

The Use of Generative AI to Provide Financial Advice: Opportunities and Risks

Artificial intelligence (AI) is rapidly reshaping the financial sector, with generative AI emerging as a particularly consequential development. The first paper in this Bretton Woods Committee series, *Governing Artificial Intelligence Across the Global Financial System*, established the foundational context for this transformation—mapping core capabilities, identifying key risk dimensions, and introducing the central concepts of explainability, governance, and systemic risk that will shape effective oversight.¹

This paper builds on that foundation by examining one of the most consequential and consumer-facing applications of generative AI: financial advice. The emergence of AI-enabled personal financial advisers has the potential to fundamentally alter how advice is delivered, shifting from episodic, transaction-based interactions to continuous, adaptive engagement. Such systems can expand access, reduce cost, and enable a level of personalization that was previously unattainable at scale. At the same time, they introduce new challenges with respect to accuracy, accountability, and consumer protection.

A comprehensive paper addressing cross-cutting regulatory frameworks and policy recommendations will be published separately as part of this series. Against that backdrop, the focus here is on understanding how generative AI is changing the nature of financial advice in practice—examining both the opportunities it unlocks and the risks it creates—and identifying the core elements required to support its responsible deployment.

AI IN FINANCIAL ADVISORY SERVICES—PERSONALIZATION AND COMMUNICATION

Generative artificial intelligence has the potential to materially reshape the provision of personalized financial advice by altering both the form and substance of advisory interactions. Unlike earlier digital tools that relied on static questionnaires, predefined decision trees, and episodic outputs, generative AI enables continuously adaptive engagement. Such systems can synthesize personal financial information, respond dynamically to user queries, and deliver guidance that evolves as circumstances change. As a result, financial

¹ Bretton Woods Committee Working Group on Artificial Intelligence in Finance, *Governing Artificial Intelligence Across the Global Financial System* (October 2025), <https://brettonwoods.org/governing-artificial-intelligence-across-the-financial-system/>.

advice increasingly resembles an ongoing dialogue rather than a series of discrete recommendations, fundamentally changing how individuals interact with financial decision-making.

The roots of AI-enabled financial advice predate generative models. Early robo-advisers such as Betterment and Wealthfront used algorithmic techniques to automate portfolio construction, rebalancing, and tax optimization based on client-provided inputs. Such platforms expanded access to low-cost investment management and demonstrated the scalability of automated advice, but their functionality remained largely transactional and modular.² The emergence of large language models marks a shift from automation toward interaction. Instead of producing bounded outputs, generative systems engage users in open-ended conversation, interpret natural-language questions, and provide contextual explanations. This transition broadens the scope of AI-enabled advice while simultaneously introducing new governance, consumer protection, and supervisory challenges.

AI-enabled financial advice now operates across three distinct models. The *first model* consists of

direct-to-consumer personal AI advisers—chat-first systems that engage directly with individuals and are often marketed explicitly as “AI advisers.” These tools answer questions such as how to rebalance a portfolio, whether a major purchase is affordable, or how financial goals interact over time, frequently drawing on linked financial accounts and personal data.³ The *second model* involves adviser copilots deployed within regulated firms, where AI systems support human advisers by drafting communications, summarizing meetings, and surfacing relevant research, while preserving the human adviser as the formal provider of advice.⁴ The *third model* includes traditional algorithmic, or robo-advice, platforms, which focus on automated portfolio management and are increasingly enhanced by generative interfaces but remain grounded in predefined models and workflows. Algorithmic, or robo-advice, platforms are described in the preceding paragraph.

This paper focuses exclusively on the first category—direct-to-consumer personal AI advisers—because it is in this domain that generative AI most directly interfaces with retail consumers and where supervisory authorities have expressed the greatest concern. European regulators, in particular, have cautioned

Figure 1. AI models in financial advice

Model	Description	Key Features
1. Personal AI advisers	Interact directly with individuals/consumers and provide advice	• Answer questions about portfolio rebalancing, major purchases, and financial goals • Use linked financial accounts and personal data
2. Adviser copilots	Support human advisers	• Draft communications and surface relevant research • Human adviser remains the formal provider of advice
3. Robo-advisers	Traditional automated platforms, now enhanced by generative AI	• Predefined models and workflows that focus on portfolio management • Increasingly use generative AI for user interaction

² Aditya Kashyap, “AI-Driven Financial Advisory: The Rise of Robo-Advisors,” available at SSRN, <https://doi.org/10.2139/ssrn.5268858>.

³ JD Power, “As More U.S. Consumers Struggle with Rising Prices, Many Turn to Artificial Intelligence for Financial Advice,” August 28, 2025, <https://www.jdpower.com/business/resources/more-us-consumers-struggle-rising-prices-many-turn-artificial-intelligence>.

⁴ Charles Schwab Advisor Services, *Schwab Advisor AI in Action: AI’s Transformative Impact on RIAs*, Guiding Principles Series (January 2026), <https://advisorservices.schwab.com/resource/advisor-AI-in-action>.

that publicly available AI tools are capable of generating highly persuasive investment guidance that may blur the line between general financial information and regulated advice.⁵ The combination of conversational delivery, personalization, and perceived authority places these systems in a uniquely influential position within the retail financial ecosystem.

To be clear, fully autonomous, direct-to-consumer AI financial advisers are not yet operational among major financial institutions. Current deployments focus almost exclusively on adviser-facing tools—Morgan Stanley’s Aladdin Auto Commentary for personalized portfolio insights,⁶ Citi Wealth’s AskWealth for adviser support,⁷ and Franklin Templeton’s AI-powered analyst productivity tools,⁸ all of which keep human professionals firmly in the advisory loop. The pure, direct-to-client adviser remains, for now, an early-stage architecture rather than a commercial reality.

Yet the direction of travel is unmistakable. Industry adoption is accelerating at a remarkable pace: A January 2026 study by Schwab Advisor Services found that 63 percent of registered investment advisers now use AI tools, more than double the rate just two years earlier, with 68 percent expecting AI to transform financial advice within three years.⁹ McKinsey estimates that generative AI alone could deliver an 8 percent efficiency gain in investment management, while a midsize asset manager could potentially cut 25 to 40 percent from its cost base through broader AI adoption.¹⁰ Most tellingly, Goldman Sachs is already building autonomous

AI agents with Anthropic for internal operations,¹¹ and Bank of America identifies agentic AI as a rapid next wave of capabilities “arriving before chatbots and copilots reach version 2.0.”¹² The infrastructure for autonomous advice is being built, piece by piece, across the industry.

FRAMING “GOOD” AI IN FINANCIAL ADVICE

Before proceeding further, we should frame what we mean by “good” AI in financial advice. Good AI is not defined solely by outputs, but by the integrity of the processes that generate them. That includes robust data governance, disciplined model development and validation, clear boundaries on use cases, and ongoing monitoring to detect drift, degradation, or unintended behaviors. At the same time, outputs must meet a baseline standard of accuracy and should avoid hallucination or fabrication of information, with clear alignment to the system’s stated purpose. In practice, such systems likely would not rely on a single model or output, but would instead draw on multiple models, prompts, and scenario analyses, with results synthesized to produce a higher degree of robustness and reliability.

Equally important, AI education should be understood as a critical complement to the deployment of AI in financial advice. Building trust is central to adoption, and trust requires that users—both clients and professionals—have a clear understanding of what the system is designed to do, its limitations, and how its outputs

5 European Securities and Markets Authority, *Using Artificial Intelligence for Investing: What You Should Consider* (investor warning factsheet) (March 25, 2025), https://www.esma.europa.eu/sites/default/files/2025-03/ESMA_Warning_on_the_use_of_AI_-_EN.pdf.

6 “Aladdin® by BlackRock—Software for Portfolio Management,” Blackrock, accessed January 17, 2026, <https://www.blackrock.com/Aladdin>.

7 “Citi Wealth Launches ‘Advisor Insights’ Pilot and ‘AskWealth,’ AI-Driven ‘Gamechangers’ for Client Communications,” Citigroup, accessed January 28, 2026, <https://www.citigroup.com/global/news/press-release/2025/citi-wealth-launches-advisor-insights-askwealth>.

8 “Franklin Templeton Scales Investment Insight,” Databricks, accessed January 12, 2026, https://www.databricks.com/customers/franklin-templeton/agent-bricks?utm_source=bambu&utm_medium=social&utm_campaign=advocacy&blaid=8213697.

9 Mike Peterson, “Schwab Study Reveals RIA AI Adoption More Than Doubles—but Most Firms Still in Early Stages,” Charles Schwab, January 22, 2026, <https://pressroom.aboutschwab.com/press-releases/press-release/2026/Schwab-Study-Reveals-RIA-AI-Adoption-More-Than-Doubles--But-Most-Firms-Still-in-Early-Stages/default.aspx>.

10 Jonathan Godsall, Philipp Koch, Piyush Sharma, Raj Bector, and Todd Pingaro, with Alena Fedorenko, “How AI Could Reshape the Economics of the Asset Management Industry,” McKinsey & Company, July 16, 2025, <https://www.mckinsey.com/industries/financial-services/our-insights/how-ai-could-reshape-the-economics-of-the-asset-management-industry>.

11 “Goldman Sachs Teams Up with Anthropic to Automate Banking Tasks with AI Agents, CNBC Reports,” Reuters, February 6, 2026, <https://www.reuters.com/business/finance/goldman-sachs-teams-up-with-anthropic-automate-banking-tasks-with-ai-agents-cnbc-2026-02-06/>.

12 Vanessa Cook, “The New Wave: Agentic AI,” Bank of America Institute, January 2025, <https://institute.bankofamerica.com/transformation/agentic-ai.html>.

should be interpreted. Taken together, strong governance, disciplined processes, high-quality outputs, and informed users form the foundation of trustworthy AI-enabled financial advice.

GENERATIVE AI CHANGES THE NATURE OF FINANCIAL ADVICE

Generative AI introduces several qualitative changes to personal financial advice. Most notably, advice becomes tailored, conversational, and continuous. Rather than delivering a one-time financial plan or periodic recommendations, personal AI advisers enable *iterative engagement*. Users can ask follow-up questions, test hypothetical scenarios, and receive step-by-step guidance that evolves over time. This conversational loop transforms advice from an episodic service into an ongoing process of financial navigation, with the AI system effectively accompanying users through successive decisions.

At the same time, personalization shifts from broad segmentation to something closer to a “financial digital


twin.” Traditional advisory models rely on generalized assumptions and periodic updates, constrained by time and human capacity. Generative AI, when paired with consent-based data aggregation, can tailor guidance to an individual’s specific cash flows, liabilities, tax position, benefits, and long-term objectives. Recommendations can adjust *dynamically* as market conditions shift or life events occur, producing advice that is not only individualized but continuously refreshed.¹³

Finally, generative AI enables increasingly *agentic* forms of advice. Beyond responding to questions, these systems can propose actions, such as portfolio adjustments, savings strategies, or planning updates, and can coordinate follow-up steps across financial workflows. While many current implementations include human approval gates in regulated contexts, the trajectory is toward systems that not only recommend outcomes but also orchestrate execution. This evolution invites closer scrutiny of issues of accountability and control, including how responsibility is divided between firms, systems, and users.¹⁴

Figure 2. Evolution of financial advice



Figure 3. Transformation of financial advice with generative AI



¹³ Alex Cosmas, Guilherme Cruz, Sebastian Cubela, Mark Huntington, Sohrab Rahimi, and Sanchit Tiwari, “Digital Twins and Generative AI,” McKinsey & Company, April 11, 2024, <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/digital-twins-and-generative-ai-a-powerful-pairing>.

¹⁴ Andrew Renzella, “Agentic AI in Financial Services: The Future of Autonomous Finance Solutions,” AWS Blogs, September 8, 2025, <https://aws.amazon.com/blogs/aws/marketplace/agentic-ai-solutions-in-financial-services/>.

UNLOCKING BENEFITS IN ACCESS, COST, AND PERSONALIZATION

The potential benefits of the generative AI model are significant. AI advisers address long-standing structural constraints in the delivery of financial advice, particularly those related to scale, cost, and access. Historically, high-quality financial guidance has been limited by human capital, pricing structures, and minimum asset thresholds. Generative AI offers a credible mechanism to transcend those constraints by enabling continuous, personalized interaction at a marginal cost approaching zero. If deployed responsibly, such systems could increase engagement in financial planning and improve outcomes for a much broader segment of the population.

Access is the most immediate benefit. Personal AI advisers remove many of the traditional barriers associated with human-delivered advice, such as minimum account sizes. In traditional advisory models the economics of human-delivered advice typically require account minimums. Because advisers are often compensated through fees tied to assets under management and must allocate time to each client relationship, small accounts are often uneconomic to serve—whereas AI advisers can operate at near-zero marginal cost. As a result, individuals who may never interact with a human adviser can nonetheless engage with planning tools that address retirement readiness, savings strategies, and investment trade-offs. In this respect, personal AI advisers function not merely as efficiency enhancements but as access-expanding infrastructure, lowering the threshold for participation in long-term financial planning and investment decision-making.^{15,16}

AI advisers also expand access by *eliminating geographic and availability constraints* associated with traditional advisory models. Human advisers are typically concentrated in major financial centers and operate within limited business hours, which can restrict access for individuals in smaller markets or those with different schedules. AI systems, by contrast, can provide financial guidance to anyone with an internet connection and are available continuously and tirelessly. As a result, individuals can access planning tools, portfolio insights, and financial advice regardless of where they live or when they need assistance.

The *reduction in cost* reinforces the gains in access. Traditional advisory services are labor intensive, with pricing structures that reflect both adviser compensation and operational overhead. Generative AI systems can automate substantial portions of the advisory workflow, including data analysis, scenario modeling, and client communication. The automation compresses the marginal cost of delivering advice and enables alternative pricing models, ranging from low-cost subscriptions¹⁷ to free access supported by other revenue streams. While cost efficiencies do not eliminate the need for human oversight—particularly for complex or high-stakes decisions—they have the potential to reshape expectations around what constitutes a reasonable baseline for financial guidance.

Personalization represents the most transformative benefit. Generative AI allows advice to be tailored continuously and at scale, incorporating real-time financial data and contextual information that would be difficult for human advisers to process consistently.^{18,19} Advice can evolve in response to changes in income, market conditions, or personal circumstances, supporting

15 Michael Reher and Stanislav Sokolinski, "Robo Advisors and Access to Wealth Management," *Journal of Financial Economics* 155: 103829, <https://doi.org/10.1016/j.jfineco.2024.103829>.

16 Stacie Calder, "How AI Is Making Financial Advice More Accessible," World Economic Forum, June 16, 2025, <https://www.weforum.org/stories/2025/06/ai-financial-advice-accessible/>.

17 Andrea Coombes and Taryn Phaneuf, "How Much Does a Financial Advisor Cost?" NerdWallet, updated March 26, 2026, <https://www.nerdwallet.com/financial-advisors/learn/how-much-does-a-financial-advisor-cost>.

18 Josipa Majic Predin, "AI-Powered Financial Planning and the Rise of Personalized Financial Independence Tools," *Forbes*, updated November 12, 2025, <https://www.forbes.com/sites/josipamajic/2025/11/06/ai-powered-financial-planning-and-the-rise-of-personalized-financial-independence-tools/>.

19 Ganesh Kumaran Ramalingam, "AI-Driven Financial Advice: Its Impact on Household Financial Decision-Making," January 20, 2025, available at SSRN, <http://dx.doi.org/10.2139/ssrn.5143410>.

more timely and relevant guidance. Over time, this shift may move financial advice away from static planning exercises toward continuous financial navigation.

STRUCTURAL VULNERABILITIES IN AI-ENABLED ADVICE

The aforementioned benefits, however, are inseparable from a set of significant risks that are amplified by the scale, autonomy, and persuasive capacity of generative AI. Because personal AI advisers interact directly with consumers—often without immediate human intermediation—errors or misaligned incentives can propagate rapidly and at scale. As a result, their risk profile differs materially from that of traditional advisory tools.²⁰

A central concern is *accuracy*. Generative AI systems are optimized to generate coherent and contextually plausible responses, not to validate factual accuracy or ensure internal consistency across complex financial assumptions. As a result, they may produce outputs that appear authoritative yet rest on incorrect premises, incomplete information, or flawed reasoning. In a financial advisory context, such failures can materially affect investment decisions, tax planning, debt management, or long-term financial strategies, with consequences that may not become apparent until well after decisions have been executed.²¹

This risk is amplified by the conversational nature of personal AI advisers. Iterative dialogue can create a sense of continuity and trust that masks underlying uncertainty or model limitations. Unlike human advisers, AI systems may not reliably recognize when a question exceeds their competence, when critical information is missing, or when assumptions embedded in prior responses should be revisited.²² When deployed at scale, even infrequent *hallucinations* can propagate broadly, particularly if users rely on AI-generated

guidance as a primary decision-making input rather than as a supplement to professional judgment.²³

In the context of generative AI, a *hallucination* refers to a situation where an AI model produces information that sounds confident and plausible but is incorrect, misleading, or entirely fabricated. Hallucinations occur because generative models predict likely sequences of words or outputs based on patterns rather than verifying facts; they are more likely when prompts are ambiguous, when the model lacks reliable information, or when it is pushed beyond its training or contextual limits. Depending on context, model hallucinations can have high-stakes consequences: For example, a legal research assistant might rely on generative AI that invents court cases or citations that appear legitimate but do not actually exist. In *Mata v. Avianca* (678 F. Supp. 3d 443 (2023)), attorneys submitted a U.S. federal court brief containing AI-generated legal citations that were later discovered to be entirely fabricated by the model. The federal judge imposed a \$5,000 sanction on the two lawyers and their law firms and ordered them to notify their client and the judges whose names appeared in the fake opinions.

Bias in model outputs is a related structural concern. Generative AI systems learn from large datasets that reflect historical financial behavior, market dynamics, and institutional practices. Such patterns may embed assumptions or correlations that influence assessments of risk tolerance, planning horizons, or recommended strategies in ways that are not economically optimal for individual users. A 2025 Cornell study found that state-of-the-art large language models (LLMs) exhibited systematic product biases in their investment recommendations, favoring stocks like AAPL from Apple and MSFT from Microsoft. Notably, the preferences persisted despite debiasing techniques, with far-ranging

20 Sarah Hammer, “Smart Algorithms, Smarter Investing: How AI-Powered Financial Advice Will Help You,” MarketWatch, updated September 7, 2024, <https://www.marketwatch.com/story/smart-algorithms-smarter-investing-how-ai-powered-financial-advice-will-help-you-ed87ac0f>.

21 Andrew W. Lo and Jillian Ross, “Generative AI from Theory to Practice: A Case Study of Financial Advice,” An MIT Exploration of Generative AI, March 27, 2024, <https://mit-genai.pubpub.org/pub/l89uu140/release/2>.

22 Robert Booth, “UK Consumers Warned over AI Chatbots Giving Inaccurate Financial Advice,” *The Guardian*, November 18, 2025, <https://www.theguardian.com/technology/2025/nov/18/warning-ai-chatbots-inaccurate-financial-advice-tips-chatgpt-copilot-uk>.

23 Lo and Ross, “Generative AI from Theory to Practice.”

implications for consumer financial security and even financial stability at large.²⁴

Model bias refers to systematic tendencies in a model's outputs that arise from patterns embedded in its training data, model architecture, or optimization process. This can be problematic because biased outputs systematically distort information, recommendations, or predictions in ways that are difficult for users to detect. In financial contexts, such bias could influence portfolio recommendations, risk assessments, or credit-related guidance if the model reflects historical data patterns or incomplete market information. Explainability techniques, such as feature attribution, model interrogation, and output analysis, can help unpack these dynamics by revealing how inputs and training data influence the model's generated responses.

Another study found that LLM-generated investment advice “exhibit[s] similar cognitive biases as human investors” such as aggressive trend following and stock picking—recommendations that are misaligned with modern portfolio theory.²⁵ Because such effects are often indirect and embedded within complex model architectures, they can be difficult to detect through traditional testing or disclosure mechanisms.²⁶ At scale, consistent directional bias in recommendations may distort financial decision-making outcomes across large user populations, even absent any intent to discriminate.²⁷

Privacy risks are also heightened. Personal AI advisers require extensive access to sensitive financial and personal information, including transaction histories, income patterns, liabilities, benefits, and behavioral

data. The aggregation and continuous analysis of this information heighten traditional risks associated with data security, unauthorized access, and improper use. Breaches or misuse in this context can expose not only static account information but dynamic insights into an individual's financial life.²⁸

Generative AI introduces additional privacy challenges through inference, as these systems can derive conclusions about user preferences, financial behavior, or future actions that were never explicitly disclosed. In a financial advisory context, such inferred attributes may influence recommendations or suggested actions in ways that are opaque to users and difficult to trace or contest. This raises important questions about consent, transparency, and user control, particularly where individuals may not understand how their data—and the inferences drawn from it—are shaping the advice they receive.²⁹

Inference refers to the process by which a trained model applies learned patterns to new inputs to generate outputs. This capability is powerful because it allows AI systems to infer relationships, produce insights, and respond to novel situations without explicit programming. At the same time, inference can create risks if the model infers sensitive attributes about individuals from limited information. For example, research shows that models trained on large datasets can predict sensitive characteristics such as health status, political affiliation, or socioeconomic status from indirect signals like purchasing behavior, writing style, or geolocation patterns. Such attribute inference can expose private information even when users never explicitly disclose it. A dangerous example

24 Yuhao Zhi, Xiaoyu Zhang, Longtian Wang, Shumin Jiang, Shiqing Ma, Xiaohong Guan, et al., “Exposing Product Bias in LLM Investment Recommendation,” March 11, 2025, available at arXiv.org, <https://arxiv.org/abs/2503.08750>.

25 Philipp Winder, Christian Hildebrand, and Jochen Hartmann, “Biased Echoes: Large Language Models Reinforce Investment Biases and Increase Portfolio Risks of Private Investors,” *PLoS One* 20 (6): e0325459, <https://doi.org/10.1371/journal.pone.0325459>.

26 Gissela Moya and Vinhcent Le, “Algorithmic Bias Explained: How Automated Decision-Making Becomes Automated Discrimination,” The Greenlining Institute, February 2021, <https://greenlining.org/wp-content/uploads/2021/04/Greenlining-Institute-Algorithmic-Bias-Explained-Report-Feb-2021.pdf>.

27 “Model bias” refers to systematic discrimination that is embedded into an AI model, often unconsciously, through its design or training data. If unchecked, AI systems can generate prejudiced outputs that reinforce harmful stereotypes or even propagate harm—see examples including the COMPAS recidivism algorithm, Amazon's hiring tool, or iTutor Group's job application software.

28 Alexandra Jonker and Alice Gomstyn, “Exploring Privacy Issues in the Age of AI,” IBM, <https://www.ibm.com/think/insights/ai-privacy>.

29 Xiaoyu Zhang, Chao Chen, Yi Xie, Xiaofeng Chen, Jun Zhang, and Yang Xiang, “A Survey on Privacy Inference Attacks and Defenses in Cloud-Based Deep Neural Network,” *Computer Standards & Interfaces* 83 (2023): 103672, <https://www.sciencedirect.com/science/article/abs/pii/S0920548922000435>.

of this can be found in AI systems' capacity to synthesize scientific knowledge in ways that could assist malicious actors. A RAND study* found that advanced AI systems can describe technical steps to recovering a live virus from synthetic DNA constructs, illustrating how AI could lower knowledge barriers for developing biological weapons if misused.

* Roger Brent and Greg McKelvey Jr., *Contemporary Foundation AI Models Increase Biological Weapons Risk* (RAND, 2025), <https://www.rand.org/pubs/perspectives/PEA3853-1.html>.

Some have also queried whether widespread adoption of similar AI systems could lead to *model convergence*, with different models arriving at the same or highly correlated recommendations. To the extent that occurs, it could contribute to more aligned investment behavior and, in certain conditions, could reinforce market trends. Work by the Bank for International Settlements notes that reliance on common data, models, and third-party providers may increase the potential for herding dynamics, with possible implications for financial stability.³⁰

Conflicts of interest remain the most complex and unresolved concern. Generative AI systems may be optimized for engagement, retention, or monetization in ways that do not align cleanly with traditional fiduciary or best-interest standards. Existing conflict-of-interest frameworks for broker-dealers³¹ and investment advisers³² were designed for human decision-makers and do not map neatly onto AI systems, whose objectives are embedded in model design and training. When personalization is driven by engagement metrics or revenue incentives, the risk is that advice may be subtly shaped to favor certain products, behaviors, or outcomes without transparent disclosure.

SECTOR-SPECIFIC REGULATION AND HORIZONTAL AI FRAMEWORKS

Industry-specific regulatory frameworks governing financial advice already provide a dense and well-developed set of obligations that apply directly to AI-enabled advisory tools. As set forth in the appendix, in the United States, for example, Financial Industry Regulatory Authority (FINRA) rules governing communications with the public, investment advice, and research—including FINRA Rule 2210³³—apply regardless of whether content is generated by a human adviser or an algorithm. These rules impose long-standing requirements that communications be fair, balanced, and not misleading; that recommendations satisfy suitability or best-interest standards; that investment advice is provided in accordance with applicable fiduciary duties; and that research and analytical tools be subject to supervision, disclosure, and conflict controls. Similar conduct-based regimes exist in other jurisdictions, including the Markets in Financial Instruments Directive II (MiFID II) in the European Union and the Financial Conduct Authority's conduct and suitability rules in the United Kingdom. In each case, the regulatory premise is technology neutral: The legal obligation attaches to the activity rather than the means by which that activity is performed. As a result, firms deploying personal AI financial advisers are already subject to an extensive set of sector-specific rules, even in the absence of AI-targeted legislation.

At the same time, regulators have begun to issue AI-specific supervisory guidance that clarifies how the existing frameworks should be interpreted in the

30 Juan Carlos Crisanto, Cris Benson Leuterio, Jermy Prenio, and Jeffery Yong, *Regulating AI in the Financial Sector: Recent Developments and Main Challenges* (Bank for International Settlements, 2024), <https://www.bis.org/fsi/publ/insights63.pdf>.

31 The U.S. regulatory framework governing retail-focused broker-dealer investment recommendations is Regulation Best Interest (Regulation BI), adopted by the U.S. Securities and Exchange Commission (SEC) and enforced in practice through examinations and supervision by the Financial Industry Regulatory Authority. Regulation BI requires broker-dealers and their associated persons to act in the best interest of retail customers when making securities recommendations, including obligations related to disclosure, care, conflict-of-interest management, and compliance. Regulation Best Interest: The Broker-Dealer Standard of Conduct, 17 C.F.R. § 240.15l-1, adopted by the SEC pursuant to the Securities Exchange Act of 1934.

32 The regulatory framework governing investment advice provided by investment advisers is established under the Investment Advisers Act of 1940 and administered by the SEC. The statute imposes a fiduciary duty on registered investment advisers when providing advice to clients, requiring them to act in the client's best interest, including duties of care and loyalty, along with obligations related to disclosure, conflicts of interest, and compliance oversight. Investment Advisers Act of 1940, 15 U.S.C. §§ 80b-1 et seq. and SEC Commission Interpretation Regarding Standard of Conduct for Investment Advisers, Release No. IA-5248 (June 5, 2019).

33 FINRA Rule 3110—Supervision, FINRA Rule 3120—Supervisory Control System, FINRA Rule 4511—Books and Records.

context of generative technologies. A leading example is the European Securities and Markets Authority's public statement on the use of artificial intelligence in investment services, which explicitly affirms that AI-driven advice must comply with MiFID II's organizational, conduct, and best-interest requirements, while highlighting additional expectations around governance, testing, and oversight of AI systems. Similar messages have emerged from global standard setters such as the International Organization of Securities Commissions (IOSCO) and from professional bodies, including the Certified Financial Planner Board of Standards (CFP Board), which has issued guidance on the ethical use of generative AI in financial planning. Together, these AI-specific statements function as interpretive overlays for existing regulation that reinforce accountability and risk management, underscoring that firms remain responsible for AI-generated outcomes. The full list of applicable statutes, rules, and supervisory guidance across jurisdictions is set out in the appendix.

BUILDING AN EFFECTIVE AI GOVERNANCE AND OVERSIGHT FRAMEWORK

Although a more comprehensive framework for AI governance will be set out in a forthcoming paper, it is useful to establish foundational principles here. Effective AI governance should be driven from the top. A clear senior sponsor—ideally at the chief executive officer level—signals that AI is a strategic priority and ensures alignment with the firm's risk appetite and business objectives. Firms should establish an executive-level AI committee with cross-functional representation (business, technology, risk, legal, and compliance) to set direction, approve use cases, and oversee key risks. This governance should be anchored in a formal responsible AI policy that defines permissible uses, establishes boundaries, and articulates expectations for accuracy, appropriate use, and ongoing monitoring. The objective is not to centralize all decisions, but to create clear accountability, escalation pathways, and consistency across the enterprise.

A strong model for AI governance is grounded in the three-lines-of-defense framework. Implementation should sit squarely within the first line, which is responsible for developing, deploying, and operating AI systems in accordance with the firm's standards, including embedded risk management practices. The second line—including compliance and model risk oversight—provides independent review, effective challenge, and ongoing monitoring of adherence to regulatory expectations and internal policies. The third line, internal audit, should periodically assess the effectiveness of the overall AI governance framework, including controls, oversight structures, and issue remediation. Taken together, a three-lines-of-defense model helps ensure that AI is governed as a core enterprise capability, with clear ownership, disciplined oversight, and continuous validation.

CONCLUSION

Generative AI has the potential to fundamentally reshape financial advice by expanding access, reducing cost, and enabling continuous, highly personalized engagement. Generative AI systems offer a credible path to democratizing financial guidance and improving decision-making outcomes at scale. At the same time, their capabilities introduce meaningful risks that are amplified by their autonomy, reach, and persuasive nature. The central challenge is not whether AI should be used in financial advice, but how it can be deployed in a manner that preserves the integrity of advice and protects end users.

An understanding of what constitutes "good" AI provides a useful framework for interpreting these tradeoffs and grounding practical takeaways. Good AI is grounded in disciplined processes, governed through clear enterprise accountability, and designed to produce accurate, bounded, and purpose-aligned outputs. It is supported by strong governance frameworks anchored at the senior-most levels of the firm, implemented through the three lines of defense, and reinforced by existing financial regulatory regimes that apply regardless of technology. Critically, trust must be actively built—through transparency, user understanding,

and ongoing oversight—rather than assumed. As AI continues to evolve from assistive tools toward more autonomous systems, firms and regulators alike will need to ensure that these foundational elements remain

intact. The trajectory of AI in financial advice is clear; the imperative now is to shape it in a way that is both innovative and worthy of the trust it seeks to earn.

The author would like to thank Allyson Ye for her assistance in completing this paper.

APPENDIX

United States

Financial Industry Regulatory Authority (FINRA) (Broker-Dealers and Associated Persons)

Communications with the public

- **FINRA Rule 2210 (Communications with the Public)**—Core standards for retail communications, including “fair and balanced” content and no misleading omissions; applies to AI-generated client-facing content and prompts/outputs used in marketing and client communications.
- **FINRA communications exam/risk guidance**—FINRA explicitly frames Rule 2210 as principles-based and intended to apply as communications technology evolves (high relevance to AI-generated messaging).

Investment advice/recommendations

- **FINRA Rule 2111 (Suitability)**—Suitability obligations for recommendations of securities transactions/strategies; directly implicated if an AI tool is making or shaping a “recommendation.”
- **FINRA Rule 2214 (Investment Analysis Tools)**—Specific requirements and constraints when offering tools that project outcomes or provide investment analysis; relevant if the “AI adviser” is packaged as a tool that generates modeled results or scenario outputs.

Research

- **FINRA Rule 2241 (Research Analysts and Research Reports)**—Applies when the AI output qualifies as “research report” (equity research) and triggers independence/conflict and disclosure frameworks.
- **FINRA Rule 2242 (Debt Research Reports and Analysts)**—Parallel framework for debt research.

Generative AI-specific FINRA guidance

- **FINRA Regulatory Notice 24-09**—Explicit reminder that existing FINRA rules apply when firms use generative AI/large language models (including supervisory systems, recordkeeping, communications).

U.S. Securities and Exchange Commission (SEC) (Investment Advisers and Broker-Dealers)

Broker-dealers

- **Regulation Best Interest (Reg BI) (Exchange Act Rule 15l-1)**—“Best interest” standard for broker-dealer recommendations to retail customers; applies when the AI system (or the firm through the AI system) is making a recommendation of a securities transaction or strategy.

Investment advisers

- **Investment Advisers Act fiduciary standard (SEC interpretation)**—SEC articulation that advisers are fiduciaries, including duty of care and loyalty; foundational if the AI adviser is offered under a registered investment adviser (RIA) model or used in delivering advice.
- **SEC Marketing Rule (Rule 206(4)-1)**—Governs adviser marketing/advertising; directly relevant if the AI adviser’s outputs are used in marketing, performance presentations, testimonials/endorsements, or “hypothetical performance” style outputs.

SEC staff guidance on digital/robo advice (relevant by analogy for “AI advisers”)

- **SEC Division of Examinations Risk Alert on electronic investment advice/robo-advisers**—Flags compliance issues around disclosures, fiduciary obligations, and controls that map cleanly to AI-driven advice experiences.

Other U.S. Agencies (Important Adjacent Constraints)

- **Federal Trade Commission (consumer protection/advertising)**—FTC actions and guidance emphasize that “AI” claims and AI-enabled conduct remain subject to deception/unfairness standards; relevant for “AI adviser” marketing, capabilities claims, and potentially harmful product design.

- **Consumer Financial Protection Bureau (CFPB) (consumer financial protection; adjacent for household finance workflows)**—CFPB has stressed there is no AI exemption from federal consumer financial protection requirements (not securities advice per se, but often implicated when AI advisers touch credit, payments, and consumer finance).

European Union

Investment services framework (MiFID II lens)

- **Markets in Financial Instruments Directive II (MiFID II) conduct and organizational requirements (including “best interest of the client” and suitability/appropriateness concepts)**—Baseline obligations for investment firms providing investment services.
- **European Securities and Markets Authority (ESMA) Public Statement on AI in investment services (May 30, 2024)**—Directly on-point guidance: AI use must comply with MiFID II, with emphasis on organizational requirements, conduct requirements, and acting in clients’ best interests.

Data protection

- **General Data Protection Regulation (GDPR) Article 22 (automated decision-making/profiling)**—Constraints and safeguards around solely automated decisions that have legal or similarly significant effects; relevant if AI advice outputs drive consequential outcomes without meaningful human involvement.

Horizontal AI regulation

- **EU AI Act (general framework)**—Not advice-specific, but can become relevant depending on system classification (e.g., high-risk categories, general-purpose AI obligations) and deployment model.

United Kingdom

Advice and suitability

- **Financial Conduct Authority (FCA) Conduct of Business Sourcebook (COBS) suitability obligations (COBS 9/9A)**—Suitability assessments for personal recommendations; relevant if an AI adviser crosses into personal recommendation territory.

Communications/promotions

- **FCA COBS 4 (Communicating with clients; financial promotions)**—“Fair, clear and not misleading” communications; highly relevant for AI-generated marketing, “explanations,” and product prompts/outputs used externally.
- **FCA Consumer Duty (PRIN 2A)**—Outcome-focused obligations for retail customers; materially relevant to AI adviser user-experience design, comprehension, choice architecture, and communications.
- **FCA guidance on social media promotions (FG24/1)**—Reinforces how Consumer Duty standards apply to communications and promotions on social media (relevant if AI advisers or AI-generated content are distributed through those channels).

UK data protection

- **UK GDPR automated decision-making and profiling (Article 22 analogue)**—Similar constraints and safeguards as EU GDPR; relevant to automated, consequential decisions.

IOSCO (Global Supervisory Expectations)

- **International Organization of Securities Commissions (IOSCO) FR06/2021: AI and machine learning (ML) by market intermediaries and asset managers**—Baseline supervisory expectations around governance, oversight, controls, and risk management for AI/ML use.

Professional Standards (Financial Planning)

- **Certified Financial Planner Board of Standards (CFP Board) Code of Ethics and Standards of Conduct (fiduciary duty)**—CFP professionals must act as fiduciaries when providing financial advice; directly relevant where CFPs use AI tools in advice delivery or where a product is marketed to/through CFP professionals.
- **CFP Board Generative AI Ethics Guide (February 2025)**—Practical checklist addressing privacy, data integrity, bias, and other professional-ethics issues when CFP professionals use generative AI.







Artificial Intelligence Working Group

CO-CHAIRS: William C. Dudley and Sarah Hammer

WORKING GROUP: Oluyomi Adebote, Paul Andrews, Beverly Barrett, Daniela Bassan, Richard Berner, Spyridon Bisidis, David Borcsok, Wesley Brown de Colstoun, Yisroel Brumer, Camille Bryan, Onorina Buneanu, F. Christopher Calabia, Christian Callender-Easby, Massimo Carassinu, Keith B. Carter, Alexander James Chesney, Aleksandar Colic, Gurunath Dasari, Alexandra de Hoop Scheffer, Jag DeSaint Dior, Qahir Dhanani, Alexandra Dimitrijevic, Mona Elliot, Anthony Elson, Jonathan R. Everhart, Omer Fein, Pilar Cristina Frank, Nitin Gaur, Samim Ghamami, Daniel L. Gleizer, Jeffrey Greenbaum, Jon Hartley, Olive Hsiao, Jason Hsu, Michael Hsu, Linda Jeng, Peter E. Johnson, Oliver Kellman, Ron Korland, Mahesh Kotecha, Abhesh Kumar, Pat LaVecchia, Fergus McCormick, Tetsuharu Minowa, Margarita Musante, Patrick Noonan, Gregory J. Peters, Caroline Pham, Maximilien Pucci-Sisti Maisonrouge, Sebastian Robertson, Gursimran Rooprai, Sunil Sabharwal, Paul R. Samson, Jason Schenker, Lemma W. Senbet, Irina Shaorshadze, Adnan Pasha Siddiqui, Michael F. Silva, Divya Singaravelu, Maggie Sklar, Danica Starks, Radek Stech, Rajesh Swaminathan, Raffaella A. Tenconi, Hilary F. Till, Siddharth Tiwari, Helen Turner, Tobin G. Van Ostern, Markus Veith, Manuela Veloso, Andreas Veneris, Akshay Vijayendiran, David Webster, Rex Wempen, Linda Yueh, Bryan Zhang



THE BRETTON WOODS COMMITTEE
1701 K St NW #950, Washington, DC 20006
www.brettonwoods.org