not just volatility. This is our approach at Winton and a stress test that looks at the impact of sudden changes in correlation found in early March that although volatility was decreasing, other risks were still elevated.

In practice, the volatility of a strategy where the manager is truly employing a multifaceted approach to risk will vary in a range. If other measures of risk are elevated, then a lower portfolio volatility may be prudent at that moment.

Conclusion

Investors in trend-following CTA strategies are presented with a vast choice of managers, all with different ways of profiting from the same effect.

Too often managers seek to differentiate their strategies with features that are, at best, inconsequential, or at worst, overfitted and detrimental for long-term performance. Alternatively, managers will seek out – sometimes spurious – weaknesses in notable trend-following innovations that they do not have the resources to implement, such as alternative markets or diversifying signals.

This paper aims to provide allocators with the information to distinguish between trend-following CTAs effectively by describing in detail the four main differentiators: 1) universe; 2) trend signal; 3) diversifying signals; and 4) risk management.

We believe that these four factors explain most of the return variation in the industry. Taking the time to understand their role and interaction will, in our view, increase the chances of identifying the best trend-following CTA for meeting an allocator’s investment objectives.

About Winton

Winton is a quantitative investment management firm, founded by Sir David Harding in 1997 and headquartered in London. We are a leader in trend following, a style of systematic investment strategy we have pioneered for more than 25 years, and we count some of the world’s largest institutional investors as clients. For more information, please email investorservices@winton.com or visit www.winton.com.

⁹ See our October 2019 article [Risk and the Cult of Volatility](#) for more on this topic.

Appendix: Hypothetical Performance

Hypothetical performance results have many inherent limitations, some of which are described below. No representation is being made that any account will or is likely to achieve profits or losses similar to those shown. In fact, there are frequently sharp differences between hypothetical performance results and the actual results subsequently achieved by any particular trading program.

One of the limitations of hypothetical performance results is that they are generally prepared with the benefit of hindsight. In addition, hypothetical trading does not involve financial risk, and no hypothetical trading record can completely account for the impact of financial risk in actual trading.

For example, the ability to withstand losses or to adhere to a particular trading program in spite of trading losses are material points which can also adversely affect actual trading results. There are numerous other factors related to the markets in general or to the implementation of any specific trading program which cannot be fully accounted for in the preparation of hypothetical performance results and all of which can adversely affect actual trading results.

Important Information

This document is prepared by Winton Capital Management Limited ("WCM" together with its affiliated companies, "Winton") which is authorised and regulated in the United Kingdom by the Financial Conduct Authority. Its registered office address is 20 Old Bailey, London EC4M 7AN, England.

Opinions expressed are those of the author and may not be shared by all personnel of Winton. These opinions are subject to change without notice, are for information purposes only and do not constitute an offer or invitation to make an investment in any financial instrument or in any product to which Winton provides investment advisory or any other financial services.

The information included is believed to be materially correct but Winton makes no representation or warranty as to its accuracy or completeness and accepts no liability for any inaccuracy or omission. This document is confidential and you should not disseminate, distribute or copy it or any of its contents. Past performance is not indicative of future results.

This document contains information sourced from S&P Dow Jones Indices LLC, its affiliates and third party licensors ("S&P"). S&P® is a registered trademark of Standard & Poor's Financial Services LLC and Dow Jones® is a registered trademark of Dow Jones Trademark Holdings LLC. S&P make no representation, warranty or condition, express or implied, as to the ability of the index to accurately represent the asset class or market sector that it purports to represent and S&P shall have no liability for any errors, omissions or interruptions of any index or data. S&P does not sponsor, endorse or promote any Product mentioned in this material.

The SG Trend CTA is designed to represent the performance of the 10 largest trend following CTA programs. To qualify for inclusion in the index, a program must be open to new investment, report returns on a daily basis, be an industry recognised Trend Follower, and exhibit significant correlation to trend following peers and the SG Trend Indicator. At the end of each year all CTA programs in the SG CTA database that meet the inclusion requirements are ranked by program assets. The 10 largest programs are selected as index constituents for the following year. At the beginning of the year a hypothetical portfolio is formed with each constituent program given an equal allocation. The index daily return is simply the daily return of this hypothetical portfolio. There is no rebalancing of allocations during the year.

The HFR index returns are shown for illustration and performance comparison purposes only. They are not formal benchmarks and do not form part of the formal investment objectives or investment policy. It is not possible to invest in the indices. The indices shown are: i) those in the sub-section category of "multi-strategy" and represent an equal-weighting of single-manager funds that report to the HFR Database and satisfy the relevant HFR criteria for inclusion. It is selected as a comparison to other multi-strategy hedge funds; ii) an index representing the sub-category "systematic diversified" which comprises an equal-weighted index of single-manager funds that report to the HFR database and satisfy the relevant criteria for inclusion. It is selected as a comparison to diversified systematic hedge funds; and iii) an asset weighted composite index of single manager funds that report to the HFR database and have a minimum of \$50 million under management or \$10 million under management and a twelve month track record of active performance. The constituents are weighted according to the AUM reported by each fund for the period month. It is selected as a general comparison to the hedge fund industry. Please refer to the HFR website for further information <https://www.hfr.com/hfri-indices-index-descriptions>.

Morningstar 2023. All rights reserved. The information contained herein: (1) is proprietary to Morningstar and/or its content providers; (2) may not be copied, adapted or distributed; and (3) is not warranted to be accurate, complete or timely. Neither Morningstar nor its content providers are responsible for any damages or losses arising from any use of this information, except where such damages or losses cannot be limited or excluded by law in your jurisdiction. Past performance is no guarantee of future results.