

Review | Received 22 November 2023; Accepted 19 March 2024; Published 28 March 2024
<https://doi.org/10.55092/let20240002>

Generative AI in finance: risks and potential solutions

Nydia Remolina^{1,2}

¹ Yong Pung How School of Law, Singapore Management University, Singapore

² SMU Centre for AI and Data Governance, Singapore

E-mail: nydiarl@smu.edu.sg

Abstract: Generative Artificial Intelligence captured the attention of academics, policymakers, the private sector, and some regulators in 2022 after the launch of ChatGPT and its widespread adoption worldwide. Then, during the World Economic Forum session in 2023, Microsoft Chairman and Chief Executive Officer Satya Nadella said that the ‘golden age’ of AI is underway and generative AI is set to play a big role in it. Accordingly, Microsoft has invested billions of dollars into OpenAI, the company behind the launch of ChatGPT. Additionally, Google and Meta have also created their own generative AI models. Given the multiplicity of potential use cases of generative AI models, the tech industry is not the only one interested in this development. Many industries were caught in the hype of generative AI, including the financial sector. Indeed, financial institutions worldwide are deploying generative AI and are starting with internal uses. However, along with that comes potential pitfalls and challenges. This paper argues that addressing the risks and challenges of generative AI requires a context and sector-specific debate. The challenges of using AI highly depend on the context in which the models are implemented and used. For example, generating content to provide financial advice has different implications for society than generating the image of a turtle. Hence, in the current debate about AI governance, and specifically in the generative AI space, it is important to analyse the risks of generative AI in finance. This paper is a first attempt to map the ethical and legal risks of generative AI in the financial sector given that previous scholarship and most policy discussions so far have focused on the risks of predictive AI in the financial sector. Based on this analysis and following a functional approach, the paper also proposes an avenue to deal with the risks created by generative AI in finance.

Keywords: generative AI; financial sector; AI governance; artificial intelligence; predictive AI; finance; AI regulation; financial regulation



Copyright©2024 by the authors. Published by ELSP. This work is licensed under Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.

1. Introduction

Generative Artificial Intelligence (Generative AI) captured the attention of academics, policymakers, the private sector and some regulators in 2022 after the launch of ChatGPT [1] and its widespread adoption worldwide [2]. Then, during the World Economic Forum session in 2023, Microsoft Chairman and Chief Executive Officer Satya Nadella said that the ‘golden age’ of Artificial Intelligence (AI) is underway and generative AI is set to play a big role in it [3]. Accordingly, Microsoft has invested billions of dollars into OpenAI, the company behind the launch of ChatGP [4]. Additionally, Google [5] and Meta have also created their own generative AI models [6]. Given the multiplicity of potential use cases of generative AI models, the tech industry is not the only one interested in this development. Many industries were caught in the hype of generative AI. Financial institutions are also deploying generative AI for use cases such as fraud detection, compliance, consumer-facing chatbots, and internal processes [7].

However, most of the debates around AI governance and AI regulation have been centred around predictive AI [8]. Such debates are often not based on a proper discussion and analysis of the risks created by generative AI models and its use in specific sectors. Consequently, the absence of a thorough discussion and analysis of the risks of generative AI may derail the development of the technology while failing to tackle the most important challenges these models present [9]. This has been the case in the European Union where the proposed AI Act considers generative AI as a *foundation model*. Providers of foundation models would have to comply with additional obligations, like transparency obligations, requirements relating to registration, training, design, and dataset standards, among others [10]. The scope of this approach results in extensive documentation obligations and audit requirements. However, these obligations are vague in terms of their meaning and how they relate to the risks created by generative AI. Likewise, Singapore’s Infocomm Media Development Authority has opened for public consultation a proposed model AI Governance Framework for Generative AI. This Framework aims to address Generative AI concerns while continuing to facilitate innovation across all private sector industries in Singapore [11].

Similarly, without first undertaking an analysis of sector-specific risks of generative AI, the United States Senate scheduled a first hearing of a series of hearings to help legislators design regulations around the ethical, legal and national security concerns of AI. In this first hearing, held on 16 May 2023, Sam Altman, the chief executive of ChatGPT’s OpenAI, testified that generative AI needs to be regulated [12]. He even proposed the formation of an agency that would license the most powerful AI systems and have the authority to take that license away and ensure compliance with safety standards [12]. However, the discussion was not context or sector-specific.

Regarding predictive AI in finance, regulators have adopted different approaches [8]. Some regulators have adopted a risk-based approach or a light touch approach to guide the industry depending on the materiality of certain use cases. For instance, credit scoring AI models are classified as high-risk models under the European Union AI Act [13]. Other regulators that have adopted a voluntary approach to AI governance have also emphasised

the importance of considering the materiality of the specific use cases in the particular sector. For example, the Monetary Authority of Singapore (MAS) launched the Veritas Initiative. Under this initiative, MAS suggested an assessment methodology to help financial institutions validate the fairness, ethicality, accountability and transparency of specific AI use cases such as credit scoring, insurance predictive underwriting, customer marketing, and fraud detection [14].

This paper employs a systematic literature review in the fields of finance, computer science and law. Following Kitchenham's approach for a systematic literature review [15] to identify the relevant literature, assessing quality, extracting and synthesizing data, the paper assesses that the scholarship on AI Governance and its intersection with Financial Regulation have not provided an analytical evaluation on the risks associated with generative AI in the financial sector. Also, using a doctrinal analysis, a method that involves a systematic examination of current regulatory approaches to understand the nature, development, and functioning of the law, rather than empirical methods or social sciences [16], the paper analyses and interprets the current approaches to AI Governance in Finance with a particular focus on generative AI. This type of analysis is typically desk-based to draw conclusions about the current regulatory approaches and its implications.

Moreover, the paper has used a comparative methodology [17] to analyse the proposed frameworks that address Generative AI concerns developed by the private industry and regulators across different jurisdictions. A comparative methodology in legal research involves systematically comparing regulatory approaches across different jurisdictions. This method seeks to understand the similarities and differences in legal approaches and to analyse why these differences exist. By comparing and contrasting, the paper makes normative assessments about the effectiveness of different legal regimes and propose improvements for regulators to consider in addressing the concerns of generative AI. Furthermore, our methodology extends to a functional approach in regulatory studies. Rather than a mere exposition of legal frameworks and institutional arrangements, we dissect how regulatory mechanisms can effectively mitigate the identified risks. This is achieved by examining the functionality of regulatory measures in achieving policy objectives and managing the complex interplay between technological innovation and financial regulation.

As a result, this paper lays the theoretical foundation by outlining the current state of generative AI in finance. It argues that a sector-specific analysis is crucial to address the heightened and novel risks that generative AI introduces to the financial sector. The paper posits that the unique challenges posed by generative AI in the financial sector necessitate a nuanced debate that is sensitive to the context and industry. The application of AI, such as in financial advice, carries distinct societal implications compared to non-financial uses. By examining the impact of generative AI from a risk-based perspective, the paper scrutinizes how these AI systems handle financial data and decision-making automation, assessing the potential threats to financial stability, security, and ethics.

This work contributes to the risk mapping exercise of generative AI in finance, distinguishing it from the existing focus on predictive AI, and suggests strategies to address these emerging risks. By identifying and classifying the heterogeneity of these risks, which

may range from data privacy breaches to the amplification of systemic risks, the paper enriches the theoretical understanding of AI's potential negative externalities in finance. It presents a nuanced perspective that, after explaining how generative AI exacerbates old risks in the debate of AI governance and creates new ones, provides a framework for academics and practitioners to consider solutions that address the challenges of generative AI integration in financial ecosystems.

In practice, the paper's significance is anchored in its actionable proposals that aim to tackle the identified risks of employing generative AI in financial services. Hence, the proposals do not rely on regulatory action by regulators, but also provide possible solutions and risk-mitigation proposals for financial institutions to consider in their ethical and AI governance journey. These proposals likely span across technical, regulatory, and ethical domains, offering a roadmap for responsible experimentation with generative AI in finance. Hence, the paper does not only present a critical reflection on the potential pitfalls of generative AI but also paves the way for its secure and efficacious adoption.

2. The gaps in the regulatory debate on AI governance

2.1. Predictive AI versus generative AI

The policy and regulatory approaches to AI often do not address the ethical, legal and governance aspects of generative AI and predictive AI separately. Thus, not many discussions about the risks of generative AI and how different these models are from predictive AI - or other types of AI - have been guiding the regulatory design. Similarly, AI regulations applicable to financial services and products – even light touch approaches such as guidelines on the ethical principles of artificial intelligence in the financial sector – do not differentiate between generative AI and predictive AI.

The European Union is trying to be a first mover in this regard by proposing a regulatory approach in the AI Act that attempts to address generative AI models [18]. However, this approach falls short. The AI Act follows a principles-based and risk-based approach [19], which means that the riskier the AI model, the stricter the requirements for the company developing the model. After months of deliberations, last-minute adjustments were introduced pertaining to generative AI [19]. Under the AI Act, generative AI is considered a foundation model. Providers of foundation models would have to comply with additional requirements, like transparency obligations, requirements relating to registration, training, design, and dataset standards, among others [10]. The scope of this approach results in extensive documentation obligations. However, some of these obligations are vague in terms of their meaning and how they logically connect to the risks created by generative AI. In practice, the latest proposal of the EU AI Act treats all foundation models as high-risk applications, which imposes additional and onerous obligations on AI developers to mitigate hypothetical risks that may never materialize [20]. Moreover, questions about enforceability and the meaning of certain obligations are yet to be settled. These questions will be discussed in the dialogue between the Parliament, the European Commission and the Council. It is noteworthy that the General Data Protection Regulation (GDPR) underwent the same

process. As a result of such a dialogue, some safeguards that would have been applicable to automated decision-making, hitherto included in previous versions of the GDPR, were excluded [21]. Hence, the proposed approach by the European Union is not only limited in reach (since it only applies to providers regulated by the EU AI Act), but its obligations and requirements might also not make it into the final version of the Act.

Hence, the current approaches in AI governance and regulation do not sufficiently address the differences in ethical, legal, and governance challenges between predictive AI and generative AI [22]. Predictive AI and generative AI are two distinct branches of artificial intelligence that have different purposes and functionalities. As this paper proceeds to explain, this is especially relevant in finance and the ongoing AI governance discussion.

On the one hand, generative AI refers to the use of AI to create new data, such as audio, images, texts, and more [23]. The AI models that are part of this subfield of AI find and learn patterns in the given data set to produce new data that is similar to the given data set. The model is not designed with the specific purpose of following some instructions to make decisions if some conditions are met, or to predict an output based on a data set. Generative AI models require large training data sets to learn the patterns and generate outcomes. A simple example of this is a generative AI model that is trained on a given data set to create an image of a turtle. The given data set in this case would contain many images of turtles from which the model can learn the patterns of a turtle to generate images that would look like a turtle. The turtles generated by the model are not contained in the given data set. These generated images are often referred to as synthetic data [24].

Figure 1 explores three prominent methodologies for creating generative AI models.

Generative AI techniques can be applied to financial sector data to develop trading strategies and investment models. By generating synthetic market data and simulating trading scenarios, these models can assist in predicting market movements, identifying patterns, and generating insights that can be used for decision-making [25].

On the other hand, predictive AI refers to systems designed to predict outcomes based on data. It uses statistical techniques to interpret past data and make predictions about future events. It analyses historical data to identify patterns or trends and makes predictions about future events [26]. Thus, predictive AI analyses existing data to make predictions, while generative AI creates new content based on learned patterns. Predictive AI in finance is often used in cases such as algorithmic credit scoring. In these use cases, a lender uses AI to predict a borrower's likelihood of repaying a loan [8].

Due to the different capabilities and applications of these technologies, the nature of the ethical, legal and governance challenges can vary considerably [27]. For instance, predictive AI models such as algorithmic credit scoring have the potential to perpetuate biases and discrimination due to societal prejudices embedded in the data. If the data has historically excluded certain groups of individuals, the use of AI to predict the creditworthiness of such individuals has the potential to perpetuate biases and exacerbate the problem of financial exclusion [8]. The use of generative AI could potentially worsen this effect because of the way in which data is created in these types of systems. The content that the model creates is increasingly difficult to distinguish from real data. If that content finds its way back into

future models as training data, next generation AI models could become even more biased, creating a snowball effect of compounding bias with potentially wide implications for access to finance [27].




Figure 1. Reprinted with permission. Copyright (2023): SMU Centre for AI and Data Governance.

2.2. AI Governance policy debate: the lack of a sector-specific approach

In the AI governance policy debate, a significant issue is the lack of a sector-specific approach. This refers to the broad, generalized soft or hard regulatory approaches that are often applied to AI systems without considering the unique challenges and requirements of different industry sectors. For instance, AI applications in healthcare, finance, transportation, or education have distinct considerations concerning data privacy, ethical implications, and operational risks. A sector-specific approach would tailor governance frameworks to address

these nuances, ensuring that regulations are effective and appropriate for the context in which the AI is deployed. The current debate should centre on finding the balance between universal principles for AI and the need for regulations that are sensitive to sector-specific demands.

Only few regulators have taken some steps in this direction [8]. For example, the European Union has adopted a risk-based approach framework in the AI Act. Under this framework, certain use cases developed in some sectors are considered high-risk. Algorithmic credit scoring is considered a high-risk use case of AI systems under this new regulatory framework because it has significant implications for individuals' financial opportunities and privacy [15]. The risks include potential biases in decision-making, lack of transparency, and challenges in contesting or rectifying incorrect decisions. Therefore, these AI systems undergo stringent assessments and conform to higher regulatory standards to protect individuals' rights and ensure fairness in financial services [28].

Other jurisdictions have adopted a softer approach to AI governance in finance. Namely, the Monetary Authority of Singapore (MAS) follows such an approach. In 2018, MAS released a set of principles that covered Fairness, Ethics, Accountability and Transparency for responsible AI adoption within financial institutions [29]. However, those principles were framed in a very general way without translating them into actionable tools for financial institutions to assess the materiality of the risks created by AI use cases in finance. Therefore, a couple of years later, MAS launched the Veritas Initiative. Under this initiative, MAS suggested an assessment methodology to help financial institutions validate the fairness, ethicality, accountability and transparency of specific AI use cases such as credit scoring, insurance predictive underwriting, customer marketing, and fraud detection [14]. Veritas equips financial institutions with a range of suggested metrics to gauge AI's impact in various use cases, underscoring the need to weigh the materiality of each specific case. Although adherence to Veritas is not mandated for financial institutions in Singapore, there is an industry expectation, backed by regulators, that these institutions will closely align their risk assessment practices with the methodologies provided by Veritas.

3. Generative AI in finance – potential benefits and stage of development

By leveraging AI, financial institutions are better equipped to really transform the decision-making process to be more accurate, efficient, and successful [30]. Accordingly, predictive AI has already been used extensively in the industry, typically with structured data for prediction and segmentation. The rise of generative AI introduces a new dimension. This form of AI excels in analysing vast and varied datasets and translating them into accessible and user-friendly formats. This capability is pivotal in bolstering operational efficiency, refining customer interactions, enhancing risk management strategies, and fortifying compliance reporting [31]. This evolution of AI promises to drive one of the most significant shifts in the financial landscape.

Generative AI has the potential to automate document processing and analysis for due diligence, offer advanced fraud detection capabilities, and simulate financial scenarios to aid in strategic planning. Additionally, generative AI can optimize operational efficiencies by

automating routine tasks and processes, thereby reducing costs and human error. For instance, generative AI could supplement traditional fraud analytics with models that can identify abnormal patterns in large volumes of financial transactions so that financial institutions can halt suspicious transactions faster [32]. Financial institutions could also use generative AI to create synthetic data that simulates fraud so they can develop more robust fraud detection algorithms [32]. Moreover, generative AI models can determine credit risk more accurately by analysing vast amounts of data, including financial statements, credit scores, transaction histories, and other relevant data, including alternative data [33]. This can lead to better lending decisions that reduce credit risk [34]. This could also lead to increase financial inclusion and access to finance given generative AI's capability to discern unconventional variables and patterns to evaluate a borrower's creditworthiness. This would allow companies to make more informed lending decisions for borrowers with minimal credit history [31].

Additionally, Generative AI could enhance algorithmic trading by creating models that can simulate and predict market behaviours, generate synthetic financial data for stress testing, and heighten trading strategies by identifying complex patterns and insights from large datasets that are not readily apparent through traditional analysis [35]. This helps in decision-making for investment strategies by considering potential market movements and sentiment analysis derived from diverse sources, such as news articles and social media [36]. By doing so, generative AI can provide traders with a competitive edge through more informed and timely decisions.

Therefore, aware of these potential benefits, most financial institutions are currently exploring use cases and researching which models and practices would benefit them the most. Banks like JPMorgan Chase, Wells Fargo, Goldman Sachs, and Singapore's DBS have been experimenting with AI for some time and are now ideating generative AI-powered systems [7]. One example of generative AI in finance is the use of Generative Adversarial Networks (GANs) for fraud detection. GANs can be trained to generate synthetic financial transaction data that mimics normal consumer behaviour. By doing so, financial institutions can then train their fraud detection models to identify anomalies or patterns that deviate from these synthetic 'normal' transactions, which may indicate fraudulent activity. For example, JPMorgan Chase is using large language models [37] to detect fraud by examining patterns in emails for signs of compromise [38]. Likewise, Goldman Sachs is using generative AI to assist software engineers in code development [39]. SouthState Bank trained an enterprise-version of ChatGPT that allows employees to use it to query about bank policies, draft emails and summarize meetings [40]. In Singapore, DBS is leveraging AI to enhance its operation too. According to publicly available data, there are currently more than 300 AI use cases across financial institutions. This resulted in a revenue uplift of \$150 million in 2022 — double that in 2021 — and an additional \$30 million from cost avoidance due to enhanced credit monitoring, fraud prevention and productivity gains [41]. Like most financial institutions leveraging AI, DBS is experimenting with generative AI and running small-scale pilots to see if it can help improve revenue and employee productivity as well as other internal use cases such as helping employees grow their careers [41].

As generative AI's applications in the financial sector continue to evolve, it is crucial to address the risks created when implementing generative AI systems in finance to ensure responsible and secure use. Those risks should be independently evaluated from those created by other artificial intelligence models currently used in finance.

4. The new and old risks of generative AI in finance

Generative AI brings renewed attention to many of the risks created by AI, such as the potential to perpetuate bias hidden in training data, while presenting new ones that require attention from policymakers.

4.1. Financial hallucinations

The problem of "hallucination" in generative AI refers to instances where the AI model generates information that is not based or is factually incorrect. This can happen when the AI produces outputs that seem plausible but are either partially or false or nonsensical [42]. Hallucinations in AI are particularly problematic when the AI is used in contexts where accuracy and truth are critical. For instance, inaccuracy in risk assessment processes has negative implications for risk-taking and management. Financial services offered to customers through generative AI-supported chatbots could give inappropriate advice or offer unsuitable products to financial consumers [43]. These outcomes not only erode public trust in AI systems, but also in financial institutions using these technologies.

Mitigating hallucinations in AI require careful model design, high-quality data for training, frequent updates with accurate information, and the implementation of safeguards like fact-checking and human oversight.

4.2. A greater explainability problem

Explainability is an AI governance principle that is relevant in any type of AI use case [44]. However, this principle becomes more complex in certain AI use cases. Generative AI models can be complex, and interpretation may become difficult. Indeed, the vast amount of varied data is essential for the effectiveness of generative AI, but it also makes it complicated to trace the origins of a particular piece of information the AI produces [45]. This is exacerbated by the risk of hallucinations, as explained above. In addition, generative AI models use multiple layers within neural networks, which work together in intricate ways [37].

These networks utilize a staggering number of parameters to weigh the likelihood of one element of its output over another. All these parameters act like a multitude of dials and switches, fine-tuning the AI's performance. As in deep-learning systems, this creates the problem often referred to as the black box [46]. As a result, it is challenging to understand how each internal decision contributes to the final output. This method of operation also makes it challenging for financial institutions to iterate and explain automated outputs to customers or regulators.

4.3. Data protection, privacy in foundation models

Foundation models are typically trained on vast amounts of data, which may include or inadvertently collect sensitive financial data, thus raising privacy issues in a highly regulated sector from a data protection perspective [47]. Moreover, the integration of publicly accessible generative AI systems, such as GPT, into the workflow of financial institutions amplifies this problem because of the risk of inadvertent disclosure of confidential financial details and personal data entered by employees of the financial institutions while interacting with the generative AI. Additionally, the terms and conditions of generative AI platforms state that the technology provider doesn't guarantee the protection or privacy of user-provided data, which further emphasizes the potential security vulnerabilities associated with their use [43].

Generative AI systems developed internally by financial institutions may mitigate this problem but could also make it difficult for small financial institutions to take advantage of the potential benefits of cost reductions due to the impossibility of deploying their own internal model.

4.4. Fairness

Generative AI models are only as good as the data they are trained on. If the data contains biases or inaccuracies, these biases could be amplified by the generative AI model and, thus, result in a wrong analysis. This could lead to discriminatory decisions or other adverse outcomes. Furthermore, the responses and content generated by the generative AI model are influenced by prompts directed to the AI system. The prompts could reflect human biases [27]. This is a significant problem because it can lead to discriminatory outcomes or reinforce societal inequalities in highly sensitive contexts such as access to finance. Some regulators have approached the risk of discrimination derived from the use of AI in finance differently, particularly in lending markets. As mentioned above, the AI Act considers generative AI models as high-risk, but fails to address the nuances of generative AI in lending. Other regulators, such as the Monetary Authority of Singapore, encourage financial institutions to assess the impact of the use of algorithmic credit scoring [14]. However, the metrics suggested by the regulators are not mandatory, and do not address the particularities of generative AI.

4.5. Systemic risk and financial stability

The widespread use of AI in finance impacts systemic risk in the financial sector in a unique way due to the inherent procyclicality of AI, the speed of AI outcomes, the new levels of interconnectedness derived from AI networks and new third-party dependencies, and the establishment of new 'too big to fail' entities. These factors, coupled with the lack of a consistent and enforceable approach to AI governance and concrete models to assess the impact of AI on financial stability, make it challenging for supervisors and financial institutions to prevent systemic risks in the era of AI [8].

This problem is also aggravated using generative AI. The potential excessive reliance on generative AI due to its convenience may increase financial stability concerns. Decisions made by employees of financial institutions based on generated economic, market, or risk reports could be susceptible to herd behavioural bias and mispricing if these reports reflect public sentiments captured by data sets and, consequently, used by AI systems [43]. Likewise, financial hallucination is a concern that could be systemically important if the misleading outputs are widely adopted in the financial system [43]. Furthermore, rising interconnectedness due to the saturated market for providers of foundations models makes the traditional financial stability issue of too big to fail more complicated. The technology markets are highly concentrated and the providers are often beyond the reach of financial supervisors and regulators [8].

4.6. Authorized push payment frauds in the era of generative AI

Generative AI is also revolutionising financial crime. AI that is both increasingly sophisticated and available is making it easier for fraud and scams to deceive financial consumers into authorizing payments. For example, fraudsters may use generative AI to impersonate customer service representatives from legitimate organizations in a more credible way due to the sophistication of the technology [48]. Generative AI systems can create such data in seconds and there are no barriers of access to the technology by fraudsters. The use of deepfakes is certainly something that should concern financial regulators when it comes to dealing with authorized push payment fraud. Generative AI puts the financial consumer at risk, and even highly educated and conservative consumers may fall prey to scammers using generative AI systems capable of creating credible content. According to authorities in some jurisdictions, AI deepfakes had been used to trick facial recognition programs by imitating the people pictured on the identity cards, according to police [49].

Even financial institutions' employees are vulnerable to these enhance scam techniques. In Hong Kong, an employee of a financial firm was tricked into paying out \$25 million to hoaxers using deepfake technology to pose as the company's chief financial officer in a video conference call [49]. The employee first received a message that was purportedly from the company's UK-based chief financial officer. Initially, he suspected it was a phishing email, as it talked of the need for a secret transaction to be carried out. However, the employee put aside his early doubts after the video call because other people in attendance had looked and sounded just like colleagues he recognized [49].

4.7. Copyright infringement

There is growing unease about the incorporation of copyrighted content in training datasets and whether permission has been obtained from the copyright holders. In the financial sector, the application of generative AI raises copyright issues when these systems are trained using proprietary market analyses, commentary, and financial reports that are copyrighted. Stakeholders, including financial institutions and market analysts, are concerned about the unauthorized use of their intellectual property. AI models might even replicate or replace

traditional financial analysis services. Currently, legal frameworks have not yet provided legal certainty on this issue. Some copyright owners, outside of the financial sector, have sued generative AI companies in the United States and the United Kingdom, but courts have yet to decide on these matters. Copyright holders have sought compensation for the use of their creative works in AI training, fearing competition and economic impact. They propose licensing arrangements for text and data mining to support machine learning and suggest opt-out mechanisms to protect their financial interests. Some contend that machine learning, which often involves text and data mining, doesn't equate to copyright violation, as it doesn't replicate the original artistic elements. Nonetheless, the practicality of securing permissions from all copyright owners and the balance between legal compliance and AI efficacy remains a contentious topic. Some jurisdictions, such as Japan, South Korea, and the United Kingdom, are also exploring self-regulatory solutions such as copyright guidelines and codes of practice for developers and end-users. However, none of these efforts has targeted the financial sector specifically.

4.8. Value alignment

The problem of value alignment in generative AI pertains to the challenge of ensuring that the AI's actions and decisions align with human values and ethics [50]. This is particularly crucial in generative AI, which has the capacity to create new, original content. If the AI's value system is misaligned with societal norms, it could potentially generate harmful, biased, or unethical content, leading to negative consequences. In financial contexts, this misalignment might manifest in generating investment strategies that are overly risky or even deceptive, which could undermine trust in financial institutions and financial stability.

5. Addressing the nuances of generative AI in finance

5.1. The importance of a sector and sub-sector specific approach

Different industries have unique characteristics, risk profiles, and operational practices. AI applications in healthcare, finance, education, and law enforcement, for instance, each have distinct implications. Tailored regulation can address the specific risks and requirements of each sector, ensuring that AI is used safely and effectively without stifling innovation. Even within the financial sector, each AI use case has different implications and impact financial regulation goals differently. For instance, AI credit scoring may have a greater impact on financial inclusion, financial consumer protection and fairness due to the possibility of the models perpetuating discrimination in access to finance [51]. In contrast, the use of algorithmic trading may impact systemic risk and the way market misconduct operates in capital markets [52]. This analysis is applicable to generative AI in finance too. The use of generative AI to provide financial advice to consumers have different implications from use cases to enhance fraud monitoring.

Hence, a one-size-fits-all approach does not address the nuanced challenges present in different industries, much less in different financial sector sub-industries. Some regulatory

approaches, such as the AI Act, adopt a risk-based approach in response to this concern. Nevertheless, sectors evolve at different paces and directions. A sector-specific approach would allow governance frameworks to adapt dynamically to the pace of AI development and deployment in each sector, facilitating timely updates to regulations.

5.2. Ex-post supervisory monitoring of AI systems

Ex-post supervisory monitoring involves the oversight and analysis of AI systems after deployment [51]. Such monitoring can reduce the risks linked to generative AI in finance by evaluating its impact and consequences within the sector. Financial regulators may mandate that institutions report both experimental and operational uses of generative AI. Continuous performance assessment of these systems enables supervisors to detect any deviations from expected outcomes, which may indicate issues like bias, errors, or other unforeseen problems, including threats to financial stability. This approach could be more effective than relying on rigid legislation for managing generative AI in finance, given the technology's current stage of development. Also, only an ex-post approach would address risks such as those related to value alignment. Ensuring value alignment not only involves complex programming but also continuous oversight to align AI outputs with human ethical standards and the risk management strategy of a financial institution.

5.3. Generative AI audits in finance

Algorithmic audits may become a common tool of AI regulation, and is set to be adopted by the EU AI Act for high-risk and foundation models. Even outside of the implementation of such regulatory proposals, AI models like ChatGPT go through various forms of auditing and evaluation, although the specifics of these processes are not fully disclosed to the public for proprietary reasons. It is not clear how AI audits should be conducted, and the private sector has yet to develop a standardised approach. This is still a nascent field in which regulators must think about the basic components before audits can reliably address algorithmic harms [53].

Moreover, generative AI audits in finance is a policy that needs to be complemented with a requirement of human accountability. Generative AI output does not consist of decisions but content. Accordingly, generative AI is mostly used to provide recommendations, advice, or analysis, with human actors making the decisions. Hence, financial regulators could clarify that financial institutions need to understand the reasons behind their decisions, and whether these actions are based on outputs generated by generative AI systems. This requires financial institutions to understand the limitations of generative AI and that it is just one component in the decision-making process.

5.4. Enhanced financial consumer protection

In an era of fraud sophistication, undertaking a cost-benefit analysis to determine if changes can be made to the rules applicable to reimbursements in online scam cases, especially in authorized push payment fraud, could be fair to financial consumers. The United Kingdom

is currently discussing a change in this policy to incentivise payment services providers to improve their fraud prevention capabilities, which would result in customers enjoying greater protections, and fraud victims having their money reimbursed with very limited exceptions [54].

5.5. Open dialogue with the industry and flexible regulatory approaches

As generative AI becomes more advanced in the financial sector, ongoing dialogue between regulators and the industry is essential. This is because generative AI's capabilities and potential applications are expanding rapidly, often outpacing current regulatory frameworks. Continuous engagement ensures that regulations remain relevant and effective, addressing emerging risks without stifling innovation. This collaboration can facilitate the development of standards and guidelines that uphold ethical practices, consumer protection, and market integrity while allowing the financial industry to harness the benefits of generative AI technology [55].

6. Conclusion

The integration of generative AI into the financial sector represents a pivotal shift with far-reaching implications. This paper has argued emphatically for a nuanced, context and sector-specific debate to address the multifaceted risks and challenges posed by generative AI in finance. The argument presented emerges from the urgent need to align AI's rapid technological evolution with the equally dynamic financial landscape. The financial sector's adoption of generative AI comes with a promise of transformational change. From the nuanced personalization of customer interactions to the development of sophisticated fraud detection systems and the automated generation of financial reports, the potential applications are vast. The benefits of the technology are counterbalanced by an array of ethical and legal risks that demand scrutiny. This paper has endeavoured to delineate some potential benefits in the use of generative AI in finance, but also the risks, underscoring the importance of a contextual understanding of AI's role in the financial sector. At the heart of the discussion is the recognition that the implications of generative AI in finance are substantially different from those in less sensitive applications. For instance, the generation of creative content for marketing purposes has different implications from the use of AI to provide financial advice. The latter involves significant fiduciary responsibilities and has the potential for far-reaching socioeconomic consequences, thus necessitating a more stringent governance framework.

This paper has highlighted the regulatory lacuna that currently exists, wherein most AI governance frameworks are heavily weighted towards predictive AI, often overlooking the distinct challenges of generative AI. The existing policy discussions and regulatory proposals, while foundational, are insufficient for the nuanced demands of the financial sector. They often lack the sector-specific granularity required to navigate the complex interplay between AI-generated outputs and the financial market's stability and integrity. The paper also analyses the heterogeneity of potential risks that generative AI creates in financial services. Building on this analysis, the paper has charted a set of proposals to confront the challenges of experimenting with AI in finance. While the paper has proposed initial

solutions to these challenges, future research should aim to refine these proposals, ensuring they are pragmatic and actionable. Further studies could develop metrics to evaluate the extent of the risks highlighted in this study. Future research could also expand on creating detailed, use-case specific guidelines and perform empirical evaluations of the proposed frameworks in real-world financial settings.

Acknowledgments

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Emerging Area Research Project Funding Initiative. For valuable comments and feedback, I would like to thank the anonymous referees. For excellent research assistance, I am grateful to Fabien Tan.

Conflicts of interests

The author declares no conflict of interest.

References

- [1] Introducing ChatGPT. Available: <https://openai.com/blog/chatgpt> (accessed on 12 June 2023).
- [2] Chat GPT: Achieving 100 Million Users in Just 2 Month — A Deep Analysis. Available: <https://ai.plainenglish.io/chat-gpt-achieving-100-million-users-in-just-2-month-a-deep-analysis-a453e6f85acf> (accessed on 13 June 2023).
- [3] 3 trends will shape the future of work, according to Microsoft's CEO. Available: <https://www.weforum.org/agenda/2023/02/future-jobs-workplace-ai-microsoft/> (accessed on 13 June 2023).
- [4] Microsoft Confirms Its \$10 Billion Investment Into ChatGPT, Changing How Microsoft Competes With Google, Apple And Other Tech Giants. Available: <https://www.forbes.com/sites/qai/2023/01/27/microsoft-confirms-its-10-billion-investment-into-chatgpt-changing-how-microsoft-competes-with-google-apple-and-other-tech-giants/?sh=a9a12fa36246> (accessed on 13 June 2023).
- [5] Discover: Generative AI. Available: <https://ai.google/discover/generativeai/> (accessed on 13 June 2023).
- [6] Reimagining Our Infrastructure for the AI Age. Available: <https://about.fb.com/news/2023/05/metasp-infrastructure-for-ai/> (accessed on 13 June 2023).
- [7] How banks can adopt generative AI. Available: <https://www.americanbanker.com/list/how-banks-can-adopt-generative-ai> (accessed 2 February 2024).
- [8] The Role of Financial Regulators in the Governance of Algorithmic Credit Scoring. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4057986 (accessed on 12 February 2024).
- [9] Generative AI at the Crossroads. Available: <https://blogs.law.ox.ac.uk/oblb/blog-post/2023/06/generative-ai-crossroads> (accessed on 13 June 2023).
- [10] Proposal for a regulation of the European Parliament and of the Council on harmonized rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts. Available: https://www.europarl.europa.eu/meetdocs/2014_2019/plmrep/COMMITTEES/CJ40/DV/2023/05-11/ConsolidatedCA_IMCOLIBE_AI_ACT_EN.pdf (accessed on 13 June 2023).

- [11] Infocomm Media Development Authority, proposed Model AI Framework for Generative AI – Fostering a Trusted Ecosystem. Available: https://aiverifyfoundation.sg/downloads/Proposed_MGF_Gen_AI_2024.pdf (accessed on 17 January 2024).
- [12] Oversight of A.I.: Rules for Artificial Intelligence. Available: <https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-rules-for-artificial-intelligence> (accessed on 13 June 2023).
- [13] Regulatory framework proposal on artificial intelligence. Available: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai> (accessed on 13 June 2023).
- [14] MAS-led Industry Consortium Publishes Assessment Methodologies for Responsible Use of AI by Financial Institutions. Available: <https://www.mas.gov.sg/news/media-releases/2022/mas-led-industry-consortium-publishes-assessment-methodologies-for-responsible-use-of-ai-by-financial-institutions> (accessed on 13 June 2023).
- [15] Kitchenham B. *Procedures for performing systematic reviews*. Keele: Keele University, 2004, 33, pp. 1-26.
- [16] Hutchinson T, Duncan N. Defining and describing what we do: doctrinal legal research. *Deakin Law Rev.* 2012, 17(1):83-119.
- [17] David Clark, *Comparative Law and Society* (Edward Elgar Publishing, 2012) 35; Horatia Muir Watt, *Globalization and Comparative Law* (Oxford University Press, 2006) 13
- [18] AI Act: a step closer to the first rules on Artificial Intelligence. Available: <https://www.europarl.europa.eu/news/en/press-room/20230505IPR84904/ai-act-a-step-closer-to-the-first-rules-on-artificial-intelligence> (accessed on 13 June 2023).
- [19] European AI Alliance - The EU AI Act's Risk-Based Approach: High-Risk Systems and What They Mean for Users. Available: <https://futurium.ec.europa.eu/en/european-ai-alliance/document/eu-ai-acts-risk-based-approach-high-risk-systems-and-what-they-mean-users> (accessed on 13 June 2023).
- [20] Regulating ChatGPT and other Large Generative AI Models. Available: <https://dl.acm.org/doi/10.1145/3593013.3594067> (accessed on 13 June 2023).
- [21] Wachter S, Mittelstadt B, Floridi L. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *Int. Data Priv. Law* 2017, 7(2):76-99.
- [22] Bale A S, Dhumale R B, Beri N, Lourens M, Varma R A, *et al.* The Impact of Generative Content on Individuals Privacy and Ethical Concerns. *Int. J. Intell. Syst. Appl. Eng.* 2024, 12:697-703.
- [23] Generative Deep Learning. Available: <https://www.oreilly.com/library/view/generative-deep-learning/9781098134174/> (accessed on 16 February 2024).
- [24] Synthetic data for deep learning. Available: <https://link.springer.com/book/10.1007/978-3-030-75178-4> (accessed on 12 February 2024).
- [25] Coletta A, Prata M, Conti M, Mercanti E, Bartolini N, *et al.* Towards realistic market simulations: a generative adversarial networks approach. In *Proceedings of the Second ACM International Conference on AI in Finance*, November 3-5, 2021, pp. 1-9.
- [26] Predictive AI vs. Generative AI: Unveiling the Distinctions. Available: <https://www.linkedin.com/pulse/predictive-ai-vs-generative-unveiling-distinctions-ishan-kapoor/> (accessed on 16 February 2024).
- [27] Humans Are Biased: Generative AI Is Even Worse. Available: <https://www.bloomberg.com/graphics/2023-generative-ai-bias/> (accessed on June 23 2023).
- [28] Algorithms, credit scoring, and the new proposals of the EU for an AI Act and on a Consumer Credit Directive. Available: <https://www.tandfonline.com/doi/abs/10.1080/17521440.2023.2168940> (accessed on 16 February 2024).

- [29] MAS introduces new FEAT Principles to promote responsible use of AI and data analytics. Available: <https://www.mas.gov.sg/news/media-releases/2018/mas-introduces-new-feat-principles-to-promote-responsible-use-of-ai-and-data-analytics> (accessed on 2 February 2024).
- [30] Shrestha YR, Vaibhav K, and Georg von K. Augmenting organizational decision-making with deep learning algorithms: Principles, promises, and challenges. *J. Bus. Res.* 2021, 123:588-603.
- [31] Generative AI in Financial Services: Revolutionizing the Finance Sector. Available: <https://www.batonglobal.com/post/generative-ai-in-financial-services-revolutionizing-the-finance-sector> (accessed on 15 February 2024).
- [32] The Benefits of Generative AI for Banking & Financial Leaders. Available: <https://www.actian.com/blog/data-analytics/benefits-generative-ai-banking-financial/> (accessed on 10 February 2024).
- [33] Yu X, Yang Q, Wang R, Fang R, Deng M. Data cleaning for personal credit scoring by utilizing social media data: An empirical study. *IEEE Intell. Syst.* 2020, 35(2):7-15.
- [34] Sargeant H. Algorithmic decision-making in financial services: economic and normative outcomes in consumer credit. *AI Ethics* 2023, 3(4):1295-1311.
- [35] Carvajal-Patiño D, Ramos-Pollán R. Synthetic data generation with deep generative models to enhance predictive tasks in trading strategies. *Res. Int. Bus. Finance* 2022, 62: 101747.
- [36] Dahal S B. Utilizing Generative AI for Real-Time Financial Market Analysis Opportunities and Challenges. *Adv. Intell. Inf. Syst.* 2023, 8(4):1-11.
- [37] A language model is a machine learning model that aims to predict and generate plausible language. A large language model has been trained on a vast corpus of text data and is a specific application of generative AI. Jakub M. Tomczak, *Deep Generative Modeling* (Springer 2022).
- [38] JPMorgan Chase using advanced AI to detect fraud. Available: <https://www.americanbanker.com/news/jpmorgan-chase-using-chatgpt-like-large-language-models-to-detect-fraud> (accessed on 15 February 2024).
- [39] AI will make workers 'superhuman': Goldman Sachs CIO Marco Argenti. Available: <https://www.americanbanker.com/news/ai-will-make-workers-become-superhuman-goldman-sachs-cio-marco-argenti>. (accessed on 15 February 2024).
- [40] 'It's worth all the hype': SouthState Bank deploys enterprise ChatGPT. Available: <https://www.americanbanker.com/news/its-worth-all-the-hype-southstate-bank-deploys-chatgpt-like-tech> (accessed on 15 February 2024).
- [41] DBS shares its secret to becoming AI-fuelled. Available: <https://www.theedgesingapore.com/digitaledge/digital-economy/dbs-shares-its-secret-becoming-ai-fuelled> (accessed on 15 February 2024).
- [42] del Campo M, Leach N, Eds. *Machine Hallucinations: Architecture and Artificial Intelligence*. Hoboken: John Wiley & Sons, 2022, pp.1-144.
- [43] Generative Artificial Intelligence in Finance. Available: <https://doi.org/10.5089/9798400251092.063> (accessed on 15 February 2024).
- [44] Taeihagh A. Governance of artificial intelligence. *Policy Soc.* 2021, 40(2):137-157.
- [45] Atkinson K, Bench-Capon T, Bollegala D. Explanation in AI and law: Past, present and future. *Artif Intell.* 289, 103387.
- [46] Rai A. Explainable AI: From black box to glass box. *J. Acad. Mark. Sci.* , 48, 137-141.
- [47] Gupta M, Akiri C, Aryal K, Parker E, Praharaj L. From chatgpt to threatgpt: Impact of generative ai in cybersecurity and privacy. *IEEE Access* 2023, 11:80218-80245.
- [48] FraudGPT: How AI is—and isn't—revolutionizing financial crime. Available: <https://resistant.ai/blog/fraudgpt-how-ai-is-revolutionizing-financial-crime> (accessed on 15 February 2024).

- [49] Finance worker pays out \$25 million after video call with deepfake ‘chief financial officer’. Available: <https://edition.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk/index.html> (accessed on 10 February 2024).
- [50] Liu R, Zhang G, Feng X, Vosoughi S. Aligning generative language models with human values. In *Findings of the Association for Computational Linguistics: NAACL 2022*, Seattle: Association for Computational Linguistics, 2022, pp. 241-252.
- [51] Gillis T. The Input Fallacy. *Minn. Law Rev.* 2022, 106: 1175
- [52] Azzutti A, Ringe WG, Stiehl HS. Machine learning, market manipulation, and collusion on capital markets: Why the "black box" matters. *U. Pa. J. Int'l L.* 2021, 43:79.
- [53] Goodman EP, Trehu J. Algorithmic Auditing: Chasing AI Accountability. *Santa Clara High Tech. L.J.* 2022, 39:289.
- [54] Fighting authorised push payment fraud: a new reimbursement requirement. Available: <https://www.psr.org.uk/publications/policy-statements/ps23-3-fighting-authorised-push-payment-fraud-a-new-reimbursement-requirement/>. (accessed on 15 February 2024).
- [55] de Marcellis-Warin N, Marty F, Thelisson E, Warin T. Artificial intelligence and consumer manipulations: from consumer's counter algorithms to firm's self-regulation tools. *AI Ethics* 2022, 2(2):259-268.