



White paper: AI and the data-driven economy

By Mark McKee and David Wilks, 2nd April 2025

Abstract

What does your financial services business need to know about the data it has, how to monetise it and, crucially, harnessing AI in a way that is sustainable and can be governed? The problems that businesses are grappling with are firstly about multiple, distributed and disparate data, so for example, getting a full picture of a customer across asset classes is hard. Secondly, AI is hugely over-hyped, so how do you get beyond the marketing to what can be useful to your business? In this brief paper, we examine 1) how to pull together disparate sources of data so you can get a unified view of your customer and be in a position to view patterns across customer segments that will help you position your products and services more effectively; 2) Whatever stage one is at in the AI adoption curve, governance is critical and needs the partnership, input and support of compliance leaders.

Background

People in financial services, especially technology leaders, have been trying to solve for a consolidated view of data across all manner of sources for years, for example, from trading systems, order management systems, time-series capture platforms like KDB, CRM sales platforms, settlement and payment systems. These firms have invested heavily in multi-year 'data lake' initiatives that have yielded very little return for all the effort and sunk costs. Often this has involved the use of buzz-word tools like Hadoop, and streaming data. People may have, at times, succeeded in setting up a data lake, but it has been so exhausting, that once in place, little has been done except check the box that all the areas requested have started submitting their data. What is missing is turning that data into intelligence.

Next up is the challenge of getting to the heart of what forms of AI can be useful to a business. If you have large amounts of data stored, what sort of use cases can be implemented to train AI on so as to review the results and determine how much to invest further in this emerging technology? We also need to separate automation over AI, as people try to hype their automation as flashy AI, which is just marketing fluff. A good example is a chat agent. That isn't AI: it is just code to execute predefined tasks that interact with you unsupervised. They are run on a set of rules that determine their behaviour. For example, a customer service chatbot on a self-invested pension plan (SIPP) app interacts with a customer who is transferring in an external pension. In this use-case the chatbot is using a set of rules to solve specific problems within well-defined tasks relating to a user ID, policy number, provider and fund value to gather what it needs from the customer so they don't have to sit on the phone. This example doesn't have the ability to generalise knowledge or learn from the experience it has had with the customer. That has simply automated a task. The same goes for KYC (know



your client). It's a relief to just need to take a photo of one's passport and a selfie for digital image processing to quickly validate that the ID and photo relate to the same person. The customer onboarding experience is greatly sped up for getting up and running with a financial app.

Fulsome use-cases for AI come along with examples such as when your data is training a software agent to furnish a wealth advisor with an investment suggestion for a client that fits within their existing risk appetite, market performance of products that they are interested in, or being able to suggest new products that a more functionally specialised financial advisor may not be aware of, but is within the regulatory jurisdiction of the client and approved by the investment entity for selling to the client.

We don't believe that the compliance function should always be playing catch-up, whether it be for regulatory, surveillance, AML, fraud and associated risks, or seeming like it is playing the role of bad cop to curtail the innovators from doing what they do best. The most successful firms have realised that compliance and innovation go hand in glove and that in order to get to market fast, a true partnership with designing, building/testing and validating new AI technology products facilitates a more confident release to production and ongoing feedback loops.

Solution

We need to start with what business use-cases would benefit from AI. What would it look like for a particular business outcome to happen? What data would you need? How do you measure success? What risks do we need to understand? Then you can work backwards to ascertain if you have the required data at hand. If not, is it somewhere that is not currently joined up to the rest of the data that would give the full picture for the use-case? You don't need a multi-year fundamental data architecture paradigm change to achieve this. Start small and prove out if you can gather the data you need for the identified use case. There has to be a financial plan in place so you can work out the investment decision. If it is too costly for the projected revenues (i.e. your NPV is less than zero) then abandon the idea and think of something else. This is about being a learning organisation and trying out new ideas that don't risk a massive failure. If you're going to fail, better to do it small scale! We recommend that implementing the AI model itself should not be done by experienced tech folks either. Why? Simply because they will bring their enterprise software engineering patterns to the task and want to build a battleship when all you want is to test out an idea. Bring interns or graduates to the problem as they not only come with a fresh perspective but have already been learning this kind of tech at university and will have been using tools that are open source, keeping costs low. They can work with experienced product managers to guide them on the specific domain, helping to understand what problem they are solving and evaluating risks such as whether it will work as intended and if they have the skills and training to do the work.

We also need to consider governance around AI, and this is of particular importance for heads of compliance functions. How do you make sure you are not trying to shut the stable door after the horse has bolted? The key here is collaboration. By getting involved



at the use-case level, you can be inputting ideas early and bringing your expertise to view the training model output to ensure compliance across various angles of risks. There is a huge amount of value to building compliance into a model than trying to retrofit it later! The result should be no surprises when it comes to reviewing the AI product prior to releases and that there are explicitly documented risks and their mitigations that are understood by all stakeholders. AI is never going to be risk free, but having guardrails in place will allow the firm to exploit the exciting business opportunities this new technology offers whilst having a clear but evolving set of governance principles for firms to draw upon.

Conclusion

Cutting through the hype surrounding AI to understanding what business cases lend themselves to AI are key. Any undertaking for funding and building an AI initiative needs to be treated the same as any other software investment: if it doesn't bring a net benefit then you need to look at a more productive way to deploy investment capital. All of this is underpinned by data that can be drawn together and also ensures that it complies with regional rules according to different jurisdictions. Finally, compliance leaders are necessary for the successful inception and iterations of strong commercial AI implementations that not only help your customers, but are great for your business and its reputation.

References

McKinsey Digital: *Next-gen banking success starts with the right data architecture*, 28 February 2025, (<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/tech-forward/next-gen-banking-success-starts-with-the-right-data-architecture>, accessed 31/03/2025).

AI Magazine, *Top 10: Risks of AI*, 21 August 2024, (<https://aimagazine.com/top10/top-10-risks-of-ai>, accessed 31/03/2025).

David Wilks, *Data in government*, 13 July 2017, (<https://dataingovernment.blog.gov.uk/author/davidwilks/>, accessed, 1/4/25)

Further information

Mark McKee is a founder of Agile Mind, based in London, UK. For more info, please contact mark.mckee@agilemind.uk