

GENERATIVE PRETRAINED TRANSFORMERS FOR INVESTOR-CENTRIC PORTFOLIO CONSTRUCTION

Dimitrios Papakyriakopoulos¹

Manolis Kritikos²

Abstract

The objective of this paper is to examine the application of generative artificial intelligence in personalized portfolio construction and evaluate its performance relative to traditional benchmarks. A generative AI model, specifically OpenAI's GPT-4o, was employed to construct investment portfolios for ten virtual investor profiles over a fixed three-month investment horizon. The methodology involved prompting the model to create portfolio allocations, followed by performance evaluation using financial metrics including total return, volatility, beta, Sharpe ratio, and maximum drawdown. All AI-generated portfolios outperformed the S&P 500 index over the investment period, demonstrating stronger risk-adjusted returns and lower drawdowns. These results highlight the potential of large language models to synthesize financial data and produce competitive investment strategies. The study contributes to the growing body of research on AI-driven decision-making in finance and provides a foundation for the development of generative models tailored to asset and wealth management

Keywords: Generative AI, wealth management, portfolio optimization, personalization

JEL Classification: G11 Portfolio Choice; Investment Decisions

1. Introduction

Wealth management oversees the strategic allocation of capital through portfolio construction on behalf of investors and has been a critical driver of the financial industry, due to its indispensable role in shaping economic growth, promoting financial stability and driving innovation through investment vehicles [14]. Wealth management's significance can be further highlighted by the trillions of dollars of managed capital. The overall finance industry is witnessing an unprecedented technological arms race to adapt, upgrade, and implement emerging technologies to gain a competitive edge, enhance operational efficiency, and meet the evolving needs of their clients [12]. This trend signifies the onset of a technological revolution that is reshaping the industry's competitive dynamics, driving the adoption of breakthroughs like artificial intelligence, which is setting a new trajectory towards the replacement of traditional mechanisms, by making financial institutions leaner,

¹Management Science Laboratory, Athens University of Economics and Business, Greece, dim.papakyriakopoulos@gmail.com; corresponding author

²Management Science Laboratory, Athens University of Economics and Business, Greece, kmn@aub.gr

more agile and digitalized [8]. AI's potential is particularly evident in the realm of wealth management, which is traditionally characterized by human intervention, meticulous analysis, and high-touch interactions. Portfolio construction is now at the precipice of a transformation, leveraging AI's capabilities to offer a more personalized, dynamic, and innovative approach to offer tailor-made solutions and client experience.

Portfolio theory, ever since it was introduced by Harry Markowitz in 1952, has continued to remain a cornerstone in the financial investment domain and portfolio management [13]. The theory provides investors with a mathematical framework for assembling a portfolio of assets in a way that maximizes the expected return for any given level of risk [15]. The modern portfolio theory not only managed to connect the risk and the return of an asset in a quantitative function, but it further paved the way for the introduction of the efficient frontier, which highlights the most efficient portfolios an investor could possibly invest at various risk levels [11]. Despite its foundational role, the modern portfolio theory struggles to capture the intricate, non-linear interrelationships within the financial markets fully. These markets are marked by their complexity and inherent volatility, which marks as a necessity the constant adaptation and improvements to existing models. While the traditional models and methods have been critical in shaping investment strategies in the past, their ability to generate optimal portfolios under the diverse market conditions of the 21st century remains limited.

Artificial intelligence, due to its vast computational power, offers the potential to minimize errors that occur from the traditional models' inability to include multi-domain external factors that influence the financial markets and the assets' performances. AI's ability to learn from large amounts of data uncovers hidden patterns and enhances its performance over time, making it a promising tool for portfolio construction [9]. AI's potential in finance is increasingly being recognized, with advancements in machine learning and predictive analytics fundamentally altering how financial professionals manage, operate, and interact with financial systems [1]. The role of AI is not just limited to automating specific data-heavy or time-consuming tasks, but it has started to evolve into decision-making, reshaping strategies, and influencing outcomes.

In October of 2022, the launch of OpenAI's ChatGPT raised dramatically the interest around "Generative AI", which is an innovative branch of artificial intelligence that is capable of producing and creating data and outputs in various formats, such as text, images, audio and 3D models, which are highly realistic and resemble human-like content and originality [10]. GenAI's ability to generate unique, original and novel data - instead of just understanding and re-creating pre-existing datasets - is a key driver that led make generative AI models stand out from other machine learning and deep learning algorithms [7]. Generative AI utilizes generative models, such as generative adversarial networks, variational auto-encoders and generative pertained transformers to create original data with similar statistical properties and attributes with their respective training data set [10]. These groundbreaking models leverage concepts that have been around for a long time, but their efficiency and potential have reignited interest and sparked curiosity in researchers, data scientists, and the broader public alike [5]. Large language models, which serve as the foundation for generative pertained transformers, have been in use for more than 50 years [4]. The first generation of these models used "n-gram" based systems to estimate the probability of a word given the previous words [5]. However, limitations arise when the computational complexity increased dramatically with higher n-values. This obstacle was

overcome with the introduction of neural networks and the advances in computational power by machines, which made it possible to calculate probabilities for longer n-grams and set the foundations for the creation of generative pertained transformers models [5]. A generative model, is trained with the purpose of understanding the joint probability distribution of the $P(x, y)$ function of the inputs x and outputs y in a training dataset, in contrast to common machine learning and discriminative models that are trained based on the conditional probability distribution $P(y | x)$, which is the probability of outcome y given x as input [10]. Generative AI -and Large Language Models (LLMs) in particular- are trained on vast volumes of unlabeled data by extracting and learning patterns from substantial datasets, which requires extensive resources and time [6].

In the financial industry, the potential of generative AI is particularly compelling in the realm of portfolio management [3]. In theory, GenAI has the potential to deeply analyze and generate an unlimited number of diverse and personalized portfolio solutions that are capable of accommodating the risk tolerance, investment objectives, financial conditions, and other preferences of individual investors. The central aim of this paper is to explore how this theoretical potential of Generative AI can be actualized in the real-world of financial markets through the lens of wealth management, and the increasing consumer trend for personalization. Our methodology presents a pilot experiment evaluating the use of OpenAI's GPT-4o model in personalized portfolio construction. Ten synthetic investor profiles were generated using Python to reflect diverse financial and demographic characteristics. Portfolios were created by prompting the model with each profile and restricting the investment universe to S&P 500 stocks and cash. Historical price data from October 2023 to January 15, 2025, was used to ensure no overlap with the model's training. Performance was assessed using return, volatility, beta, Sharpe ratio, and maximum drawdown, with all AI-generated portfolios outperforming the S&P 500. These results demonstrate the potential of generative AI to support efficient and adaptive wealth management.

2. Proposed Methodology

The purpose of this paper is to explore how generative artificial intelligence can be used to create personalized investment portfolios tailored to different types of investors. Python was the primary tool used throughout the process, with the Pandas library supporting data processing and analysis. Visualizations and performance charts were created using Matplotlib and Seaborn to better illustrate the results. Historical financial data for all S&P 500 companies was sourced through the Bloomberg Terminal to ensure high-quality and reliable inputs. Portfolio construction was performed by sending tailored prompts for each investor profile to OpenAI's API, leveraging the GPT-4o model to generate portfolio allocations. The results were evaluated over the specific time period between January 15, 2025, and April 15, 2025, with the use of key performance metrics such as return, volatility, beta, Sharpe ratio, and maximum drawdown. We aim to demonstrate the practical value of generative AI in asset management by showing how it can adapt to individual investor needs and generate data-driven portfolio strategies.

To generate realistic and diverse investor profiles, the Faker library was used, producing ten distinct individuals with varying backgrounds, financial goals, and risk tolerances. The profiles spanned a broad spectrum of demographics and psychographics. The profiles

ranged from a 20-year-old musician, with a very low risk tolerance, to a 62-year-old interpreter, with a high-risk tolerance. The dummy investors also varied in their financial knowledge, with some having extensive investment experience, while others had no investment experience. The profiles also covered a wide range of occupations, from a civil engineer to a journalist, further diversifying the group's background. Economic factors, such as income, were also varied, with the yearly income spanning from \$32,371 to \$178,896. Debt, marital status, and level of education were other factors considered to create a comprehensive, nuanced picture of each investor. This diversity enabled the creation of a robust testing ground to examine the versatility and adaptability of generative AI in constructing investment portfolios based on people with different characteristics and financial needs and objectives. In Tables 1, 2 and 3, the results from the different investor profiles generated:

Investor ID	Age	Marital Status	Children
1	23	Single	No
2	62	Single	Yes
3	44	Single	Yes
4	29	Single	Yes
5	20	Married	Yes
6	49	Married	No
7	21	Single	Yes
8	34	Married	Yes
9	22	Married	Yes
10	21	Single	No

Table 1: Personal details of investors, including demographics and family information

Investor ID	Occupation	Education	Income	Debt
1	Chemist	High School	32,371	7,463
2	Interpreter	PhD	174,374	116,260
3	Veterinarian	Bachelor's Degree	169,571	43,817
4	Artist	PhD	108,375	30,729
5	Musician	Master's Degree	145,123	77,519
6	Chef	Master's Degree	113,492	33,773
7	Civil Engineer	Master's Degree	65,792	28,907
8	Journalist	High School	178,896	15,407
9	Dentist	PhD	136,034	50,551
10	Musician	PhD	120,639	52,428

Table 2: Educational background, occupation, and financial metrics for each investor

Investor ID	Risk Tolerance	Financial Knowledge	Investment Experience
1	High	Medium	None
2	Low	Low	Extensive
3	Low	Low	Extensive
4	Very High	Low	Some
5	Very Low	Low	Some
6	Moderate to High	Medium	Some
7	Moderate	High	Extensive
8	Very High	Medium	Extensive
9	Low	Low	Extensive
10	Moderate to Low	Low	Extensive

Table 3: Investors' risk profiles and experience with financial instruments

To guide the generative AI model in producing meaningful and realistic investment portfolios, a detailed set of constraints was defined and embedded within each prompt. The objective was to construct one personalized portfolio per investor, aligned with their unique financial profile and risk tolerance. Each portfolio was composed of a combination of cash and equities drawn exclusively from the S&P 500 index. To ensure sufficient diversification without overwhelming complexity, the number of stocks in each portfolio was constrained to fall between 8 and 15. Additionally, to promote broad exposure to the overall market, the total set of portfolios was required to utilize at least 100 distinct S&P 500 companies across the ten investors. The allocation for each stock was expressed in percentage terms, and the sum of all stock weights in each portfolio had to equal 100%.

A key variable in the construction was the investor's risk tolerance, which directly influenced the cash allocation. Cash weights were set to range from 15% to 80%, with more risk-averse investors receiving higher allocations to cash and fewer equity holdings. This allowed the model to adjust the aggressiveness of each portfolio according to individual investor preferences. To preserve output structure and enable automated parsing, the response format was specified in JSON, listing only the stock tickers with their respective weights (e.g., "AAPL - 10%"), followed by 2–3 concise comments explaining the investor's profile and associated risk approach. The final personalized portfolios and the comments are presented in Table 4 and 5 respectively:

Investor ID	Risk Profile	Cash Allocation (%)	Stocks
-------------	-----------------	---------------------------	--------

1	High	21	ADI - 12.31%; NET - 9.87%; NFLX - 10.43%; AIG - 3.49%; ZTS - 4.27%; COST - 5.99%; CRWD - 10.58%; SBUX - 10.07%; PSX - 11.99%
2	Low	63	MDT - 1.49%; PG - 2.66%; BA - 2.57%; SYK - 2.51%; USB - 2.18%; F - 1.24%; DE - 3.0%; PYPL - 1.59%; TSLA - 4.96%; PEP - 1.39%; LIN - 2.3%; JNJ - 3.72%; WBA - 4.8%; V - 2.59%
3	Low	63	ROKU - 2.45%; APD - 4.11%; BDX - 2.86%; SYK - 2.88%; KO - 2.14%; REGN - 2.4%; FDX - 5.51%; PLD - 4.17%; LYFT - 3.99%; ETN - 1.96%; NOW - 4.53%
4	Very High	15	LLY - 12.25%; NOW - 5.72%; HD - 3.37%; CI - 6.75%; ADP - 6.18%; AMAT - 5.51%; MDT - 14.24%; DE - 13.45%; AXP - 6.75%; KHC - 10.78%
5	Very Low	78	PLD - 1.01%; BAC - 2.68%; CB - 1.63%; TMO - 0.73%; COST - 0.64%; MMC - 0.77%; SHOP - 1.88%; LOW - 2.24%; LIN - 1.44%; ZM - 0.67%; INTU - 1.35%; UNH - 2.67%; AAPL - 1.67%; ABNB - 2.62%
6	Moderate to High	39	C - 6.02%; DOCU - 5.55%; AMD - 4.17%; LIN - 2.59%; ORCL - 9.93%; NET - 9.91%; KHC - 12.48%; PLD - 10.35%
7	Moderate	43	GILD - 3.02%; UBER - 8.85%; CI - 8.62%; ADBE - 8.59%; AMGN - 6.24%; PYPL - 5.43%; LIN - 2.66%; ADBE - 4.5%; WBA - 9.09%
8	Very High	15	DIS - 3.7%; CAT - 7.08%; MDT - 7.81%; PLD - 5.03%; AMGN - 10.96%; ABNB - 6.58%; XOM - 4.51%; APD - 7.7%; F - 9.57%; TXN - 4.41%; KHC - 5.48%; MA - 12.17%
9	Low	63	JPM - 1.13%; TMO - 4.19%; MA - 1.35%; IBM - 1.39%; MMC - 1.23%; ZTS - 2.54%; CRWD - 3.34%; V - 2.31%; ORCL - 3.22%; BA - 3.45%; AMD - 4.26%; KO - 4.96%; MDT - 3.63%
10	Moderate to Low	58	AIG - 6.21%; COST - 1.69%; GOOGL - 5.04%; BA - 5.0%; GOOGL - 1.73%; CB - 5.46%; V - 5.85%; COF - 2.0%; TSLA - 2.46%; AMZN - 4.38%; AVGO - 2.18%

Table 4: Personalized Portfolio Asset Allocation with Weights

Investor ID	Risk Profile	Cash	GPT Model Comment
-------------	--------------	------	-------------------

Allocation			
(%)			
1	High	21	Investor has a high risk tolerance. Allocated 21% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
2	Low	63	Investor has a low risk tolerance. Allocated 63% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
3	Low	63	Investor has a low risk tolerance. Allocated 63% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
4	Very High	15	Investor has a very high risk tolerance. Allocated 15% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
5	Very Low	78	Investor has a very low risk tolerance. Allocated 78% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
6	Moderate to High	39	Investor has a moderate to high risk tolerance. Allocated 39% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
7	Moderate	43	Investor has a moderate risk tolerance. Allocated 43% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
8	Very High	15	Investor has a very high risk tolerance. Allocated 15% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.
9	Low	63	Investor has a low risk tolerance. Allocated 63% to cash due to risk preference. Stock weights are

10	Moderate to Low	58	diversified to reflect their risk profile.
			Investor has a moderate to low risk tolerance. Allocated 58% to cash due to risk preference. Stock weights are diversified to reflect their risk profile.

Table 5: GPT Model Commentary for cash allocation

3. Results

The results of this study offer strong evidence of the ability of generative AI to construct personalized portfolios that align with individual investor profiles while demonstrating resilience and competitive performance in real market conditions. Table 6 presents the key performance metrics - return, risk (standard deviation), maximum drawdown, and beta versus the S&P 500 - for each of the ten portfolios generated by the generative pretrained transformer model during the investment window from January 15, 2025, to April 15, 2025. The S&P 500 served as the benchmark for performance comparison.

Investor	Return	Risk (σ)	Max Drawdown	Beta vs S&P 500
Investor 1	-0.02	0.016	-0.16	0.79
Investor 2	-0.01	0.003	-0.023	0.12
Investor 3	-0.02	0.003	-0.0353	0.1426
Investor 4	-0.004	0.011	-0.1014	0.5223
Investor 5	-0.003	0.0009	-0.0094	0.0445
Investor 6	-0.0241	0.0088	-0.0992	0.4305
Investor 7	0.0015	0.0043	-0.0337	0.1654
Investor 8	-0.0350	0.0129	-0.12	0.5898
Investor 9	-0.002	0.0028	-0.0261	0.1373
Investor 10	-0.005	0.0033	-0.0255	0.1633
S&P 500	-0.09	0.0192	-0.189	1

Table 6: Performance Metrics of Generated Portfolios and the S&P 500

Despite the S&P 500 recording a significant negative return of -9% over the three-month investment period, all ten generative AI-created portfolios demonstrated notably stronger performance, showcasing as a result the model’s capacity to construct resilient and context-aware strategies. Returns for the AI portfolios ranged from -3.50% to +0.15%, with every single portfolio outperforming the benchmark by a meaningful margin. Investor 7’s portfolio stood out as the only one to close with a positive return (+0.15%), despite the broader market downturn. Meanwhile, portfolios belonging to Investors 2, 4, 5, 7, 9, and 10 remained close to breaking even, reflecting the model’s ability to protect capital in challenging market conditions and effectively tailor asset allocations to individual investor

risk profiles. Even the most underperforming portfolio (Investor 8) managed to outperform the S&P 500 by over 5 percentage points, illustrating the model's robustness and adaptability in designing investment strategies that balance exposure and caution. This outcome is particularly noteworthy given that no forward-looking data or financial statements were fed into the model; only static investor profiles and historical price data were used to guide the allocation.

Risk metrics further highlight the defensiveness and discipline of the AI-generated portfolios. The standard deviation of portfolio returns (used here as a proxy for volatility) was significantly lower than the S&P 500 benchmark (0.0192) across all portfolios. The lowest volatility was observed in Investor 5's portfolio (0.0009), a result that aligns perfectly with the investor's low-risk tolerance and higher cash allocation. Such consistency between intended risk preferences and realized portfolio behavior demonstrates the model's capacity to interpret and execute tailored strategies effectively. Most other portfolios also maintained low volatility, with even the highest level (Investor 1 at 0.016) still coming in below the market average. In addition, Maximum drawdown, an important measure of downside protection, tells a similarly positive story. While the S&P 500 experienced a drawdown of -18.9% during the investment period, none of the AI portfolios came close to this level of loss. The smallest drawdown was once again recorded by Investor 5 at just -0.94%, while even the highest drawdown (Investor 1 at -16%) still provided better downside protection than the benchmark. Furthermore, beta values across all portfolios were substantially lower than the market (1.0), with the majority falling below 0.2. These figures reflect a consistently lower sensitivity to market fluctuations, particularly for the more conservative investor profiles. Investors 2, 3, and 5, whose risk aversion was high by design, exhibited beta values of 0.12, 0.14, and 0.04, respectively. Such insulation from systemic market risk was primarily achieved through larger cash allocations and restrained exposure to high-volatility equities. All of the above results provide a clear indication of the AI model's ability to synthesize profile information into coherent and risk-aligned investment strategies.

A visualization of the normalized returns for all portfolios relative to the S&P 500 over the investment period is presented in Figure 1. The graph reveals a stark contrast between the sharp declines experienced by the benchmark index and the smoother, more stable trajectories of the AI-generated portfolios. While the market endured periods of heightened volatility and pronounced drawdowns, the portfolios produced by the model exhibited a strong degree of capital preservation, with many displaying gradual, stable performance curves. In some cases, particularly among the moderately risk-tolerant profiles, a slow but steady positive drift was evident—an indication that the model was not only reducing risk but also identifying profitable opportunities within its constrained investment universe.




Figure 30: Normalized AI-Generated Portfolio Returns vs S&P 500

4. Discussing the Challenges and Limitations

While the AI-generated portfolios consistently outperformed the benchmark over the investment period, the methodology employed is subject to several notable limitations and operational challenges. Foremost among these is the token constraint inherent in generative pretrained transformer models, which imposes a ceiling on the amount of information that can be processed in a single prompt. This technical limitation precluded the integration of detailed historical financial data, macroeconomic metrics and predictive modeling, thereby preventing the execution of a dynamic, day-to-day simulation that could have enabled real-time portfolio rebalancing, adaptive allocation strategies, and continuous price monitoring. Consequently, the model's ability to respond to evolving market conditions was inherently restricted. In addition, the same constraint hindered the incorporation of unstructured data, such as financial news, analyst commentary, and social sentiment, which are critical inputs in contemporary investment decision-making. The inability to perform natural language-based sentiment analysis or interpret current events surrounding individual securities diminished the contextual richness of the model's outputs and restricted its scope to a purely static investment framework.

Another critical limitation originates from the opaque, "black-box" nature of the model's architecture. The rationale behind the selection of specific securities and their respective

weightings remains undisclosed, as the model does not provide interpretable outputs or accompanying justification for its decisions. This lack of explainability poses significant challenges, particularly in the context of wealth management, where transparency, accountability, and traceability are foundational; not only for building and maintaining client trust but also for meeting evolving regulatory and compliance standards. Moreover, the stochastic nature of large language models like generative pretrained transformers means that results may vary across iterations, even when the same inputs and constraints are applied [2]. This intrinsic non-determinism introduces an additional layer of uncertainty, as portfolio structures and performance outcomes may diverge significantly from one generation to the next. As such, practitioners would be required to conduct multiple iterations and statistical aggregation to detect consistent patterns or gain actionable insights; an approach that adds computational complexity and raises concerns about reliability and repeatability.

Finally, the scope of this study was deliberately limited to equities listed on the S&P 500, in order to manage data complexity and maintain consistency across the experiment. While this constraint enabled a focused analysis, it also restricted the model's exposure to the broader financial universe, including international equities, fixed income instruments, commodities, and alternative investments. As a result, the study does not fully explore the model's potential to navigate multi-asset portfolio construction or handle heterogeneous financial instruments, an area warranting further investigation in future research.

5. Conclusion

Portfolio management has long stood as a cornerstone of financial research, evolving significantly over the past century. The pioneering work of Harry Markowitz in the 1950s laid the groundwork for what is now known as Modern Portfolio Theory (MPT), introducing a mathematical and statistical framework for optimal asset allocation. MPT marked a paradigm shift by formalizing the trade-off between risk and return, enabling investment professionals to construct efficient portfolios through diversification and quantitative analysis. As both technological capability and investor expectations have advanced, so too has the discipline of portfolio theory. Recent decades have witnessed the emergence of more personalized, dynamic, and risk-sensitive frameworks that account for a broader array of factors, including behavioral characteristics, time horizons, and individual financial goals. These developments reflect a broader movement within the industry toward customization and adaptability in wealth management. At the forefront of this transformation is artificial intelligence (AI), which has demonstrated immense utility across various domains of finance from algorithmic trading and fraud detection to credit scoring and risk assessment. AI's ability to process vast datasets in real time, identify complex patterns, and generate predictive insights offers substantial value to institutions seeking to gain a competitive edge. Among the most recent and disruptive innovations in this field is the emergence of generative AI models, most notably following the release of ChatGPT in late 2022. Unlike traditional AI systems, generative models exhibit the unique capability to produce human-like content and simulate reasoning across diverse tasks. These models are not only capable of synthesizing vast amounts of structured and unstructured data but can also generate nuanced responses, narratives, and even investment

strategies tailored to specific prompts. Their application to the finance industry holds significant promise—particularly in the domain of wealth management, where personalization and responsiveness are increasingly critical. This paper explores the application of generative AI in personalized portfolio construction. Utilizing OpenAI’s GPT-4o model, the study generated investment portfolios for ten synthetic investor profiles over a three-month period, from January 15, 2025, to April 15, 2025. Each portfolio was tailored to reflect the investor’s demographic, economic, and risk-related characteristics. Despite current limitations, such as token constraints and the inherent variability of generative outputs, the results were encouraging. All AI-generated portfolios outperformed the S&P 500 benchmark during the evaluation window and exhibited lower levels of volatility and drawdown. The findings underscore the potential for generative AI to revolutionize portfolio management by offering scalable, individualized investment strategies that replicate the work of countless human analysts. Financial institutions could leverage these models to enhance productivity, reduce operational costs, and deliver sophisticated, real-time portfolio services to clients of all types. As the AI landscape continues to evolve, it is reasonable to expect that existing technical limitations will diminish. Future iterations of generative models may be capable of ingesting and analyzing large-scale financial datasets, parsing news and sentiment data, and continuously rebalancing portfolios in response to market dynamics. The implications for the financial sector are profound: institutions that fail to adapt may risk obsolescence, while those that invest in AI capabilities stand to benefit from enhanced decision-making, client engagement, and market competitiveness. As the generative AI arms race accelerates, the financial industry stands at a pivotal moment; one where embracing technological change is not just advantageous, but essential.

References

- [1] Lee, C., Wu, J., & Gao, L. (2024). A comprehensive review of generative AI in finance. *Journal of Risk and Financial Management*, 3(3), 25. <https://doi.org/10.3390/jrfm3030025>
- [2] Desai, R., Mallya, M., Luqman, A., Ravi, V., Kota, S. M., & Yadav, D. (2024). Generative-AI in finance: Opportunities and challenges. *arXiv preprint arXiv:2410.15653*. <https://arxiv.org/abs/2410.15653>
- [3] (2025). Generative AI for finance: Applications, case studies, and challenges. *Expert Systems*, 42(1), e70018. <https://doi.org/10.1111/exsy.70018>
- [4] (2025). Generative AI solutions to empower financial firms. *Journal of Information Technology*. <https://doi.org/10.1177/07439156241311300>
- [5] Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI*. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

- [6] Frantar, E., Ashkboos, S., Hoefler, T., & Alistarh, D. (2022). GPTQ: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*. <https://arxiv.org/abs/2210.17323>
- [7] Hacker, P., Engel, A., & Mauer, M. (2023, June). Regulating ChatGPT and other large generative AI models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1112–1123). <https://doi.org/10.1145/3593013.3594084>
- [8] Jang, J., & Seong, N. (2023). Deep reinforcement learning for stock portfolio optimization by connecting with modern portfolio theory. *Expert Systems with Applications*, 218, 119556. <https://doi.org/10.1016/j.eswa.2023.119556>
- [9] Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3). <https://doi.org/10.25300/MISQ/2021/16274>
- [10] Gozalo-Brizuela, R., & Garrido-Merchan, E. C. (2023). ChatGPT is not all you need. A state of the art review of large generative AI models. *arXiv preprint arXiv:2301.04655*. <https://arxiv.org/abs/2301.04655>
- [11] Feldman, D., & Reisman, H. (2003). Simple construction of the efficient frontier. *European Financial Management*, 9(2), 251–259. <https://doi.org/10.1111/1468-036X.00226>
- [12] Hua, X., Huang, Y., & Zheng, Y. (2019). Current practices, new insights, and emerging trends of financial technologies. *Industrial Management & Data Systems*, 119(7), 1401–1410. <https://doi.org/10.1108/IMDS-08-2019-0431>
- [13] Elton, E. J., & Gruber, M. J. (1997). Modern portfolio theory, 1950 to date. *Journal of Banking & Finance*, 21(11–12), 1743–1759. [https://doi.org/10.1016/S0378-4266\(97\)00048-4](https://doi.org/10.1016/S0378-4266(97)00048-4)
- [14] Surtee, T. G. H., & Alagidede, I. P. (2023). A novel approach to using modern portfolio theory. *Borsa Istanbul Review*, 23(3), 527–540. <https://doi.org/10.1016/j.bir.2022.12.005>
- [15] Leković, M. M. (2021). Historical development of portfolio theory. *Tehnika*, 76(2), 220–227. <https://doi.org/10.5937/tehnika2102220L>

Bibliography

Lee, C., Wu, J., & Gao, L. (2024). A comprehensive review of generative AI in finance. *Journal of Risk and Financial Management*, 3(3), 25. <https://doi.org/10.3390/jrfm3030025>

- Desai, R., Mallya, M., Luqman, A., Ravi, V., Kota, S. M., & Yadav, D. (2024). Generative-AI in finance: Opportunities and challenges. *arXiv preprint arXiv:2410.15653*. <https://arxiv.org/abs/2410.15653>
- (2025). Generative AI for finance: Applications, case studies, and challenges. *Expert Systems*, 42(1), e70018. <https://doi.org/10.1111/exsy.70018>
- (2025). Generative AI solutions to empower financial firms. *Journal of Information Technology*. <https://doi.org/10.1177/07439156241311300>
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI*. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
- Frantar, E., Ashkboos, S., Hoefler, T., & Alistarh, D. (2022). GPTQ: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*. <https://arxiv.org/abs/2210.17323>
- Hacker, P., Engel, A., & Mauer, M. (2023, June). Regulating ChatGPT and other large generative AI models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1112–1123). <https://doi.org/10.1145/3593013.3594084>
- Jang, J., & Seong, N. (2023). Deep reinforcement learning for stock portfolio optimization by connecting with modern portfolio theory. *Expert Systems with Applications*, 218, 119556. <https://doi.org/10.1016/j.eswa.2023.119556>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3). <https://doi.org/10.25300/MISQ/2021/16274>
- Gozalo-Brizuela, R., & Garrido-Merchan, E. C. (2023). ChatGPT is not all you need. A state of the art review of large generative AI models. *arXiv preprint arXiv:2301.04655*. <https://arxiv.org/abs/2301.04655>
- Feldman, D., & Reisman, H. (2003). Simple construction of the efficient frontier. *European Financial Management*, 9(2), 251–259. <https://doi.org/10.1111/1468-036X.00226>
- Hua, X., Huang, Y., & Zheng, Y. (2019). Current practices, new insights, and emerging trends of financial technologies. *Industrial Management & Data Systems*, 119(7), 1401–1410. <https://doi.org/10.1108/IMDS-08-2019-0431>
- Elton, E. J., & Gruber, M. J. (1997). Modern portfolio theory, 1950 to date. *Journal of Banking & Finance*, 21(11–12), 1743–1759. [https://doi.org/10.1016/S0378-4266\(97\)00048-4](https://doi.org/10.1016/S0378-4266(97)00048-4)

Surtee, T. G. H., & Alagidede, I. P. (2023). A novel approach to using modern portfolio theory. *Borsa Istanbul Review*, 23(3), 527–540. <https://doi.org/10.1016/j.bir.2022.12.005>

Leković, M. M. (2021). Historical development of portfolio theory. *Tehnika*, 76(2), 220–227. <https://doi.org/10.5937/tehnika2102220L>