3rd International Workshop on Behavioural Financial Regulation and Policy

Behaving Fairly: Artificial Intelligence and Conduct in Wholesale Markets

Bank of Italy, Rome, 28th November 2019

Remarks by Mark Yallop, Chair FMSB

This paper discusses the use of Machine Learning ("ML") by wholesale financial market participants today and the major risks that may arise from wider adoption of ML across wholesale markets in future. It focusses specifically on whether ML will help or hinder the quest for fairer and more effective markets, how new types of risk that are different from those already identified in recent years may develop in automated market structures and describes the work that FMSB is undertaking in this area.

Good afternoon

It a real pleasure to be here with you today to talk about the impact of intelligent machines on the wholesale markets.

I don't need to dwell on the importance of the wholesale markets with this audience. Turning over $17 trillion – just about the annual GDP of the EU – every day, they are the foundation of the global economy, the transmission mechanism for global growth and increasing welfare of hundreds of countries, millions of corporations and billions of consumers worldwide. More than 95% of this turnover is in fixed income, currency and commodity markets; more than 40% of it takes place here in Europe; and 75% is conducted electronically, of which a large and increasing proportion is automated.

Automated trading is not new. Electronic equity markets have employed rules-based automated trading – so called algorithmic trading – for well over two decades. Electronic markets and algorithmic trading have been slower to develop in fixed income products, but are now as a result of changes in the past 15 years well-entrenched, and growing rapidly, in listed derivatives, spot foreign exchange, liquid government bonds, repo, and some other over the counter markets. Algorithmic trading has created new opportunities.

Rules-based trading algorithms have dramatically speeded up execution of orders, cut costs and increased volumes - and at least the appearance of liquidity - in markets. They have catalysed 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 crash" has attracted a lot of attention. And trading algorithms can be programmed to manipulate markets.

But ultimately, rules-based algorithms only codify and automate fixed trading strategies that humans have been using for decades or longer. For this reason, rules-based algo trading hasn't fundamentally changed the nature of trading or markets and I don't propose to spend more time on it now.

Much more interesting are the questions posed by machine learning in wholesale markets. By machine learning I mean algorithms that are not simple static, deterministic rules-based trading engines, forced to drive only along the rails laid by their human creator-programmers. But instead 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.

How widespread is this machine learning-driven activity and what are the implications, and particularly the hazards, of "true" artificially intelligent machines operating in wholesale markets? In particular, what are the consequences for users of markets looking for venues where they can be confident of fair treatment and effective execution and settlement?

How widespread is Machine Learning in Wholesale Markets?

Last month the Bank of England and the FCA published a survey of over 100 financial services firms in the UK to understand how they use machine learning.

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 therefore:

- despite long-time use of rules-based algorithmic 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.

The 2008 global financial crisis, and the manipulation of markets that was uncovered after that, have made us all think more deeply about the importance of fair and effective markets in which users can find counterparties who behave with integrity, reliable liquidity and dependable post trade settlement arrangements.

The fact that we are in the earliest stages of machine learning in markets means we have an opportunity to anticipate and mitigate potential future problems with the technology. So 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 more greater than we have allowed for so far. Human biases will likely creep into the programming of machine learning engines in the same way that they already can with deterministic, rules-based algorithmic machines.

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** - an inevitable result of the continuous lifecycle of machine learning - 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 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 and 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 monopolies 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:

- These skills are in short supply in the traditional 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.

- The other big question is the role of 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, 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 the technical complexity 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 a piece 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.