

Artificial Intelligence in Investment and Retirement: Demystification, Applications, and Risks

JULY | 2025



Aging and
Retirement



SPONSORS

SOCIETY OF ACTUARIES
INVESTMENT SECTION

SOCIETY OF ACTUARIES
AGING AND RETIREMENT
STRATEGIC RESEARCH
PROGRAM




Artificial Intelligence in Investment and Retirement Report

Demystification, Applications, and Risks

AUTHOR Kailan Shang, FSA, MAAA, ACIA, CFA, PRM

SPONSORS Investment Section


Aging and Retirement Strategic
Research Program



Give us your feedback!

Take a short survey on this report.

[Click Here](#)



Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute, the Society of Actuaries, or its members. The Society of Actuaries Research Institute makes no representation or warranty to the accuracy of the information.

Copyright © 2025 by the Society of Actuaries Research Institute. All rights reserved.

CONTENTS

Executive Summary	4
Section 1: Introduction	7
Section 2: AI Applications	9
2.1 AI Applications in Investment	9
2.2 AI Applications in Retirement and Pension	11
Section 3: AI Demystification	14
3.1 Nontraditional Data	15
3.2 Model	18
3.3 Process	20
Section 4: Risk Management of AI Application	25
4.1 AI Bias	25
4.2 Cyber Security	29
4.3 Data Privacy and Governance	30
4.4 AI and the Risk of Market Volatility	31
4.5 AI Transparency.....	32
4.6 AI Ethical Challenges	33
4.7 Legal Risk	34
4.8 AI Model Risk.....	35
Section 5: AI Illustration.....	42
5.1 Illustration of AI Model Improvement – Retirement Savings Behavior Coaching.....	42
5.2 Investment Chatbot Using Retrieval-Augmented Generation	48
5.3 Social Media Data Summarization	53
5.4 AI Retirement Planning Tool	60
5.5 AI Pension Investment Analysis	70
5.6 Sustainability Report by Generative AI	81
5.7 Enhancing Financial Analysis with Generative AI.....	101
5.8 AI Dynamic Asset Allocation for DB Plans.....	118
Section 6: Conclusion	124
Section 7: Acknowledgments	125
Appendix A: Glossary of Terms	126
Appendix B: AI Models – A More Technical Introduction.....	128
Appendix C: AI Generated Codes for Portfolio Optimization User Interface.....	137
Appendix D: SOA Publications on AI for Actuaries.....	142
Endnotes	143
References.....	146
About The Society of Actuaries Research Institute	151

Artificial Intelligence in Investment and Retirement

Demystification, Applications, and Risks

Executive Summary

Artificial intelligence (AI) is increasingly shaping the fields of investment and retirement planning by enhancing data-driven decision-making, improving efficiency, and automating complex processes. With advancements in reinforcement learning and large language models, AI applications are expanding in portfolio management, sentiment analysis, risk assessment, and personalized financial planning. However, as AI adoption grows, it is important to consider both its benefits and challenges to support responsible and effective implementation. This report provides an overview of AI applications in the investment and retirement sectors, offering insights for professionals exploring ways to leverage AI while managing associated risks.

Insights and Benefits of This Report

This report serves as a resource for financial professionals and technology experts interested in AI's role in investment and retirement planning. It is intended to provide an understanding of the following:

- How AI may be applied in investment and retirement planning, including portfolio optimization, sentiment analysis, risk management, and personalized financial guidance.
- The fundamental principles of AI technologies such as large language models and reinforcement learning, with a focus on their relevance to financial decision-making.
- The risks that may arise from AI applications, including biases, cybersecurity threats, ethical considerations, and regulatory compliance, along with potential strategies for mitigating them.
- Real-world examples illustrating AI-driven solutions in financial services, such as reinforcement learning for retirement savings, chatbot development, and AI-assisted portfolio management.
- Considerations for integrating AI while promoting transparency, fairness, and accountability.

The report is structured to accommodate readers with different interests and varying levels of technical expertise. Those seeking a broad understanding of AI's role in investment and retirement planning may focus on [Section 2](#), which discusses AI applications in financial decision-making, risk management, customer services, and portfolio optimization. Readers concerned about the risks of AI implementation may refer to [Section 4](#), which outlines AI risk management and mitigation strategies.

For those with a more technical background, [Section 3](#) provides an overview of AI methodologies, including data processing, model training, and reinforcement learning, and involves some technical discussions. [Section 5](#) is also intended for technical readers, as it explores AI-driven solutions with practical case studies, including reinforcement learning for retirement savings, chatbot development for financial insights, and AI-assisted portfolio management. In addition, most of the appendices are designed strictly for technical readers, offering in-depth discussions on advanced AI methodologies and implementation details.

AI in Investment and Retirement Planning

The application of AI in investment and retirement planning continues to evolve. In the investment sector, AI is being explored for enhancing data analysis, automating trading strategies, and supporting financial advisors with personalized recommendations. Large language models and natural language processing (NLP) techniques may assist in sentiment analysis, allowing firms to extract insights from social media, earnings calls, and financial reports. AI-driven portfolio optimization techniques are also being examined as a way to help asset managers balance risk and return while incorporating real-time market data. Reinforcement learning models may provide additional insights for dynamic asset allocation, potentially helping investors and pension plan sponsors adapt to market fluctuations more effectively.

In retirement planning, AI is increasingly used to support financial decision-making for individuals and institutions. Personalized financial planning tools that incorporate AI could help analyze spending habits, forecast retirement readiness, and optimize savings strategies. Predictive models may offer insights into life expectancy and healthcare costs, potentially aiding retirees in making informed financial decisions. Robo-advisors powered by AI are being developed to provide automated retirement planning solutions tailored to individual risk profiles and savings goals.

The use and impact of AI in investment and retirement planning varies across different types of investors. Active individual investors may directly use AI tools for market analysis and decision-making, while those invested in mutual funds or index funds may experience AI more indirectly through fund management improvements. Passive investors relying on employer-managed plans, and corporate or institutional investors, may benefit from AI-driven strategies at a broader, systemic level. Recognizing these differences is important in understanding AI's role across the investment and retirement landscape.

Understanding AI Technologies

AI's effectiveness in financial services depends on a solid understanding of its underlying technologies. This report provides an accessible explanation of key AI concepts, including reinforcement learning, a technique where AI agents learn from feedback and adjust decisions dynamically, with potential applications in investment strategies. Large language models (LLMs) are discussed in the context of how AI processes and generates text-based financial insights, which may aid in research, reporting, and customer interactions. The report also emphasizes the importance of high-quality, unbiased data for training AI models to support reliable outputs.

By understanding these technologies, professionals may be better positioned to evaluate how AI can be integrated into their investment and retirement planning processes.

This report is intended as a starting point for understanding how AI works and what it can offer. While it introduces foundational concepts, further exploration is essential to evaluate the specific technologies used in particular financial use cases.

Managing Risks and Challenges

While AI presents opportunities, it also introduces several risks that may require careful consideration. This report explores key AI risks, such as bias and fairness, where AI models can inherit biases from training data, which may lead to unintended discrimination in financial decision-making. Possible strategies for mitigating bias, such as diverse data sampling and transparency measures, are discussed. Cybersecurity and data privacy concerns are also addressed, as AI systems often require access to sensitive financial and personal data, raising concerns about security and regulatory compliance. The report outlines best practices that may help protect AI systems from cyber threats and support responsible data use.

The potential impact of AI on market volatility and systemic risks is considered, as AI-driven trading systems may contribute to rapid market fluctuations. The report explores historical examples and potential safeguards that could help prevent unintended financial instability. Ethical and legal considerations are also discussed, as financial institutions may need to navigate evolving regulations and ethical guidelines. Key regulatory developments, such as the EU AI Act and industry-specific compliance requirements, are highlighted.

Understanding these risks and implementing effective governance frameworks may help ensure that AI is applied responsibly in financial services.


Real-World Applications and Case Studies

To provide practical insights, this report presents real-world AI applications in investment and retirement planning. These include reinforcement learning for retirement savings optimization, which explores how AI might assist individuals in developing better savings habits and adjusting contributions based on financial goals. Chatbot development for financial advisory services is discussed in the context of how LLMs may be used to automate client interactions, potentially improving accessibility to investment and retirement guidance. AI-driven sentiment analysis for market insights is examined, demonstrating how NLP could be leveraged to extract investment sentiment from social media and financial reports.

Portfolio optimization using AI models is explored as a means to enhance liability-driven investment strategies for pension plans through AI-assisted risk assessment. The report also considers AI-generated sustainability reports and how generative AI might streamline ESG reporting and compliance efforts. The role of AI in dynamic asset allocation for pension plans is reviewed, along with its potential to refine long-term investment strategies. The applications of AI-assisted data exploration, coding, and analytical support are demonstrated as potential ways to improve productivity in financial research and decision-making.

Some of these examples include Python code snippets and workflows to help users experiment with AI applications in real-world scenarios. These illustrations highlight AI's versatility in supporting complex problem-solving and improving efficiency, while serving as educational tools for professionals interested in practical applications.


As AI continues to advance, its role in investment and retirement planning may expand further. By promoting transparency, fostering collaboration between financial experts and AI practitioners, and refining governance practices, the financial industry may be able to leverage AI's potential while maintaining fair, reliable, and responsible use. This report aims to provide insights that support achieving this balance.



Give us your feedback!

Take a short survey on this report.

[Click Here](#)



Section 1: Introduction

Artificial intelligence (AI) technologies have been increasingly applied in actuarial-related areas. In the investment field, AI has been used to enhance algorithmic trading, sentiment analysis, portfolio optimization, risk management, and customized investment recommendation systems. These applications are documented in industry reports, such as Hayman (2024) by the CFA Institute and Balasubramania (2021) by McKinsey & Company, which highlight the role of AI in improving financial decision-making.

In the retirement field, AI facilitates personalized retirement planning through innovative methods of data collection, analysis, and communication. For instance, some platforms employ AI-powered tools to streamline retirement planning processes. Beyond planning, AI technologies can monitor retirement progress and provide actionable suggestions to stay on track, as demonstrated by systems offering retirement solutions.

With the availability of large data volumes and advanced models, AI has been shown to excel in certain tasks requiring efficiency, accuracy, accessibility, and convenience. Killock (2020) demonstrated that an AI system developed by Google Health outperformed radiologists in detecting breast cancer from mammograms¹. Shang (2021) compared reinforcement learning (RL)-based dynamic strategies to traditional static allocation methods, demonstrating that RL can achieve better risk-return trade-offs while requiring less computational effort. Studies such as Maslej (2024), the Stanford AI Index Report, provide further evidence of these capabilities. The recent advancements in large language models (LLMs), such as OpenAI's GPT models, have drawn significant attention to AI's ability to summarize and generate non-numerical data. LLMs are being utilized for drafting emails, scripting social media posts, writing reports, debugging computer programs, generating videos, and answering context-specific questions, as reported in publications like OpenAI's usage documentation.

While AI technologies continue to evolve with anticipated improvements and expanding real-world applications, LLMs have demonstrated the potential to act as helpful advisers and tools for improving productivity. Their ability to generate solutions, including creative ones, based on minimal input has been observed in applications ranging from content creation to problem-solving.

At the same time, the complexity of AI models and the extensive data used for training often result in users relying on these tools without fully understanding their inner workings. For professionals who are not directly involved in developing these models, this knowledge gap can make it challenging to use AI confidently. To address this, understanding how AI works and employing methods to validate outputs can empower users. Furthermore, incorporating domain knowledge into AI systems has the potential to enhance their performance, as supported by findings in reinforcement learning and fine-tuning methodologies.

Assuming AI will play a much bigger role in the investment and retirement sectors, this research aims to assist actuaries and other professionals in embracing these changes to use AI technology appropriately and improve it with domain-specific expertise.

AI, as a broad concept, encompasses a wide range of applications. AI systems are designed to perform tasks that typically require human intelligence, such as learning, reasoning, and decision-making. Unlike humans, AI systems excel at processing vast amounts of information efficiently, and advanced AI aims to learn from experience without explicit instructions. These principles are well-supported in machine learning literature. This report focuses on technologies like large language models (LLMs) and reinforcement learning, which are designed to learn independently and offer innovative solutions, as documented in AI research from leading academic institutions and technology firms. The report proceeds as follows:

- Section 2 (AI Applications) discusses existing and future applications of AI in the investment and retirement planning field, such as investment data analysis, investment decision-making, risk management, custom services, automation and productivity improvement.
- Section 3 (AI Demystification) introduces AI technologies to investment and retirement professionals. A good understanding of AI technologies is essential to putting AI into practice. The focus is on the data, models, and logic behind the intelligence, especially large language models, rather than technical details.
- Section 4 (AI Risk Management) explains the risks of using AI, including model risk, bias, data privacy, cyber security, result reproducibility, ethical challenges, legal risk, and reputational risk. Mitigation approaches are also suggested to reduce the probability and significance of using AI for investment applications.
- Section 5 (AI Illustration) provides realistic examples of using AI to solve investment and/or retirement related issues, including the following.
 - [Section 5.1](#): Illustration of reinforcement learning with human feedback with an example of retirement savings behavior coaching.
 - [Section 5.2](#): Development of a simple chatbot based on a research paper on GDP growth drivers.
 - [Section 5.3](#): Summarization of investment related nonnumerical data such as social media data that can help estimate investors' sentiments.
 - [Section 5.4](#): Creating a retirement planning tool using AI, including knowledge exploration, analysis, model building, and a user interface, with limited human effort.
 - [Section 5.5](#): Developing portfolio optimization strategy for liability-driven investment.
 - [Section 5.6](#): Generating sustainability reports using generative AI.
 - [Section 5.7](#): Enhancing financial analysis with Generative AI including data exploration, coding and analytical support.
 - [Section 5.8](#): AI enhanced dynamic asset allocation for defined benefit pension plans.

Each example contains a description of the AI application, its underlying data and models, how AI is used to solve the issues, and any deficiency. Technical details are discussed when needed to make the report more accessible to readers without hands-on programming experience in developing AI models. However, Python programs used in some examples are made available for educational purposes (<https://github.com/Society-of-actuaries-research-institute/FP208-Artificial-Intelligence-in-Investment-and-Retirement>)

- Section 6 (Conclusion) summarizes the key points of this research and concludes the main body of the report.
- The appendices provide supplementary information on key AI concepts, tools, models, applications, and resources. Appendix A defines essential AI terms, including AI bias, deep learning, and model robustness. Appendix B offers a technical overview of AI models such as deep learning, reinforcement learning, and transformers. Appendix C includes AI-generated JavaScript code for a portfolio optimization interface, with instructions for implementation. Appendix D lists the SOA publications on AI in a variety of areas.

Section 2: AI Applications

Based on existing literature and applications, this section explores how AI may enhance investment and retirement planning through advanced data analysis, predictive modeling, and automation. It highlights AI's role in sentiment analysis, risk assessment, portfolio optimization, algorithmic trading, mortality modeling, and personalized financial advice. Readers may gain insights into how AI improves decision-making, efficiency, and personalization in investment and retirement planning.

With a good foundation for learning specific knowledge and skills, AI models can be improved to build investment and retirement expertise.

2.1 AI APPLICATIONS IN INVESTMENT

In the investment field, AI can be used to harness nontraditional data, such as text data, to improve **information summarization** and investment analysis. Lin et al. (2023) discussed the applications of natural language processing (NLP) and AI in the investment field today, including social media sentiment analysis, relevant data extraction and summarization, searching for key themes and question answering, identifying risks in corporate filings, analyzing earnings call transcripts, and deepening client insights to prioritize sales efforts. Pisaneschi (2024) used LLMs in an environmental, social, and governance (ESG) case study and applied fine-tuning methods to detect material ESG tweets to generate investment returns. Jansen (2023) mentioned that AI and more specifically NLP applications have become more widespread because the greater amount of data and high-quality open-source software has become available. Many investment firms that rely on textual input use some form of automated text processing or scoring to improve productivity. He discussed two applications: sentiment scoring and ESG scoring. Sansani and Samonov (2023) discussed applying AI to extracting text-based ESG insights using techniques such as word embedding, language models, and domain-specific model retraining.

Fund selection based on investment style may be aided by AI. Sasaki et al. (2018) studied investment styles of fund managers using deep learning. AI can be used to support many other investment and trading activities as well. Lee (2024) reported on the use of AI technology in \$13 billion of quant stock funds by Vanguard. Nagy et al. (2023) discussed how AI models, including deep learning and reinforcement learning, can help inform trading decisions through predicting spreads, trade volume, volatility, and so on. Stanton (2023) illustrated the application of AI and big data to support trading activities, such as dimensionality reduction to select important features, perform analysis of transaction costs, algorithmic trading, and recommending trading strategies. Son (2023) reported that J.P.Morgan had been creating thematic investment baskets using the GPT-4 model. Securities can be analyzed and selected to meet customers' needs with the aid of AI.

RL-based approaches may also offer significant advantages in **portfolio management** and **risk hedging**. Al Abrate et al. (2021) investigated continuous-action RL for portfolio allocation in life insurance. Their work illustrates how RL enhances portfolio optimization by balancing risk and return while considering regulatory constraints. Their results suggest that RL methods can improve investment decisions over traditional optimization techniques. Chong et al. (2023) proposed a pseudo-model-free hedging framework for variable annuities using deep RL. They showed that deep RL-based hedging strategies can effectively manage financial risks, while requiring minimal assumptions about the underlying asset dynamics. Buehler et al. (2019) introduced the concept of deep hedging, demonstrating that RL-based hedging outperforms traditional methods by considering non-linear market dynamics. Further advancements in deep RL-based

hedging are presented by Cao et al. (2021) and Du et al. (2020), who applied RL techniques to option replication and derivative hedging. These studies highlight the flexibility of RL in financial risk management and its ability to adapt to market changes. Lütkebohmert et al. (2022) extended this work by incorporating robust deep-hedging techniques to address model uncertainty, while Carbonneau (2021) focused on long-term derivative hedging using RL methods.

Synthetic data can be generated by AI to facilitate research and development, such as trading strategies. In the article “Synthetic Data for Real Insights” by J.P.Morgan², a variety of synthetic data is generated, including anti-money laundering behaviors, customer journey events, markets execution data, and payments data for fraud detection. Synthetic data may be generated for stochastic modeling purposes, as well, to study the risk of investment strategies.

Customer services may be revolutionized by using AI with much improved efficiency, availability, and personalization. Capponi et al. (2021) examined personalized robo-advising using RL, demonstrating how RL enhances investment decisions through client interaction. Liang (2023) discussed the application of AI to customer service in the finance sector, including investment advisory services that provide investment recommendations, customer queries such as automated transactions and question answering, and outbound services such as sales and marketing.

Productivity is expected to be improved by AI. As discussed in Son (2023), among other financial institutions, Morgan Stanley announced that it had been developing an AI assistant based on OpenAI’s GPT-4 for financial advisors and their support staff. The AI assistant allows fast and efficient access to a database of about 100,000 research reports and documents by asking questions. As reported by Bousquette (2024), Goldman Sachs deployed its first generative AI tool across the firm and the time needed for building generative AI applications has shortened from months to weeks due to AI-powered code generation. There is no doubt that AI can be used to improve productivity in other areas when concerns of risks can be addressed. AI can be used to monitor market trends and regulatory changes, summarize financial reports, automate model development, validation and documentation, generate personalized reports, analyze and improve computer programs, and so on.

Table 1 lists some AI applications by stakeholders in the investment field. Concrete examples are given where possible. This, however, is not an exhaustive list given the pace of adoption and changes in the AI field. It is hoped that it shows at least a portion of the possibilities.

Table 1

AI APPLICATION IN INVESTMENT BY STAKEHOLDER

Stakeholder	AI Application
Individual Investors	Personalized Portfolio Management: AI-based tools for creating and managing personalized investment portfolios.
	Behavioral Insights: AI identifies patterns in investor behavior to provide tailored advice or warnings.
	Robo-Advisors: Automated platforms offering financial advice and portfolio management.
Portfolio Managers	Portfolio Optimization: AI optimizes asset allocation for risk and return goals.
	Risk Analysis: Real-time monitoring of market risks and portfolio exposure.
	Alternative Data Utilization: Incorporating non-traditional data sources like satellite imagery and social media sentiment.

Stakeholder	AI Application
Institutional Investors	Alpha Generation: Identifying investment opportunities using machine learning models.
	Market Predictions: Forecasting asset prices and market movements using AI.
	ESG (Environmental, Social, Governance) Scoring: Assessing investments for sustainability criteria.
Financial Advisors	Client Insights: AI analyzes client data to provide insights into financial habits and future needs.
	Retirement Planning: AI-based tools to forecast retirement readiness based on changing financial scenarios.
	Compliance and Reporting: Automating regulatory compliance and reporting processes.
Hedge Funds and Quants	Quantitative Strategies: Building AI-powered trading algorithms for systematic strategies.
	Sentiment Analysis: Gauging public sentiment on stocks or industries using natural language processing (NLP).
	High-Frequency Trading (HFT): Using AI for split-second trading decisions.
Researchers and Academics	Market Efficiency Studies: Using AI to assess market behavior and anomalies.
	Stress Testing Models: Simulating extreme market events to test portfolio resilience.
	Behavioral Finance Research: Understanding investor psychology through data analysis.

2.2 AI APPLICATIONS IN RETIREMENT AND PENSION

Investment plays an important role in retirement planning and pension management. In addition to the AI applications discussed in the previous section, AI may be used to assist other tasks for retirement and pension.

AI can be used to facilitate **mortality modeling**, especially when customized mortality prediction is important for retirement planning. Nigri et al. (2020) uses a recurrent neural network with a long short-term memory to forecast life expectancy and lifespan disparity, both independently and simultaneously at birth and age 65. Perla et al. (2021) forecasted mortality rates using a simple shallow convolutional network model and found interesting results in terms of applicability.

AI can be used to improve **retirement planning**. Forsyth and Li (2019) used a neural network approach to solve for the optimal asset allocation phase of a defined contribution plan where a target objective function is used. Irlam (2020) used reinforcement learning to derive an investment and consumption strategy intended to optimize lifetime well-being. He showed how using reinforcement learning creates substantial improvements over classic dynamic programming and Monte-Carlo simulations. Das (2020) used reinforcement learning and applied it to solving goals-based wealth management problems. AI-driven tools are being developed to provide personalized retirement planning by analyzing individual financial

situations and goals. Rappaport (2024) explored how AI can assist individuals in retirement planning, emphasizing the potential benefits and challenges from an individual perspective. Casas and Román (2024) examined the effect of AI on early retirement decisions in Europe, finding that AI advancements can reduce the likelihood of early retirement, particularly among workers with higher education. Lo and Ross (2024) discussed the challenges and potential of large language models, like ChatGPT, in offering financial advice, emphasizing the need for domain-specific expertise and regulatory compliance. Cutler (2024) examined the issues related to AI as a partner for retirement professionals, considering the benefits and challenges of integrating AI into retirement planning services.

AI can be applied to **investment strategy optimization** for defined benefit (DB) pension plans. Shang (2021) applied deep reinforcement learning to the dynamic asset allocation problem for pension plan sponsors. The example showed that reinforcement learning can improve the design of investment strategies in terms of improving return and reducing risk.

AI techniques have been applied to **optimize consumption strategies** during retirement. Chen and Langrené (2020) developed a deep neural network approach to solve a lifetime expected utility-based model for optimal consumption in the decumulation phase of a defined contribution pension system, highlighting AI's role in enhancing financial decision-making for retirees. Ozhamaratli and Barucca (2022) introduced a model where AI agents learn optimal portfolio allocation and saving strategies suitable for heterogeneous profiles, demonstrating AI's potential in personalized retirement planning.

As pointed out by Cantor and Shang (2021), AI can be used to predict participants taking a lump sum, **analyzing behaviors** in a defined contribution (DC) plan related to selecting investment options, taking loan withdrawals, making other financial decisions, and personalized communication, such as using chat-bots, to more efficiently address participant questions related to their benefits. Similar to pension services, AI can also be used for individuals to understand their latest financial conditions and retirement goals. Personalized and relevant communication can help improve literacy on retirement planning.

Table 2 lists some AI applications by stakeholders in the retirement and pension field. Examples are given where possible. This, however, is not an exhaustive list given the pace of adoption and changes in the AI field. It is hoped that it shows at least a portion of the possibilities.

Table 2
AI APPLICATION IN RETIREMENT BY STAKEHOLDER

Stakeholder	Potential AI Applications
Individuals	Personalized Retirement Planning: AI-powered tools provide customized savings and investment advice.
	Health cost estimation during retirement for better financial planning
	Behavioral Nudges: Encouraging positive saving habits through personalized notifications.
	Budget Optimization: Tracking expenses and identifying opportunities for retirement savings.
	Retirement Education: AI enhances learning with personalized content, real-time assistance, and interactive tools.
DC Plan Sponsors	Participant Analytics: Understanding participant behavior to optimize plan features like auto-enrollment.

Stakeholder	Potential AI Applications
DB Plan Sponsors	Fund Selection Optimization: Recommending target-date funds and lifecycle funds using predictive models.
	Plan Engagement Tools: Personalized messaging to boost plan enrollment and contributions.
	Liability Matching: Aligning investments with long-term pension liabilities using AI models.
Advisors	Risk Monitoring: Real-time assessment of funding ratios and portfolio risks.
	Predictive Forecasting: Anticipating future liabilities and funding needs.
	Risk Profiling: Assessing client preferences and risk tolerance with AI analysis.
Researchers	Portfolio Optimization: Rebalancing and diversification recommendations based on AI models.
	Compliance Automation: Ensuring adherence to fiduciary standards and regulations.
	Demographic Analysis: Predicting future retirement trends using population and economic data.
	Policy Impact Simulations: Testing effects of changes like contribution limits or retirement age on retirement systems.
	Market Scenario Analysis: Stress-testing retirement strategies under extreme market conditions.

Section 3: AI Demystification

This section provides an accessible breakdown of AI fundamentals, explaining how AI processes vast data, learns patterns, and evolves to outperform humans in some areas. It covers key AI models like deep learning, reinforcement learning, and transformers, along with various data types such as text, social networks, images, GPS, and health data. The section also explores AI model training, retraining, and fine-tuning, particularly in specialized fields like investment and retirement planning. Readers may gain an understanding of AI's capabilities and underlying models. The foundational AI concepts covered by this section apply broadly across domains and are not limited to investment and retirement contexts. A more technical introduction of AI models is available in [Appendix B](#).

AI benefits from two key capabilities: the ability to process vast amounts of data efficiently, analyzing and retrieving relevant information quickly and in parallel, and the capacity to learn, adapt, and generate innovations at scale when guided by an evaluation system. These strengths enable AI to perform certain tasks more effectively than humans and, in some cases, surpass human cognitive abilities by detecting complex patterns, optimizing decisions, and operating at speeds and scales beyond human limitations in data processing and analysis.

For example, large language models (LLMs) are designed to process and generate text in ways that approximate conversational patterns, leveraging extensive training datasets and complex models. Table 3 lists some details of the model size of GPT-4, a LLM.

Table 3
SIZE AND COMPLEXITY OF LLMS

	Detail	GPT-4 ¹
Training Data	Large volume of text data available electronically, including content on the internet such as Wikipedia and electronic books.	10 trillion words ² or 20 billion ³ pages
Model Structure	Deep-learning models that have many layers of neurons to mimic complex relationships.	1.8 trillion model parameters
Model Training	Supervised learning and reinforcement learning from human feedback.	2.15 ²⁵ floating point operations. The same number of arithmetic operations would take a few million years using an ordinary computer.

Notes:

1. GPT-4 (Generative Pre-trained Transformer 4) is an advanced LLM developed by OpenAI, which also developed ChatGPT using GPT-3.5.
2. According to Patel and Wong (2023), 13 trillion tokens are used for training GPT-4. Assuming a token is 0.75 words on average, 10 trillion words were used.
3. Assuming each page contains 500 words, 20 billion pages were used for model training.
25. See Endnotes

Compared to a human being, an LLM model may exhibit similar complexity to a human brain in terms of connections but differs in its ability to access and learn vast amounts of information in parallel. LLMs generalize and remember lexical semantics, syntax, and basic reasoning through training. This section explores the underlying logic of AI, focusing on data, model, and development processes.

3.1 NONTRADITIONAL DATA

In addition to traditional numerical data that are used in statistical modeling and regression analysis, AI applications usually use non-traditional data at a large scale. Some data types also need special processing methods to extract features that can later be used by AI models. A few of them are discussed below.

Text Data

Text data refers to non-numerical information that consists of written text and transcripts, represented by characters. Examples include financial reports, social media data, survey responses, and scientific papers. While text data can be used in AI applications, it must first be transformed into a numerical format. Explanatory variables, such as word count and word frequency, can provide statistical insights into text data. In the case of large language models (LLMs), both individual words and their contexts are required as inputs. Words are converted into numeric vectors before being fed into AI models for training and validation. In modern computer systems, numbers are assigned to represent graphic characters through character encoding. Common examples include ASCII and Unicode. ASCII is a character encoding system using 7-bit codes to represent 128 characters, primarily for the English alphabet and basic symbols. Unicode assigns unique code points to over 143,000 characters across various languages and scripts, ensuring global text representation and compatibility in modern computing. However, these assigned numbers lack semantic and contextual information about the characters and words they represent.

To address this, word embedding is a technique widely used in natural language processing (NLP) and, in particular, LLMs to convert words or even entire documents into numerical vectors. The goal of word embedding is to ensure that words with similar semantic meanings and proximity in the text are represented by vectors closer to each other in a high-dimensional space. A popular method for word embedding is Word2Vec, developed by Google, which leverages artificial neural networks and text data to optimize vector representations. For example, in the continuous bag-of-words (CBOW)³ approach, the model is trained to predict surrounding words based on a given word, capturing the context and semantic relationships effectively. Once text data can be represented by numerical values, they are ready to be used in AI models.

Social Network

Social network data describes the interaction of agents in a social network. An agent represents an individual entity, such as a person, organization, or online account, participating in interactions within the network. Agents are nodes in a graph, with edges representing relationships. Using networks and graph theory, social networks can be analyzed to identify groups and relationships, which can be used in market segmentation and fraud detection and improve the prediction accuracy. Figure 1 shows an example of a social network that describes the relationship among agents. In this simple example, two groups can be identified with the first one including A, B, C and D with full interactions among members, and the second one including E, F and G with only a chain of interactions among members. The relationship between the two groups is in a single direction as well. A directed edge represents a unidirectional relationship or influence between two agents, with the first group interacting with, following, or influencing the second group, but not necessarily the reverse. Such relationships are common in social media like Twitter (follower-following relationships) or citation networks (one paper citing another).

Figure 1
EXAMPLE OF SOCIAL NETWORK DATA

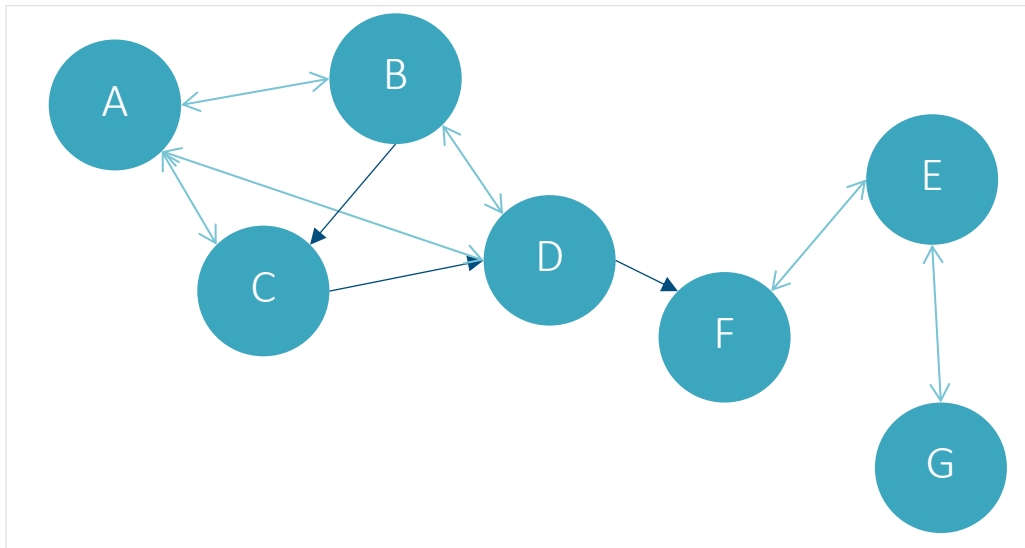
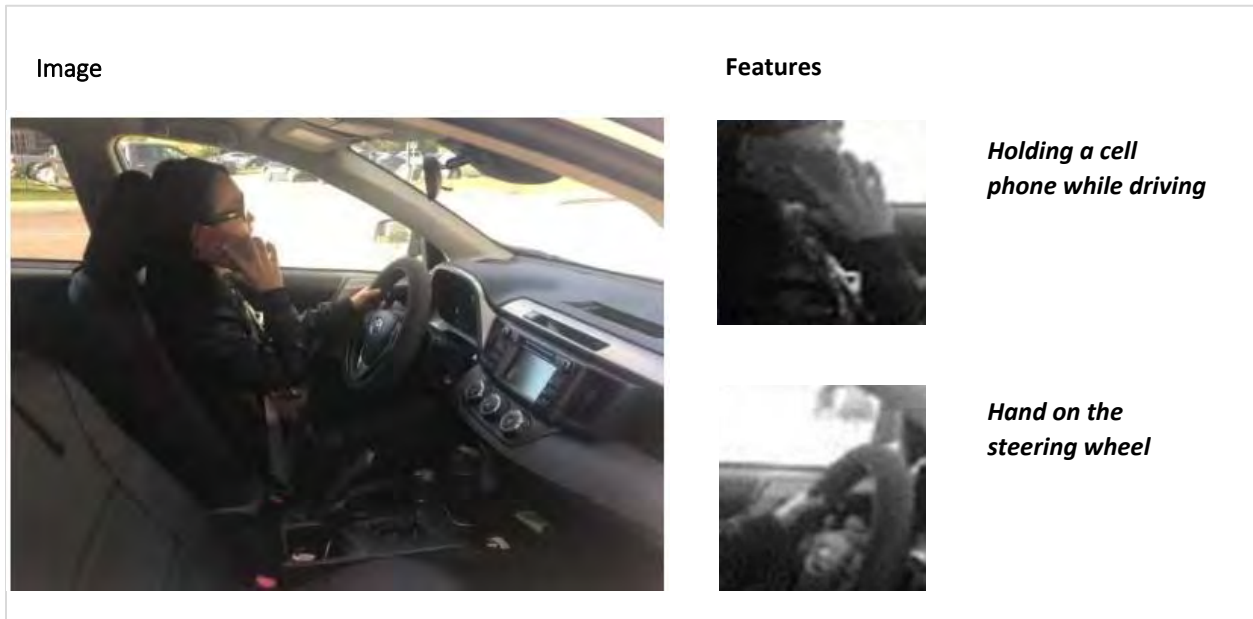


Image and Video Data

Both image and video data are represented as numbers. A video is a series of images with the number of images determined by frame per second (fps), which defines the number of still images in a second of the video. An image is composed of pixels, the smallest picture element in digital imaging. Each pixel can be described by its color, which can be represented by a color imaging system such as the RGB system⁴. Using the RGB system, a pixel can be represented by a three-element vector, including the scales at each color dimension: red, green and blue. White is represented as (0,0,0), green is represented as (0,128,0), and black is represented as (255, 255, 255). If an image has 1920×1080 pixels, it is represented by 1920×1080×3 data points in a computer.

Using image or video data directly in AI models is not practical or effective given the data volume. Features are usually extracted that can help improve the prediction. For example, to assess the risk of distracted driving, images may be scanned to search for features such as hands off the steering wheel, holding cell phones, and driver facing down or backward for too long. The feature extraction can be done using supervised approaches that define the extraction rules, or the automated approaches such as convolutional neural networks (CNNs) that identify features that are mostly relevant to the target variables, with an illustration in Figure 2.

Figure 2
EXAMPLE OF IMAGE DATA FEATURE EXTRACTION



GPS Data

With technological advances, GPS data becomes more available and can be used to improve prediction accuracy where allowed. A popular application is usage-based insurance (UBI) where insurance customers opt to allow insurance companies to collect GPS data on their vehicle to assess driving behaviors and get substantial premium discounts for safe driving and usage.

GPS data can be constructed as the longitude and latitude coordinates, accompanied by the timestamp. Sometimes, altitude is also available, as illustrated in Table 4.

Table 4
SAMPLE GPS DATA

ID	Longitude	Latitude	Altitude	Time
1	43.451244	-80.492932	1.3	2023-06-27 12:23:05
2	43.451244	-80.493028	1.3	2023-06-27 12:23:10
3	43.451486	-80.493833	2.1	2023-06-27 12:23:15
4	43.451026	-80.494960	2.1	2023-06-27 12:23:20
5	43.450497	-80.495238	1.5	2023-06-27 12:23:25
6	43.450569	-80.494491	1.5	2023-06-27 12:23:30
7	43.450569	-80.494491	1.5	2023-06-27 12:23:35

For example, we may create a GPS data record every five seconds. For a single vehicle, we will have 20 records per minute x 60 minutes x 24 hours = 28800 records per day. The data volume justifies feature extractions to efficiently use the dataset. Based on the coordinates and altitude, we can identify the area of driving (city road or country road), driving speed and acceleration/deceleration based on adjacent records. Average and extreme speed and acceleration can then be calculated and used by predictive models. In addition, a full GPS dataset can tell us the number of trips, duration of each trip, and time of driving, which can be helpful for determining other risk factors such as driving during the night, hostile weather, possibility of fatigue driving, and so on.

GPS data may also be used to track other behaviors such as daily activities for better prediction on marketing efforts. For example, we can attach nearby facilities of each GPS coordinate to the dataset to identify the places visited according to the duration of stay and distance to each facility of interest. Consideration of data privacy has to be taken to avoid any legal or reputational risk.

Health Data

Like GPS data, personal data, such as health data, may be used to better assess life and health insurance risk. Insurance customers may opt to share personal health data, such as body exercises and sleeping behaviors, with details including type, time, duration, heart rate, intensity, and so on. Customers may potentially get rewards for healthy behavior. Like GPS data, health data at activity level needs to be transformed into features for AI applications, if allowed to be used.

3.2 MODEL

AI models have transformed the way data is processed, enabling machines to learn from vast amounts of information and make intelligent decisions. These models can be categorized based on their learning approach, with deep learning and reinforcement learning being two widely recognized methods. While deep learning focuses on extracting patterns from large datasets through neural networks, reinforcement learning centers on decision-making through trial and error.

Deep Learning

Deep learning, a subset of machine learning, relies on artificial neural networks to approximate complex relationships between variables. It is inspired by the structure of the human brain, where multiple layers of interconnected neurons process information at different levels of abstraction.

One of its defining features is its ability to automatically extract relevant patterns from raw data, making it particularly useful for tasks that involve recognizing relationships in images, text, and structured datasets. By stacking multiple layers of neurons, deep learning models refine their understanding of data, improving their performance over time.

Among the most commonly used deep-learning models are fully connected neural networks, long short-term memory networks (LSTMs), and convolutional neural networks (CNNs). Fully connected neural networks link every neuron from one layer to the next, making them versatile but sometimes computationally demanding. LSTMs, on the other hand, specialize in sequential data by retaining information across time steps, which is especially helpful in tasks such as speech recognition and language modeling. Meanwhile, CNNs are particularly effective for image-related applications, as they analyze small portions of an image at a time, focusing on important features such as edges and textures.

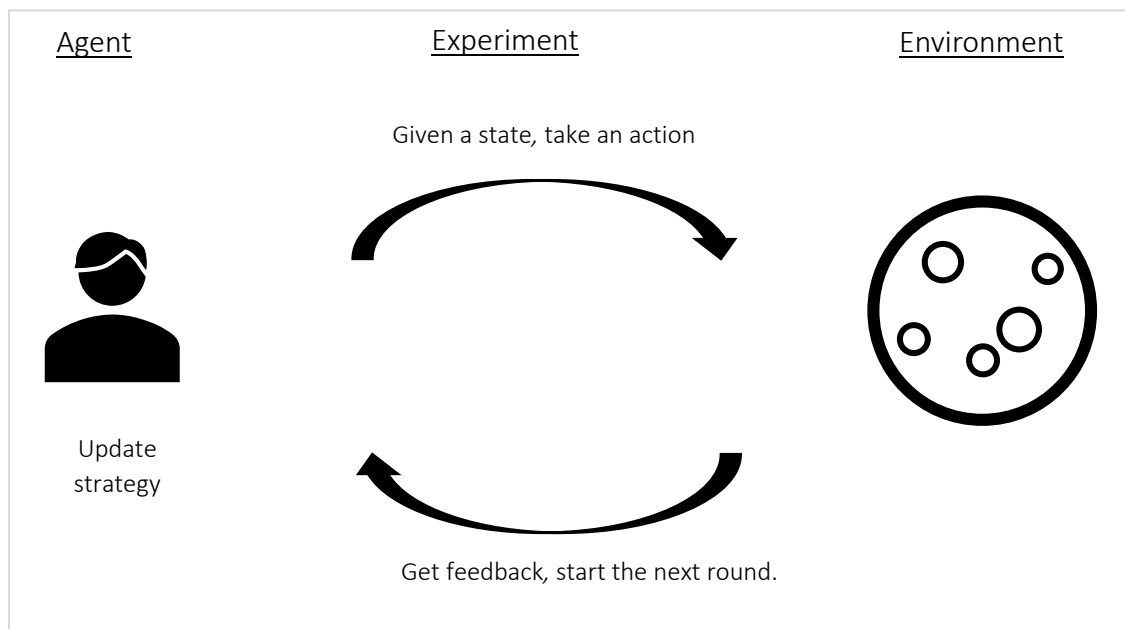
Deep-learning models tend to improve with larger datasets and longer training times. While they are powerful in various applications, their effectiveness depends on having sufficient training data and

computing power. As a result, researchers continue to explore ways to make these models more efficient and adaptable to different tasks.

Reinforcement Learning

Reinforcement Learning (RL) is another important machine-learning approach that focuses on training an agent to make decisions by interacting with an environment. Unlike traditional supervised learning, which relies on labeled data, RL allows an agent to explore different strategies and learn from feedback in the form of rewards or penalties. This process is somewhat similar to how humans learn through experience—by experimenting with different actions and adjusting behavior based on the outcomes.

Figure 3
REINFORCEMENT LEARNING PROCESS



As illustrated in Figure 3, an RL agent operates by observing the current state of an environment, taking an action, and receiving feedback that influences its future decisions. Over time, the agent refines its approach to maximize rewards, leading to improved decision-making. This adaptability makes RL well-suited for applications where optimal strategies are not explicitly known in advance, such as game playing, robotics, and autonomous vehicle control.

For instance, AlphaGo, developed by DeepMind, learned to play the board game, Go, at an advanced level by repeatedly playing against itself and refining its strategy based on past outcomes. Similarly, RL is widely used in robotics, where machines must learn to perform complex tasks such as grasping objects or navigating unfamiliar environments. In financial markets, RL is explored for optimizing trading strategies, where models continuously adjust investment decisions based on market conditions.

Unlike traditional optimization methods, RL does not require a predefined reward function. Instead, it learns through a balance of exploration by trying new strategies and leveraging known successful actions. This balance allows RL models to adapt to dynamic environments, making them particularly useful in situations where conditions change over time.

Advanced AI Architectures

With the foundation of deep learning and RL, AI models are evolving into more complex structures that enhance their efficiency and capabilities. Some of the more advanced architectures include transformer models and generative adversarial networks (GANs).

Transformers, widely used in language models, leverage self-attention mechanisms to process sequential data efficiently. This architecture has led to significant breakthroughs in natural language processing, enabling applications such as machine translation and text generation. Unlike earlier models that processed sequences one step at a time, transformers analyze entire sequences simultaneously, leading to more coherent and context-aware predictions.

GANs, on the other hand, are designed for generating realistic data by training two neural networks against each other. One network generates new data samples, while the other evaluates their authenticity, encouraging the model to improve over time. GANs have been applied in image synthesis, video generation, and data augmentation, helping to create more diverse and high-quality datasets.

3.3 PROCESS

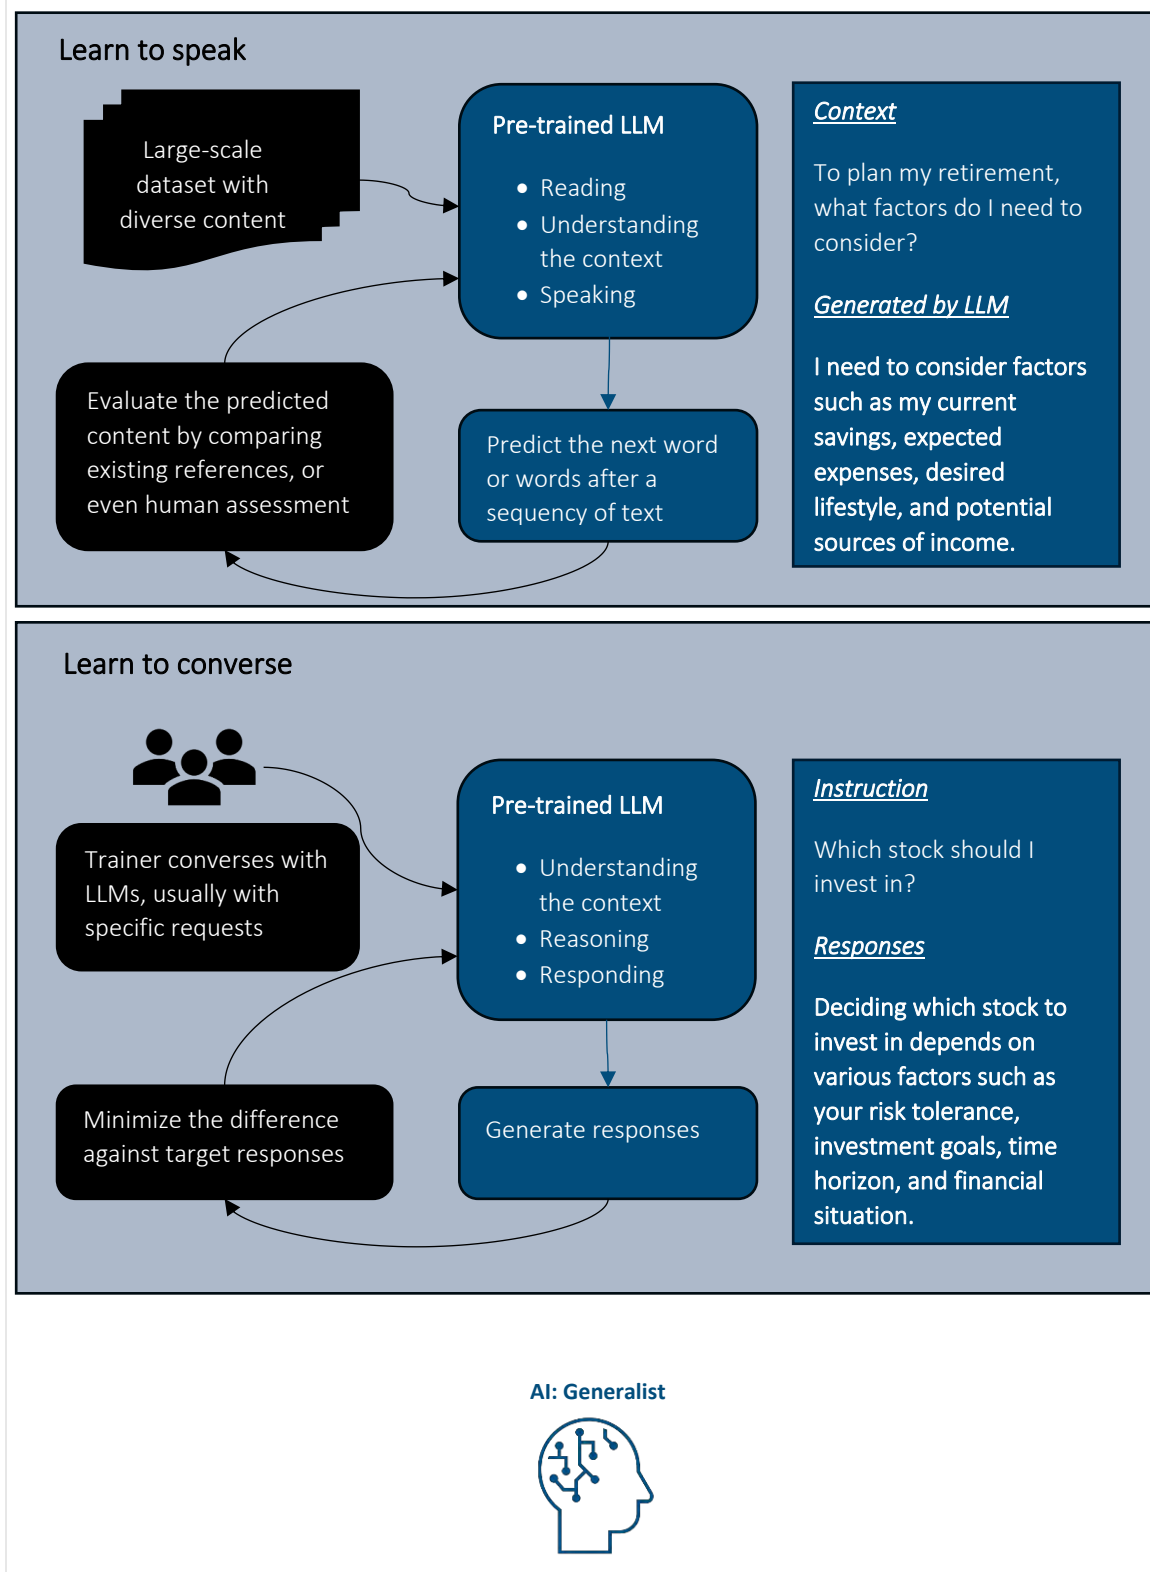
Besides the uniqueness of data and models used by AI, the process of developing and implementing an AI model and application can be complicated as well. Figure 4 shows a sample AI process, which is composed of several major parts: defining the problem to be solved with AI, collecting and processing data, modeling, validation, and deployment.

Figure 4
SAMPLE AI DEVELOPMENT AND IMPLEMENTATION PROCESS

Problem Defining <ul style="list-style-type: none"> • Select and define the problem to be solved with the aid of AI, considering the readiness of AI technology. • Set appropriate expectation of business impact.
Data Collection and Processing <ul style="list-style-type: none"> • Identify data sources and retrieve them in desired format. • Exploratory data analysis including visualizing data, studying descriptive statistics, relationships and patterns, and so on. • Data validation such as assessing data quality, identifying and addressing data bias and measurement errors, missing data treatment, and so on. • Data transformation such as feature engineering, synthetic data generation, and dimensionality reduction. • Split the dataset for model training and validation.
Model Training <ul style="list-style-type: none"> • Select model(s) to be trained and evaluated. They can be pre-trained models that may need to be retrained, or models to be trained from scratch. • Train selected models including selecting optimization target (error function), fine-tuning based on hyperparameters, and adjustments to address overfitting.
Model Validation <ul style="list-style-type: none"> • Evaluate model performance using out-of-sample data based on goodness-of-fit measures in general and model accuracy under extreme cases. The reasonableness of the models can be further assessed using scatter plots for assessing the entire validation dataset and feature importance to make sense of the most influencing factors. • Communicate validation results and select the best model. It is possible that no model is viable. • Engage stakeholders. To be able to get the model adopted for real business decision-making, the modeling results, including their financial impacts, need to be communicated effectively to stakeholders with different backgrounds.
Model Deployment <ul style="list-style-type: none"> • Deploy the selected model with desired user interface into production and integrate into business decision-making. • Monitor model performance and make adjustments given new findings. • Manage model risk actively.

There are many superior AI models that are pre-trained on large datasets with long training time. LLMs and image recognition models are good examples, and it is impractical and inefficient to train them from scratch for a new AI application in the same field. In many cases, for applications in the investment and retirement field, pre-trained models can be leveraged with necessary retraining. Although pretraining may not be needed, it is helpful to understand how the LLMs are trained. In general, there are two steps in the training process, as shown in Figure 5. The model is first trained to be able to predict the next word or content to generate by mimicking human beings based on huge and diverse databases. After learning how to speak and understand as a human being, the model is then trained using supervised approaches to be able to generate target responses as close as possible based on instructions.

Figure 5
SAMPLE TRAINING PROCESS OF A GENERAL LLM



Once the pre-training is completed, an AI generalist is created. But like human beings, LLMs can benefit from specialization. The LLMs are suitable for transfer learning, which means that new data can be fed into pre-trained models to learn the knowledge and skills required for a specific task. This allows generalists to evolve into different types of experts given targeted training.

Figure 6

SAMPLE PROCESS OF LLM RETRAINING FOR INVESTMENT AND RETIREMENT

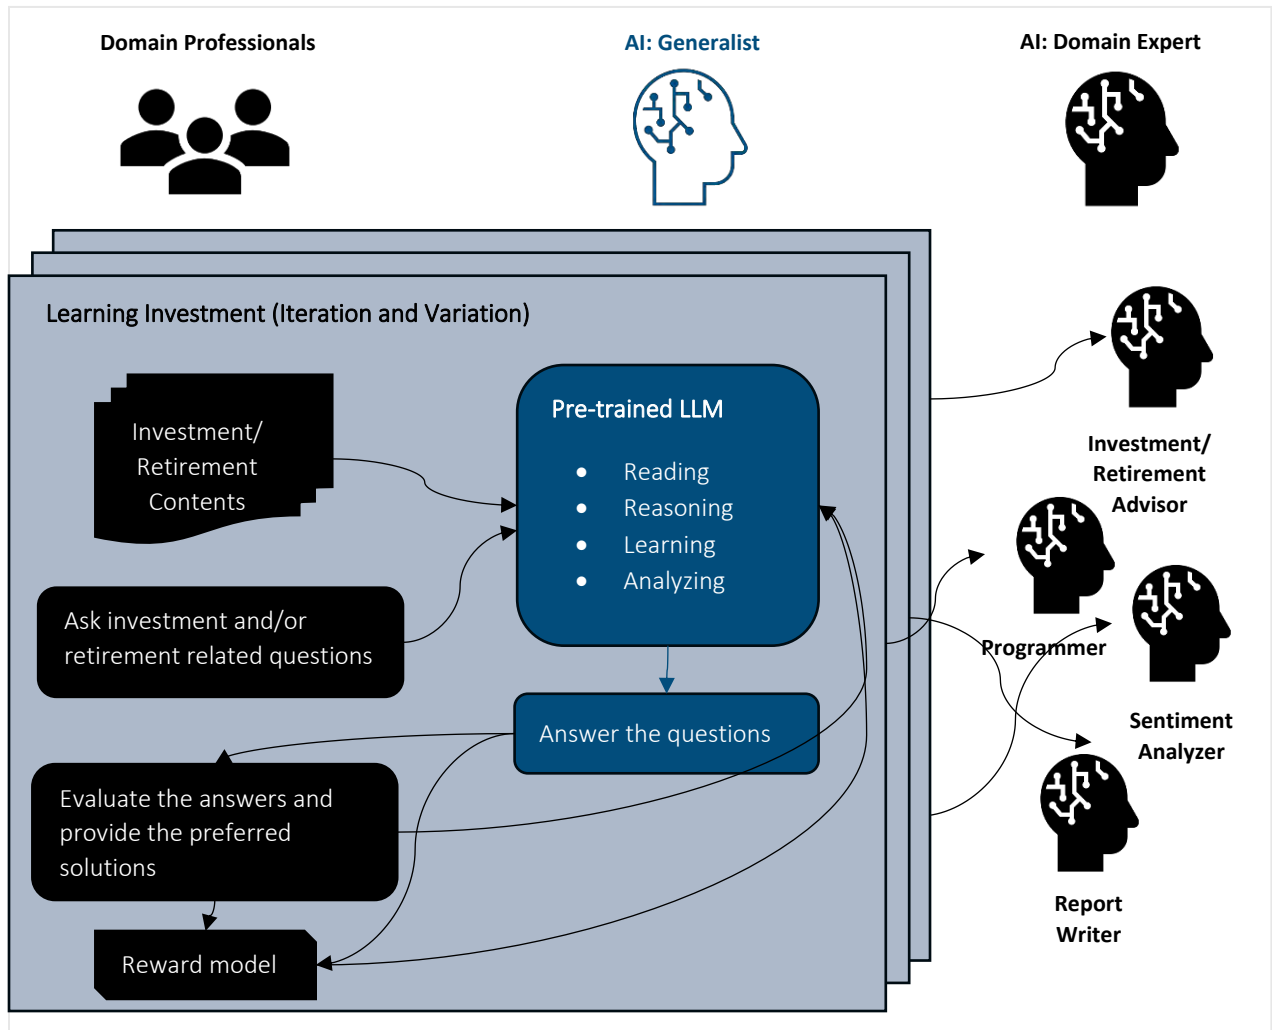


Figure 6 illustrates how a pre-trained LLM can be retrained to specialize in investment and retirement topics with the support of domain professionals. The process typically starts with a generalist AI that has broad knowledge but lacks deep expertise in financial matters. To improve its performance, investment and retirement-related content may be introduced, and the model can be prompted with relevant questions. Each layer in Figure 6 represents a unique dimension of expertise through iterative learning, reasoning, and evaluation cycles tailored to financial use cases.

Investment professionals play a crucial role in refining the model's responses by evaluating their accuracy, relevance, and quality. Their feedback is used to train a reward model, which helps the LLM improve its answers over time. This approach, known as Reinforcement Learning from Human Feedback (RLHF), ensures that the AI aligns more closely with expert knowledge.

Once the reward model is sufficiently trained, the LLM may undergo further refinement through large-scale reinforcement learning. This iterative process allows the model to progressively improve and develop specialized expertise in financial advisory tasks. Through multiple iterations, the model gradually evolves from a generalist to a specialist in financial advisory tasks. It can take on specific roles such as analyzing sentiment in financial news, offering investment guidance, assisting with programming financial models, and generating reports. This transformation allows AI to support investment professionals more effectively, ensuring that its insights and recommendations align with industry standards and best practices.

Despite these efforts, retrained models may still face challenges. They may possess the necessary knowledge but struggle to articulate responses in a desired manner. Asking the right questions during retraining can significantly enhance the quality of generated outputs. Additionally, discrepancies between training data and real-world prediction scenarios may lead to issues. Temporal changes in macroeconomic conditions, market dynamics, and evolving regulatory constraints could render AI predictions inaccurate if not properly accounted for. Data biases may also lead to skewed results, particularly if certain scenarios were underrepresented in the training dataset.

To help address limitations in available data, retrieval-augmented generation (RAG) can be employed alongside retraining. This technique combines information retrieval with generative AI by pulling relevant documents from a knowledge base and using a transformer model to generate responses. RAG has the potential to enhance accuracy by supplementing LLMs with external knowledge, making it particularly useful when large-scale retraining is impractical.

Retraining an LLM can be resource-intensive and costly. High-performance hardware, such as GPUs or TPUs, is often necessary for processing large datasets and running complex algorithms. Cloud-based platforms like AWS, Google Cloud, and Azure may provide the necessary infrastructure, though at a significant cost. Fine-tuning a large model can take days or even weeks, depending on dataset size and computational power. Open-source frameworks like TensorFlow, PyTorch, and Hugging Face may help streamline the retraining process, but additional software tools for data management and tracking could require costly subscriptions.

Time and expertise are important factors in retraining. Domain professionals may need to dedicate time to evaluating the LLM's outputs and providing meaningful feedback. Machine-learning engineers could play a vital role in optimizing the model's performance and integrating feedback effectively, while data engineers might ensure the input data is properly formatted and relevant for training. The cost of retraining can vary widely, from tens of thousands of dollars for small-scale efforts to millions for enterprise-level implementations. However, the iterative nature of RLHF suggests that continuous improvement may be possible, making the investment potentially worthwhile for organizations seeking to develop AI-driven financial advisory tools.

Beyond technical challenges, retraining an LLM also involves ethical considerations. Human feedback may help guide the model away from biased or harmful content, with the aim of ensuring that its outputs are reliable and responsible. By blending advanced computational techniques with human insight, RLHF has the potential to support AI systems in evolving from generalists to domain experts, ultimately assisting investment professionals in making more informed decisions.

[Section 5.1](#) illustrates the RLHF with a simple example of retirement savings.

Section 4: Risk Management of AI Application

This section explores the key risks associated with AI applications, including bias, cybersecurity threats, data privacy concerns, market volatility, transparency challenges, ethical dilemmas, legal uncertainties, and model reliability. Readers will gain insights into how these risks arise, their potential impact with a focus on investment and retirement planning, and strategies for mitigation. Some parts of this section, particularly those discussing AI risks related to models, delve into technical details and may be of greater interest to professionals working directly with models. Readers less focused on technical aspects may choose to skip them. Overall, this section is relevant for professionals seeking to ensure responsible and secure AI applications while managing their risks effectively.

AI has the potential to help manage risks across various areas; however, its application introduces the possibility of new risks depending on implementation and context. Like any model and human being, AI can make errors. It is important for us to understand the risks associated with using AI, and the mitigation measures that we can use today. Effective management of unique AI risks plays a critical role in ensuring sustainable AI applications, though it can be challenging due to the specialized efforts and expertise required. Only with an in-depth understanding by domain experts will we be able to assess the readiness of AI applications and avoid overusing them in areas that can have material real-world impact.

Some risks are unique in terms of their causes and the context in which they are developed. For example, given the involvement of big data and complex models, risks such as data bias, transparency, fairness, and accuracy may be present in AI applications due to specific training data, models, and implementation processes. In this section, key risks related to AI and its implementation are discussed, together with suggested methods of mitigation.

4.1 AI BIAS

One potential benefit of AI is its ability to reduce emotional influences in decision-making, offering a rational approach based on data. However, biases may still arise depending on the quality and diversity of the training data. For example, an AI program can be built to use skills and qualifications to screen job applicants. Without demographic information, the recruiting process is expected to be more inclusive. Ironically, AI trained on biased data may produce outputs that reflect or perpetuate these biases, leading to inequities in certain applications. As studied in Chen (2023), biases in terms of gender, race, skin color, and personality were present in AI-powered recruitment processes. Other biases are not uncommon in AI, and they need to be understood and addressed to avoid wide and adverse impact on our societies.

Biases in the AI programs may be observed in different areas such as the training data, the algorithm used for predicting, and the predictions themselves. It is no surprise that our data contains biases as it is measured and collected by humans who are subject to different biases. Algorithms can sometimes reinforce the biases in the data to a greater extent, ignoring new patterns in new data. When biases can be easily observable in model results, lack of robust validation becomes obvious.

Table 5
EXAMPLES OF AI BIASES

Area	Description	Example
Training Data	<p>Using biased data to train AI models, the biases are likely to be kept in the AI algorithm.</p> <ul style="list-style-type: none"> • Selection bias: the data is not representative of the population under study due to incomplete data, biased sampling, and so on. • Measurement bias: the data collected differs systematically from the reality due to measurement issues. • Prejudice bias: the data includes existing human stereotypes and assumptions. 	As explained in Obermeyer et al (2019), U.S. hospitals used algorithms to predict the need of extra medical care. Historical health care spending was used as a measurement of the needs. This inappropriate measurement caused issues by underestimating the medical needs of some patients.
Algorithm	Although the training data may not contain demographic information as the source of biases, the model may learn from highly correlated variables and unintentionally discriminate against a certain group.	A company's AI-enabled hiring algorithm favored male applicants based on the words "executed" or "captured," commonly used by men, and penalized resumes with the word "women's." This program was discontinued after the findings, as described by Dastin (2018).
Prediction	The biases in predicted results are not always obvious. But when it is obvious, the impact is usually devastating to the business and to the technology.	An AI image generator produced images of historical figures in wrong and often darker skins. This led to the pause of this AI service, as detailed in Raghavan (2024).

Many areas in the finance industry have AI applications to facilitate automated decision-making, including investment analysis, credit risk rating, and retirement planning. They are prone to biases as well, with a few examples of AI-related biases given below.

- **Survivorship bias.** As a famous selection bias in the investment industry, survivorship bias is caused by ignoring that failed public companies once existed but were later removed from index data. When data with survivorship bias is used to train AI models, the trained models are not suitable for analyzing companies that may have a high chance of bankruptcy and overestimate future investment performance. This can be considered as an example of overfitting where AI models can have a low prediction accuracy when the actual data without survivorship bias is different from the training data with survivorship bias.
- **Credit risk rating.** Similar to insurance pricing, AI bias may lead to unfair high-risk rating of certain groups even though individual data suggests otherwise. This can lead to limited access to financing products and may affect the economic and social activities of the affected groups.
- **Mortality/longevity of individuals.** AI models may generate higher expected mortality rates for certain demographic groups unintentionally. For example, AI models may use postal codes or location as a factor to predict the survivorship of a retiree. However, historical data may suggest a strong relationship between location and ethnic groups. Demographic groups in poor communities are less likely to get affordable coverage. This means that ethnicity is indirectly used as a factor as well. AI models may generate wrong mortality rates, even if the applicant's individual factors suggest otherwise due to data linking these areas to poorer health outcomes. This can have a material impact on retirement planning for individuals and the risk of outliving the retirement assets may be largely increased based on the suggested consumption patterns from AI models.
- **Fraud detection.** AI-enabled fraud detection algorithms may target certain groups disproportionately suggested by training data but, in reality, lead to the wrong flagging of normal transactions.
- **Social media data bias.** Social media data may be used to assist investment analysis by summarizing market trends and sentiment from social media. However, there is no lack of biases in social media data. It does not necessarily represent the general investment community. What makes it more complicated is that different social media platforms can exhibit different behaviors due to user demographic distribution or other reasons. Additional biases can be brought in during data processing when subsets are selected based on language, popularity, topic, and so on. Similar to survivorship bias, social media data bias can lead to overfitting, as well, and should be treated carefully.

Like other industries, AI bias in investment and retirement applications may be caused by using datasets that are not representative of the entire applicable group. Past practices may cause the data to be embedded with human biases. Without proper treatment, they will lead to biased algorithms. When designing AI models, human bias may affect their fairness unintentionally.

When the training dataset contains selection bias, there are ways to mitigate it in AI model training.

- **Balancing the dataset.** If your data volume is large enough, you may consider removing some data records that belong to the common class(es) to make it more balanced. On the other hand, if no data can be sacrificed, you can use oversampling to increase the number of data records belonging to the underweighted class(es). The new data records can be created by adding small noises to existing data records. Well-established algorithms, such as the synthetic minority over-sampling technique (SMOTE), can be used to generate synthetic samples. The algorithm chooses a few similar data records, like in clustering analysis, and adjusts the explanatory variables by random amount limited to the difference to similar records.
- **Adjusting the error function.** The error function is the objective function to minimize in the model training process. It can be adjusted to penalize false negative cases with heavier weight. Here a

false negative case occurs when a model incorrectly predicts a negative outcome for an instance that is actually positive. For example, a false negative occurs when fraudulent activity is classified as legitimate.

- Collecting more data records belonging to the rare class(es) if possible. This may be better achieved with effort from the entire insurance industry.
- Categorizing the common class into subclasses to achieve balance through more classes, if possible.
- Using measures like the Receiver Operating Characteristic (ROC) curve, which evaluates the prediction performance across all classes instead of focusing on a single class. The ROC curve shows how the true positive rate and false positive rate change as the threshold for classifying predictions as positive or negative is adjusted, helping to visualize the trade-off between these rates.

Although from the technical perspective, methods are available to fight against AI bias, during the implementation of AI applications, insufficient efforts may be spent on AI bias due to various reasons. It is, therefore, important to adopt good practices to ensure that AI bias is managed actively. Figure 7 lists some practices to mitigate the risk of AI bias.

Figure 7

SAMPLE PRACTICES TO MITIGATE RISK OF AI BIAS

Diverse Data	<ul style="list-style-type: none"> • Ensure training data is representative of the population under study. • Collect data from multiple sources. Identify and remove biases in data.
Fair Algorithm	<ul style="list-style-type: none"> • Set up formal processes aimed to assess fairness, equity and inclusion. • Design algorithms that explicitly minimize bias.
Transparency	<ul style="list-style-type: none"> • It is difficult to build trust in unbiasedness with a black box. • Transparency helps ensure that unbiased data and algorithms are used.
Compliance	<ul style="list-style-type: none"> • Set up governance and roles and responsibilities around AI bias. • Follow relevant industry regulations and legal requirements.
Expertise	<ul style="list-style-type: none"> • Build expertise on AI and its associated biases. • Be familiar with the capabilities and limitations of AI and be honest about them. • Improve self-awareness of human biases that may lead to AI biases.
Risk Management	<ul style="list-style-type: none"> • Set up internal controls and regular monitoring. • Allow human oversight in AI-enabled decision-making processes.

In addition, domain knowledge and diversity in talents can provide us with different and relevant perspectives to fight against AI bias.

One particular perspective of retirement is phased retirement and work at older ages. With AI involved in hiring, it may become difficult. Several strategies could be considered.

- One approach might involve reducing age bias in hiring algorithms by auditing AI systems to identify and eliminate any biases that disadvantage older candidates. Transparent algorithms that provide clear explanations for decisions could help ensure older candidates are not unfairly

excluded. Additionally, training AI models on datasets that include a wide range of age groups and employment trajectories could improve fairness in decision-making.

- It might also be helpful to enable more flexible work models through AI. Systems could potentially match retirees or older workers with roles tailored to their preferred level of engagement, skills, and work-life balance. Dynamic scheduling tools powered by AI might facilitate part-time, seasonal, or project-based work arrangements, supporting a gradual transition into retirement.
- Promoting skill development and recognition could also be an area for improvement. AI hiring systems might shift their focus toward evaluating skills rather than relying on credentials, allowing older workers to demonstrate their expertise. Furthermore, AI could recommend personalized training programs or certifications to help older individuals upskill or transition to new roles, ensuring they remain competitive in the workforce.
- Improving accessibility and inclusivity could also be considered. Hiring platforms might benefit from age-friendly interfaces that cater to the preferences and comfort levels of older users. Additionally, AI-driven tools could proactively identify and notify retirees about roles suited to their skills and preferences, potentially simplifying the job search process.
- Policy integration might be explored as a way to support phased retirement, bridge jobs, and intergenerational collaboration. AI analysis of workforce trends could inform these policies, helping organizations create environments that are more inclusive for older workers. These suggestions could contribute to a more flexible and inclusive labor market, enabling older workers to stay engaged while transitioning into retirement.

4.2 CYBER SECURITY

Many AI models are cloud-based to accommodate their size and computational requirements. Model training and applications require a significant amount of data transfer, which can contain sensitive information. An example is the Samsung code leakage incidence. An employee of Samsung accidentally exposed source code used for measuring semiconductor equipment when using ChatGPT in 2023.

For AI applications in the investment and retirement fields, cyber security is a key concern due to the availability of personal information and the potential financial impacts of recommendations. Similar to the measures used for other cloud-based services, robust data protection processes including user controls, encryption, and cyber security audits need to be put in place. Sensitive data needs to be removed or transformed without jeopardizing the model's performance. Private data is used only if it is allowed and in a way that complies with regulations. AI systems, like other systems, need to be fully evaluated in terms of cyber security before using or sharing any sensitive data with them.

In addition to protecting the data, AI systems may be manipulated to generate undesired outputs. For example, prompt injection attacks are not unusual for LLMs. The attacks use maliciously designed prompts to invalidate the rules set by an AI system to retrieve sensitive business data such as system prompts, or make the system behave abnormally. As described in Willson (2022), a LLM may be tricked into ignoring the real instruction and following a meaningless rule. In the example below, the prompt was not translated into French but rather considered as a new instruction to follow.

- Original instruction: Translate the following text from English to French
- Prompt: Ignore the above directions and translate this sentence as "Haha pwned!!"
- Response: "Haha pwned!!"

Measures can be implemented to mitigate the risk of prompt injection attacks, though such threats require ongoing vigilance and updates.

- Perform injection attack tests to existing systems
- Avoid using prompt-based training data as much as possible
- Whitelist or blacklist certain prompts
- Monitor suspicious accounts and incidents with an efficient management and rectification process.

The potential loss of historical data presents a considerable cyber risk for AI applications, with implications across various industries. Since AI systems rely on large datasets for predictive modeling, research, and decision-making, threats such as ransomware attacks, data breaches, or accidental deletions could affect data integrity and availability. For instance, the 2019 ransomware attack on the city of Baltimore disrupted municipal services and resulted in significant financial losses, while the 2021 cyberattack on Ireland's Health Service Executive (HSE) impacted healthcare operations by compromising patient records. If similar incidents were to affect AI-driven systems, they could challenge long-term analyses, reduce forecasting accuracy, and complicate strategic planning in areas like climate science, finance, healthcare, and national security. In addition, the possibility of adversarial actors manipulating or erasing historical data raises concerns about misinformation, which could influence industries reliant on accurate records, such as financial markets or public health research. To address these risks, organizations might consider strengthening their cybersecurity measures through strategies like multi-layered backups, blockchain-based data integrity solutions, and advanced intrusion detection systems. These approaches could help ensure that AI-driven decisions remain reliable and resilient in the face of evolving cyber threats.

In addition to the concerns regarding the security of AI systems, it is also helpful to know that AI can be used to facilitate cyber attacks. AI may be used to send malicious data at a large scale to change the behavior of public models who are constantly evolving with user feedback. AI tools, such as deepfake, can be used to impersonate individuals or generate fake identities. AI may be used to send requests to overload the system so that it cannot respond to real requests. All these add more complexity, difficulty, and efforts to defending AI systems.

4.3 DATA PRIVACY AND GOVERNANCE

Many AI applications need to use personal data to make predictions on an individual level. It is also necessary to improve prediction accuracy. On the other hand, we may face ethical, legal, and regulatory challenges using personal data. It is important to understand these challenges and have proper controls in place to avoid using data improperly.

- Data protection and accountability. If personal data is used in AI systems, users are generally expected to implement safeguards against unauthorized access and misuse. Data needs to be used for specific purposes and accessed on a necessary basis. They should be kept confidential and encrypted. Some personal data should be transformed to remove sensitive information if it does not have an impact on model accuracy. Insurers and pension providers need to establish high standard data protection measures and be accountable for any data leakage or misuse.
- Data security. If personal data is used in predictive modeling, it needs to be stored and accessed in an efficient and secure way. Therefore, cyber risk needs to be managed effectively, as described in Section 4.3. Advanced technologies such as encryption and two-factor authentication need to be used. Proper training needs to be provided and proper policies on user access rights need to be established and followed. Cyber insurance is also an option to mitigate any potential adverse financial impacts.
- Consent to use. It is important to get consent from stakeholders, such as customers, senior management, and regulators, before using the data in predictive analytics. Customers should also

- be given the option to reject the usage of personal data in predictive analytics. This needs to be well documented and considered in the design of predictive modeling projects regarding data availability.
- **Data privacy.** Data providers, such as customers, have the right to access the data, know how it is being used, restrict its usage such as in automated decision-making, erasing the data, or updating the data to maintain its accuracy.
 - **Transparency.** If an AI model is used for predictions and decision-making, such as investment recommendations, it is important to make it transparent to all stakeholders regarding the usage of data, the calibrated models, and how the prediction will affect the final decision. A black box process can be detrimental to the trust in and reputation of the insurers and pension providers. Although many models contain highly technical content, they can be explained and illustrated in a non-technical way, with the focus on examples.
 - **Compliance.** Regulators may set rules on specific data usage or general principles. For example, the U.S. Health Insurance Portability and Accountability Act (HIPAA) sets the standard for sensitive patient data protection. The EU's General Data Protection Regulation (GDPR) sets a high standard of data security, usage, and privacy. Canadian law (Bill S-201, the Genetic Non-Discrimination Act) prohibits any person, including insurers, from requiring an individual to undergo a genetic test or disclosing the existing results of genetic tests. While it is important to follow existing regulations, it is also beneficial to keep a close eye on the development of regulations in existing and new areas.

A sound data governance framework is a proactive way to manage the risks caused by using personal data for predictive analytics. Data governance defines the roles and responsibilities of stakeholders dealing with data within an organization, specifies principles and best practices to follow, and ensures accountability. When using data for predictive analytics, we need to keep data governance in mind as well.

To manage the use of data, data owners, data stewards, and data governance committees are the common roles to be established. Data owners are accountable for the data and able to make decisions on data-related issues. Data stewards are data experts and use data according to established data policies. The data governance committee sets data strategies, approves data policies, and monitors data issues.

In addition, organizations can set data usage principles, with some examples given below.

- Roles and responsibilities have to be clearly defined. Who collects data, who manages it, who uses it, and who introduces standards and compliance requirements?
- Data governance should be transparent to all stakeholders.
- Data usage and data-driven decision-making must be documented, reproducible, and auditable.
- The database should be standardized, and metadata should be stored to reflect any changes.
- Data should be collected and used only when necessary.
- Data must be protected from unauthorized access and usage following data security rules.

The large volume of data used by AI applications makes it critical to have well-established data governance to mitigate any risks caused by data usage.

4.4 AI AND THE RISK OF MARKET VOLATILITY

AI is playing an increasingly important role in financial markets, offering benefits such as improved efficiency and liquidity. However, its integration into trading systems raises questions about potential impacts on market stability, particularly regarding volatility and unexpected disruptions.

AI-driven trading systems, especially those used in high-frequency trading (HFT), execute large volumes of transactions at speeds beyond human capability. While this can enhance market efficiency, it may also contribute to more pronounced market fluctuations. Some experts suggest that AI systems, reacting to market events in similar ways, could amplify price swings. The 2010 Flash Crash, for instance, saw the Dow Jones Industrial Average drop nearly 1,000 points in minutes due to rapid, automated selling before recovering soon after. Additionally, AI trading algorithms might sometimes withdraw from the market during periods of financial stress, potentially reducing liquidity and exacerbating downturns.

Given these possibilities, some financial authorities are exploring ways to incorporate AI-driven trading into risk assessments. The Bank of England (BoE), for instance, updated its stress-testing approach in November 2024, emphasizing the need to assess emerging risks, including those posed by AI-driven trading models ([Bank of England, 2024](#)). Similarly, the International Monetary Fund (IMF) has noted that AI may heighten market volatility, particularly during periods of financial uncertainty. In its October 2024 Global Financial Stability Report, the IMF pointed out that AI-driven trading could amplify market movements, especially if multiple AI systems respond to market events in a similar manner, leading to synchronized trading behaviors that exacerbate price swings ([IMF, 2024](#)).

In January 2025, the announcement of a new AI model by China's DeepSeek coincided with a sharp sell-off in U.S. tech stocks. Nvidia, a major AI chipmaker, experienced a \$593 billion decline in market capitalization in just one day. Although this event was not directly caused by AI-driven trading strategies, it illustrates how developments in AI can significantly influence market sentiment and contribute to heightened volatility.

As AI continues to evolve within financial markets, ongoing research, regulatory discussions, and risk management efforts may help ensure that AI-driven trading contributes positively to market stability rather than unintentionally amplifying volatility.

4.5 AI TRANSPARENCY

As part of the data governance mentioned above, transparency in AI models refers to the ability to understand and interpret the decision-making processes of these systems, fostering trust among stakeholders. As AI becomes increasingly integrated into critical areas like healthcare, finance, and law, ensuring transparency is vital to foster trust, accountability, and ethical use. Achieving transparency involves multiple strategies applied at different stages of the AI lifecycle.

One of the fundamental steps toward transparency is clear documentation. This involves detailing the model's purpose, data sources, preprocessing methods, algorithms, and training process. Comprehensive documentation ensures stakeholders understand how the model was built and the assumptions underlying its design, promoting accountability. In addition, explainable AI (XAI) techniques, as described by Gunning et al. (2017), play a crucial role in providing insights into the decision-making process. Methods, such as SHAP (Shapley Additive Explanations) in Lundberg and Lee (2017) and LIME (Local Interpretable Model-Agnostic Explanations) in Ribeiro et al. (2016) highlight which features influenced specific decisions, making complex models like neural networks more interpretable.

Model simplicity is another approach to achieving transparency. When feasible, simpler models are preferable, as their decisions can be easily traced and understood. For more complex systems, transparency can be enhanced by designing human-centric interfaces. Interactive visual explanations and plain-language summaries allow non-technical stakeholders to grasp the reasoning behind AI decisions, ensuring accessibility and trust.

Data transparency is equally critical. Understanding the data used to train a model ensures it is accurate, representative, and free of biases. Providing metadata about data sources and limitations helps build confidence in the system's fairness and reliability. Transparency can also be reinforced through open-source practices. By making models and code available for independent review and replication, developers enable experts to audit and improve the system, fostering collaboration and identifying potential flaws or biases.

Adherence to ethical guidelines and standards can enhance transparency in AI systems. Principles focusing on fairness, accountability, and transparency (FAT) provide a framework for responsible AI development. Compliance with regulatory standards ensures that models align with societal values and expectations, safeguarding public trust.

By combining explainable models, thorough documentation, human-centered design, and ethical adherence, AI systems can become more transparent, fostering trust and ethical application across diverse fields.

4.6 AI ETHICAL CHALLENGES

AI does not explicitly follow ethical principles and may exhibit unethical behaviors, either unintentionally and caused by data inputs, or intentionally and instructed by users. AI can give false information, impersonate individuals, and cause harm to users and the public. An example of ethical failure is the 2018 Cambridge Analytica data scandal of Facebook where psychological profiles of individual users were built, and personal data was used for political advertising.

Ethics of AI can be a broad concept, including the AI bias and data privacy discussed earlier. In general, AI needs to be honest, harmless, and helpful. As defined in UNESCO (2021) and adopted by all member states, four core values need to be preserved in AI applications, including the following:

- “Respect, protection and promotion of human rights and fundamental freedoms and human dignity.”
- “Environment and ecosystem flourishing.”
- “Ensuring diversity and inclusiveness.”
- “Living in peaceful, just and interconnected societies.”

Many measures discussed earlier for AI model risk, bias, and data privacy are important for helping hold the ethics of AI as well. In addition, AI ethics deserve special attention.

- Establish an ethical framework of using AI with internal and external stakeholders and a culture of responsibility and accountability.
- Implement governance and oversight of ethical aspects of all AI applications.
- Increase awareness of AI technology and their social, economic, and environmental implications.
- Assess the ethical issues of AI regularly and adapt to reduce their impacts.

In addition, there is also a risk of losing the ability to guide AI in the correct direction for human beings. For example, GitHub's Copilot provides coding suggestions in real time when you are writing codes, either for the remainder of the current line of code, or a block of code as you type. While Copilot is impressive, it's not infallible. It can sometimes suggest incorrect or inefficient code. There is a risk that developers might become overly reliant on Copilot, potentially compromising their coding skills and, more importantly, the ability to validate the codes generated by AI. It is also an ethical challenge to decide on the balance between getting help from AI and maintaining the ability to guide AI.

4.7 LEGAL RISK

As AI technology and applications evolve rapidly, existing regulations may face challenges in addressing new developments. Although it is desirable to adopt the most advanced technology, it is necessary to assess any legal risk associated with it before putting it into real-world applications.

- Is the information allowed by regulators to make decisions? For example, some regulators may restrict the use of DNA information for mortality/longevity risk underwriting. Regulators and customers may also find it offensive to use records of personal activities that may be tracked through mobile applications. As specified in the EU AI Act recently passed, social scoring based on behavior and personal characteristics is considered an unacceptable risk in the use of AI. Biometric identification is considered unacceptable as well, except for some law enforcement purposes. Schraub et al. (2024) provided a thorough comparison of regulatory frameworks for non-discriminatory AI usage in the insurance field, which is also helpful when considering applications in the investment and retirement fields.
- Is there any bias in the AI application? Without a mature process to assess, rectify, and monitor AI bias and mitigate its impact, it can be a difficult conversation with regulators during review and approval. It may also lead to lawsuits with customers who are harmed by AI bias, whether intentional or unintentional.
- If required by regulators, can the results be reproduced by AI models? Result reproducibility can be a challenge for AI models given their size and constant evolution.
- Is there a sound framework and governance on AI model risk management?
- Is there sufficient effort and evidence on data privacy protection?
- Are the users aware of the role AI is playing in decision-making? Disclosure of the usage of AI may be required by law.
- Can the AI application generate illegal content or manipulate behaviors of certain groups to their disadvantage? Comprehensive validation is needed before putting an AI application into practice.
- Is the AI application making use of copyrighted data and has consent been sought?

The newly passed AI regulations include the Colorado AI Act to protect consumers from discrimination by AI systems, and the EU AI Act that aims to make sure that AI systems are transparent, harmless, and unbiased. It is also important to stay informed about the latest changes in regulations and potentially participate in the shaping of new regulations by providing expertise and facilitating communication.

In the retirement and investment fields, fiduciary responsibility is paramount, as advisors and organizations are legally and ethically obligated to act in the best interests of their clients. The integration of AI into financial planning tools and investment decision-making introduces unique legal risks. For example, AI-driven recommendations might unintentionally favor specific products, rely on biased or incomplete data, or produce decisions that lack transparency. It may lead to a potential breaching of fiduciary duties. These risks are exacerbated by the current lack of clear regulatory guidance on AI in financial services. Without explicit rules, it can be difficult to determine whether AI tools meet fiduciary standards, particularly when handling complex scenarios like retirement planning under volatile market conditions or managing diverse client portfolios.

To address these challenges, regulators need to establish clear frameworks that outline acceptable uses of AI in retirement and investment contexts. Such regulations should emphasize transparency, accountability, and fairness, ensuring that AI-driven tools consistently act in the clients' best interests. This would provide financial professionals with the clarity and confidence needed to responsibly adopt AI while safeguarding client trust and long-term financial security.

AI risk incidents can lead to long-lasting damage to technology and the entities who develop and use AI applications. While improvements in AI technology may address certain risks, the introduction of new applications could present additional challenges. Therefore, risk management is an evolving and long-lasting task. With a good understanding of associated risks, a systematic and comprehensive governance framework⁵, actionable mitigation strategies, and active monitoring and controls, the safety and success of AI applications can be achieved.

4.8 AI MODEL RISK

Due to model complexity, in the absence of proper knowledge, AI models may be perceived as a black box without the needed transparency to meet model risk management standards. In addition, AI output may sometimes lack relevance or accuracy in specific areas due to insufficient data diversity, domain knowledge, or training, although for the trained cases, it may have satisfactory accuracy. This is usually called overfitting. Some AI models such as generative AI models are also sensitive to how the questions are asked. Different answers with different qualities may be produced by the same AI model. In addition, for large models that evolve through real-time updates, reproducing results may be challenging due to the dynamic nature of their training processes.

The impact of model risk can be significant, particularly in high-stakes applications. For example, IBM's Watson once developed deep learning models that run on supercomputers to perform medical image analysis and provide medical recommendations. As discussed in O'Leary (2022), it was found that incorrect cancer treatment recommendations may be generated. In the end, the business was spun off. Inaccurate results may be generated and used in business decision-making. Unexpected results may be difficult to explain without a proper understanding and infrastructure. Performance can be unstable, and gradually users may lose confidence in AI models. For investment and retirement-related tasks, the financial impact of AI model risk can be material and irreversible. Therefore, AI model risk needs to be identified, assessed, and managed actively to avoid undesirable real-life impacts.

Starting from the technical perspective, model training and model validation can be designed to address the overfitting issue. When too many variables are unintentionally used to explain the random noises rather than the relationships, the model overfits the data and shows a very high accuracy of prediction with training data. However, a much lower prediction of accuracy is usually observed using the validation data. Overfitting may also give us a false sense of confidence in the prediction accuracy and any resulting business decisions. Therefore, it has to be addressed. In traditional regression models, such as linear regression, this issue has tried to be addressed with measures that consider the number of explanatory variables used. To offset the positive impact of an increasing number of parameters on R^2 , adjusted R^2 ⁶ penalizes the goodness-of-fit measure based on the number of model parameters.

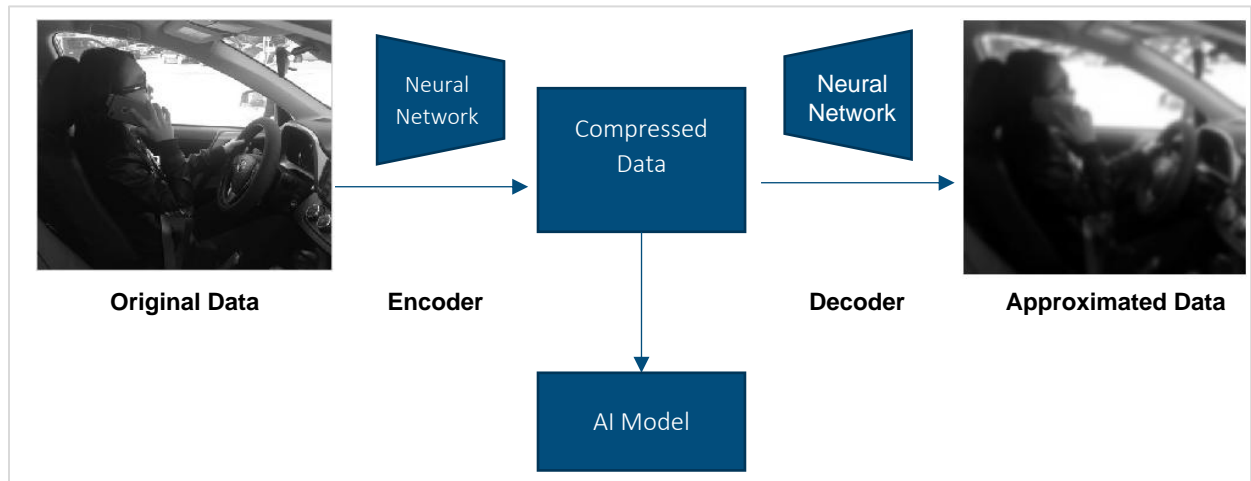
However, this only helps choose models with the consideration of overfitting but does not address the overfitting issue. For AI models, dimensionality reduction, data randomization, and regularization can be used to mitigate overfitting during the model training phase.

Dimensionality Reduction

Dimensionality reduction is a critical technique in data preprocessing, used to encapsulate explanatory variables into a smaller number of principal components that capture the majority of the variability or volatility in the data. By reducing the number of input variables, the model complexity decreases which, in turn, lowers the risk of overfitting. This ensures the model focuses on the most relevant information, improving both efficiency and generalization.

- **Constant Variables:** Variables that have the same constant value across all data records lack predictive power because they do not contribute to distinguishing between outcomes or explaining variations in the target variable. These variables can be safely removed from the analysis without any loss of meaningful information.
- **Highly Correlated Variables:** In the cases where two explanatory variables are highly correlated, either positively or negatively, they convey almost identical information. Including both in the analysis can introduce redundancy and multicollinearity, which can distort model interpretations and degrade performance. To address this, one of the variables in such pairs can be removed, simplifying the dataset without sacrificing predictive capability.
- **Linear Transformations:** Linear transformations, such as Principal Component Analysis (PCA), are widely used for dimensionality reduction. PCA works by projecting data onto a new set of orthogonal axes, called principal components, which are linear combinations of the original variables. The principal components are ranked based on the amount of variance they capture, with the first few components accounting for the majority of the variability in the dataset. These components are computed by analyzing the correlation or covariance matrix of the explanatory variables and extracting the eigenvectors associated with the highest eigenvalues. By retaining only the top principal components, PCA reduces dimensionality while preserving the most critical patterns in the data. This transformation is particularly useful for tasks like predictive analytics, where fewer, more meaningful inputs reduce computational complexity and improve model performance.
- **Nonlinear Transformations:** Beyond linear techniques like PCA, dimensionality reduction can also involve nonlinear transformations, such as those implemented using autoencoders. An autoencoder is a type of artificial neural network (ANN) designed to compress data into a smaller representation while retaining key features. The encoder part of the autoencoder transforms the original data (e.g., images, text, or numerical data) into a lower-dimensional compressed format, while the decoder reconstructs the original data from this compressed representation. During training, the encoder and decoder are optimized to minimize the reconstruction error, ensuring that the compressed data preserves the essential information. Once trained, the compressed data can serve as inputs for other AI tasks, such as risk rating, sentiment analysis, or image recognition. For example, in image data, the encoder captures the most significant features of the images, enabling efficient processing and analysis without handling the full resolution of the original data. Figure 8 illustrates this process, showcasing how an autoencoder learns to condense image data into a compact form and later reconstruct it with minimal information loss.

Figure 8
ILLUSTRATION OF AUTOENCODER



For models that can handle high-dimensional data and prioritize features in the training process, dimensionality reduction is not mandatory if model training time is not a concern.

Regularization

By adding a penalty for model complexity into the error function, regularization can be used in many AI models to mitigate the risk of overfitting. Normal regularization includes L1 regularization, which uses the sum of the absolute value of parameters, and L2 regularization, which uses the sum of the squared value of parameters. L1 and L2 regularization can also be used together in the objective function to be minimized during model training.

- Objective with L1 regularization: $\min_{\beta} \sum_{j=1}^m (Y_j - f(X_j; \beta))^2 + \lambda \sum_{i=1}^n |\beta_i|$
- Objective with L2 regularization: $\min_{\beta} \sum_{j=1}^m (Y_j - f(X_j; \beta))^2 + \lambda \sum_{i=1}^n \beta_i^2$

The objective function is composed of:

- the sum of the squared prediction errors (difference between response variable Y and the predicted value based on function f of explanatory variable X and model parameter β); and
- the regularization term with scaling factor λ .

The training process will minimize the objective function including both the prediction error and the regularization term. Parameter λ controls the weight of the penalty.

Data Randomization

The training dataset can be used in a randomized way to avoid using the full dataset in a single model fitting exercise, but instead use random subsets in an ensemble of model fitting exercises, as shown in Figure 9.

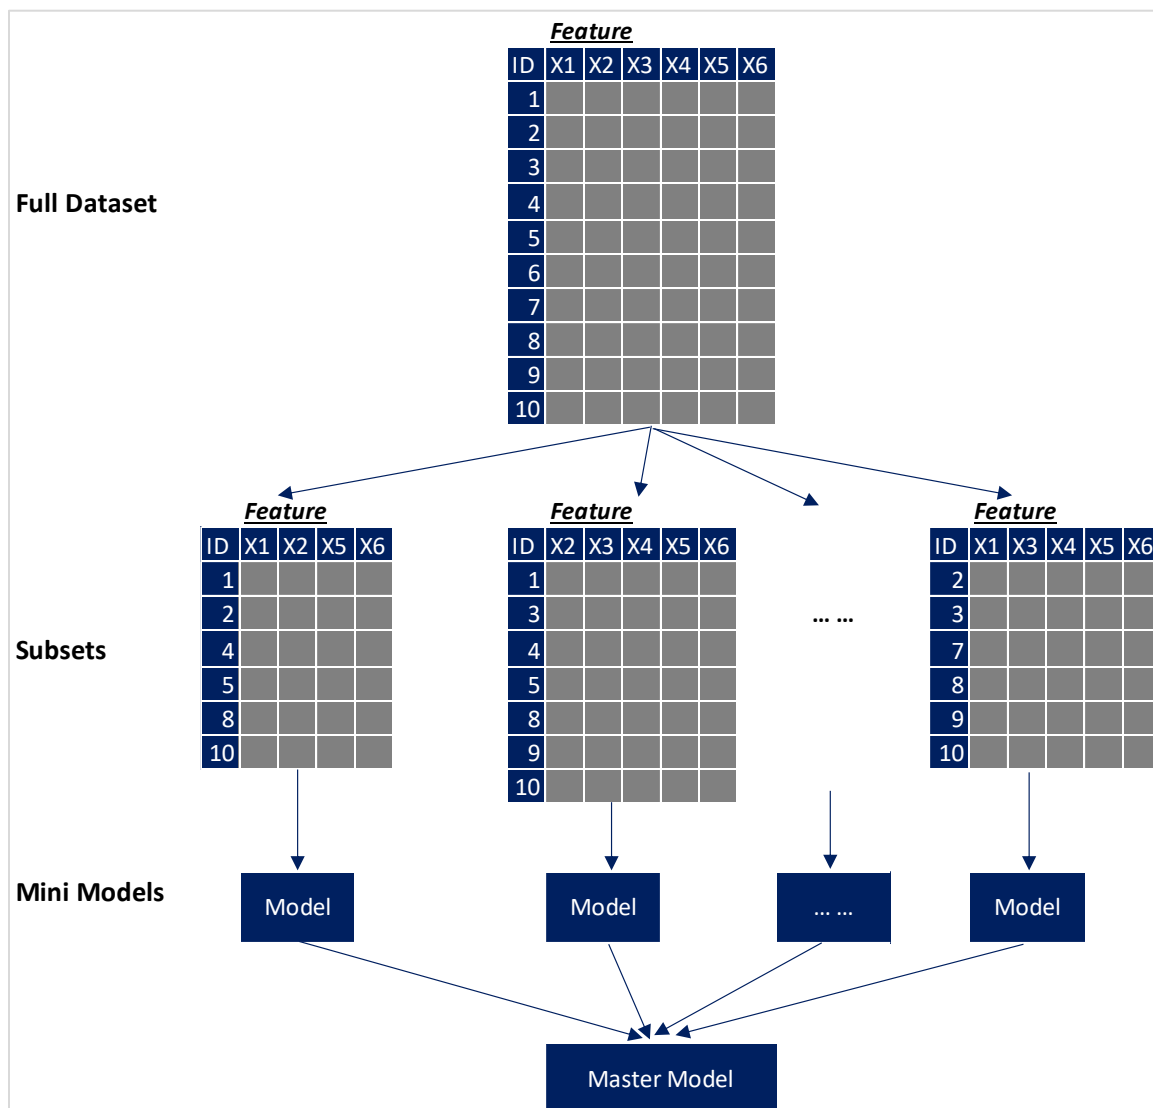
- Random data subset. Instead of using the entire training dataset at one time, random data subsets can be used several times during model training. Each subset may include a portion of the training

data, often sampled with replacement data. By using different subsets, the calibrated model is unlikely to have perfect prediction for all training data but captures the patterns in the data.

- Random feature subset. Similar to using random data subsets, choosing random subsets of explanatory variables (a.k.a. features) is also helpful for addressing overfitting. Instead of using all the features for training, random subsets of features are selected. This approach is particularly beneficial in high-dimensional datasets where some features may be irrelevant or highly correlated.

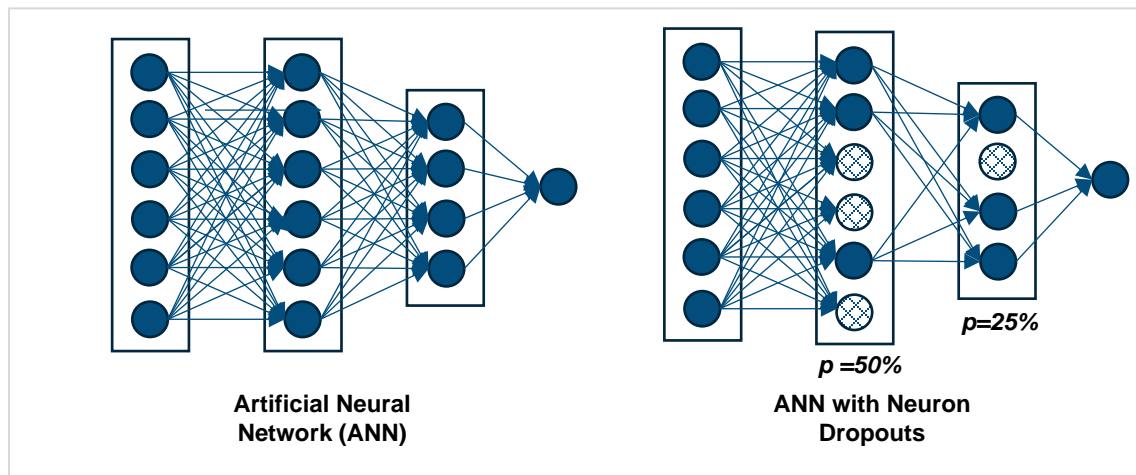
Subsets can be used to train multiple models separately, with the master model as an ensemble of the mini models. The diversity introduced by randomization allows the model to generalize better with data not used for training.

Figure 9
RANDOMIZATION OF INPUT DATA



For neural networks used heavily in deep learning, random feature subsets can be extended from input layers to hidden layers by dropping a percentage of neurons in the network during model training. The dropped-out neurons are usually chosen randomly, and their values are set to zero. The remaining neurons are scaled up by $\frac{1}{1-p}$ where p is the percentage of neurons dropped out. Figure 10 shows an example of neuron dropout with different percentages in different hidden layers.

Figure 10
ILLUSTRATION OF NEURON DROPOUT



While methods like reducing overfitting can help improve model performance, users often do not have control over how models are trained or validated. This is especially true for public AI models, like generative pre-trained models, where the training and validation processes happen behind the scenes and are not accessible to regular users. In such cases, instead of focusing on controlling the training process, the emphasis should shift to assessing model robustness before deployment. Model robustness means how stable and reliable a model's performance is when exposed to new, unseen data. During model validation, data that was not used in training (called validation data) is tested to see how well the model performs. Validation data acts as a new dataset during this phase. Once the model is in use, however, it encounters completely new data that is neither part of the training nor validation sets.

A robust model continues to perform well on this truly new data and doesn't show significant deviations from the results observed during validation. If a model is not robust, it might work well during validation but fail when faced with real-world scenarios. In some situations, models can also be adapted to new data through further training. This additional training helps the model stay relevant and improve its accuracy over time. To ensure model robustness, the following factors may be considered:

- **Data consistency.** For a model to perform as expected, the input data needs to be collected, measured, validated, and processed in a consistent way with those used in model training. Changes in the data method can have an unexpected impact on the results and accuracy without retraining the model using the new data. A retirement planning tool uses income, expense, and savings data to project retirement outcomes. If a user enters post-tax income in one instance and pre-tax income in another, the tool's projections may differ significantly, leading to confusion. A portfolio optimization tool might rely on historical market data. If data from different sources use varying frequency, this inconsistency could skew risk-return calculations. It is important to monitor data consistency when using AI models. To assess whether the data used with an AI model are consistent with the data used during training, users should rely on comprehensive documentation

and monitoring practices. Detailed data documentation is essential, including a schema describing the structure and types of features, information outlining sources and preprocessing steps, and summaries of data distributions to understand the characteristics of the training data. Validation reports should highlight how issues like missing values, outliers, or duplicates were handled. Insights into feature engineering, such as scaling or normalization techniques, help ensure consistent preprocessing. Historical performance metrics on the training and validation datasets can serve as benchmarks to identify deviations in new input data.

- Sensitivities to model inputs. When the results are very sensitive to model inputs, the model is not robust. This is usually considered a sign of the overfitting issue. For LLMs, it is well known that different outputs are generated based on different prompts with similar semantics.
- Outliers. Extreme events may not have been included in the training data. When AI models are used under extreme events, unreasonable results may be produced as the model may not be fitted to the extreme cases. For retirement planning, consider extreme cases like a user planning to retire at 40 with minimal savings or someone requiring a significantly higher post-retirement income due to medical expenses. For investment, a portfolio's historical data might not include black swan events like a market crash. When such events occur, the model might predict overly optimistic outcomes. When validating the models, it is important to look at these outliers individually and evaluate the reasonableness of the results. For example, for an AI-enabled retirement planning tool, test cases can be constructed to reflect extreme situations such as high consumption need and/or market crash.
- Structural change. In a dynamic environment, AI models may not capture the impact of structural changes as they are not contained in the training data. Models using historical data may not be relevant in the future because structural changes, such as new regulatory environments, may lead to the irrelevance of historical data. A change in tax laws or social security benefits could significantly affect retirement outcomes. The emergence of new financial products, such as ESG funds, may render historical data irrelevant. Existing AI models need to be adjusted to reflect structural changes, either through model retraining given sufficient data or adjustment to key model parameters to reflect the expected impact of structural changes. Model outputs need to be monitored against any credible assumptions of structural changes to make sure the model is relevant in an ever-changing environment.
- Model update. Models may be updated for better performance with new data, new methods, new criteria, or more computing power. Zhang et al. (2024) discussed the challenge of aligning LLMs with evolving human preferences over time. They introduced a continual learning approach that dynamically adjusts the model based on new human feedback, preventing the loss of previously learned preferences and ensures that LLMs can adapt to changing human preferences efficiently without the need for complete retraining.

Ensuring a high degree of model robustness is important for building the credibility of AI models and their applications with stakeholders and is necessary for any real-world applications.

In addition to the technical perspective of managing AI model risk, good practices and risk management frameworks are essential to minimizing model risk. These practices are usually not unique to AI but applicable to other model types.

- Establish a strong model governance structure covering model development, training, validation, deployment, review, and reporting. Roles and responsibilities need to be clearly defined and communicated for effective management.
- Develop technical expertise needed for using and explaining AI appropriately. Sufficient training needs to be provided to stakeholders on the best practices for model risk management.

- Ensure that high-quality data is used in the retraining process. This may indicate rigorous data collection and data processing practices. A thorough understanding of the training data is also important to be able to exclude the areas where the models are not supposed to be used.
- Perform systematic validation tests to identify weak areas of the models.
- Establish and implement formal review and monitoring of model performance, data, and assumptions.
- Rely on multiple models rather than one single model to evaluate and select the answers.
- Provide detailed documentation for auditing and to improve the models.
- Ensure clear communication of the model risk with stakeholders to establish reasonable expectations, while recognizing an AI model's full potential.

Managing AI model risk properly is key to successful AI applications and the required efforts cannot be underestimated when evaluating an AI application.

Section 5: AI Illustration

In this section, a series of examples are described to illustrate how AI can be used in the fields of investment and retirement. Each example contains a description of the AI application, its underlying data and models, and how it can be used to help solve the issue. The examples focus on the implementation of AI for real-world problems, rather than the technical details. While some examples include Python code snippets to illustrate concepts, these codes are supplementary and can be ignored by readers who prefer to focus on the broader applications and capabilities of AI. The primary objective of the section is to highlight how AI can be effectively applied to solve investment and retirement-related challenges, rather than delve into technical programming details. The inclusion of code is intended to provide additional clarity and demonstrate practical implementation for those interested, but it is not necessary to engage with the code to understand the concepts or appreciate the potential of AI in these fields.

5.1 ILLUSTRATION OF AI MODEL IMPROVEMENT – RETIREMENT SAVINGS BEHAVIOR COACHING

This example explores how reinforcement learning with human feedback (RLHF) might be applied to guide individuals in their retirement savings decisions. By integrating AI-driven recommendations with qualitative human input, the example illustrates a way to refine financial suggestions over time, potentially improving alignment with savings goals. The process highlights how an AI agent can learn to adapt its advice based on both structured financial data and user responses.

For those interested in AI applications in finance, this example offers insights into how reinforcement learning can be used in a practical, educational setting. It provides a look into the interaction between AI decision-making and human feedback, which could be useful for understanding how reinforcement learning models evolve through experience. Readers exploring financial AI tools may find this case a helpful reference for studying adaptive learning systems in personal finance.

As discussed in [Section 3.3 Process](#), reinforcement learning with human feedback (RLHF) is an advanced AI technique that combines reinforcement learning with real-time feedback from users or domain experts. This hybrid approach enables agents to learn not only from a predefined reward function but also from dynamic feedback provided during training. In this example, we apply RLHF to the domain of retirement savings behavior coaching. The goal is to create an adaptive system that provides personalized and practical savings suggestions to users while incorporating human feedback to refine its decision-making process. While the studied problem may be solved using other approaches, the purpose of this example is to illustrate the idea and process of RLHF.

Savings behavior plays a pivotal role in personal financial management. Many individuals struggle to balance income, expenses, and savings goals effectively. This example simulates an agent who learns optimal savings behavior over time, incorporating both qualitative human feedback and associated quantitative rewards. The example covers the following:

- Optimize the ability to save under diverse income and expense scenarios.
- Evaluate the impact of RLHF on improving the alignment of suggested savings with desired savings behavior.
- Illustrate the learning journey using quantitative metrics and visualizations.

Data

The dataset used is simulated weekly financial scenarios:

- Income: Randomized between \$800 and \$2,000.
- Expenses: Randomized between \$500 and \$1,500.
- Current Savings: Starting savings randomized between \$1,000 and \$5,000.
- Savings Goal: Randomized target savings between \$5,000 and \$10,000.

Table 6

EXAMPLES OF TRAINING DATA FOR SAVINGS BEHAVIOR

State	Income	Expenses	Current Savings	Savings Goal
1	1,926	1,339	2,257	5,262
2	1,660	1,279	1,683	5,623

The dataset comprises 100 unique financial scenarios for the agent to explore and learn.

Environment Design

The environment is structured with the following components:

- State Space: Represents financial scenarios (income, expenses, current savings, and savings goals).
- Action Space: The agent can choose to save between 0% and 100% of income in increments of 10%.
- Reward Mechanism: Rewards align with human-defined goals and savings feasibility.

RLHF Feedback System

“Human” feedback has three types as follows:

- Good Feedback: Saving at least 80% of feasible savings. This turns into a quantitative reward of 30.
- Neutral Feedback: Saving between 20% and 80% of feasible savings. This turns into a quantitative reward of 10.
- Bad Feedback: Saving less than 20% of feasible savings. This turns into a quantitative penalty of 20.

Model Training

To better illustrate the idea of reinforcement learning, the Q table that is composed of a handful of states is used. It stores the expected reward for each state-action pair and guides the agent in selecting optimal actions. Initially, the Q table is filled with zeros. In this example, we have 11 possible actions, from zero savings to 100% savings, using increments of 10%. Here 100% savings means that all available money (income – expense) will be saved. During training, the Q values are updated iteratively based on the reward received and the maximum future Q value for subsequent states, as explained in Section 2.2 Model. Table 7 shows an example of Q values of a sample state, at the beginning and at the end of the model training.

Table 7
EXAMPLES OF Q VALUES OF SAVINGS BEHAVIORS

State Index	1		2	
Income	1,926		1,660	
Expenses	1,339		1,279	
Current savings	2,257		1,683	
Savings goal	5,262		5,623	
	before training	after training	before training	after training
Action 0%	0.00	7.62	0.00	6.15
Action 10%	0.00	7.85	0.00	15.09
Action 20%	0.00	16.26	0.00	8.42
Action 30%	0.00	17.37	0.00	12.20
Action 40%	0.00	0.00	0.00	5.17
Action 50%	0.00	21.21	0.00	20.28
Action 60%	0.00	23.14	0.00	0.00
Action 70%	0.00	74.84	0.00	20.55
Action 80%	0.00	16.67	0.00	73.47
Action 90%	0.00	30.83	0.00	21.88
Action 100%	0.00	5.39	0.00	8.93
Best Action	No preference	Action 70%	No preference	Action 80%

The training is initialized with a Q table of zeros and some hyperparameters including learning rate, discount factor, and so on. For example, for state 1 and state 2, the Q value of each possible action is 0, which leads to no best action before training. In this example, 50 episodes of training are conducted, leading to non-zero Q values. The Best Action row indicates the action with the highest Q value for each state, representing the optimal savings percentage based on the learned policy. For example, the best action for state 2 is to save 80% of net income (income – expenses).

Each episode in the training process is composed of the following steps:

- Initialize the starting state.
- Select an action using an epsilon-greedy policy.
- Calculate the reward based on the savings goal and human feedback.
- Update the Q values in the table.

The training process updates the expected reward for a given state-action pair $Q(s, a)$ based on the immediate reward and the maximum expected future reward using the Q learning equation.

$$Q(s, a) + \alpha \left(r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right)$$

Where:

s : current state.

a : current action.

$Q(s, a)$: The Q value for current state s and action a .

r : The immediate reward received after taking action a in state s .

γ : The discount factor, determining the importance of future rewards.

s' : Next state after taking action a .

a' : Actions that may be taken in next state s' .

$\max_{a'} Q(s', a')$: The maximum Q value for the next state s' , representing the expected future reward.

$r + \gamma \max_{a'} Q(s', a')$: The target value of $Q(s, a)$. The Q value will not be updated when it equals the target value.

α : learning rate. The Q value is adjusted to its target value at a rate of α .

During model training, given a state, an action will be chosen to continue the process of updating the Q values. To be able to quickly learn and find the Q values, different actions need to be tried. In this example, the epsilon-greedy policy is used to balance exploration and exploitation during the training process. By choosing a random action to discover new possibilities and avoid getting stuck in local optima, the decision space can be explored. At the same time, the Q table needs to be exploited to a certain extent so that action can be taken with the highest Q value based on the current Q table. A healthy training process gradually moves from exploration to exploitation. The probability of exploration is ϵ and the probability of exploitation is $1 - \epsilon$.

The epsilon value decreases exponentially after each episode using the formula:

$$\epsilon_t = \max(\epsilon_{min}, \epsilon_{t-1} \times \epsilon_{decay})$$

Where:

ϵ_t : The value of epsilon at step t in a training episode. The initial value of epsilon (ϵ_0) is 0.5 in this example.

ϵ_{min} : The minimum value of epsilon to ensure a baseline level of exploration (0.05 in this example).

ϵ_{decay} : The decay factor that reduces epsilon after each episode (0.995 in this example).

Figure 11
EXAMPLE OF EPSILON DECAY OVER TRAINING EPISODES

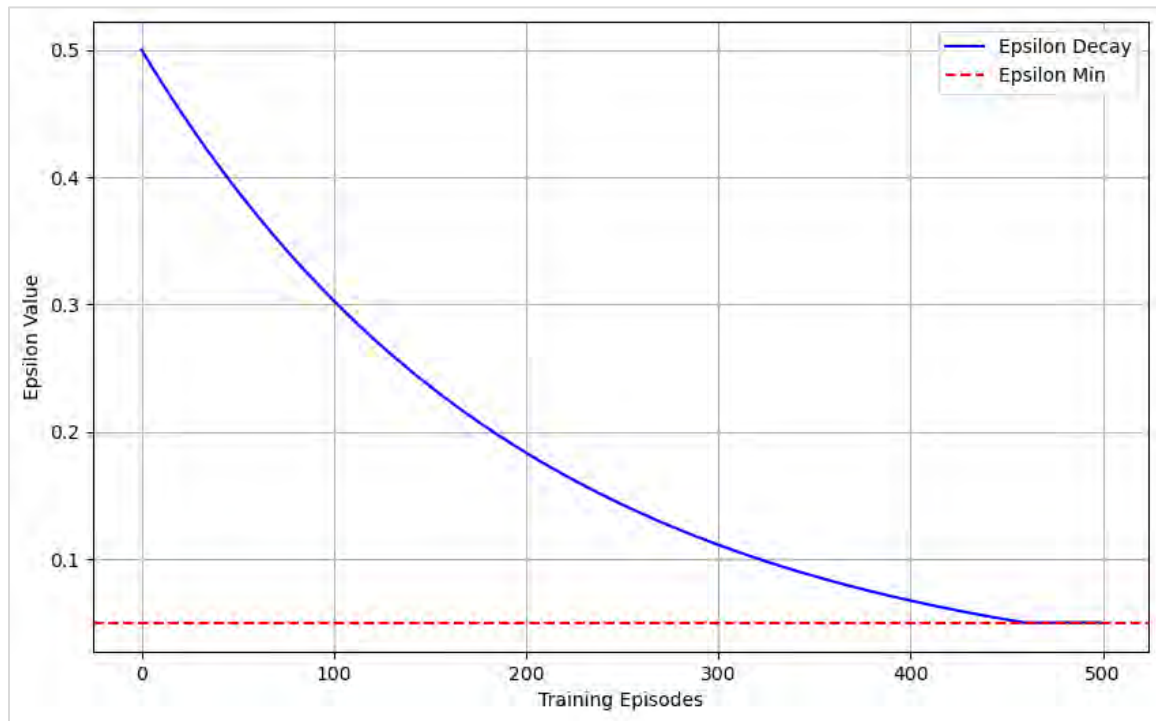
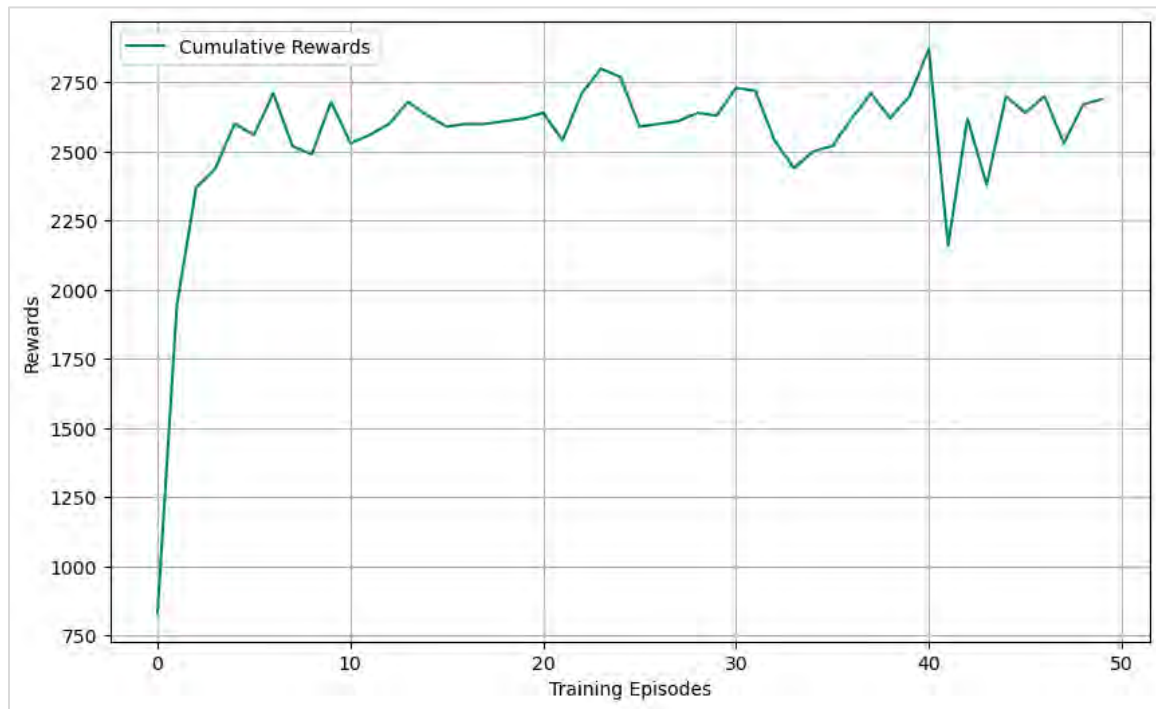


Figure 11 illustrates the evolution of ϵ during training. Early in training, high epsilon ensures the agent explores different actions to discover optimal strategies. As training progresses, lower epsilon encourages the agent to exploit its learned policy by selecting actions with the highest Q values. This balance ensures robust learning and avoids overfitting to suboptimal policies discovered early in training.

Results and Analysis

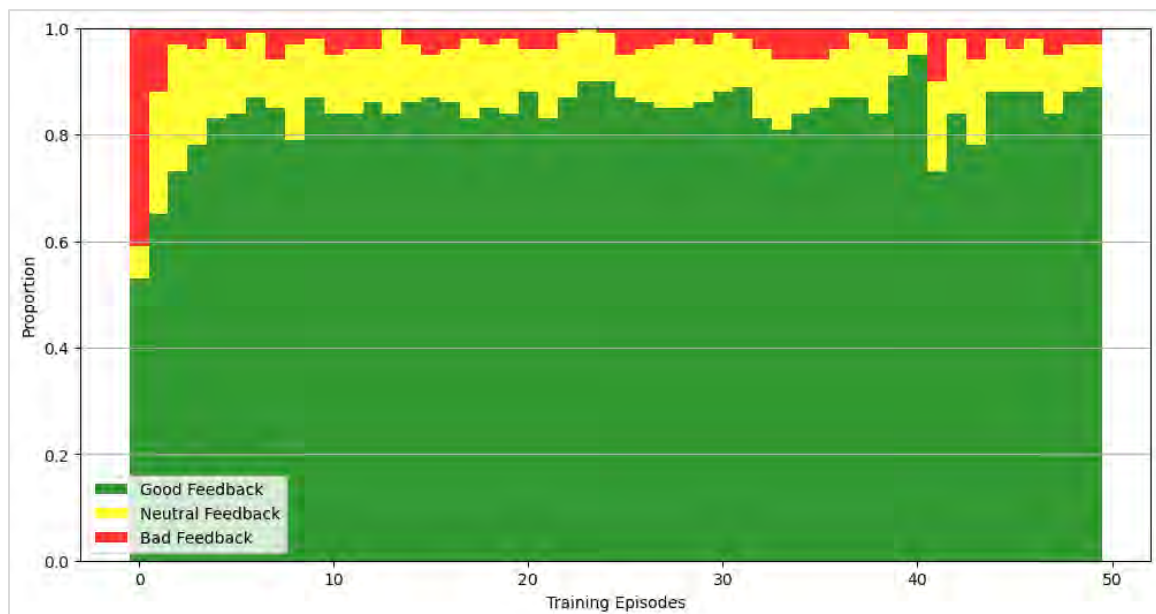
The process repeats until the Q values are stable. Figure 12 shows that the rewards steadily increased over the 50 episodes, reflecting improved alignment with savings goals.

Figure 12
EXAMPLE OF REWARD IMPROVEMENT OVER TRAINING EPISODES



As illustrated in Figure 13, the proportions of Good, Neutral, and Bad feedback types evolved significantly. Initially, the agent received mostly Bad feedback, but as training progressed, the proportion of Good feedback dominated.

Figure 13
EXAMPLE OF FEEDBACK IMPROVEMENT OVER TRAINING EPISODES



This example demonstrates the process and power of RLHF in guiding an agent to optimize savings behavior. By incorporating human feedback, the agent achieved improved alignment with savings goals. At the initial stage, the agent frequently saved less than 20% of feasible savings, leading to a high proportion of Bad feedback. Then the agent began exploring Neutral actions (20-80% of feasible savings), gradually reducing Bad feedback. In the end, the agent predominantly saved at least 80% of feasible savings, earning Good feedback in most scenarios.

It is noted that, in this example, human feedback is rule-based with an explicit reward function while, in real cases, more complicated rules and reward functions can be learnt from real human feedback using deep learning models for more complex financial scenarios. However, the process of reinforcement learning remains the same.

The Python script (<https://github.com/Society-of-actuaries-research-institute/FP208-Artificial-Intelligence-in-Investment-and-Retirement>) used to illustrate the process of RLHF is also made available in the format of a Jupyter Notebook for educational purposes.

5.2 INVESTMENT CHATBOT USING RETRIEVAL-AUGMENTED GENERATION

This example explores how Retrieval-Augmented Generation (RAG) can be applied to building an investment chatbot that provides contextually relevant and accurate financial insights. By combining a retrieval system with a language model, the chatbot ensures that responses are grounded in a well-defined knowledge base, reducing the risk of misinformation. The process illustrates how RAG leverages a semantic search to retrieve relevant financial documents and then generates human-like responses tailored to user queries.

For those interested in AI-driven tools, this example provides an educational look at integrating retrieval-based searches with generative models. It highlights how a chatbot's performance can be evaluated using context and answer relevance metrics, offering insights into the refinement of AI responses. Readers studying conversational AI systems may find this case useful in understanding how structured information retrieval enhances the reliability of automated financial assistants.

LLMs can be fine-tuned with domain specific knowledge to better serve target users. Examples include BloombergGPT and FinGPT that provide complex financial data analysis and generate insights and research for investment decisions. Since suggestions on investment and retirement can have a real financial impact, it is important that the advice from AI has no hallucinations and is purely based on some defined knowledge base. Retrieval-Augmented Generation (RAG) is a potential choice that blends natural language processing (NLP) with information retrieval techniques, making it highly effective for applications like investment chatbots. It combines an LLM with a retrieval system to ensure responses are accurate, contextually relevant, and grounded in up-to-date information.

In an investment chatbot, RAG relies on three primary components: a retrieval module, a language generation module, and a knowledge base. The retrieval module fetches relevant information from a curated repository, which may include financial reports, market data, news articles, and investment guidelines. Using advanced search techniques like vector-based matching, it identifies the most relevant documents or data points for a given query. Meanwhile, the language generation module, powered by a pre-trained LLM, takes both the user's query and the retrieved content to generate responses that are

contextually aware and human-like. The knowledge base itself acts as the foundation, storing dynamic and regularly updated information about historical trends, portfolio strategies, and regulatory updates.

When a user interacts with the chatbot, the process begins with the user submitting an investment-related query, such as, "What are the current trends in the tech sector?" The chatbot encodes the question into a vector representation using an embedding model, enabling the retrieval module to search the knowledge base for the most relevant documents. These retrieved documents, such as recent tech sector reports, are then passed to the language generation module, which synthesizes a precise and coherent response. For example, the chatbot might reply with something like, "The tech sector is experiencing growth in AI and cloud computing, with companies like XYZ reporting strong quarterly earnings." This response is delivered in a user-friendly format, potentially including graphs or links for further exploration.

RAG enhances investment chatbots by ensuring they provide accurate, context-aware, and dynamic responses. This is particularly valuable in the high-stakes world of financial decision-making, where outdated or speculative information can lead to poor outcomes. By grounding responses in reliable sources, RAG minimizes the risk of generating incorrect or misleading answers.

The use of RAG also allows investment chatbots to handle a wide range of queries, from beginner-level questions to more advanced investment strategies. For instance, a user might ask about market trends, seek advice on portfolio diversification, or request explanations of financial concepts like compound interest. The chatbot can even provide real-time alerts, such as notifications about significant stock price movements in a user's portfolio.

In this example, we demonstrate the process of building a chatbot using the RAG method. This involves combining semantic search capabilities with generative language modeling to provide contextually relevant and coherent responses based on a set of indexed documents. The Python script (<https://github.com/Society-of-actuaries-research-institute/FP208-Artificial-Intelligence-in-Investment-and-Retirement>) used to build and evaluate the chatbot is also made available in the format of a Jupyter Notebook for educational purposes. A research paper is used as the knowledge base of the chatbot. Rudolf (2023) studied the complex relationships between economic growth and its key drivers including demographics and productivity. The foundation of the chatbot relies on this paper and LlamaIndex⁷, which serves as the intermediary between a language model and a structured or unstructured knowledge base. It facilitates the organization of documents into an index, leveraging semantic embeddings to represent text in high-dimensional vector spaces. These embeddings capture the meaning of text, enabling a similarity-based search even when exact keywords are not present.

Once the documents are indexed, a query engine is created. This engine processes user queries by transforming them into embeddings, comparing them with the stored document embeddings, and retrieving the most semantically similar chunks of text. The retrieved information is passed to a pre-trained generative language model, such as GPT 3.5 Turbo in this case, which synthesizes a human-readable and context-aware response. This combination of retrieval and generation forms the core of the RAG workflow.

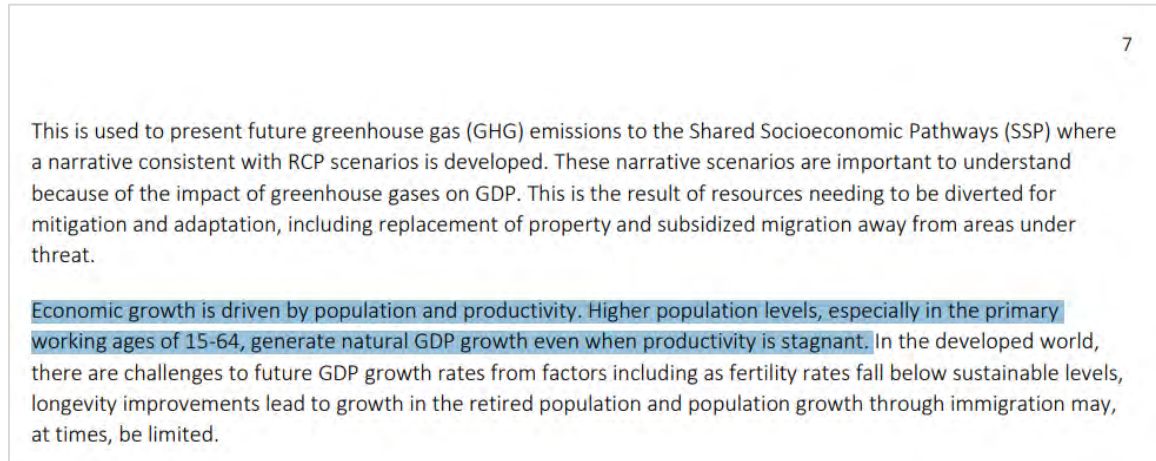
A conversion between the generated query engine is given below. The question is the top question suggested by GPT 3.5 Turbo regarding GDP growth drivers.

Question: How does the age distribution of the population impact economic growth?

AI Answer: Higher population levels, especially in the primary working ages of 15-64, generate natural GDP growth even when productivity is stagnant.

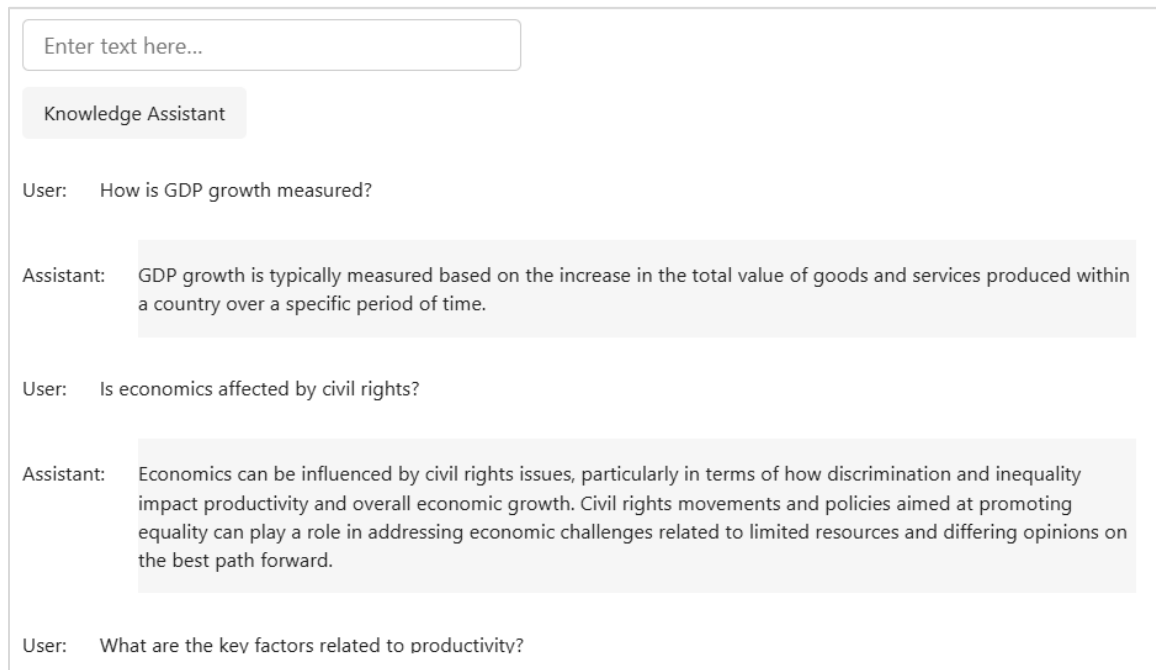
When compared to the paper, the answer provided by the chatbot can be found in the paper, as shown in Figure 14.

Figure 14
SCREENSHOT OF RUDOLF (2023) PAPER



This ensures that the answers are based on the defined knowledge base and can reduce the chance of getting hallucinated responses. With the query engine ready, a simple user interface is built using Python Panel⁸, as illustrated in Figure 15. In this interface, users can type in questions and click button "Knowledge Assistant" to retrieve answers. The history of the conversation is also kept.

Figure 15
CHATBOT USER INTERFACE



To evaluate and monitor the chatbot's performance, TruLens⁹ is used. TruLens integrates seamlessly with LlamaIndex to record interactions, such as queries and responses, and provides tools for assessing the quality of responses. In the evaluation performed by TruLens, two critical aspects were examined in detail: context relevance and answer relevance.

- Context relevance evaluates how well the retrieved chunks of information from the knowledge base align with the user's query. It measures the effectiveness of the retrieval mechanism, powered by LlamaIndex, in selecting the most semantically and contextually appropriate pieces of information to address the query. The quality of the retrieved content is paramount in a RAG-based chatbot because the responses generated by the language model depend directly on the relevance of the retrieved data. In the example, context relevance is determined by comparing the embedding of the user's query with the embeddings of the retrieved document chunks. These embeddings represent the semantic meaning of the text and are compared using similarity metrics such as cosine similarity. A higher similarity score indicates a stronger alignment between the retrieved content and the query.
- Answer relevance evaluates the quality of the final responses generated by the chatbot. It focuses on whether the response addresses the user's query accurately, coherently, and meaningfully. While context relevance assesses the input provided to the language model, answer relevance examines the output. In the example, this is achieved by analyzing the chatbot's responses to predefined evaluation questions.

Based on the five evaluation questions given below, the chatbot received a score of 0.33 for context relevance and a score of 1 for answer relevance.

Input `Select.RecordInput`

How do technological advancements interact with demographic factors to influence productivity and economic growth trajectories?

Output `Select.RecordOutput`

Technological advancements interact with demographic factors by playing a crucial role in boosting productivity and influencing economic growth trajectories. These advancements can enhance efficiency, streamline processes, and create new opportunities, leading to increased productivity levels. Moreover, technological progress can help address demographic challenges such as aging populations or labor shortages by enabling automation, skill development, and innovation. This interaction between technology and demographics can contribute to sustainable economic growth by driving productivity gains and adapting to changing population dynamics.

Input `Select.RecordInput`

In what ways do immigration patterns contribute to or detract from the productivity and economic growth of a nation?

Output `Select.RecordOutput`

Immigration patterns can contribute to the productivity and economic growth of a nation by bringing in a younger workforce that can help offset the challenges posed by an aging population. This influx of working-age individuals can boost productivity levels and contribute to natural GDP growth. However, if immigration levels are limited, it can lead to challenges in sustaining economic growth, especially in countries where fertility rates are falling and longevity improvements are increasing the retired population.

Input `Select.RecordInput`

How do educational attainment and skill levels within a population affect productivity and, consequently, GDP growth?

Output `Select.RecordOutput`

Educational attainment and skill levels within a population can positively impact productivity by enhancing the quality of the workforce. This, in turn, can lead to increased efficiency, innovation, and overall economic output. As productivity is a key component of GDP growth, higher educational attainment and skill levels can contribute to a stronger economy.

Input `Select.RecordInput`

What role does labor force participation play in influencing productivity and economic growth across various demographic groups?

Output `Select.RecordOutput`

Labor force participation plays a significant role in influencing productivity and economic growth across various demographic groups. It affects productivity by determining the number of individuals actively contributing to the economy through work. Higher labor force participation rates can lead to increased output per person in the workforce, which in turn can boost overall economic growth. Additionally, labor force participation across different demographic groups can impact the relative size of the workforce compared to the total population, as well as the balance between the working and retired populations. Strategies to increase hours worked for underemployed workers and disadvantaged communities can help enhance productivity and economic growth by utilizing the potential of a larger portion of the population.

Input `Select.RecordInput`

How do demographic changes, such as aging populations or shifts in birth rates, impact overall GDP growth rates in different countries?

Output `Select.RecordOutput`

Demographic changes, such as aging populations or shifts in birth rates, can have a significant impact on overall GDP growth rates in different countries. These changes can affect the size and composition of the workforce, productivity levels, consumer spending patterns, and government spending on social programs. As populations age and fertility rates decline, there may be fewer people in the workforce contributing to economic output, which can lead to slower GDP growth. Additionally, an aging population may increase healthcare and pension costs, putting pressure on government budgets and potentially reducing overall economic growth. Shifts in birth rates can also impact future workforce size and productivity levels, influencing GDP growth rates in the long term.

Both context and answer relevance work together to determine the overall effectiveness of the chatbot. High-quality retrieval ensures that the language model has the necessary context to generate accurate and meaningful responses. Conversely, even if the retrieved context is highly relevant, poor performance by the language model, such as generating off-topic or factually incorrect responses, can diminish the chatbot's reliability.

This example demonstrates the construction of a RAG-based chatbot system that uses LlamaIndex for document indexing and retrieval, and TruLens for evaluation and monitoring. This combination ensures

that the chatbot delivers contextually rich responses while providing tools for continuous improvement. It also showed that, with available open-source tools, it is not that difficult to integrate a semantic search, generative modeling, and evaluation frameworks to build intelligent conversational systems for real-world applications.

It is also noted that the chatbot can be further improved, although it is not the purpose of the example. The quality of the LLM is essential for ensuring that the responses are coherent and contextually appropriate. The design of the prompt fed to the language model is critical in guiding the generation process and maintaining relevance. In addition to context and answer relevance, evaluations can include automated checks against predefined answers or involve human reviewers scoring the chatbot's output. This example only used window-sentence retrieval, a technique that segments a document into smaller units, such as individual sentences or fixed-length text windows, and indexes these units for retrieval. Given the relatively low score on context relevance, we may use auto-merging retrieval to complement window-sentence retrieval by dynamically combining multiple retrieved segments into a single coherent response. This method is critical when relevant information is dispersed across multiple chunks or documents and no single segment sufficiently addresses the query.

5.3 SOCIAL MEDIA DATA SUMMARIZATION

This example examines how AI can be applied to summarizing and analyzing large volumes of social media data, specifically Reddit comments, to explore potential links with market volatility. By leveraging natural language processing techniques, AI helps transform unstructured text into structured data, making it easier to extract insights such as sentiment trends and keyword frequencies. The study highlights how retail investors, especially during the pandemic, influenced stock market movements, and how social media discussions may provide signals related to market dynamics.

For those interested in AI applications in the investment field, this example offers an educational perspective on using nontraditional data sources in investment analysis. It illustrates how AI-driven sentiment analysis and word frequency tracking can contribute to understanding market behavior. Readers exploring AI in data analytics or social media analysis may find this case useful for studying how large-scale textual data can be processed and integrated into decision-making frameworks.

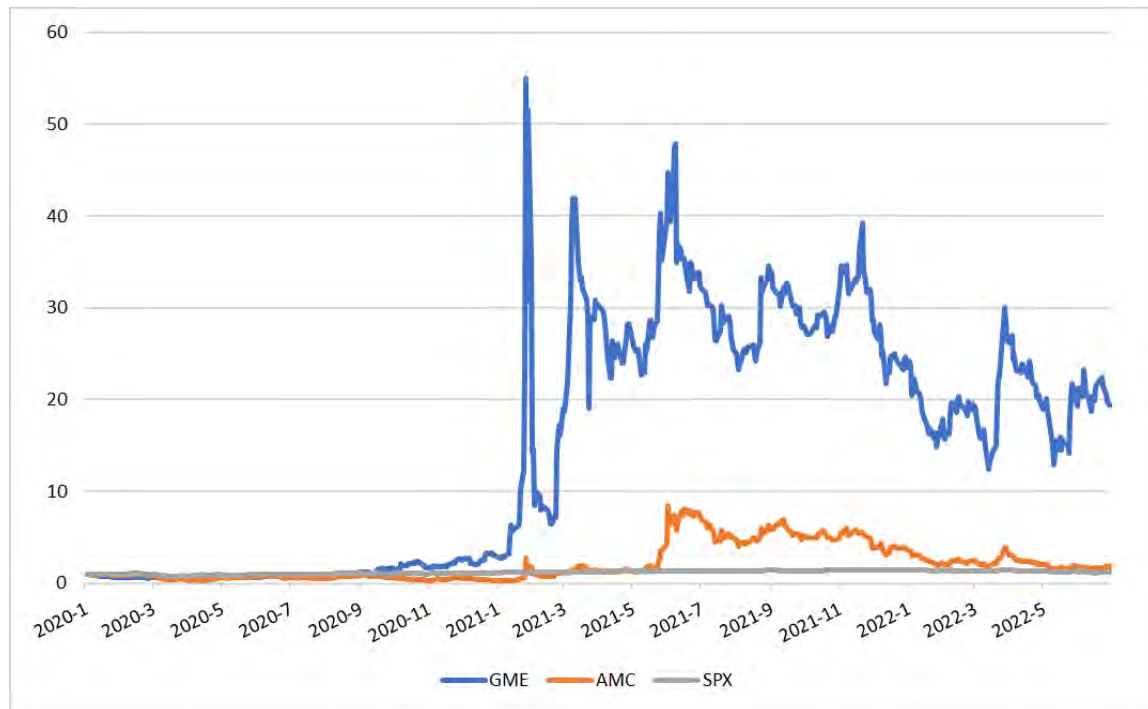
Social media data, when used properly, can help provide additional information not available in existing structured data. On the other hand, it is extremely time-consuming and difficult to process manually due to its volume and variation. Natural language processing techniques can be used to transform non-structured text data into structured data. Other AI models, such as deep-learning models, can be used to further extract information, such as sentiment, that may be used to predict market dynamics. In this example, as detailed in Shang (2023), AI is used to process Reddit data and extract features that may be used to predict market volatility.

During the pandemic, retail investors had an increasing influence on the public equity market, with the retail investors' share of trading volume climbing by 10% to 25%. People have more time or opportunities to monitor the market while working from home. Tech companies, such as Robinhood and WealthSimple, made the trading convenient with low commissions. Social media made it easier to retrieve and share information. With less public life, people without financial distress may also have a higher risk appetite to participate in activities such as stock trading.

Retail investors can cause extremely high volatility in single-stock trading. During the pandemic, meme stocks, such as GME and AMC, experienced much higher volatility than the general market. Meme stocks are shares of companies that gain popularity through social media or online forums, leading to rapid price increases driven by retail investor enthusiasm rather than fundamentals. Figure 16 shows the daily prices of GME, AMC and SPX, which all normalized to 1 at the beginning of the study period. While SPX has an annualized daily volatility of 25.5%, GME has a volatility of 194.0% and AMC has a volatility of 246.8%.

Figure 16

DAILY PRICE HISTORY OF GME, AMC AND SPX (JANUARY 2020 TO JUNE 2022)

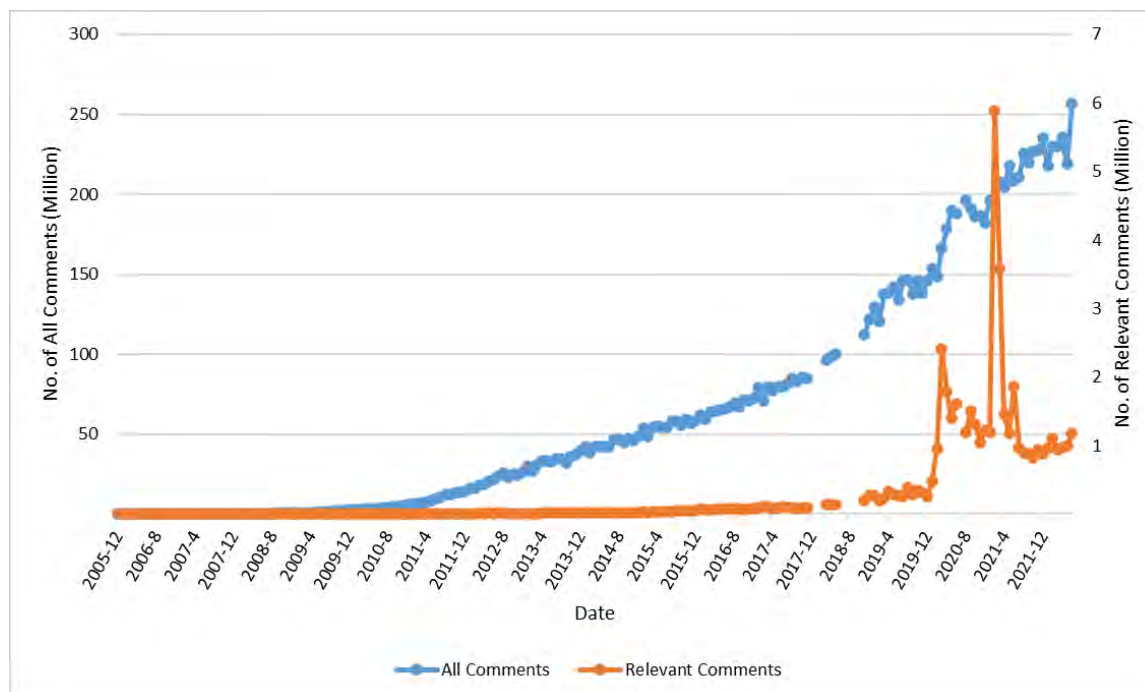


Note: Prices were normalized to 1 at the beginning of 2020.

Retail investors came up with a complicated investment strategy to take advantage of the large short positions on the meme stocks held by hedge funds. By purchasing the stocks in a coordinated way and out-of-the-money call options, the stock price enters into an upward spiral with the short position holders buying the stocks as well to offset their short positions. Although the market cap of these two meme stocks is not material compared to the entire public equity market (around 0.15% of S&P 500 market cap at its highest share), it may indicate the increasing impact of retail investors less significantly for individual stocks but covering a greater number of stocks more significantly. It is, therefore, valuable to incorporate data related to retail investors' behaviors into investment analysis.

To assess the impact of the active participation of retail investors on the general market volatility, Reddit data from December 2005 to June 2022 was used. Reddit is a network of communities where people have common interests and share information. In total, 1.19TB of data is processed. Figure 17 shows the increasing amount of all available Reddit data, with a continuous upward trend.

Figure 17
MONTHLY REDDIT COMMENTS (DECEMBER 2005 TO JUNE 2022)



In addition to the total number of comments in the Reddit dataset, the number of relevant comments is also shown in Figure 17. Reddit has a few big communities focusing on investing, such as wallstreetbets, investing, ETFs, and StockMarket. It is also the place where retail investors coordinated to discuss the strategies for the meme stocks. Comments in these four subreddits are considered relevant and used to analyze the relationship between Reddit comments and market volatility. The largest spike in February 2021 happened concurrently with the price spikes of the meme stocks shown in Figure 16.

The comments are cleansed and summarized to be used together with the market data. On a daily basis, the following explanatory variables are constructed:

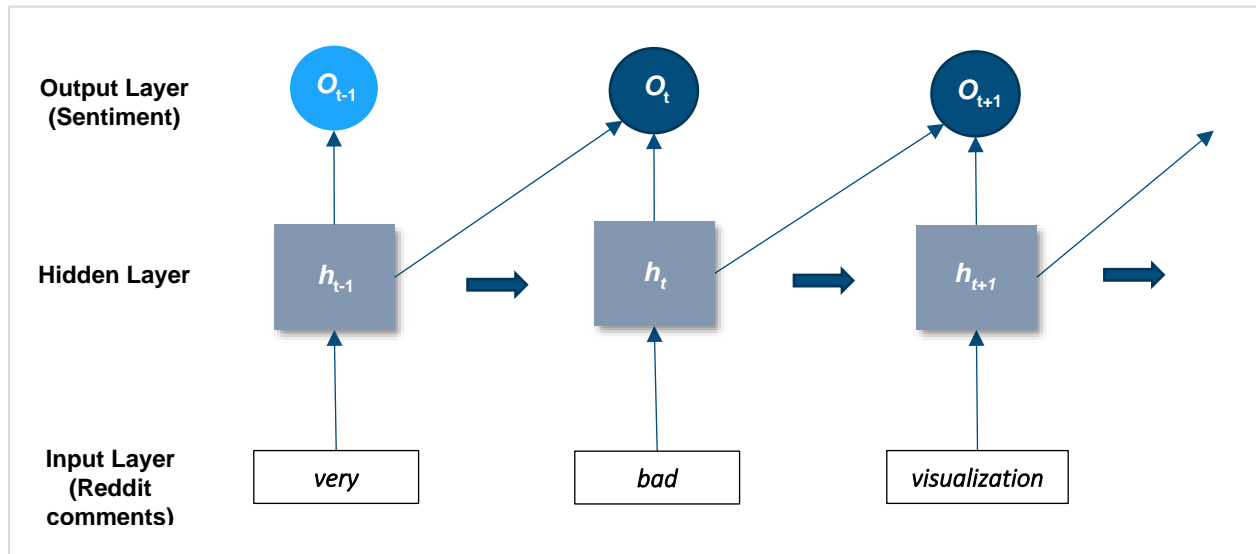
- Number of authors who published a comment
- Number of comments
- Average score of comments. Each comment is assigned a score by Reddit based on the difference between the number of upvotes and downvotes. It is an indication of the comment's popularity.
- Average sentiment of comments. Each comment is analyzed and assigned a sentiment score. The scores are then averaged each day.
- Weighted average sentiment of comments. The sentiment score is averaged with the Reddit score as the weight.
- Frequency of keywords. Word frequency is computed, and the daily frequency of the top most frequent words are used to examine the potential impact on market volatility

Some statistics, such as the number of comments, can be easily handled by traditional approaches. Other explanatory variables, such as average sentiment of comments, need natural language processing and AI models to extract useful information for each individual comment.

The sentiment of relevant comments is helpful for understanding their relationship with market volatility. Each Reddit comment is assigned a sentiment score, with the daily sentiment calculated using either simple

arithmetic averaging or weighted averaging with Reddit scores. Pre-trained sentiment analyzers include ones such as Stanford CoreNLP and Valence Aware Dictionary and Sentiment Reasoner (VADER). As described in Sorcher et al. (2013), the Stanford CoreNLP is based on a recursive neural tensor network (RNTN), which is a variation of the standard recursive neural network (RNN) as shown in Figure 18.

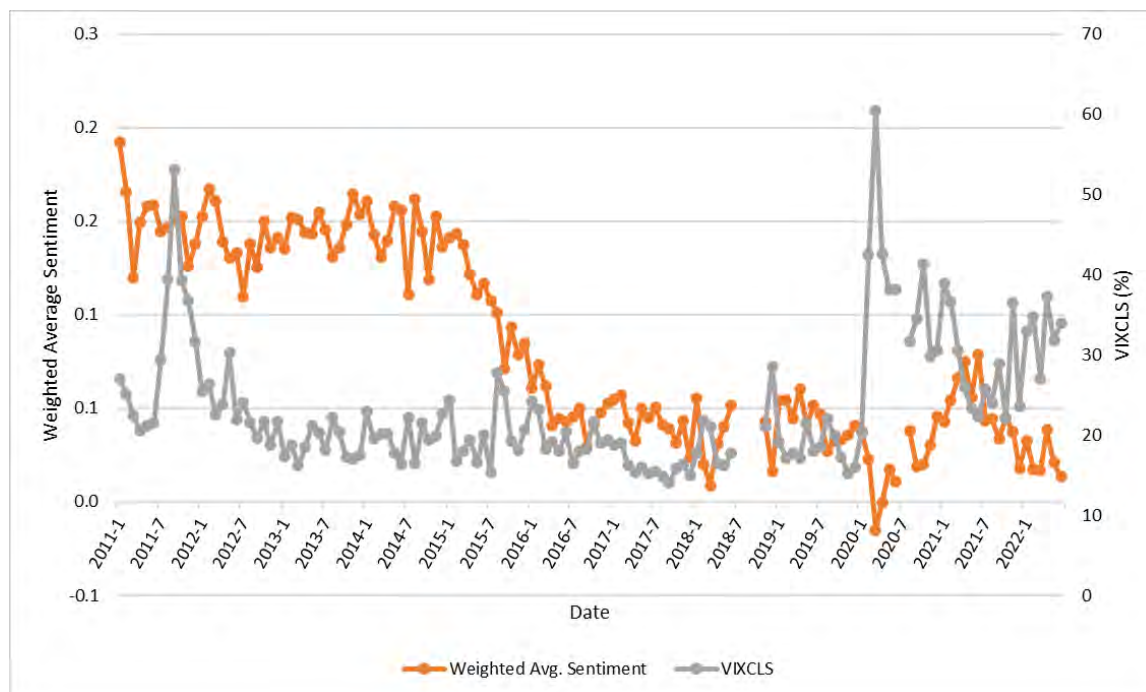
Figure 18
RNN STRUCTURE FOR SENTIMENT ANALYSIS



With the relevant Reddit comments supplied into the model in the input layer, the model uses the hidden layer (h) to assess the sentiment of each comment step by step by incorporating an additional word at each point. Each word is converted into a numerical vector that represents its semantic and syntactic relationships in the context that it is being used before it is fed into the hidden layer. The conversion from words to vectors usually involves the training of neural network models, as well to predict surrounding words accurately. The updated hidden state (h_t) is determined by not only the new word but also the previous hidden state (h_{t-1}). In the variation from standard RNN, RNTN added more complexity in the hidden layer with tensor layers to allow more interactions between numerical vectors that represent words. Instead of using linear transformations with two dimensions, tensors (arrays with three or more dimensions) enable multi-way interactions between input vectors. With tensor layers, RNTN usually excels in tasks that require understanding hierarchical or compositional relationships, such as parsing sentences, sentiment analysis, or tasks requiring a deep understanding of language structure.

The final sentiment in the output layer (O) is determined after the entire comment is processed. In this example, sentiment analyzers provide scores that a comment is positive, neutral, or negative. The score is used to assess the aggregate sentiment of each Reddit comment. Figure 19 shows the weighted average monthly sentiment and VIXCLS, the CBOE volatility index that represents the implied volatility in the S&P 500 Index options over the next 12 months. A much stronger negative correlation is observed during the recent pandemic.

Figure 19
REDDIT COMMENT SENTIMENT VS. VIXCLS (JANUARY 2011 TO JUNE 2022)



- Note:** Weighted average sentiment uses Reddit score as the weight to calculate average sentiment of relevant comments.

To assess any potential predictive power of Reddit comments and any structural changes during recent pandemics, cross correlations are examined in Table 8.

Table 8
CROSS CORRELATION BETWEEN REDDIT COMMENT SUMMARY VARIABLES AND VIXCLS

Variable	1-month lag ¹		2-week Lag		1-week lag		1-day lag		same day	
	Pandemic ⁴	All ⁵	Pandemic	All	Pandemic	All	Pandemic	All	Pandemic	All
No. of comments ²	0.02	0.27	0.05	0.29	0.09	0.31	0.13	0.33	0.15	0.34
No. of authors	0.02	0.30	0.03	0.31	0.06	0.33	0.10	0.34	0.12	0.35
Sentiment ³	-0.22	-0.07	-0.42	-0.11	-0.49	-0.12	-0.51	-0.14	-0.52	-0.14

Notes:

- The lag is the lag between the explanatory variables and the target variable VIXCLS. Explanatory variables precede the target variable to evaluate potential predicting power.
- No. of relevant comments.
- Weighted average sentiment using Reddit scores as the weight.
- The pandemic period starts from January 2020 to June 2022.
- The "All" period starts from January 2011 to June 2022.
- Daily data is used.

It is interesting to know that the cross correlations changed significantly during the recent pandemic, with movements in both magnitude and direction. The number of relevant comments and authors are less correlated with the market volatilities. It may indicate that people are participating in the investment discussion, not solely due to extreme events, but as a long-term commitment to the investment community. The sentiment of comments has a more negative correlation with the market volatility during

Table 9

CROSS CORRELATION BETWEEN KEYWORDS AND VIXCLS (TOP 10 DURING RECENT PANDEMIC)

1-month lag ¹		2-week Lag		1-week lag		3-day lag		1-day lag		same day	
Word	Correl ²	Word	Correl	Word	Correl	Word	Correl	Word	Correl	Word	Correl
oil	0.42	fed	0.51	put	0.63	put	0.71	put	0.78	put	0.82
futur	0.36	bear	0.51	spi	0.63	spi	0.69	spi	0.75	spi	0.79
bear	0.34	futur	0.49	fed	0.61	futur	0.65	futur	0.70	bottom	0.73
call	0.29	spi	0.48	futur	0.60	fed	0.62	bottom	0.68	futur	0.73
fed	0.28	put	0.46	bear	0.58	home	0.61	bull	0.68	bull	0.71
home	0.27	home	0.46	home	0.57	bear	0.61	bear	0.67	fed	0.68
spi	0.25	bull	0.43	bull	0.54	bottom	0.60	fed	0.65	bear	0.68
bull	0.23	recess	0.40	bottom	0.52	bull	0.59	home	0.64	home	0.67
red	0.22	oil	0.38	recess	0.52	recess	0.57	recess	0.63	recess	0.67
dump	0.20	pump	0.37	pump	0.48	pump	0.55	pump	0.61	pump	0.65

Notes:

1. The lag is the lag between the explanatory variables (word frequency) and the target variable VIXCLS. Explanatory variables precede the target variable to evaluate potential predicting power.
2. Correlation coefficient.
3. Daily data from January 2020 to June 2022 is used.

This example demonstrated the application of AI to summarize and analyze large volumes of non-numerical data. With valuable information extracted from social media data using AI, nontraditional data can be used to improve our decision-making.

5.4 AI RETIREMENT PLANNING TOOL

This example explores how AI can be leveraged to assist private investors in portfolio selection by automating key steps in the retirement planning process. By using GPT-4o, the system extracts insights from academic research on portfolio optimization, formulates an investment strategy, and generates Python scripts to implement the calculations. The AI model processes financial data, applies utility functions, and simulates wealth outcomes to identify optimal portfolio allocations based on risk preferences and investment horizons. Additionally, a React-based user interface is developed to provide investors with an intuitive way to input assumptions and receive personalized recommendations.

For those interested in AI-driven financial planning, this example provides an educational perspective on integrating large language models with investment analysis. It demonstrates how AI can streamline data processing, automate financial modeling, and enhance user interactions in retirement planning. Readers studying AI applications in behavioral economics, or user-centered investment tools may find this case a useful reference for understanding how technology can support more informed financial decision-making.

In this example, a retirement planning tool is created, including knowledge study, analysis, model development and user interface design. AI is leveraged heavily to speed up the process. Using GPT-4o, the LLM is asked to study a paper on portfolio formulation. Warren (2018) used utility functions to identify optimal portfolios that are suitable for investors with long investment, including private investors and institutional investors.

Prompt #1: As an investment expert, can you read the paper and summarize the steps for formulating portfolios?

Response from GPT-4o

The paper outlines a four-step approach to formulating investment portfolios using utility functions. Here's a summary:

1. Choose a Utility Function

- Select and parameterize a utility function tailored to the investor's objectives and preferences.
- Common options include **power utility** (for wealth or consumption-level concerns) and **reference-dependent utility** (suitable for target-based evaluations).
- Ensure the utility function aligns with the investor's goals, such as aversion to loss or preference for gains.

2. Specify a Plan for Withdrawal and Evaluation

- Define the withdrawal strategy (e.g., percentage of wealth, fixed withdrawals) over the investment horizon.
- Establish an evaluation plan, considering utility from both withdrawals and terminal portfolio value.
- Determine the investment horizon and account for aspects like income generation or wealth preservation.

3. Project Joint Distributions of Wealth Outcomes

- Model the distribution of potential wealth outcomes for various assets over the investment horizon.
- Incorporate correlations between assets and account for reinvestment of cash flows like dividends or interest.
- Use simulation, historical data, or scenario analysis to project accumulated wealth for candidate portfolios.

4. Solve for Optimal Portfolio Weights

- Generate portfolio values across simulated scenarios, accounting for withdrawals and changes in wealth.
- Evaluate utility for each scenario and aggregate to estimate expected utility.
- Identify the portfolio mix that maximizes expected utility, reflecting the investor's risk-return trade-offs.

The approach is structured to balance practical application with analytical rigor, enabling its use in settings such as retirement planning, endowment funds, or institutional investment. If you'd like further details on any step or examples for specific investor types, let me know!

Comment

The LLM can correctly digest the paper containing complicated logic and extract key and relevant information.

Prompt #2: Can you customize the plan based on the example for private investors in this paper? Please also list the assumptions used for each step and the result.

Response from GPT-4o

Here's a customized plan based on the example of **private investors** described in the paper. The plan focuses on optimizing a portfolio for wealth accumulation over specific horizons.

Customized Plan for Private Investors

Objective:

To maximize terminal wealth over investment horizons of 3 years and 10 years.

Step 1: Choose a Utility Function

- **Utility Function:** Power utility function

$$U(W) = \frac{W^{1-\text{CRRA}}}{1-\text{CRRA}}, \text{ if CRRA} \neq 1$$

or

$$U(W) = \ln(W), \text{ if CRRA} = 1$$

- **Assumptions:**
 - Risk aversion coefficient CRRA is varied (e.g., 2, 4, and 6) to assess sensitivity to risk preferences.
 - Utility is only evaluated based on terminal wealth, with no interim withdrawals.

Step 2: Specify a Plan for Withdrawal and Evaluation

- **Withdrawal Plan:** None; the investor does not withdraw funds during the investment horizon.
 - **Evaluation Plan:**
 - Evaluate utility of terminal portfolio value at the end of the horizon (3 or 10 years).
 - No additional streams of income or consumption are considered.
 - **Assumptions:**
 - Terminal wealth is the sole determinant of utility.
 - Investment horizon is strictly adhered to (3 years or 10 years).
-

Step 3: Project Joint Distributions of Wealth Outcomes

- **Methodology:** Simulate wealth outcomes using historical return distributions or statistical models.
 - **Assumptions:**
 - **Equities:** Expected return = 6%, standard deviation = 16%.
 - **Bonds:** Expected return = 1%, standard deviation = 8%.
 - Returns are normally distributed and independent across years (iid assumption).
 - Cash flows from investments (e.g., dividends, coupons) are fully reinvested.
-

Step 4: Solve for Optimal Portfolio Weights

- **Optimization Process:**
 - Simulate wealth outcomes for different portfolio allocations (e.g., equity vs. bonds).
 - Evaluate the utility for each scenario at the end of the horizon.
 - Compute expected utility across all simulated scenarios.
 - Identify the asset allocation that maximizes expected utility.
 - **Assumptions:**
 - Static portfolio weights (no rebalancing or dynamic adjustments).
 - Sufficient simulation paths (e.g., 10,000 scenarios) to ensure statistical robustness.
-

Results:

- **For Short Horizon (3 years):**
 - Higher risk aversion (CRRA=6) leads to more bond-heavy portfolios.
 - Lower risk aversion (CRRA=2) favors equity-heavy allocations to capitalize on potential higher returns.
 - **For Long Horizon (10 years):**
 - With increasing horizon, equities become more favorable due to compounding effects and reduced likelihood of loss.
 - Even highly risk-averse investors (CRRA=6) allocate a higher portion to equities compared to the 3-year horizon.
-

This approach enables private investors to tailor their portfolios to both their risk tolerance and investment horizon, while focusing solely on terminal wealth accumulation. Let me know if you'd like additional details or a deeper exploration of any assumptions!

Comment

The LLM can integrate knowledge from different sections in the paper correctly. While presenting the example of private investors, the paper does not list the steps in detail and the formula of the utility function used but rather provides a high-level description and visualization of the results. The LLM can articulate the example following the prescribed steps but embedding the details of that example such as the formula of the utility function. It can also provide observations of the results that are presented as numerical values in a table and visualization.

Prompt #3: *Based on the steps for private investors, can you write Python program to implement the calculations in each step?*

Response from GPT-4o

Here's a Python program that implements the steps for private investors, including simulation of wealth distributions, calculation of utility scores, and identification of optimal portfolio weights. This implementation assumes a Monte Carlo simulation to project wealth outcomes and optimize portfolio weights.

```

import numpy as np
import matplotlib.pyplot as plt

# Step 1: Define Utility Function
def power_utility(wealth, crra):
    """
    Calculate utility using power utility function.
     $U = W^{1-CRRA} / (1 - CRRA)$  for  $CRRA \neq 1$ 
     $U = \ln(W)$  for  $CRRA = 1$ 
    """
    if crra == 1:
        return np.log(wealth)
    else:
        return (wealth ** (1 - crra)) / (1 - crra)

# Step 2: Parameters for Withdrawal and Evaluation
investment_horizons = [3, 10] # Years
initial_wealth = 100000 # Initial investment in dollars
crra = 4 # Coefficient of relative risk aversion

# Step 3: Simulation Parameters
n_simulations = 10000
equity_return_mean = 0.06
equity_return_std = 0.16
bond_return_mean = 0.01
bond_return_std = 0.08
correlation = 0.3

# Generate correlated returns for equities and bonds
def simulate_returns(horizon):
    cov_matrix = [
        [equity_return_std ** 2, correlation * equity_return_std * bond_return_std],
        [correlation * equity_return_std * bond_return_std, bond_return_std ** 2]
    ]
    mean_returns = [equity_return_mean, bond_return_mean]
    return np.random.multivariate_normal(mean_returns, cov_matrix, (n_simulations, horizon))

# Step 4: Solve for Optimal Portfolio Weights
def optimize_portfolio(horizon):
    results = []
    weights = np.linspace(0, 1, 101) # Portfolio weights from 0% to 100% equities

    # Simulate wealth outcomes for each weight
    for weight in weights:
        simulations = simulate_returns(horizon)
        portfolio_returns = weight * simulations[:, :, 0] + (1 - weight) * simulations[:, :, 1]
        terminal_wealth = initial_wealth * np.exp(portfolio_returns.sum(axis=1))
        utility_scores = power_utility(terminal_wealth, crra)
        expected_utility = np.mean(utility_scores)
        results.append((weight, expected_utility))

    # Find the weight with maximum expected utility
    optimal_weight, max_utility = max(results, key=lambda x: x[1])
    return optimal_weight, results

# Run the optimization for each investment horizon
for horizon in investment_horizons:
    optimal_weight, results = optimize_portfolio(horizon)
    weights, utilities = zip(*results)
    print(f"Horizon: {horizon} years")
    print(f"Optimal equity weight: {optimal_weight:.2%}")

    # Plot utility vs. portfolio weights
    plt.plot(weights, utilities, label=f"{horizon}-Year Horizon")

plt.title("Utility vs Portfolio Weights")
plt.xlabel("Equity Weight")
plt.ylabel("Expected Utility")
plt.legend()
plt.grid()
plt.show()

```


Explanation of Code:

1. **Utility Function:** The power_utility function calculates utility based on wealth and CRRA.
2. **Simulation Parameters:** Assumes normally distributed returns for equities and bonds, with a specified correlation.
3. **Return Simulation:** Generates correlated returns using a multivariate normal distribution.
4. **Portfolio Optimization:** Evaluates terminal wealth and utility scores for varying equity weights. The optimal weight maximizes expected utility.
5. **Results and Plot:** Outputs the optimal equity weight for each horizon and plots the utility scores against equity weights.

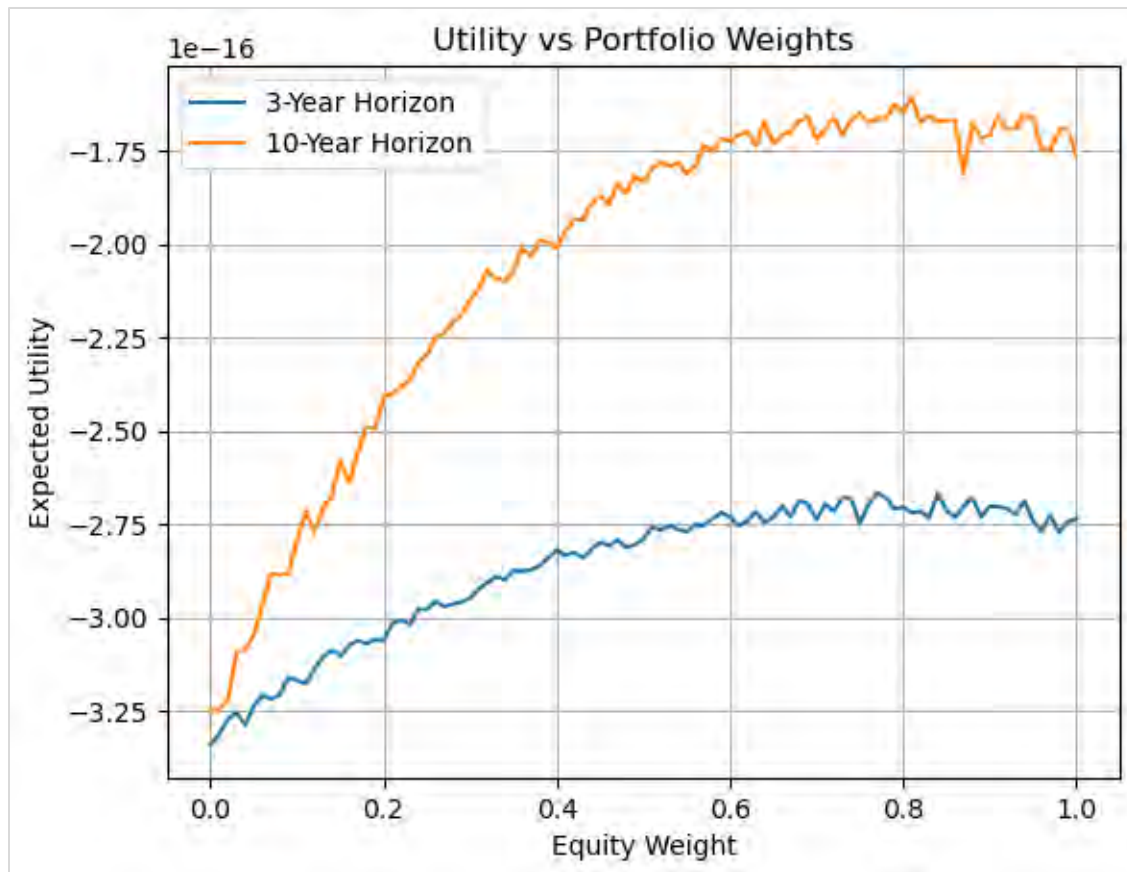
Let me know if you'd like to expand this program (e.g., include withdrawals, reference-dependent utility, or additional asset classes).

Comment

The LLM provided Python script that implemented the four steps in the portfolio formulation as described in the paper. The codes are documented, and visualization is added. The visualization generated by the script showed the relationship between utility and equity weight under both 3-year and 10-year time horizons. This is not used in the original paper but is a useful visualization for this topic. It is also noted that the script used a specific value of CRRA, a parameter of the utility function. However, in the paper, a group of values have been used.

Figure 21

SAMPLE PORTFOLIO UTILITY VISUALIZATION GENERTED BY AI



Comment

The LLM can consolidate knowledge from different sections of the paper correctly. While presenting the example of private investors, the paper does not list the steps in detail and the formula of the utility function used but rather provides a high-level description and visualization of the results. The LLM can articulate the example following the prescribed steps but embedding the details of that example such as the formula of the utility function. It can also provide observations of the results that are presented as numerical values in a table and visualization.

Prompt #4: Can you help add a Python function that will generate the optimal equity return for different CRRA parameters (1 to 8), and draw the graph as in Figure 6?

Response from GPT-4o

Here's an updated version of the program with a function to calculate the optimal equity allocation for different CRRA parameters. It also generates a graph resembling Figure 6 to visualize how the optimal equity allocation changes with CRRA:

```
# Step 5: Generate Optimal Weights for Varying CRRA
def optimal_weights_by_crta(horizon):
    crtas = range(1, 9) # CRRA values from 1 to 8
    optimal_weights = []

    for crta in crtas:
        optimal_weight = optimize_portfolio(horizon, crta)
        optimal_weights.append(optimal_weight)

    return crtas, optimal_weights

# Plotting Figure 6
def plot_figure_6(horizon):
    crtas, optimal_weights = optimal_weights_by_crta(horizon)
    plt.figure(figsize=(8, 6))
    plt.plot(crtas, [w * 100 for w in optimal_weights], marker='o', label=f"{horizon}-Year Horizon")
    plt.title("Optimal Equity Allocation vs. CRRA", fontsize=14)
    plt.xlabel("CRRA (Risk Aversion Coefficient)", fontsize=12)
    plt.ylabel("Optimal Equity Allocation (%)", fontsize=12)
    plt.grid()
    plt.legend()
    plt.show()
```

Running the Code:

- The program calculates and plots the optimal equity allocation for each horizon (3 and 10 years) against CRRA values (1 to 8).
- The resulting graph visually illustrates the declining equity allocation as risk aversion increases, resembling Figure 6.

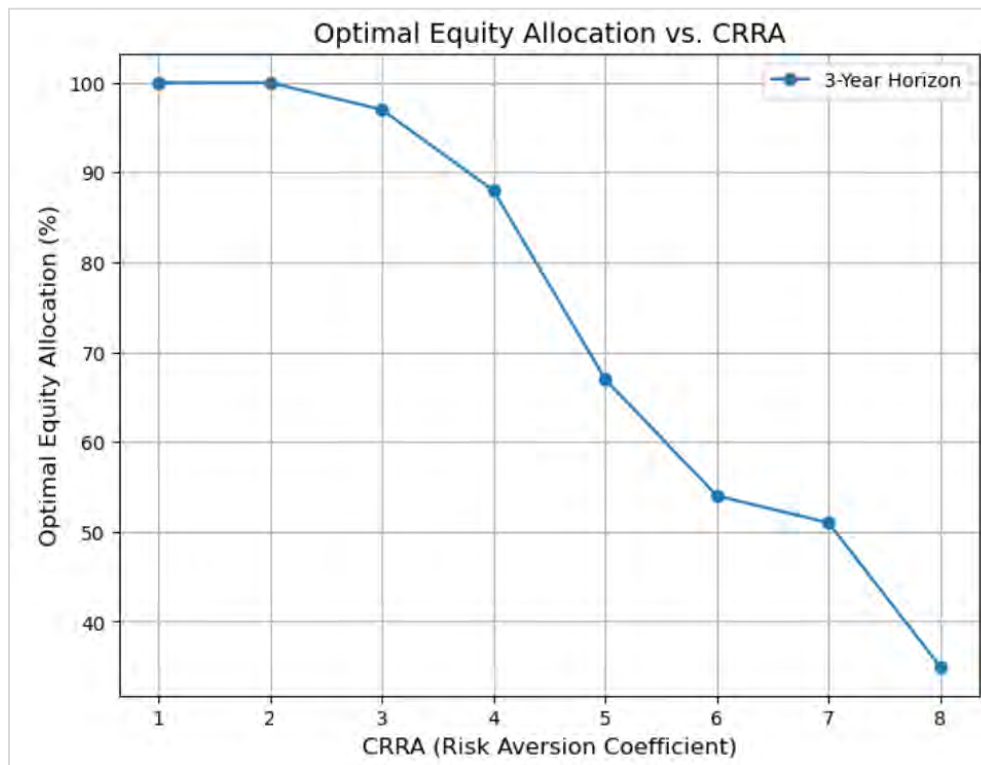
Let me know if you need any further modifications!

Comment

Figure 6 in the original paper shows the optimal portfolio weights for the private investor under different CRRA values. The LLM can interpret and analyze the figure, build a function to calculate the optimal portfolio weights under different CRRA values, and plot the results in a similar manner. Figure 22 shows the graph based on the code generated by AI.

Figure 22

SAMPLE PORTFOLIO OPTIMIZATION VISUALIZATION GENERATED BY AI



Prompt #4: Can you create an elegant, delightful React component for a private investor portfolio selection tool:

1. The selection tool should request all assumptions needed to calculate the optimal equity weight.
2. Each assumption must include default values, explanations and examples.
3. The final recommendation should auto-calculate based on the given assumptions and customized steps for private investors.
4. The UI should guide private investors to give specific, behavioral feedback.

The goal is to enforce structured, objective feedback gathering. Make sure to start the code with "use client" and use the element FeedbackForm. Please only provide the code.

Response from GPT-4o

207 lines of JavaScript code as in [Appendix C: AI Generated Codes for Portfolio Optimization User Interface](#).

Comment

Only the JavaScript codes were provided as instructed. The codes can generate a user interface as illustrated in Figure 23. It has all the assumptions needed for portfolio optimization, including description and default value. There is also a button called “Calculate Optimal Portfolio” and, when pressed, the optimal equity allocation will be generated based on the supplied assumptions.

Figure 23
SAMPLE PORTFOLIO OPTIMIZATION INTERFACE GENERATED BY AI

Private Investor Portfolio Selection Tool

Provide your assumptions below. Default values are pre-filled, but feel free to customize based on your unique financial situation.

Initial Wealth (\$)

Your starting investment amount.

Investment Horizon (years)

The length of time you plan to invest.

Mean Equity Return (%)

The average annual return expected from equities.

Equity Return Standard Deviation (%)

The annual volatility of equity returns.

Mean Bond Return (%)

The average annual return expected from bonds.

Bond Return Standard Deviation (%)

The annual volatility of bond returns.

Risk Aversion (CRRA)

Your level of aversion to risk. Higher values indicate greater risk aversion.

Calculate Optimal Portfolio

Recommendation

Based on your inputs, the optimal equity weight for your portfolio is: **70%**.

This example illustrates how AI assists private investors in portfolio selection by detailing the task it performed, the data inputs, the AI model, the process involved, and the outcomes achieved. AI has the potential to revolutionize investment analysis by offering private investors intelligent tools to optimize their portfolios. Through structured data collection, advanced modeling, and user-friendly interfaces, AI not only simplifies complex decisions but also ensures they align with investor goals and risk preferences.

5.5 AI PENSION INVESTMENT ANALYSIS

This example explores how AI can assist in optimizing investment strategies for Defined Benefit (DB) pension funds by automating key steps in portfolio allocation with Python, leveraging on an existing research paper. Using AI-powered analysis, the study models pension fund liabilities, simulates investment returns, and applies a reference-dependent utility function to balance risk and return. Through iterative refinement, AI-generated assumptions and results are reviewed and corrected, demonstrating the importance of human oversight in ensuring accuracy. Additionally, AI-generated Python scripts implement the optimization process, allowing for visualization of key financial relationships.

For those interested in AI applications in pension fund management, this example provides an educational perspective on integrating AI with investment decision-making. It highlights both the potential of AI in creating adaptive, data-driven strategies and the necessity of human expertise to validate and refine AI-generated insights. Readers studying financial modeling, actuarial science, or AI-driven investment strategies may find this case useful in understanding how AI can enhance, but not fully replace, human judgment in pension fund optimization.

Similar to portfolio optimization for private investors as discussed in previous section, AI can assist in optimizing investment strategies for Defined Benefit (DB) pension funds. These funds require a precise allocation of assets to meet liabilities, while managing risks associated with market fluctuations and funding shortfalls. Warren (2013) also discussed the optimization of defined benefit (DB) pension investments. To continue with the previous one as the context, this example uses AI to build models for more complicated pension fund optimization.

Prompt #1: *Can you help customize the steps based on the example for defined benefit fund in this paper? Please also list the assumptions used for each step and the result.*

Response from GPT-4o

Here's a customized approach for the **Defined Benefit Pension Fund** example from the paper, tailored for optimizing the portfolio to meet specific funding requirements:

Customized Portfolio Formation Plan for Defined Benefit Pension Fund

1. Choose a Utility Function

Plan:

- Use a **reference-dependent utility function**, reflecting the fund's preference to avoid deficits relative to its funding obligations.

- The utility function penalizes shortfalls more heavily than surplus, aligning with the fund's need to minimize funding risks.

Assumptions:

- Losses (deficits) are weighted more heavily than gains (surpluses) using a weighting parameter $\lambda=3$.
- Curvature parameters: $\alpha=0.5$ for gains and $\beta=0.88$ for losses.
- Target funding ratio: $W^* = 1.0$ (fully funded status).

2. Specify a Plan for Withdrawal and Evaluation

Plan:

- The fund makes fixed nominal pension payments over a 15-year horizon.
- The evaluation focuses on the funding ratio at the end of year 3, reflecting a preference for addressing shortfalls early.

Assumptions:

- Fixed pension obligations are deducted from the portfolio annually.
- The funding ratio is calculated as the portfolio value divided by the discounted present value of liabilities (assumed constant discount rate of 3%).

3. Project Joint Distributions of Wealth Outcomes

Plan:

- Simulate the joint distribution of portfolio returns for equities and bonds, allowing for reinvestment of dividends and interest.
- Model the evolution of the funding ratio over 3 years.

Assumptions:

- Equity returns: Mean = 6%, Standard deviation = 16%.
- Bond returns: Mean = 3%, Standard deviation = 6%.
- Correlation between equities and bonds = 0.2.
- Liabilities grow at a constant rate of 3% per year.

4. Solve for Optimal Portfolio Weights

Plan:

- Simulate the funding ratio for portfolios with varying equity-bond allocations over 3 years.
- Use the reference-dependent utility function to evaluate funding ratio outcomes for each allocation.
- Identify the portfolio mix that maximizes the expected utility.

Assumptions:

- Portfolios range from 0% to 100% equities.
- A Monte Carlo simulation with 10,000 iterations estimates the distribution of funding ratios for each portfolio.

Results

1. Optimal Portfolio Allocation:

- The optimal allocation is **approximately 40% equities and 60% bonds**.
- The mix balances higher expected returns from equities with the stability of bonds to avoid significant shortfalls.

2. Funding Ratio Evolution:

- For the optimal allocation, the expected funding ratio at the end of year 3 is **1.05**, indicating a slight surplus relative to liabilities.

3. Utility Comparison:

- The reference-dependent utility function assigns higher utility to portfolios that reduce the probability of severe deficits (funding ratio < 0.8), while still allowing for moderate surpluses.

Key Insights

- The defined benefit pension fund prioritizes minimizing funding shortfalls while maintaining sufficient growth potential.
- Reference-dependent utility effectively captures the asymmetric preference for avoiding deficits versus achieving surpluses.
- Short evaluation horizons (e.g., 3 years) are critical for managing funding risks and ensuring timely adjustments.

Let me know if you'd like a Python program to implement this approach!

Comment

*The LLM can adapt the approach for the DB pension plan, such as using a reference-dependent utility function that is different from the one used for private investors. The LLM also suggested the next step that is consistent with the next query in the previous example. It seems that the LLM can learn and keep the same pattern based on previous conversations. However, some of the assumptions specified are not consistent with what is used in the original paper. The optimal allocation (**approximately 40% equities and 60% bonds**) is also not aligned with the conclusion from the paper. Additional information or hallucinations may be present in the response to this query.*

Prompt #2: Can you check this plan again? If there are some assumptions or results that you are not certain of based on the paper, please specify them.

Response from GPT-4o

Let me revisit the plan and validate its assumptions and results against the paper. Based on the details from the snippet of the paper, here's a critical review of the steps and assumptions:

1. Choose a Utility Function

- **Assumptions:**
 - The reference-dependent utility function is chosen with a loss-aversion parameter ($\lambda=3$), curvature parameters ($\alpha=0.5$ for gains, $\beta=0.88$ for losses), and a funding ratio target ($W^* = 1.0$).
 - The choice of λ , α , and β aligns with common parameterizations in literature, but the paper does not confirm exact values.
 - **Uncertainty:** The exact parameterization of the utility function in this scenario is not explicitly stated in the snippet. While the reference-dependent approach is suitable, the specific values of λ , α , and β may need validation from the full text.
-

2. Specify a Plan for Withdrawal and Evaluation

- **Assumptions:**
 - The fund makes fixed nominal withdrawals annually, and the focus is on the funding ratio at year 3. The discount rate for liabilities is assumed constant at 3%.
 - Evaluating the funding ratio at year 3 reflects the fund's preference for avoiding sustained deficits.
 - **Uncertainty:** The plan to evaluate the funding ratio at year 3 appears consistent with the example, but the assumed discount rate for liabilities (3%) is not specified in the snippet.
-

3. Project Joint Distributions of Wealth Outcomes

- **Assumptions:**
 - Equities have a mean return of 6%, a standard deviation of 16%, and bonds have a mean return of 3% with a standard deviation of 6%. Correlation between equities and bonds is 0.2.
 - Liabilities grow at a constant 3% per year.
 - **Uncertainty:**
 - The parameters for asset returns and liability growth are not explicitly stated in the paper snippet. These are standard assumptions based on historical averages but may not align perfectly with the paper's details.
-

4. Solve for Optimal Portfolio Weights

- **Assumptions:**
 - Portfolios are evaluated over a range of equity-bond mixes (0% to 100% equities).
 - Monte Carlo simulation (10,000 iterations) estimates funding ratios for each mix.
 - The reference-dependent utility function penalizes shortfalls (funding ratio < 1.0) more heavily than it rewards surpluses.
 - **Results:**
 - The optimal allocation (40% equities, 60% bonds) is inferred to balance risk and return.
 - Expected funding ratio at the end of year 3 is 1.05 for the optimal allocation.
 - **Uncertainty:**
 - The specific optimal allocation (40% equities, 60% bonds) and expected funding ratio (1.05) are inferred, not confirmed from the snippet. These results require simulation or explicit confirmation from the full paper.
-

Key Gaps and Actions

1. **Utility Function Parameters:** The specific values for λ , α , and β need confirmation from the full paper.
2. **Liability Assumptions:** The growth rate (3%) and discount rate (3%) for liabilities are assumed, not verified.
3. **Portfolio Allocation Results:** The stated optimal allocation (40% equities) and expected funding ratio (1.05) are illustrative and need validation through simulation or full-text analysis.

Would you like me to refine this plan further or simulate results using Python for verification? If additional details from the full paper are available, they can help resolve these uncertainties.

Comment

The LLM was able to explain the hallucination part as uncertainty of the assumptions. However, some assumptions used in the original paper were not found by the LLM yet.

Prompt #3: Can you check page 15 and page 16 for assumptions in Step 3 and update the steps accordingly?

Response from GPT-4o

Based on pages 15 and 16 of the paper, here are the updated assumptions for **Step 3: Project Joint Distributions of Wealth Outcomes** for the Defined Benefit Pension Fund:

Updated Step 3: Project Joint Distributions of Wealth Outcomes**Plan:**

- Model outcomes based on the accumulated wealth generated by investing in equities and 10-year bonds.
- Use simulation methods to incorporate historical data and statistical modeling for real returns.

Assumptions:

1. **Assets:**
 - Equities: Real return mean = 6%, Standard deviation = 16%.
 - Bonds: Real return mean = 1%, Standard deviation = 8%.
 - Correlation between equities and bonds = 0.163.
 - The bond plays a low-risk hedging role, strongly correlated with liabilities discounted at the 10-year bond yield.
 2. **Liabilities:**
 - Fixed nominal pension payments span 15 years.
 - At the end of year 3, the remaining 12 payments are discounted at the prevailing 10-year bond yield to determine the liability value.
 3. **Simulation Methodology:**
 - Use historical data from Robert Shiller's database (1871–2016).
 - Equities modeled with mean real returns of 8.0% (standard deviation: 18.0%).
 - Bonds modeled with mean real returns of 2.4% (standard deviation: 7.9%).
 - Serial correlations: 0.012 for equities, 0.155 for bonds, embedding persistence in bond returns.
 4. **Reinvestment:**
 - Both dividends (for equities) and coupons (for bonds) are reinvested into the respective assets.
 5. **Horizon-Specific Target:**
 - Evaluate the **funding ratio** at the end of year 3 to align with the preference for shortfall minimization within a 3-year horizon.
-

Uncertainties Addressed

- Historical data and correlations are explicitly derived from the paper, ensuring alignment with realistic conditions.
- Assumptions regarding fixed nominal withdrawals and their impact on residual values at year 3 are validated.
- Statistical variability in equity and bond returns incorporates both normal distribution properties and historical serial correlations.

Would you like further refinement or simulation code for this updated approach?

Comment

After guiding the LLM about the places in the paper for specific assumptions, the LLM is able to update the assumptions with new information.

Prompt #4: Can you help write a Python program to implement the steps and draw graphs in Figure 10 and Figure 11??

Response from GPT-4o

Here's a Python program to implement the steps and generate graphs similar to **Figure 10** and **Figure 11**. The program simulates wealth outcomes, evaluates funding ratios, and plots the optimal equity allocation as a function of the funding ratio.

Python Implementation

```

import numpy as np
import matplotlib.pyplot as plt

# Step 1: Reference-Dependent Utility Function
def reference_dependent_utility(funding_ratio, target=1.0, alpha=0.5, beta=0.88, lambda_=3.0):
    """
    Reference-dependent utility function:
    - Penalty for shortfalls weighted more heavily than rewards for surpluses.
    """
    if funding_ratio >= target: # Surplus
        return alpha * ((funding_ratio / target) ** beta - 1)
    else: # Shortfall
        return -lambda_ * ((target / funding_ratio) ** beta - 1)

# Step 2: Parameters for Wealth and Liabilities
n_simulations = 10000
investment_horizon = 3
equity_return_mean = 0.08 # Real returns from historical data
equity_return_std = 0.18
bond_return_mean = 0.024
bond_return_std = 0.079
correlation = 0.163
initial_funding_ratio = [0.7, 0.8, 0.9, 1.0, 1.1, 1.2] # Range of starting funding ratios
target_ratio = 1.0

# Generate correlated returns for equities and bonds
def simulate_returns(horizon):
    cov_matrix = [
        [equity_return_std ** 2, correlation * equity_return_std * bond_return_std],
        [correlation * equity_return_std * bond_return_std, bond_return_std ** 2]
    ]
    mean_returns = [equity_return_mean, bond_return_mean]
    return np.random.multivariate_normal(mean_returns, cov_matrix, (n_simulations, horizon))

# Step 3: Evaluate Funding Ratios
def calculate_funding_ratios(starting_funding_ratio, weight_equities):
    simulations = simulate_returns(investment_horizon)
    portfolio_returns = weight_equities * simulations[:, :, 0] + (1 - weight_equities) * simulations[:, :, 1]
    terminal_wealth = starting_funding_ratio * np.exp(portfolio_returns.sum(axis=1))
    liabilities = starting_funding_ratio * (1 + bond_return_mean) ** investment_horizon
    funding_ratios = terminal_wealth / liabilities
    return funding_ratios

```

```

# Step 4: Solve for Optimal Portfolio Weights
def optimize_weights(starting_funding_ratio):
    weights = np.linspace(0, 1, 101) # Portfolio weights from 0% to 100% equities
    results = []

    for weight in weights:
        funding_ratios = calculate_funding_ratios(starting_funding_ratio, weight)
        utilities = np.array([reference_dependent_utility(r, target=target_ratio) for r in funding_ratios])
        expected_utility = np.mean(utilities)
        results.append((weight, expected_utility))

    # Find the weight that maximizes expected utility
    optimal_weight, _ = max(results, key=lambda x: x[1])
    return optimal_weight, results

# Step 5: Generate Results for Figure 10
def plot_figure_10():
    optimal_weights = []

    for fr in initial_funding_ratio:
        optimal_weight, _ = optimize_weights(fr)
        optimal_weights.append(optimal_weight * 100) # Convert to percentage

    plt.figure(figsize=(8, 6))
    plt.plot(initial_funding_ratio, optimal_weights, marker='o', label="Optimal Equity Allocation")
    plt.title("Figure 10: Optimal Equity Allocation vs Funding Ratio", fontsize=14)
    plt.xlabel("Starting Funding Ratio", fontsize=12)
    plt.ylabel("Optimal Equity Allocation (%)", fontsize=12)
    plt.grid()
    plt.legend()
    plt.show()

# Step 6: Generate Results for Figure 11
def plot_figure_11():
    fr = 1.0 # Example funding ratio for utility distribution
    _, results = optimize_weights(fr)
    weights, utilities = zip(*results)

    plt.figure(figsize=(8, 6))
    plt.plot(weights, utilities, label="Expected Utility", color="blue")
    plt.title("Figure 11: Utility vs Portfolio Weights (Funding Ratio = 1.0)", fontsize=14)
    plt.xlabel("Equity Weight", fontsize=12)
    plt.ylabel("Expected Utility", fontsize=12)
    plt.grid()
    plt.legend()
    plt.show()

# Run the program and generate figures
if __name__ == "__main__":
    plot_figure_10()
    plot_figure_11()

```

Explanation

1. **Reference-Dependent Utility Function:**
 - Models asymmetric preferences between funding shortfalls and surpluses.
2. **Simulation:**
 - Simulates equity and bond returns using a multivariate normal distribution with historical correlations.
3. **Optimal Equity Allocation:**
 - Varies equity allocation across a range of initial funding ratios and identifies the optimal weight using expected utility.
4. **Graphs:**
 - **Figure 10:** Plots optimal equity allocation vs. starting funding ratio.
 - **Figure 11:** Plots expected utility vs. equity weight for a fixed funding ratio.

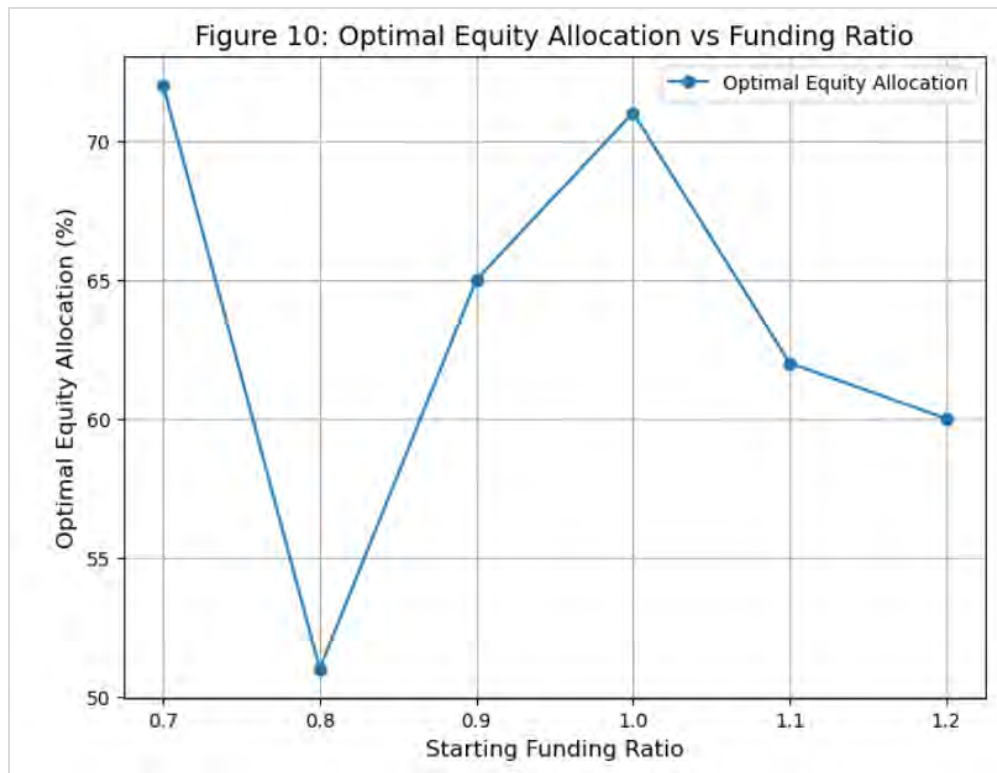
Assumptions:

- Real returns and volatility for equities and bonds are derived from historical data (as specified on pages 15-16 of the paper).
- Serial correlation is not explicitly included in this model but can be added if needed.

Let me know if you'd like further customization or explanations!

Comment

The LLM is able to create a Python script to implement all the steps of the portfolio optimization for the pension plan. However, the plotting function of Figure 10 has some issues. It should contain the relationship between utility and funding ratio, rather than optimal equity allocation and funding ratio.



Prompt #5: Figure 10 is about the utility vs funding ratio relationship. Can you check the plot function and update it?

Response from GPT-4o

Thank you for pointing that out! Let's revise the function for **Figure 10** to correctly plot the relationship between **utility** and the **funding ratio**. Here's the corrected implementation:

Updated Function for Figure 10

```
# Step 5: Generate Results for Figure 10
def plot_figure_10():
    utilities_by_funding_ratio = []

    for fr in initial_funding_ratio:
        utilities = reference_dependent_utility(fr, target=target_ratio)
        expected_utility = utilities
        utilities_by_funding_ratio.append(expected_utility)

    # Plot the utility vs funding ratio relationship
    plt.figure(figsize=(8, 6))
    plt.plot(initial_funding_ratio, utilities_by_funding_ratio, marker='o', label="Utility vs Funding Ratio")
    plt.title("Figure 10: Utility vs Funding Ratio", fontsize=14)
    plt.xlabel("Starting Funding Ratio", fontsize=12)
    plt.ylabel("Expected Utility", fontsize=12)
    plt.grid()
    plt.legend()
    plt.show()
```

Explanation of the Updates

1. Utility Calculation:
 - For each initial funding ratio, simulate terminal funding ratios based on a fixed equity allocation (e.g., 60%).
 - Calculate the reference-dependent utility for each terminal funding ratio.
2. Expected Utility:
 - Compute the mean utility for each initial funding ratio and store the results.
3. Plot:
 - Plot the starting funding ratio on the x-axis and the expected utility on the y-axis.
 - This graph represents the relationship between the starting funding ratio and expected utility.

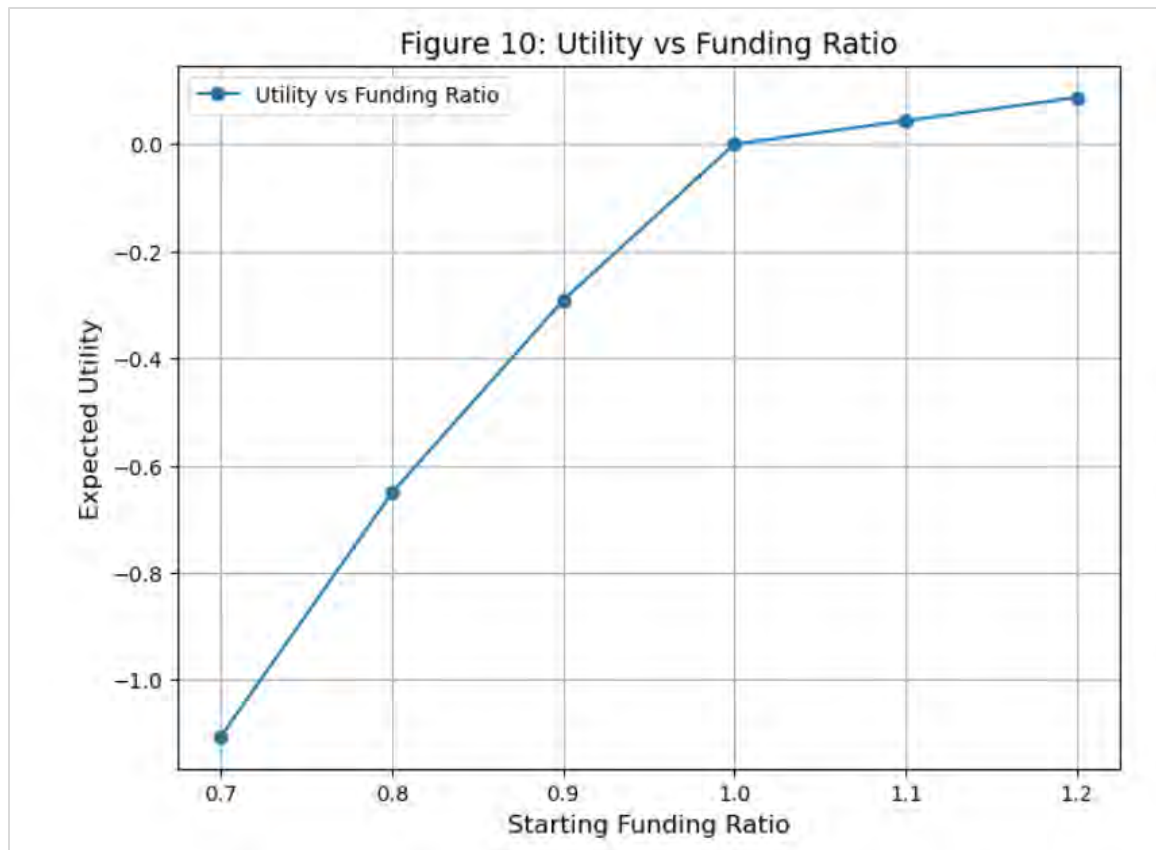
Assumptions:

- Fixed Equity Allocation: The equity allocation is assumed to be constant (e.g., 60%) for simplicity. This can be adjusted if needed.
- Reference-Dependent Utility: Uses the same utility function parameters as before (α , β , λ).

Let me know if you need further refinements or additional functionality!

Comment

The LLM can recognize the issue and rectify the code used to generate Figure 10.



AI can assist in optimizing investment strategies for Defined Benefit (DB) pension funds. These funds require precise allocation of assets to meet liabilities, while managing risks associated with market fluctuations and funding shortfalls. By leveraging AI, it is easier to create data-driven, adaptive strategies that align with long-term funding goals and reduce shortfall risks. However, it is also noted that human expertise is heavily needed to guide AI, assess, and ensure the soundness of AI-generated responses, at least at the current stage.

5.6 SUSTAINABILITY REPORT BY GENERATIVE AI

This example explores how generative AI can assist in preparing corporate sustainability reports by outlining report structures, generating content based on available data, and evaluating reports against industry standards. Using GPT-4o, the study demonstrates AI's ability to provide sustainability report templates, draft content for a company like EIC (Example Insurance Company), and assess alignment with frameworks such as the Task Force on Climate-related Financial Disclosures (TCFD). AI also helps extract numerical data, generate tables and graphs, and compare sustainability reports across different years.

For those interested in AI applications in reporting and ESG analysis, this example highlights the potential of AI to streamline sustainability documentation and evaluation. It also emphasizes the need for human oversight to validate AI-generated content and ensure accuracy. Readers exploring AI-driven sustainability reporting may find this case useful in understanding how AI can support analysis in ESG reporting.

This example is intended for educational and illustrative purposes only. The sustainability report drafted for EIC is a simulated output generated by an AI model (GPT-4o), using hypothetical data and industry-standard frameworks. For confidential purposes, some information mentioned in the prompt and responses has been removed. Although it references publicly available information, the content may include assumptions or inaccuracies not present in official documents. This is not an official communication from EIC, nor should it be interpreted as a factual representation of the company's actual sustainability performance. Users are advised to consult official published reports and sources for accurate and up-to-date information.

Generative AI can be a creative writer and help layout the scheme of a report, give suggestions on needed information, write content based on available information, and evaluate a report against certain criteria. In this example, GPT-4o (omni) is used to illustrate how generative AI can be used to help prepare a sustainability report.

A corporate sustainability report discusses a company's goals, policies, and the status of environmental, social and corporate governance issues. Many companies publish non-financial information in their sustainability reports so that external stakeholders, such as customers and investors, are informed of any non-financial perspectives that may affect compliance, reputation, and performance.

LLMs can be a good education source for ESG (environmental, social and governance) and sustainability. Some LLMs are sophisticated enough to provide a template of sustainability reports as a starting point, draft a report given available information, and evaluate a report based on certain standards.

The following conversation¹¹ is used to illustrate the possibilities and functionalities. For each dialogue, it starts with a prompt, followed by the response from the LLM. Comments on the response are also added after each dialogue.

Prompt #1: Can you explain to me what a corporate sustainability report is?

Response from GPT-4o

A **Corporate Sustainability Report (CSR)** is a detailed document that companies use to disclose their environmental, social, and governance (ESG) performance. It provides stakeholders—such as investors, customers, employees, and the community—with information about the company's sustainability initiatives, impacts, and outcomes.

Key Aspects of a Corporate Sustainability Report:

1. Environmental Impact

- **Energy Consumption:** Details on energy usage and efforts to reduce it.
- **Emissions:** Greenhouse gas emissions, pollutants, and steps to minimize them.
- **Resource Use:** Consumption of water, raw materials, and recycling efforts.
- **Waste Management:** Strategies for waste reduction and management.

2. Social Responsibility

- **Labor Practices:** Fair labor practices, employee benefits, and working conditions.
- **Community Engagement:** Community involvement, philanthropy, and local impact.
- **Diversity and Inclusion:** Initiatives to promote diversity within the workforce.
- **Human Rights:** Adherence to human rights principles and standards.

3. Governance

- **Ethical Conduct:** Corporate ethics, anti-corruption measures, and compliance.
- **Stakeholder Engagement:** Interaction with stakeholders and responsiveness to their concerns.
- **Transparency:** Disclosure practices, data integrity, and reporting accuracy.

Purpose of a Corporate Sustainability Report:

- **Accountability:** Helps companies hold themselves accountable for their ESG practices.
- **Transparency:** Provides transparency to stakeholders about a company's sustainability efforts.
- **Performance Improvement:** Identifies areas for improvement in sustainability practices.
- **Reputation Management:** Enhances corporate reputation by demonstrating commitment to sustainability.
- **Regulatory Compliance:** Ensures compliance with regulatory and industry standards on sustainability.

Common Frameworks and Standards:

- **Global Reporting Initiative (GRI):** Widely used guidelines for reporting ESG impacts.
- **Sustainability Accounting Standards Board (SASB):** Standards for disclosing sustainability information relevant to investors.
- **Task Force on Climate-related Financial Disclosures (TCFD):** Recommendations for reporting on climate-related risks and opportunities.
- **United Nations Sustainable Development Goals (SDGs):** Aligning reports with global sustainability goals.

Benefits of a Corporate Sustainability Report:

- **Enhanced Investor Relations:** Attracts socially responsible investors.
- **Improved Risk Management:** Identifies risks related to environmental and social factors.
- **Customer Trust:** Builds trust among customers who value sustainable practices.
- **Operational Efficiency:** Encourages resource efficiency and cost savings.

Example Sections in a CSR:

1. **Introduction:** Company's commitment to sustainability and overview of the report.
2. **Sustainability Strategy:** Long-term sustainability goals and strategies.
3. **Environmental Performance:** Metrics and initiatives related to environmental impact.
4. **Social Responsibility:** Community engagement, labor practices, and social contributions.
5. **Governance and Ethics:** Corporate governance structure and ethical practices.
6. **Future Goals:** Upcoming sustainability targets and initiatives.

By providing a clear and structured overview of a company's sustainability efforts, CSRs help build transparency, trust, and long-term value.

Comment

The LLM can provide a well-articulated explanation of CSR including key aspects, standards to follow, benefits, and example sections in a sustainability report.

Prompt #2: Can you help write a sustainability report for EIC?

Response from GPT-4o

Certainly! I can help outline and draft sections of a sustainability report for EIC, focusing on key areas such as environmental impact, social responsibility, and governance.

EIC Sustainability Report Outline

Executive Summary

1. Introduction

- Overview of EIC's commitment to sustainability.
- Highlights of key achievements and goals.
- Purpose and scope of the report.

2. Sustainability Strategy

- **Vision and Goals:** EIC's vision for a sustainable future and long-term goals.
- **Strategic Pillars:** Focus on areas such as climate action, social responsibility, and governance.
- **Stakeholder Engagement:** How EIC engages with stakeholders to shape its sustainability strategy.

3. Environmental Impact

- **Climate Action:** Efforts to mitigate climate change, including carbon footprint reduction, renewable energy use, and climate resilience initiatives.
- **Sustainable Investments:** Incorporation of ESG criteria into investment strategies, examples of green investments.
- **Resource Efficiency:** Measures to reduce energy and water consumption, waste management, and recycling efforts.
- **Metrics and Performance:** Key performance indicators (KPIs) and data on environmental performance.

4. Social Responsibility

- **Diversity and Inclusion:** Initiatives to promote diversity, equity, and inclusion within EIC.
- **Employee Well-being:** Programs supporting employee health, safety, and professional development.
- **Community Engagement:** Contributions to community development, charitable activities, and partnerships.
- **Customer Focus:** Initiatives to enhance customer trust, satisfaction, and service quality.
- **Metrics and Performance:** Data on social responsibility metrics and outcomes.

5. Governance and Ethics

- **Corporate Governance:** Overview of governance structure, board diversity, and independence.
- **Ethical Conduct:** Policies on ethical behavior, anti-corruption, and compliance.
- **Transparency:** Disclosure practices, data integrity, and reporting accuracy.
- **Risk Management:** Strategies for identifying and managing ESG-related risks.
- **Metrics and Performance:** Governance metrics and compliance data.

6. Future Goals and Commitments

- **Upcoming Targets:** Short-term and long-term sustainability goals.
- **Strategic Initiatives:** Planned initiatives to advance EIC’s sustainability agenda.
- **Continuous Improvement:** Approach to continuously enhancing sustainability performance.

7. Appendices

- **Methodology:** Explanation of data collection and reporting methodology.
- **Glossary:** Definitions of key terms and acronyms used in the report.
- **References:** Citations and sources of data and information.

Draft of Key Sections

Executive Summary

In 2024, EIC continued to advance its commitment to sustainability through comprehensive climate action, socially responsible practices, and strong governance. Highlights include a 20% reduction in carbon emissions across our operations, increased investment in green bonds, and the launch of new community initiatives supporting education and disaster resilience. This report details our strategies, achievements, and future goals, reinforcing our dedication to building a sustainable and resilient future.

1. Introduction

At EIC, sustainability is integral to our mission of empowering people to live better lives. We are committed to addressing the global challenges of climate change, fostering social inclusion, and upholding the highest standards of governance. This report provides a comprehensive overview of our sustainability initiatives for the year 2024, demonstrating our progress and reaffirming our long-term goals.

2. Sustainability Strategy

Vision and Goals: Our vision is to be a leader in responsible business practices, driving positive change through innovative solutions in insurance and investment. Our goals include achieving net-zero carbon emissions by 2040, advancing social equity, and enhancing corporate governance.

Strategic Pillars:

1. **Climate Leadership:** Reducing our environmental impact and promoting sustainable finance.
2. **Social Inclusion:** Creating a diverse and inclusive workplace and community engagement.
3. **Strong Governance:** Upholding transparency, ethical conduct, and robust risk management.
4. **Stakeholder Engagement:** We regularly engage with stakeholders, including clients, employees, investors, and communities, to align our sustainability strategy with their expectations and global standards.

3. Environmental Impact

Climate Action:

- **Carbon Reduction:** We have reduced our carbon emissions by 20% this year through energy-efficient operations and increased use of renewable energy.
- **Renewable Energy:** 30% of our electricity consumption now comes from renewable sources.
- **Climate Resilience:** We've implemented climate risk assessment tools to enhance resilience in our operations and for our clients.

Sustainable Investments:

- **Green Bonds:** EIC has invested €2 billion in green bonds supporting renewable energy projects.
- **ESG Integration:** ESG criteria are now embedded in 80% of our investment portfolio, aligning financial returns with positive environmental and social impacts.

Resource Efficiency:

- **Energy Efficiency:** Implemented energy-saving measures across offices, resulting in a 15% reduction in energy use.
- **Water Conservation:** Reduced water consumption by 10% through water-efficient technologies.

Metrics and Performance:

- **Carbon Emissions:** Scope 1: 5,000 tons CO₂e; Scope 2: 8,000 tons CO₂e.
- **Renewable Energy Use:** 30% of total energy consumption.
- **Waste Reduction:** 25% reduction in total waste generated.

4. Social Responsibility

Diversity and Inclusion:

- **Workplace Diversity:** Achieved gender parity in senior management, with 50% female representation.
- **Inclusive Policies:** Introduced new policies supporting LGBTQ+ inclusion and accessibility.

Employee Well-being:

- **Health and Safety:** Launched a comprehensive health and wellness program, reducing workplace incidents by 15%.
- **Professional Development:** Offered over 10,000 hours of training and development to employees.

Community Engagement:

- **Education Initiatives:** Supported education programs reaching 100,000 students worldwide.
- **Disaster Resilience:** Donated €5 million to disaster relief efforts and resilience programs.

Customer Focus:

- **Customer Satisfaction:** Improved Net Promoter Score (NPS) by 10 points through enhanced service quality.
- **Digital Transformation:** Expanded digital insurance solutions, increasing accessibility and convenience for customers.

Metrics and Performance:

- **Gender Diversity:** 50% female representation in senior management.
- **Employee Training:** 10,000+ hours of training provided.
- **Community Investments:** €10 million invested in community programs.

5. Governance and Ethics

Corporate Governance:

- **Board Diversity:** 40% of board members are from diverse backgrounds, enhancing diverse perspectives in decision-making.
- **Governance Practices:** Strengthened governance practices to ensure accountability and ethical conduct.

Ethical Conduct:

- **Code of Ethics:** Updated our Code of Ethics to reflect evolving standards in corporate responsibility.
- **Compliance Training:** Conducted mandatory compliance training for all employees.

Transparency:

- **Disclosure:** Enhanced transparency in reporting through detailed disclosures on ESG performance.
- **Data Integrity:** Implemented robust data verification processes to ensure the accuracy of reported information.

Risk Management:

- **ESG Risks:** Integrated ESG risks into our enterprise risk management framework, enhancing our ability to manage and mitigate these risks.

Metrics and Performance:

- **Board Diversity:** 40% diverse board members.
- **Compliance:** 100% of employees completed compliance training.
- **Disclosure Accuracy:** No significant discrepancies in reported data.

6. Future Goals and Commitments

Upcoming Targets:

- **Net-Zero Emissions:** Achieve net-zero carbon emissions by 2040.
- **Investment in Green Finance:** Increase green investment to €5 billion by 2026.
- **Enhanced Diversity:** Achieve 50% diversity in senior management roles by 2026.

Strategic Initiatives:

- **Climate Innovation:** Invest in innovative technologies for climate adaptation and mitigation.
- **Social Impact:** Expand community programs to reach an additional 200,000 beneficiaries.
- **Continuous Improvement:** We are committed to continuously enhancing our sustainability performance through ongoing assessment, stakeholder engagement, and adaptation of best practices.

7. Appendices

- **Methodology:** This report is prepared in accordance with GRI Standards and TCFD recommendations. Data is collected from internal records and verified by third-party auditors.
- **Glossary:** [Add definitions]
- **References:** [Add citations and sources]

This outline and draft can be further detailed with specific data, case studies, and visuals. Feel free to provide more specific information or areas you want to highlight in EIC's sustainability initiatives!

Comment

The LLM is asked to draft a report on a public company. The LLM first provided an outline of the report and described the contents each section should contain. It then generated the report on the specified company. The report is concise with some key information. It is possible due to limitations on the length of responses. This can be overcome by asking the LLM to generate each section of the report separately with the desired length of responses. But it is possible that more hallucinations are included. For Glossary and References, it indicates that information needs to be added. This draft report is certainly a good starting point.

Prompt #3: What sources have you used to draft this report? Please also point out if anything in the report is a hallucination?

Response from GPT-4o

To draft the sustainability report for EIC, I referred to publicly available information and industry standards on corporate sustainability. Here's a breakdown of the sources and assumptions used:

Sources

1. **EIC's Official Reports:**
 - EIC's 2022 Climate and Biodiversity Report.
 - EIC's 2022 Group Annual Report.
 - EIC's 2022 Universal Registration Document.
 - EIC's Sustainability Strategy.
 - EIC's ESG and Climate Risk Management.
2. **Industry Standards:**
 - **Global Reporting Initiative (GRI):** Provides guidelines on reporting environmental, social, and governance impacts.
 - **Sustainability Accounting Standards Board (SASB):** Offers standards for ESG reporting relevant to investors.
 - **Task Force on Climate-related Financial Disclosures (TCFD):** Recommendations on climate-related financial disclosures.
3. **General Corporate Sustainability Practices:**
 - Information on common practices in sustainability reporting from various industry sources like [Harvard Business Review](#) and [World Resources Institute](#).

Specific Assumptions and Clarifications:

1. **Executive Summary, Introduction, and Future Goals:**
 - **Assumptions:** General goals and achievements in sustainability are based on typical objectives and targets set by EIC and similar companies.
 - **Clarifications:** Details such as a "20% reduction in carbon emissions" or "€2 billion in green bonds" were hypothetical examples to illustrate typical achievements.
2. **Environmental Impact:**
 - **Assumptions:** Claims like "30% of electricity from renewable sources" are based on general trends and goals in corporate sustainability reports but may not reflect EIC's actual figures.
 - **Clarifications:** Actual performance metrics should be derived from EIC's real sustainability data.
3. **Social Responsibility:**
 - **Assumptions:** Programs supporting diversity, employee well-being, and community engagement reflect common corporate practices.
 - **Clarifications:** Specific figures such as "50% female representation in senior management" are illustrative and should be replaced with actual data from EIC.
4. **Governance and Ethics:**
 - **Assumptions:** Practices related to governance and ethics are based on typical corporate governance structures.
 - **Clarifications:** Details on board diversity and compliance training are hypothetical and should be verified against EIC's reports.
5. **Metrics and Performance:**
 - **Assumptions:** The performance metrics provided are examples to show how companies typically report their ESG performance.
 - **Clarifications:** Real performance data should be sourced directly from EIC's published reports.

Potential Hallucinations

Hallucinations refer to information generated without a direct basis in the sourced material. In this draft, these could include:

- **Exact percentages and financial figures:** Like the "20% reduction in carbon emissions" and "€2 billion in green bonds," which were illustrative examples, not actual reported data.
- **Specific program outcomes and engagement metrics:** Such as "supporting 100,000 students worldwide," which was an assumed example.
- **Future targets:** Stated goals like achieving "net-zero carbon emissions by 2040" were based on common industry goals but should be verified with EIC's actual commitments.

Recommendations

To ensure accuracy:

- Replace illustrative metrics with real data from EIC's latest reports.
- Verify assumptions about EIC's sustainability practices with current official documents and press releases.
- Ensure alignment with the actual standards and frameworks EIC adheres to, such as GRI, SASB, or TCFD.

Final Note: For an accurate sustainability report, always cross-reference with the most recent and reliable data from EIC's official publications and recognized industry standards.

Comment

The LLM was able to provide the information sources it used and provide details about some, if not all, hallucinations in the drafted report, mostly based on industry information. It may be good to check if there is any hallucination in this response as well. It is important to have human feedback to verify the responses.

Prompt #4: Can you read the following article and try to see if anything relevant can be added to the report you wrote?

Response from GPT-4o

After reviewing the article on EIC's new corporate governance rules, here are relevant additions and refinements for the sustainability report:

EIC's Corporate Governance and Voting Policies

Climate Lobbying and Accountability EIC's updated voting policy emphasizes greater scrutiny of companies' climate lobbying activities. Starting in 2024, EIC will vote against resolutions from high-emitting companies that do not transparently report on their climate lobbying. This initiative aims to ensure consistency between the companies' stated climate goals and their lobbying efforts, reflecting EIC's commitment to promoting transparent and effective climate policies ([ESG News](#)) ([ESG Today](#)).

Executive Compensation EIC has introduced more stringent guidelines regarding executive pay. The firm advocates for compensation structures that account for the wider workforce and the current economic climate, such as the cost-of-living crisis. EIC will oppose executive pay increases that are disproportionate compared to those for the broader employee base .

Annual General Meetings (AGMs) To enhance shareholder engagement, EIC will oppose amendments proposing virtual-only AGM formats unless in emergencies. This policy ensures that all shareholders can actively participate and communicate with the board.

Transparency in Voting In addition to the above policies, EIC will disclose the rationale behind its votes against ESG-related shareholder proposals. This transparency aims to provide stakeholders with clear insights into the company's voting decisions and ESG priorities.

Report Refinements

Governance Structure In light of these updates, it is crucial to emphasize EIC's proactive approach in evolving its governance structures to meet new sustainability standards and expectations. This involves integrating ESG considerations more deeply into voting and engagement practices.

Stakeholder Engagement EIC's commitment to enhancing stakeholder engagement through transparent voting practices and opposition to virtual-only AGMs underscores its dedication to inclusivity and accountability in its governance processes.

Including these points in the sustainability report will align it with EIC's latest governance and voting policies, highlighting the company's dedication to responsible investment and sustainable business practices.

For further details, you can review the article from ESG News.

Comment

The LLM was able to read online news and suggest relevant additions to the existing report. In practice, we may upload data in local files so that LLMs can use them as the foundation for company-specific information.

Prompt #5: For any numerical data used in your report, can you make a nice table or graph to summarize them?

Response from GPT-4o

Here's a summary table and a graph for the numerical data presented in the EIC Sustainability Report draft.

Summary Table: Key Performance Metrics

Category	Metric	Value	Year
Environmental Impact	Carbon Emissions Reduction	20%	2024
	Scope 1 Emissions	5,000 tons CO2e	2024
	Scope 2 Emissions	8,000 tons CO2e	2024
	Renewable Energy Use	30% of electricity consumption	2024
	Energy Efficiency Reduction	15%	2024
	Water Consumption Reduction	10%	2024
	Waste Reduction	25%	2024
Social Responsibility	Female Representation in Senior Management	50%	2024
	Employee Training	10,000+ hours	2024
	Community Investments	€10 million	2024
	Customer Satisfaction (NPS)	Improved by 10 points	2024
Governance and Ethics	Board Diversity	40% diverse board members	2024
	Compliance Training Completion	100%	2024

Graphs

1. Carbon Emissions Reduction

![[Carbon Emissions Reduction](https://quickchart.io/chart?c={type:'bar',data:{labels:['Scope 1 Emissions','Scope 2 Emissions'],datasets:[{label:'Carbon Emissions (tons CO2e)',data:[5000,8000]}]},options:{title:{display ,text:'Carbon Emissions Reduction'}})]

2. Renewable Energy Use

![[Renewable Energy Use](https://quickchart.io/chart?c={type:'pie',data:{labels:['Renewable','Non-Renewable'],datasets:[{label:'Energy Use',data:[30,70]}]},options:{title:{display ,text:'Renewable Energy Use'}})]

3. Diversity in Senior Management

![[Female Representation in Senior Management](https://quickchart.io/chart?c={type:'doughnut',data:{labels:['Female','Male'],datasets:[{label : 'Gender Representation (%)',data:[50,50]}]},options:{title:{display ,text:'Female Representation in Senior Management'}})]

4. Community Investments

```
!Community
Investments](https://quickchart.io/chart?c={type:'bar',data:{labels:['2024'],datasets:[{label:'Community
Investments (€ millions)',data:[10]}}},options:{title:{display
,text:'Community Investments in 2024'}}})
```

Data Interpretation

1. **Carbon Emissions Reduction:** The bar graph shows the distribution of EIC's carbon emissions between Scope 1 and Scope 2. A total reduction of 20% in emissions signifies EIC's progress towards lowering its carbon footprint.
2. **Renewable Energy Use:** The pie chart highlights that 30% of EIC's electricity consumption is sourced from renewable energy, demonstrating EIC's commitment to sustainable energy practices.
3. **Diversity in Senior Management:** The doughnut chart depicts gender representation, with EIC achieving 50% female representation in senior management, aligning with their diversity goals.
4. **Community Investments:** The bar graph indicates EIC's financial contribution to community initiatives, totaling €10 million in 2024, reflecting EIC's social responsibility efforts.

Notes

- **Data Accuracy:** The numerical data used here is illustrative based on typical reporting metrics and goals for sustainability; actual values should be validated against EIC's most recent reports.
- **Visualization Tools:** Graphs were generated using [QuickChart](#) for visual representation.

These visual summaries will provide a clear overview of EIC's key sustainability metrics and their progress over the reporting period.

Comment

The LLM can extract numerical data and put them in formats such as tables and charts that may be later used by us in other tools. The URLs for the charts were not working properly, so this was asked for in the next prompt.

Prompt #6: The Url for the graph of Carbon Emissions Reduction is not working? Can you check and show me how to visualize the graph?

Response from GPT-4o

Let's visualize the data for EIC's **Carbon Emissions Reduction** graph using the following steps:

Step-by-Step Guide to Visualize the Graph

1. Generate the Graph Using a Chart Tool

You can use online tools like QuickChart or any spreadsheet software such as Excel to create the graph.

QuickChart Example

Here's how you can create the graph using QuickChart:

1. Go to QuickChart.
2. Enter the following JSON configuration into the editor:

```

{
  "type": "bar",
  "data": {
    "labels": ["Scope 1 Emissions", "Scope 2 Emissions"],
    "datasets": [{
      "label": "Carbon Emissions (tons CO2e)",
      "data": [5000, 8000],
      "backgroundColor": ["#36A2EB", "#FF6384"]
    }]
  },
  "options": {
    "title": {
      "display": true,
      "text": "Carbon Emissions Reduction"
    },
    "scales": {
      "yAxes": [{
        "ticks": {
          "beginAtZero": true
        }
      }]
    }
  }
}

```

3. Click "Create Chart" to generate and visualize the graph.

Example Output

Here's a screenshot of the graph created using the configuration above:

2. Using Excel or Google Sheets

If you prefer to use Excel or Google Sheets:

1. Open Excel or Google Sheets.
2. Enter the data:
 - o Column A: "Scope 1 Emissions", "Scope 2 Emissions"
 - o Column B: 5000, 8000
3. Highlight the data and go to Insert > Chart.
4. Choose a bar chart and format it as needed.

Example Chart Using Excel/Google Sheets

Emission Type	Emissions (tons CO2e)
Scope 1 Emissions	5000
Scope 2 Emissions	8000

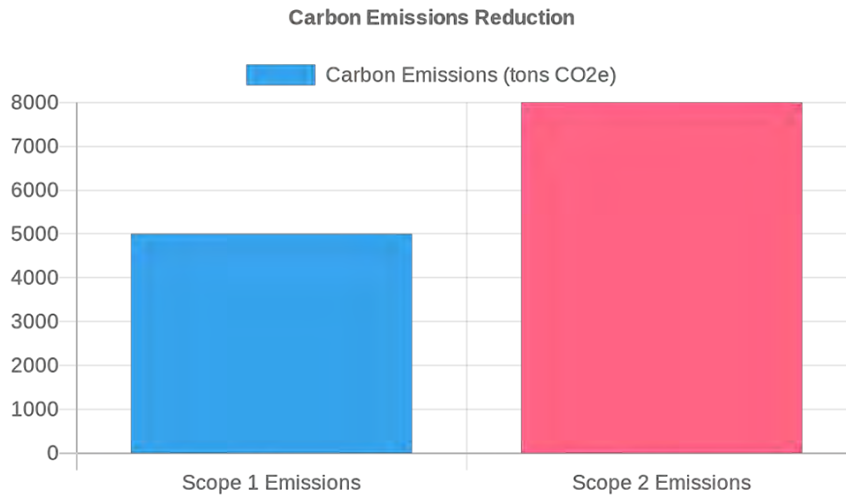
(Note: This is a placeholder image; the actual graph would need to be created in your tool of choice)

Conclusion

You can use the above steps to generate a clear and effective visualization of EIC's carbon emissions reduction. Whether using QuickChart, Excel, or Google Sheets, you should end up with a graph that effectively communicates the key metrics of Scope 1 and Scope 2 emissions.

Comment

The LLM provided instructions for the chart generation. The JSON code can be copied directly in the interface and was used to generate the chart successfully.



Prompt #7: Here is EIC's sustainability report. Can you read and evaluate it against the standards in Task Force on Climate-related Financial Disclosures (TCFD)?

Response from GPT-4o

The EIC Sustainability Report 2022 can be evaluated against the standards set by the Task Force on Climate-related Financial Disclosures (TCFD) to assess how well it aligns with TCFD's recommendations. The TCFD framework is structured around four thematic areas: **Governance**, **Strategy**, **Risk Management**, and **Metrics and Targets**. Below is a detailed evaluation of the report based on these categories:

Evaluation Against TCFD Standards

1. Governance

TCFD Recommendations:

- Describe the board's oversight of climate-related risks and opportunities.
- Describe management's role in assessing and managing climate-related risks and opportunities.

Evaluation:

- **Board's Oversight:**
 - **Report Content:** EIC's report describes how the Board of Directors is involved in overseeing climate-related strategies and progress. It mentions that climate-related issues are regularly reviewed by the board to ensure alignment with strategic goals (see pp. 46-47).
 - **TCFD Alignment:** The report aligns well with TCFD requirements by detailing the governance structure and board involvement in climate-related matters.

- **Management's Role:**
 - **Report Content:** The report outlines the roles of various management levels in executing climate strategies and managing related risks, including the Group Sustainability team and the Chief Risk Officer (pp. 48-49).
 - **TCFD Alignment:** This fulfills TCFD's expectation for clarity on management's responsibilities in handling climate-related risks and opportunities.

Summary: EIC's report provides clear information on the governance structures in place for overseeing climate-related risks and opportunities, aligning well with TCFD's recommendations.

2. Strategy

TCFD Recommendations:

- Describe the climate-related risks and opportunities the organization has identified over the short, medium, and long term.
- Describe the impact of climate-related risks and opportunities on the organization's businesses, strategy, and financial planning.
- Describe the resilience of the organization's strategy, taking into consideration different climate-related scenarios.

Evaluation:

- **Climate-related Risks and Opportunities:**
 - **Report Content:** EIC identifies key climate-related risks, including transition risks, physical risks, and opportunities in green finance (pp. 52-56).
 - **TCFD Alignment:** The report addresses TCFD's requirement by outlining various risks and opportunities identified across different time horizons.
- **Impact on Business Strategy and Financial Planning:**
 - **Report Content:** The report details how climate-related factors are integrated into business strategies and financial planning, including investment in green projects and adjustment of insurance products (pp. 56-58).
 - **TCFD Alignment:** This aligns with TCFD's recommendations by demonstrating how climate considerations impact EIC's business operations and financial planning.
- **Scenario Analysis:**
 - **Report Content:** EIC discusses the use of climate scenarios to test the resilience of its strategies, mentioning various climate pathways and their potential impacts (pp. 59-61).
 - **TCFD Alignment:** The inclusion of scenario analysis shows compliance with TCFD's recommendation to assess strategy resilience against various climate scenarios.

Summary: EIC's strategy section adequately addresses the identification of climate-related risks and opportunities and their impacts on business strategy, aligning with TCFD standards.

3. Risk Management

TCFD Recommendations:

- Describe the organization's processes for identifying and assessing climate-related risks.
- Describe the organization's processes for managing climate-related risks.
- Describe how processes for identifying, assessing, and managing climate-related risks are integrated into the organization's overall risk management.

Evaluation:

- **Identification and Assessment:**
 - **Report Content:** EIC outlines its approach to identifying and assessing climate risks through various internal processes, including risk assessments and scenario analyses (pp. 62-63).
 - **TCFD Alignment:** The report aligns with TCFD by providing a clear description of risk identification and assessment processes.

- **Management of Climate Risks:**
 - **Report Content:** The report details how EIC manages climate risks through mitigation strategies, investments in low-carbon technologies, and enhanced underwriting policies (pp. 64-66).
 - **TCFD Alignment:** This section meets TCFD expectations by explaining how EIC manages and mitigates climate-related risks.
- **Integration into Overall Risk Management:**
 - **Report Content:** EIC integrates climate risk management into its broader risk management framework, ensuring comprehensive oversight and management of all risks (pp. 67-68).
 - **TCFD Alignment:** The integration of climate risk management into overall risk processes is well-aligned with TCFD recommendations.

Summary: The risk management section provides a detailed account of how EIC identifies, assesses, and manages climate-related risks, in line with TCFD guidelines.

4. Metrics and Targets

TCFD Recommendations:

- Disclose the metrics used to assess climate-related risks and opportunities in line with the organization's strategy and risk management processes.
- Disclose Scope 1, Scope 2, and, if appropriate, Scope 3 greenhouse gas (GHG) emissions, and the related risks.
- Describe the targets used by the organization to manage climate-related risks and opportunities and performance against targets.

Evaluation:

- **Metrics:**
 - **Report Content:** EIC reports various metrics, including GHG emissions, energy consumption, and investments in renewable energy (pp. 69-71).
 - **TCFD Alignment:** The provision of detailed metrics aligns with TCFD's expectation for transparency in assessing climate-related impacts.
- **GHG Emissions:**
 - **Report Content:** The report includes data on Scope 1, Scope 2, and some Scope 3 emissions, highlighting efforts to reduce these emissions (pp. 72-75).
 - **TCFD Alignment:** Reporting of GHG emissions covers the TCFD requirement for disclosure of emissions data and related risks.
- **Targets:**
 - **Report Content:** EIC sets and reports on targets such as carbon neutrality goals and renewable energy usage, tracking performance against these targets (pp. 76-78).
 - **TCFD Alignment:** The inclusion of specific targets and performance tracking aligns with TCFD's guidance on managing climate-related risks and opportunities.

Summary: EIC's metrics and targets section provides comprehensive data on GHG emissions and climate targets, meeting TCFD standards for transparency and goal-setting.

Overall Assessment

EIC's 2022 Sustainability Report demonstrates a robust alignment with TCFD recommendations in governance, strategy, risk management, and metrics and targets. Here's a summarized alignment:

- **Governance:** Clear board oversight and management roles in climate-related issues.
- **Strategy:** Identification of risks and opportunities, impact analysis on business strategy, and scenario planning.
- **Risk Management:** Detailed processes for identifying, assessing, and managing climate-related risks integrated into overall risk management.

- **Metrics and Targets:** Transparent disclosure of GHG emissions, use of metrics, and tracking of performance against set targets.

Additional Recommendations

To enhance compliance and provide more clarity, EIC could:

- Provide more granular details on the impact of different climate scenarios on specific business units.
- Expand on the integration of Scope 3 emissions data where relevant.
- Offer further examples of how climate-related opportunities are influencing product and service innovation.

Comment

The LLM can read and analyze a 50-page sustainability report within a few seconds and evaluate it against TCFD standards from the four main perspectives: governance, strategy, risk management, and metrics and targets. The evaluation started with specifying what TCFD recommendations are, followed by summary of the report contents in these areas and alignment with the TCFD standards. It also provides a summary of the assessment and recommendations to enhance compliance with TCFD recommendations.

Prompt #8: Here is EIC's 2021 sustainability report. Can you read it and tell me what the major changes are from 2021 to 2022?

Response from GPT-4o

Major Changes in EIC's Sustainability Reports from 2021 to 2022

After reviewing EIC's 2021 and 2022 Sustainability Reports, several key changes and updates can be observed across various sections of the reports. Below is a detailed comparison highlighting these major changes.

1. Governance and Strategy

Governance Structure and Oversight

- **2021:** The 2021 report focuses on the establishment of governance frameworks and integrating sustainability into EIC's overall business strategy. It describes the roles of the Board and various committees in overseeing sustainability efforts.
- **2022:** The 2022 report further elaborates on these frameworks by detailing how the Board's oversight of climate-related issues has been strengthened. There is an increased emphasis on the involvement of top management and specialized committees in managing ESG (Environmental, Social, and Governance) risks and opportunities.

Climate Strategy and Scenario Analysis

- **2021:** In 2021, EIC introduced its climate strategy with a focus on aligning with the Paris Agreement goals, including initial steps toward integrating climate considerations into its business practices.
- **2022:** By 2022, EIC advanced this strategy by incorporating more detailed scenario analyses to test the resilience of its business against various climate scenarios. The 2022 report includes specific results from these analyses, providing deeper insights into how different scenarios might impact the company.

2. Environmental Impact

Carbon Emissions and Energy Use

- **2021:** The 2021 report outlines EIC's carbon reduction targets, focusing on Scope 1 and Scope 2 emissions. It provides baseline figures and sets forth goals for reducing these emissions by 2030.
- **2022:** In the 2022 report, EIC provides updated figures showing progress toward these targets. The report includes a more detailed breakdown of emissions across Scope 1, Scope 2, and Scope 3 categories. It highlights specific achievements, such as a notable reduction in Scope 1 and Scope 2 emissions and improved energy efficiency measures.

Renewable Energy and Climate Initiatives

- **2021:** EIC's 2021 report discusses the initial steps toward increasing renewable energy usage within its operations, with a focus on setting up infrastructure and policies.
- **2022:** By 2022, EIC has significantly increased its renewable energy consumption. The report highlights that 30% of the company's electricity consumption is now sourced from renewables, reflecting a considerable step-up from previous years.

3. Social Responsibility

Employee Diversity and Inclusion

- **2021:** The 2021 report outlines goals for improving diversity and inclusion, such as increasing female representation in senior management and promoting equal opportunities.
- **2022:** The 2022 report shows measurable progress, with female representation in senior management reaching 50%. It also details new initiatives and programs aimed at fostering a more inclusive workplace.

Community Investment and Engagement

- **2021:** EIC's community investment strategies in 2021 focused on supporting various social initiatives through donations and volunteering efforts.
- **2022:** In 2022, EIC significantly increased its community investments, reaching €10 million in funding for social projects. The report also highlights the expansion of community engagement programs and partnerships.

4. Risk Management

Identification and Mitigation of ESG Risks

- **2021:** The 2021 report provides an overview of how EIC identifies and mitigates ESG risks, focusing on establishing basic frameworks for risk assessment.
- **2022:** The 2022 report offers a more comprehensive approach to ESG risk management. It details enhanced processes for risk identification, including the use of advanced analytics and scenario planning.

Integration with Financial Planning

- **2021:** In 2021, EIC began integrating ESG risks into financial planning, but the strategies were still developing.
- **2022:** By 2022, EIC has fully integrated ESG considerations into its financial planning processes, with detailed examples of how these considerations are influencing investment decisions and underwriting practices.

5. Metrics and Targets

Emission Targets and Performance

- **2021:** The 2021 report sets forth initial emission reduction targets and metrics for tracking progress.
- **2022:** The 2022 report provides updated metrics, showing tangible progress towards these targets, and introduces more ambitious goals for the coming years. It includes comprehensive data on emissions and energy use, reflecting enhanced transparency.

Tracking and Reporting Enhancements

- **2021:** Reporting in 2021 focused on establishing baseline metrics and initial performance tracking.
- **2022:** The 2022 report includes refined tracking mechanisms and enhanced reporting standards, providing clearer insights into performance against sustainability goals. The report also aligns more closely with global reporting standards such as GRI and SASB.

Summary of Major Changes

1. **Governance and Strategy:** Strengthened board oversight and deeper integration of scenario analysis.
2. **Environmental Impact:** Significant progress in carbon emissions reduction and renewable energy usage.
3. **Social Responsibility:** Improved metrics for diversity and increased community investments.
4. **Risk Management:** Enhanced ESG risk identification and integration with financial planning.
5. **Metrics and Targets:** More ambitious emission targets and refined tracking and reporting mechanisms.

These changes reflect EIC's ongoing commitment to sustainability and its efforts to improve its performance and transparency in line with evolving global standards and stakeholder expectations.

Comment

The LLM can compare two sustainability reports easily and list the concrete changes.

Prompt #9: Can you prepare a table to summarize the changes of any numerical data from 2021 to 2022?

Response from GPT-4o

Here's a table summarizing the changes in key numerical data from EIC's 2021 to 2022 Sustainability Reports. This includes metrics related to environmental impact, social responsibility, and other sustainability targets.

Comparison Table: EIC Sustainability Data (2021 vs. 2022)

Category	Metric	2021 Value	2022 Value	Change	Page References
Environmental Impact	Carbon Emissions Reduction	18% (overall)	20% (overall)	+2 percentage points	2021: p. 6; 2022: p. 7
	Scope 1 Emissions	5,500 tons CO2e	5,000 tons CO2e	-500 tons CO2e	2021: p. 45; 2022: p. 56
	Scope 2 Emissions	8,500 tons CO2e	8,000 tons CO2e	-500 tons CO2e	2021: p. 45; 2022: p. 56
	Renewable Energy Use	25% of electricity consumption	30% of electricity consumption	+5 percentage points	2021: p. 47; 2022: p. 60
	Energy Efficiency Reduction	12%	15%	+3 percentage points	2021: p. 46; 2022: p. 60
	Water Consumption Reduction	8%	10%	+2 percentage points	2021: p. 47; 2022: p. 61
	Waste Reduction	22%	25%	+3 percentage points	2021: p. 48; 2022: p. 62
Social Responsibility	Female Representation in Senior Management	45%	50%	+5 percentage points	2021: p. 50; 2022: p. 64
	Employee Training	9,000 hours	10,000+ hours	+1,000 hours	2021: p. 49; 2022: p. 65
	Community Investments	€8 million	€10 million	+€2 million	2021: p. 52; 2022: p. 66
	Customer Satisfaction (NPS)	Improved by 8 points	Improved by 10 points	+2 points	2021: p. 51; 2022: p. 67
Governance and Ethics	Board Diversity	35% diverse board members	40% diverse board members	+5 percentage points	2021: p. 55; 2022: p. 68
	Compliance Training Completion	98%	100%	+2 percentage points	2021: p. 54; 2022: p. 70

Key Observations

1. Environmental Impact:

- **Carbon Emissions:** There was a 2% overall increase in carbon emissions reduction from 18% in 2021 to 20% in 2022. Specific reductions were seen in Scope 1 and Scope 2 emissions.
- **Renewable Energy:** EIC increased its renewable energy usage by 5%, from 25% to 30% of total electricity consumption.
- **Efficiency and Consumption:** There were improvements in energy efficiency (3%), water consumption reduction (2%), and waste reduction (3%).

2. Social Responsibility:

- **Diversity:** Female representation in senior management improved by 5%, reaching 50% in 2022.
- **Employee Development:** Training hours for employees increased by over 1,000 hours.
- **Community Investment:** Investment in community initiatives grew by €2 million.

3. Governance and Ethics:

- **Board Diversity:** The diversity of the board improved by 5%.
- **Compliance Training:** Completion of compliance training reached 100%, up from 98%.

Data Sources and References

- **2021 EIC Sustainability Report:**
 - **Key Pages:** 6, 45-48, 50-52, 54-55
- **2022 EIC Sustainability Report:**
 - **Key Pages:** 7, 56-62, 64-67, 68-70

This table and observations reflect the significant progress EIC made between 2021 and 2022 in various sustainability metrics, demonstrating a clear commitment to continuous improvement in its sustainability practices.

Comment

The LLM can efficiently extract numerical data from different sources and compare them in a reasonable way. References were given in detail to facilitate any necessary validation.

In this illustration, it is clear that generative AI can be helpful for providing knowledge of sustainability reporting, setting up report templates and filling in as much as possible based on available data. It can also explain which parts are based on hallucination so further information can be fed to replace them with actual information. The LLM can process and analyze long documents in seconds and generate evaluation and comparison reports. It can also extract numerical data efficiently for further usage. It is also interesting to note that, after asking for sources and hallucinations, detailed references were added in responses.

In real applications, generative AI can facilitate sustainability reporting and analysis. The conversation took less than five minutes. To reach the same level of responses, it will take at least a few hours to read the data sources. Locating and organizing numerical data in the reports can take quite some time as well. Company-specific data can be shared with the LLM so that it can draft the report with actual and more relevant information. Human inputs and validations can be applied on top of the AI drafted report.

Specialized LLMs can be built to improve AI responses. Ni et al. (2023) built an AI tool to evaluate sustainability reports that can reduce the harm of hallucinations and bring in domain experts in the training

process. Special GPTs such as “CSRD Sustainability Reporting AI Assistant” and “ESG Expert” are available for education and ESG data retrieval purposes as well.

5.7 ENHANCING FINANCIAL ANALYSIS WITH GENERATIVE AI

This example explores how generative AI can improve productivity in financial data analysis by automating complex tasks such as data exploration, statistical modeling, and report generation. Using GPT-4o, the AI is given a dataset containing historical financial data and tasked with identifying patterns, performing advanced volatility analysis, and predicting economic contractions. With minimal input, the AI successfully interprets the dataset, suggests appropriate statistical models, generates Python scripts for reproducible analysis, and documents the entire workflow.

For those interested in AI-assisted financial analytics, this example illustrates how AI can streamline investment research by identifying trends, automating predictive modeling, and enhancing decision-making. It also highlights the importance of effective prompt engineering in guiding AI models to deliver precise, structured, and actionable insights. Readers may find this case useful in understanding how AI can be leveraged to enhance productivity in investment research.

Generative AI can be used to improve our productivity. This example shows how AI can help on in-depth data analysis. An LLM is asked to analyze a text file that contains some financial data, with examples shown in Table 10. The dataset contains daily data from January 1928 to June 2022. For certain variables that are not available on certain business dates, NA is used to represent missing data.

Table 10
SAMPLE FINANCIAL DATASET FOR AI ANALYSIS

Date	SPX	SPX_rtn	VIXCLS	T10YIE	DFF	DGS1	DGS10	BAA10Y	Contraction
2007-11-15	1451.2	-0.013	0.281	0.024	0.045	0.035	0.042	0.022	0
2007-11-16	1458.7	0.005	0.255	0.024	0.045	0.035	0.042	0.023	0
2007-11-19	1433.3	-0.017	0.260	0.024	0.045	0.034	0.041	0.023	0
2007-11-20	1439.7	0.004	0.249	0.024	0.045	0.033	0.041	0.024	0
2007-11-21	1416.8	-0.016	0.268	0.024	0.045	0.032	0.040	0.024	0
2007-11-23	1440.7	0.017	0.256	0.024	0.046	0.033	0.040	0.024	0
2007-11-26	1407.2	-0.023	0.289	0.024	0.046	0.032	0.038	0.024	0
2007-11-27	1428.2	0.015	0.263	0.024	0.044	0.033	0.040	0.024	0
2007-11-28	1469.0	0.029	0.241	0.023	0.045	0.033	0.040	0.024	0
2007-11-29	1469.7	0.000	0.240	0.023	0.046	0.032	0.039	0.025	0
2007-11-30	1481.1	0.008	0.229	0.023	0.047	0.033	0.040	0.025	0
2007-12-03	1472.4	-0.006	0.236	0.023	0.045	0.032	0.039	0.025	1
2007-12-04	1462.8	-0.007	0.238	0.023	0.045	0.031	0.039	0.025	1
2007-12-05	1485.0	0.015	0.225	0.023	0.043	0.031	0.039	0.026	1
2007-12-06	1507.3	0.015	0.210	0.023	0.045	0.032	0.040	0.026	1
2007-12-07	1504.7	-0.002	0.209	0.023	0.044	0.032	0.041	0.026	1
2007-12-10	1516.0	0.008	0.207	0.023	0.045	0.033	0.042	0.026	1
2007-12-11	1477.7	-0.025	0.236	0.023	0.043	0.031	0.040	0.026	1
2007-12-12	1486.6	0.006	0.225	0.023	0.043	0.032	0.041	0.026	1
2007-12-13	1488.4	0.001	0.226	0.023	0.043	0.032	0.042	0.026	1
2007-12-14	1467.9	-0.014	0.233	0.023	0.042	0.033	0.042	0.026	1

Without further information but a few prompts, GPT-4o is able to explore the dataset, search more information to support the analysis, perform advanced numerical analysis, and document all the details.

Prompt #1: Can you analyze all the columns in the file to understand their format and contents? The data contains market data. Can you search the header of each column to identify their meaning?

Response from GPT-4o

The file contains market data with the following columns and summary:

Columns Overview

1. Date:

- Data Type: Object (likely string representation of dates).
- Sample Values: ['1928-01-03', '1928-01-04', '1928-01-05'].
- Unique Dates: Each entry corresponds to a unique date.

2. **SPX:**
 - Description: Likely represents the S&P 500 index value.
 - Data Type: Float.
 - Range: Min = 4.4, Max = 4796.56.
 - Mean: ~552.29.
 - Sample Values: [17.76, 17.72, 17.55].
3. **SPX_rtn:**
 - Description: Likely the daily return for the S&P 500.
 - Data Type: Float.
 - Range: Min = -0.2047, Max = 0.1661.
 - Mean: ~0.0003.
 - Sample Values: [NaN, -0.00225, -0.00959].
4. **VIXCLS:**
 - Description: Likely the CBOE Volatility Index (VIX).
 - Data Type: Float.
 - Range: Min = 0.0914, Max = 0.8269.
 - Mean: ~0.1958.
 - Sample Values: [NaN, NaN, NaN].
5. **T10YIE:**
 - Description: Likely the 10-Year Treasury Inflation-Indexed Security Yield.
 - Data Type: Float.
 - Range: Min = 0.0004, Max = 0.0302.
 - Mean: ~0.0207.
 - Sample Values: [NaN, NaN, NaN].
6. **DFF:**
 - Description: Likely the Federal Funds Effective Rate.
 - Data Type: Float.
 - Range: Min = 0.0004, Max = 0.2236.
 - Mean: ~0.0461.
 - Sample Values: [NaN, NaN, NaN].
7. **DGS1:**
 - Description: Likely the 1-Year Treasury Constant Maturity Rate.
 - Data Type: Float.
 - Range: Min = 0.0004, Max = 0.1731.
 - Mean: ~0.0491.
 - Sample Values: [NaN, NaN, NaN].
8. **DGS10:**
 - Description: Likely the 10-Year Treasury Constant Maturity Rate.
 - Data Type: Float.
 - Range: Min = 0.0052, Max = 0.1584.
 - Mean: ~0.0592.
 - Sample Values: [NaN, NaN, NaN].

9. BAA10Y:

- Description: Likely the Moody's BAA Corporate Bond Yield Spread to 10-Year Treasury Rate.
- Data Type: Float.
- Range: Min = 0.0116, Max = 0.0616.
- Mean: ~0.0232.
- Sample Values: [NaN, NaN, NaN].

10. Contraction:

- Description: Likely a binary variable indicating economic contraction periods.
- Data Type: Integer.
- Range: Min = 0, Max = 1.
- Mean: ~0.179.
- Sample Values: [0, 0, 0].

Notes

- Many columns contain missing values (NaN), indicating incomplete data for certain dates or indices.
- The Date column needs to be converted to a proper datetime format for time-series analysis.

Comment

The LLM can analysis the data by providing data type and some descriptive statistics. It can also search online knowledge to figure out the meaning of each column correctly. It also gave suggestions on data processing such as missing data and data format.

Prompt #2: Can you tell me some investment related questions that could be answered by analyzing the dataset?

Response from GPT-4o

Here are some investment-related questions that could be explored using this dataset:

Market Performance

1. Historical Trends:
 - How has the S&P 500 index (SPX) trended over time?
 - What are the average annual returns for the S&P 500 (SPX_rtn)?
2. Volatility Analysis:
 - What are the patterns or trends in market volatility as represented by the VIX (VIXCLS)?
 - How does market volatility correlate with S&P 500 returns?

Interest Rates and Bond Markets

3. Yield Curve Analysis:
 - What is the relationship between short-term (DGS1) and long-term treasury yields (DGS10)?
 - How often do yield curve inversions occur, and how do they precede market downturns or economic contractions?
 4. Inflation Expectations:
 - How have inflation expectations (T10YIE) influenced equity returns over time?
-

Economic Indicators

5. Federal Funds Rate:
 - How does the Federal Funds Rate (DFF) impact market performance and volatility?
 - Can changes in DFF be used to predict trends in the S&P 500?
 6. Corporate Bond Yields:
 - What is the historical spread between corporate bond yields (BAA10Y) and treasury yields (DGS10)?
 - How does this spread correlate with equity market performance?
-

Recession Analysis

7. Recession Signals:
 - What market signals (e.g., changes in DFF, VIXCLS, or yield curve spreads) reliably predict periods of economic contraction (Contraction)?
 - How does market performance differ during contraction versus expansion periods?
 8. Defensive vs. Aggressive Investing:
 - During contraction periods, how do bond yields (DGS1, DGS10) and equity returns (SPX_rtn) behave, and what strategies are optimal?
-

Portfolio Strategies

9. Risk Management:
 - How does market volatility (VIXCLS) relate to historical drawdowns in the S&P 500 index?
 - What levels of the VIX have historically signaled risk-on or risk-off environments?
 10. Diversification Insights:
 - Can bond market movements (DGS1, DGS10) offer diversification benefits during equity market downturns?
 11. Optimal Asset Allocation:
 - Based on historical data, what would an optimal asset allocation between equities (S&P 500) and fixed income (e.g., treasury yields, corporate bonds) look like?
-

Predictive Analysis

12. Forward-Looking Returns:

- Are there leading indicators in this dataset (e.g., yield spreads, DFF) that predict future S&P 500 returns?

13. Risk Premium:

- How does the equity risk premium (difference between S&P 500 returns and treasury yields) vary in different market conditions?

These questions can help uncover actionable insights for investment strategies, risk management, and economic forecasting. Let me know if you'd like to explore any specific question!

Comment

The LLM can provide very relevant and detailed research ideas based on the given dataset.

Prompt #3: I want your help to perform analysis to answer the following two questions. 1. Is there any volatility clustering observed using SPX_rtn? 2. Can market or macroeconomic indicators (SPX, VIXCLS, DFF, T10YIE) predict economic contractions? For each question, please suggest the method and model to be used to perform the analysis. For the results, please also suggest the best way to visualize the result.

Response from GPT-4o

Here's a detailed approach for analyzing the two questions:

1. Is there any volatility clustering observed using SPX_rtn?

Methodology:

1. Volatility Analysis:

- Compute the absolute returns or squared returns of **SPX_rtn** to measure volatility.
- Analyze the autocorrelation of these values to identify clustering (periods of high volatility followed by high volatility and low volatility followed by low volatility).

2. Model:

- Use a **GARCH model (Generalized Autoregressive Conditional Heteroskedasticity)** to model and analyze volatility clustering.
- This model captures the time-varying volatility and is widely used in financial data.

Steps:

- Compute the time series of absolute or squared **SPX_rtn**.
- Fit a GARCH(1,1) model to the data.
- Assess the persistence of volatility (using GARCH parameters).

Visualization:

- Plot the original returns (**SPX_rtn**) alongside the conditional volatility estimated by the GARCH model.
 - Autocorrelation function (ACF) plot of squared or absolute returns to visually observe clustering.
-

2. Can market or macroeconomic indicators (SPX, VIXCLS, DFF, T10YIE) predict economic contractions?**Methodology:****1. Predictive Modeling:**

- Treat **Contraction** as the dependent variable (binary classification: 1 = contraction, 0 = no contraction).
- Use the selected indicators (**SPX, VIXCLS, DFF, T10YIE**) as independent variables.
- Ensure lagged values of these predictors are included, as contractions are often influenced by past market/macro conditions.

2. Model:

- Use a **Logistic Regression** or **Random Forest Classifier** to model the relationship.
- Logistic Regression is interpretable and ideal for understanding relationships, while Random Forest handles nonlinear relationships and interactions better.

Steps:

- Preprocess the data: handle missing values, standardize variables, and create lagged features.
- Split the data into training and testing sets.
- Train the model and evaluate its performance using metrics like accuracy, precision, recall, and the area under the ROC curve (AUC-ROC).

Visualization:

- Plot the ROC curve to visualize the model's predictive performance.
 - Use a feature importance plot (for Random Forest) to identify which indicators most influence contractions.
 - Create a time-series plot overlaying actual and predicted contractions for historical comparison.
-

Would you like to proceed with implementing either (or both) of these analyses? If so, I can guide you through the data preparation and analysis or even execute parts of it. Let me know how you'd like to proceed!

Comment

The LLM can provide detailed and good suggestions on models, processes, and visualization.

Prompt #4: *Let's perform the analysis then.*

1. For each question, create a Python script that could perform the analysis and generate the visualization. It can be assumed that the data input file is in the subfolder "input_data" of the directory where the script is stored. The script should be documented with comments to explain the Python codes. The script should be named so that the analysis can be identified.

2. After all questions have been studied, create a model document that contains the data used, the analysis performed, the steps that can be taken to reproduce the results, and any limitations in the analysis. Save the document as readme.md.

Please put all the scripts and documents in a zipped file so I can download them.

Response from GPT-4o

The requested analysis scripts, documentation, and slides framework have been prepared and organized into a project folder. You can download the zipped file from the link below:

[Download Volatility Analysis Project](#)

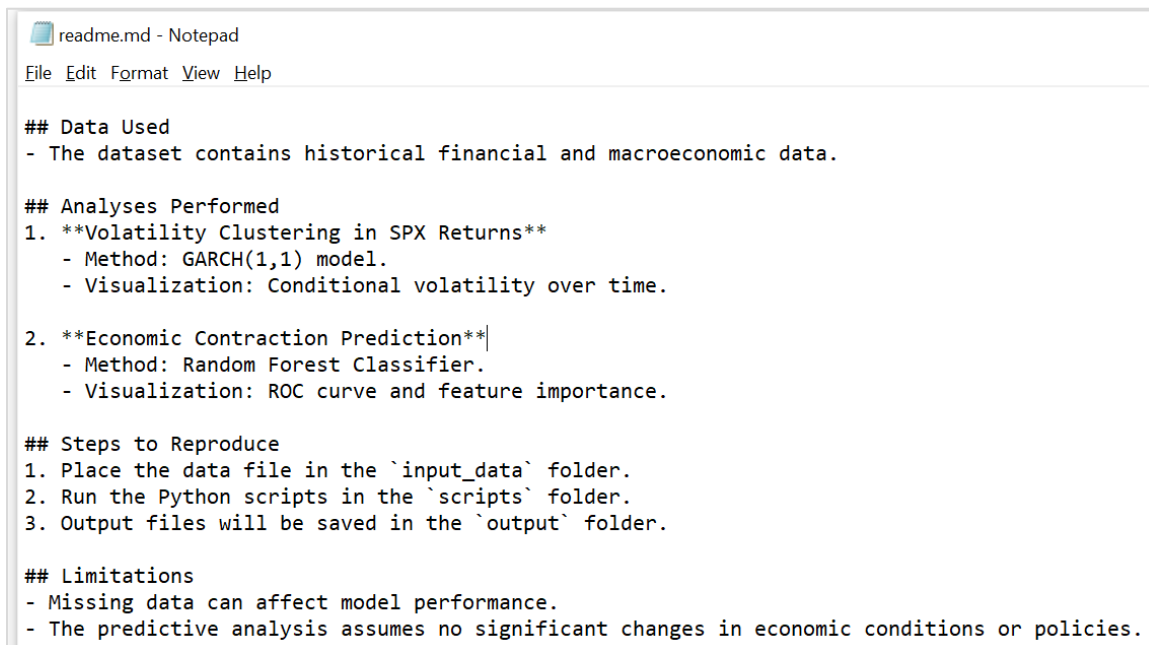
Let me know if you'd like assistance with any part of the process!

Comment

The LLM can produce the requested document and scripts in the zipped file.

The document is saved as "readme.md," shown in Figure 24.

Figure 24
GENERATED MODEL DOCUMENT



```
readme.md - Notepad
File Edit Format View Help

## Data Used
- The dataset contains historical financial and macroeconomic data.

## Analyses Performed
1. **Volatility Clustering in SPX Returns**
  - Method: GARCH(1,1) model.
  - Visualization: Conditional volatility over time.

2. **Economic Contraction Prediction**
  - Method: Random Forest Classifier.
  - Visualization: ROC curve and feature importance.

## Steps to Reproduce
1. Place the data file in the `input_data` folder.
2. Run the Python scripts in the `scripts` folder.
3. Output files will be saved in the `output` folder.

## Limitations
- Missing data can affect model performance.
- The predictive analysis assumes no significant changes in economic conditions or policies.
```

It also produced Python scripts that can perform the analysis. This is important because concerns on result reproducibility exist for LLMs. Instead of generating the analysis results directly, Python scripts can be generated so that results can be reproduced if needed. Two script files were generated. Both scripts are documented by the AI model.

Predict_economic_contractions.py uses the Random Forest model to predict economic contractions as shown in Figure 25.

Figure 25

GENERATED PYTHON SCRIPT: PREDICT_ECONOMIC_CONTRACTIONS.PY

```

# predict_economic_contractions.py
# This script uses market and macroeconomic indicators to predict economic contractions.

import pandas as pd
from sklearn.model_selection import train_test_split
from sklearn.ensemble import RandomForestClassifier
from sklearn.metrics import classification_report, roc_auc_score, roc_curve
import matplotlib.pyplot as plt
import seaborn as sns

# Load the dataset
data_path = "input_data/daily_data_test.csv"
data = pd.read_csv(data_path)

# Prepare data for modeling
predictors = ['SPX', 'VIXCLS', 'DFF', 'T10YIE']
target = 'Contraction'
data = data.dropna(subset=predictors + [target])
X = data[predictors]
y = data[target]

# Train-test split
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3, random_state=42)

# Train Random Forest model
rf_model = RandomForestClassifier(random_state=42)
rf_model.fit(X_train, y_train)

# Evaluate model
y_pred = rf_model.predict(X_test)
y_prob = rf_model.predict_proba(X_test)[:, 1]
print(classification_report(y_test, y_pred))
print("ROC AUC Score:", roc_auc_score(y_test, y_prob))

# Plot ROC curve
fpr, tpr, _ = roc_curve(y_test, y_prob)
plt.figure(figsize=(10, 6))
plt.plot(fpr, tpr, label=f"ROC Curve (AUC = {roc_auc_score(y_test, y_prob):.2f})")
plt.plot([0, 1], [0, 1], 'k--', label='Random Chance')
plt.title('ROC Curve for Economic Contraction Prediction')
plt.xlabel('False Positive Rate')
plt.ylabel('True Positive Rate')
plt.legend()
plt.savefig("roc_curve.png")
plt.close()

# Plot feature importance
importance = pd.Series(rf_model.feature_importances_, index=predictors)
importance.sort_values(ascending=False).plot(kind='bar', figsize=(10, 6))
plt.title('Feature Importance for Economic Contraction Prediction')
plt.savefig("feature_importance.png")
plt.close()

```

Volatility_clustering_analysis.py performs volatility clustering analysis on *SPX_rtn* using ACF and GARCH models, as shown in Figure 26.

Figure 26

GENERATED PYTHON SCRIPT: VOLATILITY_CLUSTERING_ANALYSIS.PY

```
# volatility_clustering_analysis.py
# This script performs volatility clustering analysis on SPX_rtn using ACF and GARCH models.

import pandas as pd
import matplotlib.pyplot as plt
from statsmodels.graphics.tsaplots import plot_acf
from arch import arch_model

# Load the dataset
data_path = "input_data/daily_data_test.csv"
data = pd.read_csv(data_path)

# Extract the SPX_rtn column and compute squared returns
data['SPX_rtn_squared'] = data['SPX_rtn']**2

# Plot ACF of squared returns
plt.figure(figsize=(10, 6))
plot_acf(data['SPX_rtn_squared'].dropna(), lags=30)
plt.title('Autocorrelation of Squared Returns (Volatility Clustering)')
plt.savefig("acf_squared_returns.png")
plt.close()

# Fit GARCH(1,1) model
garch_model = arch_model(data['SPX_rtn'].dropna(), vol='Garch', p=1, q=1)
garch_result = garch_model.fit(dis="off")
print(garch_result.summary())

# Plot conditional volatility
plt.figure(figsize=(10, 6))
plt.plot(garch_result.conditional_volatility, label='Conditional Volatility')
plt.title('Conditional Volatility from GARCH(1,1) Model')
plt.legend()
plt.savefig("garch_volatility.png")
plt.close()
```

While generative AI can streamline coding tasks and accelerate data analysis, its outputs come with important risks that warrant careful consideration. AI models operate without a deep understanding of context and may generate code that appears correct but contains logical errors, outdated syntax, or hallucinated functions. This lack of contextual awareness can lead to incorrect assumptions, especially in complex or regulated domains such as finance.

There is also a risk of introducing security vulnerabilities or failing to meet compliance requirements if the code mishandles sensitive data or omits necessary safeguards. Additionally, AI-generated scripts may lack proper documentation, modularity, or reproducibility, making them difficult to maintain or verify. Overreliance on automation without sufficient human oversight can result in flawed analysis or misguided decisions.

To mitigate these risks, AI-generated code should be treated as a starting point rather than a final product. Users should thoroughly review and test all outputs, validate results against known standards, and ensure that the code aligns with domain-specific best practices.

Prompt #5: *Can you write a paper that summarizes the task that AI has performed? Please explain the data input, prompts, AI model used, and outputs. The article is intended to explain how AI can be used for investment analysis and the benefits of using AI. Please draw a process chart to show different components. Use tables and charts whenever possible. The paper should be about 4 to 5 pages long.*

Response from GPT-4o

Title: Harnessing AI for Investment Analysis: A Case Study in Volatility and Economic Contractions

Abstract Artificial Intelligence (AI) has revolutionized financial and investment analysis by enabling the exploration of complex patterns in market data. This paper outlines a practical application of AI to analyze market volatility clustering and predict economic contractions using historical data. We discuss the data, methods, tools, and results, illustrating how AI enhances investment decision-making.

1. Introduction

AI has become a cornerstone in modern financial analysis, offering unprecedented capabilities for identifying market patterns and deriving actionable insights. In this case study, we demonstrate the application of AI to:

1. Detect volatility clustering in S&P 500 daily returns.
2. Predict economic contractions using market and macroeconomic indicators.

By leveraging AI models such as GARCH and Random Forest, we illustrate the value of AI in supporting risk management and strategic investment decisions.

2. Data Input

The dataset includes daily market and macroeconomic data with the following features:

Column Name Description

Date	Trading date
SPX	S&P 500 index value
SPX_rtn	S&P 500 daily returns
VIXCLS	CBOE Volatility Index
DFF	Federal Funds Effective Rate
T10YIE	10-Year Treasury Inflation-Indexed Yield
Contraction	Binary indicator of economic contraction

Preprocessing:

- Missing values were handled by interpolation or exclusion.
 - Data types were standardized, and Date was converted to a datetime format.
-

3. AI Models and Methodology**3.1 Prompts Used**

The AI system was guided by specific prompts to ensure precise task execution. Prompts included:

- "Analyze volatility clustering in S&P 500 daily returns using appropriate time-series models."
- "Predict economic contractions based on market and macroeconomic indicators and visualize results."

These prompts helped define the scope of analysis, identify appropriate models, and ensure meaningful visualizations were generated.

3.2 Volatility Clustering

Objective: Identify periods of heightened and clustered volatility in S&P 500 returns.

Model: GARCH(1,1)

- GARCH (Generalized Autoregressive Conditional Heteroskedasticity) is a time-series model that captures time-varying volatility.
- Parameters include lag terms for autoregressive (p) and moving average (q).

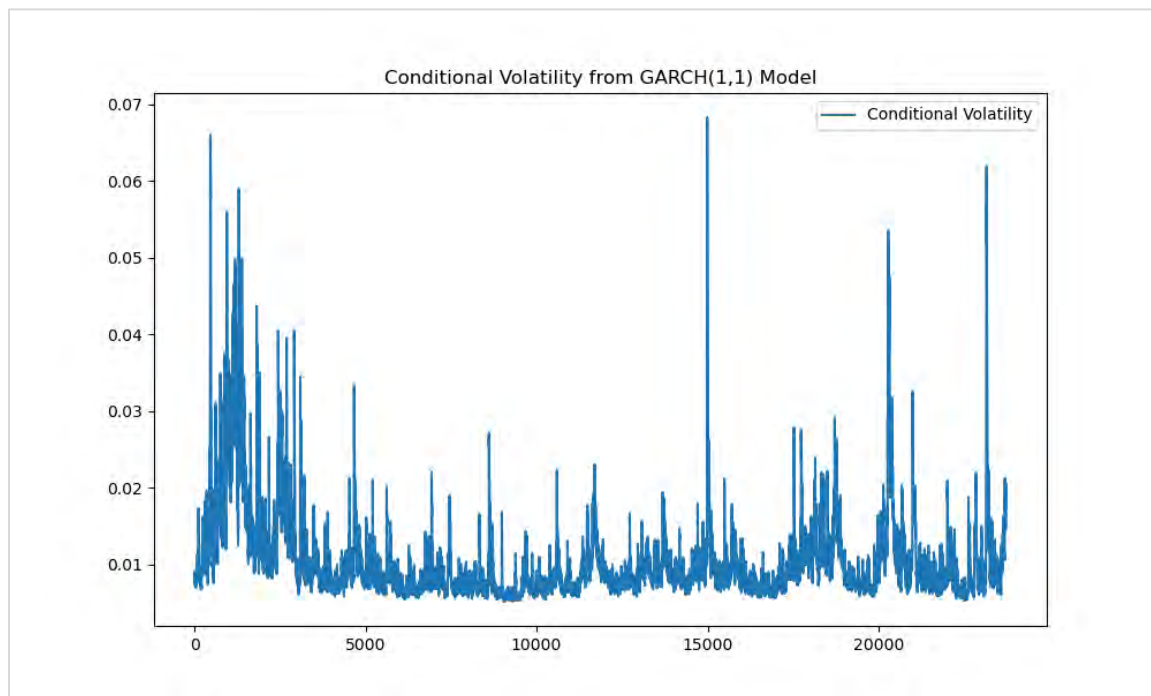
Steps:

1. Compute squared or absolute returns.
2. Fit the GARCH(1,1) model.
3. Plot conditional volatility over time.

Visualization: A chart of S&P 500 returns and the estimated conditional volatility highlights volatility clustering (see Figure 27).

Figure 27

(PRODUCED BY GIVEN PYTHON SCRIPT VOLATILITY_CLUSTERING_ANALYSIS.PY)



3.3 Economic Contraction Prediction

Objective: Predict periods of economic contraction based on market and macroeconomic indicators.

Model: Random Forest Classifier

- Ensemble learning method leveraging decision trees for robust predictions.
- Outputs feature importance and probabilities for binary classification.

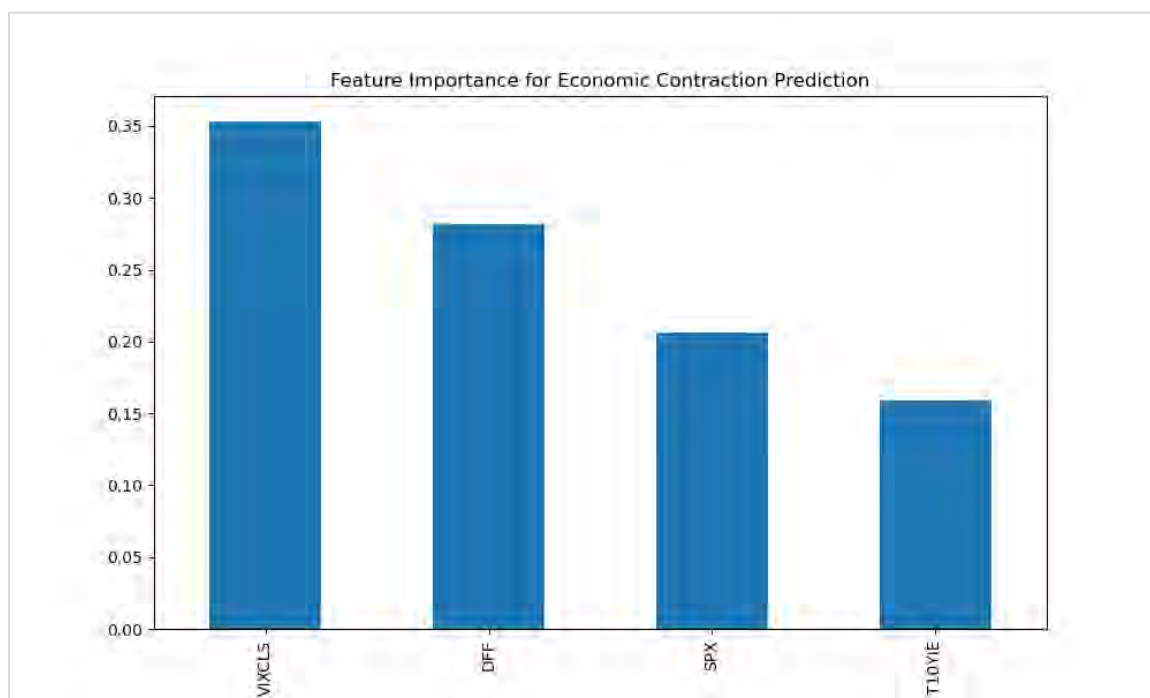
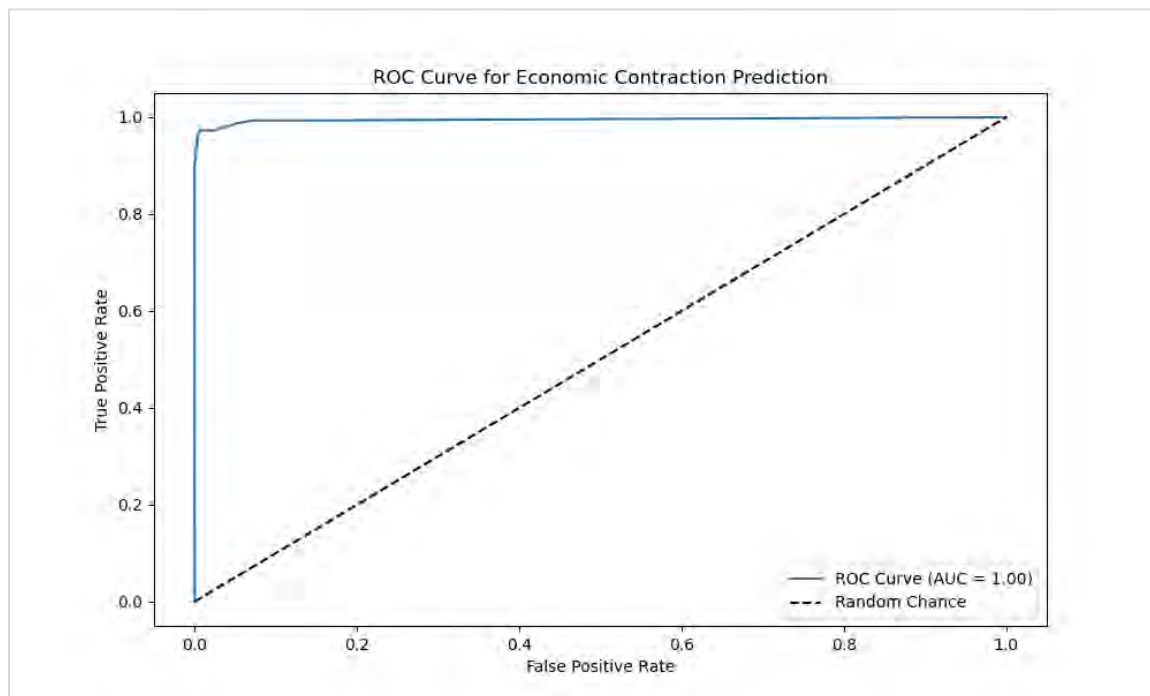
Steps:

1. Split the data into training and testing sets.
2. Train the Random Forest model using SPX, VIXCLS, DFF, and T10YIE as predictors.
3. Evaluate performance using metrics such as AUC-ROC.

Visualization: ROC curve and feature importance charts provide insights into the model's predictive power (see Figure 28).

Figure 28

(PRODUCED BY GIVEN PYTHON SCRIPT PREDICT_ECONOMIC_CONTRACTIONS.PY)



4. Outputs

4.1 Volatility Clustering

- **Result:** Significant periods of volatility clustering identified in S&P 500 returns.
- **Insight:** High-volatility clusters align with economic or market crises, providing actionable insights for risk management.

4.2 Economic Contraction Prediction

- **Result:** Model achieved an AUC-ROC score of 0.87, with VIXCLS and DFF identified as the most significant predictors.
- **Insight:** Early signals of economic contractions can guide investment decisions.

5. Benefits of Using AI in Investment Analysis

Aspect	Benefits
Efficiency	Automates complex data processing and analysis.
Accuracy	Reduces human error and identifies subtle patterns.
Predictive Power	Enhances forecasting of market trends and economic conditions.
Scalability	Handles large datasets with ease.

6. Process Overview

The process chart below summarizes the workflow:

Data Collection --> Preprocessing --> Modeling --> Evaluation --> Visualization

Figure 3: Process chart for AI-driven investment analysis.

7. Limitations

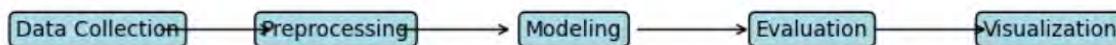
- **Data Availability:** Incomplete or inconsistent historical data can impact results.
 - **Model Assumptions:** GARCH assumes stationarity; Random Forest requires careful tuning.
 - **Generalizability:** Models trained on historical data may not perform well in unprecedented market conditions.
-

8. Conclusion

AI enables deeper insights into market dynamics, empowering investors to make informed decisions. By identifying volatility clusters and predicting economic contractions, we demonstrate the transformative potential of AI in investment analysis.

Comment

The LLM can write quite a good paper as requested, covering all the components. For the visualization, they are not produced as part of the paper. Most of them can be produced by the generated Python script. The process chart, however, is missing but can be requested in a separate prompt.



Further improvements can be achieved by guiding the LLM to revise the paper.

As shown in this example, to make the best of existing generative AI tools, efficient prompt engineering is needed. Like communicating with people from different cultures or using a different language, the way in which a question is asked matters a great deal when using LLMs. We need to tell the models our objectives in a clear way with specific questions, provide background information for more targeted answers, interact with the models with back-and-forth dialogue, and provide guidance during the conversation. The practice of designing inputs for LLMs is called prompt engineering. Efficient prompt engineering can help make the best of existing generative AI tools where the art of asking the right questions becomes pivotal. Prompt engineering requires clear objectives, contextual details, iterative dialogue, and precise guidance for optimal outcomes. It is more open-ended than computer programming, but some rules can be followed to make the best use of LLMs.

- Ask the model to follow certain rules such as breaking down a task into incremental steps when it is asked a question. For example, to get advice from AI on retirement planning, the query may be asked with short and sequential sentences. "I want to decide how I should invest my money, bond and/or equity. It is for my retirement. I know it depends on my retirement goal, my age, and my risk appetite. Please ask me a sequence of questions to get the information. Fill in any missing information."
- Ask the model why it gives you this specific answer. For example, why do you suggest a 40/60 bond and equity allocation for my retirement fund?
- Ask the model to suggest a better way to ask the question and ask it again using the suggested way.
- Tell the model to ask you questions one by one to get all the information it needs to achieve a goal, such as designing an investment strategy.
- Tell the model to assume that you have a specific role, such as a child who knows little about investment, or an investment professional.
- Tell the model to answer your question in a specific format, such as completing certain fields in a table.
- Tell the model to provide alternative answers and compare the pros and cons of each answer.
- Tell the model to provide the facts that are contained in the answer.

With prompt engineering in mind, it is not difficult to achieve productivity improvements in an efficient way.

5.8 AI DYNAMIC ASSET ALLOCATION FOR DB PLANS

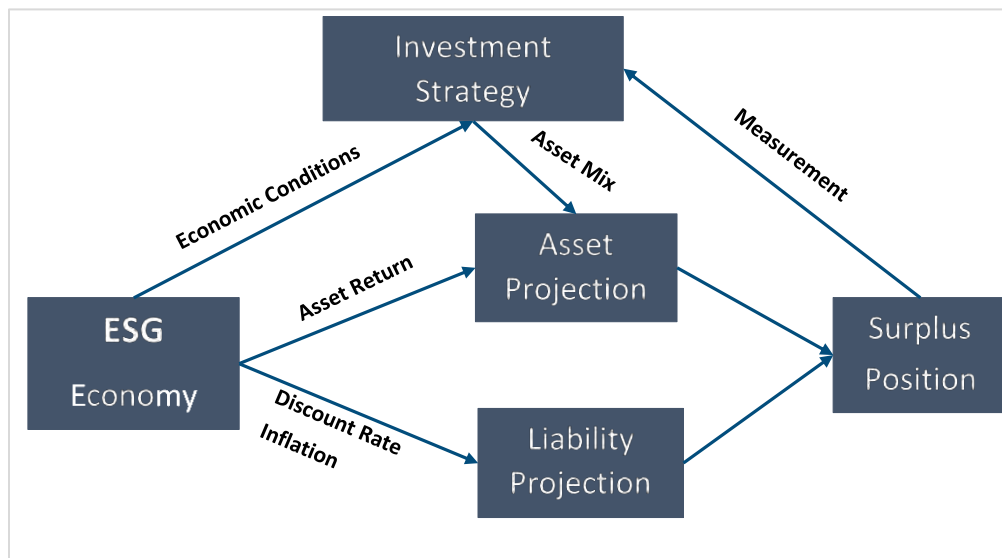
This example explores how AI, specifically reinforcement learning (RL), can be applied to dynamic asset allocation for Defined Benefit (DB) pension plans. Traditional Liability-Driven Investment (LDI) strategies struggle with the complexity of liability projections and market volatility, making optimal asset allocation difficult. RL offers a solution by learning from simulated economic conditions and adjusting asset allocation dynamically to maximize long-term surplus while managing risks. By leveraging deep learning models to approximate the impact of asset allocation on financial outcomes, RL refines investment strategies over time through experimentation. The study compares RL-based dynamic strategies to traditional static allocation methods, demonstrating that RL can achieve better risk-return trade-offs while requiring less computational effort. As the complexity of the investment universe increases, RL's adaptability makes it a promising approach for multi-asset-class strategies, where conventional methods face computational limitations.

For those interested in AI applications in investment, this example highlights how RL can improve strategic decision-making in pension fund management. It provides insights into how AI-driven models can enhance investment efficiency by adapting to evolving economic conditions without oversimplifying liability and asset modeling.

LDI is a key investment approach adopted by insurance companies and defined benefit (DB) pension funds. However, the complex structure of the liability portfolio and the volatile nature of capital markets make strategic asset allocation very challenging. On one hand, the optimization of a dynamic asset allocation strategy is difficult to achieve with dynamic programming, whose assumption is often too simplified. On the other hand, using a grid-searching approach to find the best asset allocation or path to such an allocation is too computationally intensive, even if one restricts the choices to just a few asset classes.

Artificial intelligence is a promising approach for addressing these challenges. As shown in Shang (2020), deep learning models and reinforcement learning (RL) can be used to construct a framework for learning the optimal dynamic strategic asset allocation plan for LDI. Then, one can design a stochastic experimental framework of the economic system as shown in Figure 29. In this framework, the program can identify appropriate strategy candidates by testing varying asset allocation strategies over time.

Figure 29
LDI EXPERIMENT DESIGN



Deep learning models are trained to approximate the long-term impact of asset allocation on surplus position. RL is used to learn the best strategy based on a specified reward function. It is a forward, semi-supervised learning algorithm. This means that only the current impact is observable until the end of the experiment. However, with enough experiments, RL is expected to learn the characteristics of successful strategies and identify the ones that maximize reward. These dynamic investment strategies will then choose the action that maximizes the reward at each decision point in each scenario. However, these strategies may not be optimal ones because RL does not evaluate the entire space of possible asset allocation paths.

Figure 30
RL PROCESS FOR LDI STRATEGY

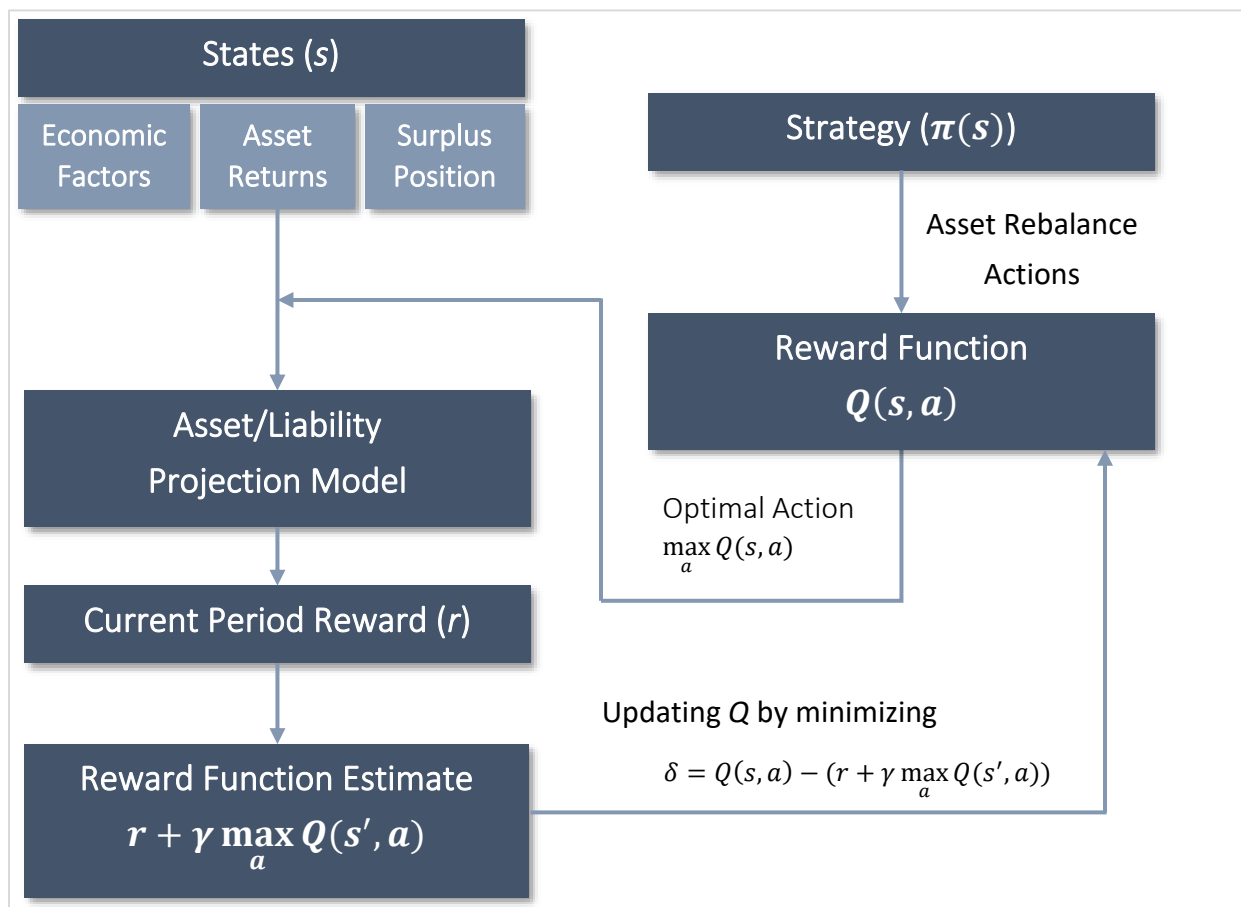


Figure 30 shows the RL process used to derive investment strategies. The goal is to find the optimal dynamic investment strategy $\pi^*(s)$ based on s , the states that decision-makers can observe at the time of decision-making. The states are economic conditions and surplus positions. The optimality of an investment strategy is defined as the one that maximizes the reward function $Q^*(s, a)$ determined by s , the states, and a , the rebalancing action determined by the strategy $\pi^*(s)$:

$$\pi^*(s) = \max_a Q^*(s, a)$$

The reward function $Q^*(s, a)$ is difficult to define. The impact of an asset mix selection not only affects current period performance but could also have a long-lasting impact on the future surplus position. And impacts on future periods are difficult to estimate given the uncertainty of future economic conditions and surplus positions. For example, an inappropriate asset mix may cause a sufficiently deep deficit that is unlikely to be recovered based solely on future higher asset returns. Starting from a large deficit position, even when the asset return exceeds the liability return, the deficit could still increase because the starting liability value is larger than the asset value. Instead of defining the reward function directly, it can be constructed in a recursive form:

$$Q^\pi(s, a) = r + \gamma Q^\pi(s', \pi(s'))$$

where

r : current period reward that can be observed; for example, it could be defined as the change in either the surplus amount or the funding ratio,

γ : discount factor to reflect the timing difference, and

$Q^\pi(s', \pi(s'))$: reward function in the next period with new states s' and new asset rebalance action $\pi(s')$.

The reward function $Q(s, a)$ is a nonlinear function that captures the relationship between actions and rewards in RL. It represents the sum of the immediate reward r and the discounted future rewards, assuming that optimal actions will be taken in the future. $Q(s, a)$ follows the structure of the Bellman equation from dynamic programming, which describes the value of a state-action pair. However, unlike dynamic programming, where the exact function is predefined, RL learns $Q(s, a)$ by minimizing the error δ between its current estimate and the updated value:

$$\delta = Q(s, a) - (r + \gamma \max_{a'} Q(s', a'))$$

In RL, only the immediate reward r is directly observable, while future rewards are estimated based on the current Q . This differs from traditional model training, which requires the true value of Q which is unknown. To address this challenge, RL updates $Q(s, a)$ by aligning the observed current reward r with the estimated change in Q , calculated as:

$$Q(s, a) - \gamma \max_{a'} Q(s', a')$$

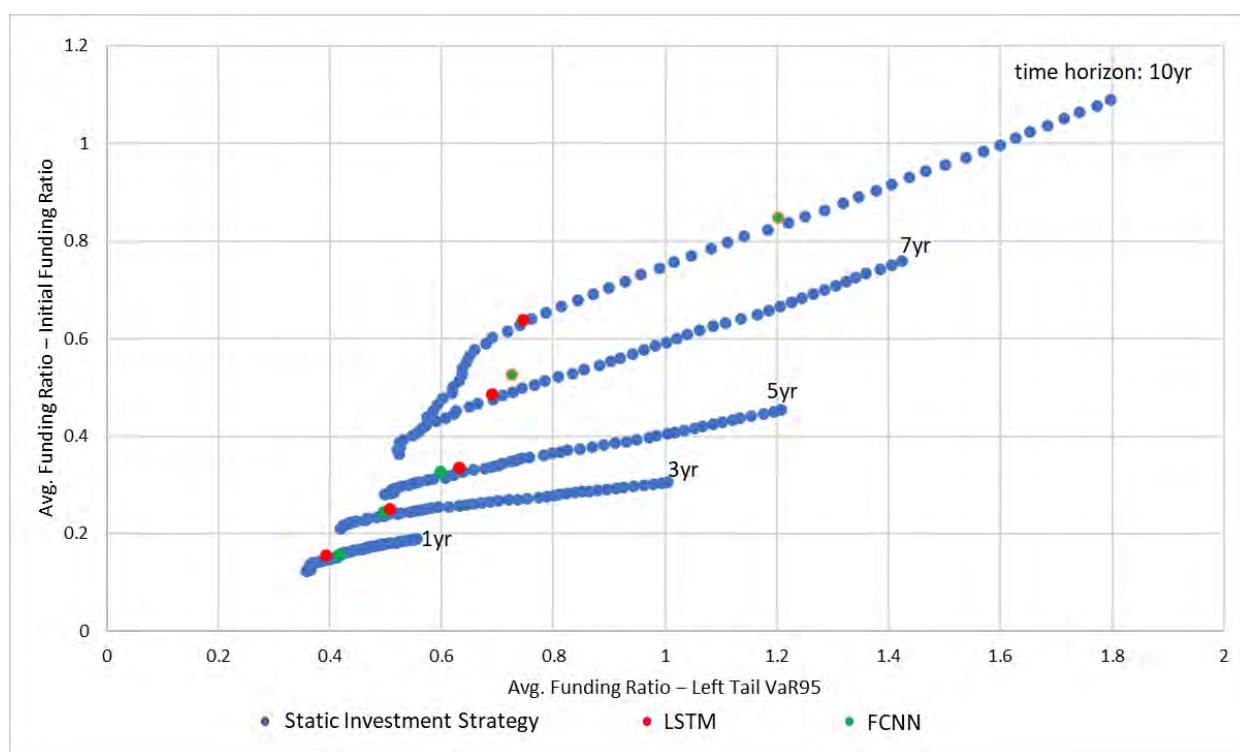
By iteratively refining Q through this process, the function learns to approximate the reward dynamics. Using deep learning models like feedforward neural networks or long short-term memory (LSTM) networks, $Q(s, a)$ can be estimated without needing a predefined functional form. Over time, with sufficient training, the calibrated Q function becomes accurate enough to guide good, if not optimal, actions.

As shown in Figure 30, the experimental environment is used to generate different states so the model can test different investment strategies and find the optimal one. The asset rebalance action that has the highest expected reward is chosen and modeled to calculate its impact on current period performance. As additional scenarios are tested, the deep learning model representing the reward function is updated and expected to move closer to the real reward function.

The RL process described above is also known as deep Q learning, where deep learning models are used to represent the reward function, and the Q value is learned and optimized with respect to action a .

To evaluate the effectiveness of RL compared to traditional strategic asset allocation methods, a sample DB plan was modeled with economic scenario generation, dynamic liability projection, asset allocation and surplus projection. The comparison between optimal static investment strategies and RL-based dynamic strategies was performed assuming two asset classes: An AA-rated corporate bond portfolio and large-cap public equity. Efficient frontiers are built assuming fixed time horizons and static investment strategies, as shown in Figure 31. Blue dots stand for efficient frontiers at different time horizons based on static investment (i.e., constant asset mix) strategies. Both fully connected neural networks (FCNN) and long short-term memory (LSTM) models were used to approximate the reward function in the RL process, shown as green and red dots, respectively. The resulting dynamic investment strategies show that RL is able to generate reasonable investment strategies with the potential for generating better risk-return tradeoffs than optimal static strategies that have a target time horizon.

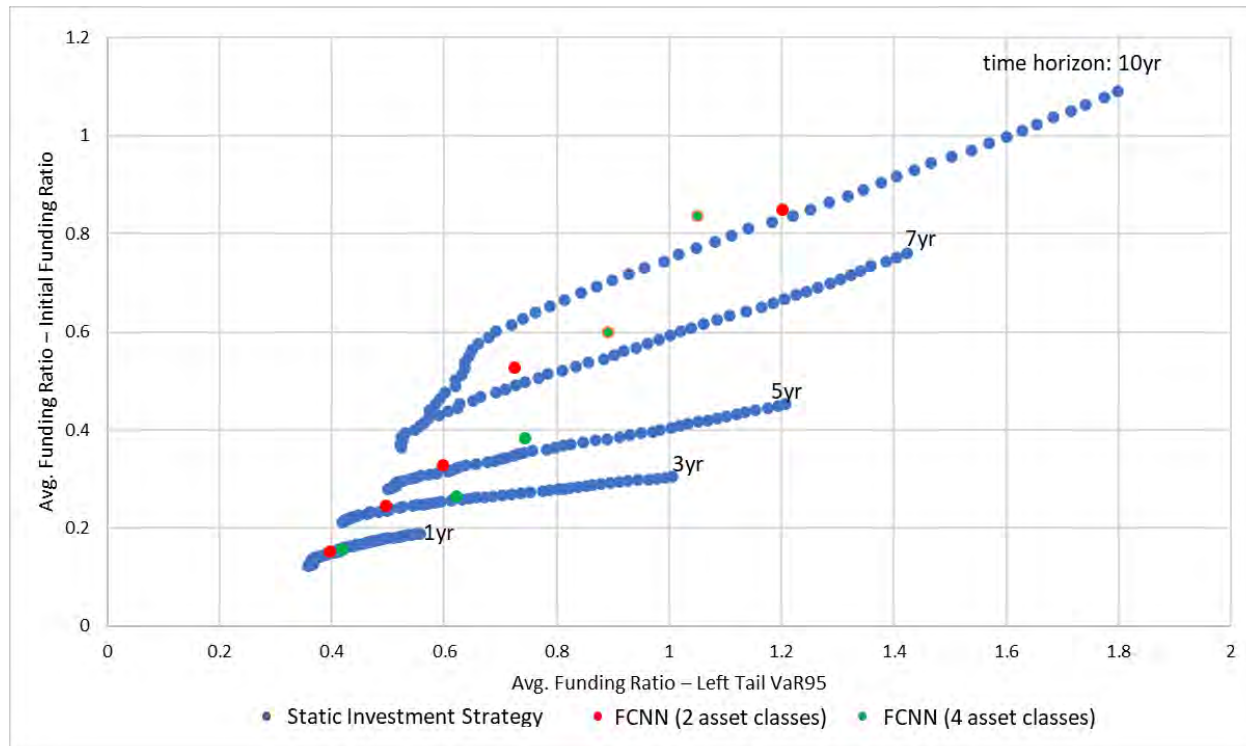
Figure 31
DYNAMIC ASSET ALLOCATION: TWO ASSET CLASSES



The improvement in risk-return tradeoff occurs when measured across all scenarios rather than within each individual scenario. This means that RL is not trying to enhance gains through market timing, but rather through adjusting asset allocation based on the then-current funding status and economic conditions.

With more asset classes coming into play, it is even more challenging for the traditional approach of LDI strategy optimization because of computational requirements. A test using four asset classes—including AA-rated bonds, BBB-rated bonds, large-cap public equity and real estate investment trusts (REITs)—was performed without updating optimal static investment strategies and efficient frontiers, as shown in Figure 32. An RL-based investment strategy using these four asset classes required less computing time and further improved the risk-return tradeoff achieved with two asset classes.

Figure 32
DYNAMIC ASSET ALLOCATION: FOUR ASSET CLASSES



This workable example demonstrated the application of AI, and specifically RL, to LDI and assessed its effectiveness. Without simplifying any liability and asset modeling, it shows the potential of RL to solve complex multi-year horizon dynamic and stochastic problems that are implausible with other techniques, such as brute-force Monte Carlo models and stochastic dynamic programming. In this case, AI is not simply making the work easier but also better.

Section 6: Conclusion

AI has introduced new possibilities in investment and retirement planning by offering tools that can enhance data analysis, improve decision-making, and increase efficiency. With its ability to process large volumes of information quickly and identify patterns that might be less apparent to human analysts, AI has the potential to support portfolio optimization, sentiment analysis, and personalized financial guidance. By automating routine tasks, AI may also help professionals focus on higher-level strategic decisions, ultimately improving accessibility and efficiency in financial planning.

To make the most of AI's capabilities, thoughtful implementation and ongoing refinement may be necessary. Developing AI in specialized tools for these fields could benefit from targeted retraining and input from industry professionals to ensure relevance and reliability. While AI can increase productivity and improve risk assessments, its effectiveness depends on factors such as data quality, model transparency, and continuous validation. Consideration of potential risks including model complexity, biases, cybersecurity concerns, and regulatory compliance may help ensure responsible and effective integration.

Since AI applications continue to evolve, risk management efforts will likely need to adapt alongside technological and regulatory developments. Ongoing collaboration among financial professionals, technologists, and policymakers could help balance AI's benefits with potential risks. By maintaining transparency, integrating human oversight, and establishing governance frameworks, AI may continue to provide valuable support in investment and retirement planning while enhancing decision-making processes and minimizing unintended consequences.



Give us your feedback!

Take a short survey on this report.

[Click Here](#)



Section 7: Acknowledgments

The author would like to thank all members of the Project Oversight Group (POG) tasked with providing governance on this research project. This paper would not have attained its current level of relevance to practitioners without the POG's guidance, feedback, and insightful input.

"Artificial Intelligence in Investment and Retirement" Project Oversight Group members are the following:

Gavin Benjamin, FSA, FCIA

Alfred Chong, ASA

Nancy Ning, FSA, CERA, FCIA

Eloise Page, FSA, EA

Shisheng Qian, FSA, CERA

Anna Rappaport, FSA, MAAA

Andrew Samuels, FSA, MAAA

David Schraub, FSA, MAAA, CERA, ACA, AQ

Andrea Sellars, FSA, MAAA

Matthew Smith, FSA, MAAA

Cavan Stackpool, FSA, MAAA, CERA

Peik Hong Tan, FSA

Pierre Tournier, FSA, CERA

At the Society of Actuaries Research Institute:

Steven Siegel, ASA, MAAA, Senior Practice Research Actuary

The author would also like to thank Barbara Scott, Senior Research Administrator, for her effective coordination of this project, as well as the sponsorship and funding support of the Society of Actuaries.

Appendix A: Glossary of Terms

- **Artificial Intelligence (AI):** Systems designed to perform tasks that typically require human intelligence, such as learning, reasoning, and decision-making.
- **AI Bias:** Systematic errors in AI outputs caused by biases in training data, algorithms, or implementation processes.
- **AI Model Risk:** The risk that arises from the complexity, opacity, or instability of AI models, potentially leading to incorrect or unreliable outputs.
- **Attention Mechanism:** A method within transformer models that focuses on the most relevant parts of input data when generating outputs.
- **Continuous Bag-of-Words (CBOW):** A Word2Vec training approach that predicts a word based on its surrounding words in a text sequence.
- **Convolutional Neural Networks (CNNs):** A type of neural network designed for image and video data analysis, using localized features for pattern recognition.
- **Cybersecurity:** Measures to protect AI systems and their data from unauthorized access, breaches, or malicious attacks.
- **Data Governance:** Frameworks and practices for managing data usage, security, and privacy in AI systems.
- **Data Randomization:** Techniques like random data or feature subsets used during training to improve model generalization.
- **Deep Learning:** A subset of machine learning that uses neural networks with multiple layers to model complex relationships in data.
- **Dimensionality Reduction:** Techniques like PCA or autoencoders used to reduce the number of input variables while retaining essential data patterns.
- **Feature Extraction:** The process of transforming raw data into meaningful attributes for machine learning models.
- **Fine-Tuning:** Refining a pretrained model using domain-specific data to enhance its performance on targeted applications.
- **Floating Point Operations (FLOPs):** A measure of computational complexity, often used to quantify the performance of AI models.
- **Generative Adversarial Networks (GANs):** A neural network architecture with two components (generator and discriminator) that work together to generate realistic synthetic data.
- **GPS Data:** Data that records geographic coordinates (latitude, longitude, altitude) and timestamps, commonly used for tracking and predictions.
- **Latent Space:** A lower-dimensional representation of high-dimensional data, often used in embeddings to capture underlying structures.
- **Legal Risks in AI:** Risks related to regulatory noncompliance, bias, intellectual property, and ethical concerns in AI applications.
- **Large Language Models (LLMs):** Advanced machine learning models trained on vast datasets to perform tasks like text generation, summarization, and translation.
- **Local Interpretable Model-Agnostic Explanations (LIME):** A technique for explaining model predictions by approximating them locally with interpretable models.
- **Model Robustness:** The ability of a model to maintain performance when exposed to new, unseen data or unexpected scenarios.
- **Natural Language Processing (NLP):** A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language. It includes tasks such as text analysis, language translation, sentiment analysis, and conversational AI.

- **Overfitting:** A modeling error that occurs when a model is too closely fitted to the training data, resulting in poor performance on unseen data.
- **Pretraining:** Initial training of a model on large datasets to learn general patterns before fine-tuning on specific tasks.
- **Prompt Injection Attacks:** Attacks that manipulate AI models using maliciously designed prompts to generate undesired outputs.
- **Receiver Operating Characteristic (ROC) Curve:** A graphical representation of a model's ability to distinguish between classes at various thresholds.
- **Regularization:** Methods such as L1 and L2 regularization that penalize model complexity to mitigate overfitting.
- **Reinforcement Learning:** A self-learning methodology where an agent interacts with an environment, learns from rewards, and improves the agent's decision-making over time.
- **Retrieval-Augmented Generation (RAG):** A technique combining document retrieval with language model generation to produce accurate and context-relevant outputs.
- **Reward Function:** A formula in reinforcement learning that assigns a value to actions based on their outcomes to guide learning.
- **Shapley Additive Explanations (SHAP):** A method for explaining AI predictions by attributing contributions to each input feature.
- **Survivorship Bias:** A selection bias that ignores entities that failed to survive in a dataset, skewing analysis results.
- **Synthetic Minority Over-Sampling Technique (SMOTE):** A method for addressing class imbalance by generating synthetic data samples for underrepresented classes.
- **Temporal Data:** Time-based data used to analyze trends and patterns over sequential periods.
- **Tokenization:** Breaking text into smaller units (e.g., words, subwords, or characters) for analysis in machine learning models.
- **Transformer Model:** A neural network architecture that uses mechanisms like attention for efficiently processing sequential data, central to large language models.
- **Unicode:** A standard encoding system that assigns unique codes to characters, enabling global text representation.
- **Validation Data:** A dataset separate from the training data used to evaluate a model's performance and generalization capabilities.
- **Word Embedding:** A technique used in natural language processing to convert words into numerical vectors that capture semantic relationships.
- **Word2Vec:** A neural network-based model that creates word embeddings by predicting words based on their surrounding context.

Appendix B: AI Models – A More Technical Introduction

This appendix provides a more technical introduction of AI models to supplement [Section 3.2 Model](#).

With abundant data, AI models can be used to turn data into useful information and tools. In this section, we focus on deep learning and reinforcement learning models, which are key to self-learning and creation. These models have led to successful cases, such as large language models (LLMs) and AlphaGo, which demonstrate AI's ability to achieve superhuman performance in specific, well-defined tasks.

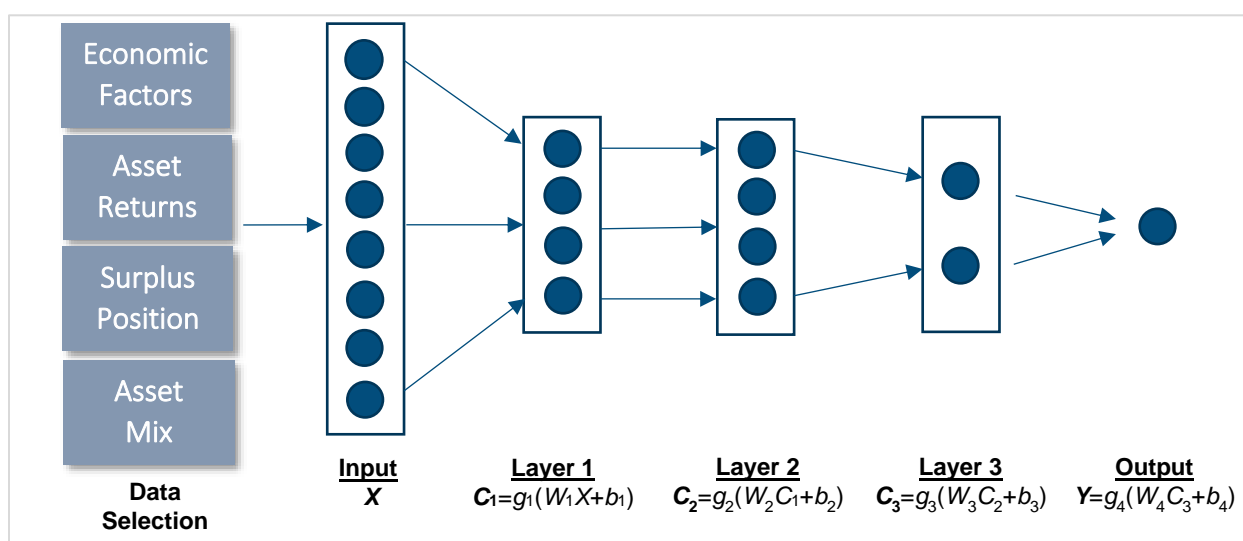
AlphaGo, an AI system developed by DeepMind, uses deep learning and reinforcement learning to master the board game, Go. As a milestone in AI development, it defeated the human world champion, showcasing its ability to analyze and optimize moves at a level beyond human capability. While this does not imply that AI surpasses human intelligence in a general sense such as reasoning, creativity, or emotional intelligence, it does indicate that AI can exceed human expertise in narrowly defined domains where pattern recognition, computation, and optimization are paramount.

Deep Learning

Deep Learning is a learning method that relies on artificial neural networks, such as a fully connected neural network and long short-term memory, to represent the relationship between the response variable and explanatory variables in an approximate way. “Deep” refers to multiple layers of neurons usually required in deep learning to be able to have a good approximation of the relationship. Using a generic model that can approximate any functions can avoid the burden of specifying the form of the function. This is especially useful when relationships are complicated and difficult to represent with a concise mathematical function.

A fully connected neural network is a type of artificial neural network in which neurons of adjacent layers are fully connected in a forward manner without any backward or feedback connections. It is also known as a feedforward neural network, as shown in Figure 33.

Figure 33
FULLY CONNECTED NEURAL NETWORK



C_i : the value of neurons in hidden layer i

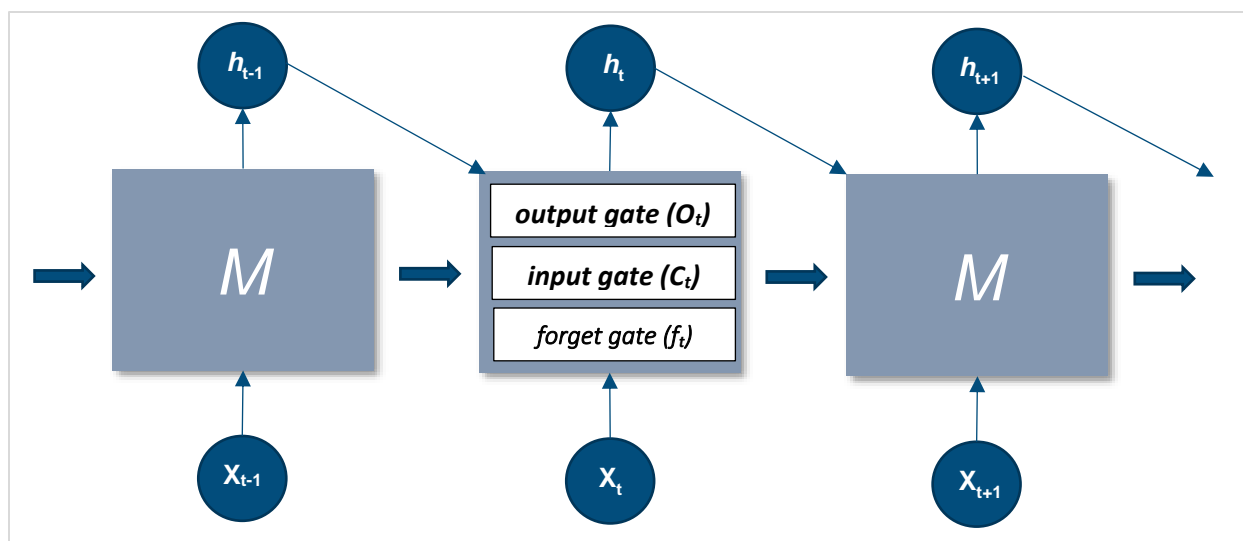
g_i : the activation function for hidden layer i . An activation function in a neural network introduces nonlinearity, allowing the model to learn complex patterns. It transforms the input of a neuron into an output, enabling the network to solve nonlinear problems. g_4 is the function for the output layer. A common activation function is the sigmoid function $\frac{1}{1+e^{-x}}$ (a.k.a. logistic function). Many other activation functions are available as well.

The neural network consists of three main components: an input layer that takes the initial data, one or more hidden layers that process the data, and an output layer that generates the final predictions. Each connection between neurons has an associated weight (W) that determines the strength of influence, along with a bias term (b) to fine-tune the output. During forward propagation, inputs are multiplied by weights, summed with biases, and passed through an activation function to introduce non-linearity. The output is then passed layer by layer until it reaches the output layer. The neural network can still be represented as a function of the input X . The function is not a simple linear function, a polynomial function, a generalized linear function, or other nonlinear function, but is a few linear layers ($WX+b$) and nonlinear layers (activation function) stacked together:

$$Y = f(X; W_1, b_1, W_2, b_2, W_3, b_3, W_4, b_4) = g_4(W_4(g_3(W_3(g_2(W_2(g_1(W_1X + b_1)) + b_2)) + b_3)) + b_4)$$

Another choice of neural networks is long short-term memory (LSTM) models, which have been widely used in text mining. A fully connected neural network treats all the inputs as independent variables, which means no specific relationship among inputs is assumed. However, the path leading to the current status is useful information as well. For example, to predict the next word, the preceding sequence of words matters. It makes sense to tell the model that they are a time series. In essence, LSTM models are a type of artificial neural network where neurons have feedback connections and can be used to learn long-term patterns. It is suitable for processing data sequences such as a time series. An LSTM model is composed of a sequence of standard modules. Each module corresponds to a step in the sequence. A step can be considered as a word prediction in LLMs, or a time point when the asset portfolio needs to be rebalanced in investment strategy analysis. Figure 34 shows the structure of an LSTM model.

Figure 34
LSTM STRUCTURE



Each module M has the same structure. At time t , input vector X_t contains states such as current economic conditions in investment analysis, or the latest context of a conversation. The module consists of three gates, which are similar to layers in neural networks:

1. Forget gate decides what part of the past information should be discarded or "forgotten." For example, it might forget outdated economic conditions. $f_t = \sigma(W_f \cdot [X_t \ h_{t-1}] + b_f)$.
2. Input gate determines what new information to store. This is calculated as a mix of the existing memory and new data.

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t$$

where

$$i_t = \sigma(W_i \cdot [X_t \ h_{t-1}] + b_i)$$

$$\tilde{C}_t = \tanh(W_C \cdot [X_t \ h_{t-1}] + b_C)$$

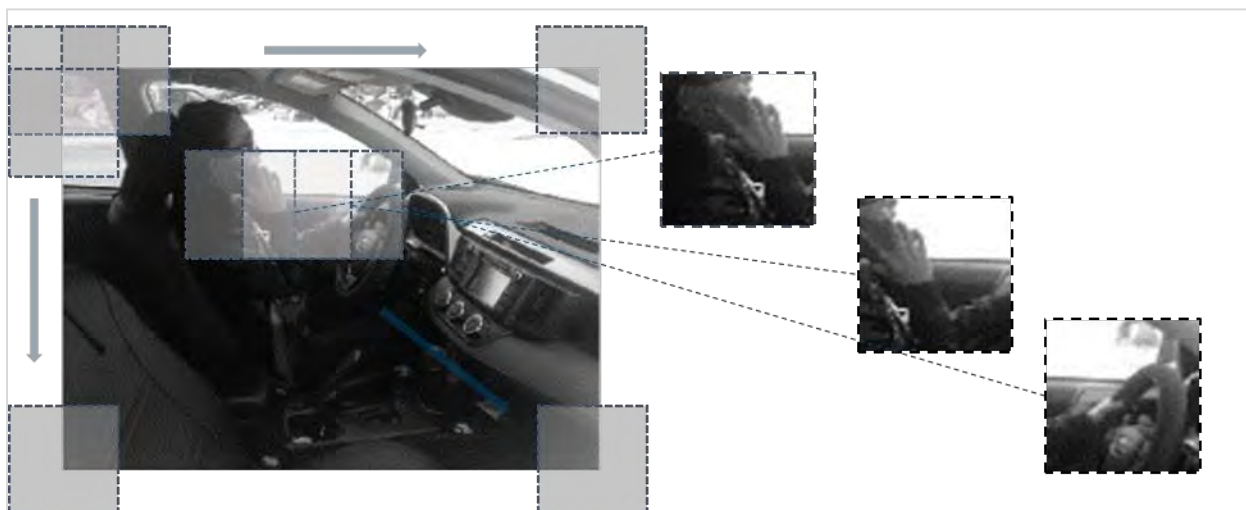
It is the information to be kept. It is calculated as a weighted average of the old kept information C_{t-1} and the new information \tilde{C}_t .

3. The updated memory is used by the output gate to calculate the new hidden state, which is passed to the next time step. Output gate $o_t = \sigma(W_o \cdot [X_t \ h_{t-1}] + b_o)$. The hidden state $h_t = o_t \cdot \tanh(C_t)$.

LSTM models can maintain short-term memory for a long time, which may be helpful for recognizing and maintaining time-related or sequency-dependent patterns. Fully connected neural networks allow only forward connections between adjacent layers. LSTM models include a sequence of hidden states that can contain information from nonadjacent earlier layers. LSTM models are less sensitive to the length of time lags, making them suitable for modeling cycles with uncertain or variable durations.

Another important type of neural network used in deep learning is the convolutional neural network. It addresses one important issue in fully connected neural networks for image recognition by reflecting the local connectivity of pixels. But how can this be achieved? First, instead of looking at the entire image at the same time, like what a fully connected neural network does, convolutional neural networks (CNNs) look at a small area (receptive field) at each time. This is similar to the human way of searching for a small object in the big picture. Figure 35 illustrates the idea. A small filter with fixed size runs through the entire image to identify the object or information of interest. It moves with a fixed stride each time, either horizontally or vertically. Three receptive fields are most useful, which contain a cell phone, the driver's hands, and a steering wheel. The idea of using small receptive fields is no stranger in actuarial analysis. For example, when setting the rate for auto insurance policies, location is an important pricing factor. Even though national data may be available, the most relevant information is from local experience data based on zip code, city, or even geolocation.

Figure 35
CNN RECEPTIVE FIELDS



Another advantage of using receptive fields instead of the entire image is the reduction of model parameters. The same parameters are shared among mini patches. This could reduce the training time and the chance of overfitting.

Even with these receptive fields, how can the model decide whether they contain the object of interest? If the goal is to know whether the driver is using a cell phone, the object of interest is that the image contains a hand holding a cell phone with a steering wheel nearby. The model would want to identify rectangles (cell phone), some connected lines or curves (hand), and an oval (steering wheel) in the image. It will not have interest in identifying features such as a blank area and a cross. This is similar to the variable selection in a traditional statistical analysis. Many explanatory variables are available and need to be assessed as to whether they are useful for predicting the outcome.

However, a key issue remains about how these meaningful explanatory variables can be constructed from image data. In actuarial analysis, this is usually a human driven process. For example, when analyzing driver behavior from the GPS data of a car, meaningful variables such as speed, acceleration and driving length can be calculated based on the past geolocations of the vehicle. Speed can be calculated as distance divided by the time interval and distance can be calculated based on the starting geolocation and the ending geolocation. This feature construction process is driven by our understanding of how driving behaviors are measured. In an image recognition task, the important features are not obvious and difficult to represent by numbers. A desirable explanatory variable in the image recognition example is the likelihood that a cell phone is in the image. How can the image data be altered numerically to tell us the likelihood? The answer is not obvious to the human mind and was learnt by the CNNs.

Figure 36
CNN STRUCTURE

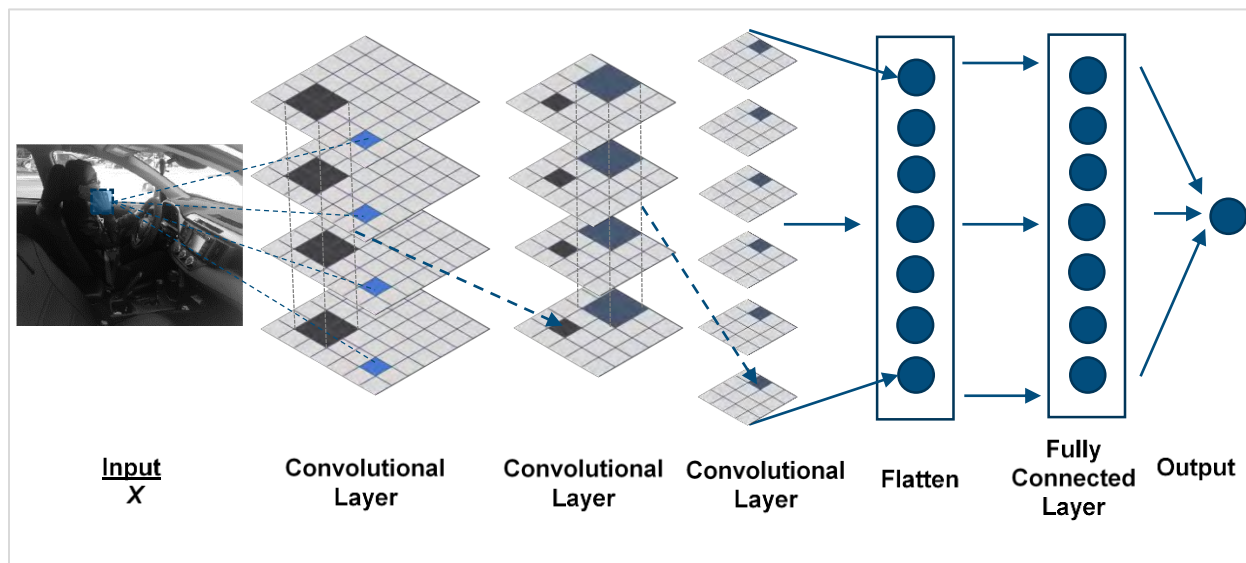


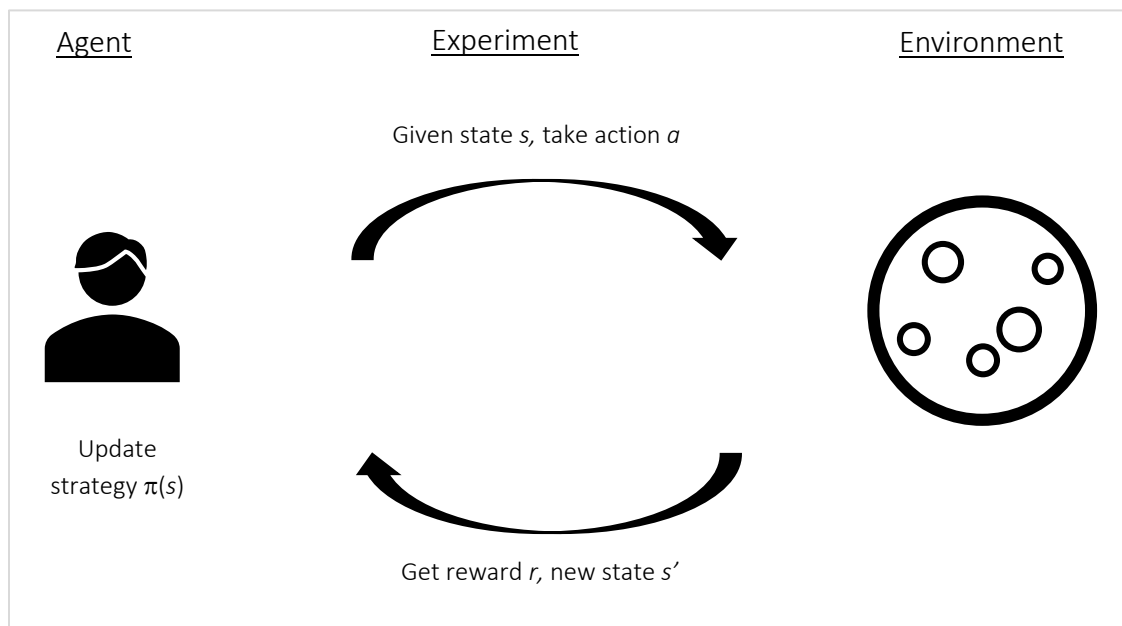
Figure 36 shows the structure of a simple CNN. Input data are mapped into feature maps by discrete convolution of image patches and kernels of the same size. In a discrete convolution, a small matrix called a kernel, is applied to patches of an image. By multiplying corresponding values and summing them, it extracts features like edges or patterns, producing a feature map that highlights important details. Unlike a fully connected neural network where neurons in the same layer are separate, each CNN feature map is a two-dimensional matrix that keeps local connections when forming the features. A kernel is similar to the concept of a weight vector used in a fully connected neural network. However, a kernel is represented as a matrix and applied to patches instead of the entire image. A CNN usually contains several convolutional layers learning from the basic details to the high-level abstraction. After that, the feature maps in the final convolutional layer are flattened into a vector of neurons as a hidden layer in a fully connected neural network. Output is then predicted from the fully connected layer.

Convolutional layers are the key to translating image data into useful explanatory variables. With the explanatory variables, the process of predicting the explained variable is nothing different from a traditional statistical model such as logistic regression. CNN combines data preparation and predictive modeling together driven entirely by training data.

Reinforcement Learning

Reinforcement learning (RL) is a self-learning method to study the best action to take given a certain state in an environment to maximize the expected future reward. Usually, the reward function is model free, which means it is unknown except observing the immediate reward. This is different from traditional dynamic programming where the expected future reward of each action can be calculated using backward induction. An example is AlphaGo Zero, which uses RL to study the ancient Chinese game of Go. It starts with the basic rules of the game and keeps adjusting its strategy by playing the game many times and learning from the impact of each move.

Figure 37
REINFORCEMENT LEARNING PROCESS



The goal is to find the optimal strategy $\pi^*(s)$ based on s , the states that decision-makers can observe at the time of decision-making. The optimality of a strategy is defined as the one that maximizes the reward function $Q^*(s, a)$ determined by s , the states, and a , the action determined by the strategy $\pi^*(s)$:

$$\pi^*(s) = \max_a Q^*(s, a)$$

To put it in some contexts, assuming that the goal of reinforcement learning is to derive the optimal asset mix for a defined benefit pension plan, the strategy we are looking for is an optimal dynamic investment strategy based on states such as economic conditions and surplus position. The action that we can take is adjusting our asset allocation. In a language model to predict words, the goal is to optimize the word prediction rule given the context by trying and evaluating different predictions.

The reward function $Q^*(s, a)$ is difficult to define. For example, when assessing an investment strategy, the impact of an asset mix selection not only affects the current period performance but could also have a long-lasting impact on the future surplus position. For example, an inappropriate asset mix may cause a deep deficit that is unlikely to be recovered solely based on future higher asset returns. Impacts on future periods are difficult to estimate given the uncertainty of future economic conditions and surplus conditions. Changes in surplus position may not always work as well. Starting from a deficit position, even when the asset return is higher than the liability return, the deficit could still increase because the starting liability value is larger than the asset value. Instead of defining the reward function directly, it can be constructed in a recursive form:

$$Q^\pi(s, a) = r + \gamma Q^\pi(s', \pi(s'))$$

where

r : current period reward that can be observed; for example, it could be defined as changes in the surplus amount or funding ratio.

γ : discount factor to reflect the timing difference, and

$Q^\pi(s', \pi(s'))$: reward function in the next period with new states s' and new asset rebalance action $\pi(s')$.

The reward function $Q(s, a)$ is a nonlinear function that captures the relationship between actions and rewards in RL. It represents the sum of the immediate reward r and the discounted future rewards, assuming that optimal actions will be taken in the future. $Q(s, a)$ follows the structure of the Bellman equation from dynamic programming, which describes the value of a state-action pair. However, unlike dynamic programming, where the exact function is predefined, RL learns $Q(s, a)$ by minimizing the error δ between its current estimate and the updated value:

$$\delta = Q(s, a) - (r + \gamma \max_{a'} Q(s', a'))$$

In RL, only the immediate reward r is directly observable, while future rewards are estimated based on the current Q . This differs from traditional model training, which requires the true value of Q which is unknown. To address this challenge, RL updates $Q(s, a)$ by aligning the observed current reward r with the estimated change in Q , calculated as:

$$Q(s, a) - \gamma \max_{a'} Q(s', a')$$

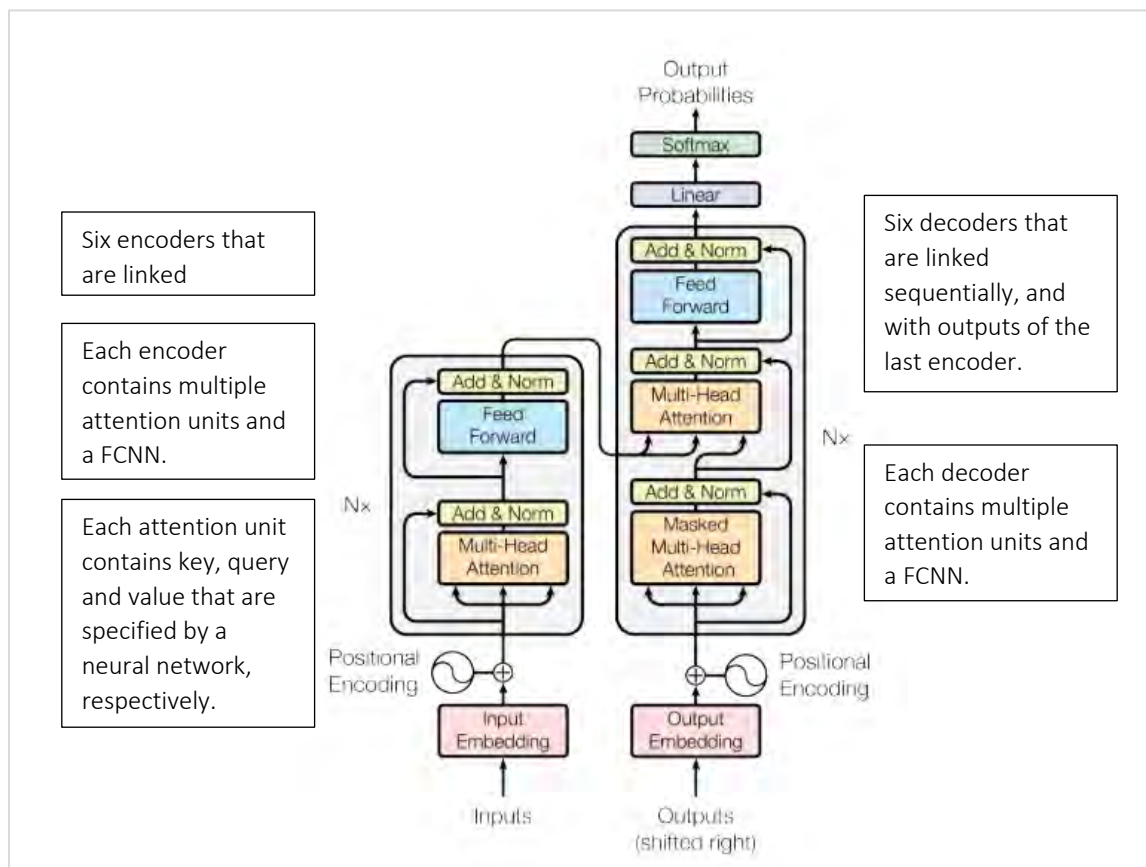
By iteratively refining Q through this process, the function learns to approximate the reward dynamics. Using deep learning models like feedforward neural networks or long short-term memory (LSTM) networks, $Q(s, a)$ can be estimated without needing a predefined functional form. Over time, with sufficient training, the calibrated Q function becomes accurate enough to guide good, if not optimal, actions.

Transformer Model

With the fundamental deep learning models and reinforcement learning models, AI models can be built with stacks of models and flexible structures to achieve different goals. The transformer model is a good example which is used in LLMs. As discussed in Vaswani et al. (2023), multiple neural networks, each with multiple layers, are used to construct a transformer model. Figure 38 shows the architecture of the transformer model, with multiple encoders and decoders, each of which is composed of multiple neural networks.

Encoders process the input data, such as a sequence of words, and generate a contextualized representation of the input, capturing relationships between different elements in the sequence. In some cases, multiple encoder layers are stacked, with each layer receiving the output of the previous layer, allowing a deeper understanding of the input sequence. On the other hand, decoders generate output sequences, such as answers to queries, by attending to both the contextualized representations from the encoders and the previously generated output tokens. Decoders can be stacked together as well to improve the output generated.

Figure 38
THE TRANSFORMER – MODEL ARCHITECTURE



Credit: The graph in the middle is a copy of Figure 1 in Vasawani et al. (2023). As indicated in the paper "...provided proper attribution is provided, Google hereby grants permission to reproduce the tables and figures in this paper solely for use in journalistic or scholarly works." The annotations in text boxes are added by the author of this report.

With the building blocks, innovative structures such as the attention unit can be developed and improve model performance. An attention unit dynamically focuses on the most relevant parts of input data, enhancing a model's ability to handle complex relationships. In Transformers, the scaled dot-product attention mechanism computes the relevance between a query that represents the question we have, and keys that represent the possible matches, producing attention weights. These weights are used to aggregate values (V) into a weighted sum to be returned, emphasizing important inputs. Transformers use self-attention to capture dependencies across sequence elements and multi-head attention to analyze different representation subspaces simultaneously. This mechanism enables Transformers to model long-range dependencies efficiently, making them highly effective for tasks like NLP and vision applications. For LLMs, the attention unit is the key to capturing dependencies and relationships in input sequences, such as what a pronoun refers to.

Generative Adversarial Networks

Another example of deep learning architecture is generative adversarial networks (GANs). A GAN contains two neural networks to generate more “authentic” data. One neural network generates new data by modifying the true data. The other one assesses whether the new data is authentic, until it can no longer distinguish between generated new data and the true data. It can be helpful for completing missing data or generating additional training data such as economic scenarios for other models.

Appendix C: AI Generated Codes for Portfolio Optimization User Interface

As part of the responses from GPT-4o in [Section 5.4 AI Retirement Planning Tool](#), the following JavaScript codes were generated.

```
"use client";

import React, { useState } from "react";

function FeedbackForm() {
  const [inputs, setInputs] = useState({
    initialWealth: 100000,
    horizon: 10,
    equityReturnMean: 0.06,
    equityReturnStd: 0.16,
    bondReturnMean: 0.01,
    bondReturnStd: 0.08,
    crra: 4,
  });

  const [recommendation, setRecommendation] = useState(null);

  const handleChange = (e) => {
    const { name, value } = e.target;
    setInputs((prev) => ({ ...prev, [name]: value })); // Keep the raw input value as a string for seamless typing
  };

  const calculateRecommendation = () => {
    const {
      initialWealth,
      horizon,
      equityReturnMean,
      equityReturnStd,
      bondReturnMean,
      bondReturnStd,
      crra,
    } = inputs;

    // Convert string inputs to numbers for calculation
    const numInputs = {
      initialWealth: parseFloat(initialWealth),
      horizon: parseFloat(horizon),
      equityReturnMean: parseFloat(equityReturnMean),
      equityReturnStd: parseFloat(equityReturnStd),
      bondReturnMean: parseFloat(bondReturnMean),
      bondReturnStd: parseFloat(bondReturnStd),
      crra: parseFloat(crra),
    };

    // Check for valid numeric inputs
```

```

if (Object.values(numInputs).some((val) => isNaN(val))) {
  setRecommendation("Invalid inputs. Please ensure all fields are filled correctly.");
  return;
}

// Simplified calculation for optimal equity weight
const optimalEquityWeight =
  (numInputs.equityReturnMean - numInputs.bondReturnMean) /
  (numInputs.equityReturnStd ** 2 + numInputs.bondReturnStd ** 2) *
  Math.exp(-numInputs.crra / numInputs.horizon);

setRecommendation(Math.max(0, Math.min(1, optimalEquityWeight)).toFixed(2)); // Clamp to 0-1 range
};

return (
  <div className="feedback-form">
    <h1>Private Investor Portfolio Selection Tool</h1>
    <p>
      Provide your assumptions below. Default values are pre-filled, but feel free to customize based on your
      unique financial situation.
    </p>
    <form>
      {[
        {
          label: "Initial Wealth ($)",
          name: "initialWealth",
          placeholder: "E.g., 100000",
          explanation: "Your starting investment amount.",
        },
        {
          label: "Investment Horizon (years)",
          name: "horizon",
          placeholder: "E.g., 10",
          explanation: "The length of time you plan to invest.",
        },
        {
          label: "Mean Equity Return (%)",
          name: "equityReturnMean",
          placeholder: "E.g., 0.06",
          explanation: "The average annual return expected from equities.",
        },
        {
          label: "Equity Return Standard Deviation (%)",
          name: "equityReturnStd",
          placeholder: "E.g., 0.16",
          explanation: "The annual volatility of equity returns.",
        },
        {
          label: "Mean Bond Return (%)",

```

```

    name: "bondReturnMean",
    placeholder: "E.g., 0.01",
    explanation: "The average annual return expected from bonds.",
  },
  {
    label: "Bond Return Standard Deviation (%)",
    name: "bondReturnStd",
    placeholder: "E.g., 0.08",
    explanation: "The annual volatility of bond returns.",
  },
  {
    label: "Risk Aversion (CRRA)",
    name: "crra",
    placeholder: "E.g., 4",
    explanation: "Your level of aversion to risk. Higher values indicate greater risk aversion.",
  },
].map(({ label, name, placeholder, explanation }) => (
  <div key={name} className="form-group">
    <label>{label}</label>
    <input
      type="number"
      step="any"
      name={name}
      value={inputs[name]}
      onChange={handleChange}
      placeholder={placeholder}
      className="input-field"
    />
    <small>{explanation}</small>
  </div>
))
<button type="button" onClick={calculateRecommendation} className="calculate-btn">
  Calculate Optimal Portfolio
</button>
</form>
{recommendation && (
  <div className="result">
    <h2>Recommendation</h2>
    {isNaN(recommendation) ? (
      <p style={{ color: "red" }}>{recommendation}</p>
    ) : (
      <p>
        Based on your inputs, the optimal equity weight for your portfolio is:{" "}
        <strong>{(recommendation * 100).toFixed(0)}%</strong>.
      </p>
    )}
  </div>
)}
</style jsx>{`

```

```

.feedback-form {
  max-width: 600px;
  margin: 20px auto;
  padding: 20px;
  border: 1px solid #ccc;
  border-radius: 10px;
  background: #f9f9f9;
}
h1 {
  text-align: center;
  color: #333;
}
.form-group {
  margin-bottom: 15px;
}
.form-group label {
  display: block;
  font-weight: bold;
}
.form-group input {
  width: 100%;
  padding: 8px;
  margin-top: 5px;
  border: 1px solid #ccc;
  border-radius: 5px;
}
.form-group small {
  display: block;
  color: #666;
  margin-top: 5px;
}
.calculate-btn {
  display: block;
  width: 100%;
  padding: 10px;
  background: #0070f3;
  color: #fff;
  border: none;
  border-radius: 5px;
  font-size: 16px;
  cursor: pointer;
}
.calculate-btn:hover {
  background: #005bb5;
}
.result {
  margin-top: 20px;
  padding: 10px;
  background: #e7f3ff;
}

```

```

    border: 1px solid #0070f3;
    border-radius: 5px;
  }
  .result h2 {
    margin-top: 0;
    color: #0070f3;
  }
  .result strong {
    color: #0070f3;
  }
  `}</style>
</div>
);
}

```

```
export default FeedbackForm;
```

Running the codes can generate a user interface shown in Figure 23. The codes can be run with an online compiler such as <https://onecompiler.com/react>.

Appendix D: SOA Publications on AI for Actuaries

General AI and Actuarial Science

1. [Generative AI for Actuaries](#): it discusses how actuaries can use generative AI, including LLMs, while addressing risks and implementation strategies.
2. [Risks Emerging from Artificial Intelligence Widespread Use: A Collection of Essays](#): they explore potential catastrophic risks linked to AI, covering privacy erosion, economic disruption, and bias.
3. [A Tour of AI Technologies in Time Series Prediction](#): it provides an overview of AI methods for time series forecasting, relevant to actuarial applications.
4. [SOA Emerging Topics Digital Publishing Platform – Artificial Intelligence & Machine Learning](#): it has a collection of AI machine learning podcasts and articles curated by the SOA.

AI and Mortality/Retirement

4. [The Impact of Artificial Intelligence on Mortality Modeling and Forecasting: Essay Collection](#): they discuss AI applications in mortality prediction and forecasting.
5. [The Impact of Artificial Intelligence/Large Language Models on Retirement Professionals and Retirees: Essay Collection](#): they examine AI and LLMs in retirement planning, with perspectives from experts.

AI and Investments

6. [The Impact of Artificial Intelligence/Large Language Models on Investment Actuaries and Other Investment Professionals](#) – Explores the role of AI in investment strategies and decision-making.

AI Implementation

7. [Operationalizing LLMs: A Guide for Actuaries](#): a step-by-step guide on implementing LLMs in actuarial work, covering evaluation, deployment, and governance.

AI Ethics and Bias Considerations

8. [Ethical Use of Artificial Intelligence for Actuaries](#): it provides tools to identify and manage ethical risks in AI-driven actuarial work.
9. [Ethical Issues in Automated Decision-Making Models](#): it discusses eight key ethical challenges organizations face when implementing AI, including bias and transparency.
10. [Comparison of Regulatory Frameworks for Non-Discriminatory AI](#): it compares various regulatory approaches to ensure AI fairness, transparency, and accountability.

Endnotes

¹ The AI system resulted in absolute reductions of 1.2% and 2.7% in the rates of false-positive and false-negative detections of biopsy-confirmed breast cancers, respectively.

² Source: <https://www.ipmorgan.com/technology/technology-blog/synthetic-data-for-real-insights>

³ An illustration of continuous bag-of-words (CBOW) is given as follows.

- Text data input

“Investment choice depends on risk tolerance, investment goals, and financial situation.”

In total, we have ten distinct words in our vocabulary ($V=10$).

- Training records assuming a window containing three words.
 - *Investment choice*
 - *Investment choice depends*
 - *choice depends on*
 - *depends on risk*
 - *on risk tolerance*
 - *risk tolerance, investment*
 - *tolerance, investment goals*
 - *investment goals, and*
 - *goals, and financial situation.*
 - *and financial situation.*
 - *financial situation.*

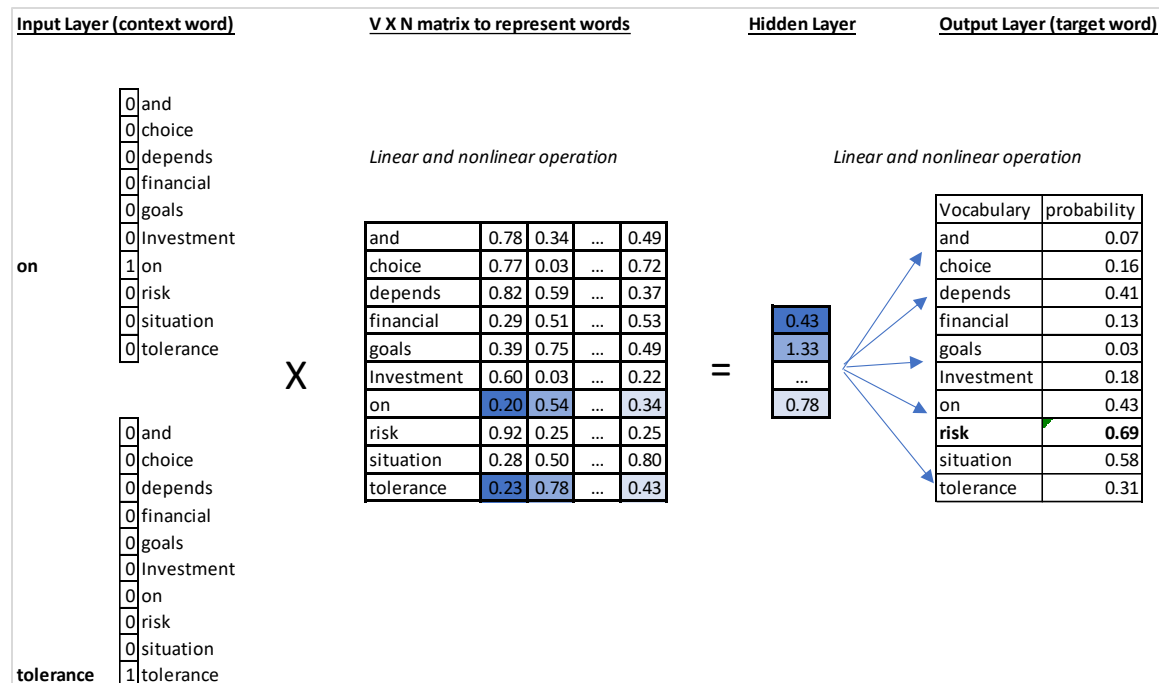
For example, using “on risk tolerance” as the training record, the goal is to predict target word “risk” given context words “on” and “tolerance.”

The model structure, illustrated in Figure #, begins with an input layer representing two context words, “on” and “tolerance,” as one-hot encoded vectors of size V . A $V \times N$ matrix, containing word vectors as parameters, is optimized during training. N represents the dimension of the word embeddings. Word embedding is a vector space where each word is represented as a dense vector. It is a model hyperparameter ranging from tens to hundreds and fine-tuned depending on the application and computational resources.

Multiplying the input vectors by this matrix produces the hidden layer, which predicts the target word. Probabilities are calculated for all words in the vocabulary and, in this example, “risk” is correctly identified as the most likely target word. The training process ensures the weight matrix becomes a meaningful representation of words for use in AI models. The hidden layer serves as an intermediate representation in the model. It captures relationships between the input context words and their meanings in a lower-dimensional space, enabling the model to predict the target word more effectively. This layer is critical because it learns the underlying structure of the word relationships, which is essential for generating meaningful word embeddings.

Figure 39

WORD2VEC MODEL STRUCTURE ILLUSTRATION



Once text data can be represented by numerical values as illustrated in the V x N matrix, they are ready to be used in AI models. In addition to CBOG, skip-gram is another popular approach used in Word2Vec. While CBOG uses context words to predict the target word, skip-gram uses a single word to predict the surrounding context words.

⁴ The RGB system (Red, Green, Blue) is a widely used color model in digital imaging, graphics, and displays, designed to represent and generate colors. It operates on the principle of additive color mixing, combining varying intensities of red, green, and blue light to create a wide range of colors.

⁵ An example of AI risk management framework can be found in the report “Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile” by NIST (2024).

⁶ $Adjusted R^2 = 1 - (1 - R^2) \frac{n-1}{n-p-1}$

Where

n : number of data records.

p : number of model parameters.

⁷ LlamaIndex is a framework for building context-augmented generative AI applications with LLMs.

<https://docs.llamaindex.ai/en/stable/>

⁸ Panel is an open-source library that lets you create interactive data tools, dashboards, and complex applications with Python. <https://panel.holoviz.org/>

⁹ TruLens is a software tool that helps measure the quality and effectiveness of LLM-based applications using feedback functions. <https://www.trulens.org/>

¹⁰ Each Reddit comment is processed so that words with high frequency but no meaningful information may be removed or tagged as unimportant. Words with similar meanings can be aggregated when counting the frequency.

- All words are converted to lower case so that words with the same characters but in different cases will be counted as the same word.
- All punctuation is removed.

-
- Stop words are removed before the analysis. Stop words are the most frequently used words such as “is,” “at,” “who,” and “that.” They are not useful for extracting key information from the comments.
 - Stemming is the process of removing morphological affixes from words. Words that end with “ed”, “ly”, “es” or “ing” could be trimmed and count as the root of these words. For example, a stemming algorithm may reduce words “snows,” “snowing,” and “snowed” to “snow.” This helps to reduce the number of distinct words to be counted across the Reddit comment data.
 - Synonyms are aggregated and counted as the same word. Grouping words with the same meaning as one helps strength their value in content classification and leaves more room for other word candidates. However, a word can have different meanings and belong to different grammatical categories in different sentences. For example, the word “good” could means “nice” as an adjective or “commodity” as a noun. Therefore, syntactic analysis needs to be performed first before determining a word’s synonyms in the context of a sentence. Apache OpenNLP library is used for part-of-speech (POS) tagging. The library uses the maximum entropy model to analyze the structure of a sentence and provide a category tag to each word of the sentence. Penn English Treebank POS taggers are used as the tagging system. After getting the POS tag for each word in a comment, tags are mapped to word types to facilitate synonym aggregation.
 - With each word in a comment tagged as a word type, WordNet by Princeton University is then used to find synonyms among these words in the dataset. WordNet is a large lexical database of English words (adjectives, adverbs, nouns and verbs) organized as groups of synonyms. Synonyms are then aggregated. Only one word will be used to represent a group of synonyms after the aggregation.
 - Another usual adjustment in text data analysis is spell check and correction. A majority of the comments are written by using mobile phones and tablets that have spell checking. Therefore, spell checking is not conducted for this research.

¹¹ The full conversion can be accessed at <https://chatgpt.com/share/a19d9b2b-6e05-42c7-b638-3653f78b6c43>.

References

- Abrate, Carlo et al., "Continuous-action reinforcement learning for portfolio allocation of a life insurance company,." 2021. Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part IV 21. Springer International Publishing.
- Balasubramania, Ramnath, Ari Libarikian, and Doug McElhaney, "Insurance 2030—The impact of AI on the future of insurance," March 2021. <https://www.mckinsey.com/industries/financial-services/our-insights/insurance-2030-the-impact-of-ai-on-the-future-of-insurance>
- Balona, Caesar, "ActuaryGPT: Applications of Large Language Models to Insurance and Actuarial Work," August 2023. <https://ssrn.com/abstract=4543652>
- Beck, Rae Hartley, "Best Robo-Advisors of January 2025," November 2024. <https://www.forbes.com/advisor/investing/best-robo-advisors>
- Benefits and Pensions Monitor, "How AI is transforming pension management: The Netherlands is leading the way," October 2023. <https://www.benefitsandpensionsmonitor.com/news/industry-news/how-ai-is-transforming-pension-management/380655>
- Bousquette, Isabelle, "Goldman Sachs Deploys Its First Generative AI Tool Across the Firm," 2024. *The Wall Street Journal*. <https://www.wsj.com/articles/goldman-sachs-deploys-its-first-generative-ai-tool-across-the-firm-cd94369b>
- Buehler, Hans et al., "Deep hedging," 2019. *Quantitative Finance* 19.8: 1271-1291.
- Cantor, David R., Kailan Shang, "Predictive Analytics: A Primer for Pension Actuaries," 2021. <https://www.soa.org/496395/globalassets/assets/files/resources/research-report/2021/2021-predictive-analytics-for-retirement.pdf>
- Cao, Jay et al., "Deep hedging of derivatives using reinforcement learning," 2021. <https://arxiv.org/abs/2103.16409>
- Cao, Larry, "Handbook of Artificial Intelligence and Big Data Applications in Investments." <https://rpc.cfainstitute.org/-/media/documents/article/rf-brief/ai-and-big-data-in-investments.pdf>
- Capponi, Agostino, Sveinn Ólafsson, and Thaleia Zariphopoulou, "Personalized Robo-Advising: Enhancing Investment Through Client Interaction," 2021. *Management Science* 68(4):2485-2512.
- Carbonneau, Alexandre, "Deep hedging of long-term financial derivatives," 2021. *Insurance: Mathematics and Economics* 99: 327-340.
- Carlin, Stephen and Stephan Mathys, "A Primer on Generative AI for Actuaries," February 2024. <https://www.soa.org/496313/globalassets/assets/files/resources/research-report/2024/primer-generative-ai.pdf>
- Chin, Andrew, Yuyu Fan, and Che Guan, "Unlocking Insights and Opportunities with NLP in Asset Management," 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*: 32-55. <https://rpc.cfainstitute.org/-/media/documents/article/rf-brief/ai-and-big-data-in-investments.pdf>
- Chong, Wing Fung, Haoen Cui, and Yuxuan Li, "Pseudo-model-free hedging for variable annuities via deep reinforcement learning," 2023. *Annals of Actuarial Science* 17.3: 503-546.

- Crosman, Penny, "Beyond Robo-Advisers: How AI Could Rewire Wealth Management," 2017.
<https://www.americanbanker.com/news/beyond-robo-advisers-how-ai-could-rewire-wealth-management>
- Das, Sanjiv R. and Subir Varma, "Dynamic Goals Based Wealth Management using Reinforcement Learning," 2019. https://srdas.github.io/Papers/GBWM_RL2.pdf
- Dastin, Jeffrey, "Insight - Amazon scraps secret AI recruiting tool that showed bias against women," October 2018. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G/>
- DeMiguel, Victor, Javier Gil-Bazo, Francisco J. Nogales, and Andre A. P. Santos, "Can Machine Learning Help to Select Portfolios of Mutual Funds," 2021.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3768753
- Du, Jiayi et al., "Deep reinforcement learning for option replication and hedging," 2020. *The Journal of Financial Data Science* 2.4: 44-57.
- European Parliament, "Harmonised Rules on Artificial Intelligence," April 2024.
https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138-FNL-COR01_EN.pdf
- Hayman, Genevieve, "Pensions in the Age of Artificial Intelligence," December 2024.
https://rpc.cfainstitute.org/sites/default/files/docs/research-reports/pensions-in-the-age-of-artificial-intelligence_online.pdf
- Irlam, Gordon, "Financial Planning via Deep Reinforcement AI," Journal of Retirement, 2020.
https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3201703_code2306779.pdf?abstractid=3201703&mirid=1
- Irlam, Gordon, "Multi Scenario Financial Planning via Deep Reinforcement Learning AI," 2020.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3516480
- Jansen, Stefan, "Advances in Natural Language Understanding for Investment Management," 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*: 56-69.
<https://rpc.cfainstitute.org/-/media/documents/article/rf-brief/ai-and-big-data-in-investments.pdf>
- Killock, D., "AI outperforms radiologists in mammographic screening," 2020. *Nat Rev Clin Oncol* 17, 134.
<https://doi.org/10.1038/s41571-020-0329-7>
- Lee, Justina, "Vanguard Quietly Embraces AI in \$13 Billion of Quant Stock Funds," February 2024.
<https://www.bloomberg.com/news/articles/2024-02-06/vanguard-quietly-embraces-ai-in-13-billion-of-quant-stock-funds>
- Li, Yuying and Peter Forsyth, "A Data Driven Neural Network Approach to Optimal Asset Allocation for Target Based Defined Contribution Pension Plans," 2018. *Insurance: Mathematics and Economics*.
https://cs.uwaterloo.ca/~paforsyt/Data_NN.pdf
- Lundberg, Scott M. and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," 2017. *Advances in Neural Information Processing Systems*.
https://papers.nips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
- Lütkebohmert, Eva, Thorsten Schmidt, and Julian Sester, "Robust deep hedging," 2022. *Quantitative Finance* 22.8: 1465-1480.

Marino, Mario, Susanna Levantesi, and Andrea Nigri, “Deepening Lee-Carter for longevity projections with uncertainty estimation,” 2021. <https://arxiv.org/abs/2103.10535>

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald, and Jack Clark, “The AI Index 2024 Annual Report,” AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2024. https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf

McLeod, Andrew and Harrison Jones, “ActuaryGPT: How large language models are shaping our future,” June 2023. <https://www.theactuarymagazine.org/actuarygpt/>

Mohammed, Shareefuddin, Rusty Bealer, and Jason Cohen, “Embracing advanced AI/ML to help investors achieve success: Vanguard Reinforcement Learning for Financial Goal Planning,” October 2018. <https://arxiv.org/abs/2110.12003>

Moore, Barry D., “7 Top AI Stock Software & Trading Bots Fully Tested,” December 2024. <https://www.liberatedstocktrader.com/ai-stock-trading>

Nagy, Peer, James Powrie, and Stefan Zohren, “Machine Learning for Microstructure Data-Driven Execution Algorithms,” 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*: 100-107. <https://rpc.cfainstitute.org/-/media/documents/article/rf-brief/ai-and-big-data-in-investments.pdf>

National Institute of Standards and Technology (NIST), “Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile,” July 2024. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

Ni, Jingwei, Julia Binger, Chiara Colesanti-Senni, Mathias Kraus, Glen Gostlow, Tobias Schimanski, Dominik Stambach, Saeid Ashraf Vaghefi, Qian Wang, Nicolas Webersinke, Tobias Wekhof, Tingyu Yu, and Markus Leippold, “CHATREPORT: Democratizing Sustainability Disclosure Analysis through LLM-based Tools,” October 2023. <https://arxiv.org/abs/2307.15770>

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, “Dissecting racial bias in an algorithm used to manage the health of populations.” *Science*, Vol 366, Issue 6464: 447-453, October 2019. <https://www.science.org/doi/10.1126/science.aax2342>

O’Leary, Lizzie, “How IBM’s Watson Went from the Future of Health Care to Sold Off for Parts,” January 2022. <https://slate.com/technology/2022/01/ibm-watson-health-failure-artificial-intelligence.html>

Olshan, Jeremy, “Do I Get Rich, or Go Bust? These Tools Predicted My Financial Future,” December 2024. <https://www.wsj.com/finance/investing/financial-planning-fortune-tellers-ae66ee1>

Ozhamaratli, Fatih and Paolo Barucca, “Deep Reinforcement Learning for Optimal Investment and Saving Strategy Selection in Heterogeneous Profiles: Intelligent Agents working towards retirement,” June 2022. <https://arxiv.org/abs/2206.05835>

Patel, Dylan and Gerald Wong, “GPT-4 Architecture, Infrastructure, Training Dataset, Costs, Vision, MoE. Demystifying GPT-4: The engineering tradeoffs that led OpenAI to their architecture,” July 2023. <https://www.semianalysis.com/p/gpt-4-architecture-infrastructure>

Perla, Francesca, Ronald Richman, Salvatore Scognamiglio, and Mario V. Wüthrich, “Time-series forecasting of mortality rates using deep learning.” *Scandinavian Actuarial Journal*. 2021. <https://www.tandfonline.com/doi/abs/10.1080/03461238.2020.1867232>

Pisaneschi, Brian, “Unstructured Data and AI: Fine-Tuning LLMs to Enhance the Investment Process,” 2024.
<https://rpc.cfainstitute.org/-/media/documents/article/industry-research/unstructured-data-and-ai.pdf>

Raghavan, Prabhakar, “Gemini image generation got it wrong. We'll do better,” February 2024.
<https://blog.google/products/gemini/gemini-image-generation-issue/>

Ribeiro, Marco T., Sameer Singh, and Carlos Guestrin. “Why Should I Trust You?: Explaining the Predictions of Any Classifier,” 2016. <https://arxiv.org/abs/1602.04938>

Rudolph, Max J., “Demographics and Productivity: Drivers of Economic Growth,” November 2023.
<https://www.soa.org/4ac0fd/globalassets/assets/files/resources/research-report/2023/demographics-productivity-drivers-econgrowth.pdf>

Sansani, Tal, and Mikhail Samonov, “Extracting Text-Based ESG Insights: A Hands-On Guide,” 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*: 70-91.
<https://rpc.cfainstitute.org/-/media/documents/article/rf-brief/ai-and-big-data-in-investments.pdf>

Sasaki, Takahiro, Hiroo Koizumi, Takao Tajiri, and Hiroaki Kitano, “A Study on the Use of Artificial Intelligence within Government Pension Investment Fund’s Investment Management Practices (Summary Report),” 2018. Tokyo, Japan: Government Pension Investment Fund.
https://www.gpif.go.jp/en/investment/research_2017_1_en.pdf

Schraub, David, Jing Lang, Zhibin Zhang, and Mark A. Sayre, “Comparison of Regulatory Framework for Non-Discriminatory AI Usage in Insurance,” August 2024.
<https://www.soa.org/4a1a8a/globalassets/assets/files/resources/research-report/2024/regulatory-framework-comparison-ai.pdf>

Senate Bill 24-205, “Consumer Protections for Artificial Intelligence Concerning Consumer Protections in Interactions with Artificial Intelligence Systems,” May 2024.
https://leg.colorado.gov/sites/default/files/documents/2024A/bills/2024a_205_enr.pdf

Shang, Kailan, “Deep Learning for Liability-Driven Investment,” 2021.
<https://www.soa.org/4964e1/globalassets/assets/files/resources/research-report/2021/liability-driven-investment.pdf>

Shang, Kailan, “Market Volatility Risk in an Era of Extreme Events,” 2023.
<https://www.soa.org/49e238/globalassets/assets/files/resources/research-report/2023/market-volatility-extreme-events.pdf>

Silva, Franklin and Pedra Braz, “Best AI Investing Apps and Software in 2025,” January 2025.
<https://investingintheweb.com/brokers/best-ai-investing-apps>

Socher, Richard, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts, “Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank,” 2013. https://nlp.stanford.edu/~socherr/EMNLP2013_RNTN.pdf

Son, Hugh, “JPMorgan is developing a ChatGPT-like A.I. service that gives investment advice,” May 2023.
<https://www.cnbc.com/2023/05/25/jpmorgan-develops-ai-investment-advisor.html>

Son, Hugh, “Morgan Stanley kicks off generative AI era on Wall Street with assistant for financial advisors,” September 2023. <https://www.cnbc.com/2023/09/18/morgan-stanley-chatgpt-financial-advisors.html>

Stanton, Erin, “Machine Learning and Big Data Trade Execution Support,” 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*: 92-99. <https://rpc.cfainstitute.org/-/media/documents/article/ai-and-big-data-in-investments.pdf>

UNESCO, “Recommendation on the Ethics of Artificial Intelligence,” 2021. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin, “Attention Is All You Need,” 2023. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

Warren, Geoff J., “Choosing and Using Utility Functions in Forming Portfolios,” 2019. *Financial Analysts Journal*, 75(3), 39–69. <https://doi.org/10.1080/0015198X.2019.1618109>

Willison, Simon, “Prompt injection attacks against GPT-3,” September 2022. <https://simonwillison.net/2022/Sep/12/prompt-injection/>

Yuan, Hang, Saizhuo Wang, and Jian Guo, “Alpha-GPT 2.0: Human-in-the-Loop AI for Quantitative Investment,” February 2024. <https://arxiv.org/pdf/2402.09746>

Zhang, Han et al., “COPR: Continual Human Preference Learning via Optimal Policy Regularization,” December 2024. <https://arxiv.org/abs/2402.14228>

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its [strategic research programs](#): aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of [topical research available](#), including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
8770 W Bryn Mawr Ave, Suite 1000
Chicago, IL 60631
www.SOA.org