Generative Artificial Intelligence in FinTech Services: Use Cases, Value Creation and Emerging Regulatory Risks

Jialin Guo

School of Slavonic and East European Studies, University College London, London, The United Kingdom

zczqjg2@ucl.ac.uk

Abstract. The convergence of Artificial Intelligence (AI) and Financial Technology (FinTech) is reshaping the architecture of modern financial services, introducing new capabilities in automation, decision-making, and user interaction. In particular, recent advances in generative AI—most notably Large Language Models (LLMs)—are driving a paradigm shift in how information is produced, communicated, and acted upon within financial institutions. This paper investigates the integration of generative AI into FinTech, focusing on four principal areas of application: automated financial content generation, customer service systems, investment advisory tools, and regulatory compliance reporting. Drawing on recent empirical studies and industry practices, the analysis demonstrates that LLMs are increasingly embedded in financial workflows to enhance speed, scalability, and user interaction. While these systems offer operational advantages, including real-time responsiveness and reduced labour costs, they also introduce substantive challenges related to accuracy, interpretability, and regulatory alignment. The paper evaluates both the functional potential and the structural limitations of generative AI in finance, highlighting the necessity of human oversight, domain-specific fine-tuning, and governance mechanisms. Overall, the findings suggest that generative AI will continue to expand its role in FinTech, particularly through hybrid frameworks that combine automation with expert control.

Keywords: Generative AI, fintech, large language models.

1. Introduction

FinTech is changing how financial services are delivered. It introduces digital tools into payments, lending, investment and regulatory tasks. This shift has made services faster, more accessible, and increasingly personalised. In recent years, AI has added a new layer of innovation. These models, especially LLMs, can produce financial content that resembles human writing. As a result, they are now being used to automate financial reports, generate regulatory documents, and support client communication.

Generative AI refers to models that create new content based on patterns in large datasets. These models can write coherent text, generate code, or respond to user prompts in real time. Because they adjust to input, they are useful for automation, summarisation, and decision support. FinTech includes a wide range of technology-based services in payments, digital banking, investment, insurance, and regulation. Combining generative AI with FinTech creates new opportunities. These include cost savings, faster operations, and more personalised services. For instance, generative AI can help write investment summaries, guide users through onboarding, or answer client queries with natural language.

Generative AI systems are more flexible than traditional rule-based tools. They can respond based on context and adjust their output to match tone or purpose. Lee et al. show that firms already use generative AI for compliance support, internal knowledge sharing, and customer service [1]. Barde and Kulkarni discuss similar applications in robo-advisory, credit assessments, and generating financial narratives [2]. As use of these systems grows, it is important to understand both the value they bring and the risks they pose.

So far, most academic research on AI in finance has focused on predictive models. These include machine learning systems used for fraud detection or credit scoring. While generative AI is gaining attention in industry, several studies have begun to raise concerns about its practical deployment.

Researchers point to risks such as low transparency, unreliable outputs, and limited regulatory clarity. By contrast, real-world applications of generative AI in finance are still less studied. Andronie et al. argue that many current uses lack transparency and are difficult to evaluate [3]. Jain et al. highlight challenges with using AI in high-stakes areas like finance, where mistakes can lead to serious consequences [4]. Generative models may produce inaccurate or biased outputs. They also raise concerns about data privacy and unclear regulations. These risks make it necessary to assess both their benefits and their limitations.

There is a clear need to understand how generative AI is used in financial services, and what impact it has. It is important to ask whether these models actually improve performance and whether they are reliable in practice. This paper contributes to that discussion by studying how generative AI is applied in FinTech. It explores the value these models create and the risks they bring. The goal is to provide a balanced, evidence-based view that supports both innovation and responsible oversight. As highlighted by Gowda [5], achieving this balance requires understanding not only the operational benefits of generative AI but also its epistemic and systemic risks, particularly in fast-moving financial contexts.

2. Generative AI Integration in FinTech: Key Use Cases

2.1. Automated Financial Content Generation

Generative AI is increasingly applied to automate the production of financial texts, including earnings summaries, investment briefs, ESG reports, and customer-facing documents. Compared with manual drafting, AI-generated content enables faster output, lower operational costs, and greater scalability. Recent industry examples show that LLMs are now generating draft earnings summaries, ESG disclosures, and investment memos with minimal human editing (Lee et al., [1]).

Zhang, Feng, and Dong demonstrate how generative components can support real-time decision-making through the automatic generation of structured financial updates [6]. Their LAMDA architecture, originally designed for anomaly detection, has been applied to generate cross-market signals and summary explanations at low latency, enabling users to receive synthesised financial insights without manual interpretation. While not focused on narrative generation alone, the framework incorporates automated textual outputs tailored to live financial data streams, suggesting how generative AI can be embedded within decision-support interfaces.

These examples reflect how generative AI is being integrated into content pipelines for both internal analysis and external communication. Automated generation of market commentary, performance updates, and briefing materials allows financial institutions to respond more quickly to changing market conditions, while also standardising communication quality across teams. Beyond speed and consistency, the readability and interpretability of AI-generated financial reports can directly influence how such content is scored by ESG agencies. Shimamura, Tanaka, and Managi demonstrate that report clarity, often influenced by generative output, significantly affects ESG scoring outcomes [7].

2.2. Customer Service and Conversational Agents

Generative AI is transforming customer service in FinTech by enabling more adaptive and interactive communication tools. One of the most common applications is in AI-powered chatbots or virtual assistants. These tools allow financial institutions to respond to client queries in real time, using natural language that mimics human interaction. Compared to traditional scripted systems, generative models can handle open-ended requests and follow complex dialogue flows.

Desiraju and Khan highlight how generative AI has been adopted in banking environments to improve customer support [8]. Their study notes that such tools reduce waiting times, provide 24/7 service, and help financial firms handle high query volumes without increasing staff costs. In particular, they show that AI agents are effective in routine tasks such as balance checks, transaction history explanations, and guiding users through account setup processes. These applications free up human agents to focus on more complex cases.

Saha, Shukla, and Kumar provide broader survey evidence from global financial institutions, confirming a growing trend in deploying generative AI for front-end services [9]. Their findings suggest that over 60% of surveyed firms had implemented or piloted generative AI tools in customer-facing roles.

In practical terms, some banks have integrated generative chat agents directly into mobile applications to offer instant support on payments, card blocking, and loan eligibility assessments. Beyond basic query handling, these systems are also being used to pre-fill forms, automate responses to common regulatory queries, and provide multilingual support, which is especially valuable in global markets. The ability to personalise interactions at scale is becoming a key differentiator for firms aiming to improve digital customer experience.

2.3. Investment and Advisory Applications

Generative AI is being increasingly explored in investment and wealth advisory services. Its ability to process large volumes of financial data and generate user-friendly summaries has made it an attractive tool for both institutional and retail investors. These systems can assist with portfolio explanations, simulate market scenarios, and deliver personalised, real-time responses to client portfolio questions.

Lakkaraju et al. examined how LLMs perform when used as financial advisors in decision-making contexts [10]. Their findings suggest that while LLMs can generate coherent investment advice, their recommendations are not always optimal unless paired with financial reasoning modules or expert oversight. This highlights both the promise and limitations of using generative AI for personal finance.

Han et al. explored multi-agent generative AI systems designed to support investment analysis [11]. Their approach uses different AI agents for specific tasks—such as analyzing fundamentals, market sentiment, and portfolio risk. The study found that combining outputs from multiple specialized agents leads to better performance than using a single general-purpose model, especially during periods of market volatility, where quick interpretation of diverse data is critical.

In practice, financial platforms are deploying generative AI across a range of advisory applications. Robo-advisors use LLMs to generate customised investment recommendations. These are often accompanied by narrative rationales for portfolio changes, tailored to user profiles and risk parameters. These systems improve client engagement while automating explanation generation at scale.

Generative AI also enables users to interact with AI agents that simulate macroeconomic events—such as inflation shocks or interest rate hikes—and receive plain-language summaries of potential portfolio impacts. Such scenario modelling supports more informed and proactive decision-making.

At the institutional level, asset managers leverage generative AI to streamline the creation of analyst notes and thematic reports, drawing from market data and earnings calls to produce timely insights. Additionally, some firms integrate AI into advisor workstations, where it assists in drafting reports, client communication, and performance summaries, thereby enhancing productivity and advisory bandwidth without increasing operational overhead.

2.4. Regulatory Compliance and Risk Reporting

Generative AI is being increasingly explored in regulatory compliance and risk reporting within financial institutions. Tasks such as Know-Your-Customer (KYC) documentation, Anti-Money Laundering (AML) narratives, audit reporting, and transaction monitoring have traditionally relied on structured rule-based systems or human analysts. The emergence of Large Language Models (LLMs) has introduced a new layer of flexibility and automation to these processes.

Khanvilkar and Kommuru propose a novel system that integrates regulatory graph structures with generative AI to create real-time explanations of suspicious transactions [12]. Their model uses a combination of graph-based financial representations and language generation modules to produce audit-ready summaries that align with compliance requirements. Importantly, their framework demonstrates the potential of LLMs to support explainable compliance reporting with high precision and interpretability.

Beyond transaction-level analysis, financial firms are also beginning to deploy generative models for streamlining internal documentation workflows related to regulatory submissions. For instance, AI-generated narratives are being used to automate periodic compliance reports required by supervisory authorities, helping to reduce manual workload and standardise reporting language across departments. In KYC and onboarding procedures, LLMs assist in drafting due diligence summaries by synthesising structured client data into textual formats suitable for compliance review.

The International Monetary Fund notes that generative AI systems are being piloted in areas such as automated suspicious activity reports (SARs), model risk documentation, and regulatory change monitoring. One practical application includes summarising lengthy regulatory updates and identifying their relevance for specific business units, enabling compliance teams to respond more quickly to evolving requirements.

Another emerging use case lies in stress testing and risk scenario documentation, where generative AI helps convert quantitative model outputs into regulator-facing narratives that explain assumptions, outcomes, and mitigation strategies. This supports transparency and speeds up the process of preparing materials for internal audit committees or external regulators.

Overall, generative AI is expanding its footprint in compliance-related tasks, especially in contexts that benefit from the model's natural language capabilities. While oversight remains essential, these tools are increasingly embedded within operational workflows to augment speed, consistency, and adaptability in regulatory reporting. As these applications mature, the role of generative AI in compliance functions is shifting from experimental tools to embedded infrastructure.

3. Challenges and prospects

3.1. Automated Financial Content Generation

The automated production of financial texts—such as earnings summaries, ESG disclosures, and market briefings—faces multiple operational barriers. The primary concern lies in the factual accuracy of generated output, particularly when models operate without grounding in real-time financial databases. Andronie et al. highlight that even marginal deviations from factual correctness in financial communication can result in material misstatements, posing reputational and regulatory risks [3]. Furthermore, ensuring stylistic consistency with institutional communication standards remains non-trivial. As noted by Lee et al. [1], post-editing of AI-generated content is still required to meet the editorial and legal standards imposed on financial institutions. These additional oversight processes reduce some of the purported efficiency gains and constrain full automation.

However, structured integration of generative models with live financial data sources and editorial style guides presents a promising pathway forward. Zhang, Feng, and Dong illustrate how low-latency generative frameworks, when linked to structured market inputs, can generate real-time summaries with improved reliability [6]. Over time, as retrieval-augmented generation methods and firm-specific fine-tuning improve, the production of internal and client-facing texts is likely to become increasingly automated, consistent, and cost-effective. In this setting, human review may shift from a verification role to a strategic oversight function.

3.2. Customer Service and Conversational Agents

Despite rapid deployment, generative conversational systems remain limited in handling multi-turn dialogues, regulatory nuances, and intent ambiguity in financial queries. Saha, Shukla, and Kumar observe that customer-facing LLM agents often falter when required to interpret context-specific requests that go beyond predefined workflows [9]. Additionally, where user interactions pertain to sensitive personal or financial information, the lack of native interpretability in generated responses poses a risk to transparency and compliance. These concerns are exacerbated in crossjurisdictional settings, where multilingual support must also comply with region-specific financial regulations.

As dialogue management frameworks mature and firms implement model-specific guardrails, generative systems are expected to deliver increasingly personalised and compliant support. Desiraju and Khan find that financial firms integrating LLMs within defined query boundaries (e.g., transaction lookups, account status) have achieved measurable gains in response time and service availability [8]. Moreover, the use of hybrid models—where rule-based validation layers oversee generative outputs—offers a scalable architecture for balancing flexibility and accuracy. In high-volume service environments, such systems could significantly reduce human workload while enhancing consistency in service delivery.

3.3. Investment and Advisory Applications

The deployment of generative models in investment decision support raises concerns around both technical reliability and fiduciary appropriateness. Lakkaraju et al. demonstrate that LLMs, while capable of producing syntactically coherent recommendations, lack embedded financial reasoning and may suggest suboptimal or noncompliant strategies if used in isolation [10]. Furthermore, when used in retail advisory settings, such systems must conform to suitability and disclosure regulations that require justification based on objective investment criteria—a requirement LLMs cannot consistently satisfy. These limitations have restricted the role of generative models to content augmentation rather than autonomous recommendation.

Recent developments in modular multi-agent systems offer a structured response to these constraints. Han et al. propose dividing investment analysis into discrete tasks—such as macroeconomic synthesis, risk assessment, and performance attribution—each handled by specialised generative components. This approach allows greater control over output scope and interpretability [11]. In practice, such systems can support human advisors by drafting scenario narratives or simulating client-specific investment outcomes, thereby enhancing client engagement and planning capacity. The potential for improved personalisation at scale is particularly relevant for digital wealth management platforms seeking to extend advisory services without linear increases in staffing.

3.4. Regulatory Compliance and Risk Reporting

In high-stakes compliance domains, generative AI introduces both technical opacity and regulatory complexity. The IMF notes that outputs lacking provenance or consistency are incompatible with regulatory expectations for auditability and explainability. This is especially critical in functions such as suspicious activity reporting or KYC, where documentation must not only be accurate but also traceable to verifiable data sources. Jain et al. caution that uncontrolled deployment of generative models in compliance settings may result in inadvertent regulatory breaches, particularly in environments with strict reporting timelines and content constraints [4].

Targeted integration of LLMs with structured compliance ontologies and rule-based validation layers offers a feasible path forward. Khanvilkar and Kommuru present a hybrid framework where regulatory graph structures guide generative output, ensuring alignment with reporting standards [13]. This approach enhances both interpretability and consistency across outputs. In practice, LLMs may soon serve as drafting tools that assist compliance officers in generating reports, summarising audit trails, and identifying anomalous patterns—without removing human oversight. As regulatory agencies begin to issue clearer guidelines on AI auditability and documentation standards, such systems could play a central role in reducing compliance workloads and improving process transparency. Additionally, as Shabsigh and Boukherouaa point out, the systemic risks introduced by generative models in regulatory contexts—such as model drift, misinterpretation of legal language, and regulatory capture—remain underexplored and raise critical oversight challenges [12].

4. Conclusion

This paper examined how generative AI is being integrated into FinTech across four domains: automated content generation, customer service, investment advisory, and compliance reporting. As discussed in the introduction, these models offer a new layer of flexibility and scale beyond traditional automation. The analysis shows that generative systems enhance speed, personalisation, and operational efficiency. Yet, challenges persist—particularly around reliability, explainability, and regulatory alignment. Many applications still rely on human oversight to ensure accuracy and accountability. While technology is still in development, its role in financial services is growing. Future efforts should focus on establishing model evaluation standards and investigating hybrid systems that integrate generative outputs with expert validation, ensuring both innovation and control.

References

[1] Lee D K C, Guan C, Yu Y, Ding Q. A comprehensive review of generative AI in finance. FinTech, 2024, 3 (3): 460–478.

[2] Barde K, Kulkarni P A. Applications of generative AI in fintech. Proceedings of the Third International Conference on AI-ML Systems, 2023, October: 1–5.

[3] Andronie M, Blazek R, Iatagan M, Skypalova R, Uță C, Dijmărescu A, et al. Generative artificial intelligence algorithms in Internet of Things blockchain-based fintech management. Oeconomia Copernicana, 2024, 15 (4).

[4] Jain Y, Gupta S, Yalciner S, Joglekar Y N, Khetan P, Zhang T. Overcoming complexity in ESG investing: The role of generative AI integration in identifying contextual ESG factors. SSRN, 2023. Available at SSRN: 4495647.

[5] Gowda P G A N. Benefits and risks of generative AI in FinTech. Journal of Scientific and Engineering Research, 2024, 11 (5): 267–275.

[6] Zhang S, Feng Z, Dong B. LAMDA: Low-latency anomaly detection architecture for real-time cross-market financial decision support. Academia Nexus Journal, 2024, 3 (2).

[7] Shimamura T, Tanaka Y, Managi S. Evaluating the impact of report readability on ESG scores: A generative AI approach. International Review of Financial Analysis, 2025, 101: 104027.

[8] Desiraju K, Khan A. Enhancing customer support services in banking using generative AI. International Working Conference on Transfer and Diffusion of IT, 2023, December: 288–295.

[9] Saha B, Rani N, Shukla S K. Generative AI in financial institution: A global survey of opportunities, threats, and regulation. arXiv, 2025. arXiv:2504.21574.

[10] Lakkaraju K, Vuruma S K R, Pallagani V, Muppasani B, Srivastava B. Can LLMs be good financial advisors? An initial study in personal decision making for optimized outcomes. arXiv, 2023. arXiv:2307.07422.

[11] Han X, Wang N, Che S, Yang H, Zhang K, Xu S X. Enhancing investment analysis: Optimizing AI-agent collaboration in financial research. Proceedings of the 5th ACM International Conference on AI in Finance, 2024, November: 538–546.

[12] Khanvilkar K, Kommuru K. Regulatory graphs and GenAI for real-time transaction monitoring and compliance explanation in banking. arXiv, 2025.

[13] Shabsigh M G, Boukherouaa E B. Generative artificial intelligence in finance: Risk considerations. International Monetary Fund Working Paper, 2023.