Rise of the robots: AI in financial market regulation

Phil Taylor
5 February 2024
#ReglexWrites

It seems hard to believe that it's only been 14 months since ChatGPT was released, bringing the capabilities and potential of large language models and generative Artificial Intelligence (AI) to public attention. AI as a concept is of course much older, dating back to the golden age of science fiction, and work on the technologies which reached maturity in ChatGPT had been underway for many years. But it was the user-friendly interactivity of the chat-based tool which sparked a wave of excitement and interest from companies keen to keep up with technology and use it to gain a competitive edge.

In the financial services (FS) sector, well resourced firms have been quick to embrace the new era of AI technology. They are perhaps more comfortable with the idea of automation in their systems than companies in other sectors, given that algorithms (affectionately known as "algos") have been traders' friends - and sometimes enemies - for decades.

As Investopedia explains, algorithmic trading "attempts to strip emotions out of trades, ensures the most efficient execution of a trade, places orders instantaneously and may lower trading fees". These tools are based on old-fashioned hard coded software technologies, however, making use of historic data and mathematical models to try to predict future market movements. This makes them notoriously susceptible to so-called black swan events - the appearance of the unexpected. There is a distinct lack of flexibility, and this is where modern AI technologies, including generative AI (algorithms which can generate new content rather than simply consuming data) come into their own.

"I always tell people that AI is just another tool like your calculator, but the difference is you typically expect a machine to be deterministic: if you enter 1+1, the answer today is the same as tomorrow," explains Calvin Chan, an AI product manager with a data scientist background, working for Roche and speaking in his personal capacity.

"AI is different - there's a certain margin of error that can be tolerated - it's more flexible like a human," Chan continues.

This is what makes new forms of AI ideal for applications involving analysis of, and making predictions based on, large quantities of data. This includes healthcare research, image processing – and FS.

Almost two years ago, an Economist Intelligence Unit (EIU) survey of IT executives in banking found 85% "had a clear strategy for adopting AI in the development of new products and services", while four out of five senior banking executives "agreed that unlocking value from AI will distinguish winners from losers".

Fraud detection has become a sub-sector where AI has already proved its worth, drawing on and leveraging the flexibility of the technology. The EIU survey also revealed that this was the top application of AI by banks in 2022. We have now, very rapidly following the ramp-up of AI's technical capabilities, reached the point where it would seem inconceivable for a large financial institution not to make use of AI tools to support its anti-fraud and financial crime analytic processes.

FS regulators are well aware of the potential applications of AI, and the incumbent risks, and - taking Hong Kong as an example - have not been slow to get involved in the subject. In November 2019, the Hong Kong Monetary Authority (HKMA) conducted a survey on the use of AI which found that many were adopting or planning to adopt AI applications, and that the scope of AI usage was "expanding from customer-facing services (e.g. chatbots and personalised marketing) to internal processes and risk management areas (e.g. operational automation, cyber and fraud risk management)."

At the same time, the HKMA issued its Circular on Consumer Protection in respect of Use of Big Data Analytics and Artificial Intelligence by Authorized Institutions in response to calls from the sector for a set of guiding principles on consumer protection aspects in respect of the use of "big data analytics and artificial

intelligence", or BDAI. The HKMA's Circular drew on guidance from the Organisation for Economic Co-operation and Development (OECD) in its [Updated Effective Approaches for Financial Consumer Protection in the Digital Age](#), and set out guiding principles focusing on four areas (governance and accountability, fairness, transparency and disclosure, and data privacy and protection), while continuing to emphasise the need for a risk-based approach. A year later, the HKMA published a report titled [Alternative Credit Scoring of Micro-, Small and Medium-sized Enterprises](#), which studied the use of AI in SME loan applications.

"The HKMA has been quite positive around innovation and fintech," says Padraig Walsh, a partner at Tanner de Witt law firm in Hong Kong. "There have been positive statements in various places where they speak about the adoption of tech themselves, for example in respect of 'suptech' as a subset of regtech. They have looked at incubating certain kinds of fintech generally as a way of making sure they're doing what they're encouraging the market to do, and looking at how they can innovate themselves."

Although AI is now firmly embedded, albeit to varying extents, within firms operating in the FS sector, and regulators are monitoring this area and taking steps to ensure it is appropriately regulated, conversations are at a much earlier stage when it comes to the use of AI by those who supervise the sector.

"It's interesting to consider not just how [the regulators] view AI in use in the market, but how they themselves would develop AI tools to help their own role, and actually, I think this is one of the most exciting areas of development in the sector," comments Philip Robinson of Tigersolv, a UK-based end-to-end business software solutions company.

"Artificial intelligence possesses immense potential to reshape the process of financial regulation. By employing AI systems, regulatory agencies can enhance their ability to interpret and enforce complex financial regulations," wrote Barry Quinn, Fearghal Kearney and Abhishek Pramanick in [a recent paper for The Economics Observatory](#).

As those authors point out, the number of rules faced by FS firms has increased dramatically over time. By studying levels of investment in fintech and regtech companies, they conclude that "the complexity of the regulatory landscape requires additional capital investment."

On top of this, FS institutions are themselves complex, adding to the burden of supervising them. So we are looking at a landscape with many complex rules and many complex players, meaning huge volumes of unconnected and intricate data. This is perhaps the perfect setting to fully exploit the potential of AI tools which, as we have seen, are capable of analysing large datasets, identifying patterns, and responding with a degree of flexibility and adaptability.

"What if the same principles of applying tech to free up the human experts so they can cover more firms, more effectively, with less effort were applied to regulators in the same way the companies in the sector apply it to themselves? What if the regulator could have their cake and eat it, too?" asks Robinson.

The Economics Observatory authors suggest that machine learning algorithms could "assist in identifying irregularities or discrepancies in financial transactions" which "could be indicative of non-compliance, fraud or other illicit activities". This would allow for quicker and more efficient regulatory interventions.

"The kinds of applications you could see [the regulators] looking to deploy AI for are … looking for anomalies in the datasets they are looking at, in terms of the market and institutions they are supervising, as a way of trying to eliminate things they don't need to be concerned about and looking at outliers," says Walsh. "Using AI for that type of analytics makes sense - it gives a degree of direction to their supervision."

AI can be effectively harnessed to look at entire, large datasets rather than resorting to random or selective sampling, which could miss important issues. Robinson goes on to suggest some other areas of application including: AI coaches, who "could help new firms reach necessary standards that they may not be familiar with"; AI data analysis examining client and transactional data to ensure growing firms are not slipping on quality while they take on more clients; and monitoring of telephone calls with human agents to "detect when clients (or agents) have become distressed, upset, or even more subtle moods such as confused, or unsure".

Walsh's view is that take-up of AI by regulators is inevitable.

"In terms of the volume of transactions and the scale of supervision they need to engage in, unless they use the tools available they won't be able to effectively

supervise," he says. "I don't think any regulator could take the view that they will reject AI because of the perceived risks, because they don't really have that luxury with respect to the scale of the tasks."

Here, Walsh hints at some of the hurdles which may have to be overcome before AI could be applied to regulatory activities. Some of its innate features could also pose significant challenges for this application.

High-profile media stories about AI inventing legal case law, "hallucinating" in general, or potentially breaching commercial or personal data privacy, will not help imbue the sector with confidence.

"Of course regulators need to be careful that they themselves don't put a foot wrong," says Robinson. AI is far more fuzzy than the tech we're used to dealing with." There is also the so-called black box problem: a lack of transparency and interpretability as to exactly what goes on inside the AI 'brain'. As Associate Professor Samir Rawashdeh of University of Michigan at Dearborn explains, "deep learning algorithms are trained much the same way we teach children. You feed the system correct examples of something you want it to be able to recognize, and before long, its own trend-finding inclinations will have worked out a 'neural network' for categorizing things it's never experienced before."

However, similar to our own human intelligence, "we have no idea of how a deep learning system comes to its conclusions," Rawashdeh says.

This could clearly be a big problem in the FS sector, where supervisors exercise a quasi-public function and have a duty to the public. They must remain fully accountable for their actions. It is difficult to imagine any responsible regulator using AI technology without strict supervision, so that the final output did not have the full fingerprints of the regulator on it.

"The regulator needs to be able to say 'We came to a conclusion or outcome in respect of an approach or decision by virtue of various things'. And if AI is part of that, they must be able to explain how they used it." says Walsh.

In other words, regulators must be able to show their workings. The key, says Chan, is to see AI as an assistant to, rather than a replacement for, key processes.

"With Tesla cars, you're still having your hands on the wheel – you have responsibility for what happens. This is the same as all AI products. Microsoft's 'Copilot' is exactly it: it needs a human there for when the margin of error occurs," Chan explains.

Walsh picks up on this point.

"I can see using AI as a support tool to something which the people within the regulator take the analysis from, and move forward. Whatever output comes from AI is not the determination of the regulator but is the starting point for the regulator to take forward and look at," he says.

Commentators agree that – despite the level of clamour from the wider market to jump on the AI bandwagon – a gradual introduction of the technology would be the only reasonable way forward for supervisors to take.

"You look at things capable of being analysed and that are replaceable [with AI]. If it can be done in a way which is reasonably secure, under appropriate supervision, then it's a better quality of regulation and supervision," says Walsh. "But the core public function remains – that doesn't go away just because you are using AI or other tech."

"The regulators would be slower than most, but this is by dint of them trying to control things in terms of error as much as they can. However, they are not constrained by trying to be first in the market so they can be more prudent," he adds.

Introducing innovation into the workflow of regulators is notoriously difficult given the tension between their public functions and duties, and the need to keep up with market trends. This is where 'sandboxes' and other virtual/test environments can be so useful. Walsh says we can expect to see this approach taken in terms of the adoption of AI, too.

In Robinson's view, although there can be value in moving quickly, it is worth sacrificing some speed in the name of stability. "In addition to rollouts being, perhaps, slow – applying AI only in one new area at a time and giving it time to bed in and spot side-effects – we also might understand that regulatory officers can only handle so many new firms of products or systems at once."

For all the hype surrounding it, AI is, at its core, just another form of technology. It brings with it all the usual concerns as well as advantages, and proper processes must be followed in rolling it out. This applies as much to private firms as to regulatory bodies. The consensus appears to be that regulators do not need to shy away from embracing AI as a way to help them supervise increasingly large and complex markets. But care must be taken, and careful choices made.

"It's not about trying to avoid innovation, more about making sure it's the right innovation," concludes Walsh.

reglex·io

Democratising Compliance
www.reglex.io | sales@reglex.io

Access | Control | Empower
www.reglex.io