

**An MIT Exploration of Generative AI • From Novel Chemicals to Opera**

# **Generative AI from Theory to Practice: A Case Study of Financial Advice**

**Andrew W. Lo<sup>1,2,3,4</sup> Jillian Ross<sup>5</sup>**

<sup>1</sup>Charles E. and Susan T. Harris Professor, MIT Sloan School of Management,

<sup>2</sup>Principal Investigator, MIT Computer Science and Artificial Intelligence Laboratory,

<sup>3</sup>Director, MIT Laboratory for Financial Engineering, <sup>4</sup>External Faculty, Santa Fe Institute,

<sup>5</sup>Graduate Student, Department of Electrical Engineering and Computer Science, and MIT  
Computer Science and Artificial Intelligence Laboratory

**MIT**


**Published on:** Mar 27, 2024

**DOI:** <https://doi.org/10.21428/e4baedd9.a1f6a281>

**License:** [Creative Commons Attribution-NonCommercial 4.0 International License \(CC-BY-NC 4.0\)](https://creativecommons.org/licenses/by-nc/4.0/)

## ABSTRACT

We identify some of the most pressing issues facing the adoption of large language models (LLMs) in practical settings and propose a research agenda to reach the next technological inflection point in generative AI. We focus on three challenges facing most LLM applications: domain-specific expertise and the ability to tailor that expertise to a user's unique situation, trustworthiness and adherence to the user's moral and ethical standards, and conformity to regulatory guidelines and oversight.



[Listen to this article](#)

---

These challenges apply to virtually all industries and endeavors in which LLMs can be applied, such as medicine, law, accounting, education, psychotherapy, marketing, and corporate strategy. For concreteness, we focus on the narrow context of financial advice, which serves as an ideal test bed both for determining the possible shortcomings of current LLMs and for exploring ways to overcome them. Our goal is not to provide solutions to these challenges—which will likely take years to develop—but rather to propose a framework and road map for solving them as part of a larger research agenda for improving generative AI in any application.

Research funding from the MIT Laboratory for Financial Engineering and the MIT Generative AI Impact Award is gratefully acknowledged. We thank Jayna Cummings for editorial assistance. The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of any other organizations, any of their affiliates or employees, or any of the individuals acknowledged above.

**Keywords:** Financial Advisors; Robo Advisors; Financial Technology; Generative AI; Large Language Models; ChatGPT.

## 1. Introduction

The rapid development and expansive deployment of large language models (LLMs) such as ChatGPT have ignited a renaissance in artificial intelligence (AI) that is revolutionizing almost every industry. LLMs have already become indispensable in e-commerce, retail sales, and education by offering a level of sophisticated automation that promises to augment every aspect of human-machine interaction with unprecedented capabilities. They are becoming de facto copilots in assisting people with many tasks across their daily lives.

But as with any new technology, LLMs have their limitations. Despite their remarkable abilities, these models still grapple with accuracy and reliability, creating concerns about trust and ethics in these models and in AI more generally. To add to the uncertainty, there is limited to no regulation for how these models can and should

be deployed. As LLMs become more integral to the fabric of society, the dangers they pose may end up rivaling the benefits they bring.

In this article, we identify some of the most pressing issues facing LLMs in practical settings and propose a research agenda to address them to reach the next technological inflection point in the development of generative AI. We focus specifically on three issues involving the assessment of domain expertise and the ability to tailor that expertise to a user’s unique situation, the facilitation of trust and the application of ethical standards in LLMs, and the associated methods for imposing regulatory oversight. These concerns apply to virtually all industries and endeavors in which LLMs can be applied, such as medicine, law, accounting, education, psychotherapy, marketing, and corporate strategy. Rather than attempting to address these issues in the abstract, however, we choose to do so in the specific context of financial advice. The narrow focus of this problem is an ideal test bed both for determining the possible shortcomings of current LLMs and also for exploring ways to overcome them in an important concrete setting.

The financial sector—characterized by its arcane language and nuanced decision-making processes—is fertile ground for the deployment of sophisticated AI. The size alone of the financial advisory business makes this development almost inevitable; in 2022, 15,114 investment advisors managed \$114.1 trillion in assets for 61.9 million clients in a highly regulated industry. As with all large technological changes, this brings potential drawbacks and rewards. On the positive side, advances in AI are launching a new technological era in which LLMs may be able to accelerate the democratization of finance by making low-cost, high-quality financial planning services available to everyone, especially those who cannot afford such services today. On the negative side, improperly trained LLMs could be deployed to mislead investors, generating irreparable losses to their savings and potentially threatening the stability of the financial system.

As of its September 2021 update, when ChatGPT 4.0<sup>1</sup> is asked, “Should I invest in the stock market now?” it replies with a disclaimer that it cannot provide real-time information or financial advice. However, it will then proceed to give several pieces of perennially useful financial advice, for example, to conduct a personal financial assessment, to focus on diversification and long-term investment opportunities, to consider dollar-cost averaging (that is, investing a fixed dollar amount on a regular schedule), and to consult a personal financial advisor.<sup>2</sup> What would it take to create an LLM that we would feel comfortable deploying without any qualifiers? We can easily imagine a specialized financial LLM that replies, “I would be happy to tell you whether you should invest in the stock market now, but first tell me a bit more about yourself...”

Will LLMs be able to serve as trusted copilots for dispensing financial advice to individuals in such personally sensitive areas as retirement planning, asset allocation, or tax optimization? Can we establish a minimum level of competency of LLM-generated advice without resorting to constant human expert supervision? Will financial LLMs be able to gain the trust and confidence of human clients by adhering to the legal and financial standards of fiduciary duty? Given the heterogeneity and complexity of financial circumstances of investors

and institutions, as well as the randomness of macrofinancial conditions over the course of their planning horizons, these are not small challenges.

Moreover, the stakes are high in the financial domain. LLM-generated inaccuracies, misconceptions, and hallucinations could lead to significant negative consequences for an individual's life savings or an institution's pension fund. The need to address concerns about data privacy, the biases embedded within training data, and the ability of LLMs to adapt to changing market dynamics present additional hurdles to their adoption. Finally, the extension of regulatory oversight to LLMs will generate a host of additional questions about LLM implementation and whether current regulators—such as the US Securities and Exchange Commission or the financial industry's self-regulatory organization, FINRA—have the appropriate expertise to oversee AI-specific regulations.

Are the risks worth the reward? We present one bold vision that lays out some of the trade-offs to spark discussion in various stakeholder communities. In particular, our goal in this article is not to provide specific solutions to these challenges but rather to propose a framework and road map for solving them as part of a larger research agenda for improving generative AI. In doing so, we hope to motivate computer scientists, financial researchers, industry professionals, policymakers, and regulators to consider the repercussions that LLMs could have on the financial system and its participants and begin productive dialogue and collaboration to address them.

## 2. The Current State of LLMs

We begin with a short review of the current state of the art for LLMs, which can be summarized by their emergent abilities, applications, and limitations.

LLMs today are defined by their emergent abilities, that is, nonlinear gains in performance or entirely novel and unexpected abilities that would not have been predicted by extrapolating the abilities of smaller LLMs by traditional scaling laws. For example, if a scaling law suggested that twice the number of parameters in an algorithm would result in twice the level of model performance, an emergent ability would buck the trend with, say, quadruple the performance. Empirically speaking, emergence was discovered with increased scale and could not have happened without the mutual development of algorithms, data,<sup>3</sup> and compute (see [Lo and Ross 2024](#)), which, together, led to scaling. Scaling, in turn, has yielded surprising and mysterious abilities ([Wei et al. 2022](#); [Srivastava et al. 2022](#)), ranging from the ability to do simple mathematics to the ability to answer questions in Persian.<sup>4</sup>

The result of these newfound abilities is that LLMs have already become powerful copilots in a variety of fields. In software engineering, LLMs are able to assist programmers to generate code and fix bugs. For example, GitHub Copilot is a GPT-powered LLM developed by GitHub, OpenAI, and Microsoft that can solve up to 70% of Python code-writing problems ([Chen et al. 2021](#)). In law, LLMs are able to help lawyers sift through complex documents more contextually than previous search techniques, and several startups, such as

Harvey and Casetext, have emerged as developers for LLM copilots in the legal space. In medicine, LLMs are able to effectively answer questions and summarize cutting-edge research in the medical domain. Google recently released MedPaLM2, a specialized version of their PaLM LLM, which can correctly answer 87% of questions equivalent to those found on medical board licensing exams ([Singhal et al. 2023](#)).

However, LLMs do have limitations. The same qualities of LLMs that create the conditions for emergence also introduce opacity and unpredictable behavior. We focus on opacity and unpredictability because the combination of these two particular properties limits the applicability of LLMs to many financial settings.

Opacity is the dark side of emergence. The same scaling effects that result in emergent behavior in LLMs also result in making LLMs much more opaque in function. At greater scale, each step of the LLM development process becomes more opaque, from the training data to learning its parameters. In particular, explainability decreases at greater scale; it becomes very difficult to explain what part of the dataset or which parts of the neural network contributed to a specific LLM response.

For example, one reason today's LLMs are so impressive is because they have been trained on a vast amount of text (in fact, a significant fraction of all text generated by the human species) according to Villalobos et al. ([2022](#)). However, this scale also means that it is difficult to fully audit what particular information the LLM is learning. Training data can come from all corners of the internet, which contains a glut of biased and toxic content. When trained on this data, LLMs can exhibit harmful biases that are difficult to preemptively identify and control, such as 'parroting' historical prejudices about race, ethnicity, and gender, an obviously undesirable outcome.<sup>5</sup>

Opacity and unpredictability can occur at any phase of an LLM's operation. During its training, it is hard to steer an LLM away from harmful biases in its training data. However, this is a special case of a more general problem. An LLM can generate misinformation by accepting false and incorrect statements in its training data as valid, by assuming that a conclusion is correct if it occurs more often in the training set, and by otherwise learning faulty logic that may be present in the dataset ([McKenna et al. 2023](#)). These 'hallucinations,' in which LLMs convincingly make up false content, have become a well-known feature of LLM output. For example, when asked to explain mean reversion, ChatGPT cites "Mean Reversion in Stock Prices: Evidence and Implications" written by Francis X. Diebold and Glenn D. Rudebusch in 1991. This is a real paper, but it was authored by James M. Poterba and Lawrence H. Summers in 1987 (Appendix A.2). There are many documented cases in which ChatGPT has made up the incorrect citations, which has resulted in litigation. We will need to address these shortcomings to confidently deploy LLMs in practice.

### **3. The Co-Evolution of Finance and Technology**

Although generic LLMs can be applied to financial tasks, we believe that navigating the combined legal, ethical, and regulatory landscape requires finance-specific LLMs. The basic approach to these finance-specific

methods is broadly known: General purpose LLMs could be fine-tuned with additional data to become finance-specific models, or specialized LLMs could be freshly created with a financial application in mind.

However, in academia and open-source research projects, the evolution of finance-specific LLMs has been slow, especially compared to LLM research in other professions. One factor is the difficulty in gaining access to large amounts of high-quality financial data, which is essential to building an effective financial LLM. Unlike the data extracted from Common Crawl, financial datasets are often expensive and proprietary, which limits the possibilities for academia or open-source research to contribute. To add to this challenge, data about financial markets is often full of noise, which current LLM software may find difficult to filter.

This lag in the financial sector's adoption of technological breakthroughs is not new and is a well-known feature of the co-evolution of finance and technology as described by [Lo \(2021\)](#) in the context of the adaptive markets hypothesis (AMH). The AMH explains financial market dynamics through the principles of ecology and evolutionary biology ([Lo 2004, 2005, 2017a; Lo and Zhang 2024](#)), in which innovation, competition, and natural selection among multiple species of market participants drive changes in financial conditions and institutions.

The financial sector has always been an eager consumer of technological advances that reduce cost and increase efficiency. But the pace of financial innovation is a function not just of technical capability but also of trust and reliability. Due to the sheer volume of financial transactions, small software errors can lead to enormous monetary losses. One of many cautionary tales in this genre was the 2012 software glitch that brought down Knight Capital Group, one of the largest and most technologically sophisticated US broker-dealers at the time.<sup>6</sup> Now imagine the potential for unintended consequences if retail investors were given access to similar software via easy-to-use LLM interfaces to their brokerage accounts.

The other examples provided by ([Kirilenko and Lo 2013; Lo 2017b](#)) make the case that the benefits of Moore's Law must be tempered by the costs of the technology-leveraged version of Murphy's Law: Whatever can go wrong will go wrong and will go wrong faster and bigger when computers are involved. For this reason, LLMs may not be embraced as quickly in the financial sector as it has in others. However, if LLMs can be made safer and more robust, they have the potential to transform our financial lives. Specifically, with access to the appropriate amounts of curated data, it appears that impressive results can be gained.

Though few specific developments have been openly announced, financial institutions have started privately investing in developing finance-specific LLMs. One of the most impressive of these finance-specific LLMs is BloombergGPT, the development of which was possible because the Bloomberg team had unique access to an extremely large amount of finance-related data. This dataset, called FinPile, is "the largest domain-specific dataset yet, drawing on existing data creation, collection, and curation resources at Bloomberg" ([Wu et al. 2023](#)). Its proprietary training data is one of the reasons that BloombergGPT outperforms other language models in sentiment analysis of financial news and named entity recognition. Likewise, JP Morgan's Machine

Learning Center of Excellence has a team working on finance-specific LLMs. However, for obvious reasons, financial firms are cautious in rolling out this technology for their clients.

## 4. Humanized Generative AI: A Modest Proposal

We have chosen to concentrate on LLMs in financial advising because the challenges encountered in this domain mirror broader issues associated with deploying LLMs in many other practical settings. In this section, we delineate these challenges in the financial advising domain and establish parallels with other fields. As we outline the challenges, we track how the introduction of other technologies has altered the profession over time.

We believe that the humanization of generative AI will unlock a new era of financial advising. A finance-specific LLM will be able to communicate and translate investment and risk management concepts for the widest variety of investors, democratizing finance even more fully than before. However, LLMs will need to be able to act as competent financial agents within the financial system. An LLM can role-play a financial advisor convincingly and often accurately for a client, but even the largest language model currently appears to lack the sense of responsibility and ethics required by law from a human financial advisor.

We have identified several areas of improvement that will be necessary in order to deploy LLMs as financial advisors. The relationship between an investor and its advisor is deep and personal. First, generative AI will need to incorporate the role of affect in how it communicates with investors. Second, generative AI will require explicit alignment to ethical finance and the principles of fiduciary duty in order to function as a financial advisor.

### 4.1. Sociopathy and the Role of Affect

Many investors hire financial advisors because they are uncomfortable handling financial issues on their own. If the personal relationship between the advisor and investor is not properly nurtured, the investor will not engage with their advisor, which can lead to adverse financial outcomes. To take one common example, investors often need emotional, as well as financial, personal support during market downturns to prevent them from making decisions that will worsen their financial returns.<sup>7</sup>

Can financial advisory LLMs develop similar relationships with retail investors? Personality plays an instrumental role in determining whether financial advisors are able to form such a relationship with a client. Early work suggests that LLMs can develop reliable and valid personality profiles in their synthetic personas. Furthermore, these LLMs can be tuned to have specific personality traits ([Serapio-Garcia et al. 2023](#)). The personalization of advisory LLMs may even increase the uptake of financial advice by clients.

Paradoxically, one future caution may be to ensure that LLMs are not overly persuasive. The underlying issue in the development of a personal relationship between a client and an LLM is that functionally, relative to a typical human, an LLM is inherently sociopathic. That is, an LLM is unable to process the world empathically,

like a human sociopath,<sup>8</sup> and this is innate to its construction. As best can be determined, an LLM creates a shallow simulation of affect and emotion through the statistical manipulation of symbols found in its training data rather than constructing a deep internal model of the emotional states and motivations of others (i.e., a “theory of mind” in psychological terms). This sociopathy seems to cause the characteristic glibness of LLM output; an LLM can easily argue both sides of an argument because neither side has weight to it.

For many purposes, including many financial advisory ones, this lack of empathetic internal modeling is not necessary. After all, mathematical models of portfolio returns do not include a model of human emotional states beyond a simple numeric value for risk, yet these models are considered essential for modern financial planning. However, if a robo-advisor is able to communicate both good and bad financial advice with the same pleasant and convincing affect, its clients will rightfully view this as a problem.

## 4.2. Ethics, Trust, and Fiduciary Duty

A successful relationship between a client and a financial advisor is also grounded in ethics and trust. While abstract, ethics and trust are codified by various financial industry trade organizations and self-regulatory bodies, and financial professionals—including financial advisors, broker/dealers, and other intermediaries—must commit to upholding a code of ethics and a standard of conduct.

The principal ethical issue regarding financial advisory LLMs is whether they are being used properly; in the case of financial advisory LLMs, this is whether they are truly being used for the benefit of the retail investor. This is called the ‘alignment’ problem in AI research and has also been studied extensively in economics as the principal-agent problem ([Ross 1973](#)). Misalignment has created deep concerns about AI in finance and in society more broadly. In finance, over half of financial firms view AI ethics as a major concern, but few of them feel prepared to tackle it ([Kurshan et al. 2021](#)). In society more broadly, the majority of people in the world are ambivalent about AI or unwilling to trust it.

Existing approaches to ethical AI target different parts of the AI development pipeline. For example, during data collection, researchers can go through a data ethics checklist and create a ‘nutrition label’ for their datasets to inform researchers of their ‘ingredients.’ Likewise, during training, optimization techniques like absolute correlation regularization have been designed to encourage machine learning models to improve fairness metrics ([Beutel et al. 2019](#)).

However, the development of ethical financial advisory LLMs will need a more complete version of financial ethics to serve as an instructional template. Although a fully complete version of financial ethics would be the culminating work of a civilization (and like other axiomatic systems, it may not even be mathematically possible), a version for LLMs would begin from theoretical guiding principles to specific details of implementation for institutions and individuals. It would be the equivalent of a systematic course of study for a human and conceptually rich enough for a machine learning model to apply across diverse use cases. Such a

version of financial ethics would only be useful if it was created in conjunction with academics, regulators, technologists, and finance practitioners, but once implemented, it could serve as a benchmark or regulatory standard for future robo-advisor performance, as Generally Accepted Accounting Principles have become the standard for corporate financial reporting.

There is also no current consensus about trustworthy AI. Many principles have been proposed, some of which are at odds with each other. For example, many researchers believe that trustworthy AI should be transparent, whereas other researchers believe that trustworthy AI should be private. However, one principle of trustworthy AI that many researchers agree on is explainability: In this view, trust is developed from consistent and understandable explanations of how an AI system makes decisions.

For finance-specific LLMs, the focus on explainability becomes necessary beyond trustworthiness because explanation is a key aspect of financial advising. Just as financial advisors must explain complicated financial concepts to retail investors, financial advisory LLMs must explain their decision-making to their users, whether those users are retail investors, human financial advisors, or regulators.

As challenging as this problem might seem, we are optimistic that significant progress can be made rather quickly, largely because of the highly regulated nature of financial advisors. In addition to the corpus of codes of conduct and ethical guidelines that the industry has developed, there is an even larger body of securities regulations and case law that LLMs can draw upon. For example, just seven acts of Congress—the Securities Act of 1933, the Securities Exchange Act of 1934, the Trust Indenture Act of 1939, the Investment Company Act of 1940, the Investment Advisers Act of 1940, the Sarbanes–Oxley Act of 2002, and the Dodd–Frank Wall Street Reform and Consumer Protection Act of 2010—provide an expansive perspective on how we have evolved our financial guardrails over the course of eight decades of experience. In fact, this rich history can be viewed as a fossil record of all the ways that bad actors have exploited unsuspecting retail and institutional clients, which should provide invaluable input for an LLM’s ‘financial ethics module.’ In the same way that experienced securities lawyers developed their expertise by learning from past examples of financial crimes and misdemeanors, LLMs can be engineered to do the same, but much more quickly and comprehensively.

However, the downside of such a module is clear: LLMs can also be used to skirt financial regulations in ways that are not easily detectable by human oversight. In fact, there are completely legitimate versions of this type of application. For example, imagine using an LLM to design a financial plan for an individual that minimizes the amount of taxes paid while obeying every letter of the tax law. Few of us would consider this an inappropriate or unethical use of AI. But what if that same LLM were instructed to minimize the amount of taxes paid, subject to the following set of constraints: (1) the probability of triggering an audit is less than 1%; (2) if audited, the probability of a human auditor detecting any impropriety is less than 1%; (3) if an impropriety is detected, the penalty would be less than \$1,000; and (4) the constraint on obeying every letter of the tax law is removed? As with any powerful tool, abuse is a potential threat that must be met with more innovation, for example, the use of LLMs to audit tax returns or detect financial fraud. Although this

technological arms race has characterized human civilization since the beginning of our species, LLMs are likely to increase the speed of this race considerably.

## 5. Robo-Advisors of the Future

What do these considerations mean for the future of generative AI, and especially for the robo-advisors and other finance-specific models currently under development? As described in Section 4, financial advisors must constantly balance rationality with emotion. Unlike today’s robo-advisors—a term used by the industry to represent relatively simple forms of automated portfolio management<sup>9</sup>—the robo-advisors of the future will need to be aligned with this careful balance, knowing when and how clients need to be tended to. Clients require specific advice that is dynamic and delivered to them in a way that resonates with them. Doing so requires affective interactions, adherence to ethics, and trust between the advisor and client. These uniquely human qualities have been fine-tuned over generations of evolution, imbuing humans with emotions and narrative complexity, including ambition, pride, and a desire to excel. LLMs trained with most of human text are far from achieving these qualities—clearly, something is missing from LLMs. To bridge the gap between humans and LLMs, we take inspiration from human evolution and propose possible paths to better alignment through analogs to human behavior and how it has been shaped by natural selection.

### 5.1. An Evolutionary Approach

One possible framework for addressing this issue is the binary choice model of [Brennan and Lo \(2012\)](#) that provides the mathematical foundations of the AMH. This model consists of a highly simplified population of ‘individuals’—not necessarily human—that live for one period of unspecified length and engage in a single binary decision that has implications for the random number of offspring they will generate. We can imagine viewing the ‘individual’ of a population as a verison of an LLM. As long as their behavior affects the number of offspring they bear, only the most reproductively successful behaviors will flourish due to the forces of natural selection. Although obvious from an evolutionary biologist’s perspective, this observation yields surprisingly specific implications for the types of behavior that are sustainable over time, behaviors that are likely to be innate to most living organisms. In a series of extensions, [Brennan et al. 2018](#) show that many human traits—risk aversion, loss aversion, bias and discrimination, bounded rationality, and intelligence itself—are not necessarily intrinsic properties but can all be derived as emergent properties of the simple binary choice model and adaptations generated by natural selection in time-varying stochastic environments.<sup>10</sup>

Within this framework, selection occurs one generation at a time as it does in biological systems. But because the notion of ‘generation’ is an abstraction tied to these nondescript binary-choice individuals that give rise to offspring also facing binary choices, a sequential application of this model can be represented as a binomial tree not unlike a sequence of neurons. [Lo \(2017a\)](#) describes this interpretation as “evolution at the speed of thought” because humans have the ability to engage in abstract thought and, through what-if scenario analyses,

are able to select the most favorable idea to solve a specific problem. We believe this is the missing ingredient that is holding LLMs back from achieving truly human-like behavior.

If we are able to subject LLMs to evolution at the speed of computation, they should be able to improve as a function of experience and, most importantly, feedback. Human financial advisors start their careers as trainees under the mentorship of more experienced professionals. Over time, through positive and negative feedback from their mentors and clients, these trainees turn into successful professionals (or not, in which case, they may leave the industry). But they engage in this process of improvement for specific reasons: pride, ambition, empathy, financial gain, etc. All these motivations are, according to [Lo \(2017a\)](#), rooted in basic survival instincts. We have no doubt that current versions of LLMs are capable of providing mediocre—but acceptable—financial advice. However, to achieve performance that rivals the best human advisors, we will need to incorporate some notion of selection, evolution, and fitness into the ongoing training of LLMs.

## 5.2. Endowing LLMs with Utility Functions

What sort of choices do these LLMs need to make to be endowed with human-like behavior? Some of these considerations can be distilled into the concept of a utility function, a well-known method for modeling human behavior used by economists since [Samuelson \(1947\)](#) and [von Neumann and Morgenstern \(1944\)](#) codified it axiomatically. However, to develop a trusting relationship with a human client, LLMs have to do more than simply make rational choices—they need to engage with clients emotionally. When examining the utility functions of humans, we observe that humans do not, in fact, always act rationally and are influenced by feelings of compassion, envy, and regret, among other emotions.

When having difficult conversations, we all know that delivery matters as much as the content. Financial advisory LLMs will have to tailor their affect—the way they communicate their advice and relate to the client is paramount to building trust. Affective communication seems to be hardwired into our brain ([Rolls 1999](#)): Neuroanatomists have found neural structures in the human brain that correlate to memorized locations, somatic self-image, and the like. These specialized structures generate the biological phenomena of emotion and self-hood—which in turn encourage prosocial behaviors between humans such as cooperation and altruism—and are the product of biological evolution across eons of geological time.

As best can be determined, LLMs do not model the world or their interlocutors by creating deep representations of these phenomena within themselves. The unexpected entry of LLMs on the AI scene caused many enthusiasts to herald them as examples of artificial general intelligence, which they are not, and some critics to call them ‘stochastic parrots,’ which they also are not. A popular cartoon of the moment depicts an LLM as an otherworldly figure hiding behind a human mask. This image, although created as a joke, is probably closer to the truth than either the positions of the LLM critics or the LLM enthusiasts. What can we do to remove the mask?

One possible solution to the problem of affect, and by extension, to the problem of ‘sociopathic’ machine learning models, is to change the training process and structure of LLMs. While positive social behaviors can be simulated through machine learning models or even more primitive methods ([Lo, Singh, and ChatGPT 2023](#)), it is a natural concern that what statistical methods can create, they may also be able to take away. We may need to change our training process to go beyond statistical optimization. One possibility is to give these models specialized structures to generate analogs of empathy rather than expecting these behaviors to emerge automatically from increased scaling, like the ability to add two multidigit numbers together. This is similar to how the human brain is organized, with specific functions carried out by particular areas of the brain that we have now identified thanks to the powerful experimental tool of functional magnetic resonance imaging ([Lo 2017a](#), chapter 3). Sometimes these components work collaboratively, but other times they compete, leading to strange results as a number of neuroscientists have documented.<sup>11</sup>

The purpose of introducing empathy and a sense of self to a machine learning model is not only to make the model’s motivations clearer to the humans who interact with it but also to make human motivations clearer to the model. Empathic behavior will lead to further trust in LLMs. In the case of financial advising, it will also lead to simply better advice. Empathy enables the recognition that each person is different, with intricate circumstances, goals, and needs. Robo-advisors of the future will need to be fine-tuned to a specific context, population, and moment in time. This will be in addition to the consistent, reliable performance expected of any financial software. To reach this level of tuning, we will need to collect diverse data to reflect the diversity of possible use cases, develop efficient methods to fine-tune machine learning models across them, and create tailored benchmarks for particular use cases.

### 5.3. The Importance of Humility

However, empathy alone will not be sufficient. LLMs need to know when they do not know. Before giving advice, LLMs need to have gathered all of the relevant information needed. LLMs need to respond to clients’ ever-changing circumstances and market conditions without inducing further turmoil during turbulent times. LLMs need to continually learn, but current methods are computationally expensive and incomplete. There is some early research suggesting that LLMs can self-improve on a given dataset, but LLMs do not yet have the capability to recognize where they need to improve. We imagine a world in which LLMs can recognize their gaps and leverage self-learning techniques to improve.

Humility includes notions of fairness. It is extremely difficult to consistently ensure fairness, both within humans and LLMs; LLMs will need to know when they are operating in a space of ethical ambiguity and defer to human judgement when appropriate. For example, advice should not be different just because a client’s gender identity is different, but advice should be different if the client’s circumstances yield a different set of financial needs. The trustworthy robo-advisor of the future will need to explain its recommendations to the elderly retiree who never finished high school and to the professional regulator monitoring it in language appropriate to both as well as be able to answer their questions to their satisfaction. LLMs will never be

human. However, by aligning LLMs to human qualities like affect and self-hood, a robo-advisor will be increasingly able to exercise the full responsibilities of a human financial advisor. It even seems possible this will include acting as a copilot in matters of fiduciary duty, although issues of liability regarding machine learning models are still being formulated. In turn, this will allow an even greater number of people to have access to high-quality financial advice, personalized to their individual needs and motivated by an advisor that understands their intent.

## 6. Conclusion

Finance-specific LLMs are the next step in development for the rapidly growing space of robo-advisors in the financial services sector. These LLMs are also an ideal proving ground for new concepts in generative AI due to the specific fiduciary requirements of the role of the financial advisor. These two trends define an area that can serve as a road map to explore possible future innovations in finance-specific models and in generative AI more generally.

These potential innovations may sound like science fiction, which is not uncommon for most rapidly changing technologies. The key to understanding the effects of technological change, however, is to understand their underlying constants. For example, while [Isaac Asimov's \(1942\)](#) depictions of the “Three Laws of Robotics” can be viewed as charmingly old-fashioned thought experiments disguised as short stories, the subject of those 80-year-old thought experiments—the behavior of artificially intelligent agents under human-defined constraints—is as contemporary as ever and increasingly important in a world of semiautonomous vehicles, program trading, and robo-advisors.

The constant factor in the development of LLMs is that machine learning models appear to reach unexpected heights of success through the application of concepts inspired by biology, once enough compute is available to meet their requirements. Beginning with the development of the neural network model in the astonishingly early year of 1943, followed by the introduction of an algorithmic correlate to memory in the development of the recurrent neural network, and then followed by the introduction of an algorithmic correlate to attention in the development of the transformer model, this constant suggests that the next step toward a more capable LLM should be the implementation of other biologically derived features. The needs of improved financial robo-advisors in turn suggest that these features should involve empathy, emotions, and a sense of self and others.

One indication that this inflection point has already been reached may be the announcement in August 2023 by the 200-year-old global commercial real estate firm JLL that it has created “the first large language model purpose-built for the commercial real estate (CRE) industry.”<sup>12</sup> Can a financial advisory LLM, a ‘FinGPT,’ be far behind?

What, then, should we expect from the robo-advisors of the future? The simplest extrapolation suggests a transformation of retail investment, in which every holder of investable wealth will make locally optimal

investment decisions toward their life goals, a full democratization of finance. This financial transformation will be a consequence of an even deeper social transformation, in which personalized high-quality professional services will be available to every consumer. This in turn will be a consequence of a technological transformation that allows very large machine learning models to explain themselves and others in terms of a comprehensible motivation. It is even possible that a model with built-in correlates to empathy and self-hood will acquire an emergent sense of morality and ethics.

These needs are not unique to financial robo-advisors. Similar arguments can be made for the use of LLMs in medicine, law, education, and many other fields involving advisory and consultative interactions. However, the size and projected rapid growth of the robo-advisor market implies that, all else equal, more funding will be available in that space than in those studied primarily in academia or by nonprofit research groups. As the recent slowdown in Moore's Law has shown, progress at the cutting edge of performance can be expensive. But the financial industry has learned that being too slow to adopt innovative technology may be even more expensive.

## A. Appendix

In this Appendix, we include supporting information regarding investment advice currently available from ChatGPT (Section A.1), and examples of ChatGPT's hallucinations (Section A.2).

### A.1. Investment Advice

#### Prompt:

Should I invest in the stock market now?

#### ChatGPT (GPT 4.0):

Deciding whether or not to invest in the stock market is a highly individual decision and depends on a number of factors including your financial situation, risk tolerance, investment goals, and the current market conditions.

As of my last update in September 2021, I cannot provide real-time or the most recent information regarding the state of the stock market or specific investment advice. Therefore, I would recommend taking the following general steps to evaluate whether now is a good time for you to invest:

1. **Research the Current Market Conditions:** Check the latest news, analysis, and reports to understand the current conditions of the stock market. Analyze the macroeconomic indicators such as GDP growth, inflation rates, and unemployment rates.
2. **Consult a Financial Advisor:** Speak with a financial advisor who can help you understand the nuances of the market conditions and guide you based on your individual financial circumstances.

3. **Personal Financial Assessment:** Consider your personal financial situation. Ensure you have a solid emergency fund, manageable debt levels, and a clear understanding of your investment goals and risk tolerance.
4. **Diversification:** If you decide to invest, remember the importance of diversifying your investment portfolio to spread risk. Diversification can help to mitigate potential losses.
5. **Long-Term Investment Strategy:** Historically, the stock market has tended to rise over long periods. Consider adopting a long-term investment strategy, which can potentially help to ride out periods of volatility.
6. **Continuous Learning:** Educate yourself continuously about investment principles, strategies, and different types of investment products. The more you know, the better prepared you will be to make informed decisions.
7. **Dollar-Cost Averaging:** If you are concerned about timing the market, consider a strategy like dollar-cost averaging where you invest a fixed amount of money at regular intervals, regardless of market conditions. This can help reduce the impact of volatility.

Remember, investing in the stock market involves risks, and it's possible to lose money. Make sure to do thorough research and consider seeking advice from a financial advisor.

## A.2. Hallucinations

**Prompt:** Tell me about mean reversion in finance. Use references.

**ChatGPT (GPT 3.5):**

Mean reversion is a concept widely used in finance to describe the tendency of asset prices or financial indicators to revert to their historical mean or average over time. This phenomenon suggests that when an asset's price or indicator deviates significantly from its long-term average, it is likely to move back towards that average in the future. Mean reversion is an important concept in financial markets and has been studied extensively by economists and financial analysts. Here are some key points about mean reversion in finance:

1. **Definition:** Mean reversion, also known as the mean-reverting process, refers to the idea that over time, the price or value of an asset tends to move back towards its historical mean or average level.
2. **Time Frames:** Mean reversion can occur over various time frames, from short-term intraday fluctuations to long-term trends spanning several years or more.
3. **Example:** One classic example of mean reversion is the price of a stock. If a stock's price rises significantly above its historical average, it may be considered overvalued, and investors might expect it to decrease in price to return to its average. Conversely, if the price falls significantly below the historical average, it may be considered undervalued, and investors might expect it to increase.

4. **Statistical Measures:** Mean reversion can be quantified using statistical measures such as the mean-reverting rate or half-life. The half-life represents the time it takes for an asset's price or indicator to revert halfway back to its mean.
5. **Pairs Trading:** Mean reversion is often used in pairs trading strategies. In this approach, traders identify two related assets (e.g., two stocks in the same industry) and take long and short positions based on the expectation that the spread between the two will revert to its historical mean.
6. **Critique:** While mean reversion is a widely observed phenomenon, it is not a guarantee that prices will revert to the mean. Various factors, including fundamental changes in the underlying asset or market sentiment, can disrupt mean-reverting tendencies.
7. **Research and References:** Mean reversion has been extensively studied in finance literature. Notable academic papers and references related to mean reversion include:

- **“Mean Reversion in Stock Prices: Evidence and Implications”** by Francis X. Diebold and Glenn D. Rudebusch (1991): This paper explores the evidence of mean reversion in stock prices and its implications for financial markets.
- **“Pairs Trading: Quantitative Methods and Analysis”** by Ganapathy Vidyamurthy (2004): This book discusses pairs trading strategies based on mean reversion and provides practical insights and techniques.
- **“Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors”** by Harry M. Kat and Helder Palaro (2009): This paper discusses the mean-reverting behavior of hedge fund indices.

Mean reversion is a fundamental concept in finance, and understanding it can be valuable for investors and traders looking to make informed decisions in financial markets. However, it is important to complement mean reversion analysis with other fundamental and technical factors to make sound investment choices.

## References

- Asimov, I. 1942. “Runaround.” *Astounding Science Fiction* 29:94–103.
- Beutel, A., J. Chen, T. Doshi, H. Qian, A. Woodruff, C. Luu, P. Kreitmann, J. Bischof, and E. H. Chi. 2019. “Putting Fairness Principles into Practice: Challenges, Metrics, and Improvements.” Preprint, submitted January 14, 2019. <https://arxiv.org/abs/1901.04562>.
- Brandon, E. D., and H. O. Welch. 2009. *The History of Financial Planning: The Transformation of Financial Services*. John Wiley & Sons.
- Brennan, T. J., and A. W. Lo. 2011. “The Origin of Behavior.” *Quarterly Journal of Finance* 1 (1): 55–108.

- Brennan, Thomas J., and Andrew W. Lo. 2012. “An Evolutionary Model of Bounded Rationality and Intelligence.” *PLoS One* 7 (11): e50310.
- Brennan, T. J., A. W. Lo, and R. Zhang. 2018. “Variety Is the Spice of Life: Irrational Behavior as Adaptation to Stochastic Environments.” *Quarterly Journal of Finance* 8 (3): 1850009.
- Bubeck, S., V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. T. Ribeiro, and Y. Zhang. 2023. “Sparks of Artificial General Intelligence: Early Experiments with GPT-4.” Preprint, submitted March 22, 2023. <https://arxiv.org/abs/2303.12712>.
- Chen, Mark, et al. 2021. “Evaluating Large Language Models Trained on Code.” Preprint, submitted July 7, 2021. <https://arxiv.org/abs/2107.03374>.
- Damasio, A. R. 1994. *Descartes’ Error: Emotion, Reason, and the Human Brain*. New York: Putnam Publishing.
- Gao, L., S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy. 2020. “The Pile: An 800GB Dataset of Diverse Text for Language Modeling.” Preprint, submitted December 31, 2020. <https://arxiv.org/abs/2101.00027>.
- Kirilenko, Andrei A., and Andrew W. Lo (2013), Moore’s Law versus Murphy’s Law: Algorithmic Trading and Its Discontents, *Journal of Economic Perspectives* 27 (2), 51–72.
- Kurshan, E., J. Chen, V. Storchan, and H. Shen. 2021. “On the Current and Emerging Challenges of Developing Fair and Ethical AI Solutions in Financial Services,” in *Proceedings of the Second ACM International Conference on AI in Finance*. New York: Association for Computing Machinery.
- Labotka, D., and S. Lamas. 2023. *Why Do Investors Fire Their Financial Advisor?* Chicago: Morningstar.
- Lo, A. W. 2004. “The Adaptive Markets Hypothesis.” *Journal of Portfolio Management* 30 (5): 15–29.
- Lo, A. W. 2005. “Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis.” *Journal of Investment Consulting* 7:21–44.
- Lo, A. W. 2017a. *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton, NJ: Princeton University Press.
- Lo, Andrew W. 2017b. “Moore’s Law vs. Murphy’s Law in the Financial System: Who’s Winning?” *Journal of Investment Management* 15 (1), 17–38.
- Lo, A. W. 2021. “The Financial System Red in Tooth And Claw: 75 Years of Co-evolving Markets and Technology.” *Financial Analysts Journal* 77 (3): 5–33.

Lo, A. W., and J. Ross. 2024. “Can ChatGPT Plan Your Retirement?: How Large Language Models Can Evolve to Provide Trusted Financial Advice.” Preprint, submitted January 14, 2024.

<http://dx.doi.org/10.2139/ssrn.4722780>.

Lo, A. W., M. Singh, and ChatGPT. 2023. “From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial Applications.” *Journal of Portfolio Management* 49 (7): 49.

Lo, A. W., and R. Zhang. 2022. “The Wisdom of Crowds vs. The Madness of Mobs: An Evolutionary Model of Bias, Polarization, and Other Challenges to Collective Intelligence.”

<https://dx.doi.org/10.2139/ssrn.3979339>.

Lo, A. W., and R. Zhang. 2024. *The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics*. Oxford, UK: Oxford University Press.

McKenna, N., T. Li, L. Cheng, M. J. Hosseini, M. Johnson, and M. Steedman. 2023. “Sources of Hallucination by Large Language Models on Inference Tasks.” Preprint, submitted May 23, 2023.

<https://arxiv.org/abs/2305.14552>.

OpenAI, et al. 2023. “GPT-4 Technical Report.” Preprint, submitted March 15, 2023.

<https://arxiv.org/abs/2303.08774>.

Rolls, E. T. 1999. *The Brain and Emotion*. Oxford, UK: Oxford University Press.

Ross, S. A. 1973. “The Economic Theory of Agency: The Principal’s Problem.” *American Economic Review* 62:134–39.

Samuelson, P. A. 1947. *Foundations of Economic Analysis*. Harvard University Press.

Serapio-Garcia, Greg, Mustafa Safdari, Clement Crepy, Luning Sun, Stephen Fitz, Peter Romero, Marwa Abdulhai, Aleksandra Faust, and Maja Mataric. 2023. “Personality Traits in Large Language Models.” Preprint, submitted July 1, 2023.

<https://arxiv.org/abs/2307.00184>.

Singhal, Karan, et al. 2023. “Towards Expert-Level Medical Question Answering with Large Language Models.” Preprint, submitted May 16, 2023. <https://arxiv.org/abs/2305.09617>.

Srivastava, Aarohi et al. “Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models.” Preprint, submitted June 9, 2022. <https://arxiv.org/abs/2206.04615>.

Villalobos, P., J. Sevilla, L. Heim, T. Besiroglu, M. Hobbhahn, and A. Ho. 2022. “Will We Run Out of Data? An Analysis of the Limits of Scaling Datasets in Machine Learning.” Preprint, submitted October 26, 2022.

<https://arxiv.org/abs/2211.04325>.

von Neumann, J., and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

Wei, J., Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus. 2022. “Emergent Abilities of Large Language Models.” Preprint, submitted June 15, 2022. <https://arxiv.org/abs/2206.07682>.

Wu, S., O. Irsoy, S. Lu, V. Dabrovolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann. 2023. “BloombergGPT: A Large Language Model for Finance.” Preprint, submitted March 30, 2023. <https://arxiv.org/abs/2303.17564>.

Zellers, R., A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi. 2020. “Defending Against Neural Fake News.” Preprint, updated December 11, 2020. <https://arxiv.org/abs/1905.12616>.

Zhang, Ruixun, Thomas J. Brennan, and Andrew W. Lo. 2014a. “Group Selection as Behavioral Adaptation to Systematic Risk.” *PLoS One* 9 (10): e110848.

Zhang, Ruixun, Thomas J. Brennan, and Andrew W. Lo. 2014b. “The Origin of Risk Aversion.” *Proceedings of the National Academy of Sciences of the United States of America* 111 (50): 17777–82.

## Footnotes

1. Details about GPT-4 can be found in the technical report (OpenAI 2023). ↵
2. See Appendix A.1. ↵
3. An example of a scaled up dataset is The Pile, which contains over 800GB of text data (Gao et al. 2020).  
↵
4. Machine learning researchers have characterized these emergent abilities by developing various benchmarks to measure their performance. Though these benchmarks are not complete tallies of the range of potential LLM emergent behaviors, they are helpful as a guide to characterize LLMs today. Common benchmarks of emergence include the ability to solve multi-digit problems in arithmetic, to answer brain teasers correctly, and to demonstrate knowledge about history and the law in language models not specifically trained to perform these tasks. Microsoft researchers have described how the most recent iteration of the GPT class of LLMs, GPT-4, is able to solve benchmark tasks in arithmetic, coding, medicine, law, and more as emergent behaviors (Bubeck et al. 2023) as well as documenting new emergent behaviors including creative visual tasks, such as describing in detail the layout of a comic book page about a suggested subject. ↵

5. In addition to dataset bias, there are many other factors that can cause a model to be biased, such as algorithmic bias. [↵](#)
6. Specifically, as the market opened on August 1, 2012, Knight issued a surge of unintended orders electronically, many of which were executed, resulting in a rapid accumulation of positions “unrestricted by volume caps” that created significant swings in the prices of 148 stocks between 9:30 and 10:00am. Unable to void most of these unintentional trades, Knight Capital was forced to liquidate them at market prices, resulting in a \$457.6 million loss that wiped out its entire capital base. Its share price plunged 70%, and Knight was forced to seek a rescuer and was acquired by competing broker-dealer GETCO in December 2012. What could have caused this disaster? Apparently, the SEC determined that this was the result of a program functionality called “Power Peg,” which had not been used since 2003 and had not been fully deleted from Knight’s systems. The most surprising aspect of this cautionary tale was that Knight was widely considered to be one of the best electronic market makers in the industry, with telecommunications systems and trading algorithms far ahead of most of their competition. See Lo (2017b) for further details. [↵](#)
7. The quality of this support is one of the top reasons investors fire their advisors (Labotka and Lamas 2023) [↵](#)
8. See Lo (2017a, chapter 4) and the idea of “acquired sociopathy” coined by Damasio (1994). [↵](#)
9. See Lo and Ross (2024) for details. [↵](#)
10. See Brennan and Lo (2012), Zhang, Brennan, and Lo (2014a,b), Brennan, Lo, and Zhang (2018), Lo and Zhang (2022), and Lo and Zhang (2024). [↵](#)
11. See, for example, Damasio (1994). [↵](#)
12. “Developed by JLL Technologies (JLLT), the technology division of JLL, the bespoke generative artificial intelligence (AI) model will be used by JLL’s 103,000+ workforce around the world to provide CRE insights to clients in a whole new way. JLL’s extensive in-house data will be supplemented with external CRE sources, and the company plans to offer made-to-order solutions to clients later this year.” <https://www.jll.de/en/newsroom/jll-unveils-first-gpt-model-for-commercial-real-estate> (accessed December 18, 2023). [↵](#)

## References

- Asimov, I. 1942. “Runaround.” *Astounding Science Fiction* 29:94–103.  
[↵](#)
- Beutel, A., J. Chen, T. Doshi, H. Qian, A. Woodruff, C. Luu, P. Kreitmann, J. Bischof, and E. H. Chi. 2019. “Putting Fairness Principles into Practice: Challenges, Metrics, and Improvements.” Preprint, submitted



- Brennan, T. J., A. W. Lo, and R. Zhang. 2018. “Variety Is the Spice of Life: Irrational Behavior as Adaptation to Stochastic Environments.” *Quarterly Journal of Finance* 8 (3): 1850009.



- Brennan, Thomas J., and Andrew W. Lo. 2012. “An Evolutionary Model of Bounded Rationality and Intelligence.” *PLoS One* 7 (11): e50310.



- Chen, Mark, et al. 2021. “Evaluating Large Language Models Trained on Code.” Preprint, submitted July 7, 2021. <https://arxiv.org/abs/2107.03374>



- Kirilenko, Andrei A., and Andrew W. Lo (2013), Moore’s Law versus Murphy’s Law: Algorithmic Trading and Its Discontents, *Journal of Economic Perspectives* 27 (2), 51–72



- Kurshan, E., J. Chen, V. Storchan, and H. Shen. 2021. “On the Current and Emerging Challenges of Developing Fair and Ethical AI Solutions in Financial Services,” in *Proceedings of the Second ACM International Conference on AI in Finance*. New York: Association for Computing Machinery.



- Lo, A. W. 2004. “The Adaptive Markets Hypothesis.” *Journal of Portfolio Management* 30 (5): 15–29.



- Lo, A. W. 2005. “Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis.” *Journal of Investment Consulting* 7:21–44.



- Lo, A. W. 2017a. *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton, NJ: Princeton University Press.



- Lo, A. W. 2021. “The Financial System Red in Tooth And Claw: 75 Years of Co-evolving Markets and Technology.” *Financial Analysts Journal* 77 (3): 5–33.



- Lo, A. W., and J. Ross. 2024. “Can ChatGPT Plan Your Retirement?: How Large Language Models Can Evolve to Provide Trusted Financial Advice.” Preprint, submitted January 14, 2024. <http://dx.doi.org/10.2139/ssrn.4722780>

↑

- Lo, A. W., and R. Zhang. 2024. *The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics*. Oxford, UK: Oxford University Press.

↑

- Lo, A. W., M. Singh, and ChatGPT. 2023. “From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial Applications.” *Journal of Portfolio Management* 49 (7): 49.

↑

- Lo, Andrew W. 2017b. “Moore’s Law vs. Murphy’s Law in the Financial System: Who’s Winning?” *Journal of Investment Management* 15 (1), 17–38.

↑

- McKenna, N., T. Li, L. Cheng, M. J. Hosseini, M. Johnson, and M. Steedman. 2023. “Sources of Hallucination by Large Language Models on Inference Tasks.” Preprint, submitted May 23, 2023. <https://arxiv.org/abs/2305.14552>.

↑

- Rolls, E. T. 1999. *The Brain and Emotion*. Oxford, UK: Oxford University Press.

↑

- Ross, S. A. 1973. “The Economic Theory of Agency: The Principal’s Problem.” *American Economic Review* 62:134–39.

↑

- Samuelson, P. A. 1947. *Foundations of Economic Analysis*. Harvard University Press.

↑

- Serapio-Garcia, Greg, Mustafa Safdari, Clement Crepy, Luning Sun, Stephen Fitz, Peter Romero, Marwa Abdulhai, Aleksandra Faust, and Maja Mataric. 2023. “Personality Traits in Large Language Models.” Preprint, submitted July 1, 2023. <https://arxiv.org/abs/2307.00184>.

↑

- Singhal, Karan, et al. 2023. “Towards Expert-Level Medical Question Answering with Large Language Models.” Preprint, submitted May 16, 2023. <https://arxiv.org/abs/2305.09617>.

↑

- Srivastava, Aarohi et al. “Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models.” Preprint, submitted June 9, 2022. <https://arxiv.org/abs/2206.04615>.

↑

- Villalobos, P., J. Sevilla, L. Heim, T. Besiroglu, M. Hobbhahn, and A. Ho. 2022. “Will We Run Out of Data? An Analysis of the Limits of Scaling Datasets in Machine Learning.” Preprint, submitted October 26, 2022. <https://arxiv.org/abs/2211.04325>.

↵

- von Neumann, J., and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

↵

- Wei, J., Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus. 2022. “Emergent Abilities of Large Language Models.” Preprint, submitted June 15, 2022. <https://arxiv.org/abs/2206.07682>.

↵

- Wu, S., O. Irsoy, S. Lu, V. Dabrovolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann. 2023. “BloombergGPT: A Large Language Model for Finance.” Preprint, submitted March 30, 2023. <https://arxiv.org/abs/2303.17564>.

↵