2019 Toronto Refinitiv Summit, 2nd December 2019

Behaving Fairly: Artificial Intelligence and Conduct in Wholesale Markets

Remarks by Mark Yallop, Chair FMSB

Good morning

Two interesting events caught my eye this past month.

The Bank of England and the Financial Conduct Authority in London published a survey showing how financial services firms in the UK use machine learning.

And Lee Se Dol announced his retirement.

For those of you who missed this event, and its significance, Lee is the 18-times world champion Go player and the only human ever to beat the AlphaGo program developed by DeepMind, a Google sister company and specialist in artificial intelligence.

The sophistication and complexity of the ancient Chinese board game Go have long made it a target for AI developers: it has about 250^{150} possible sequences of moves on a 19 x 19 board and is, in the words of one observer: "*a challenging decision-making task in an intractable search space with an optimal solution so complex it appears infeasible to directly approximate*."

AlphaGo is the standout Go-playing machine, having defeated other Go-playing computers more than 99.8% of the time in machine-to-machine matches. AlphaGo studied 5 million games and played a further 20 million games against itself to train for the matches against Lee. While Lee won one match against AlphaGo, he lost the series 4-1 to the machine, in a landmark moment for artificial intelligence. That was in 2016. Last week, explaining his retirement, he declared: *"I am not at the top even if I am the number one. There is an entity that cannot be defeated."*

Bear in mind, as you think about this, two further facts: AlphaGo progressed from playing at the level of a novice child to crushing the world champion in less than two years; and further, AlphaGo is not hand-crafted to evaluate only the rules of Go – rather, it is trained directly by game-playing, using very efficient *general purpose* learning techniques which clearly can and will be applied in many other situations.

I don’t think these two events are connected. But they did cause me to think:

- How much is artificial intelligence being used in wholesale markets – the part of the financial system that FMSB is interested in; and

- What might be the implications of the use of artificial intelligence in wholesale markets, if this is, or might become, widespread?

Of course, we know that automated trading is not new.

Rules-based automated trading – so called algorithmic trading – has been prevalent in electronic equity markets for well over two decades – and it is now well-entrenched, and growing rapidly, in fixed income products.

Algorithmic trading has created new opportunities.

- speeding up execution of orders, cutting costs and increasing volumes - and at least the appearance of liquidity;

- catalysing changes in market structure, enabling firms with the best technology to compete with banks who have historically had a stranglehold on information about supply and demand.

Algorithmic trading has also introduced new hazards.

- latency arbitrage - the process whereby algos that can get faster access to matching engines gain information advantage over slower algos and manual traders;

- the occasional “flash crashes” attract a lot of attention; and

- trading algorithms can and have been programmed to manipulate markets.

But rules-based algorithms only codify and automate fixed trading strategies that humans have been using for centuries. They haven’t fundamentally changed the nature of trading or markets.

True machine learning, of the sort demonstrated by AlphaGo, is different.

Machine learning algorithms are not simple static, deterministic rules-based trading engines, forced to drive only along the rails laid by their human creator-programmers.

Rather they are algorithms that, using neural networks and other "deep learning" techniques, access to massive data sets and enormous computational power, are able to recognize patterns, train themselves, optimize in unique ways, and make decisions about when and how to trade without being explicitly programmed by a human.

It is this transformative type of machine learning that I want to talk about today.

How widespread is Machine Learning in Wholesale Markets?

Back to that Bank of England survey.

About half the 50 banks and capital markets firms in the sample use machine learning today; and most firms expect to increase their use very significantly – some by up to 3 times – in the next few years.

In the wholesale arena, machine learning is most commonly used in second-line functions such as anti-money laundering, fraud detection and credit risk management.

This may be due to the fact that pattern recognition and natural language processing capabilities are some of the most developed machine learning techniques and are well-suited to compliance, risk management and other second line activities.

Some firms are also deploying machine learning in first-line activities such as trade pricing and execution:

- to increase speed and accuracy of processing orders – for example by using natural language processing to decipher requests from clients, speeding up response times;
- to combine very large numbers of market data time-series for pricing and to evaluate venue, timing and order size choices;
- to calculate the probability of orders being filled given the characteristics of the order and prevailing market conditions;
- to determine order routing logic, including the evaluation of venue, broker and execution algorithms, as well as optimal timing, price and size of particular orders.

In summary, the picture in 2019 is:

- despite long-time use of rules-based algorithms for trading, and growing use of machine learning in second-line control functions, machine learning is *not yet* widely deployed in first-line trading functions in wholesale markets;

- *but this is likely to change materially* in the next few years.

I think the unstated implication of this survey is this: standing on the verge of a breakthrough in machine learning in markets, and the many benefits this may deliver, we have an opportunity and the responsibility to anticipate and mitigate potential future risks with the technology.

And – given the problems uncovered in wholesale markets in recent years - the important question is: do we think machine learning will help, or hinder, the quest for fairer and more effective markets?

What risks might wider adoption of machine learning in wholesale markets bring?

Machine learning amplifies some risks that are already familiar from technology developments in the past 20 years. For example, traders will need to gather, cleanse, store and sort very large amounts of data to train their optimization engines.

Clever suppliers of data will anticipate this demand and likely increase the price and reduce the supply of data for users, possibly quite significantly. Indeed, the rush to build "data driven businesses" that we see in finance today anticipates this very trend; and it will only be accentuated by the importance to machine learning of new forms of unstructured data.

The data warehouses created to hold all this data will be vulnerable to cyber threats; and the consequences of a successful cyber-attack may be far greater than we have allowed so far.

Human biases will likely creep into the data and programming of machine learning engines in the same way that they already can with deterministic, rules-based algorithmic machines.

All these we already know about.

But machine learning will also create new risks that we haven't had to consider before.

Let me focus on just four of these, which I think are the most important.

The first is **model drift** – which is an inevitable result of the continuous lifecycle of machine learning – and which doesn't feature with traditional linear, rules-based algorithms.

- When trading engines are powered by machine learning, the relationship between data inputs and price outputs is much more obscure.

  The very low signal-to-noise ratio in the data, combined with the very large amounts of data that are mined and the complex often multi-layered decision trees that underlie the machine learning models themselves, mean that price formation is inevitably opaque.

  Unlike deterministic rules-based algos, where price formation is always performed in the same way with set inputs and steps, machine learning trading engines "learn for themselves" how to create prices by repeated, constantly evolving, experimentation and it becomes very hard, or impossible, to trace how "decisions" are made in the optimization process;

- The difficulty of tracing how "decisions" have been "made" by the machine make it very difficult to prevent in advance, or to correct afterwards, undesirable model outcomes. For example, the machine may discover complex, non-linear "hidden" correlations that it is difficult or impossible for the programmer to anticipate or discover;

- And it is impossible to predict how a machine, trained on known historical data but "making its own decisions" will react when it is live in the market with a much larger dataset and it encounters events that haven't been seen before in the data that was used to train it.

Concerns about these transparency problems lie behind the present regulatory focus on explainability, model risk management and software validation as well as how management and Boards can satisfy themselves that they understand, at some level, what is going on inside the "black box".

The second hazard is **bias**:

- Ultimately, machine learning is all about discrimination, and unpredictable discrimination during the optimization process, when an enormously wide range of factors are analysed, may carry a greater risk of unforeseeable, harmful or biased outcomes, particularly when using unstructured as well as structured data – for example a mixture of pricing, satellite images and social media.

These biases could include unexpected or unfair changes in pricing or liquidity to certain types of market users, or even to individual customers, as a result of factors that are impossible to uncover because they lie effectively undiscoverable in the heart of the optimization engine;

- Another type of bias also needs to be considered: the risk that a machine optimizing on its own will "discover" that unethical, manipulative trading practices are more profitable than ethical trading.

  Indeed this is virtually a certain outcome, if the machine doesn't have an "ethical governor" that tests the optimization process against ethical benchmarks and rejects trading tactics that fall short of these standards.

  Of course these ethical benchmarks are much more complex to describe than formal laws and regulations; and developing reliable "governor" functions may be the trickiest part of the whole process.

The third risk is **market concentration and correlation**:

- The rules-based algorithmic trading developed in the past couple of decades have fostered competition, allowing non-bank market makers and traders to develop successful businesses and grow market share at the expense of banks.

  But we know that network effects create winner-takes-all monopolistic or oligopolistic market structures. So, the way in which machine learning models improve by accessing more data is likely to create data network effects which may well in turn create very high barriers to entry for new firms.

- It remains to be seen whether these barriers will entrench the power of today's large banks and financial services firms or, alternatively, allow technology-based competitors to create new oligopolies at the expense of today's financial sector.

  Either way, unless they are carefully managed, concentrated market structures may disadvantage market users by unfair rationing of liquidity, skewed pricing, and other non-price based discriminatory barriers.

- As algos optimize big data from new sources, they may well inadvertently increase, or create new, correlations between macro-economic or other input variables. Hungry algos will over time arbitrage the profit potential – a machine learning version of the "crowded trade" phenomenon – but they may also make markets more fragile to unforeseen shocks and more interconnected, as multiple users depend on a limited number of underlying data relationships.

Such vulnerabilities may be further exacerbated by the risk that bad actors could manipulate the underlying data streams to create false signals from which their own trading strategies would profit.

The fourth challenge is **resources**. There is a big skills gap for the expert programmers, data scientists and risk managers who can safely develop, test and implement machine learning in financial markets:

- At a personal level, these skills are in short supply in the private financial services sector and among central banks and market regulators. And they contribute to a quite significant knowledge gap among senior management, in the boardrooms of financial services firms and at policy makers about the hazards of AI. This knowledge gap needs to be filled, soon.

- At a geopolitical level, the resources that do exist sit overwhelmingly in China and the US. These two countries are far ahead of the rest of the world in terms of artificial intelligence capabilities, both see leadership in AI as a vital, strategic, national priority and both are investing heavily, and rapidly extending their lead.

  They have the capability to dominate the machine learning business in wholesale markets if they so choose. And the implications of this for market structure, regulation and the operation, fairness and effectiveness of global capital markets have hardly been considered yet. This must also be addressed.

How can these challenges be addressed?

The FICC Markets standards Board was established in 2016 to address risks in wholesale markets. We are a private sector, global, body whose goal is to raise standards of behaviour and conduct in markets with the goal of making them fairer and more effective for users across the world.

The UK authorities who called for FMSB to be created, and have been strong advocates for our work since then, recognized that wholesale markets pose some special challenges for regulators: the global scope, the information asymmetry between private sector and regulators, the pace of innovation and product development, among other factors, make it very difficult for regulation to stay ahead of determined private sector firms.

Instead it is much more powerful to engage the private sector in the process and encourage it take responsibility for identifying and fixing problems in market behaviour and structure that damage the functioning and reputation of markets and their businesses. Indeed, they saw that rebuilding trust in wholesale markets required the private sector to be seen to take a lead in reinforcing orderly, fair and effective markets.

With 50 firms as members today, 14 recommendations on market behaviour published over the past 3 years and 5 more in preparation, this is what FMSB has been doing.

Regulation - and not only financial services regulation - has an essential role to play in what happens as machine learning is deployed in financial markets. New legislation may be required to create a framework for the safe exploitation of the huge opportunities that machine learning offers.

The challenges posed by concentrated market structure, for example, can probably only be addressed through competition policy and law, informed by public policy considerations.

But other challenges – transparency, explainability, model risk management, governance, bias and correlation - that I described earlier cannot be solved by regulation alone. Scarcity of skills, information asymmetry, cross-jurisdictional trading and the hectic pace of change all make this clear.

Machine learning increases both the technical and the structural complexity of markets; and it will be essential that private sector expertise, risk management and controls keep pace.

FMSB has a critical role to play in these areas. We have already done work on governance for algorithmic trading and the operation of electronic trading venues; and we will shortly publish work on the role that data plays in wholesale markets. Machine learning will be a central theme of our work over the next 2-3 years, as it will be for the industry at large.

I believe we can and will make a very significant contribution to the safe deployment and realization of the huge benefits that machine learning can, and should, deliver for the users of financial markets. And I hope that we can do this with many of you in the room here today.

Ladies and gentlemen, thank you for your attention.