How Reliable is Crisis Alpha?

Robert Hillman[1]

17 March 2017, London

There has been a rise in the use of trend following to provide portfolio tail protection. Marketing exploits the empirical observation that trend following has tended to be profitable during times of crisis. But this so-called crisis alpha is controversial. Some established managers embrace the concept, others eschew it. This paper explores what lies beneath the caution. We propose a simple method for quantifying some of the risks around crisis alpha. And we present daily conditional stress tests for an industry proxy CTA. We focus on trend following but our methodology would suit other popular strategies like risk-parity or volatility-targeted portfolios.

The term crisis alpha was coined by Kathryn Kaminski and refers to the empirical observation that during times of market crisis trend-following has tended to be profitable, see Chart 1 and Kaminski (2011) and Greyserman and Kaminksi (2014). The performance of trend follower funds in October 2008 is often cited.

These types of chart frequently appear in hedge-fund marketing materials. While often shrouded in compliance-friendly but painfully worded caveats and disclaimers, they are there to suggest, but most definitely not promise, that the strategy being marketed is helpful in times of stress.

But the idea of crisis alpha is controversial. It is an empirical observation not a guarantee. There is no theory that says that trend-following will always deliver during a crisis. A clear example of how the concept splits the industry is that two established managers, who share the same DNA, have taken different views. AHL appear relaxed about using the term (Hamill, 2016) while David Harding of Winton is somewhat less enthusiastic, to put it mildly (Harding, 2016).

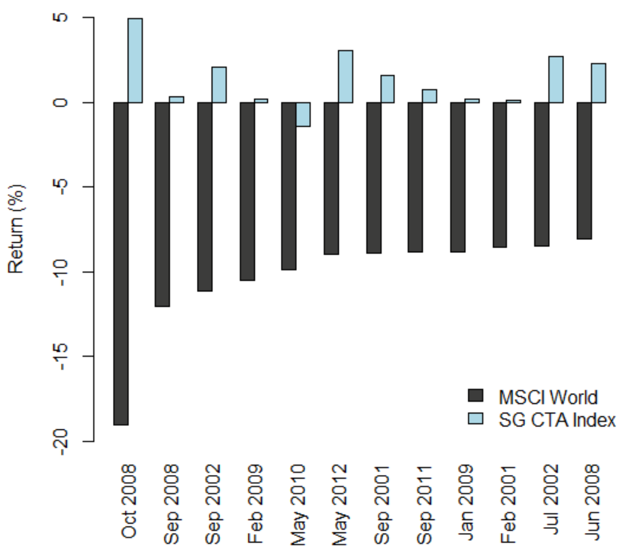

Chart 1. Crisis Alpha for the SG CTA Index

Source: Neuron. Notes: SG CTA Index returns and MSCI World returns for the worst 12 months for the MSCI since 2000[2]. Past performance is not a guarantee, or indication of future performance.

---

[1] Robert is CIO of Neuron Advisers LLP. Updated 22 March to include Appendices 1 & 2.

[2] The SG CTA Index is an equally-weighted average of the performance of major CTAs. For more info see here https://cib.societegenerale.com/fileadmin/indices_feeds/SG_CTA_Index_Methodology.pdf

Yet despite this there are signs that the use of trend-following to provide crisis alpha is on the rise. Indeed, several large US public pensions plans have been allocating capital to trend-followers explicitly for this purpose, articulating other concepts like 'crisis-risk offset' and 'risk-mitigation'[3].

But imagine the following exchange. The investor, a manager of a multi-billion US pension plan, calls up their investment consultant ABC Consulting:

Investor: 'We have a problem. We've just been through two months of market stress, our equity portfolio is down 10%, our bond portfolio is down 20%, and our 'crisis-risk offset' portfolio that you designed for us is down 13%! The press are calling this Lehmans 2.0, and in your marketing docs that I have in front of me you showed us that this strategy delivered 21% over the Lehman crisis. How am I to explain this to my board, let alone my stakeholders?'

Consultant: 'Well yes, I can see it doesn't look good. But we did say that past performance was no guarantee to future performance. It was in the footnote under……'

Let's face it neither party wants to be in that situation. In this note we aim to demonstrate what we think are the two key properties of trend following that lie behind the caution that surrounds crisis alpha. We also propose a simple way to use these properties constructively and provide a daily stress test to complement the traditional risk metrics managers and investors rely on.

A toy model

To make our key points up front, we consider first a toy model trading one market, the S&P future. It is a purely technical model as found in countless articles and books[4]. It uses a single moving-average crossover rule to determine whether to be long or short[5]. We have calibrated it to produce a return volatility of 15% per year.

Chart 2 shows the returns to this trading rule during the worst 12 months for the MSCI since 2000. It has all the hallmarks of crisis alpha. In all but 2 of the 12 worst months the rule appears to post positive returns. In October 2008, when the S&P was down around 200 points (17%), it posted nearly 5%. Note it wasn't much use in May 2010, the month of the Flash Crash.

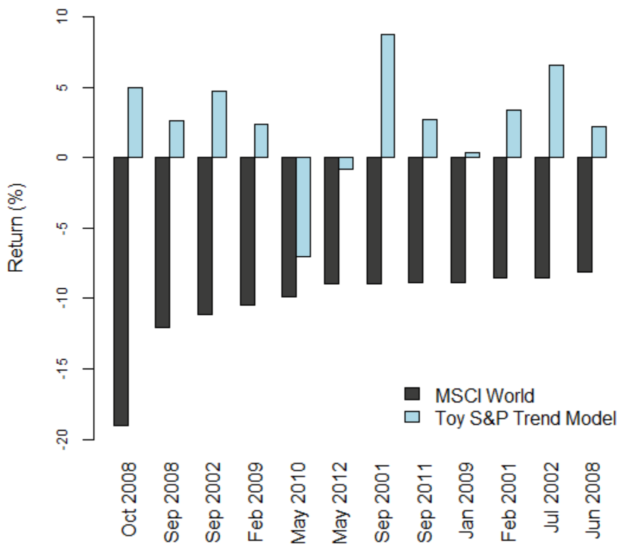

Chart 2. Crisis Alpha for a toy S&P model
Source: Neuron

But what if the S&P market had followed a different path in October 2008? Suppose it started at the same point, and ended at the same point but had got there differently?

Chart 3 shows 100 plausible paths. We created them by randomly reshuffling the 23 daily price changes that occurred in October 2008. We only show 100, but there

---

[3] See several references listed below and our paper 'Rediscovering Portfolio Insurance'. We note there are other factors at play here, one being the commoditization of trend following strategies and the intense fee pressure partly driven by a flood of competition and aided by investor pressure.

[4] It is identical in form to the EMA2 model found in Martin and Zou (2012).

[5] We first calculate a volatility-adjusted price difference series. E.g. vol_adj_diff(t) = [p(t)-p(t-1)]/volatility(t). For volatility we use a standard Riskmetrics style (EWMA) model with decay of 0.94. We pass this through a EWMA twice i.e. signal=sign(EWMA(EWMA(vol_adj_diff,0.96),0.875)). The position is then signal/volatility. We assume positions are rebalanced daily. We have ignored transactions costs, slippage and brokerage. The model 'P&L' has been scaled ex-post to produce 15% per year.

are in fact 25,852,016,738,884,976,640,000 possible paths we can create by reshuffling (23 factorial). Mathematically inclined readers might recognise this way of generating sample paths as akin to a Brownian bridge, a quasi-Monte Carlo method for calculating pricing and risk for certain derivatives.

In fact, there is a connection between our work and some of the early academic attempts to understand hedge fund characteristics. In what is now a seminal paper, Fung and Hsieh (2001) showed that trend follower returns are similar to those delivered by look-back straddle options. Although look-back options don't actively trade, in the relatively small literature that does exist, Brownian bridges feature as a means of dealing with the path dependence inherent in these options.

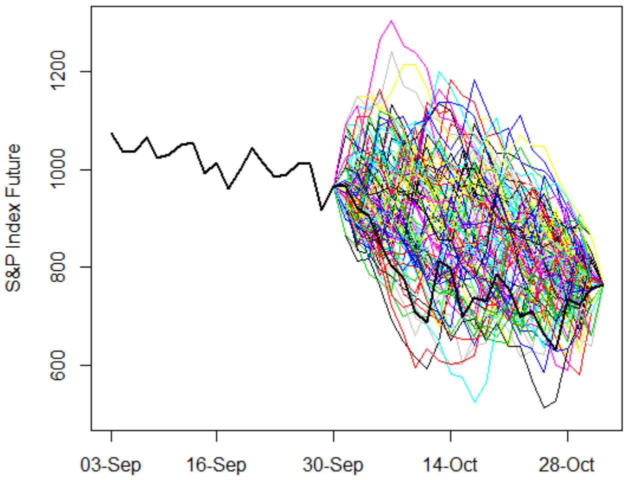

Chart 3. Various possible paths for the S&P
Source: Neuron/Bloomberg

We now evaluate our trading rule by running it over many price histories. Each shares the same price data up to the end of September 2008, but differs in how October plays out.

Chart 4 shows the histogram (and smooth density fit) of all these possible returns. Now we see that in some paths the trading rule loses money, but it is profitable in more cases than not. The range of the performance is between down 4% and up 7%. The grey dashed vertical line shows the mean return (across all paths) is just below 3%. The solid vertical line shows the return using a replay of the unshuffled path. The difference between the replay return and average return over the shuffled paths indicates that the actual path the S&P took in 2008 was favourable to our trading rule, *relative to all the other paths we examined*. This is possibly not surprising. By randomly reshuffling the daily returns we have potentially destroyed valuable (to the model) information within the price data. We will return to this point later.

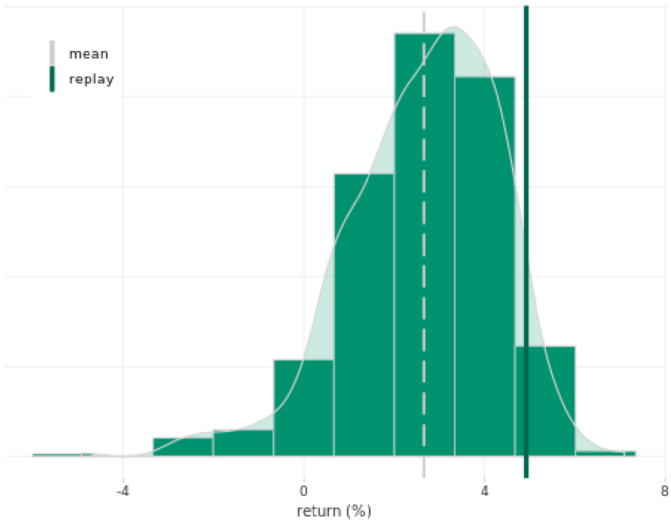

Chart 4. Sensitivity of trading rule to price path
Source: Neuron

Chart 5 examines how our toy model would have performed had it experienced the single actual October 2008 path from different starting points in the last 20 years[6]. Now we see the performance range is huge. The rule can generate anything from down 40% to up 30%. The intuition for this is simple. At the end of September 2008, the trading rule was already positioned short. It was poised to benefit from an equity market sell-off. But had it been long, it is obvious that a sudden correction would have caught the rule off-side. Some of the losses would have been considerable.

[6] We have evaluated the replay scenario at 500 equally spaced starting points between 2001 and March 2017.

So clearly the precise way in which a crisis unfolds can be important, but the starting point even more so. There are two main factors that we would expect to determine the relative importance of these uncertainties. First, the crisis horizon matters. In general, the shorter the evaluation horizon the more important is the starting point. Second, the model characteristics matter. A slow CTA that adjusts its positions slowly would be more immune to the precise path than would a faster CTA.

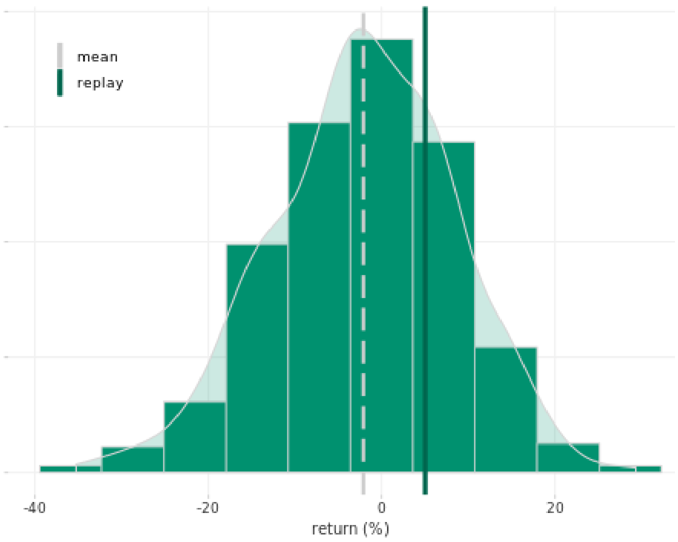

Chart 5. Sensitivity of trading rule to the starting point
Source: Neuron

A brief comment on reshuffling

There is a simple, almost trivial point we will make about the reshuffling idea. It is easy to think that by reshuffling we destroy the autocorrelation in the data. It logically follows that this action would be bad for a trend follower. But in our exercise, we can just as easily make sample paths in which we *create* autocorrelation, and thereby produce paths that are more favourable for a trend follower not less. Imagine a path that originally is +1,-1,+1,-1,+1,-1,+1,-1, and we reshuffle to get +1,+1,+1,+1,-1,-1,-1,-1. Also in many of the stress periods we are going to look at, we know that many markets underwent major moves. By reshuffling, the price is pinned at the start and the end of the path. We cannot destroy that overall trend.

Note also that in replaying historical scenarios starting from today's price levels we need to be careful about producing dubious prices paths. For example, if we allowed oil prices to drop $60 (as they did in our Lehman stress scenario) from today's prices they would go negative. If we allow interest rates to drop 400 basis points they too would go deeply negative. To avoid this we take the proportionate changes, so for oil if it dropped $60 from $120, we would assume today that oil drops from $50 to $25. We don't think there is a 'right' way to do this. A few years ago we would have probably stopped interest rates going negative. We now know that would have been wrong. In our defence we simply point out that it is very easy for us to play with other assumptions.

How relevant are these issues to CTAs?

For research purposes only, we have created a CTA proxy model that we believe offers insights into the performance of large trend-follower CTAs. It trades 80 different futures markets covering financials like equity indices and bond futures, commodities and FX. Each market's position is determined by the average of four different signals (of the same type as used in our toy S&P futures model) reflecting different time-frames. We have set it to try and deliver an annual return volatility of 15% a year. We have reflected some real-world features such as assumptions over transaction costs and slippage, and applied a 1.5% flat management fee.

Although the replica model has an 85% correlation (on monthly returns since Jan 2000 to Feb 2017, see Appendix 1 for more charts) with the SG CTA Index we used in Chart 1, we didn't try to replicate the SG Index. There is no claim of skill here. It is well known that collections of simple technical trading rules can produce simulated returns that have high correlations with benchmark trend indices, see Dao et al (2016) for a recent example, and Burghardt et al (2011) for more evidence.

The main weakness in this type of simulation exercise is the assumption about being able to trade without occurring major slippage, or indeed being able to trade at all. Given the entire purpose of the exercise is to try and evaluate the model over periods of market stress, these concerns are very real. But we have to start somewhere and the point of this exercise is not to present a specific model, but demonstrate certain features common to a lot of models.

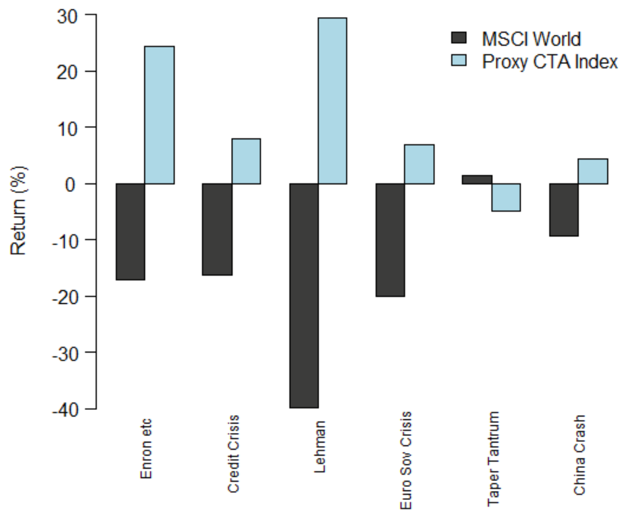

Chart 6. CTA Proxy returns during stress episodes
See footnote for the periods[7]. Source: Neuron

Chart 6 shows the typical crisis alpha style chart of our industry replica model over various stress periods. We have chosen them a little arbitrarily and included the 'Taper Tantrum' episode. See Appendix 2 for details.

Focusing on the Lehman crisis we now apply the same exercise we did with the S&P model. Chart 6 shows the CTA Proxy would have been up 29% in the simulation.

Once again, we run the model starting at the same point in time (12 Sep 2008) but now each path of the crisis is different. Chart 7 shows that in no paths did the model lose money, and in fact the worst the model generated was up 10%.

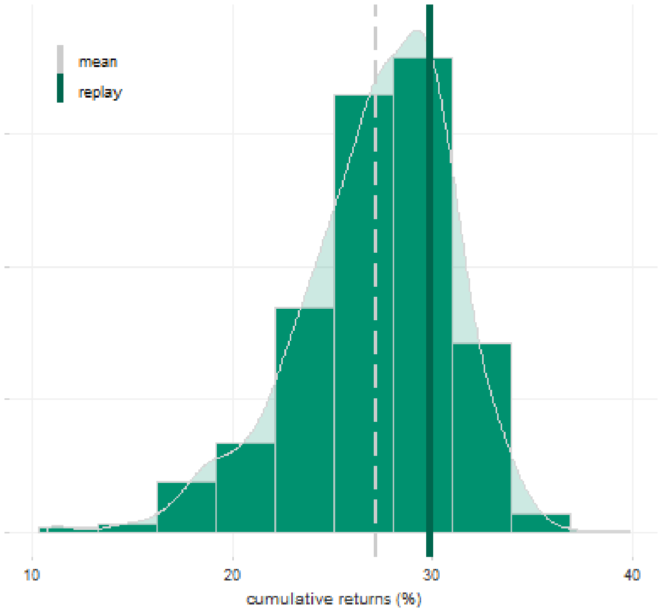

Chart 7. Stress test distribution for Lehman Crisis
As of 12 Sep 2008. Source: Neuron

The replayed performance is towards the top end of the distribution demonstrating that the actual path that prices took was quite favourable for the model.

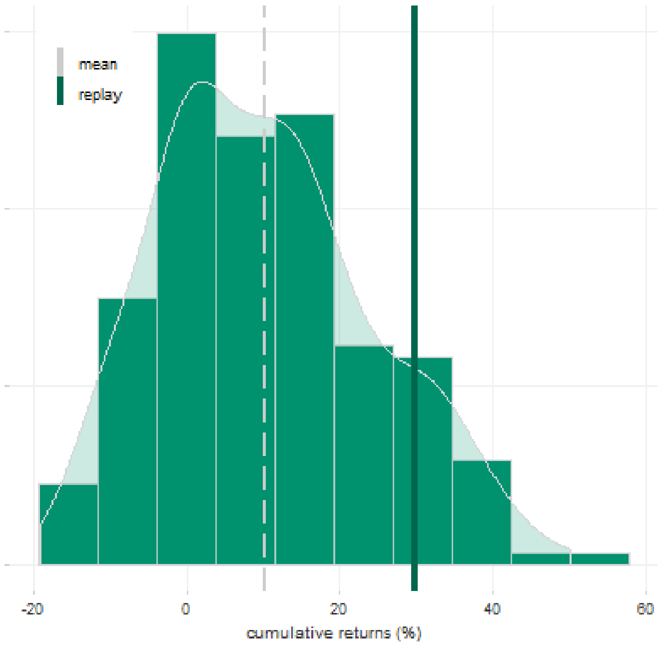

Chart 8. Stress test distribution for Lehman Crisis
if it unfolded at any point since 2000. Source Neuron.

[7] The dates for the stress periods chosen are: Enron etc (15 May 2002 to 10 Sep 2002; Credit Crisis 31 Oct 2007 to 22 Jan 2008; Lehman 15 Sep 2008 to 20 Nov 2008; Euro Sov Crisis 22 Sep 2011 to 4 Oct 2011; Taper Tantrum 01 May 2013 to 5 Sep 2013; Chinese Crash 21 May 2015 to 29 Sep 2015.

In the next exercise, in Chart 8 we explore how the model might have coped with a Lehman crisis had it begun unfolding at any point in the last 20 years. We pick 200 equally spaced starting points. Similar to the S&P exercise we see a dramatic increase in the possible range of outcomes, although the distribution appears to be positively skewed and the majority of outcomes are positive. The starting point really matters.

What is the current risk outlook for CTAs in March 2017?

We now briefly apply this idea to the current market. In the beginning of March 2017 our CTA Proxy has long positions in stocks, it is short US Treasuries and short oil. For brevity we focus on 3 scenarios. We employ our reshuffling method to get a grip on how sensitive the model might currently be to the path that prices might take in each scenario.

Case Study 1: Lehman

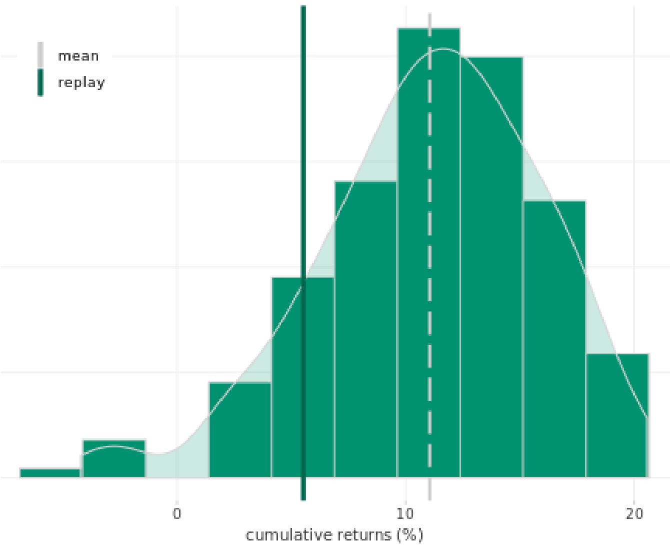

Chart 9. Stress test distribution for Lehman Crisis As of 16 March 2017. Source: Neuron

If the Lehman scenario begins replaying tomorrow the CTA proxy would be up 5% (Chart 9). Chart 7 showed us the return would have been up 29% at the time. The weaker projected performance starting today is not surprising, given the model is short US bonds and long equities. But the fact that the overwhelming number of paths do deliver positive results at the portfolio level indicates that the model may be able to adapt reasonably quickly to be able to capture some of the trends that developed late in 2008.

Drilling deeper, at a sector level we see that in no paths does the equity index sector produce a profit over the horizon. The bond sector overall produces a small profit, with gains from long positions in short maturity European bonds outweighing losses from short positions in US bonds. The energy sector does rather well because the CTA Proxy is currently short oil. During October 2008 oil fell 35%.

Case Study 2: Credit Crisis

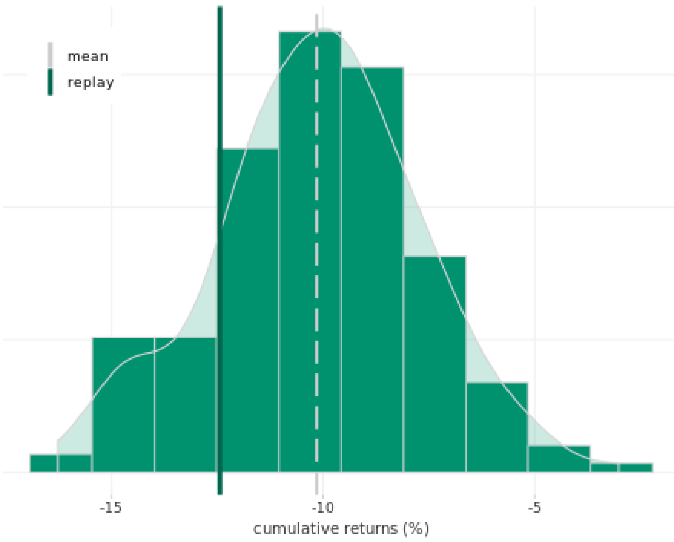

Chart 10. Stress test distribution for Credit Crisis As of 16 March 2017. Source: Neuron Advisers

The results in Chart 10 look far less rosy for the Credit Crisis stress scenario (31 October 2007 to 22 January 2008). In none of our scrambled paths does the model make money and if there is an action replay the CTA Proxy estimated return is -12%. Again, current positioning is *less* favourable than it was at the time. Our crisis alpha chart showed a return of +8%. Even though the stress scenario plays out over 3 months the exercise shows that the model would be unable to profit.

Case Study 3: European Sovereign Crisis

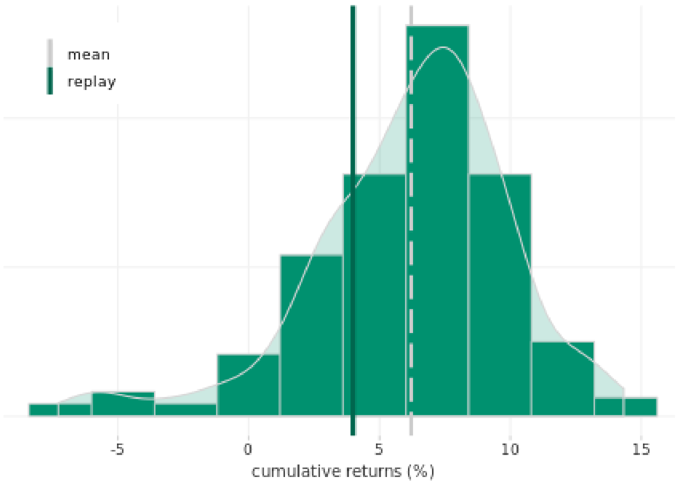

Chart 11. Stress test distribution for European Sovereign Crisis
As of 16 March 2017. Source: Neuron Advisers

Our European Sovereign Crisis covers the period 22 July to 4 October 2011. Should that replay today the projected performance is +4%. But the distribution in Chart 11 also shows that the model might do even better should it play out in a different way. The average scrambled path return is +6%. There are only a few paths in which the model loses money. Drilling deeper reveals the model is unable to make money in any of the potential equity index paths, but it makes strong gains in the bonds and the energy sector. Under the unshuffled path this scenario would see oil falling from around $50 to $35.

Concluding points

We've published this exercise in order to try and enrich the debate on the use of trend followers, particularly given the recent rise in the use of trend following to mitigate crisis risks. An increasing number of popular investment products such as strategy-indices, risk-parity and smart beta attempt to deliver a stable portfolio volatility, and in doing so exhibit a sensitivity to the path of prices. This is a risk factor that more passive and buy-and-hold strategies do not face. Within the limited confines of our experiments we hope to have brought some quantitative feel for how important these issues are.

Systematic approaches to trading are not without flaws, but one of the key strengths is that we can foresee how a portfolio will adapt over time as new market data arrives.

One practical use might be for an investment committee that has a strong view about the market outlook, and is contemplating their allocation to trend following or risk-parity. Dynamic stress testing (or more generally forward simulation) may help reveal useful information about the short-to-medium term relative expected returns from different strategies.

Finally, it is trivial to discover what types of scenarios would cause the most pain for a strategy. Knowing only the current risk exposures will offer very limited information. A reverse-stress test might help committees identify potential risk scenarios that they hadn't thought of, and that have not occurred in the past.

Appendix 1: Comparisons of the CTA Proxy with the SG Indices

SG produce two relevant indices for comparison with our CTA Proxy. The SG CTA Index that we use in the main paper, and a SG Trend Index. For full info see here https://cib.societegenerale.com/en/prime-services-indices/.

The CTA Index currently has 20 members, many of which use trend following, though this is not a necessity for inclusion in the index. The Trend Index only has 10 members but is designed to be more dedicated to trend. They are highly correlated to each other, around 97% on monthly data between Jan 2000 to Feb 2017. The CTA Proxy has a correlation of 85% to each of them. Chart A1 below shows the last 12 months returns for each index. The SG CTA Index exhibits a lower volatility (9%) than the Trend Index (14%), but all three have behaved similarly recently.

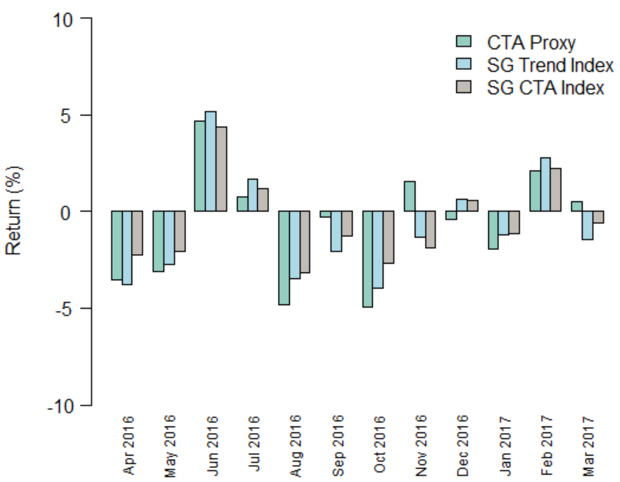

Chart A1. The CTA Proxy and two SG Indices. Source Neuron, Societe Generale.

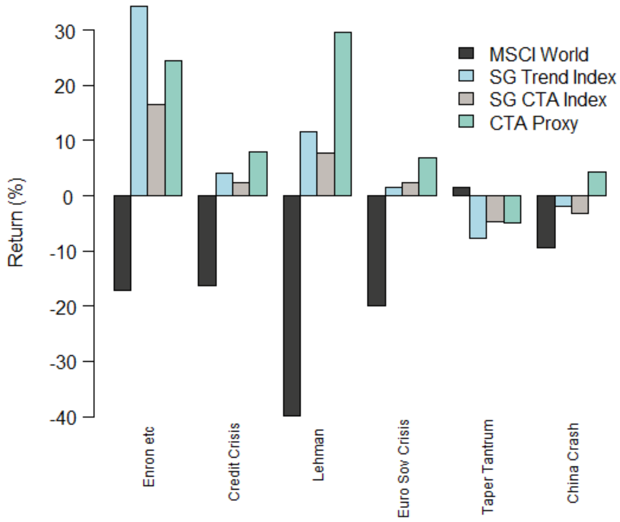

Chart A2. The CTA Proxy and two SG Indices over our stress episodes. Source Neuron, Societe Generale.

Chart A2 shows each index versus the MSCI World for our 6 stress episodes. It shows that the SG Trend index 'outperformed' the CTA Proxy in the Enron stress, but underperformed it in the Lehman episode. All three seem to behave similarly during these stress episodes. A possible exception is the China Crash, but here the returns are all low so the sign differences are not particularly exciting. Overall these charts give us confidence that our proxy model offers some insights into the behaviour of major trend followers.

Appendix 2. Market Behaviour During Our Stress Periods

Table A2. Reference Information For Market Behaviour During Each Stress Period. Source Bloomberg/Neuron.

	Enron etc	Credit Crisis	Lehman	Euro Sov Crisis	Taper Tantrum	China Crash
Start	17/05/2002	31/10/2007	15/09/2008	22/07/2011	01/05/2013	21/05/2015
End	10/09/2002	22/01/2008	20/11/2008	04/10/2011	05/09/2013	15/09/2015
# Days	82	59	48	52	91	83
S&P	-17.1%	-14.8%	-40.5%	-17.1%	3.8%	-7.2%
Nikkei	-20.7%	-25.0%	-36.9%	-15.6%	0.9%	-11.6%
Euro Stoxx 50	-25.7%	-15.3%	-31.9%	-24.6%	3.7%	-12.8%
Eurodollar	1.2%	2.0%	1.1%	-0.2%	-0.1%	-0.2%
UST	7.4%	6.5%	4.1%	5.6%	-8.2%	-0.4%
JGB	1.7%	2.0%	1.4%	0.7%	-1.1%	0.4%
WTI Oil	9.8%	-1.3%	-51.2%	-23.7%	16.0%	-23.7%
Gold	3.0%	13.0%	-2.0%	1.7%	-6.7%	-8.9%
Copper	-4.0%	-8.2%	-50.5%	-29.2%	1.8%	-14.2%
EUR	6.8%	1.2%	-11.4%	-7.7%	-0.2%	1.7%
JPY	6.9%	7.5%	11.9%	2.2%	-2.5%	0.7%
GBP	6.8%	-5.1%	-17.2%	-5.6%	0.4%	-1.2%
CHF	6.5%	5.6%	-7.3%	-11.3%	-1.3%	-3.8%

Notes: For each market we use the percent change in the future or fx price from the start to end of each period. Underlying FX rates are quotes as USD per foreign currency. To be clear, during the 'Enron etc' crisis the euro appreciated against the US Dollar by 6.8%. During 'Lehman' the euro depreciated versus the US Dollar by 11.4%.

References

Burghardt, G. and Walls, B. 'Managed futures for institutional investors'. Bloomberg. 2011.

Dao, Tung-Lam and Nguyen, Trung-Tu and Deremble, Cyril and Lemperiere, Yves and Bouchaud , Jean-Philippe and Potters, Marc, Tail Protection for Long Investors: Trend Convexity at Work (May 9, 2016). Available at SSRN: https://ssrn.com/abstract=2777657

Fung, W. and D. A. Hsieh. 2001, "The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers." Review of Financial Studies, Summer 2001, Vol. 14, No. 2, pp. 313-341. Available here https://www.grahamcapital.com/TheRiskinHedgeFundStrategies.pdf

Greyserman, A. and Kaminsky, K. (2014) "Trend Following with Managed Futures: The Search for Crisis Alpha", Wiley.

Hamill, C., Rattray, S., Hemert, O. (2016) "Trend Following: Equity and Bond Crisis Alpha". Man AHL. https://www.ahl.com/trend-following-equity-and-bond-crisis-alpha

Harding, David (2016). 'Alternative Beta, Positive Convexity, Crisis Risk Offset and Tail Risk Hedging'. Available here https://www.winton.com/Winton/files/52/528b583e-2ec1-4124-95fc-21927a6165f7.pdf

Hillman, R (July 2015) 'Don't Forget 1987 – Portfolio Insurance, Trend Following and QE' http://neuronadvisers.com/Documents/Dont%20forget%201987%20-%20portfolio%20insurance%2c%20trend%20following%20and%20QE

Hillman, R (November 2016) 'Rediscovering Portfolio Insurance' http://neuronadvisers.com/Documents/Algos%20and%20Egos%20-%20Rediscovering%20portfolio%20insurance

Kaminski, K (2011) 'Diversify Risk with Crisis Alpha' Futures Magazine. February 2011. http://www.futuresmag.com/2011/01/31/diversify-risk-crisis-alpha

Martin, R. and Zou, D. (2012) 'Momentum trading: 'skews me'. Risk Magazine. August 2012. Available here https://www.man.com/documents/download/4qJLx-CxxGb-mYDxE-3FMNN/Man_AHL_Analysis_Archive_Momentum_Strategies_offer_a_positive_point_of_skew_English_31-08-2012.pdf

Pension Consulting Alliance website, (June 2016) 'Crisis Risk Offset Class: Multiple Mandate Search' http://www.pensionconsulting.com/

Pension Consulting Alliance website, (August 2016) 'Systematic Trend Following Manager Search' http://www.pensionconsulting.com/

Pensions & Investments online (18 April 2016) 'CalSTRS Preps Risk-Mitigation Portfolio'.

San Joaquin County Employees Retirement Association (January 22, 2016) CRO Systematic Trend Following Manager Search. – Financial Meeting, San Joaquin County Employees Retirement Association Board of Retirement Agenda. http://www.sjcera.org/Pages/content/agendas/past/financial/f20160122.pdf

Disclaimer

This research note is published by Neuron Advisers LLP ('Neuron'). You have been provided with this material only upon your acceptance of these Terms and Conditions herein.

Neuron Advisers LLP is authorised and regulated by the Financial Conduct Authority, with firm reference number 563919. Neuron Advisers LLP is registered with the Commodities Futures Trading Commission under NFA ID 0439462.

Investments or investment services mentioned in this research note are not available for investment and are not being marketed in any jurisdiction. The contents of this research note are not intended to be read by any persons in any jurisdiction other than the United Kingdom.

This material is presented for information purposes only. It is intended for your personal, non-commercial use. No information or opinions contained in this material constitute a solicitation or offer by Neuron to buy or sell securities or to furnish any investment advice or service. Neuron does not provide investment advice, tax advice or legal advice through this material and you agree that this material will not be used by you for such purposes.

This material is intended as a general introduction to Neuron and the Neuron research blog by which means Neuron can express ideas and opinions. The material contained herein is the sole opinion of Neuron.

The information provided in this research note is intended for institutional investors and Professional Clients and Eligible Counterparties as defined by The Financial Conduct Authority and for those who are considered as qualified eligible persons as defined by Commodities Futures Trading Commission Regulation 4.7. It is not intended for retail investors.

The contents of this research are not intended for distribution to, or use by, any individual or entity in any jurisdiction where their distribution or use would be contrary to local law or regulation or which would subject Neuron Advisers LLP to registration with the jurisdiction. You should be aware that any rules and/or regulations applicable to providing financial services (and the resultant investor protections that may be available), may not apply to persons who obtain information from the internet and its various applications, of which this material forms part.

Neuron Advisers LLP assumes no responsibility for access to this material by any person located within a country or jurisdiction where such access would be contrary to any law or regulation in that country.

We try to ensure that the information in this research note is correct, but we do not give any express or implied warranty as to its accuracy, timeliness or completeness, nor is Neuron under any obligation to update such information. Any data supplied has not been audited and is provided for information purposes only.

We are not liable for any damages (including, without limitation, damages for loss of business or loss of profits) arising in contract, tort or otherwise from the use of or inability to use this research note, or any material contained in it, or from any action or decision taken as a result of using this research or any such material.

This research note may provide links to other research and websites. We do not control the linked sites or research and we are not responsible for the contents of any linked site or research, any link in a linked site or research, or any changes or updates to such sites or research. We provide these links to you only as a convenience, and the inclusion of any link does not imply our endorsement of the site or research.

Neuron Advisers LLP is a limited liability partnership registered in England & Wales with number OC367248 and registered office address at 2nd Floor 11 Albemarle Street London W1S 4HH.