Extreme Weather and Extreme Markets

Computer Simulation Meets Machine Learning

Robert Hillman[1]

7 September 2017, London

In this note I present some parallels between financial and weather forecasting underpinned by pressing challenges both disciplines are facing and argue that embracing computer simulation is the way forward. I give examples from weather forecasting and discuss how similar techniques are being used in finance. There are implications for those hoping to find economic value in alternative data and using data-driven machine learning.

Forecasting financial markets and forecasting the weather might seem like quite different tasks. After all, with the weather at least we have physics to guide us. Given a few equations, some atmospheric measurements and sufficient computing power it is just a question of crunching numbers isn't it? By contrast in finance it appears we can't even agree on the equations. As the old joke goes, the only function of economic forecasting is to make astrology look respectable[2].

But two recent developments in both economic and weather forecasting indicate there may be more similarities between the two tasks than at first meets the eye. Both disciplines share similar pressing challenges.

Forecasting the weather without physics

The first recent development in weather forecasting is the combination of statistical extrapolation and machine learning (ML) to make predictions over short time scales of minutes and hours[3]. The idea is easy to explain. If you track a rain cloud across the sky its path is effectively linear. You can watch the weather approaching from over the horizon, exactly as sailors have done for hundreds of years[4]. Today this process has been largely automated. Satellite imagery provides near real-time digital weather charts. Image recognition techniques help distinguish and spatially delineate weather phenomenon. Put simply, they identify where it is raining. Neural network based image recognition technology plays a role here in helping to identify the edge of rain clouds, a task that can confused by background clutter. Algorithms track movement within these images and so forecasting becomes an extrapolation task, forward in time and across space. Mobile technology puts these forecasts in the hands of the consumer. Sadly, as I write many more people are seeing the results of this technology as they track Hurricane Irma.

Apparently, there is no need for any physics at all. To improve forecasts beyond a simple linear pixel-by-pixel extrapolation we can apply machine learning (ML) to help classify types of

---

[1] Robert is CIO of Neuron Advisers LLP.

[2] According to Wikiquote this is often misattributed to J.K. Galbraith but originally belonged to Ezra Solomon.

[3] For more on this see the Met Office website, in particular http://www.metoffice.gov.uk/research/foundation/precipitation-nowcasting. For a more commercial application see the blog by DarkSky here https://blog.darksky.net/cleaning-radar-images-using-neural-nets-computer-vision/

[4] Weather forecasters call this nowcasting. In economics nowcasting also developed in recent years as researchers try to improve upon official data releases. As alternative data has become available (e.g. satellite imagery of economic activity) this process is has seen new impetus.

weather front for example, or learn local (spatial) features.

Probabilistic forecasts naturally fall out of these ML models. Instead of a binary output forecasting rain or no rain, we can get a probability that it is going to rain. Probabilities are useful when there are costs involved in a decision.

Limits to forecasting

While nowcasting can offer very short-term forecasts, there is of course a natural limit to any short-term extrapolative predictability determined by the chaotic nature of the underlying physics. As Chart 1 shows improvements in more traditional short-term forecasts (1-day ahead) have levelled off in the last few years[5].

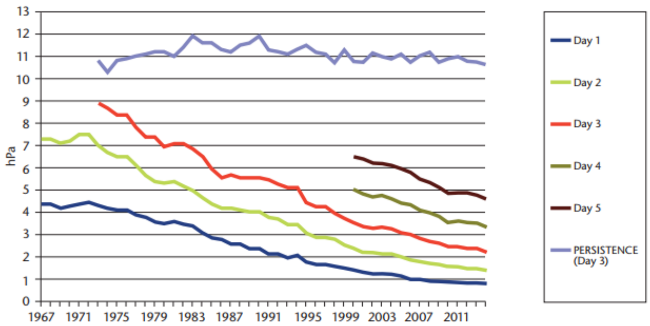

Chart 1. Short term weather forecasting improvements over the years. The chart shows the root mean square errors of sea level pressure forecasts for different horizons. The 'persistence' forecast shows the error (for a 3 days horizon) assuming the forecast is the same as todays measurement. It is analogous to the random-walk forecast seen in economic contexts. Source: Met Office (2016).

Short-term forecasting in finance

In finance, we have also seen a significant increase in data resolution, but we have not, and may never see similar public evidence of improved short-term forecasting. This is because traders who forecast financial prices affect future prices. That prices should adapt instantly to expectations about them is what underpins the efficient-markets-hypothesis. This does not mean no one can never claim success in forecasting markets, but it does suggest they had better keep it a secret.

This reflexive characteristic of financial markets – the act of forecasting the future changes the future – is not unique to finance. Proponents of climate change policy hope that their forecasts will influence behaviour and thus alter the future.

So much data, so little time

To further understand why more data in finance may not lead to forecasting improvements (or what finance people call alpha) it is worth briefly discussing the two sources of new data. Firstly, as trading has become more electronic the time series granularity of financial data has increased. Quotes and trades are time-stamped to micro-second granularity. Measuring and extrapolating data at this frequency relies on statistical techniques.

The other source of increased data resolution is alternative or Big Data. This is providing a dramatic increase in the cross-section of data. It is also providing a dramatic increase in marketing and sales efforts by data vendors and associated parties. But optimism about the potential for forecasting improvement should be tempered by the fact that the time-series of data is not expanding at anything like the cross-section. Finding the signal in the noise may be getting harder not easier. This is related to the curse of dimensionality. The training data set cannot expand fast enough to keep pace with the number of potential features hidden in the data.

In my opinion this curse is ignored, underplayed or not recognised by many of those involved in the supply of new data. I recently saw a beautifully ironic strap-line for an alternative data event. Presumably intended to impress upon potential attendees the urgency of the gold-rush for data, the strap-line read "So Much Data, So Little Time". True, but for the wrong reasons.

[5] The forecast errors in this chart here are from numerical weather prediction (NWP) models. Unlike the purely statistical nowcasting methods these use a combination of governing equations (e.g. Newton's laws) and statistical representations.

Learning or fitting?

You often hear that using ML for short-term forecasting in finance is desirable because the underlying subject matter is constantly changing. An ML approach that allows models to adapt in an ever changing world sounds great. Why wouldn't you?

I suspect it is not as well known as it might be that the word 'learning' in ML often simply means fitting. You can do learning with ML without having any sense of learning over time whatsoever. For example, if we were given a set of equity data with no knowledge of what industry sector each stock belonged, we could apply a classification algorithm to 'learn' some structure in the data. To all intents and purposes this is a static task. The word learning derives some of its use from the *gradual* numerical optimization techniques that are used in estimating the parameters for nonlinear models. Because the parameters cannot be solved for instantly (as would be possible in ordinary-least-squares linear regression) it makes sense to think of them being gradually learned. The word learning makes additional sense in the context of the cognitive inspired model architectures like neural-networks.

A real dynamic learning example would be if firms were changing their focus (and thus industry classification) over time. As time passes, more information is revealed, and the algorithm or model can improve its performance by learning more. But genuinely new data accrues at a rate constrained by the passing of time, and so dynamic learning will be slow and possibly irrelevant compared to the static learning rate from sourcing more historical data. In ML terminology this is essentially the difference between online and batch learning.

In finance the reality is that the potential rate of adaption for models that operate over times scales that are economically meaningful (> micro-seconds) is severely constrained by the rate of arrival of new data.

So to summarise this half of my note, like weather forecasting we have seen an increase in the availability of data and learning techniques in finance in the last few years. But unlike weather forecasting, we should probably be realistic that this is not likely to translate into meaningful improvements in forecasting edge, or its economic counterparty, alpha.

"Since records began…"

It is a related concern about the shortness of the history of available data that motivates me to write about a second, more recent development in weather forecasting and its links to some of today's issues in markets. The rainfall and floods experienced in the UK, in South Asia and now hurricanes Harvey and Irma are providing a timely, if disturbing backdrop.

In 2013/2014 the UK experienced worse floods than had been seen (in the classic, yet all too familiar phrase) 'since records began'. Notably "in January 2014 south east of England experienced unprecedented rainfall 30% higher than any previous January for over a century"[6].

Now this *is* a phenomenon that may feel all too familiar in finance. Despite widespread lip-service that extreme events seem to happen more often than conventional finance models predict, almost daily it is reported that something is at a historical high or historical low, or something has moved to a greater or lesser degree than it ever has before. These observations trigger Pavlovian responses from market gurus.

Fooled by randomness

In part our propensity to spot extremes and anomalies is a function of how many things there are to look at and our natural human tendency (or need) to see structure and coincidence in randomness[7]. But we are also frequently surprised by things because our memory is very short. Although meteorological records stretch back many decades, the scale and probability of extreme weather events are hard to estimate

[6] From the opening paragraph of Thompson et al (2017).

[7] Taleb (2004) literally wrote the book on this.

without a lot more data. The clue is in the word extreme.

If this wasn't problematic enough the earth's climate system today is different to the past. Over thousands of years external forcings like solar activity and variations in the earth's orbit matter. On shorter time scales the Greenhouse effect matters. So even if we could uncover more records from further back in time, we may not actually want to use them. Historical data has both a sufficiency as well as a relevancy problem.

Creating parallel universes

To address the challenge of insufficient relevant historical data the Met Office have taken to simulating their climate models (Thompson et al, 2017)[8]. Each simulation run creates a different potential realisation. Variation across simulations may come from varying different things. First, varying the initial conditions that go into the model to reflect measurement uncertainty. Second, perturbing some of the equations that are less well determined. I said rather casually at the outset that meteorologists have their physics, but in truth some bits of that physics are less well understood than other bits. For example, there is uncertainty over how quickly ice crystals form at different levels of the atmosphere.

Adding these two sources of variation into simulations help transmit meteorologists' knowledge uncertainty into forecast uncertainty.

Because of the chaotic nature of the physics that govern the weather, uncertainty over inputs can feed through into varying amounts of output uncertainty at different times. This sensitivity to initial conditions has long been known by weather forecasters, and by creating multiple forecasts, each run from slightly different starting conditions and perturbed physics, estimates of the uncertainty around short term forecasts can be delivered. Producing forecasts built on an ensemble of individual forecasts has transformed the way forecasts are formed, and how they are communicated.

A different type of alternative data

Finally, but by no means least importantly, to reflect the possibility that the climate is nonstationary, the Met Office simulations also reflect assumptions and measurements relevant to today not the past. For example, in reflecting the amount of CO_2 in the atmosphere.

Simulating the climate model many times generates a set of plausible alternative weather records. These alternative records provide a complementary data set to the single and much tortured historical path that we experienced. They call the model approach UNSEEN reflecting the nature of the data it generates[9].

Using this approach the Met Office researchers have argued that we should expect to see more weather extremes in the next few years than a simple analysis of the historical record would suggest. Figure 2 below (gratefully reproduced under license) shows an example. The grey bars show the range of monthly rainfall as contained in the historical record. The grey blob just above 200mm for January is the 2014 actual observation. It is way outside the range observed in the historical set of previous Januarys. The orange box next to it shows the range and outliers of measurements the simulations produce. It indicates that what was observed in January 2014 was within the range what might have been anticipated at the time, and that even greater rainfall can be expected.

Dealing with these uncertainties about modelling reveals a greater range of possibilities than the historical record contains. One of the headlines from the Met Office work was that there is a one in three chance of a new monthly rainfall record in at least one region each winter (Thompson et al 2017).

[8] This section draws heavily from Thompson et al (2017) but largely reflects my interpretation of their work, and in no way should be attributed to them.

[9] It also stands for UNprecedented Simulated Extremes using ENsembles.

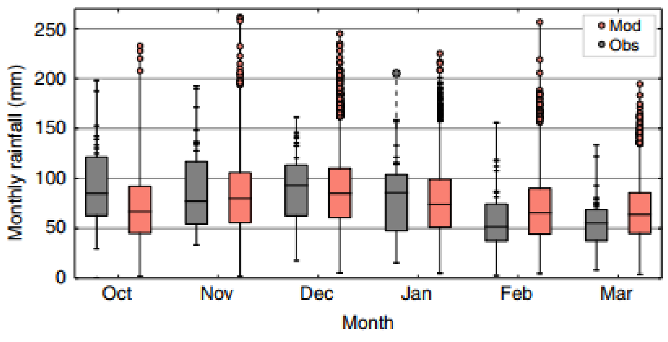

Fig. 3 Unprecedented monthly rainfall in all winter months. South east England monthly rainfall totals from observations (*grey*) and the model (*red*) for October to March. The *box* represents the interquartile range and the range of the *whiskers* represents the minimum and maximum monthly rainfall totals. *Red dots* indicate model months with greater total rainfall than has yet been observed and *ticks* on the upper observations line indicate values in the upper quartile of events. For January the *ticks* on the model line indicate months above the observed record prior to 2014 and the *grey dot* above the observations indicates the record observed monthly rainfall of January 2014

Figure 2. Climate Simulation Model Output. Source Thompson et al (2017). Reproduced under Creative Commons license: https://creativecommons.org/licenses/by/4.0/

Creating unseen financial histories

A growing area of research in economics is addressing issues that have parallels with the Met Office work. Researchers are turning to simulated market models in order to get insight into the chance and scale of risks that may have not yet been realised. The modelling approach also has the capacity to reveal risks that have not even been conceived of.

A couple of examples come from research at the Bank of England. In one study (Braun-Munziger et al, 2016) they explored how the corporate bond market might behave in the event of a redemption shock to mutual funds. It has all the hallmarks of the Met Office problem. Data on the corporate bond market only goes back a few decades at best and the relevancy of this historical data is questionable. The market has changed greatly over that period. Bond issuers have come and gone. Products and the market players have changed. We have seen a rise of mutual funds, a growth of passive index trackers and of hedge funds. The market infrastructure has changed in terms of the roles of investment banks and brokers, and in term of plumbing with the rise of exchanges and electronic trading. There is little historical data from which to try and learn.

We can however make educated guesses as to how the market works, and simulate it in a computer. We can turn to empirical measurements to paint a picture of the types and numbers of participants. For example, we can measure the size of the overall corporate bond market, and the quantity of bonds held by mutual funds versus hedge funds for example. We can glean evidence on how these participants behave by measuring how investor flows have been correlated with price moves in the past. We can posit the heuristics or explicit rules by which certain participants operate. For example, we know many traders extrapolate from past prices. Others trade with reference to some sense of fair value.

The broad label for the type of model that is designed to explore the interaction between these disaggregated microeconomic components of the markets is an agent-based-model (ABM). One way of thinking about them is by contrast to an alternative approach of modelling bond prices as driven by a few variables like GDP and inflation.

Armed with a virtual version of the real market – a computer program – we can then simulate it to create multiple histories.

A key source of the variation across these multiple histories are random shocks in the model. We need shocks because there are many things we cannot explain, model precisely, or would prefer to leave outside the model for convenience. For example, while we may have a good idea about some types of fund operate, others like hedge-funds we may be less sure about and so we 'model' them simply as generators of random orders.

Predictability may rise in crises

I believe a powerful use of ABMs is exploring how markets might behave under stress. In such times people and institutions, via codified risk management responses, become predictable. Increased connectedness propagates shocks through the financial network (Bookstaber 2017). For example, we can ask what might happen if mutual investors redeem 10% of their capital. In modelling the reaction to a redemption shock we don't necessarily have to explain how the redemption came about, we just treat it as a random shock imposed on the model. This can make the models manageable and more focused.

In a more recent paper BoE researchers explored the potential for second-round effects to amplify an initial shock to the European corporate bond market (Baranova et al 2017). Rather than a full-blown ABM the researchers use a pared down version that focuses only on two steps (that covers two weeks) that follow on after an exogenous redemption is assumed to take place. One reason they don't push the model further into the future is that as successive redemptions take place it is harder to assume what investors might do[10].

Exploring changes to the environment

The way in which shocks propagate through the system are also affected by external forcings. Once a volcano has erupted there is likely to be a local effect on weather. But this local effect may be different depending on the background level of greenhouse gases.

Similarly, a 10% redemption shock that happens when overall leverage is very high, is likely to have greater second-round effects than when it is low. Effects of other background forcings are less obvious. One of the findings from the 2016 BoE paper is stated in their abstract:

"We also explore the impact of the growth in passive investment, and find that it increases the tail risk of big yield dislocations after shocks, though, on average, volatility may be reduced."[11]

This is a topical issue. As more and more money appears to be shifting away from active towards passive trading many are asking whether this has any implications for how we might expect markets to behave.

I discussed a related concern in my 'Death-Spiral' note which asked if risk-parity and volatility-control mechanisms would at some point lead to a deleveraging driven unwind[12].

There is a fairly public spat ongoing between certain well known participants who disagree over whether these strategies pose a risk. One side says these strategies are a significant size in the market. The other side says they are not. Each side has made assumptions about the relative size of these strategies compared to average market behaviour and come to different conclusions because their inputs are different[13].

A problem with these comparative-static exercises is that they are biased to explore average behaviour. After all, it is average statistics that get fed into the calculus. They don't get directly at the behaviour that might be seen in the tails, when markets are not operating close to their average. To do this you need to examine a much wider range of paths.

We may find is that there are a just a few paths in which the nasty catastrophic dynamics develop. In highly nonlinear and stochastic systems like financial markets, path and state dependency is rife. Computer simulations seem necessary to get a grip on these issues.

But how good are these simulation models?

There is a very real question of how much faith to put into simulation models[14]. Not least there is

[10] See their discussion on page 16.
[11] I should stress this was a finding relevant to the specific model they examined and not a general statement from them on the effects of passive trading.
[12] a.k.a. 'Towards a new framework for today's markets'.
[13] See Hurst et al (2015).
[14] There is a nice parallel in how the Met Office researchers attempt to validate their model and how researchers with ABM and heterogeneous agents approach it. The Met Office use the term 'model fidelity'. They compare statistics calculated from their simulated data with real data. For example, seeing if the simulations reproduce the same mean, standard deviation, skewness and kurtosis of the real data. But this is tricky. We expect our models to potentially produce more extreme events than we have seen. We don't want things to match too closely. In finance applications there are many parameters that are not easily pinned down or calibrated. In cases such as these researchers have taken to estimating these parameters by explicitly optimising the model's ability to reproduce these empirical statistics. For econometricians this is akin to the method of simulated moments and indirect inference (Gourieroux and Monfort 1996). A newer technique

an intrinsic difficulty. In both our finance and climate examples we have reason to believe that models ought to generate behaviour we have not seen in the past. Back-testing is therefore hard.

To validate their 2017 model, the BoE researchers try and replicate the circumstances preceding the October 2008 shock, and see if their two-stage model can create the same sort of follow-through that was observed at the time[15]. But because they had calibrated and tuned their model to be relevant to today's environment they had to change some background parameters (external forcings) to adjust for the fact the world was different then. In their case the investment fund sector was smaller than now. This is like adjusting a climate model to have more CO_2 in the atmosphere today than in the past.

Forward testing is also hard. Not least because it takes time. Up until extremely recently finance ABMs have mainly seen use as proof-of-concepts: examples of models and mechanisms capable of generating the kind of volatility and behaviour we observe in real financial markets. They have not been used seriously for forecasting[16].

A role for humans

There is also a clear role for humans. Each path offers a narrative as to how different paths might develop[17]. This might be more important than a precise quantitative estimate of the probability of it occurring.

This is a positive side-effect that the Met Office researchers also claim. Using a dynamical model (as opposed to a purely statistical one) allows "investigation of the mechanisms of extreme events and remote precursors, which can aid prediction"[18].

Ultimately how much faith we put in the model output is a very human decision. A least if we have a narrative in front of us that describes how events play out, we have a chance of considering whether it seems plausible or not.

Summing up

I have argued that short data series, changing background conditions, and a need to estimate the chance and scale of extreme events has been motivating a different approach to modelling risks. I believe the key to the new approach is simulation. One could argue that this is all fine in an area like climate where we have a pretty solid understanding of the physics, but it is not relevant in economics where we have much weaker theory. I would turn this on its head. Precisely because we have weak theory as to how markets ought to operate, we should instead attempt to replicate how they do operate. Large parts of them are already in effect simply computer programs. This trend is continuing as markets become more automated and algorithmically driven[19].

As the raging debate about the risks of a deleveraging death-spiral testifies, there is genuine concern that should a metaphorical asteroid hit financial markets, the short-term reaction may well be highly predictable, and unpleasant. Simulation methods present a powerful means by which to explore this idea.

has proposed combining ML methods to help calibration (Lamperti et al, 2017). The idea is that ML can create a surrogate model that captures the behaviour of the ABM in a more efficient (computationally) way. I think of it as indirect inference on steroids. Further progress is needed but could ultimately lead to ABMs for real-time forecasting.

[15] Page 16 of the 2017 Baranova et al paper. 'Back-testing simulation outputs'.

[16] An exception is Hommes (2013). Strictly speaking he uses a heterogeneous agent model, as opposed to a (more disaggregated) agent based model. But that's a moot point, the same issues regarding forecasting affect both.

[17] The role of agent-based-models as a source of alternative narratives is a major plank of Rick Bookstaber's book, The End of Theory (2017).

[18] Thompson et al (2017) p. 4.

[19] Hillman (2017). Science Fiction Becomes Fact.

References and Sources

Baranova, Y, Coen, J, Lowe, P, Noss, J and Silvestri, L. (2017). 'Simulating stress across the financial system: the resilience of corporate bond markets and the role of investment funds'
http://www.bankofengland.co.uk/financialstability/Documents/fpc/fspapers/fs_paper42.pdf

Bookstaber, Richard. (2017), 'The End of Theory: Financial Crises, the Failure of Economics, and the Sweep of Human Interaction', Princeton University Press.

Braun-Munzinger, K, Liu, Z and Turrell, A (2016), 'An agent-based model of dynamics in corporate bond trading', Bank of England Staff Working Paper No 592, available at
http://www.bankofengland.co.uk/research/Pages/workingpapers/2016/swp592.aspx

Franke, R and Westerhoff, F 'Structural stochastic stability in asset pricing dynamics: Estimation and model contest' Journal of Economic Dynamics and Control 36(8): 1193-1211.

Gourieroux,C and Monfort, A (1996) 'Simulation-based econometric methods'. CORE Lectures. Oxford.

Hillman, Robert (2016) 'Science Fiction becomes Fact' Available here
http://www.neuronadvisers.com/Documents/Algos%20and%20Egos%20-%20Science%20Fiction%20Becomes%20Fact

Hillman, Robert (2017) 'Why investors need agent-based-models' Available here
http://www.neuronadvisers.com/Documents/Why%20Investors%20need%20agent%20based%20models

Hillman, Robert (2017) 'Towards a new framework for today's markets' Available here
http://www.neuronadvisers.com/Documents/Towards%20a%20new%20framework%20for%20todays%20markets

Hommes, Cars (2013). 'Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems'. Cambridge University Press.

Hurst, B., Huss, J., Mendelson, M., and Ooi, Y. (2015) 'Dog Bites Man: In August, Equity Selling in Risk Parity Was a Tiny Fraction of Market Volume'
https://www.aqr.com/library/aqr-publications/dog-bites-man-in-august-equity-selling-in-risk-parity-was-a-tiny-fraction-of-market-volume

Lamperti, F, Roventini, A, Sani, A. (2017) 'Agent-based model calibration using machine learning surrogates'
http://www.amirsani.com/files/ABMML.pdf

Met Office (2016) 'Met Office Science Strategy: 2016 – 2021. Delivering Science with impact'.
http://www.metoffice.gov.uk/binaries/content/assets/mohippo/pdf/migrated/met_office_science_strategy_2016-2021.compressed.pdf

Taleb, N. (2004) 'Fooled by randomness'. Penguin.

Thompson, V., Dunston, N., Scaife, A., Smith, D., Slingo, J., Brown, S., and Belcher, S. (2017) 'High risk of unprecedented UK rainfall in the current climate' Nature Communications. 8:107.
https://www.nature.com/articles/s41467-017-00275-3

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.