

USE CASES: GEN AI

The genie is out of the bottle

Generative artificial intelligence has dominated headlines and captured the imagination of the global business community with its promise of enhancing human creativity, boosting productivity and revolutionising communication. But what use cases have been implemented in the FS sector to date, and what does the future hold for this fascinating technology?

“There is much hype around generative AI [gen AI] at the moment. It’s the number one topic that financial services firms discuss with us,” says Dr James Bowden, Senior Lecturer, Accounting and Finance, Strathclyde Business School.

“Firms want to know how to incorporate it into their systems to help simplify compliance, and engage in horizon scanning,” he explains.

Despite significant enthusiasm for gen AI’s potential, there are few current use cases in the FS sector, and those that do exist are relatively narrow in scope. “There are hundreds of gen AI projects going on, but the overwhelming majority are still in the development phase,” he continues. “The current number of use cases is actually quite small.

“Most are low-risk use cases used in internal systems, to help improve operational efficiencies, extracting and summarising key information from the wealth of qualitative financial information that exists. But we’re talking about gen AI as a decision-support tool rather than a decision-making tool. I suspect that’s going to be the case for quite a while to come,” he adds.

Searching for use cases

A recent report by Goldman Sachs, entitled *Gen AI: too much spend, too little benefit?* highlighted that more than \$1tn is set to be spent on gen AI in the near future, but questioned whether this investment would pay off given the high costs involved and the lack of a decisive use case for the technology.

However, Bowden believes that gen AI can make a significant impact on the FS sector. He is a co-investigator at the Financial Regulation Innovation Lab research programme and has been exploring applications for simplifying compliance using AI, working in collaboration with FinTech Scotland and the University of Glasgow.

“We’re working on the development of a prototype of a large language model [LLM] that’s integrated with a large vision model [LVM]. It gives the consumer a kind of a visual avatar they can communicate with that helps them explore and summarise complex terms and conditions,” he explains.

Such an application has the potential to help FS firms comply with regulation such as the Financial Conduct Authority’s (FCA’s) Consumer Duty act. “This would vastly improve the consumer experience. Given the emphasis the FCA has placed on the Consumer Duty, you can see how this could be potentially extremely beneficial,” says Bowden.

“When it comes to readiness to implement AI technologies, banks have multiple advantages.”

Wendy Redshaw,
Chief Digital Innovation Officer, Retail,
NatWest Group

Overcoming barriers to adoption

The rise of AI-generated deepfakes, a particular risk in the FS sector, is just one example of several factors that can potentially undermine trust in gen AI if the issue is not brought under control. However, gen AI also has the potential to combat financial fraud. “We’re also building models for detecting financial fraud such as deep fakes, because the genie’s out of the bottle now,” adds Bowden.

Another challenge is the question of the reliability of gen AI models, which are prone to the phenomenon of ‘hallucinations’, where apparently factual information is, in fact, a total fabrication. This can include references and citations pointing to non-existent sources (a particular problem in academia, Bowden observes), invented names or fictional biographical data, to name but a few examples. A recent real-world example of an AI hallucination is that of a New York attorney who used an AI chatbot to perform legal research. The result was that they accidentally cited bogus legal precedents that did not exist.

“Hallucinations are an inherent limitation in these gen AI models,” continues Bowden. “There are ways to address them to some extent, but they are always going to exist,” he warns. One step that can be taken to overcome the potential unreliability of gen AI outputs is to engage in a process known as ‘grounding’, whereby gen AI outputs are constrained by pre-defined information.

Bowden points to the example of Morgan’ Stanley’s new gen AI tool, announced in September 2023, which enables financial advisers to rapidly consult the bank’s database of more than 100,000 financial documents.

“Morgan Stanley using gen AI to query its internal documents is an example of ‘grounding’, which involves connecting the outputs of gen AI models to verifiable data sources, which limits the potential for hallucinations,” explains Bowden.

One of the reasons hallucinations can pose a problem for the use of gen AI in the FS sector is that it raises the risk of falling foul of a growing range of AI regulation, such as the UK’s proposed Artificial Intelligence (Regulation) Bill and guidelines set out in its 2023 White Paper as well as wider legislation such as the EU’s Artificial Intelligence Act, which came into force on 1 August.

“The issues of explainability, trust and fairness are featured quite heavily in the UK’s AI framework,” says Bowden. “These issues need to be addressed if we are going to see these models used in more external-facing processes and client-facing environments. I think that’s why FS firms are treating these matters with caution.”

Human error vs machine error

Some of the challenges presented by gen AI, such as hallucinations, may be technical in nature, but accountability could ultimately rest with the humans using these tools. “The FCA’s Consumer Duty places an emphasis on fairness, explainability and traceability,” says Bowden. “But, ultimately, it’s the human that is responsible for these issues rather than the model, so there needs to be accountability there.

“And we’re not just talking about regulation, there’s also the potential for reputational damage caused by gen AI. We’ve seen instances before where machine learning [ML] has been used for

lending decisions and there have been biases in the models. And that’s led to vast amounts of negative publicity.”

However, it’s worth considering that the risks posed by gen AI may be similar to those posed by other technologies. “There are comparisons here with technologies such as self-driving cars, for example,” says Bowden. “There are so many accidents on the road every day that are a result of human error, but it takes only one error from a self-driving car and it’s all over the news.”

This tendency for negative publicity to stifle innovations – rightly or wrongly – may be another reason gen AI use cases are still relatively uncommon. “Human error takes place every day, but we don’t read about it as much as if an algorithm makes a mistake,” points out Bowden.


While the barriers to the widespread uptake of gen AI are real, he doesn’t believe that they are necessarily insurmountable. “If you think about the amount of investment that is going into this space, if it’s going into the right areas and is addressing the deficiencies in these models, I can see a scenario where these models are incorporated into decision-making roles within the next 10 years or so,” he predicts. “If we can get it right, I can’t see why we won’t have a pipeline of extremely useful gen AI tools coming through.”

Let’s see what Cora says...

One example of a bank at the forefront of implementing gen AI is NatWest, whose Cora chatbot is getting a gen AI upgrade in the form of Cora+. So how is the bank leveraging gen AI for this new iteration and how can the tech deliver even more responsive services to customers?

“In 2017, we introduced Cora as a 24/7 AI-powered digital assistant to answer day-to-day banking queries and carry out simple instructions,” explains Wendy Redshaw, Chief Digital Innovation Officer, Retail, NatWest Group.

“Over time, Cora has evolved through natural language processing [NLP] and machine learning [ML], and has incorporated more advanced capabilities, including voice recognition, to enhance the customer experience,” she explains.



USE CASES: GEN AI

- ▶ “In 2023, Cora helped 10.8 million customers resolve their banking queries – so this constant evolution of our ‘digital companion’ made the iteration to include gen AI alongside ‘traditional AI’ as an obvious one to explore,” she adds.

“With gen AI, customers can ask questions and receive responses in a more natural, human style,” she continues. “For example, previously, when a customer asked Cora about a mortgage or lending product, a link would be provided to a helpful but general page, and the customer would then perform a degree of research themselves from the different options within the page. Now, Cora+ will be able to understand the context and nuances of each customer query and will provide more accurate and personalised responses for certain journeys.”

“Human error takes place every day, but we don’t read about it as much as if an algorithm makes a mistake.”

Dr James Bowden,
Senior Lecturer, Accounting and Finance,
Strathclyde Business School

Balancing opportunity with risk

Despite the need to take the development of gen AI at a cautious pace, Redshaw thinks banks could be well placed to benefit from this technology. “When it comes to readiness to implement AI technologies, banks have multiple advantages,” she says. “The laser focus on customer benefit alongside data, risk, governance, and a highly proactive cadence of releases means that we’re well set to learn quickly and deploy at pace,” she explains.

“There’s been a huge step forward in terms of readiness, and much of that has come from building on existing competencies – for example extending existing data and ethics principles to include AI.”

NatWest has also been focusing its gen AI applications on areas where such implementation is appropriate and well considered. “There is also the ability and need to be discerning on appropriate

use cases in banking due to the rigour of our governance,” says Redshaw. “For example, some customer journeys need to be demonstratively consistent from a regulatory perspective whereas others can have a degree of flexibility.”

While NatWest is aware of the potential benefits of gen AI, the bank is also mindful of the need to balance opportunity with risks. “As with any new opportunity, the introduction of a new approach, model or technology is subject to formal processes to evaluate and manage risk well within a bank,” she stresses.

“However, it is vital that these processes are enhanced by a strong culture – one that has properly internalised what customer benefit looks like, and is capable of creatively scenario-planning, debating risks constructively, and then working transparently and collaboratively across business lines to mitigate them.”

As well as putting risk management processes in place, Redshaw believes that employees should be supported to better understand the uses and limitations of gen AI. “As with any emergent technology, we must be mindful of privacy, ethics, and the risks associated, ensuring we have the right safeguards in place for our customers and our colleagues,” she says.

“This must be front-and-centre in our employee education too. With gen AI appearing so ‘human-like’ in its responses, it can be easy to make assumptions about the level of intelligence and understanding that it has,” she cautions.

“Education is an important aspect of helping people to understand the technology – and, specifically, what it can, and cannot, be expected to do,” she adds. “Testing assumptions, thinking outside the box, and exploring ‘what-if’ scenarios can all help to guide understanding and appropriate use of the technology.”

Caution, care and consideration

Given the complexity of applying gen AI within a banking context, how long will it be before other banks follow in the footsteps of NatWest’s gen AI journey?

“We’re happy to be leading the industry on this, but we’re mindful that the pace and breadth of development in this field is fast, and probably will get faster,” says Redshaw. “It feels like the industry is still at the stage of getting all the cards out on the table – looking at what ‘could’ be, before getting to a more focused view of what ‘should’ be, rolled out. And this makes good sense.

“My expectation is that other banks will follow in our footsteps; however, my hope is that any gen AI roll-out is undertaken cautiously, and with appropriate consideration – not just in terms of appropriateness of use case, and protection of customers or colleagues, but also in terms of the potential environmental impacts of gen AI, with its higher energy requirements.”

While the gen AI revolution may not be upon us just yet, Redshaw believes that more use cases are soon to follow. “In 2023, we identified more than 100 priority use cases for AI – and in particular, gen AI.”

She concludes: “We anticipate that this innovative technology will offer a few transformative ‘hit shows’ and yield many others that focus on improving the way that existing products and processes provide enhanced capability or more efficient support for our colleagues and customers.” **CB**

