17 96

LOMBARD ODIER
INVESTMENT MANAGERS

AI in financial services

– For professional investor use only –

09/23

September 2023

Key points

- One should not be blinded by the hype surrounding AI but look at the real opportunities ahead and the implementation of it in our daily lives. It is not a mythical wonder, it is technology which has been developed over decades and gets better at executing tasks as a result of greater computing power and deeper datasets.
- AI within financial services must be viewed very differently from other sectors. The “job to be done” is a difficult interaction between objective and subjective inputs and outputs and around those variables, regulatory requirements need to be incorporated too.
- Generative AI, which is the latest hype, is not likely to have a disruptive impact on the financial sector, rather will be used to achieve efficiency gains. These models are known for their tendency to “hallucinate” which drastically limits its usability in applications that require fact-based outcomes, such as those within financial services.
- We expect that this new enthusiasm for AI will spur data-collection solutions. Some financial institutions haven’t sorted out the data layer yet, let alone the tools to use on that data.
- The entire AI market, including services and hardware, is expected to reach USD 900 billion by 2026, which would imply a CAGR of 19%. Financial institutions represent between 20 to 25% of this total market. Machine learning is by far the biggest driver of growth within the financial sector and the most likely places of deployment are still expected to be fraud prevention and customer services.
- We do not invest in loss-making hyper-growth stocks just because they put the term “AI” in their quarterly update 20 times. We look deeper at how AI is truly impacting company fundamentals.
- AI-exposed companies are now trading near 1999 bubble valuations. There is a big disconnect between AI and the rest of the technology space, which is nowhere near dotcom-era bubble valuations.

Opportunities for AI in finance are abundant

p.09

The financial industry has been using AI for many years already and is implementing this in a wide array of use-cases.

For professional investor use only. Please read important information at the end of this document.

This document has been prepared by:

Jeroen van Oerle
Portfolio manager, Global FinTech

Christian Vondenbusch
Portfolio manager, Global FinTech

For further information please visit
Professional | Lombard Odier Investment Managers

Introduction

Artificial intelligence (AI) is more commonplace than many people realise and exists as technological enhancements of everyday items. AI handles the face ID process that unlocks smartphones. It curates suggested content for streaming and social media. Social media, Netflix, and Spotify[1] suggestions are one big AI continuously trying to keep the user engaged to the platform by learning from their behavior to show new content that might keep the user engaged. Emails and Google searches use natural language processing (NLP) to correct mistakes and find out the context of words. Digital voice assistants (Siri, Alexa, Google) are all AI and there are many more examples of our daily interactions with these tools.

Within financials services, there is also a lot of AI behind the scenes. Every payment transaction made is checked by an AI for fraud. Opening or logging into a digital bank account deploys tons of Ais that monitor behavior and 'learn' preferences. Credit score providers use AI based on proprietary datasets, insurers use computer vision and deep learning for automatic claims handling and every asset manager that has quant capabilities is an active AI user and has been ever since programmable software was used to engineer data turned into investment factors.

Computer scientist John McCarthy, one of the founding fathers of AI, argued "as soon as it works, no one calls it AI anymore." And that holds very much true for all the examples provided above. This is not the first time AI has dominated the news agenda and it won't be the last. Rather than take the hype at face value, it is important to consider the real opportunities ahead and think about how AI manifests in our daily lives. It is not a mythical wonder that appeared from the skies, it is technology which has been developed over several decades and that gets better and better at executing tasks as a result of greater computing power and deeper data-sets. This is not an attempt at being dismissive of the tremendous progress that has been made in this field, but rather a down-to-earth view on the technology and its implications for users.

This report takes a broad approach to a diverse topic. We first introduce the concept of AI and examine its disruptive potential by means of the disruptive innovation framework created by professor Clayton Christensen. We then focus on its applications within the financial service sector. AI within financial services must be viewed very differently from other sectors, given the difficult interaction between objective and subjective inputs and outputs and around those variables, as well as the regulatory requirements.

Generative AI, which is currently generating headlines, is not likely to have a disruptive impact on the financial sector and will rather be used to achieve efficiency gains. AI in the broader sense, however, has the power to determine long term winners and losers. Identifying these winners means looking at company fundamentals instead of marketing slide-decks. Valuations are approaching dotcom bubble territory and extreme expectations are setting the stage for another AI winter when expectations inevitably disappoint. Investors who want to benefit from this trend in the long run need to know how to sort hype from reality.

FIG. 1 AI BUBBLES

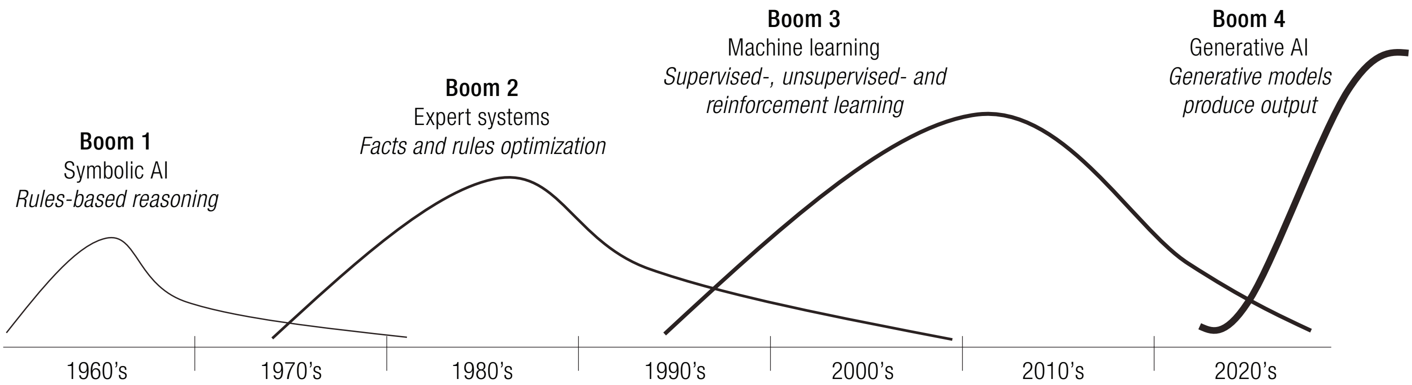

Source: LOIM, 2023.

[1] Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

AI defined

"AI systems are machine-based systems with varying levels of autonomy that can, for a given set of human defined objectives, make predictions, recommendations or decisions using massive amounts of alternative data sources and data analytics referred to as 'big data'"

- **OECD, 2021**

A good definition of AI consists of two important elements. The first relates to the technology itself and its decision-making objectives. These range from supervised learning models, whereby humans determine the inputs and judge the outputs, to completely autonomous reinforcement learning with deep learning capabilities that only require the user to set parameters to accept or decline the outcome. The second part of the OECD definition is related to the input source that allows the technology to be developed. We will discuss the characteristics of good data by looking at the four Vs, as developed by IBM,[2] in the next section. AI and big data go hand in hand. Without big data, there would be no AI because there would be no possibility to learn from the data to apply those learnings out of sample. Vice versa, there would also be no possibility to interpret big datasets without the help of AI.

Generative AI is a subset of artificial intelligence.
Generative AI is the latest offshoot to capture the public's imagination, in the form of stories about ChatGPT[3] or Deepfake. Generative AI is a subset of AI, as shown in figure 2. It differs from traditional AI which is cognitive in nature and uses analytics and strict boundaries in an analytical context to generate a desire doutput.

Generative AI is perceptive. It includes a creative process where the output is not limited to a set of opportunities but is created by a set of inputs. This technique is mostly used in creative applications like music, poetry, audio, video etc. and, ultimately, in the communication to humans and programming.

Generative AI uses Generative Adversarial Networks (GANs) to generate content. In this case, there is a generator node that comes up with the output and a discriminator node that evaluates each output. The discriminator is trained in a fixed dataset and can distinguish between real data and AI generated data. The discriminator node will 'approve' an outcome if it can no longer make that differentiation, hence new content is created. Very importantly, this new content is not necessarily based on facts. It is a creative outcome with 'true' data as a starting point.

Generative AI models are known for their tendency to 'hallucinate' (generated content that is nonsensical or unfaithful to the source). This is a very important factor to keep in mind, as it drastically limits its usability in applications that require fact-based outcomes, such as those within financial services. This also differs substantially from other AI techniques which do not suffer from this issue.

AI has been around since the late 1950s and consists of many sub-fields. Robotics, Natural Language Processing, perception: computer vision, planning, social intelligence, and machine learning.

Figure 2 shows how generative AI uses elements from deep learning, natural language processing and computer vision. The implications for software models are significant.

"Software 1.0 was deterministic code written by humans, function by function. Software 2.0 is machine learnings, where the human part is to check and label the data. Software 3.0 is generative AI which is out of the box where the basis data is already provided and the starting point is not a new dataset, but rather an expansion of an existing training model in order to sophisticate the software. That saves time and brings costs down."

- **Sarah Guo, founder of Conviction, an AI focused venture capital firm**

[2, 3] Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

FIG. 2 AI LANDSCAPE

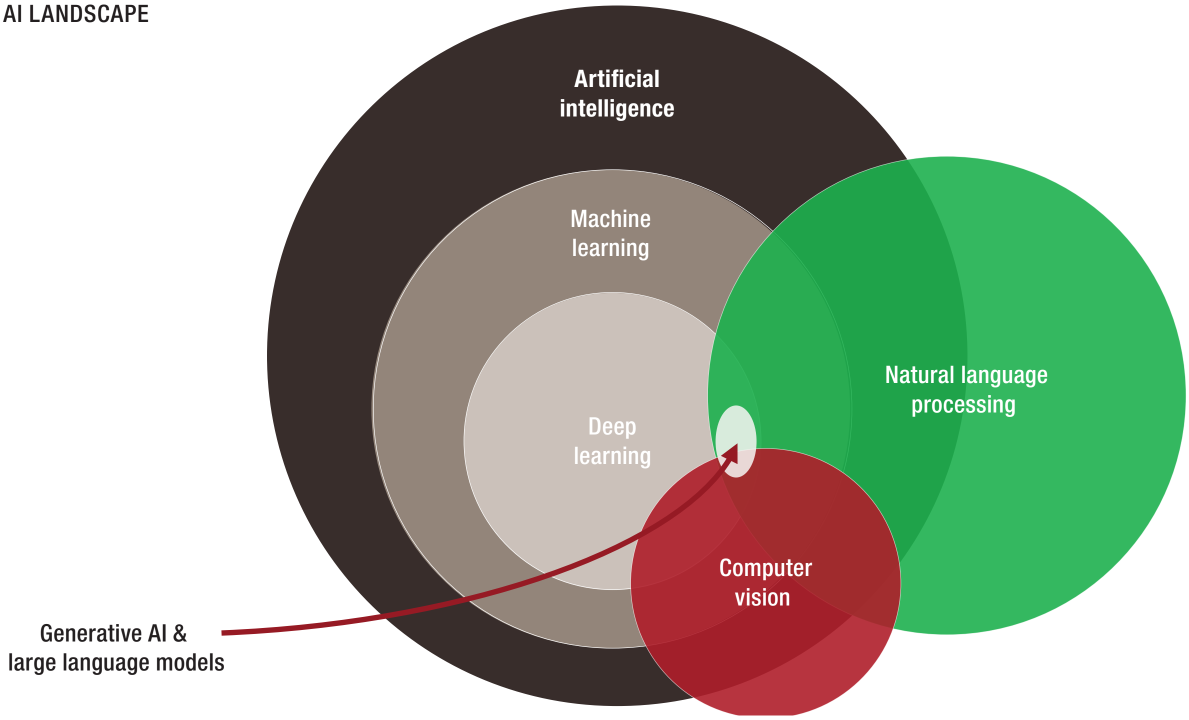

Source: Col Jung, 2023.

Generative AI is a useful tool for creative industries and has the potential to simplify programming. However, this is not useful in a situation where it is important to receive a precise, error-proof output, which is a requirement in the financial industry. For example, a suggestion by a generative AI that circumvents investor protection rules within asset management is unacceptable. Similarly, wrong advice (by generating non-proven outputs) can have very damaging impacts on users and the reputation of those who provide the advice. Manipulative transactions, initiated by autonomously operating generative AI, would result in the loss of trading permits by its users. The tendency for these models to hallucinate is not to be underestimated, as the model can confidently state, for example, that revenues of an analysed company increased, while in fact those revenues decreased. Or in case of credit scoring, a generative AI model could hallucinate by making up home ownership of a customer, thereby influencing the credit score.

It is not only the output of generative models that require attention, but also the input . Data-input rules prevent using open-source generative AI tools within financial services, as that client-data ends up with the owners of the tools, which is prohibited in most countries due to privacy rules.

Consequently, we do not believe generative AI will be as disruptive in finance as it can be in creative industries. This paper considers other AI tools which are already having an impact in the financial industry and have the potential to generate long term winners.

Big data and computing power forms the base of all AI techniques

Without data, there would be no use for AI. ~~The~~ Figure 3 shows the amount of data in existence. Just for reference, a Zetta bye equals 1024 exabyte, which equals 1024 petabyte, which equals 1024 tera byte, which equals 1024 gigabyte, a unit of data which most people are familiar with. 1 Zettabyte equals roughly one trillion Gigabytes. The world generated 97 of those in 2022 and it is estimated that this number will grow to 181 zettabytes by 2025. In terms of storage, that number is much bigger, as the numbers in figure 3 refer to annual data generation, while data centres store decades of data.

FIG. 3 GLOBAL DATA GENERATED ANNUALLY

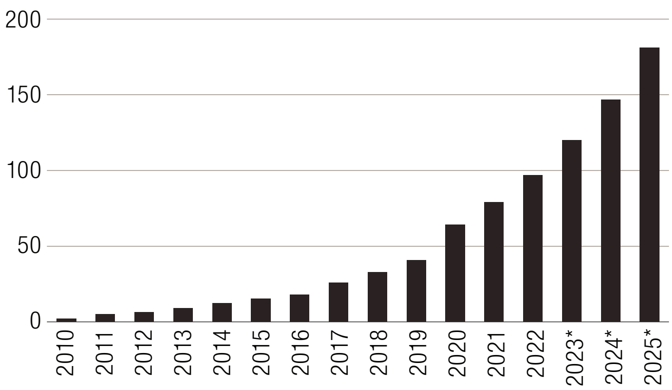

Source: explodingtopics.com, 2023

FIG. 4 THE 4 Vs OF BIG DATA BY IBM

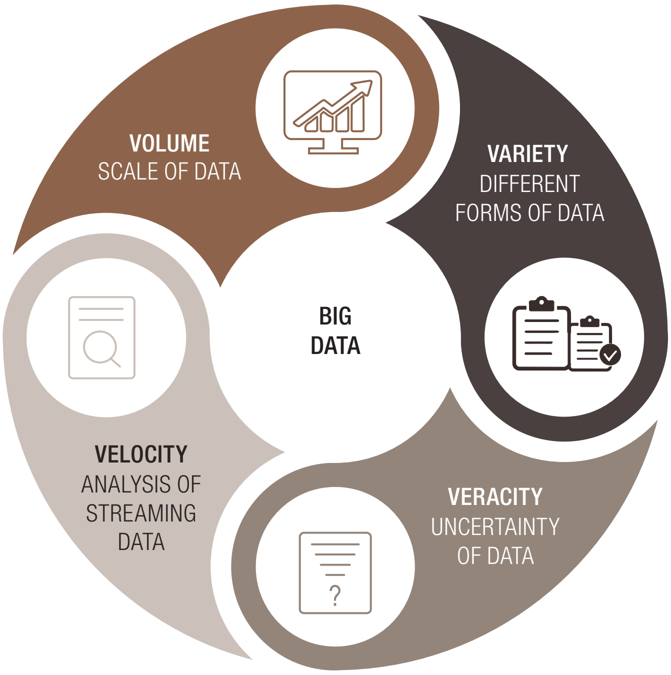

Big data 4 Vs. Source: [QAsource, 2017](#).

This amount of data requires what IBM[4] has termed the 4 Vs (figure 4).

1. The first component is volume. The **volume** of data grows every year and needs to be collected, stored, cleaned, secured and analysed. The storage alone has proven to be an important hurdle for financial institutions, as most companies do not have a data-lake solution available today. Data is compartmentalised and scattered across the organisation and its many IT systems. The first step in working with AI is to make sure that the internal data is handled efficiently. We expect that this new enthusiasm for AI will spur data-collection solutions. This is where we see big opportunities for consultants. Companies like Accenture, Tata consultancy and Capgemini[5] derive a substantial part of their revenue growth from data-management projects in the financial sector. This also shows that some companies are still in the preliminary stages of AI because they haven't sorted out the data layer yet, let alone the tools to use on that data.
2. **Velocity** is also important. This refers to the speed at which new data is generated. Within the financial sector this differs substantially. Data updates regarding new mortgages, for example, is less frequent than ticker data for stocks.
   The higher the volume, the more data is generated and the larger the training-set for AI tools.
   Different techniques require a different amount of data. Reinforcement learning techniques, such as deep learning, require more data than supervised learning techniques with labeled data, for example. Generative AI also requires a lot of data to train itself. The higher the velocity, the faster new techniques can be implemented and tested in real-life situations. This implies that it will take a long time before these techniques can be used in datasets with limited data spanning over multiple decades, for example macro predictions.
3. **Variety** is another important determinant and is classified as unstructured, semi-structured or structured data. Structured data refers to clear labels, categories, storage locations and selection criteria for the data in a dataset.
   The opposite of structured data is unstructured data which just consists of all the data pooled together in a data-lake without labels or further context.
   Semi-structured data is the final form of big data, where the data is clustered, but within that cluster there is no further categorisation of the data. The algorithms used to process data varies based on the variety.
   Within the financial industry, most of the data used so far has come in structured or semi-structured format. Payments, stocks, excel sheets full of customer data, underwriting data etc. has mostly been structured. The inclusion of large unstructured datasets is now quickly entering the financial sector .
   Generative AI can work with unstructured data and there are many different reinforcement learning algorithms that work with unstructured data. In an area like customer service, this is going to make a big impact. It differs if a customer needs to type an exact phrase with a labeled answer, for example, "I forgot my password" or a text from which the AI model infers this concerns a forgotten password, for example, "I cannot access my account and I have no clue why because the username is correct." In the latter part, there has been no mention of the word password, but since the username is correct, the issue might be related to a forgotten password. This unstructured data (since there is no record of this exact phrase with a pre-defined answer) is used by, for example, generative AI models to then add context and provide a solution to the problem. The financial sector is, and always has been, structured in nature. The fact that unstructured data can now be used is important, but also requires caution.
4. Finally, **veracity** is a crucial part of big data. This denotes the trustworthiness of the data. Not all the data gathered can be used instantly. There is no standardised approach to measuring the veracity of data, but the impact on businesses in cases of poor data quality is substantial.

[4,5] Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

FIG. 5 FROM CPU TO QPU

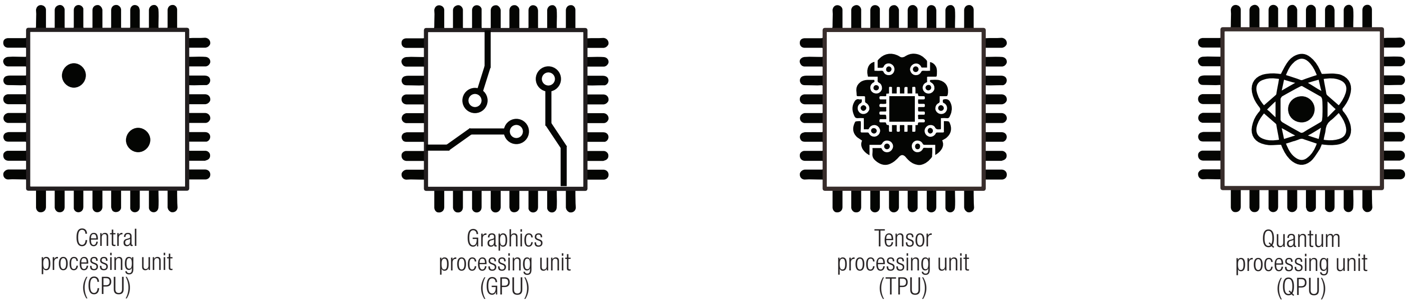

Source: CITI, 2023.

According to Gartner, the average cost of poor data per organisation is USD 9.7 million per year. IBM[6] estimates this adds up to USD 3.1 trillion in the US alone. If decisions are made based on the wrong data, it can lead to reputational damage, losses or fines. Veracity refers to biasness, noise and abnormality in data. It also refers to incomplete data, outliers or missing values. As we will elaborate further in the section on risk, veracity is one of the largest origins of risks when using AI models. This is also one of the reasons why a generative approach is only partially useable and should always be checked. Hallucination is a key issue known in generative AI outcomes and that, on its own, creates veracity issues which can spiral out of hand when there are no humans checking and correcting the outcome anymore.

Computing power innovation is the other component, next to big data, which fuels AI innovation. Advancements in computing power through clever usage of processing capabilities has been a massive enabler of AI. If we only had large datasets, but the training of AI models is slowed due to low processing speeds, it would make the AI capabilities much less usable.

Figure 5 shows the different stages of computing power over the last couple of decades. In the beginning, all computations and graphics were performed via central processing units (CPUs). In the 90s, the first graphics processing unit (GPU) was introduced in game consoles and in 1999 Nvidia[7] brought their first GPU to the market for usage in personal computers. CPUs and GPUs work together to increase the number of calculations in an application. The main difference between CPUs and GPUs is that the former is designed to handle a wide range of tasks (the speed at which it does so is referred to as the clock speed), but this cannot be done sequentially. A GPU is designed to perform parallel operations on multiple datasets, but for a narrower range of tasks than a CPU.

This is why CPU and GPU are combined in most applications, where in the classical case the GPU would take care of the graphics (single task, parallel input/output), and the CPU would be the muscle running the rest of the application.

However, in the early 2010s GPUs started to be used to accelerate calculations involving massive amounts of data. This was such a large improvement for AI techniques that it made a whole range of AI applications available which were previously not possible. Although not designed for AI purposes (rather for graphical purposes), the GPU proved to be a gamechanger for AI. After the AI industry started to use GPUs more for machine learning purposes, Google developed the Tensor Processing Unit (TPU), which is specifically designed for machine learning tasks. TPUs are tailored to perform tensor operations, which are the core building blocks of neural network computations (deep learning). Despite the fact that TPUs have a narrower range of tasks, the performance in neural networks is much better than GPUs. Another benefit of TPUs over GPUs is that TPUs are much more energy efficient. When scaling up machine learning capabilities, that is an extremely important component to take into consideration (more on this in the risk section).

The next step within processing power would be quantum computing. Although this is still far from being readily available, quantum computing holds great promises in the field of statistics, in particular massive parallel computations within scenario testing (advanced Monte Carlo). This could, theoretically, enable yet another range of AI applications.

[6, 7] Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

AI in finance is not disruptive in nature

Technology alone cannot be disruptive. The business model utilising that technology needs to be deployed in a disruptive way in order to make it so. In this case, it is important to differentiate between three groups of innovation in terms of their disruptive impact on an industry. Figure 6 illustrates these groups.

Let's take the example of the asset management industry in combination with AI. If the technology enhances the existing services (like customer reporting, KYC/AML or more efficient back-office operations) of the asset manager, we talk about efficiency innovation. When the technology creates new functionalities within the existing asset management business model that were not possible without the use of this new technology, we talk about sustaining innovation. These could, for example, be new trading algorithms or new portfolio construction algorithms. When the technology serves clients outside of the asset management industry (for example, those who are financially excluded now) by non-traditional asset management providers, we talk about disruptive innovation if that technology ultimately makes it into the traditional asset management service offering and starts to compete with the traditional asset management business model. Robo-advice and cryptocurrencies started out like this. Regulatory requirements, though, pulled these new business models back within the traditional asset management functionalities. This is very specific for the financial industry. It turns out that this holds true for many technologies deployed within the setting of the financial industry.

FIG. 6 CATEGORISING THE TYPES OF INNOVATION

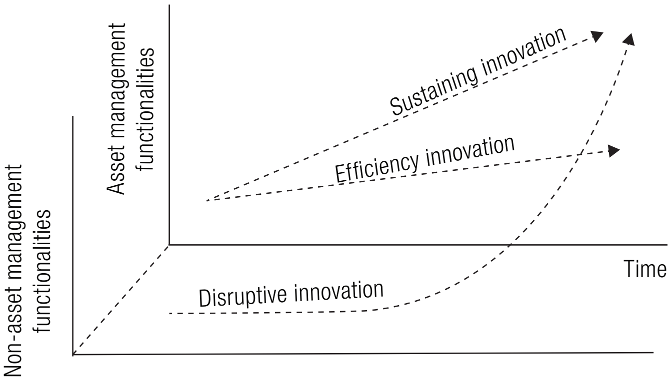

Source: C. Christensen, LOIM, 2023.

Incumbents are the most likely long-term winners when it comes to sustaining and efficiency innovation because they already have a market position and simply need to integrate the new technology into their offering. Newcomers are the most likely long-term winners of disruptive innovation, because incumbents pay no attention to the new technology (because it is outside of their traditional functionalities), until it is too late and the disruptive innovators start to target their markets. A clear example of a successful disruptive model is Apple's introduction of the iPhone. Apple was not in the mobile phone market (which at the time was dominated by Nokia). Their iPhone model served customers outside of the traditional feature phone market, as the functionalities of this new technology were very different from the traditional product offering. Internet access from the phone and an application ecosystem (the app store) turned out to be exactly what customers wanted and that new technology disrupted the entire feature phone market, with the newcomer, Apple as the long term winner.

We view AI in finance as non-disruptive. We would categorise most AI tools used in finance as either efficiency innovation or sustaining innovation. This is profoundly different from other industries, whereby the disruptive elements are much more visible. For example, within the music industry, it is very likely that non-musicians are now using the tools available to create business models that disrupt the traditional music industry/record labels/platforms. There is similar potential for many other creative industries, such as arts, writing and film. The costs associated with creating content have drastically decreased, thereby allowing new entrants to use the technology to create new business models which can disrupt incumbents in those industries. In an even more extreme case, the programming industry is being disrupted (at least at the low end) as non-programmers can now write code by engaging with generative AI models.

Within the financial service sector, the most important element preventing similar conditions to the creative industries is regulation. There are little to no rules (besides copyright) which prevent the creation of artistic content. There are a lot of rules, though, on financial outputs and on those who are allowed to provide that output to customers. For example, a fully AI generated asset manager would still have to get a license to manage money. Obtaining such a license for people outside of the industry is extremely hard to do. Therefore, the technology is much more likely to be integrated within the industry itself, by those who have the licenses already.

A good example of this point applies to crypto. Crypto started outside of the asset management industry and had very disruptive elements for the asset management industry. Slowly but gradually regulators are drafting rules to curb the freedom of those who offer these disruptive technologies, and as a consequence, crypto is likely non-disruptive as incumbents use the technology and integrate it into their existing service offering. If it wasn't for regulations, it is very likely that the disruption of crypto, and other technologies, would have been much larger in the financial industry.

AI is currently used within the financial industry to improve efficiency and to implement sustaining innovations. Efficiency improvements are not likely to be proprietary, rather than off the shelf. This implies the consequences in terms of winners and losers are not evident. We believe generative AI falls into this category. It has potential to improve efficiency, but that can easily be copied by peers in the industry. The moat around these new technologies is very narrow and if, for example, bank A is able to improve margins by integrating a generative AI model into its customer service, the management team of bank B will make sure to quickly adopt that model and implement tools in their own service offering. The costs of doing so fall within the operating expense domain and are the responsibility of respective managers to implement in their part of the business.

Sustaining innovation has potential to create winners and losers. The innovations in this category have a wider moat and shift from the operating expense line to capital expense. In the case of AI, this implies proprietary models and datasets. For example, within payments, large payment network operators like Visa and Mastercard[8] have their own models to prevent payment fraud and have trained those models on their own datasets. In the past there have been many attempts to disrupt those networks. Most notably, the Merchant Customer Exchange (MCX), which was launched in 2012. Ultimately that alternative failed for several reasons, the most important one being that fraud prevention on the payments was much less efficient than that offered by the incumbent networks in the US. Due to the higher fraud rate, the cost savings from operating a payment network that does not require interchange payments to Visa and Mastercard were nullified by the write-offs coming from fraudulent payments.

It is clear now that this technology is about much more than just the generative AI hype that started with the announcement of ChatGPT. The financial industry has been using AI for many years already and is implementing this in a wide array of use-cases. We view AI in finance as sustaining innovation and efficiency improvement. In the next section, we will take a look at current use-cases of AI in finance and the potential use-cases going forward.

8 Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

Opportunities for AI in finance are abundant

AI is used in the financial sector already. Not many people talk about it and the companies are not continuously marketing it with every earnings-release, but it has an impact on current business from fraud reduction to customer service. This is not related to generative AI, which is estimated to be roughly 0.40% (USD 1 billion) of total AI investments,[9] but to the broad concept of AI in general.

FIG. 7 AI USE CASES IN THE FINANCIAL SECTOR

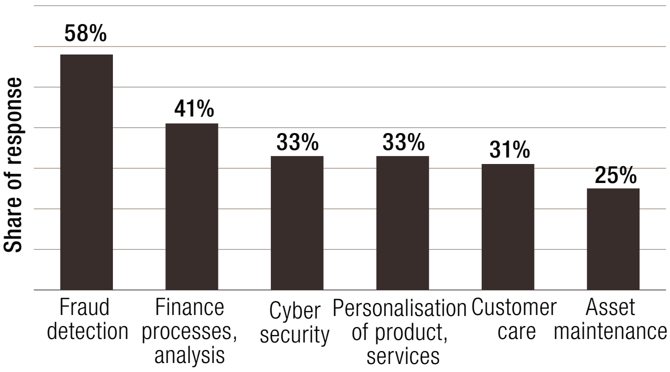

Source: AI In Finance: Business Applications, Benefits & Future Outlook 2021.

Figure 7 shows the response rate of financial institutions in 2021 when asked how they deploy AI capabilities. Generative AI was not even a topic two years ago, yet the industry was already implementing AI techniques within fraud, data analytics and cybersecurity. Personalisation and customer care are elements which can be further approved upon by implementing generative AI techniques in the years ahead.

AI consists of classical techniques (logic, planning, expert models, NLP) and modern techniques (machine learning, deep learning, data mining, and quantum computing). There are many techniques in use today in the financial industry as can be seen in figure 8. The potential impact of AI in finance going forward is twofold. First, it stands to increase efficiency by streamlining processes and secondly, it has the capacity to create winners and losers based on who has access to the best datasets and the best AI tools to utilise those datasets for revenue generating purposes.

FIG. 8 AI APPLICATIONS IN FINANCE

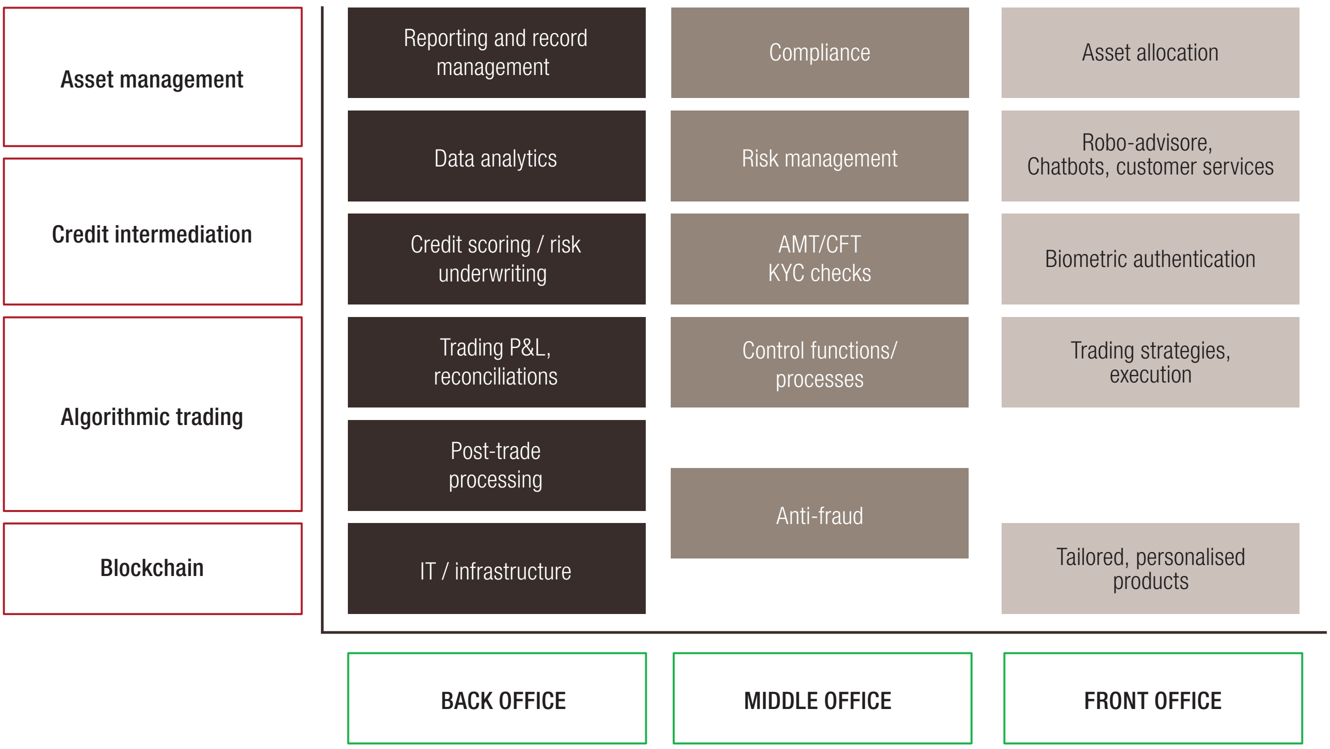

Source: OECD, 2021.

[9] Source: appquipo.com, 2023.

To provide the reader with an example of how advanced the integration of AI already is within financial services, we would like to focus on Visa,[10] a global payment network provider. Visa is clearly not alone in terms of how it integrates AI into the payment process, but it serves as a good example. The company has invested over USD 10 billion in cybersecurity over the past five years, with about 10% of that budget allocated to data analytics and AI. Visa, as at the end of 2021, has 60 petabytes of data, which is a huge source of analytics. On top of this big dataset, the company deploys more than 60 AI services to spot and block fraud on its network. One of those services is called the Visa Advanced Authorization (VAA) score. Machine learning is used to determine whether a transaction is legitimate or fraudulent within 300 milliseconds. The maximum payment transaction capacity for Visa is about 65,000 transactions per second. Going over these transactions and coming back with a fraud check within 300 milliseconds is impressive. According to Visa, VAA prevented USD 26 billion worth of fraud in 2021.

As can be seen in figure 9, Visa uses data to analyse patterns in order to determine if a transaction fits with the behavior of the user. Fraud prevention is a trade-off between risk and return. False positives (such as a transaction wrongly declined due to suspected fraud) are extremely costly. Aite Group estimates that the e-commerce industry experiences false-positive losses of USD 443 billion in 2021, a much larger number than the projected fraud losses of USD 6.4 billion. An AI model in payments, therefore, needs to balance fraud prevention and loss of business. If the model is too strict, the loss of business will increase. If the model is too loose, the fraud losses will increase. This requires very big datasets and innovative machine-learning models to optimise. Hundreds of millions of authentication requests are verified for millions of unique users per year. This data consists of behavioral data, location data, device data, social data and many more relevant datapoints.

There is a lot of value in payment data to all participants in the payments ecosystem. This data is very detailed and allows precise risk profiles and behavioral analytics. Although fiercely denied by Visa, it is often theorised that the company would be able to predict a divorce well before it actually happens just by analysing payment data. Late-night cab drives, gifts ordered online or last-minute hotel bookings can provide hints if such behavior would not fit the normal patterns gathered over the last couple of years. Visa does not track marital status, but it can be inferred from the data if needed. The fact that payment data is extremely detailed has also caught attention in big payment deals. When Alibaba[11] (a Chinese payment company) wanted to acquire MoneyGram in 2018, the US government blocked this deal. Payment data is often not allowed to be stored outside of the respective countries/continents where it was gathered.

Without the use of machine learning techniques, it would not be possible to go through all this data to build systems like VAA. New payment and privacy regulations in the EU (PSD3 and GDPR) and the US (ADPPA) create an interesting situation which is to the advantage of big tech companies. Upon the request of the data-owner (which is the client of a bank or the user of a big tech product or service), data must be shared. This implies big tech companies can get access to payment data if the user wants to. An example of where this is applicable is Apple Pay. If the user wants to use this service, the payment data is shared with Apple.[12] This provides big tech with a huge competitive advantage over financial institutions in the battle for the end-client. Big tech can build extremely detailed customer profiles by combining their own internal user data with the payment data insights. This can be used to serve those customers better, but also in a competitive way to win market share by integrating financial services in the product offering. We have previously written about embedded finance, but it is clear that the more data, the better the training set for an AI and the bigger the competitive moat.

The fact that large payment companies are very advanced when it comes to the use of AI for fraud prevention also shows in terms of the revenue profile. To stay with the example of Visa, financial statements show it earns about one-fifth of its revenues from "value-added-services." This includes fraud prevention. This is also the reason why it is very difficult (if not impossible) to replace these large payment networks. The fraud services alone would make it worthwhile to use and provides the payment networks with years of head start versus challengers. Even central bank digital currency (CBDC) requires fraud overlays to function.

[10, 11, 12] Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

FIG. 9 EXAMPLE OF DATAPOINTS USED BY VISA IN THE PAYMENT AUTHORISATION PROCESS

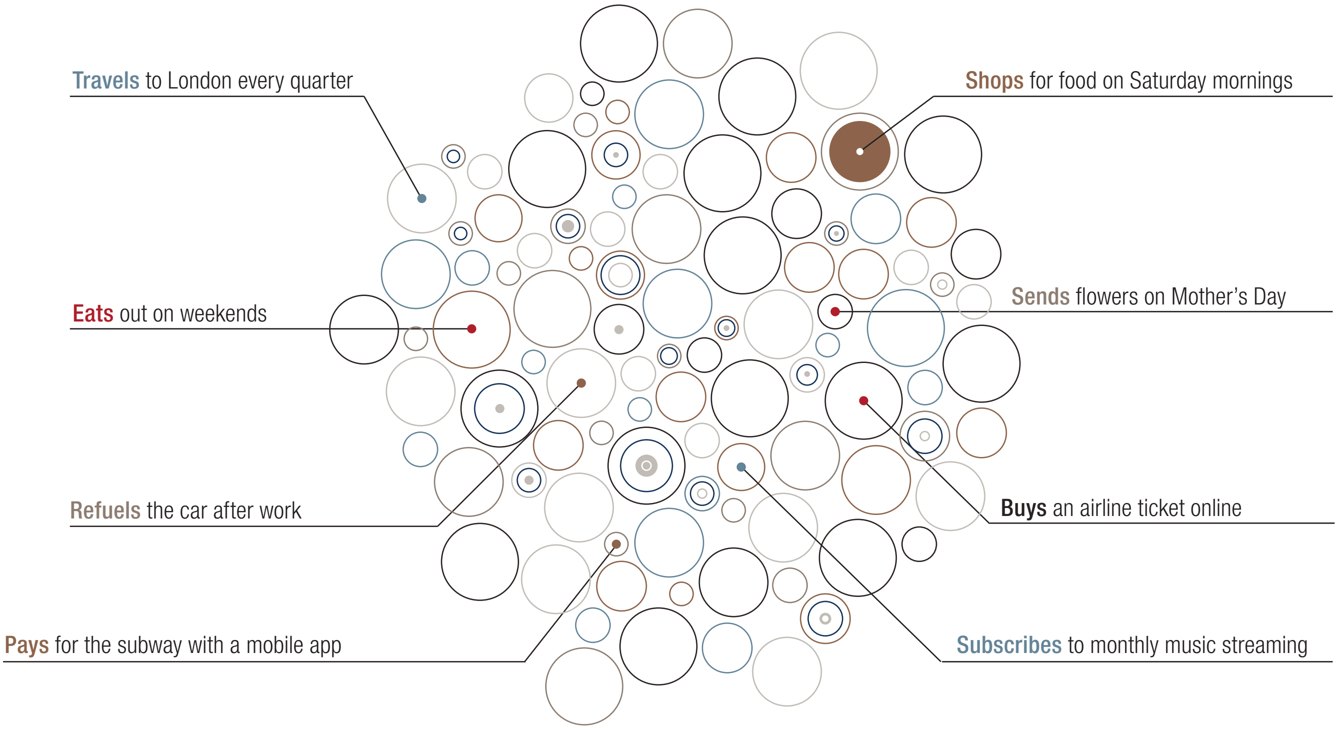

Source: VISA, 2023.

According to the Preventing Financial Crimes Playbook, Visa's investments in AI match those made by the overall financial services sector. This implies there are many other examples of financial institutions, as well as the payment industry, using AI in advanced settings. Financial institutions spent more than USD 217 billion on AI applications in fraud and risk-assessment in 2020. This figure has likely increased substantially with the integration of generative AI techniques in the customer-engagement process.

It is very clear that AI integration in the financial sector is much broader than just fraud prevention in payments. Banks also use AI techniques to prevent fraud in general (payments, deposit opening, money laundering, lending, FX, etc.). KYC/AML processes[13] are enriched by big data and AI techniques and insurance companies use the technology for underwriting purposes and claims handling. Think of, for example, computer vision and deep learning tools used in applications that allow photos to be automatically transformed into damage reports.

Returning to figure 8 and it is clear that the integration of AI techniques in financial services will continue to expand. Generative AI is an efficiency innovation which will likely be used to optimise customer engagement as well as programming, by means of allowing for interactive reports generation where the algorithm creates custom dashboards via interaction with the end-user.

Beyond generative AI, there is much more room for proprietary AI techniques and datasets. The example of Visa is a clear one. These big payment networks have accumulated a large amount of data over the past decades. These provide the perfect training set to optimise processes and create proprietary services with large defensive moats, as competitors would not be able to easily copy those insights. Similarly, large banks, asset managers and insurance companies have big datasets of their own as well. These will also be utilised to improve and personalise the service offering. However, to many financial institutions there is a long way to go still, as they do not have the basics required to make full use of AI techniques. That basis consists of data and methods of cleaning, labeling, storing and extracting that data. The first step for many financial institutions is to create data lakes which allow AI techniques to utilise all internal and external data required to provide the best output.

Risks in using AI in finance

It is important to note that AI is a man-made tool and not a god-like mythical bringer of the ultimate truth. This implies that, despite the massive progress made over the years, there are risks involved. Our focus is on those within the scope of AI in finance. Social consequences and the discussions around job-loss are relevant topics as well, but beyond the scope of this paper. Figure 9 is a perfect summary of some of the key risks associated with AI in finance.

[13] AML refers to all regulatory processes in place to control money laundering, fraud, and financial crime, while KYC is the risk-based approach to customer identification and verification that forms part of AML requirements.

FIG. 10 RISKS FROM THE DEPLOYMENT OF AI IN FINANCE

Governance and accountability

- Model governance arrangements
- Accountability and lines of responsibility
- Outsourced models of infrastructure

Policy frameworks

- AI complexity challenges technology-neutral approach (e.g. explainability, self-learning, dynamic adjustment)
- Potential incompatibilities with existing legal/reg frameworks
- Risk of fragmentation of policies (across sectors)
- Skills and employment

Non-financial risks (data, fairness)

- Biases, unfair treatment and discriminatory results (inadequate use of data or\poor quality data)
- Data privacy, confidentiality

Explainability

- Why and how the model generates results
- Inability to adjust strategies in time of stress
  - Amplify systemic risks, pro-cyclicality
- Incompatible with regulatory/supervisory frameworks and internal governance
  - Difficult to supervise AI algorithms/machine learning models

Robustness and resilience

- Unintended consequences at firm/market level
- Overfitting, model drifts (data, concept drifts)
- Correlations interpreted as causation
  - Importance of human involvement

Source: OECD, 2021. For illustrative purposes only.

The first risks stem from the data used to train the models. As discussed, data is the key input for AI learning, but poor quality data can allow biases in the data and discriminatory results, thereby potentially harming financial consumers. Credit agencies are often challenged by politicians and the media on the fairness of their data. Forbes argued in 2021 that current credit scoring models are based on incorrect/incomplete data and suffer from inherent racial bias. The arguments are historical in nature, whereby most of the data is derived from rich, white populations and their respective behavioral characteristics, resulting in a particular creditworthiness. Data shows, however, that 1 in 5 Black consumers and 1 in 9 Hispanics have poor credit scores today, which compares to only 1 in 19 for white people. Low home-ownership amongst those with poor credit scores due to the rejection of mortgage applications by financial institutions is the result. It should be asked whether that same result would have been achieved if the training set would not have been skewed to white home-owners that dominated the training datasets.

Missing data can have large consequences for the conclusions drawn by AI models. This is why explainability is an important risk. For example, a couple of years ago, Dutch research organisation TNO researched city crime to advise municipalities on what measures needed to be taken to make neighborhoods more secure. After an early version of the AI had modelled the solution, one of the suggestions was to no longer provide licenses to fish stalls in order to reduce criminality. This suggestion, for obvious reason, was met with skepticism. The researchers working on the model started to dig into the data to find out what had caused this conclusion to be drawn by the model. It appeared this was a classic case of 'missing data.' Fish stalls were a defined factor in the data, but communal spaces - 'squares' - were not. Therefore, fish stalls became the equivalent of squares, as that was the actual source of criminal activity rather than fish stalls being located at those exact same places. This just shows how carefully the output of an AI should be studied before taking action. It also shows that data-input is one of the biggest risks associated with AI. Therefore, the data must be checked, and the outcome must be explainable by those using it as a basis for their decision making.

Similarly, robustness and resiliency of the model are important. The difference between correlation and causality in data is an important one which is often overlooked (or conveniently neglected). Figure 11 shows the correlation between the number of master's degrees and box office revenues in the US between 1910 and 2020. As can be seen, these two are almost perfectly correlated. AI models are, in their basics, statistical tools. Correlation is good in most models, as there is a data-relationship between two sets. However, that does not say anything about the causality between the data sets. AI uses more data, as we have discussed above. This implies that the chances of finding perfectly correlated datasets increases as well. These so called 'spurious correlations' (correlations by chance, not by nature) should be treated carefully. In this case it is very clear to anyone that there is no causal relationship between the number of master's degrees and the number of people going to see a movie.

FIG. 11 SPURIOUS CORRELATIONS

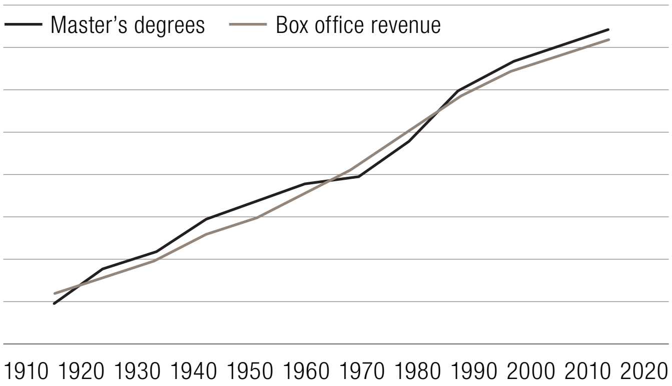

Source: Statology, 2022.

However, if figure 11 would have shown a near perfect correlation between the number of master's degrees in AI and the average outperformance of AI stocks, most would have assumed causality, because it sounds plausible. That might be a perfectly spurious correlation as well though, as other variables could have caused the average outperformance of AI stocks instead.

Finally, there are risks associated with policy frameworks and governance . In the end, companies are responsible for the output of their models and the actions taken on that basis. This implies the outcome needs to be within the regulatory requirements and there must be clear responsibilities for those who use these tools. One of the issues with generative AI, as previously discussed, is the tendency for these models to hallucinate. Within the governance domain, that has the potential to be risky. For example, if a chatbot at a bank advises clients to buy certain products or services which, in retrospect, turn out to be incompatible with the client's risk profiles or have outcomes which are completely different from the ones suggested by the chatbot, that has enormous legal and reputational consequences. This is why we expect supervised learning models to dominate. If deep learning techniques are used, the outcome would need to be triple checked and properly ringfenced before deploying the algorithm autonomously.

From a regulatory perspective, autonomous actions in finance, or worse DeFi (decentralised finance), would never be allowed. The lack of governance is central to that argument and a clear example can be found in trading. Humans tend to adopt herd behaviour, but there are always ways to break that pattern and to be counter-cyclical. The more we are driven by models though, the more we will emphasise those patterns. 'Spoofing', a technique for market manipulation that involves placing bids to buy or offers to sell securities with the intent of cancelling the bids or offers prior to the deal's execution - is an example of how patterns can be manipulated. These can, in fact, be executed by autonomous AI-powered algorithms which simply maximise profits by acting in illegal ways. If not properly governed, this can quickly spiral out of hand with many unintended consequences.

Regulators will demand more oversight in making sure the AI models used are not biased, that there is copyright protection of content and that the cybersecurity around the protection of (personal) data is ok. If these conditions are not met, regulators will enforce them. This is where big regional and sectoral differences come from. Asia has very relaxed data rules and privacy is not an issue. That means the datasets are large and AI models can be trained more quickly. In Europe the rules are very strict, which is good for citizens, but that can make AI progress much less efficient and quick as elsewhere in the world. The US is a bit in the middle, but the training set for many applications is less developed than in China. Similarly, because of all the regulations, the data which can be used within financial services differs from that used in, for example, the technology sector. As previously discussed, this has consequences for the use cases of AI in finance.

Cybersecurity is a different topic, but still related to risks in AI as well. AI is used in hacks and AI techniques to break cybersecurity at companies are being continuously developed. Companies will need to invest in their software to protect key data from hacks, but they also must prepare for scenarios in which case their data, AI models or other crown-jewels are compromised. We have written extensively on the subject of cybersecurity and our approach towards measuring what we call "cyber-hygiene." Unfortunately, 20% of listed companies do not comply with basic cyber hygiene in the sense that they have so-called "known exploited vulnerabilities" in their software. AI attacks focusing on vulnerabilities will single out these targets faster and it becomes even more important for companies to update their software and patch vulnerabilities. Hacking via 'phishing' techniques (where social engineering is used to enter proprietary company data) becomes easier as AI can be used to create profiles and gather data required to break security.

Finally, energy usage is a risk to consider. Just looking at generative AI, the size of large language models (LLMs) is growing exponentially. In 2019, GPT-2 had 1.5 billion parameters, which increased to 175 billion in GPT 3. The next version will have about one trillion parameters. This implies the processing power demand of neural networks is on course to become almost insatiable. The more computing power used, the more problems can be solved. However, energy usage can increase faster than model size.

> "If you look at the amount of energy taken to train a model two years back, they were in the range of 27-kilowatt hours for some transformer models. If you look at the transformers today, it is more than half a million kilowatt hours. The number of parameters went from maybe 50 million to 200 million. The number of parameters went up four times, but the amount of energy went up over 18,000X."
>
> - **Ian Bratt, ARM, 13D research 2023.**

FIG. 12 AI ENERGY CONSUMPTION VERSUS TOTAL ENERGY SUPPLY

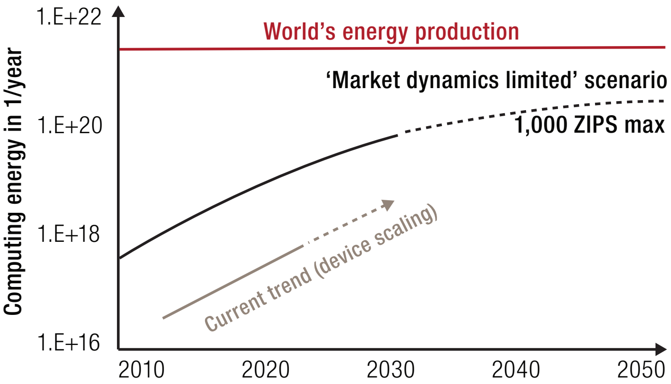

Source: AMD/Semiconductor Engineering, 2023.

Energy consumption will increase as AI models become more integrated in our day-to-day lives. After the models have been trained, they will be distributed amongst all available devices and 'consumed', which will keep energy demand elevated, even beyond the training period. Data centres will need to increase capacity to keep up with demand once AI services are fully deployed via the cloud. Luckily, chip design is improving drastically in order to get this footprint down. TPUs, as previously explained, are already more energy efficient than GPUs and chip designers are committing significant resources to improving chip design. At some point though, the capacity issue around energy usage by AI could stall progress. In the near turn this is not considered to be a large risk, but especially in a world where AI is embedded in everything, this could be a determining factor to keep in mind.

Estimated market size

As we now know, generative AI is a small subset of the larger AI market. Goldman Sachs estimates the generative AI software total addressable market (TAM) stand to be USD 150 billion, versus the global software industry TAM of USD 685 billion by 2025. Research by BAML estimates that generative AI currently represents about USD 70 billion in revenues, which amounts to roughly 3% of the total technology spend.

A reliable forecast of AI TAM is not so straightforward. It is important to separate the marketing stories from the reality. Exacerbated expectations have led to AI winters four times since the 1950s. We have no doubt AI will make progress in the years ahead, but we also know that markets tend to over-react. Especially since most market participants are, ultimately, not the people who have to develop and deploy these technologies in practice. The latest AI hype (which started around the end of 2022) could well be a situation of over-expectations followed by a realisation that reality does not unfold in one quarter after the announcement.

Generative AI in finance makes up about 1.5 to 2% of the total generative AI market. This shows that the financial sector is not necessarily the biggest user of generative AI. Customer engagement and programming would be the most important use-cases within finance.

Figure 15 shows the global AI software market is much bigger. Research by Bank of America from 2023 estimates the 2022 global AI software market stood at around USD 400 billion in 2022, and is on track to reach USD 787 billion by 2026. The entire AI market, including services and hardware, is expected to reach USD 900 billion by 2026, which would imply a CAGR of 19%. Financial institutions represent between 20 to 25% of this total market. Machine learning is by far the biggest driver of growth within the financial sector and the most likely places of deployment are still expected to be fraud prevention and customer services.

FIG. 13 GENERATIVE AI TAM AND AS PERCENTAGE OF TOTAL TECHNOLOGY SPEND

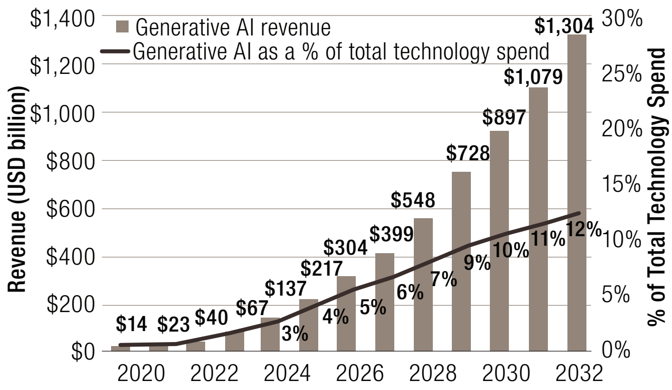

Source: Bloomberg intelligence, 2023.

FIG. 14 GENERATIVE AI TAM WITHIN THE FINANCIAL SECTOR

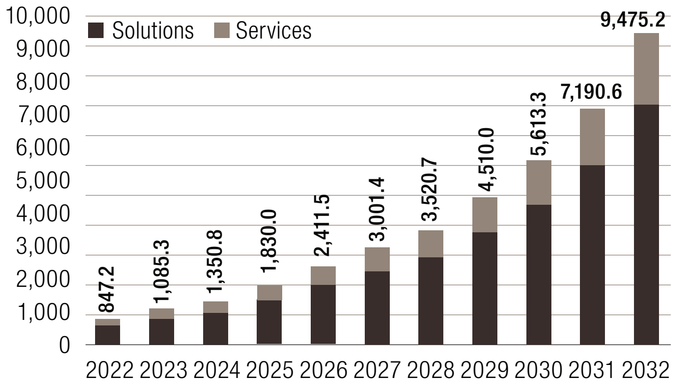

Source: Generative AI in Banking and Finance Sector , 2023

FIG. 15 GLOBAL AI MARKET SIZE

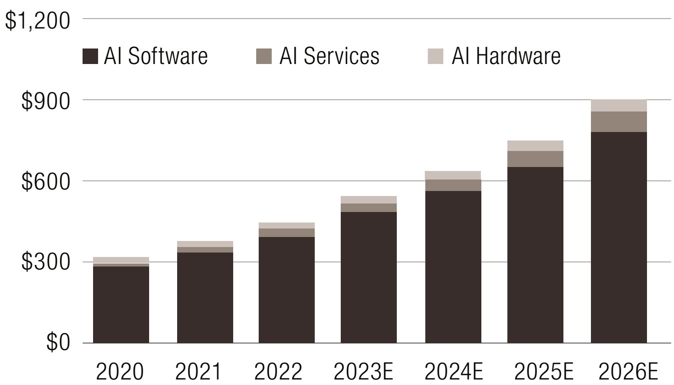

Source: BAML, 2023.

Portfolio positioning

Ultimately, everything discussed in this paper needs to be applied to portfolio positioning. We do this by adhering to the principle of ‘quality at a reasonable price.’ We do not invest in loss-making hyper-growth stocks just because they put the term “AI” in their quarterly update 20 times. We examine how AI is truly impacting company fundamentals. We prefer the picks and shovels approach instead of concentrating in a few deemed-to-be “winners.” We also look for those high-quality companies which integrate AI into their offering in order to drive revenues and improve margins.

The main reason for this approach relates back to the beginning of this paper. We think AI in finance is not disruptive but is rather sustaining innovation and efficiency improvements. This implies there will be winners and losers, but it is unlikely that newcomers take centre-stage. Those who have access to proprietary databases which can be used for AI training purposes have a natural advantage over those who start from scratch. Within FinTech, we expect to find those long term winners within the payments sub-segment and in software providers. Additionally, those who deliver specialised consulting services in order to implement AI and big data capabilities at financial institutions stand to win over the next couple of years. Especially within AI lending and insurance underwriting, we are skeptical about newcomers. First of all, there has been no real credit-cycle test so far, hence the training data is too skewed to a benign environment. Secondly, we think you can install all the AI you want, but it is all about the underwriting skills and risk management checks that determine if a company makes money in lending. Providing loans is not difficult, as everyone can do it. Getting the money back is the true challenge, and a few marketing slides on AI implementation are not going to change those basic fundamentals.

Judging from recent market behavior, we are one of the few who approach AI like this. Most participants follow a momentum-driven approach. This becomes very evident from an analysis of Credit Suisse HOLT data on stocks widely held in AI ETFs. As can be clearly seen in figure 16, AI exposed companies are now trading near 1999 bubble valuations. There is a big disconnect between the AI space and the rest of the technology space, which is nowhere near the dotcom era bubble valuations. Within FinTech we see the same. Companies which market themselves as AI and which are in the AI ETFs have performed extremely well over the past couple of months. However, there is no fundamental reason why these companies should trade where they do. Momentum, marketing and hype are setting the scene and it is very likely we will see a correction.

FIG. 16 HOLT VALUATION DATA FOR AI STOCKS VERSUS THE DEVELOPED MARKETS TECH SECTOR

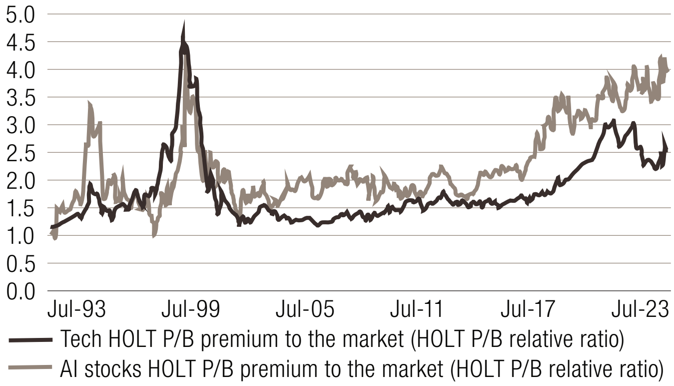

Source: HOLT, 2023

In the meantime, many investors don’t even realise that the biggest users of AI technology within the financial sector are available at extremely low valuations. These can be found within the payment segment. Proprietary databases and even proprietary chip design are enabling long term winners. The large payment network providers, such as Visa and Mastercard, are well positioned, but also lesser-known payment processing companies can drive margins by increasing their investments into AI. This is not something which will be done in the next quarter. This is something which is part of the long-term strategic positioning of a company. Finally, efficiency improvements will help increase margins or offset wage-pressure for the entire FinTech space. This differentiation between short term hype and long-term implementation is what separates hype-investors and trend investors.

Figure ~~16~~ shows our investible universe for FinTech. Roughly speaking, this universe can be split in one third payments, one third efficiency providers (which deliver software solutions to the financial industry in order for traditional financial institutions to make the transition to digital services), and one third for the blocks in the middle of the figure. Disruptive technologies capture innovations like blockchain, big data and AI. This is the investible universe. For inclusion in the portfolio, we only select the highest quality companies from this investible universe.

FIG. 17 SNAPSHOT OF LOIM FINTECH UNIVERSE

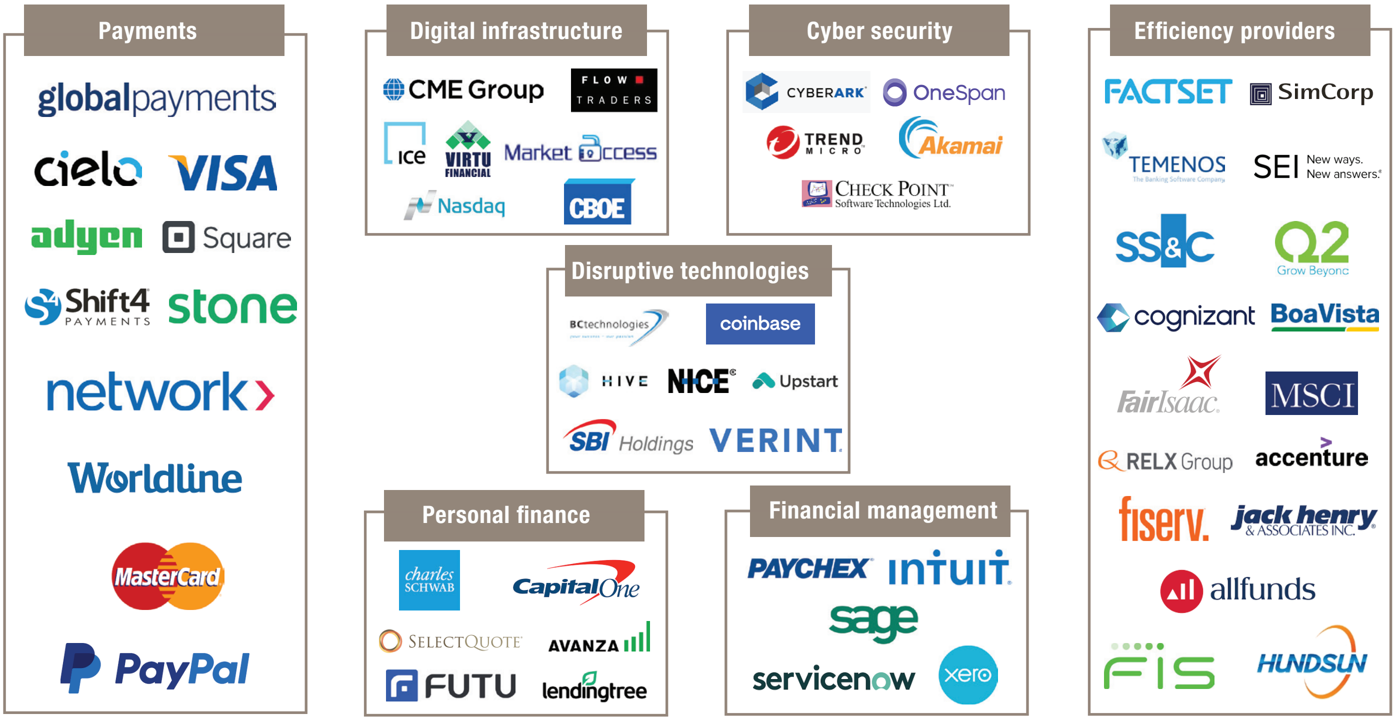

Source: LOIM, 2023. Any reference to a specific company or security does not constitute a recommendation to buy, sell, hold or directly invest in the company or securities. It should not be assumed that the recommendations made in the future will be profitable or will equal the performance of the securities discussed in this document.

With respect to AI, the exposure within our portfolio can be summarised as being over 75%. Most payment companies implement AI in fraud reduction. Companies like Visa and Mastercard are even able to offer proprietary fraud services which diversify their revenues away from classical infrastructure deployment. If this trend continues, over half of the revenue for these payment networks could be derived from value added services over the coming years.

Companies like Nice and Verint implement AI for financial services in their customer service offering. The large IT services providers, like Capgemini, Accenture, Endava and Epam are used throughout the AI implementation phase in the financial industry, ranging from the creation of database solutions to the actual development of AI techniques. Credit score providers, like FICO, use their proprietary database to train models in order to provide credit scores as accurately as possible, which are used in the underwriting process of banks (for mortgages, credit cards or personal loans) and insurance companies. Finally, there is whole range of efficiency providers which simply integrate AI techniques into their service offering. Not as stand-alone products or services, but as part of their overall service offering.

This brings us back to John McCarthy's observation that "as soon as it works, no one calls it AI anymore." This holds very much true for AI in FinTech. It is embedded in almost every part of the value-chain, but no-one calls it AI anymore. Until the generative AI hype started, and the term appeared in all sorts of marketing materials. To us, the marketing stories don't matter. We want to see fundamentals improve by implementing the latest technologies. The companies which manage this sustaining innovation best are long term winners and we try to capture as many of those as possible in our portfolio.

Important information

For professional investor use only.

This document is issued by Lombard Odier Asset Management (Europe) Limited, authorised and regulated by the Financial Conduct Authority (the "FCA"), and entered on the FCA register with registration number 515393.

This document is approved at the date of the publishing.

Lombard Odier Investment Managers ("LOIM") is a trade name.

This document is provided for information purposes only and does not constitute an offer or a recommendation to purchase or sell any security or service. It is not intended for distribution, publication, or use in any jurisdiction where such distribution, publication, or use would be unlawful. This material does not contain personalized recommendations or advice and is not intended to substitute any professional advice on investment in financial products. Before entering into any transaction, an investor should consider **carefully** the suitability of a transaction to his/her particular circumstances and, where necessary, obtain independent professional advice in respect of risks, as well as any legal, regulatory, credit, tax, and accounting consequences.

UK regulation for the protection of retail clients in the UK and the compensation available under the UK Financial Services Compensation scheme does not apply in respect of any investment or services provided by an overseas person.

A summary of investor rights and information on the integration of sustainability risks are available on https://am.lombardodier.com/home/asset-management-regulatory-disc.html.

This document is the property of LOIM and is addressed to its recipient exclusively for their personal use. It may not be reproduced (in whole or in part), transmitted, modified, or used for any other purpose without the prior written permission of LOIM. This material contains the opinions of LOIM, as at the date of issue.

Neither this document nor any copy thereof may be sent, taken into, or distributed in the United States of America, any of its territories or possessions or areas subject to its jurisdiction, or to or for the benefit of a United States Person. For this purpose, the term "United States Person" shall mean any citizen, national or resident of the United States of America, partnership organized or existing in any state, territory or possession of the United States of America, a corporation organized under the laws of the United States or of any state, territory or possession thereof, or any estate or trust that is subject to United States Federal income tax regardless of the source of its income.

Source of the figures: Unless otherwise stated, figures are prepared by LOIM.

Although certain information has been obtained from public sources believed to be reliable, without independent verification, we cannot guarantee its accuracy or the completeness of all information available from public sources.

Views and opinions expressed are for informational purposes only and do not constitute a recommendation by LOIM to buy, sell or hold any security. Views and opinions are current as of the date of this presentation and may be subject to change. They should not be construed as investment advice.