Algos and Egos

MANAGING RISK THROUGH HUMAN GUIDANCE OF AI

Robert Hillman and Antonio Del Favero[1]

3 February 2017, London

The success of DeepMind's AlphaGo and the power of AI as an analytical tool have raised the possibility of achieving better results by outsourcing risk management to non-human algorithms. In this note we evaluate the comparative strengths and weaknesses of human and machine intelligence in the investment process, how they can work most effectively together, and the challenges and opportunities faced by organisations embracing these innovations.

In 2014 a book was published called "Computational Intelligence Techniques for Trading and Investment". A look in the index offers no reference to deep learning or to Big Data. But by 2016 it felt like you could not pick up a lifestyle magazine without being told how deep learning and Big Data are poised to change our lives forever. This is no criticism of the book but rather a reflection of how recently these terms have entered into the investment world.

We see there being two associated trends, artificial intelligence (AI) and Big Data. Big Data is easier to get a grip of and is simply the availability of much more data than we have had before. In thinking about Big Data it is helpful to divide it into two parts. Data on things that existed before but was not easily measurable, and data on brand new things. Examples of the former are satellite images enabling remote sensing of both oil below ground, and that stored above ground, or of human activity across time and space as recorded by mobile phone data. An example of brand new data are tweets and other social media output. Oil has always existed, tweets are new. Measuring the economy with increasing granularity and closer to real-time has obvious attractions analogous to increasing the granularity of atmospheric measurements that led to improved short-term weather forecasting. Unlike weather forecasting however, we do not have a reliable set of equations into which we can plug microeconomic data. Instead we need to combine human expertise and data science. The hope for investors who can access and process such data before others is a profitable information edge.

The arrival of data on brand new phenomena is more challenging. There is a lot of it, it is messy, and we do not really know how it might help us do anything useful. So we need to use computational techniques such as machine learning to convert it into something useful and to propose actions we can take with it. There is a role for *artificial* intelligence in this process because we do not know what and if any structures might be hidden in the data. Big Data therefore has a need of, and will foster further improvements in, AI and machine learning (ML).

AI in a broader sense is hard to define precisely and means different things to different people. For the purposes of framing a discussion around risk management and investing we are thinking about the use of computational techniques that process data with little or no human guidance. Popular examples tend to have two features. Firstly, there are AI systems that try and produce actions based on the discovery of patterns in data, and secondly there is often an emphasis on the fact these systems are not static. The idea is that systems adapt over time. This addresses a common belief that

[1] Robert is CIO of Neuron Advisers LLP. Antonio is Senior Portfolio Strategist at Rubrics and has been collaborating with Robert on research. This note was written as a background paper for the U.S. Institute CIO Management Roundtable 'A World Disrupted' in New York, February 8-9th 2017.

because financial markets are always evolving, so should be the ways by which we engage with markets. This sounds eminently sensible but we will return later to discuss why it may be hard, if not foolish, to elevate this adaptiveness goal too high.

The rise and rise of systematic trading

Nearly all of the AI trading applications we hear about tend to be associated with a systematic approach to trading. By this we mean implementing trading rules with little or no human day-to-day contact. By construction, if an AI system has discovered something valuable it will have done so by demonstrating success when tested – systematically - on past data. If the system can be seen to work in this artificial (back-test) setting it is natural to implement it in a real-time setting. So, somewhat obviously, the process of model discovery and systematic implementation are linked.

So before discussing some of the ways in which AI can aid and abet a risk manager we will briefly cover some of the things we know about risk management and systematic trading. Systematic trading is on the up, and many traditionally discretionary asset managers are looking to incorporate it. We do not think it will be long before it will make little sense to speak of firms being either systematic or discretionary. Industry definitions are lagging this shift.

Skin in the game

To think about how risk management might work in a systematic firm consider first a discretionary firm. A typical set-up would comprise portfolio managers supported by analysts, middle-office and back-office staff. On a day-to-day basis a risk manager ought to be independent and analyse the risks in the business as a whole. On a lower frequency basis, the risk manager might well engage with the CIO or PMs in helping to frame limits and controls which form part of the day-to-day process. In this set up the PM provides much of the risk management themselves. The more linked is their remuneration to the P&L, the stronger this connection. Put simply, skin-in-the-game is a powerful source of devolved risk management.

Risk management in a systematic business

In a systematic business things tend to work differently. To make our points we begin with an extreme position and imagine an investment process has been completely systematized and can run hands-free all the way from data collection and processing through to execution, trade-booking reconciliation and reporting. What role does risk management, and more broadly humans, play?

There may be a CIO or figure-head, but it is unusual for there to be a star-trader. Instead the organizational structure is much more likely to be a collegiate research team who collectively work on 'the model'. There can be specialization with scientists focusing on different parts of the process e.g. data collection, cleaning and management, signal generation, risk management, model validation and research. Peer review is used to obtain quality control and this acts as a form of pre-emptive risk management.

People are tied together via a common infrastructure, sharing code, data and computational resources provided by the business. By contrast many discretionary PM shops charge infrastructure costs back directly to traders. This enforces skin in the game but can mean an organisation misses out on the public good of a pooled infrastructure. No one wants to subsidise anyone else. R&D investment may be less than optimal which restricts innovation and makes such firms vulnerable to technological advances.

There is a trade-off here. While a systematic firm may prioritise R&D, the collegiate organisational culture may not align the interests of individuals with the P&L as directly as in a pure multi-PM shop. The result may be little sense of "ownership" and an inability to benefit from the powerful risk-management incentives that stem from having skin-in-the-game[2].

[2] These problems of ownership and responsibility are by no means exclusive to systematic funds, and crop up in any team environment. What we find interesting is the almost paradoxical result that the more a company moves to automation the more it may have to think about human relationship issues.

Keeping humans in the loop

On the other hand, detachment from P&L can have advantages. After all, one of the most commonly cited arguments for using a systematic process is precisely that is removes emotion and personal issues from the investment process. Systematic processes are reliable, repetitive, and consistent. As it is sometimes put, systematic models won't turn up to work with a hangover. The challenge for business managers is to be able to balance the advantages of taking humans out of the investment process, without removing them so far that they don't care about the process and can't fix it when the unpredictable happens.

This balancing act occurs in other situations where automation has taken over much of the human role. A good analogy is with flying an aircraft. You want the pilot to be there in case of some unpredictable failure. If it was predictable you would have probably built in a systematic response. But you don't want the pilot actually flying the plane because in almost all situations the computer can do it more effectively. As the head of a London systematic fund execution desk put it "99.9% of the time I want a robot. 0.1% of the time I want a sh*t-buster". There is another role for humans here too which is to provide a sense of security. An apt historical analogy is with the introduction of elevators, described wonderfully by Planet Money (2015). During the early years there were accidents and people died. It took a while for sufficient safety features to emerge. But the elevator operator persisted for decades, long after the safety issues were solved. In reality they performed little more than create assurance.

Temptations of risk managers

Although risk managers may be heavily involved in helping to set overall risk parameters, controls and limits, practical situations can arise that are hard to solve. One classic problem is when there is a powerful urge to intervene in the model. This can occur if there is a loss of faith in the ability of the model to work. Sometimes this loss of faith is warranted. An example would be when a central bank announces a change to a policy (like the SNB who announced a cap on the Swiss franc in January 2015). This might induce an entirely sensible response to "turn the model off", although the intervention process would not stop there. Decisions would need to be made about knock-on effects, perhaps the reallocation of risk for example. Protocols and procedures can greatly alleviate the stresses that these surprises can produce[3].

Trapped by your own process

Sometimes a trader may find themselves suffering from the opposite problem, when they are trapped by their own process and frozen, unable to make even a simple intervention. This can happen if the culture of being "systematic" is too strong, and doesn't admit the possibility of failure or the need for flexibility. This is when faith in a systematic approach to trading turns to dogma. The outside view that a risk manager detached from the modelling can offer is a valuable counterbalancing factor in these situations.

Another problem is that there may not be an easy way to practically implement a risk management intervention, regardless of whether everyone involved buys into it or not. For example, if trading is completely automated, it could actually be quite difficult to close out a position manually without destabilising the overall portfolio risk allocations, especially if driven by an optimization process. You do not want to wait for the implementation team to release a new model that copes with the change in process. So processes need escape routes and flexibility built-in, to borrow a term from software development practice, agility. Too strong a faith in the process can result in too rigid a system.

Irreversibility problems and algorithm aversion

A systematic business can also suffer from a different type of rigidity best explained by a hypothetical example. Imagine in a discretionary fund a PM has a bad run. They may have had their risk reduced and a discussion with the risk manager. It might be suggested

[3] Many of these types of practical issues and solutions are discussed by Robert Carver in his blog http://qoppac.blogspot.co.uk/ and book (2015).

some time off is in order, and the trader takes a break. This kind of process does not tend to apply to an algorithm. When an algorithm or trading model goes through a bad patch it is much harder to respond with a decision to give it a break for a while. The net result of this is that poorly performing strategies might be allowed to survive longer than they should. And on the flipside, strategies that have been "retired" rarely see the light of day again.

It is almost paradoxical that traders might find it harder to turn off an algorithm than to fire or "offer" a human trader a break. But there is evidence that as humans we exhibit strange behavioural biases towards algorithms. While we may bring algorithms into the investment process in order to help overcome our own behavioural weaknesses we need to be aware that we may react differently, and surprisingly counterintuitively to evidence produced by an algo than from a human, see Dietworst et al work on algorithm-aversion (2015).

AI in the investment process

As mentioned earlier a lot of the discussion of AI today is based around the idea of being able to build systems that learn to adapt to an evolving financial ecosystem.

Adaptive strategies sound attractive but they are likely to cause problems for a risk manager. One problem is that by their nature they assume that the "optimal" model is changing, or at least they have to allow for this possibility. This makes them prone to fitting the recent past. This is particularly problematic when the system has been discovered by a data-driven process in which there has been little or no human insight. These systems used to be pejoratively known as black-boxes. Cynics (like ourselves) sense a certain amount of rebranding going on. The problems of black-boxes are legion. Firstly, it will be hard to know whether the pattern is just luck or real. Secondly, without interpretation it will be hard to know if the pattern will persist or what it is linked to. Thirdly, it will be very hard to have a conversation with an investor in the event the strategy fails, especially if the manager has no idea why. The first wave of enthusiasm into AI investing in the 1990s fell flat on its face when it became apparent early claims of success were not robust out-of-sample (Hillman 2015).

Inadequacy of standard risk measures

Standard measures of risk are inadequate, and possibly dangerous when applied to machine learnt strategies. To give a simple but contemporary example our research suggests that in certain US equity markets short-term buy-the-dip strategies have been profitable in recent years. A machine learning process looking for profitable price patterns will have identified this, and an adaptive portfolio will end up moving risk away from more poorly behaving strategies towards this new one. How does a risk manager assess the risk in this strategy?

The traditional approach would be to calculate the value-at-risk[4], or some variant thereof, by measuring the potential losses that might have occurred over the last few years. But we know these losses will be low. They must be because otherwise the AI/ML system wouldn't have selected the new strategy in the first place. So immediately there is an intrinsic conflict. When dealing with adaptively selected strategies traditional measures of market risk which use recent data to measure vulnerability are likely to be biased downwards. It is akin to a form of survivorship bias.

What the risk manager could do it test the strategy against a much wider set of data, looking further back in the past for example, expanding the historical data set. But here is where politics come in. The AI system's advocate would likely argue that this is irrelevant because the system would not have selected the model back in the past in which it might have not worked. It is a vicious circle.

To see how these problems might arise imagine the hypothetical situation.

Risk Manager: When we back-tested your strategy it had a Sharpe of 3 and in October 2008 it made 12%, didn't you call it crisis-alpha or something? We've just had a similar month, volatility has jumped and the S&P

[4] There are many criticisms with value-at-risk. But our critique is robust to any metric that only relies on recent data.

is down 17%, but your model is down, what's gone wrong?

Trader: It was the wrong type of volatility.

Risk Manager: Go on…

Trader: Well my system's performance is path dependent. The S&P can drop 17% in many different ways. In fact, if I reshuffle the 23 daily price changes we saw in Oct 2008 I can get 25,852,016,738,884,976,640,000 paths that all end up with the S&P down 17%[5]. When I redo my back-test running over every permutation I find the range of performance is quite wide.

Risk Manager: How wide?

Trader: It can be anywhere between being down 3% and up 15%.

Risk Manager: But you were down 10%!

Trader: Ah yes, well that's because my system isn't just path dependent, it's state dependent too.

Risk Manager: It's certainly in a bit of a state I'll give you that, but what do you mean?

Trader: Well just as an October 2008-like scenario can play out in many different ways, my system was positioned very differently a few weeks ago to how it was in September 2008. So I have re-done my back-test but this time I ran through every possible path that the S&P might take like before, but now starting at every possible point in the last 20 years. It took some time I can tell you but it proves the performance of my system is well within expectations!

Risk Manager: And pray tell what are those expectations?

Trader: I would expect it to be deliver anything between -19% and up 31%.

Risk Manager: So what you're saying is that the fact your back-test showed a positive performance in October 08 tells us virtually nothing about its likely performance in the future in a similar situation? I think we need to talk……..

Warnings about the perils of adaption

These problems of path and state dependency of rules-based trading systems arise for static systems. When the system itself evolves it is not hard to see how it gets harder to glean much from an analysis of performance over past data. The analogy to strategic games like AlphaGo is straightforward. In that case the rules of the game are fixed. But what if the rules are constantly changing? The frequency over which the rules change versus the system's learning speed will determine how successfully the algorithm can adapt.

The problem of over-adaption has become recognised across a number of areas in recent years. In the investment industry more broadly people have been steering away from complex portfolio optimization techniques for some time. More narrowly within the systematic industry the conventional wisdom is that it is better to balance exposures across different types of strategies over time than chase the holy grail of an adaptive approach.

Rick Bookstaber described an evolutionary biological analogy of the risks of over-optimization or specialization in his book "A Demon of Our Own Design". Gigerenzer has argued that the way humans have evolved is directly at odds with an optimizing approach, preferring instead simple heuristics[6]. Mervyn King, ex-Governor of the Bank of England, has also stressed heuristics and the folly in thinking risk can be precisely measured[7]. Instead we should think of a much wider sense of "radical uncertainty". And within AI itself, and its less sexy related disciplines like statistical forecasting and econometrics, the trend is towards solving prediction problems by combining simple models that may not be "optimal" in a conditional right-here-right-now sense, but are far more robust in the event of sudden and unanticipated regime change.

[5] We are using some poetic licence here. It is unlikely the trader will be able to articulate this number.

[6] Gigerenzer, Gerd. (2014). 'Risk Savvy'.

[7] King, Mervyn. (2016) 'The End of Alchemy'.

We admit that this problem is in principle no different to that one faced by humans. Why should AI machines be any worse than humans in developing robust strategies? In our experience financial market dynamics change because of factors that can be hard to quantify, even after the event, let alone before.

Teaching algorithms to behave

Many challenges will arise around governance and compliance risks. To give an example in the recent success of AI applications in poker, computers learnt how to beat players who were bluffing, and to bluff themselves. This was a form of behaviour that some had thought unlikely to emerge, but it did. In a trading application it is highly probable that an AI machine would learn behaviours that are equally cunning but judged unsavoury by our own standards. An example would be the illegal practice of spoofing, the submission of fake orders in order to create a false sense of liquidity. It is illegal when there is no intent to honour the orders hence regulators would need a way of examining and verifying this intent. It would not be wise to try and counteract this activity by a blanket ban on the cancelling of orders.

There is a conflict here because we know that in human learning contexts, playing, bluffing, and making mistakes are very much part of the process. It is possible to push machine versus human learning analogies too far but the equivalent situation might be like trying to educate a child while enforcing a zero tolerance for mistakes and banning play. That is going to be quite hard work for all concerned. Humans will need to think of ways of reflecting these legal and ethical concerns in machine learning processes.

Building laboratories

Before unleashing a poorly educated AI system onto financial markets we will need to find ways of training them in safe environments. The solitary history of market prices we have is inadequate, so we need to find other means of creating fantasy but credible data. As it happens the creation of data has been seen to help accelerate learning itself, and has been behind many of the recent successes in AI. In a seminal paper Geoffrey Hinton wrote (2007) that "To recognise shapes, first learn to generate images'. The idea is simple and is in tune with how children learn, and also carries over into the more strategic learning context. Take AlphaGo and poker. The machine can generate games and play itself, thereby bootstrapping its own intelligence.

We believe these more generative approaches to implementing AI will become more popular in investing, and have a more permanent impact on the industry than the get-rich-quick alpha promising headlines suggest. Risk managers would benefit from better simulation environments. As argued previously it is not enough to take the current set of strategies and look at their performance over recent data. Risk estimates are likely to be biased downwards. Instead risk managers need tools that allow the entire *process* to be tested, and this requires much more than the solitary history. This recognition that the broader research process needs to be taken account of when evaluating the reliability of models, has recently seen attention within the systematic industry. For example, see Bailey et al (2015).

The generation of artificial data can be done in numerous ways, and in just the same ways a risk manager today develops relevant stresses and scenarios, they will need to apply judgement and experience in guiding the production of these artificial test and risk measurement environments.

Human experience has a lot to add. For example, how many of us have seen back-tests showing a Sharpe ratio of 3 only to see the strategy perform abysmally in the real world. Often this is because the back-test performance is inflated because it misses out awkward inconveniences such as the fact spreads might have been extremely wide in the past, there was insufficient volume at the time, or that there were no dealers picking up the phone.

Taking informed risk

Part of the risk manager's job in an AI driven fund should be to try and identify the source of risks identified by AI. Risk-premia for example have been used to explain the existence of many market anomalies picked up by systematic strategies and increasingly commoditized. Profit opportunities emanating from demand and supply imbalances often have regulatory or structural drivers at their source[8]. These factors can appear and disappear without warning. Again, the risk manager should seek to identify the underlying drivers of returns, even when, or perhaps especially when, they have been detected by a machine learning process that is silent on such matters.

Generating and presenting future risks

Besides relying on experienced risk managers another way of generating future risk scenarios is via the use of computer simulations of markets like agent-based-models[9]. By simulating the interaction of automated traders within an electronic marketplace (which is exactly how many real markets can today be described) a risk manager can analyse unlimited amounts of data from which to flesh out the characteristics of machine-learnt trading strategies. The scale of data that these systems can generate is massive and so machine learning methods will be needed. The increased granularity of economic measurements that Big Data offers potentially gives us better ways of calibrating these models to the real world than has been possible up to now.

Learning reflexive risk systems

Risk systems also need to be much more reflexive or in simple terms 'alive'. A proper test would explore how a strategy might influence the market it is operating in. This is particularly true for shorter-term strategies that are liquidity-sensitive. And it can be true for slower strategies via systemic risk. It is possible (and commonplace) to stress-test strategies by altering assumptions over costs and spreads and so on, but we believe that for many strategies the path and state dependencies discussed earlier will swamp these sensitivities.

So what we really need are systems with which we can explore the wider risks of strategies, in a forward-looking environment, taking into account the interaction of the trading strategy with the market it is trading in. And we need ways of interrogating, exploring and presenting the risks to a level of succinctness that will allow humans to make decisions. In our view this is where the next and more permanent coalescence of AI and investing will take place, rather than on the alpha side.

One of the most powerful benefits of using more forward looking approaches is that it should help reduce surprises, and thus more closely align investor and manager expectations. It is the unanticipated surprise that causes stress and sometimes forces suboptimal responses within organisations or by investors. Sole reliance on historical data to measure risk almost guarantees that surprises will be inevitable. Standard ways of designing risk scenarios are prone to behavioural biases, leaving us likely to be solving the last few crises not looking to the next. More efforts will need to be made in presenting and communicating these wider risks. At the end of the day the survivability of an investment business, be it ego based or algo based, comes down to human perceptions and relationships.

---

[8] For example, distortions in relative bond prices created by QE. Imbalances such as these may have no historical precedent and could disappear in an instant.

[9] See Hillman (2016), and references Battiston (2016, with a regulatory angle), Bookstaber and Paddrik (2015) and Braun-Munziger et al (2016).

References

Bailey, David H., Borwein, Jonathan M., Lopez de Prado, Marcos, and Zhu, Qiji Jim (2015). 'The Probability of Backtest Overfitting'. Journal of Computational Finance (Risk Journals), Forthcoming, 2015. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2326253

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

Carver, Robert (2015). Systematic Trading. Harriman House.

Dietvorst, B., Simmons, J., Massey, C. 'Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err', Journal of Experimental Pyschology, 2014. http://opim.wharton.upenn.edu/risk/library/WPAF201410-AlgorthimAversion-Dietvorst-Simmons-Massey.pdf

Dunis, C., et al. (2014) 'Computational Intelligence Techniques for Trading and Investment', Routledge.

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

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

Hillman, Robert. (November 2016). 'Science Fiction Becomes Fact'. Neuron Advisers. http://neuronadvisers.com/Documents/Algos%20and%20Egos%20-%20Science%20Fiction%20Becomes%20Fact

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

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/

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

Planet Money. (2015) Podcast 'The Big Red Button' Episode 642. 29 July 2015.

http://www.npr.org/sections/money/2015/07/29/427467598/episode-642-the-big-red-button

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.

 Any questions about the contents of this material should be directed to: enquiries@neuronadvisers.com.