



euromoney institutional investor
thought leadership

**Baker
McKenzie.**

GHOSTS

IN THE MACHINE

Artificial intelligence, risks and regulation in financial markets

www.euromoneythoughtleadership.com/ghostsinthemachine

All systems Go

2

On the 15th March 2016, an artificially intelligent (AI) software programme called AlphaGo, defeated the world champion of an ancient board game called Go. The game is immensely complex, with a total combination of possible moves numbering several hundred orders of magnitude more than the number of atoms in the universe. Winning the series four-to-one, AlphaGo's victory was emphatic. It also showcased significant advances in AI's ability to recognise obscure patterns, learn new ones and adapt strategies to changing circumstances.

Yet, just two weeks after AlphaGo's impressive victory, a new chatbot called Tay, exposed a darker side to AI. Designed to engage in friendly conversation with people online and assist them with Microsoft services, Tay's unique design feature was that "she" learns from her online interactions. Upon Tay's public release a coordinated barrage of abuse and incessant trolling by Twitter users, taught Tay the wrong lessons. The programme was corrupted into spewing racist, sexist and xenophobic comments, revealing the potential for flaws in the design and programming of AI, as well as the uneasy interaction between artificial intelligence and the natural kind.

Both events expose a tension underlying the introduction of AI. Programmes like AlphaGo demonstrate how AI can analyse vast amounts of information, recognise sophisticated patterns and empower humans with new analytical capabilities. Conversely, Tay's malicious malfunction serves as a reminder that the technology is far from infallible, particularly when interacting with humans.

After conducting a global survey of 424 senior executives from financial institutions and fintech companies, as well as interviewing leading experts in the field, this tension is also apparent as AI is pioneered across financial markets.

Many see AI as a tool that will help improve financial institutions' risk management, for example through more in-depth assessment of risk in portfolios and more incisive, comprehensive and informed credit-risk assessment. In these applications, AI promises not reckless speed or loss of control, but rather an unprecedented depth and breadth of insight, and the ability to act on information and learn from its actions.

"AI promises not reckless speed or loss of control, but rather an unprecedented depth and breadth of insight, and the ability to act on information and learn from its actions."


However many experts also acknowledge a degree of risk surrounding the use of AI. This stems partly from uncertainty – it is, after all, still at experimental stages in many applications including trading, portfolio management and credit assessment. As a result, the risk of malfunctioning algorithms and concerns surrounding the security, privacy and quality of data, has led to calls for new regulation.

There is an even greater unease about the regulatory response to AI. Participants in the study express a distinct lack of confidence that regulators have the adequate knowledge and skills to stay abreast of new financial technologies. Indeed survey participants suspect that regulators are only just beginning to understand the potential implications of AI for financial markets and companies. For now much of their attention is still focused on fighting the last war, identifying compliance breaches by humans directly abusing technology. Their attention is beginning to

turn to the integrity of algorithms, and any rule-writing on machine learning in the next few years will focus here.


It may also not be surprising, given how nascent the use of AI is in the sector, that a large number of financial institutions in the survey are not confident that all AI related legal risks have been understood by their organisation. For example, data and privacy risks will increase by virtue of the much larger volumes of data AI-driven models will collect and analyse. Intellectual property disputes are also likely to increase, as the ownership of algorithms causes friction between companies and regulators. Finally contract and litigation risk may also emerge, in the likely event of AI malfunction and programming errors.

or substantial change to their own jobs over the next 15 years. Even in trading, where automation is already widespread, human roles will remain critical in areas such as algorithm validation and monitoring, as well as compliance. At this point, few believe machine learning models can or should drive financial-market operations completely independently of human control.



7 IN 10
believe AI will
bring complete
or substantial
change to their
own jobs in
15 years

AI and machine learning will undoubtedly alter both the headcount and the nature of skills required in the industry. A significant minority of survey respondents fear the effects on the workforce will be negative within the next few years. But wholesale displacement of humans is for the longer term – nearly seven in ten believe AI will bring complete



FROM THINKING FAST TO THINKING SMART

4

Owners of the June 1987 issue of the Wall Street Computer Review, will know that talk surrounding AI in financial markets is nothing new. Sporting a front cover which reads, “Teaching Computers to Emulate Great Thinkers”, and picturing a Socratic figure preaching to a crowded audience of computers, even 30 years ago there were plans for AI-based trading applications. Many of these early applications proved more theoretical than material.

Despite previous bouts of hype, however, a number of commentators believe that renewed interest in AI is justified. Continual and rapid advances in computing power, as well as dramatic declines in the cost of computing have made AI applications more practical. The growth of social networks, mobile phones and wearable consumer devices has also led to an explosion in the amount and availability of data – all of which becomes fodder to optimise AI algorithms.

Renewed interest in AI is evident from increasing investments by major financial institutions, as well as technology and fintech companies. Fund managers such as BlackRock, Two-Sigma and Renaissance Technologies have been busy poaching the best data scientists from around the world. They compete and collaborate with a growing batch of technology companies including Context Relevant, Sentient Technologies and Kensho, as well as the giants of AI, such as Google, Facebook and Microsoft. In 2015 alone, these companies spent over US\$ 8.5 billion on AI research, acquisitions and talent.

Within trading and investment management, firms such as Aidiya and Sentient Technologies are pioneering AI trading programmes. They employ a combination of machine learning techniques and evolutionary algorithms to crunch huge amounts of data, in order to recognise obscure patterns, which others have not identified. As opposed to traditional forms of quantitative trading, which employ algorithms updated by human hand, many

of the AI software programmes learn and update their models automatically and independently of human interference.

Another characteristic of AI trading programmes is the importance of differentiation. As Saeed Amen, Co-founder of The Thalesians, a financial consultancy, argues “the benefit of machine learning is that it enables traders to find relationships that are not immediately obvious and hence much more difficult to find, and potentially not as crowded with other market participants.”



“The benefit of machine learning is that it enables traders to find relationships that are not immediately obvious and hence much more difficult to find, and potentially not as crowded with other market participants.”

.. Saeed Amen , The Thalesians ..

This push towards differentiation, distinguishes AI from other forms of algorithmic trading, such as high-frequency trading (HFT). If HFT, for example, is about speed, machine learning is about depth and breadth of insight. “The machine learning revolution is about making superior decisions by identifying sophisticated patterns from the ever expanding data set or information that is available to you – on

any horizon,” according to Peter Hafez, Chief Data Scientist of RavenPack, a provider of news and analytics tools to the financial industry. “The market is moving away from being faster to being smarter.”



“The machine learning revolution is about making superior decisions by identifying sophisticated patterns from the ever expanding data set or information that is available to you – on any horizon.

The market is moving away from being faster to being smarter.”

** Peter Hafez, RavenPack **

The potential for innovation is consequently significant, not only in trading but in other parts of the financial industry such as investment advice and lending. Change will not come instantly, but it will come. In accordance with Bill Gates’ famous aphorism, “we always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten”, our survey suggests that AI will cause a similar series of transformations in financial markets.

AI IN ACTION

6

How disruptive will AI and machine learning be, and where?

Over the next three years, the most dramatic changes will be felt in the areas of trading, financial analysis and IT, according to 64%, 60% and 60% of respondents respectively (see chart 1). Large numbers also expect machine learning to materially affect risk assessment (59%), credit assessment (57%) and investment portfolio management (52%). Risk assessment and financial research are the areas where companies are most likely to experiment with machine learning applications in the next three years.

Andrew Lo, Director of the Laboratory for Financial Engineering at the MIT Sloan School of Management (and himself the founder of a quantitative investment management firm) believes the impact will be wide-ranging: "I suspect that it's going to transform all aspects of the financial industry because there are so many parts of it that can be automated using these kind of algorithms and access to large pools of data."

Peter Hafez believes machine learning will, in addition to trading and research, greatly benefit consumer credit scoring as well as the compliance function in different types of financial institutions. He notes, for example, that compliance managers are beginning to use unstructured content such as news feeds to alert them about suspicious trading.

Machine learning techniques have already found application in retail investment advisory. "Robo-advisers" – investment management websites providing automated advice to investors – form an area of AI coming under active regulatory scrutiny now, according to John Price, who serves as a commissioner in the Australian Securities and Investment Commission (ASIC). In their latest Financial Advice Market Review (FAMR) the UK's

Chart 1

Which FS functions will be changed the most by AI over the next three years?

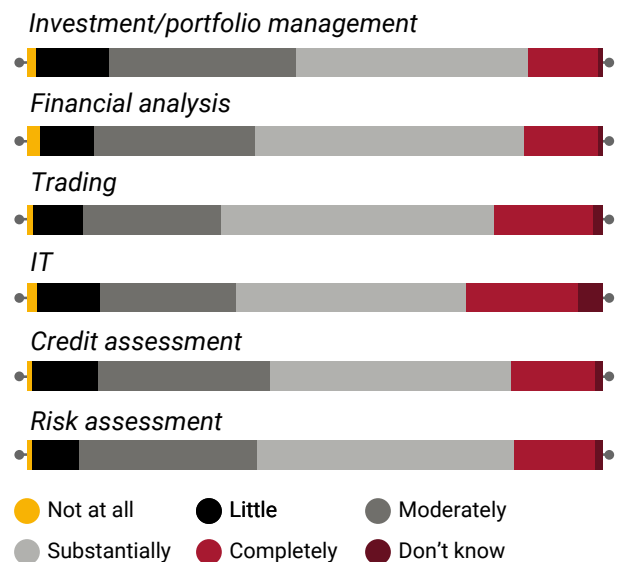



Chart 2

Where do you expect AI/machine learning technology to be introduced in your organisation in the next three years?



Financial Conduct Authority has even gone so far as to recommend robo-advice as a cost effective way for financial institutions to “streamline advice” to their customers.

Disruption on the horizon

Experts interviewed for this report believe machine learning will be mildly disruptive to financial-market operations, at least in the next three to five years. When it comes to trading, Babak Hodjat, Chief Scientist at Sentient Technologies, a firm which develops AI software, predicts there will be a sizeable impact but not fundamental disruption. “The reason is that you’re moving from people and quants coming up with the strategies and deploying them to the machines coming up with the strategies and deploying them. You can churn through that faster but the end result will be similar.”

Most of the experts agree, moreover, that the technology will have many positive applications. Paul Ebner, a senior portfolio manager within BlackRock’s Scientific Active Equity Unit, believes trading will benefit from the depth of analysis machine learning tools enable over a wide breadth of companies. “It’s being able to go a couple of steps deeper than you could just by using, say, data in a spreadsheet. Speed matters but it’s a different kind of speed than high-frequency trading. For us it’s being able to process a lot of data very quickly and coming up with the right answer that the markets will eventually discover.”

BlackRock’s Scientific Active Equity unit – a roughly 100-member team which includes data scientists and machine learning specialists as well as more traditional financial-industry ‘quants’ – is putting machine learning techniques to work in different ways to forecast share price movements. According to Mr Ebner: “We’re applying tools to analyse data about companies and using that data to forecast



“It’s being able to go a couple of steps deeper than you could just by using, say, data in a spreadsheet. Speed matters but it’s a different kind of speed than high-frequency trading. For us it’s being able to process a lot of data very quickly and coming up with the right answer that the markets will eventually discover.”

.. Paul Ebner, BlackRock ..

Chart 3

How positive or negative an impact will AI have on financial services?



the fundamentals, and then ultimately to forecast their stock returns and construct portfolios around that.” An example is the team’s analysis of companies’ currency exposure to gain more in-depth insight into their expected earnings well before their quarterly announcements (see forex exposures case study).

Survey respondents are clear about another benefit machine learning will bring: 64% believe its use will have a positive effect on competitiveness in financial markets (see chart 3). At first glance, this finding seems counter-intuitive. The high costs of the best talent and the most advanced technologies should make AI accessible to only those with the deepest pockets. This is true to an extent – only a firm like Bridgewater Associates could afford to hire the chief engineer behind IBM’s legendary supercomputer, Watson.

However smaller companies, and even individuals, have proven that they can be at the forefront of new innovations. This is partly because open-source software empowers smaller organisations to experiment with advanced algorithms and code. A recent example was seen in March 2016 when two retired hedge fund “quants”, with no prior experience of working with AI, managed to design an algorithm that diagnoses heart disease from MRI images. These individuals were relatively inexperienced in AI and yet they were able to design a transformative application by downloading software from the open-source site, GitHub.

What this means for competition across financial organisations, only time will tell. However, whilst industry giants, such as BlackRock, will be doing a lot of the heavy lifting of research and pilot work, AI applications and services could become widely available for a range of small and medium sized organisations.

AI Insight

What is clear is that advances in AI and data analytics are leading to a great expansion in the quantity and type of data being used to inform decision making. Whereas before investment decisions were being made on traditional metrics such as market prices, interest rates or earning figures, AI can factor events and sentiments into the asset-price prediction process. For example, the crunching of unstructured data is helping to advance sentiment analysis. But Mr Hafez says sifting facts is just as important as interpreting opinions: “It isn’t just about sentiment in the traditional sense, such as guidance, but also about the facts that can be extracted from unstructured content and that can be delivered in a machine-readable format.”

Machine-trading models today typically analyse earnings statements and company reports. In time they will be able to capture much more through the ability to analyse news about product releases and



“It isn’t just about sentiment in the traditional sense, such as guidance, but also about the facts that can be extracted from unstructured content and that can be delivered in a machine-readable format.”

recalls, regulatory approvals, acquisitions and other market events. Mr Hafez also expects machine learning models to be fed insights gleaned from images, video footage and live streaming. There are companies, he says, that use satellite images to track the number of cars in the parking lots of large retailers' stores to try and understand the likely direction of their sales revenue and earnings. This suggests that AI could change the parameters by which financial institutions make investment decisions. Traditional metrics will decline in importance as the subject of analysis, as financial institutions gather huge amounts of unstructured data – which can only be made intelligible through AI and machine learning.

Analytics tools are getting better at understanding context – another critical differentiator of machine learning. Paul Ebner explains that the machine learning tools his team uses are now able to use context to understand the nuances of word use. “For example, the word ‘garbage’ in an SEC filing probably refers to waste management, but the same word on an investor blog is probably a critical term for the stock or the firm’s management. We’re able to build dictionaries that learn and evolve based on the environments we’re pulling the language from.”

Similar technology is being introduced to assist retail customers with their complaints and queries. The Royal Bank of Scotland (RBS) has recently announced that it will be introducing an AI programme called Luvo, to help with customer complaints. The programme has been designed with an artificial personality that can mimic human characteristics such as friendliness, empathy and reason. According to the programmers, Luvo can also learn from its mistakes and gauge a person’s mood. In this light, AI is being cast as more friend than foe.



“For example, the word ‘garbage’ in an SEC filing probably refers to waste management, but the same word on an investor blog is probably a critical term for the stock or the firm’s management. **We’re able to build dictionaries that learn and evolve based on the environments we’re pulling the language from.**”

.. Paul Ebner, BlackRock ..

CASE STUDY: Understanding forex exposures to predict earnings

Companies with high volumes of export sales are naturally subject to the whims of foreign exchange markets. In 2014, according to BlackRock's Paul Ebner, currency movements – not least a strengthening American dollar – had a larger than normal impact on corporate earnings. BlackRock's Scientific Active Equity team was not caught out by it, thanks in part to machine learning.

About three years ago, Mr Ebner recalls, the team started mapping high-profile companies' sales exposure to exchange-rate movements, partly using machine learning techniques. So in early 2014, he says, during sustained periods of dollar appreciation, "on certain days we already knew as the dollar was moving which companies were going to be affected, and we could trade around that and make we sure we were either selling or buying the right companies." Eventually, says Mr Ebner, the market figured it out when the companies announced that they had missed their earnings targets, explaining that "the dollar had strengthened and we were too expensive".

Thanks to its model – and the three years' worth of data it had fed into it – BlackRock had good estimates of several companies' earnings figures three or four weeks ahead of their published reports, according to Mr Ebner. Compare this, he says, to the few fractions of a second advantage that high-frequency traders typically enjoy.



FROM MOORE'S LAW TO MURPHY'S LAW

11

What could go wrong? Plenty, according to Andrew Lo. He believes the markets may be in for more flash crashes, for example, or for other negative developments about which neither the industry nor regulators currently have a clear understanding. “The nature of these strategies makes them very difficult to understand. That means that the interactions are going to be hard to predict, in the same way that nobody predicted the flash crash of May 6, 2010 – even today we still don’t really understand what happened.” He also points to the demise of Knight Capital, a major US trader whose software glitch in 2012 caused it US\$440m in trading losses and sent it to the brink of bankruptcy: “I don’t think we’ve nearly fixed those kinds of issues, because ultimately you’re dealing with a mismatch between human ability and technology. It’s Moore’s law meets Murphy’s Law.”

Technology will not be able to remove the risks inherent in some financial activities, such as making bets on future events. These are likely to persist, regardless of whether humans or algorithms do the work. As Arun Srivastava, Partner at law firm Baker & McKenzie argues, “Financial institutions have been fined billions of dollars because of illegality and compliance breaches by traders. A logical response by banks is to automate as much decision-making as possible, hence the number of banks enthusiastically embracing AI and automation. But while conduct risk may be reduced, the unknown risks inherent in aspects of AI have not been eliminated.”

All in the algorithms

Regulators are uncertain as yet what the risks of machine learning are but are focusing on algorithms as an area where problems could occur. Victoria Pinnington, Senior Vice President of Market Regulation of the Investment Industry Regulatory Organization of Canada (IIROC), says her greatest concern currently is around the crafting of algorithms,

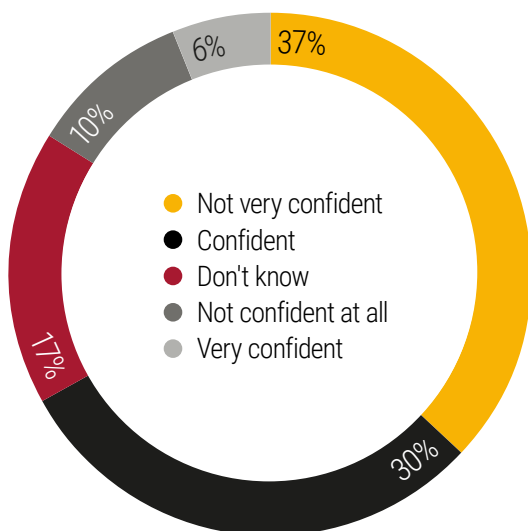
in both machine learning and broader systematic trading contexts. “If there is a problem with the algorithm,” she says, “the impact on the markets could be considerable.”

Algorithms can malfunction in a variety of ways. One of the most common errors is known as “overfitting”, which usually occurs when an algorithm is overly complex and lacks a suitable hypothesis. In such cases they fail to distinguish useful correlations (signals) from the mass of irrelevant data (the noise) and instead identify “phantom” factors or spurious correlations. Imagine trying to record a classical music concert with a sensitive microphone. Overfitting would equate to picking up the surrounding background noises, rather than the sound of the orchestra. As Babak Hodjat warns, “people could mismanage machine learning and not do the validation. If you take a machine learning algorithm and not sufficiently validate it, you might have something that’s overfitting, that might look pretty good right now but might fail miserably tomorrow. More scrutiny is required there.”

The risk of programming error increases with the pressure to launch new programmes. As mentioned earlier, differentiation is crucial to the success of an AI trader. Being the first successful application gifts an organisation a unique and profitable opportunity – an uncontested marketplace. Consequently there is a risk that organisations will rush their strategy to market. As Nate Soares, a former Google engineer and current research fellow at the Machine Intelligence Research Institute recently estimated in an interview with The Financial Times, “there is only a 5% chance of programming sufficient safeguards into advanced AI.”

Chart 4

How confident are you that all material legal risks associated with new financial technologies have been properly understood by your organisation?



Data, liability and legal risk

There is a great deal of uncertainty among survey respondents as to whether organisations understand the legal risks associated with new financial technologies: 47% are not confident that they do (see chart 4). Mr Price is not surprised. “The reason is that this technology is at a nascent stage, and it is evolving. The fact that people are cautious and a little unsure about what some of the risks might be, reflects that nascent stage of where the technology’s at.”

One risk is corporate liability. Flawed investment decisions could be made as a result of poor data, erroneous analysis about company performance, or malfunctioning algorithms, which could cause investors significant losses. Liability could also

arise should machine learning models make flawed decisions about credit risk: financial losses could occur to lenders, or alternatively borrowers’ reputations could be damaged. There is also a lack of clarity about which parties would bear liability should such situations occur –the financial institutions themselves, the writers of the algorithms, the exchange platforms, data providers or other parties.

The intelligent, data-crunching properties of machine learning may also take data protection and privacy risk to another level. Personal investor data or sensitive company data falling into undesired hands, whether by accident (to hackers) or design (to marketers and governments), is by now an all too familiar risk of the Internet age. This risk will grow simply by virtue of the much larger volumes of data that machine learning models will gather in. Organisations will increasingly need to understand how data privacy is entwined with laws on consumer protection, as well as related pieces of legislation such as the EU Cookie Directive. Adrian Lawrence, Partner at Baker & McKenzie expects access to data to play a central role in the scope and impact of AI systems, noting that “Data, and the various rules and processes which both enable and regulate access to and use of that data, stand at the heart of disruptive fintech businesses. Even the most advanced and intelligent algorithms and models are useless without efficient, secure and legal access to detailed, accurate and up-to-date data sets.”

Beyond legal risks, the survey respondents clearly lack confidence that the impact of AI is fully understood by their organisations. 49% of respondents are not confident that their organisation understand the other material risks associated with AI. Only, 32% are confident. Given the early-stage development of applications, this finding indicates that AI will present organisations with a set of risks, most of which are still to be defined.

AI Over-reliance

The biggest risks, according to some experts, lie less in machine learning techniques themselves, but rather in humans' misuse of technology, or misplaced confidence in technology to achieve goals by itself without human guidance. "If we have blind faith in technology," warns Babak Hodjat, "things will go wrong. If the success of AI means more use of technology in an uncontrolled and non-principled way, then we're risking more."

Saeed Amen worries that the industry will use machine learning as a sort of black box. Should this be the case, he says, "They'll end up creating a trading model that they don't really understand the ins and outs of. That is a dangerous scenario but it's the same with any systematic model. You really need to understand what's going on in the trading strategy."

Just like humans - programmes, computers and machines have the capacity to be stupid. The danger is that they can act at a far greater scale and speed. Examples such as the Knight Capital disaster, serve to illustrate the importance of maintaining human oversight, comprehension and control of emerging AI systems.



6 IN 10
believe machine learning will "greatly enhance" their risk-assessment processes

Herein lies the contradiction at the core of the technology. When confronted with inherently risky tasks – such as making investment decisions and bets on unknown future events – over-reliance on AI can magnify systemic risks. Yet the same technology can improve the depth and quality of financial institutions' due diligence of companies. Through their powerful data-crunching capabilities, such applications can also help identify fraud, money-laundering, bribery and other corrupt practices that more conventional methods would struggle to uncover.

With these examples in mind, the survey respondents appear hopeful that machine learning will help them to minimise risks, in some cases. Nearly six in ten (58%) believe it will "greatly enhance" their risk-assessment processes. Machine learning techniques can, for example, be used to alert fund managers about emerging weaknesses in invested companies (see box "Understanding forex exposures to predict earnings"). Astrid Raetze, Partner at Baker & McKenzie also expects consumer credit risk assessment to be enhanced through more thorough risk profiling of customers suggesting that, "AI should also reduce risk in some areas if deployed properly. Market misconduct and anti-money laundering (AML)/Know Your Customer (KYC) processes are areas where regulators could harness AI to improve regulatory oversight and scrutiny."

Machine learning-based analytics can also identify patterns in client activity that may point to some form of malfeasance. This helps explain why respondents point to risk assessment, ahead of other areas of operation, where they expect machine learning to be implemented over the next three years.

AI vs. IP

Survey respondents appear to agree that algorithms need more regulation. Over half of them (56%) say that affording regulators' access to examine trading algorithms would help to keep the financial system safe. This is a telling result that comes at a time when a number of financial regulators are planning to make the source-code of algorithms open to examination by authorities.

For example the US Commodity Futures Trading Commission (CFTC) is currently trying to push forward the regulation of Automated Trading (otherwise known as "REGAT"). One of the most controversial features of this proposal is enabling the CFTC and the US Department of Justice access to a financial firms' trading algorithm.

REGAT's most vocal opponent, Chris Giancarlo, (a Republican member of the US CFTC) has argued that giving regulators this degree of control will impose extra compliance costs on smaller market participants and hamper innovation in the futures market.

Moreover, Mr. Giancarlo, is primarily concerned that this regulation would mark an unprecedented invasion of private intellectual property rights by public authorities. As he recently expressed in a statement, "I am unaware of any other industry where the federal government has such easy access to a firm's intellectual property and future business strategies. Other than possibly in the area of national defense and security..." Similar mechanisms of oversight are expressed in MiFID market regulations in Europe and could lead to similar conflicts over the intellectual property of algorithms.

Indeed, John Flaim, Partner at Baker & McKenzie predicts a significant change in the nature of intellectual property disputes, arguing, "Intellectual property ownership will be key. There has been a

demonstrable increase in the filing of e-commerce patent applications since the turn of the century. However the U.S. Supreme Court issued a decision which has led to the invalidation of many e-commerce patents. Many other jurisdictions, like those in Europe, do not permit patenting of software. The financial industry is thus grappling with balancing their increased research and development efforts with a higher bar for patenting in the e-commerce space."

Before source code repositories are handed over, regulatory agencies also need to demonstrate competency on data and cyber security. In March 2014 a group of Chinese hackers managed to hack the US Office of Personnel Management. They stole the records of 21 million US federal employees, including senior members of the US CFTC. Weaknesses like these would need to be remedied before authorities gained access to algorithms. The big question is whether regulatory authorities are in a position to keep up with rapid changes in technology?

Are regulators up to speed?

When asked if financial regulators are "keeping pace with advances in technology", an overwhelming 76% of survey respondents say no. Nearly seven in ten express little or no confidence that "regulators have sufficient understanding of financial technologies and their impact on the financial services sector today." One respondent comments that "Regulators are woefully under-skilled in AI and need to boost their understanding or risk being marginalised."

Regulators are certainly at a disadvantage vis-à-vis large financial institutions in the competition for data scientists and other professionals with knowledge of machine learning. This makes it difficult for them to remain completely up to date on technology developments in this area. Regulators are, though, beginning to explore the role and implications of machine learning in financial markets. As seen with

Chart 5

Most important step the regulator should take to address the impact of new technologies

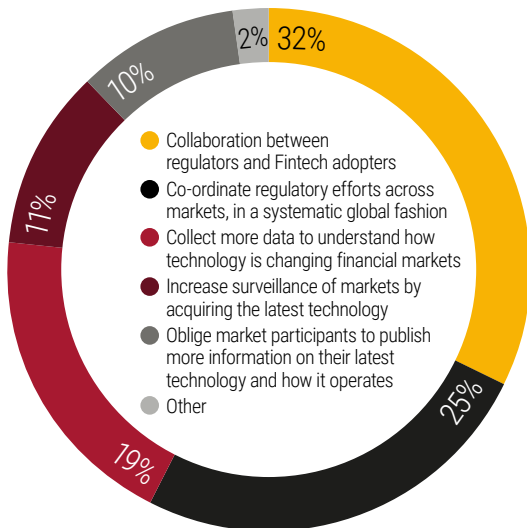
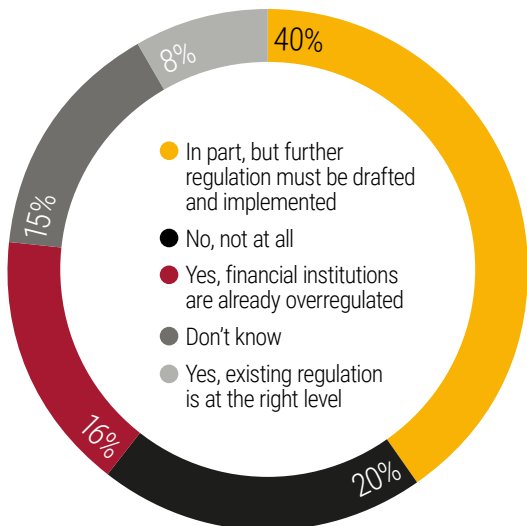


Chart 6

Is existing regulation sufficient to address issues posed by AI?



REGAT and MiFID II, much of the exploration is taking place in the context of systematic trading and not specifically of machine learning. Nevertheless, led by the Securities and Exchange Commission and the Financial Industry Regulatory Authority in the US, the UK's Bank of England and the Monetary Authority of Singapore, regulators are starting to learn about the role of AI and machine learning in financial markets.

In ASIC, John Price heads an innovation hub set up in early 2015 that is examining different areas of machine learning application in financial markets, and it is already providing advice to organisations using such techniques. Victoria Pinnington is spearheading a similar initiative within Canada's IIROC. Both officials say their organisations are exchanging the results of their research with other regulators.

Such interaction is in the spirit of recommendations made to regulators by the survey group. When asked what single step regulators should take to manage the risks of new technologies, most respondents (32%) suggested collaboration between regulators and fintech companies. The second largest group (25%) suggested co-ordinating regulatory efforts across markets, in a systematic global fashion (see chart 5).


Most industry executives in the survey believe that some form of new regulation will be required to deal with AI and machine learning. 60% of those surveyed believe that existing regulation needs to be improved and that current regulation is not sufficient (see chart 6). But regulators do not anticipate rules specific to AI to be written anytime soon. Those that emerge will focus on algorithms themselves or to the broader field of systematic trading. In Australia, says Mr Price, any rule-writing is likely to be principles-based rather than prescriptive. "Any new rules will not say 'do X, Y and Z'. Instead they will stipulate that firms must, for example, have adequate risk management procedures in place."

Pilot and autopilot

Over time, machine learning will almost certainly push some human beings – traders, analysts and other industry employees – out of their existing roles. Within 15 years, 68% of survey respondents expect to see complete or substantial change to their own jobs. Four in ten respondents fear it will have a negative effect on the structure of the workforce sooner – within the next three years.

In most occupations, however, including trading, humans are unlikely to fade from the scene anytime soon. According to Mr Hodjat, the individual trader's role is going to diminish somewhat, but not entirely. He points out that certain types of trading expertise cannot be displaced, and that talented professionals will be needed, for example, to set up and validate the algorithms. This may frustrate the predictions of one Microsoft executive, who claimed in 2014 that, "robots will be running the City within 10 years, rendering investment bankers, analysts and even quants redundant."

Mr Amen also believes that human intuition and influence is crucial to guide successful AI applications: "You still need to use a modicum of



4 IN 10
fear it will have a negative effect on the structure of the workforce sooner – within the next three years.



"There's structure around us and we may be on autopilot most of the way, but we enter the details into the navigation system and we decide when to engage the autopilot and when to fly manually.

We're in control of the plane."

.. Paul Ebner, BlackRock ..

market understanding and intuition when you use machine learning. It's not the case that you just put in a system and leave it for ten years; you constantly want to be coming up with new ideas which are correlated as the market changes, and that still requires humans at the end of the day."

Mr Ebner likens the portfolio manager's role in the age of machine learning to that of an airline pilot. "There's structure around us and we may be on autopilot most of the way, but we enter the details into the navigation system and we decide when to engage the autopilot and when to fly manually. We're in control of the plane."



CONCLUSION

17

At the beginning of 2016 a group of the world's leading entrepreneurs, including Peter Thiel and Elon Musk, announced that they would put US\$ 1 billion into creating an organisation called OpenAI. The sole purpose is to help protect humanity from Artificial Intelligence. In an open letter, the founding members summarised the tension lying at the core of this technology, writing, "It's hard to fathom how much human-level AI could benefit society, and it's equally hard to imagine how much it could damage society if built or used incorrectly."

A similar sentiment underlies the feelings surrounding AI's application to financial markets. All recognise that there is much to learn about how transformative machine learning will be. There is also much to learn about its potential downsides.

Most of our survey respondents are cautiously optimistic about AI's future role in financial markets. The optimism derives from the recognition of the great opportunity that awaits successful applications. However, like with all technology, it will largely depend on how it is wielded that will ultimately determine the risk and reward.

Defining our terms

AI is an umbrella term encompassing several fields of research in computer science, all of which seek to enable computer systems to perform tasks normally requiring human intelligence, such as visual perception and decision-making. Machine learning is a branch of AI that provides computer systems with the ability to learn and adapt independently, based on algorithms and the analysis of data. Machine learning is beginning to be deployed in several corners of the financial industry, most prominently in trading and financial research, but also in other areas such as investment advisory.

The research

In conducting the research for this report, Euromoney Institutional Investor Thought Leadership surveyed 424 senior executives from financial institutions around the world. Over one-quarter of respondents, or 26%, work in asset management firms, 16% in investment banks and the balance in banks, insurance firms, hedge funds and brokers. The majority of respondents – 57% – hold C-level positions in their companies; the others are senior managers in a variety of roles including data, technology, legal and compliance. A mix of large, midsize and small firms are represented, with 51% having 100 or more employees. Finally, the survey sample is global: one-third of respondents are based in Europe, one-third in North America, 16% in Asia and the remainder in Latin America, the Middle East and Africa.

In addition, in-depth interviews were conducted with eleven senior industry executives and independent experts.

They are:



Saeed Amen

Managing Director and Co-founder,
The Thalesians



Paul Ebner

Senior Portfolio Manager,
Blackrock Scientific Active Equity



John G. Flaim

Global Head of Intellectual Property,
Baker & McKenzie



Peter Hafez

Chief Data Scientist, RavenPack



Babak Hodjat

Chief Scientist, Sentient Technologies



Adrian J. Lawrence

Partner – Media & IT, Baker & McKenzie



Andrew Lo

Professor of Finance and Director,
Laboratory for Financial Engineering,
MIT Sloan School of Management



Victoria Pinnington

Senior Vice President of Market Regulation,
Investment Industry Regulatory
Organization of Canada (IIROC)



John Price

Commissioner, Australian Securities and
Investment Commission (ASIC)



Astrid Raetze

Partner – Banking & Finance,
Baker & McKenzie



Arun Srivastava

Partner – Financial Services Regulatory,
Baker & McKenzie



APPENDIX


The survey data yielded many interesting findings and insights into how AI technology is being introduced, managed and perceived by executives from around the world. Below are some of the other interesting findings of the survey.

Africa as an AI Outlier?

When asked which three sectors A.I. and machine learning will disrupt the most within the next three years, FS executives from nearly every corner of the world chose credit provision, asset management and stock and trading exchanges as those to be the most effected.

The exception were those respondents from the Middle East and Africa. Interestingly they predicted that payment systems and virtual wallets would be the first to be disrupted. As a number of mobile banking platforms were popularised in rural Africa, can we expect to see new AI applications further revolutionize retail banking in Africa?

Which sector will AI disrupt the most within the next three years?



- **Provision of credit**
(eg. credit cards; unsecured loans; car finance)
- **Asset management**
- **Stock and trading exchanges**
- **Payment systems** (eg. PayPal)
- **Wallets** (eg. Apple Pay, Google Wallet)

C-suite see disruption

We analysed the responses from our C-suite participants, an influential group that constituted 57% of the overall survey demographic. Alarmingly, these decision-makers believe the most negative effects of AI will be in the structure of the human workforce across financial services.

39% of respondents believe that the impact of A.I. on the structure of the human workforce, will either be very negative or negative. They also recognised negative effects of AI on market stability, with 38% believing the technology will have either a negative or very negative influence.



Collaboration and co-ordination over surveillance and intrusion

When we asked respondents to suggest the most important step regulators should take to address the impact of new technologies, the overall majority suggest that collaboration between regulators and Fintech adopters is the most important (32%). Respondents in Asia, however, identified co-ordination of regulatory efforts across markets, in a systematic global fashion, as the most important step for the

regulator to take. (38%) Unsurprisingly respondents did not believe that increased market surveillance by regulators or obliging market participants to publish more information on their technology were the best solutions.

Confidence in regulators

Confidence in regulators is low across the world but is particularly pronounced in North America. This is an unusual finding. Financial regulators in the US have introduced some of the more advanced institutions, including the Office for Financial Research; an organisation tasked with developing a more granular understanding of financial markets, through advanced data science.


AI/machine learning and regulation

A large minority of respondents believe that further regulation does need to be drafted and implemented to address issues posed by AI/machine learning. Respondents that specialise in legal, compliance and regulatory functions are making the loudest calls for new regulation, along with those respondents in data and technology specific functions.




DEEP-DIVE INTO THE DATA

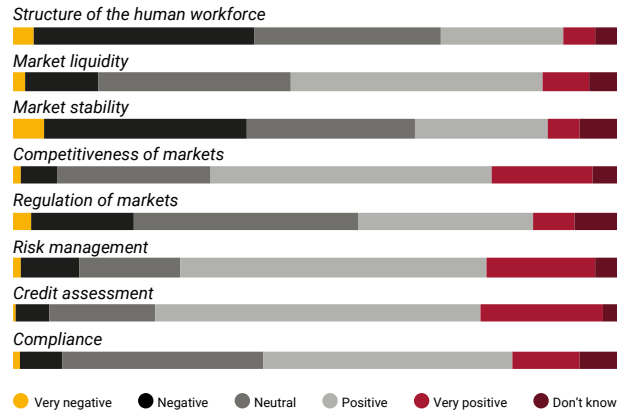
Q1: How much do you think the following financial service functions will be changed by AI and machine-learning technology over the next 3 years?




Q2: In which FS sectors do you expect AI and machine learning to have the most disruptive impact over the next 3 years? (Select up to three)




Q3: What impact will AI and machine learning have on the following aspects of the financial markets over the next three years?



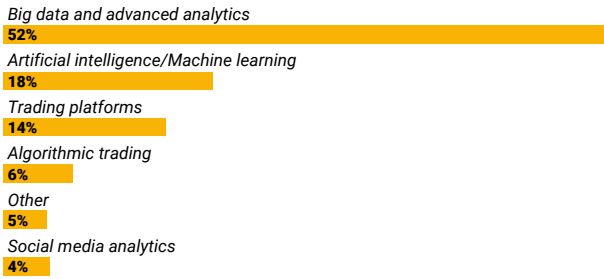
Q4: Which of the statements below most closely resembles your predictions on the impact AI and machine learning will have on the structure of financial markets?



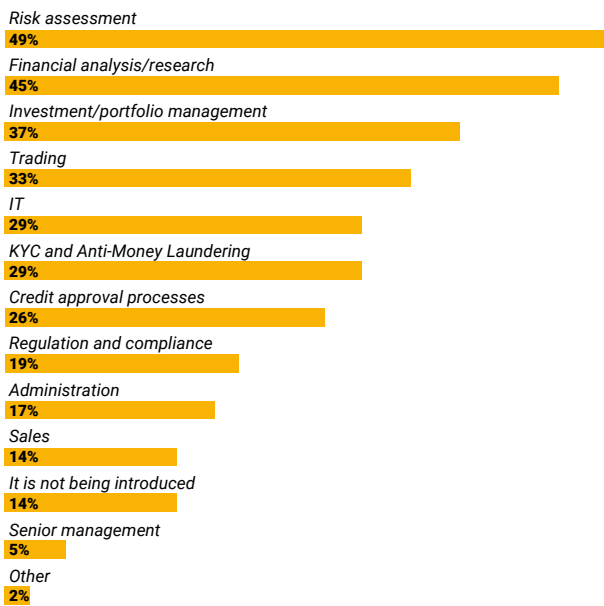
Q5: How much do you think your own job will be changed by AI and machine learning technology over the medium and longer terms?



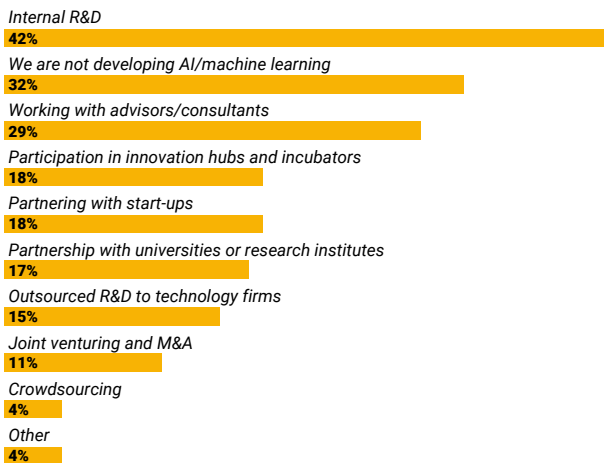
Q6: Which of the technologies in the list below is the most important to your organisation over the next three years?



Q7: Where do you expect AI/machine learning technology to be introduced in your organisation in the next three years? (Select all that apply)



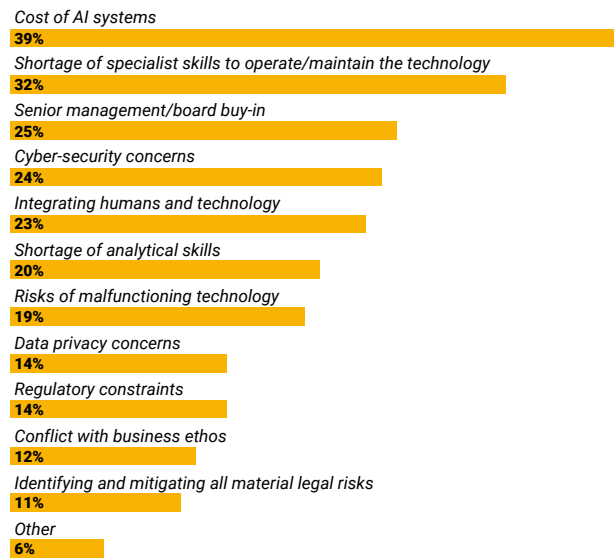
Q8: By which means is your organisation developing its capabilities in AI/machine learning? (Select all that apply)



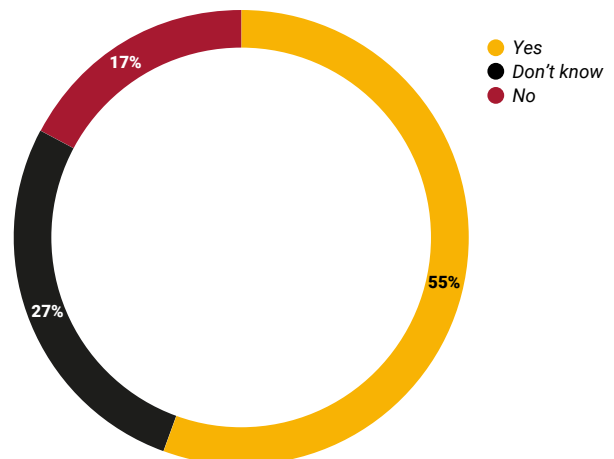
Q9: What are the most important benefits your organisation hopes to obtain from introducing AI/machine learning technologies? (Select up to three)



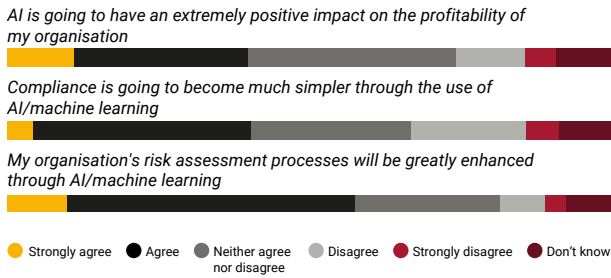
Q10: What are the toughest obstacles your organisation faces in seeking to introduce AI /machine-learning technology? (Select up to three)



Q11: Does the board of your organisation understand the wider impact of new technologies on its business?



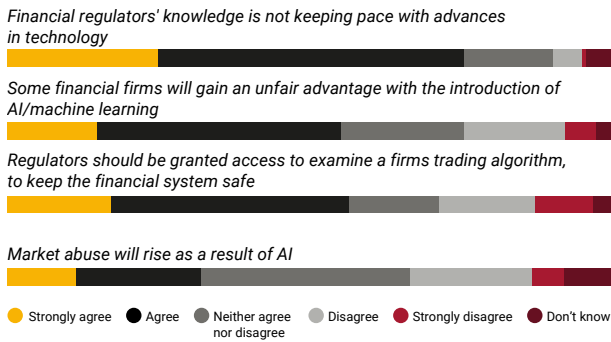
Q12: Do you agree or disagree with the following statements?




Q13: Which of the following is the single most important step the regulator should take to address the impact of new technologies on financial markets?




Q14: Do you agree or disagree with the following statements?




Q15: How confident are you that regulators have sufficient understanding of financial technologies and their impact on the current financial services sector?




Q16: Where do you think regulators should prioritise the adoption of AI technology, to reduce regulatory risk?




Q17: Do you think existing regulation is sufficient to address the issues posed by AI / machine learning?



Q18: How confident are you that all material legal risks associated with new financial technologies have been properly understood by your organisation?



Q19: How confident are you that all other risks associated with new financial technologies have been properly understood by your organisation?



GHOSTS IN THE MACHINE:

Artificial intelligence, risks and regulation in financial markets

Editor: Tom Upchurch

Writer: Denis McCauley

Designer: Claire Boston

[www.euromoneythoughtleadership.com/
ghostsinthemachine](http://www.euromoneythoughtleadership.com/ghostsinthemachine)

Ghosts in the Machine: Artificial intelligence, risks and regulation in financial markets is an Euromoney Institutional Investor Thought Leadership report, commissioned by Baker & McKenzie. The research was conducted by Euromoney Institutional Investor Thought Leadership. The findings and views expressed in this report are those of Euromoney Institutional Investor Thought Leadership alone and do not necessarily reflect the views of the sponsors.