Algos and Egos – Science Fiction Becomes Fact

Robert Hillman, 15 November 2016

This is part of a series of notes on a rapidly developing theme in financial markets and investing – namely the collision of algos and egos. I will cover issues like the revival of portfolio insurance, the replacement of star-traders by automated trading strategies, the disillusionment of institutional investors with active managers, and the hopes pinned on machine learning and technology.

The rapid growth of algorithms in the world of finance is giving regulators, economists and investment professionals plenty to think about. They may need to look toward approaches more in common with Minecraft than with traditional methods of economic analysis.

In the six years since the May 2010 'flash crash', multiple theories have been proposed as to the source of that and subsequent crashes. No simple explanations are readily available partly because conventional research methodologies are not well suited to analysing these events. Economists typically take a two pronged approach to analysis, collecting and analysing historical data, and building toy models of the situation at hand. But while there is no shortage of data in terms of sheer volume of information (because the events take place at such high-frequency that there can be millions of records per minute), there are very few distinct events to study and from which to generalise.

In terms of building models, economists are well versed in modelling human decision making, but few have until recently considered the implications of non-human trading that takes place so fast that European regulators have had to recently propose a synchronization of clocks to within a billionth of a second in order to reduce ambiguity about the order and sequence of trading orders[1].

With conventional methodology falling short, some economists are looking towards an approach developed in the late 80s and 90s in an effort to build an alternative way to study financial markets via the simulation of people and institutions (often called 'agents') with computer programs. Often labelled 'agent-based computational economics' or ACE[2], it built on 1950s work by pioneers like Herbert Simon who studied human behaviour as computational processes and vice versa. These models offered new explanations for concerning phenomena like bubbles and crashes, but they had little impact on economics or investment management professions at the time[3].

Do I really have to work with 'them'?

In a recent paper Andy Haldane at the Bank of England bemoans the lack of willingness to engage in interdisciplinary research as a factor in why economics has been a reluctant adopter of ACE. Structures do not help either. The evolution of university academia into subject-silos fighting over funding doesn't make interdisciplinary research easy[4]. But this doesn't explain

[1] See ESMA MifID II guidelines, link in references.
[2] I will use the terms ACE and ABM fairly interchangeably.
[3] For expediency purposes (ok, laziness) I will provide few references except where specifically significant, and instead point the interested reader to Leigh Tesfatsion's encyclopaedic resource on all things ACE. http://www2.econ.iastate.edu/tesfatsi/amulmark.htm. Andy Haldane's paper 'The Dappled World' also contains more macro and policy related references. Cars Hommes (2013) book gives a slightly different approach based on heterogeneous agent models.
[4] There are many exceptions of course but in general it feels interdisciplinary research is the exception rather than the rule. The Santa Fe institute in the US, in which ACE was accelerated in the 1990s, was a successful attempt to tackle this issue by building the entire institute on the foundation of interdisciplinary research.

why ACE hasn't taken off outside of universities in commercial environments where profit drives resource allocations. In fact, in the very environment in which ACE was originally focussed – that of modelling the non-normal dynamics of financial markets – we see economists, physicists, statisticians and engineers working together almost as a rule rather than as an exception. There's often the odd classicist knocking around for good measure too. This environment is that of hedge funds and the trading and research arms of investment banks and large investment managers.

So the question goes beyond one of blaming reluctance or difficulties in collaboration across disciplines. In hedge funds resources move quickly to what works and what makes money. So why haven't we seen ACE take-off in the real world of investment management?

Great expectations

The expectation that ACE would by now be playing a major role in investment management was articulately made by one of the discipline's founders. In 2001 Doyne Farmer wrote in an article entitled 'Towards Agent Based Models of Investment' that:

> *"When practical use of agent-based models becomes possible (perhaps within the next five years), their effectiveness will cause securities prices to change."*

Farmer's first sentence is 'As far as I know, no one uses agent-based models for investment'. It is 15 years since he wrote it. As far as I know still no one uses agent based models for investment. Actually that's not quite true, Sonia Schulenberg has pushed this field for many years now[5]. But to put it another way, in discussions with hundreds of small and large investors in recent years, almost without exception if I mention the term 'agent-based' the response is a blank look[6]. To be fair, this is not the only subject with which I am able to provoke blank looks, but it is probably the most reliable. In 2016 agent-based models for investing – for all practical purposes - do not exist.

Science fiction has become fact

One of the reasons why economists were initially reluctant to adopt is because ACE approximates human investor behaviour by modelling agents as simple algorithms, and twenty years ago economic research that lacked so called micro-foundations was deemed 'invalid' for want of a less-polite term. But twenty years of market evolution has brought us to a situation where today real markets closely resemble the algorithmic markets of ACE.

The last few years have seen the spread of algorithms right across the investment management landscape. Artificial markets where traders follow price trends, and others seek patterns in historical data describes exactly the growing sector of machine learning and systematic rules-based trading. Fund managers may as well be described as designers and guardians of rules-based strategies, as opposed to the egos and big-swinging-dicks of old. To me, the proliferation of exchange traded products alongside the commoditization of dynamic portfolio management techniques like risk parity, factor investing and smart beta are only accelerating this trend.

Extrapolate to accumulate

Recent research suggests that the use of simple extrapolation techniques and heuristics is rife among investors more widely[7]. There is no criticism left that artificial worlds are unrealistic because their 'agents' are somehow irrational or non-optimizing because like it or not these agents are conspicuously present in today's market. But maybe part of the reason why hedge funds

[5] See https://www.schulenburgcapital.com/index for example.
[6] Or the response 'Asian based?'. 'No, Agent-based. Let's move on..'.

[7] See the work by Greenwood, Barberis and Li in the references.

haven't yet discovered value in agent-based models is that the shift in the real trading world towards one more sympathetic to a computational environment has not been around for that long. The acceleration towards automated trading products has really only gotten underway in the last five years. And of course there was a small matter of a financial crisis (the sort quite easily generated by ACE models by the way, although hard to produce from orthodox economic models) along the way that has meant there might have been higher priorities in the research groups of hedge funds and investment banks than fundamental research into modelling methodologies.

Forecasting is hard especially about the future[8]

Another reason why ACE has not taken off in economics or investment management is down to the criteria on which we judge success. In economics, like many sciences, the ability to predict the future has been elevated above almost everything else. You can understand why. In a world of little data but competing theories, one of the most attractive ways of assessing and discriminating between theories is by testing them out in the real world. Can they predict something that is outside of what is already baked into the theory?

This perfectly laudable aim has however had unintended consequences. The desire to relentlessly measure and quantify and test has become self-perpetuating, and the requirement to publish 'significant' results demonstrating forecasting ability has become an end in itself. Unfortunately, a consequence of this is that reporting biases and exaggerated claims of statistical significance are rife in published research.

So with forecasting ability as the ultimate criterion of success, the easiest criticism to level at ACE is that while it is excellent at showing us how interesting economic phenomena *can* arise, until now ACE models have failed to provide forecasts that indicate when interesting phenomena *will* arrive. This pushes it immediately into the 'interesting but irrelevant' file, and meanwhile other methods like econometric and statistical models, or even heuristic techniques like technical analysis dominate the industry[9].

Alpha male seeking alpha

Investment management presents a similar hurdle. In a structurally alpha-male world, it is far more alpha to be seen to be searching for alpha than it is to be managing beta. And in investment management alpha is measured as the incremental forecasting ability (as measured by excess return) of a strategy or trader over and above some benchmark model. Or simply does it make money. If investors only assess ACE in terms of its ability to produce better return forecasts, then it is no wonder it hasn't got people excited. Producing better forecasts is not likely to be where ACE excels. And better is very hard to demonstrate in anything but the extreme long term.

The reasons forecasting from ACE models is hard and that the results are not likely to set the world alight is because our ability to forecast financial markets is inherently limited. But our egos refuse to accept this and drive us to keep searching for the holy grail of predictive edge[10].

ACE models demonstrate the limitations of predictability in two key ways. Firstly, ACE models show that the interaction of simple algorithms can lead to complex dynamics. In many complex systems only short-term forecasts are achievable, and even they are prone to being occasionally catastrophically wrong. Secondly, ACE models teach us that macro level behaviour can be

---

[8] A variant on a term generally attributed to Niels Bohr although it has a distinct tinge of Yogi Berra.

[9] There are a few attempts to use agent-based-like-models to forecast. For example, the heterogeneous agent models in Hommes (2013). These models effectively collapse to a reduced-form that looks like a nonlinear econometric model like a regime-switching model or a smooth-transition-autoregressive model. These models definitely have promise but by their nature they are likely to outperform simpler models only occasionally. Most forecast testing metrics tend to be global rather than local, so can wash out the episodic value of nonlinear models.

[10] As indeed does our ability, almost willingness to be fooled by randomness (Taleb 2001).

impossible to predict from an analysis of the components in isolation – the whole is greater than the sum of the parts.

Big data to the rescue

But there is an exciting development which could provide fresh impetus to ACE, and that is the advent of big data. Big data means different things to different people, but one aspect of it is that new data is becoming available at both a micro scale (disaggregated) and on an almost real-time basis. Big data means more algorithms in the investment process because of its scale and often unstructured nature.

And big data might also help alleviate some of the limits to forecasting that complex systems present. In most complex systems short term predictability is affected by the accuracy with which we can measure the state of the system at the point of making the forecast. For chaotic systems it is critically dependent on the accuracy of those measurements.

Big data could help here. In principle it means we can get closer to calibrating ACE models to reflect real world detail and on a timely basis. It is important to understand that these possibilities big data give us are not just nice-to-haves but are critically necessary given the characteristics of complex systems like financial markets.

This ability to calibrate agent-based investment models from disaggregated data and in near-real-time is new, we simply didn't have a chance to do this even five years ago. This suggests an intriguing possibility. It is plausible we are in the early stages of a fast-accelerating process where more data begets more algorithms which begets the increasing viability and relevance of ACE.[11]

Econometricians look away now

The second reason why ACE may become increasingly relevant is that we have entered an era where our egos are becoming more accepting of the idea of radical uncertainty[12], and ACE are a means of generating artificial worlds in which radical uncertainty is ubiquitous.

What I mean by radical uncertainty is the type of uncertainty that means it is futile trying to describe the statistical distribution of financial markets by probability distributions, and where it is futile to begin modelling by assuming that the data we observe has been produced by a data-generating-process (DGP) and our task is mainly one of 'discovering' the DGP, usually a simple equation and probability distribution. At this point econometricians reading this will be jumping up in the air (to be fair not many will have got this far) crying foul! So in slightly more formal terms what I mean is that the data we observe is subject to such high degrees of non-stationarity and non-constancy, stable statistical models are pretty useless[13].

One way of demonstrating that the pursuit of a stable DGP is futile is to use an agent-based model to generate data and then proceed as if it is real data. Note that by construction there is no single equation driven by a random error process involved here. The data produced by an ABM is unlikely to be characterised by a simple equation because it is the output of a complex system of interacting agents and randomness. Yes, it might be possible to describe certain broad brush features of the data by the use of quantitative measures, but it will be

[11] I hesitate to call this a virtuous process because I have severe doubts as to the welfare benefits of the increasing use of algorithms, but that is a subject for another day.

[12] There are many related concepts. King (2016) introduces the term with the sentence 'Radical uncertainty refers to the uncertainty so profound that it is impossible to represent the future in terms of a knowable and exhaustive list of outcomes to which we can attach probabilities'. A related concept is Knightian uncertainty.

[13] I am being a bit unfair but only a bit. Yes, there are many models that attempt to allow for nonlinearity and regimes breaks and so on, but most of these still ask far too much from the data in terms of ability to produce reliable parameter estimates. I consider myself an econometrician and accept my share of guilt. I think the analogy that an economist behaves like someone who has dropped their keys down a dark alley but walks to the end of the road to look for them under the nearest street lamp is extremely apt.

next to impossible to estimate a simple model for forecasting, based on price data alone.

A stark lesson from carrying out this sort of exercise is how hard it is to accurately estimate parameters given the short sample histories we have in the real world. Experience with ABM shines a spotlight on our ignorance. Only when you can simulate millions of market realisations do you realise how inadvisable it would be read too much into just one path. Yet that is of course what every investor is forced to do when presented with a back-test or even an audited track record.

Protecting ourselves from ourselves

To further understand why taking a more computational approach to modelling financial markets is sensible it is helpful to think about why (besides big data) algorithms are becoming more and more influential in our investment processes.

One explanation is the growing self-awareness that behavioural biases and psychology influence decision making. Some investment algorithms can be seen as mechanisms to protect oneself from one's own flaws, and others can be seen as designed to exploit the systematic mistakes of others. On the former the systematic trading industry has long explained a major benefit of its approach is to take out the emotion and provide consistency of process. Some managers may simply explain it as a means of avoiding the risks that might emanate from your trader (or indeed yourself) turning up to work late with a hangover.

Exploiting the mistakes of others

On the exploitation of the systematic mistakes of others, the economics profession has been active in influencing the investment industry. Academic finance has proposed tens if not hundreds of 'anomalies' produced by these behavioural effects[14]. There is a well-trodden path here, particularly in US academia. Having discovered a statistical regularity that beats the market, it is pronounced an anomaly, and a behavioural explanation is readily supplied. Am I cynical and suspicious of this process? Yes[15]. But there is now a burgeoning industry offering investors a smorgasbord of 'smart beta' exposures.

Product innovation and knowledge dissemination has led to something of a democratization of these strategies that were previously the secret-sauce of hedge funds. The more this happens the more the trading environment resembles a giant computer simulation.

Coping is hard enough

There is another line of argument that supports the idea that humans behave like algorithms, and thus supports the use of ABMs to study humans.

Gerd Gigerenzer (2014) and Daniel Kahneman (2011) each provide slightly different angles on the idea that humans do and should appeal to algorithmic ways of thinking. It is also present in Mervyn King's recent (2016) book in his proposal that we should think more about coping-strategies as opposed to optimized rules of behaviour. And it was present in Richard Bookstaber's (2007) work on coarse-behaviour rules wonderfully described in his book 'A Demon Of Our Own Design'. The common thread is that we can't and shouldn't bother trying to optimise rules of behaviour too precisely, and it is better to follow simpler rules of thumb instead.

One reason offered as to we shouldn't optimise is because we can't because there are no stable relationships and probability distribution parameters to estimate. Secondly, we shouldn't because when we do we harm our longer-term ability to survive. The presence of inherently unpredictable regime shifts and sudden changes in circumstances means that optimising to the

[14] Notwithstanding my earlier comments about the spurious significance of data-driven anomalies. Cards on table - I have major doubts as to the strength and resilience of many of the known anomalies we hear about.

[15] Related to this but slightly off topic is the idea that we should exploit people's propensity to make mistakes, and 'nudge' them in the right direction.

present leaves us fatally exposed to changes in the environment[16].

Blurred lines

To recap, so far I have argued that ACE is likely to become increasingly relevant because markets have evolved to resemble giant computer simulations. As big data, and mass digitization proceeds, the more this process becomes self-perpetuating.

If this process continues soon the distinction between a model and reality will be completely blurred. In the past models were seen as simple abstractions of reality, often involving just a few equations that captured the main features of the problem at hand. Like building a toy version of a car before the full-scale version. In investment management, models have tended to be things like econometric forecasting equations, or no-arbitrage pricing models linking different financial products[17]. But today models can actually be full-blown *parallel* copies of the real thing. And the 'real thing' for investment management is the financial market itself. This suggests a fundamentally different way of thinking about models within the research process. It is more akin to what is happening in warfare where computer gamers and drone pilots face situations that are to all intents and purposes indistinguishable. You could put a drone pilot in front of a screen and they could have no clue as to whether they are engaged in the real world or a virtual world. You could put a trader or an investment committee in front of trading screens and data sets and they too could be none the wiser[18].

Is machine learning going against the trend?

Although no one in investment management is really talking about ACE yet, almost everyone wants to talk about machine learning[19]. Like big data, machine learning means different things to different people. But one example of its use in financial markets often heard is the idea that machines can discover patterns in data that humans can't see and adapt to a fast evolving environment. Sounds great doesn't it?

Unfortunately, this idea that machine learning will help us discover self-adaptive trading strategies goes against the grain of most of what I have discussed. If Bookstaber and King are correct, then trying to look for quickly adapting strategies could be precisely the wrong thing to be aiming for in a world of radical uncertainty. And the other red flag that ACE waves is that the solitary historical path that we have at our disposal is virtually meaningless in representing the true potential risks we face. Purely historical-data-driven modelling is therefore limited to the point of being futile if not ultimately self-defeating. Big data needs machine learning. But machine learning needs to be performed with an ACE environment in mind in order to keep it on the straight and narrow.

What we can do today

Some are taking ACE seriously again, in part because policymakers have asked them to. Perhaps the best example is the request by Trichet for more investment in ACE, in a speech given in November 2010.

"[In] the face of crisis, we felt abandoned by conventional tools. ... The key lesson ... is the danger of relying on a single tool, methodology or paradigm. The atomistic, optimising agents underlying existing models do not capture behaviour during a crisis period. Agent-based modelling ... allows for more complex interactions between agents. ... we need to better integrate the crucial role played by the financial system into our macroscopic models."

But ABMs are seeing interest from more commercial purposes as well. To end with three examples.

[16] There is a huge wider literature here on related concepts like bounded rationality, satisficing, and robust control. And there is a growing literature on how to forecast under regime shifts and unstable probability distributions.

[17] The capital asset pricing model for example.

[18] Andrew Lo (2010) and co-authors have explored related concepts in a fascinating paper offering a kind of Financial Turing test.

[19] Hillman (2015).

Researchers have looked to ACE to try and understand the dynamics of market liquidity. For example, within artificial markets, regulators can explore what types of algorithms, external shocks and exchange rules might be prone to generate flash crash type behaviour[20]. Exchanges can explore the effects of mechanisms such as circuit breakers, and on the other side of the fence, investors can explore what the effect of circuit breakers might be on their operations. For these experiments, computer simulation is not merely conveniently aligned to the reality of modern trading technology, it is vital.

In another example research at the Bank of England (Braun-Munziger et al, 2016) has explored a calibrated ABM to understand liquidity in the corporate bond market. They have been able to ask what the effect of a change in the redemption policy of funds might be on the underlying market. This type of exercise has practical implications not only for regulators but also for the funds themselves. This same approach could be of use in thinking about the liquidity terms of trend follower strategies, a strategy seeing renewed investor interest as US institutional investors and public pension fund seek ways to mitigate their exposure to falling equity markets[21]. Competitive pressures have forced down notice periods from monthly to daily across many different types of funds, and in the last few years a new sector called liquid alternatives has emerged.

This example of exploring the effect of different fund liquidity policies demonstrates perfectly why ACE is potentially so useful and historical data analysis alone is so limited. We have not lived through an environment in which hedge funds and other actively managed funds have existed with the liquidity terms they have today.

Learning by doing

Finally, another area in which ACE can contribute is in the use of simulation to help train and build experience. Think of how pilots continuously train on flight simulators. Redington (2016) recently argued that pension fund trustees could benefit from engaging in experimental situations in which they individually or collectively practise making decisions. ABM models are ideally suited to providing the simulation environment. But this simulation approach to help humans gain experiences can be taken a step further. Why not use ABMs to help algorithms themselves gain experience?

While this may sound crazy and the start of some sort of infinite regress, this very same principle was behind the success of DeepMind's AlphaGo in beating Lee Sedol in Go in March 2016. One of the key technological breakthroughs was to combine supervised learning from human expert games with rounds of reinforcement learning in which the computer in effect played itself. The use of simulated games was critical in boosting the intelligence of the computer program. Without this simulated experience the potential intelligence is constrained by the limited quantity of actual historical games on which it could learn. More generally, the use of machines to create the same kind of data on which they are attempting to learn from marked a crucial step in the recent revival of interest in machine learning and AI. To abuse the title of a seminal paper in this field, to invest safely in markets, first learn how to generate markets[22].

Maybe in the next five years…?

To summarise, there are a host of new reasons to think that ABM is about to see a resurgence in economics and finance and finally begin to influence the practise of investment management. Big data, computational advances, the rise of algorithms throughout financial markets, and the growing acceptance of radical uncertainty are four intertwined factors combining to create the conditions for this revival.

What accelerated the interest in deep learning ten years ago were spectacular breakthroughs in image, speech and text recognition. It is unlikely we will see such tangible

[20] See the article by Battiston et al (2016), and Bookstaber and Paddrik (2015). Bookstaber explained his attraction toward ABMs in a NY Times article in 2013 (Norris, 2013).

[21] See Hillman (November, 2016)

[22] As exemplified in Hinton (2007).

breakthroughs with ABM in finance. And people looking to see a world beating AI hedge fund appear will probably be disappointed. But more likely within the practice of risk management, execution and exchange technology, the design and implementation of automated strategies, and in financial education we will see ABM-like systems emerge as the core modelling platform.

References

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. (2015). 'X-Capm: An Extrapolative Capital Asset Pricing Model'. Journal of Financial Economics 115 (1): 1-24 http://scholar.harvard.edu/files/shleifer/files/x-capm_published.pdf

Battiston et al, (19 Feb 2016) 'Complexity theory and financial regulation'. Science. Vol 351, Issue 6275, pp. 818-819.

Bookstaber, Richard (2007). 'A Demon Of Our Own Design'. Wiley.

Bookstaber, Richard, Paddrik, Mark, (2015) 'An Agent Based Model for Crisis Liquidity Dynamics'. Office of Financial Research Working Paper 15-18.

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

ESMA (2016) European Securities and Markets Authority, 'Guidelines on transaction reporting, order record keeping and clock synchronization'

https://www.esma.europa.eu/sites/default/files/library/20161451_final_report_on_guidelines_mifid_ii_transaction_reporting.pdf

Farmer, Doyne (2001), 'Toward Agent-Based Models for Investment' AIMR. Downloaded from http://tuvalu.santafe.edu/~jdf/papers/aimr.pdf

Gigerenzer, Gerd. (2014). 'Risk Savvy'. Penguin.

Greenwood, Robin, and Andrei Shleifer. (2014). 'Expectations of Returns and Expected Returns'. Review of Financial Studies 27 (3): 714-746. http://scholar.harvard.edu/shleifer/publications/expectations-Returns-And-Expected-Returns

Hasanhodzic,J., Lo, An., Viola, E. (2010) 'Is It Real, or Is It Randomized?: A Financial Turing Test'. Downloaded from https://arxiv.org/pdf/1002.4592v1.pdf

Haldane, Andrew (2016) 'The dappled World' Speech given 10 November 2016. http://www.bankofengland.co.uk/publications/Pages/speeches/2016/937.aspx

Hillman, Robert. (June 2015). 'Machine Learning in Finance – 25 Years On'. Neuron Advisers. http://neuronadvisers.com/Documents/Machine%20Learning%20In%20Finance%20-%2025%20Years%20On

Hillman, Robert (13 November 2016). 'Algos & Egos – Rediscovering Portfolio Insurance'. Neuron Advisers. http://neuronadvisers.com/Documents/Algos%20and%20Egos%20-%20Rediscovering%20portfolio%20insurance

Hinton, G. E., (2007) 'To recognize shapes, first learn to generate images', in P. Cisek, T. Drew and J. Kalaska (Eds.) Computational Neuroscience: Theoretical Insights into Brain Function. Elsevier. http://www.cs.toronto.edu/~hinton/

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

Kahneman, Daniel. (2011). 'Thinking Fast and Slow'. Farrer, Straus and Giroux.

LeBaron, Blake (2014) 'Heterogeneous Agents and Long Horizon features of Asset Prices' http://people.brandeis.edu/~blebaron/wps/longrange.pdf

Li, Jiacui . 'Extrapolative Return Expectation and Trend Chasing Behavior in Commodity Markets' (December 15, 2013) mimeo. http://www.jiacui-li.com/uploads/1/0/5/7/10578946/expretcmdty_rfs.pdf

King, Mervyn. (2016) 'The End of Alchemy'. Little, Brown.

Norris, Floyd. (10 Jan 2013). ' Clouds Seen in Regulators' Crystal Ball for Banks'

Redington Ampersand Institute. (15 June 2016). 'Four seasons in one day: Learn to make better decisions' http://ampersand.redington.co.uk/2016/06/15/four-seasons-one-day-learn-make-better-decisions/

Taleb, N. (2001) 'Fooled By Randomness'. Random House

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.