Finding our feet in a world of AI

Christian Hull

If you were in AI in Financial Services in 2023 it was a wild ride. Whereas AI was a tough sell in previous years; executives didn't see the value, and readily cited explainability, or untested outcomes as a reason not to pursue the technology. Machine learning was quietly being used by quants and hedge funds, but mainstream financial services were asleep at the wheel. Then, OpenAI launched ChatGPT, and the world literally changed overnight, suddenly budgets were open, and expectations were sky high!

Even now at the end of 2023 many executives are unclear on what Artificial Intelligence (AI) is all about, so in this article I will try to pull back the veil and explain what the value is, which types of AI are delivering value, and how to build an AI capability at your firm.

WHAT IS AI?

If only to dispel misbeliefs, it is always best to start with some basic concepts. So, what is AI? At its simplest AI is a black box with data as an input and probability as an output. When you dig a little deeper into the world of data science you soon discover new terms like machine learning, deep learning, supervised learning, classification, regression, etc. It can be daunting, but if you are not building actual models yourself it is safe to lump all the complexity into a single category called AI. It is, however, important to know the different types of outcomes an AI could derive from your data. The Turing Institute published a paper[1] that examines the techniques and use cases in greater technical detail that I would recommend if you wish to dive deeper into the subject.

- Predictive Analytics – Simply put this is the ability to predict outcomes, think of portfolio valuation trends, settled trade positions, cash balances, asset prices, etc. Doing this manually is often near impossible, but AI models can happily process a wide range of factors needed to calculate a future outcome. This is hugely important in financial services, and I'd argue the lion's share of AI models built so far are predictive.
- Anomaly Detection – "One of these things is not like the other". This is useful for fraud detection, anti-money laundering, or simple data quality checking.
- Simulations – AIs can be built to run complex scenario models like the famous Monte Carlo simulation. Futurists also know that this is one of the main abilities of quantum computing, and so we'll see more of this emerging in the coming years.
- Automation – Most people think this is the primary use case for AI, but I'd argue that most of the automation use cases have already been executed using existing technologies, and now we have a long tail of small tasks that don't justify the expense of AI. However, if you have a process that is expensive and was not possible to automate with existing technology, AI could be the innovation you have been waiting for. Just ensure the business case is sound.
- Large Language Models – This is the category that contains ChatGPT, and other foundation models. The novelty of these generative AI models is that they are a new technology that allows complex language-based tasks that were previously impossible to achieve economically using natural language processing.

Christian Hull
Ex Head of Innovation, EMEA, BNY Mellon

WHAT WILL AI DO FOR ME?

A recent survey[2] conducted by Oliver Wyman on UK Finance members showed that 91% of firms had deployed Predictive AI versus only 22% who had deployed any generative AI solutions in any way. This is typical of the trend seen across all geographies. Generative AI models, typically large language models that generate new content based upon the data they have been trained on, are attracting all of the hype, but the real work is being done by more traditional narrow-focus AI models. In my experience, predictive analytic models are immediately valuable to financial services firms; many others can add real value but are harder to embed into the current state process. The simple ability to predict an outcome before it happens can easily be added to existing process flows to enhance the value of the data.

It has been amazing to see how Generative AI has galvanized the executive suite into action. I see clear evidence of this in the amount of hiring that firms have been doing. The Evident AI Index[3] noted that there has been a 10% increase in AI talent in banks in the period May to September 2023, clearly showing that a significant investment is underway. Also, UK Government research[4] cites Finance and Insurance as the industry with job roles most exposed to AI vs all sectors. In summary, all the evidence is indicating that AI is a huge disruptor in the finance industry.

> ARTIFICIAL INTELLIGENCE IS A HUGELY DISRUPTING TECHNOLOGY THAT FINANCIAL SERVICES FIRMS NEED TO QUICKLY ADOPT

So, everybody is doing it, but what are they actually using AI for? A better question would be "*what should I use AI for?*". Again, as a business leader, you don't need to understand how AI works, it's more important to focus on outcomes that can be addressed with AI. Use case examples range across front, middle, and back office:[5]

(I)   Credit tasks like assessing credit risk, predicting default, or assigning a loan approval probability are all being worked on by banks already.

(II)   Sales and Marketing teams can be augmented with AI to help identify new clients, or product opportunities, and to uncover insights about your existing clients' purchase behaviours that will allow you to optimize your products and services.

(III)   Enhancing the customer experience with AI is becoming easier and easier, and I expect banks to start using generative AI-based chatbots for client interaction in 2024.

(IV)   Trading is a pure data industry, so lends itself perfectly to AI. Almost every step can be enhanced with AI. You can understand the market better and predict price movements;[6] algo-traders have been doing this for years, and the technology is getting better and better. Specialist chatbots can eradicate trade support by capturing trader discussion and converting it into formatted instructions for the traders to agree. Error-checking deals, issuing SSIs, etc. The list goes on and on.

(V)   Automation and operation efficiency are emerging as areas where generative AI will be able to help a lot, but for now you'll need to focus on use cases where the pay off is enough to cover the cost of developing a bespoke model. That said, the toolkit is expanding daily with low code, RPA, co-pilot support, etc., already making a big difference. A recent paper presented at ICAIF 23 showed very promising progress going beyond the prescriptive behaviour of RPA to a more spontaneous approach currently only possible with human operatives.[7]

(VI)   So much of risk management is finding a needle in the proverbial haystack, and human minds are not well optimized to this task. However, AI never gets bored or distracted, so can offer much better outcomes, if well managed. Systems can detect anomalies in KYC data, or fraudulent behaviours, detect sanction busting schemes, etc. This technology is already relatively mature in cyber risk, where practitioners have always needed to trawl vast amounts of data quickly. Companies like Chainalaysis have been using AI to examine blockchain data and identify criminality, and their techniques are just as relevant to the traditional finance space.

(VII)   There are many self-serving use cases where AI can be deployed to assist in the preparation of data for use in model building, for example by cleaning data. Synthetic data is emerging as a powerful tool that allows you to create a synthetic set of data that has the same characteristics as your production data but is anonymised and cannot be traced back to identifiable clients. This data can be used to test external vendors' capabilities, ensure that no production data is used in your development regions, or to allow academic collaborations that don't need strict NDAs.

(VIII) Although the return on investment may not be as high, many firms are building AI-based services to serve their internal use cases as this is a safer way to build experience with the technology. Generative AI fronted chatbots that allow semantic search of the firms' proprietary data, policies, procedures, etc., is a popular use case; as are more hidden use cases aimed at insider threat or cyber risk behaviours.

A general trend is that AI use cases assist a process, not replace it. The art of successful AI deployment is being able to make a mental leap from a business process that may be sub-optimal, and to then realise that AI could be deployed to enhance it. When you consider the return on your investment, consider if a 10% improvement in the outcome would make a meaningful difference. 10% optimization should be easily achievable for most AI-augmented processes, and many will do significantly better than this. Over time and with greater experience you can move to full automation, but starting with process assistance is the right way to start.

There is not one particular domain of finance that has an advantage over another when it comes to AI. There are advantages to reduce risk, identify alpha, increase efficiency, reduce cost, increase revenue, etc. All things that improve the way financial services works, ultimately providing a better outcome for the end consumer.

> THE IMPACT WILL BE FELT IN ALL ASPECTS OF OUR BUSINESSES, NOT JUST IN THE QUANT TEAMS WHERE IT HAS BEEN USED FOR YEARS

Dream big but accept that AI is not a magic bullet. Some use cases are not solvable. Your data may not be large enough to train a model, or may not be clean enough – there could be irreparable bias issues that will produce unwanted outcomes. Some outcomes, like predicting security prices, are beyond AI's ability to predict, as there are just too many variables. Consider, however, if you really need to predict a spot price, or would it still be valuable to predict the direction of price movement and maybe even the magnitude. In many cases, understanding the nature of the output probability is valuable in itself.

HOW DO I INTRODUCE AI TO MY FIRM?

Now you know what AI is, and understand that it is strategically important, let's discuss what it takes to build an AI capability in your firm.

There are different operating models for AI that can be adopted, depending on your firm's strategy and goals. In financial services most firms don't have widespread AI experience yet, so it is advised to build a centralised multidisciplinary capability to lead your AI effort. The cost of these teams should be considered too, where larger firms may be able to build the full capability in house, other firms may decide to outsource aspects of the function to third parties. It doesn't matter if your operating model is in-house, outsourced or hybrid, the same components need to be considered.
- Strategy – First, strong executive sponsorship is vital; as building an AI capability is a significant challenge, and is likely to face many headwinds before it is successful, the sponsor should be able to provide air cover to the team while they focus on the capability.
- Data Scientists – These are the specialists with the technical knowledge to build AI models from scratch. These specialists are highly trained and are therefore also highly sought after. Ideally you would have some data scientists in-house, but there are many firms offering data science as a service. 'Try before you buy' may be a good strategy to get your AI capability started.
- Risk, Legal and Compliance – AI presents new challenges that need to be understood so that you can build within your agreed risk appetite. Model Risk specialists have existed for

years generally looking at the work of quant teams or central planning teams, but their experience is vital to AI success. Likewise, data risk people are needed to ensure that the data that feeds your models is being used appropriately, and that the output does not have unacceptable biases. Ethics is another significant risk where AI brings specific concerns. Depending upon your budget you will decide if you hire new staff, or upskill existing staff, but the costs of ignoring AI risks will be very significant as regulators are paying very close attention to how AI is being used by firms.
- Technologists – Your existing IT team have a large part to play to ensure that hungry AI models have access to the data they need and have the compute available to run. They will decide if an 'on prem' or cloud architecture is right for your needs; both work, and there are distinct advantages to both too. One oft-ignored role is that of data engineer, as these people are expert at wrangling the data into a format that a data scientist can gain value from. Many data scientists report that a significant part of any project is taken preparing the data. As you scale it is prudent to consider assigning this task to lower paid data engineers rather than letting expensive data scientists do it for you.
- Business Experts – The last set of people you need are business experts who understand what AI technology is capable of. These experts then engage with your business leaders to understand their challenges and scout for opportunities where AI can add value. Finding use cases that are valuable to your firm is the most important task that any AI capability must achieve, without this you run the real risk of wasting time doing AI theatre.

> THERE ARE NEW RISKS TO CONSIDER, BUT THE ROUTE TO ADOPTION IS NOT INSURMOUNTABLE

A popular operating model is a 'hub and spoke' operating model, where you build expertise centrally and define standards so that teams in your business can build on that to solve problems. This gives you consistency, risk management, scalable tooling, but also business flexibility to focus on the challenges most important to each team. Obviously, this is not the only operating model that you could consider. It could be more appropriate to embed AI capability in the product lines, or to focus on one aspect of your business that will generate the most gains by leveraging AI.

The US National Institute of Standards and Technology (NIST)[8] does a good job of describing what good AI risk management looks like. Standards are coalescing around the world: the EU AI Act will form a set of crucial guidelines in Europe, and other regions are following a similar path. NIST looks at risks through three lenses: risks to people, organisation and to ecosystem. AI risk goes deeper to consider these categories in terms of data risks including privacy and bias,

legal risks covering who owns the data and what contractual protection you have for using it, model risks covering if the model does what you think it does, cyber risks, cost risks, strategic risks, ethics risks, etc. AI risk is a complex space, and it is well worth upskilling your risk management function so they can help you achieve the strategic outcomes you desire safely. For a deep dive into AI Risk, an OECD report[9] published in 2021 covers the space very well.

'Garbage in, garbage out' is a phrase often repeated when discussing an AI data pipeline. The good news is that the finance industry has been doing a lot of work over the last few years improving the quality of its data and AI is the payoff. Having good data governance practices converges with the data science talent and techniques to suddenly make the impossible possible. This is why AI technology that has been around since the 1960s feels so new. Cloud computing is also an accelerating factor, as it allows you to manage your data at scale and connect to data science solutions more easily. Well-managed data could be 'on prem' just as easily but there are advantages to cloud computing, for example scaling and cost. However, it's easy enough to build an AI model that uses 'on prem' data as it is to use cloud-stored data.

> THE SKILLS AND KNOWLEDGE OF EXISTING EXECUTIVES IS CRITICAL AS THEY WILL IDENTIFY THE MOST IMPACTFUL USE CASES TO WORK ON

Undoubtedly your firm already has experienced technologists who can leverage your data and deliver AI solutions, but it is likely you'll need to do some work to adjust your risk governance and run some education both in technology and within the business teams. One big mistake that happens too often is that leaders delegate AI to teams who are not equipped to implement it. Technology tends not to have deep understanding of the business goals or close enough relationships. Risk and compliance people are often too cautious with a technology they don't understand. Data people naturally obsess over data structures and governance, but don't tend to have the magic touch needed to build an AI capability.

Building a better data architecture and growing deeper understanding of data and AI in your risk teams will lead to better cyber risk too. There is a naïve view that AI will 'get away from us', yet I believe that the very well-regulated financial services industry is well positioned to adopt this technology,

because we already have the muscle strength of good governance. Robust risk practices and strong oversight are traits that are aligned with AI success. When we pair that strength with a desire for meaningful business outcomes it becomes a powerful driver for success.

It is hard to define what it takes to build this AI capability; however, the leadership needs to be aligned very closely to the business. To identify good AI use cases, you need people who can discuss business challenges articulately with your business leaders and clients, but who also have a deep enough understanding of AI to identify when a challenge should leverage AI to build a solution. Typically, at this point the AI capability engagement team will bring their data scientists to the discussion to dive deeper and validate the proposed solution. People who can stand in both data science and the business are the ones you should prioritise identifying as these are the people who will make or break the AI capability at your firm.

Hopefully you now understand what AI is, its value to your business and how you may go about creating an AI capability at your firm. If you already have a capability, you can use this article to assess if your firm is on track. AI is likely the largest single disrupter to financial services in a long while, both revolutionising how we run our firms, but also how our clients interact with our products and services. AI-enabled firms will naturally focus on the quality of their data, which will lead to further efficiencies, and they will develop more agile change practices that when paired together with AI will make them very hard to compete with. This is therefore a critical time to invest for the future health of your firm.

Notes
1   Maple et al, The AI Revolution: Opportunities and Challenges for the Finance Sector, The Alan Turing Institute, 2023
2   The impact of AI in financial services: opportunities, risks and policy considerations | Policy and Guidance | UK Finance), 2023
3   Evident AI Index, Banks Key Finding Report, November 2023
4   GOV.UK Impact of AI on UK jobs and training (publishing.service.gov.uk), November 2023
5   Evident AI Outcomes Report, October 2023
6   Gu, S., Kelly, B., Xiu, D. Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2233-2273, 2020. DOI: 10.1093/rfs/hhaa009
7   ICAIF '23: Proceedings of the Fourth ACM International Conference on AI in Finance November 2023 Pages 73–81
8   AI Risk Management Framework | NIST, January 2023
9   OECD (2021), Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers, https://www.oecd.org/finance/artificial-intelligence-machine-learningbig-data-in-finance.htm.