Generative Artificial Intelligence

WHITEPAPER
FEBRUARY 2024

CONTRIBUTORS
COLIN BEBEE

Key takeaways

- Potential use cases of AI across both personal and commercial levels are being broadly explored. There is likely both "overselling" and "underselling" as it pertains to capabilities, but one thing is sure – the space is evolving rapidly.
- Research and exploration around artificial intelligence has been occurring for roughly seventy years, with various successes and failures over those decades. The current generative AI wave (i.e., ChatGPT-like models) represents a leap in capabilities compared to the approaches of just a few years ago.
- Most developed economies face demographic headwinds due to aging populations that may eventually detract from their economic standing. Generative AI represents a transformational technology that may help fill certain holes left by an aging population.
- From an investment standpoint, generative AI is unlikely to alter the fundamental drivers of the cost of capital. If widespread utilization proves successful, we expect broad economic growth to experience a moderate tailwind over the coming decades due to both productivity gains in existing fields and new products/services that have yet to be invented.
- From an investment management perspective, we expect generative AI to increase the potential for alpha in the short term, as some early adopters find innovative ways to make use of AI. Likewise, we expect that broad adoption and copycat strategies will gradually deteriorate this alpha, and markets will likely remain fairly efficient as alpha becomes harder to generate over the long term.

Introduction

ChatGPT was released to the public on November 30, 2022. Almost immediately, artificial intelligence ("AI") has become one of the leading topics of discussion across industry and social circles. This paper seeks to provide a high-level overview of the most recent era of artificial intelligence (i.e., generative artificial intelligence) and potential implications as it pertains to global capital markets.

Generative artificial intelligence ("generative AI" or "gen AI") refers to the emerging technology of artificial intelligence models that can produce content

based on the data on which they were trained. Despite origins that date back over half a century, the most recent developments of this concept represent a leap forward in technological capabilities.

With any new technology, society typically seeks to extrapolate the potential uses and implications as it pertains to certain domains (e.g., capital markets). One conclusion from our research is that predictions about the future of AI should be viewed cautiously. Although the current technology already appears expansive, it is still extremely early in its development and advances are occurring at a rapid pace.

Our primary conclusion regarding generative AI's implications for investors is that it may improve efficiencies across many industries, and it may eventually become as commonplace as a spreadsheet. It also has the potential to alter economic growth/development and how investors work on a day-to-day basis.

Definitional overview

This section seeks to outline common terms, subjects, and approaches as it relates to the broad category of artificial intelligence.

Artificial Intelligence

Artificial intelligence ("AI") is the broadest category – all other items we define represent various subfields or approaches within this massive domain. Simply defined, AI is the ability of electronic systems to mimic human intelligence and human cognitive functions such as problem solving and response.[1] A very important consideration is that because we still do not fully understand how the human brain works, AI systems are more focused on emulating human intelligence rather than mirroring the exact processes by which human cognition works. This is an important system design decision that has led to some of the most recent advancements in AI.

[1] Source: https://www.ibm.com/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks/ (July 6, 2023 – IBM Data and AI Team)

Machine Learning

Machine learning ("ML") is the second broadest category and is a subset of AI. Machine learning is focused on the ability of computers to systematically analyze data associated with a given phenomenon and develop associated algorithms and/or models that represent, summarize, and/or replicate the underlying process. Alternatively, ML focuses on the use of data and algorithms to imitate the way humans learn, gradually improving its accuracy.[2] Most importantly, ML is the technique by which society is developing AI.

[2] Source: https://www.ibm.com/topics/machine-learning

Neural Network

A neural network is one methodology, approach, or structure by which machine learning can be conducted. There are many other approaches by which machine learning is conducted, however, the most novel applications currently are neural networks (e.g., ChatGPT and its competitors). Neural networks get their name and structure from the desire to imitate the human brain, and in particular, how biological neurons send signals to one another. Neural networks consist of different layers of nodes (i.e., artificial neurons) that are connected to one another. Data is passed through the input layer with different connection weights and paths among the layers of nodes determining what is ultimately sent to the output layer (see Figure 1).

FIGURE 1
Deep Neural Network

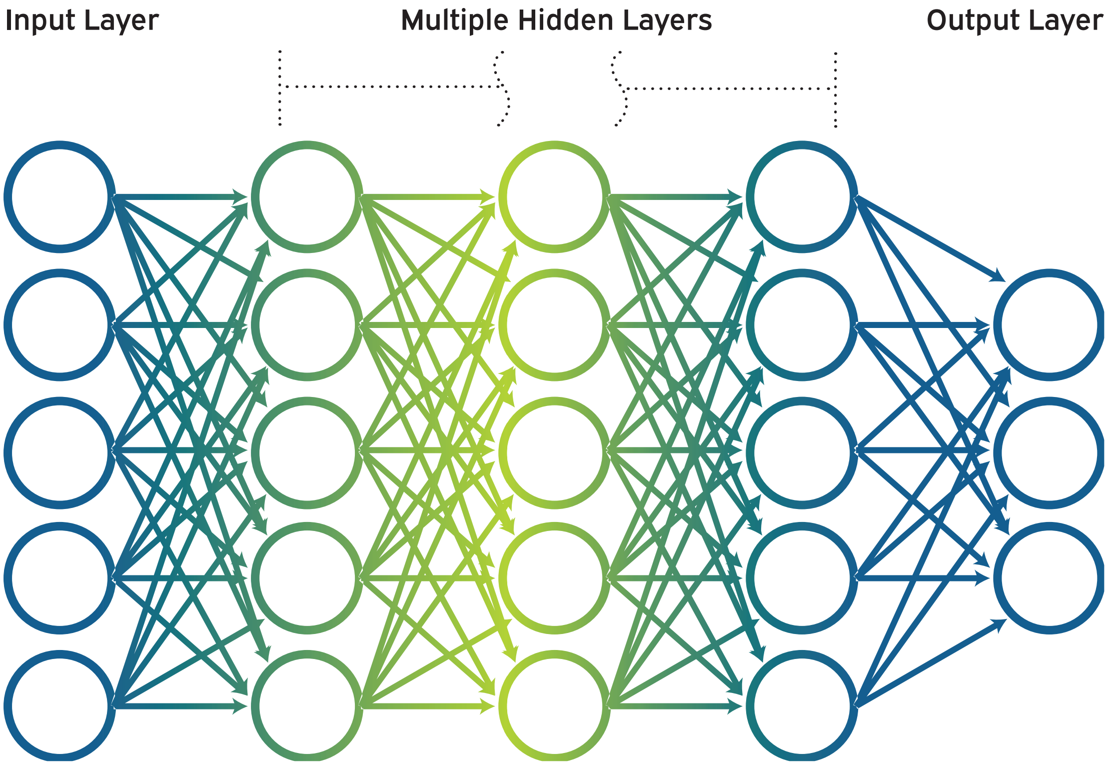

Source: Meketa Investment Group, 2024.

The process of "training" a neural network consists of determining the parameters of the network, with the weights (i.e., strengths) of the connections between nodes being one of the most important parameters.[3] As indicated in the graphic above, there is an input layer and output layer. The process of training a neural network is a machine learning process that calibrates the parameters of the neural network such that, given an arbitrary input, the output is the more accurate. For example, given millions of pictures of cats and other items, what parameters result in the neural network identifying those images with cats in them with the greatest accuracy. Most modern generative AI models (e.g., ChatGPT) use multiple iterations of training before the model is released.

[3] A related parameter is a "bias" which is effectively a constant that is added to the calculation at the node level.

Deep Learning

Several terms within the broad AI field either overlap or are used interchangeably. Deep learning and neural networks are two of those overlapping concepts. At the highest level, "deep" refers to the notion of a neural network with multiple hidden layers, and "learning" refers to actual parameterization of the neural network. In other words, deep learning is an approach by which complex neural networks are trained.

Figure 2 seeks to represent the basic relationships among the most common terms:

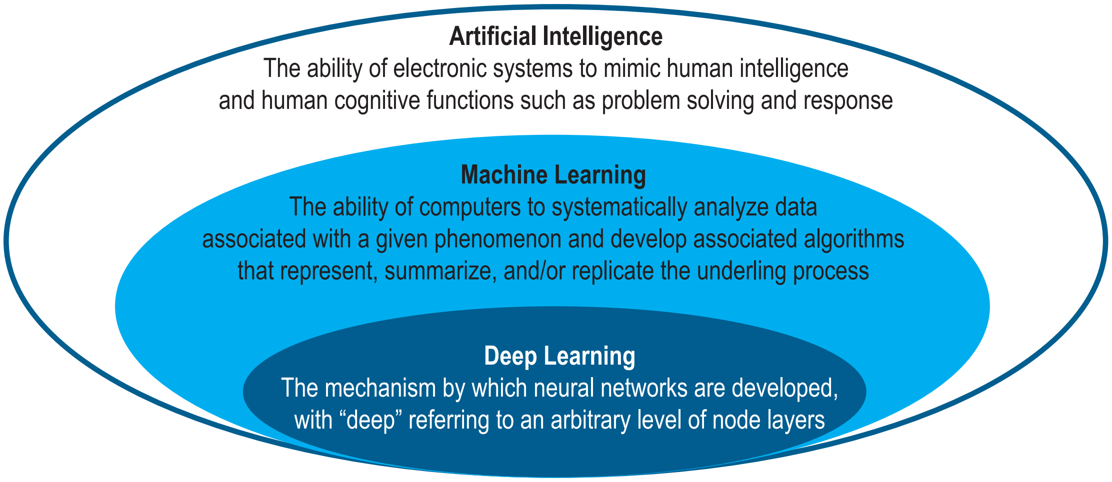

FIGURE 2
Terminology Relationships

Source: Meketa Investment Group, 2024.

Historical evolution

Given that artificial intelligence ("AI") is a subset of computer science, its development timeline is parallel to that of electronic computers. The origins of both fields trace back to the mid-20th century, around World War II.[4] In 1943, Warren McCulloch and Walter Pitts, a psychologist and logician/mathematician, respectively, proposed that the brain operates through simple operations within a complex network of neurons, illustrating how neural networks can execute complex calculations.[5] About a decade later, the term 'artificial intelligence' was coined at the Dartmouth Summer Research Project conference, organized by John McCarthy and Marvin Minsky. From this point forward, artificial intelligence has been riding the tailwinds of computational development as computers, specifically hardware up until the 1990s, saw revolutionary improvements as vacuum tubes gave way to transistors, which were eventually combined with other components on integrated circuits.[6]

Similar to other emerging fields, there was no clear consensus on the optimal direction for AI. Specifically, artificial intelligence has generally been categorized into two camps: Symbolic AI and Connectionist AI.[7]

[4] The ENIAC, or Electronic Numerical Integrator and Computer, was the first programmable, general-purpose electronic computer and was publicly released in 1946. https://www.seas.upenn.edu/about/history-heritage/eniac/

[5] Source: *A Logical Calculus of the Ideas Immanent in Nervous Activity* – Warren McCulloch and Walter Pitts, Bulletin of Mathematic Biophysics, Volume 5, 1943.

[6] For a detailed history of computational development, please see: Miller, C. (2022). Chip War: The fight for the world's most critical technology. Scribner Book Company.

[7] Source: Dominique Cardon, Jean-Philippe Cointet et Antoine Mazières, Neurons Spike Back. The Invention of Inductive Machines and the Artificial Intelligence Controversy, Réseaux 2018/5 (n° 211), pp. 173-220.

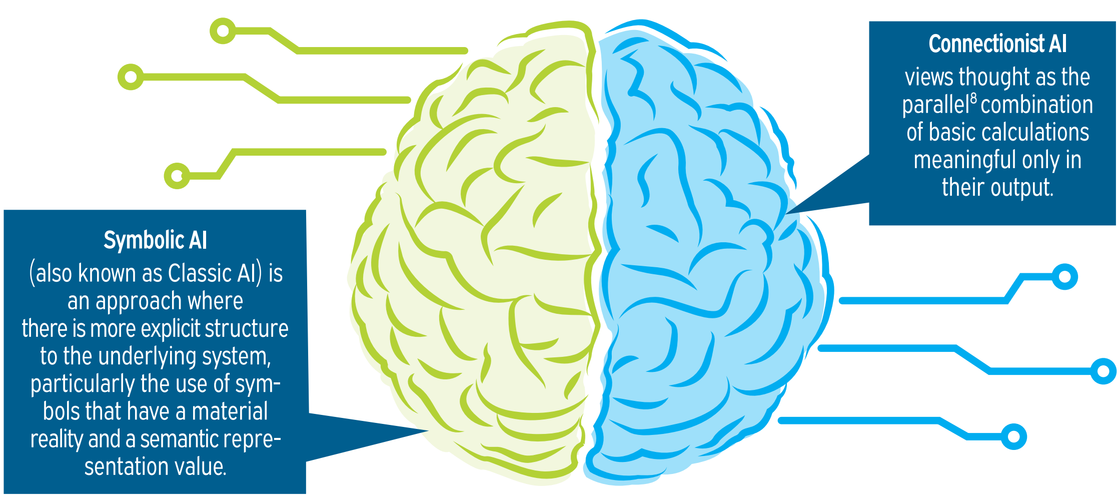

FIGURE 3
AI Schools of Thought

Source: Meketa Investment Group, 2024.

[8] The concept of parallelism is important to the recent advancements in AI, as the growth of graphics processing units ("GPUs") was a significant event given their inherent parallel calculation abilities compared to traditional central processing units ("CPUs").

The most basic examples of the different approaches would be expert systems for Symbolic AI, in which a computer is programmed to emulate the decision-making of an expert, generally through proposition logic and rules, and for Connectionist AI, artificial neural networks are the primary embodiment.

One of the key features of Symbolic AI is that it requires less data for training/learning. Instead, it requires more explicit structure and "knowledge" embedded in the system. As the original wave of Connectionist AI proved underwhelming, Symbolic AI began to dominate both research and application development. However, a material issue with this approach was the creation of the actual knowledge bases.

"The knowledge is currently acquired in a very painstaking way that reminds one of cottage industries, in which individual computer scientists work with individual experts in disciplines painstakingly to explicate heuristics. In the decades to come, we must have more automatic means for replacing what is currently a very tedious, time-consuming, and expensive procedure. The problem of knowledge acquisition is the key bottleneck problem in artificial intelligence."

Edward Feigenbaum, 1980, Knowledge Engineering: The Applied Side of Artificial Intelligence

With material headwinds for Symbolic AI across both research and commercial adoption, the industry began to shift again to more Connectionist AI approaches – specifically a shift away from rule-based AI to probabilistic methods.[9] The success of this shift was possible largely because of significant developments made in computational processing and data availability. In particular, the utilization of GPUs[10] (graphics processing units) for AI/ML purposes combined with the improved accessibility of data (largely due to the internet and decreased cost/efficiency of data storage/access) made feasible large-scale AI training processes that aligned with the connectionist approach. Furthermore, one of the most important developments in AI history occurred in 2017 with the advent of the "transformer architecture"[11] for neural networks, as introduced in the "Attention is All You Need" paper from Google Brain.[12] It was the combination of capability (due to improved computer hardware) and methodology (e.g., transformer architecture) that created the recipe for success we are now seeing across the AI industry.

Neural networks, as discussed, comprise multiple layers of interconnected nodes. The actual structure and process by which data is fed through the neural network (and how the network is trained on data) is determined by the architecture of the network. While there are various neural network architectures, and certain architectures work better on different modalities, the transformer architecture that was introduced in 2017 (originally for text) has taken over the industry and is the leading architecture behind the most prominent AI models of 2023 (e.g., ChatGPT, LLaMA, etc.).

Surpassing the original authors' expectations, the transformer architecture has spurred rapid advancements in AI. This architecture has proven to exhibit high degrees of flexibility, speed, and accuracy compared to alternatives.[13] There is a reason why the

[9] One of the more well-known papers was produced by the IBM Research Division in 1988 titled *A Statistical Approach to Language Translation*.

[10] GPUs allow for massive amounts of simple computations to be executed in parallel. Because neural networks are simply mathematically connected nodes, the associated computations are very basic and can be executed in parallel at large-scale. CPUs (central processing units) by contrast are not explicitly designed for parallel computations, but rather are designed for more complex computations executed in a sequential order.

[11] "Architecture" refers to the structure/design of a neural networks, such as its main parameters (e.g., number of nodes, layers, connections, etc.) and how data flows through the network.

[12] Source: Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention is all you need." Advances in neural information processing systems 30 (2017).

[13] Source: *Emergent and Predictable Memorization in Large Language Models*, Biderman, Stella, Prashanth, Sai, Sutawika, Lintang, Schoelkopf, Hailey, Anthony, Quentin, Purohit, Shivanshu, and Raff, Edward, 2023.

term is embedded in many of the model acronyms released across the industry:

- → OpenAI's ChatGPT = Chat Generative Pretrained Transformer
- → Google's BERT = Bidirectional Encoder Representations from Transformers

Current day capabilities

Historically, 2022 is set to be recognized as the year marking the onset of widespread consumer adoption of the next generation of artificial intelligence ("AI") systems. While elements of AI have infiltrated many technologies with which humans interact (e.g., movie recommendation systems, spam filtering, text completion, etc.), the pace at which technologies such as ChatGPT have spread has never before been seen.[14] While ChatGPT is the most well-known of the recently released products, many others have seen broad adoption, including those across other modalities than basic text/language.

[14] It is widely reported that ChatGPT took only five days to achieve one million users, a level that has typically taken months or years for products to achieve. Source: https://www.digitalinformationworld.com/2023/01/chat-gpt-achieved-one-million-users-in.html

FIGURE 4

Popular Generative AI Products/Companies

DALL-E	A generative AI system from OpenAI that generates images from text.
Stable Diffusion	A generative AI system from Stability AI that generates images from text.
Midjourney	A generative AI system from Midjourney that generates images from text.
Runway	An art (e.g., image, video, etc.) generative AI company and suite of products.
BARD	A generative AI system from Google[15] that engages in conversations based on their PaLM 2 model.
ChatGPT	A generative AI system from OpenAI that engages in conversations based on their GPT3.5 or GPT 4 models.

Source: Meketa Investment Group, 2024.

[15] Like many items from Google, there is a proliferation of similar yet different AI models/products including BERT, LaMDA, PaLM, Imagen, and MusicLM. They may be designed for specific use cases, modalities, or they may represent updated approaches.

A key concept of all the products/companies above is that of "Foundation Models": AI neural networks trained on massive, unlabeled datasets to handle a wide variety of tasks. Originally outlined in a 2021 paper from Stanford,[16] Foundation Models are forming the new paradigm of AI. A key element of these models is that they form the foundation for more task-specific products and applications. For example, ChatGPT is technically an application (or system) based on the GPT-3, GPT-3.5, and GPT-4 models (depending on the release date and free vs. paid version of ChatGPT). Similar to other foundational technologies (e.g., internal combustion engines), additional technologies or products can work side-by-side or on top of the foundation model.

[16] Source: *On the Opportunities and Risks of Foundation Models*, Stanford Institute for Human-Centered Artificial Intelligence, https://arxiv.org/pdf/2108.07258.pdf

Large Language Models ("LLM") represent the next evolution in natural language processing. LLMs are the most general, yet relevant, models as it pertains to potential economic and investment implications.[17] These models are trained on enormous amounts of textual data with the ultimate goal of discovering explicit and implicit structural relationships.

[17] Additionally, many systems/models are becoming "multimodal" in that they are explicitly designed to utilize inputs/outputs other than text (e.g., audio and visual) or they are an LLM that explicitly leverages other tools to access other modalities.

At their core, LLMs such as the GPT series that power ChatGPT seek to statistically determine what word should be next – each time adding a word.[18] To accomplish this, LLMs develop probabilistic estimates of words and language constructs by analyzing

[18] More precisely, they are seeking to find the next "token" which can be just parts of a word.

an incredible amount of data. To be specific, GPT-3 (the original LLM for ChatGPT) was trained on 570 gigabytes of textual data across a wide range of sources.[19] This equates to roughly 300 billion words.[20] As a comparison, the Oxford English Dictionary includes more than 500,000 words and phrases, many of which are no longer in use.[21] This implies that the average word is seen hundreds of thousands of times in the training data. Utilizing the transformer architecture that was discussed earlier, GPT-3 (and subsequent versions GPT-3.5 and GPT-4) were able to represent "language" as a neural network defined by 175 billion (GPT-3) or more parameters (i.e., weights and biases of the network nodes/connections).

[19] Source: https://openai.com/research/improving-language-model-behavior

[20] The training data included languages other than English including programming languages which has been one of the powerful elements of utility.

[21] Source: https://www.oed.com/information/about-the-oed

One of the most amazing, yet often underappreciated, elements of the LLM training process is that this is believed to be close to how humans learn. For millennia, the transfer of knowledge was almost exclusively done through written and spoken language. If an LLM is simply a statistical representation of the training data, then this resembles the human learning process. Furthermore, similar to human knowledge, it is not fully understood how much is represented as memorization vs. modeled inference or reasoning.[22]

[22] While model designers do understand the mathematical operations that occur at each step of a model, they still do not fully understand exactly how the combinations of parameters interact to produce the model outputs.

It is important to recognize that LLMs are effectively "lossy compression" technologies – that is they seek to represent the training data in smaller/summary form, but they are unable to convert back to the original data (as opposed to lossless compression).[23] What this means is that there is no guarantee that the output from an LLM perfectly matches the original input (i.e., training data). This is further complicated by the fact that the training data may contain contrary or overlapping information. LLMs are dimension reducing – taking all the complexities and relationships of the training set and representing the information as a probabilistic model. This has several important implications (some of which we will expand on in later sections). For example, LLMs do not know where they "learned" anything. That is, they cannot cite the exact source of the information they may be reiterating.[24] In cases where sources are cited, this is typically due to the broader AI system having access to additional tools, such as a web browser or domain-specific text, where this information can be further explored.

[23] The larger the models get, such as GPT-4, the less "loss" likely occurs, which theoretically improves their factual accuracy.

[24] Many generative AI systems now have access to other tools (e.g., the web browser) that seek to improve fact retrieval.

"They cannot tell whether they know about Harry Potter because they read about him in Joanne Rowling's books or because they read about Joanne Rowling's books in movie reviews. They cannot tell how popular Harry Potter is based on how many reviews about him they read. They don't know if they have read any specific document. They don't know how many times they read some document. They cannot write an excerpt from some book you name by title. They cannot know for sure if they read a book you name by title."

Andriy Burkov, Machine Learning Lead at TalentNeuron,
Author of The Hundred-Page Machine Learning Book

Limited by their explicit design as a "probabilistic model", LLMs, by themselves, cannot be relied upon for factual information with absolute certainty. This is further compounded by the training data itself, which very likely contains an abundance of erroneous information.[25] Moreover, most LLMs contain a "randomness" parameter (referred to as temperature in ChatGPT) that inherently modifies the output's

[25] The utilization of information found on the internet as a core component of the training data relies on the concept that, in aggregate, factual information will be present at a higher frequency than incorrect information on a given subject.

creativity and variation. This combination of features in LLMs often manifests itself in what are referred to as "hallucinations", which is the phenomenon whereby output is incorrect, nonsensical, or not real.[26]

Strengths, weaknesses, and unknowns of current LLMs

The most profound strengths of current LLMs pertain to their ability to understand and respond to textual information. This can come in the form of knowledge retrieval, summarization, classification, and follow-on execution of tasks. Because language is a foundational form of human communication, and thus drives the way humans work and interact, LLMs potentially offer a turn-key package for tasks that involve processing textual information and guidance. This has resulted in LLMs possessing the ability to pass a wide variety of standardized tests, as shown in the figure below.[27]

[26] A well-cited event occurred in 2023 when two lawyers were fined for the utilization of ChatGPT to submit fictious legal findings. Source: https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/

[27] Source: GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models, OpenAI, OpenResearch, University of Pennsylvania, August 22, 2023. https://arxiv.org/pdf/2303.10130.pdf

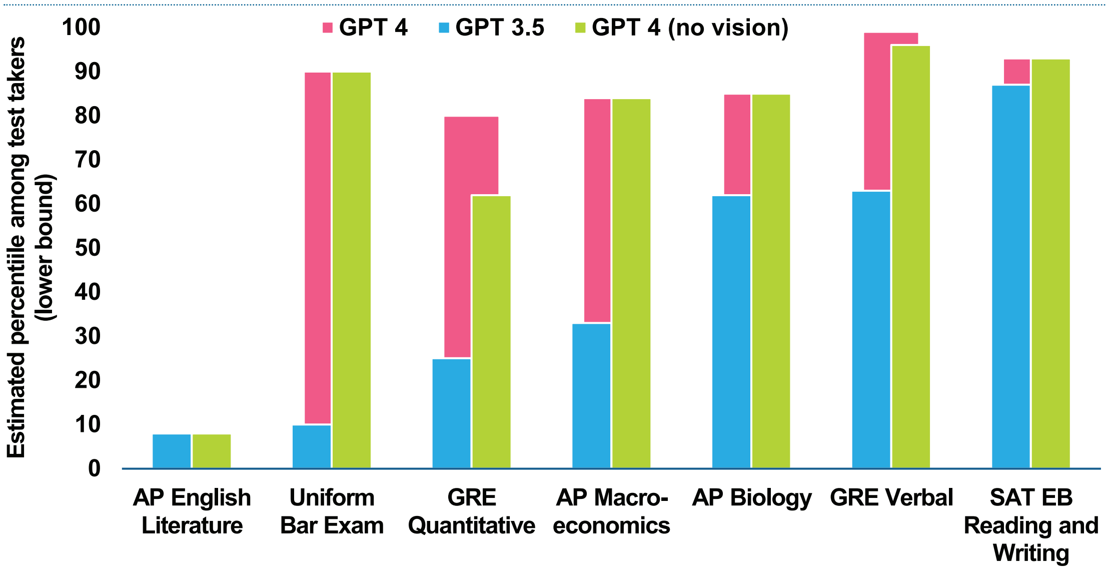

FIGURE 5
Exam Results (ordered by GPT 3.5 performance)

Source: OpenAI

Based on such test-taking results, LLMs (in their base foundation model form) are effectively equivalent to an above-average college student. This is a takeaway that we will come back to later when we describe economic/investment implications in more detail.

The weaknesses of LLMs[28] generally relate to three categories: training data, controllability and reliability, and computational ability.

Training data

LLMs are only as good as the data on which they are trained. This data may be flawed in a variety of ways, including (but not limited to) biases (e.g., stereotyped or prejudiced content), relevance (e.g., historical data that is outdated), and accuracy (e.g., false information). Relatedly, when data on a given topic is relatively sparse in the dataset (e.g., data in languages other than English), the quality of LLM output will suffer. Moreover, as a foundation model, LLMs do not necessarily contain domain-specific information, and the accuracy of their output (even when ignoring hallucinations) will vary across topics.[29] As part of a fine-tuning process, LLMs can improve upon all of these shortcomings to varying degrees and potentially become even more useful/efficient to their user.[30]

[28] For the purposes of this paper, we ignore weaknesses/cons that relate to the costs of developing, training, and using LLMs from an enterprise/commercial perspective. Instead, we focus on their functionality.

[29] While we will discuss this in minor detail later, the process of "fine-tuning" models for specific tasks and/or domains is a significant area of research and empirical progress.

[30] Fine-tuning essentially initializes the model's weights to those from the large-scale pre-training process and then modifies/optimizes the weights based on the fine-tuning data set.

Controllability and reliability

Despite knowledge of all the specific parameters within a given LLM, it is impossible to fully understand exactly how output from an LLM is generated. We do not fully understand all their capabilities and why they fail at certain tasks. Moreover, unlike traditional computational systems that are deterministic (i.e., a given input will produce the same output every time), LLMs have various feedback processes and elements of randomness that severely dampen reproducibility. Furthermore, the tendency to "hallucinate" provides an additional headwind to output reliability.

Computational Ability

Large language models are just that – language models. They represent pattern recognition systems that derive their abilities by uncovering explicit and implicit structures and relationships among language elements. While they are considered to be "general purpose computers,"[31] not all general-purpose computers are created equal. For example, it is well-documented that LLMs not only struggle with mathematical computations, but that their abilities, or quality thereof, may vary over time.[32] More broadly, one of the challenges that LLMs face is that of "computational irreducibility." Computational irreducibility refers to the idea that there are certain processes that cannot be predicted or simplified – knowing the result requires performing the actual process/computation in full. To be fair, this is a challenge that all electronic systems face to certain degrees, but it is particularly challenging for systems such as LLMs because they are not designed to perform specific calculations/computations. Given this conundrum, there are other systems that are specifically designed to address basic and complex calculations. As such, LLMs by themselves will always be limited with respect to their computational abilities and may best be used as a component of a broader system.[33]

Nearly all of the unknowns of LLMs relate to their human-like attributes. For example, there remains a large amount of uncertainty and debate as to LLMs' ability to exhibit consciousness, intuition, or reasoning. However, a growing amount of research indicates that LLMs emulate these attributes to varying degrees.[34] This begs the question of whether it matters if LLMs explicitly have human-like attributes or if they simply emulate them. At this point, the answer is that it likely does not matter. After all, it is still not fully understood how the human brain executes these attributes, although we do generally understand what portion of the brain is responsible for various attributes.

As we examine the major strengths, weaknesses, and unknowns of LLMs, one significant takeaway regarding these systems is that we are presumably still early in their development and deployment. Furthermore, these findings are leading to the idea that entire computational systems may need to be redesigned.[35]

> *"You couldn't just put the electric motor where the steam engine was and leave everything else the same, you had to rewire the entire factory."*[36]
>
> Satya Nadella, Microsoft CEO

[31] A general purpose computer is simply a computer capable of performing a large number of tasks with varying efficiencies.

[32] Source: *How is ChatGPT's Behavior Changing over Time?* Stanford University, UC Berkeley, https://arxiv.org/pdf/2307.09009.pdf

[33] This is exactly how broader systems (such as ChatGPT) accomplish most of their tasks – they use other tools (e.g., web browser, python interpreter, etc.) as needed.

[34] Source: *Thinking Fast and Slow in Large Language Models*, Hagendorff, Thilo, Fabi, Sarah, and Kosinski, Michael, https://arxiv.org/ftp/arxiv/papers/2212/2212.05206.pdf

[35] Source: https://venturebeat.com/ai/inside-the-race-to-build-an-operating-system-for-generative-ai/

[36] Source: https://www.wired.com/story/microsofts-satya-nadella-is-betting-everything-on-ai/

As it relates to new system designs, LLMs themselves can potentially serve various functions but other applications would be required for those areas where LLMs lack efficiency and accuracy (e.g., computation). As it stands right now, however, LLMs are largely being used as a "plug-in" or augmentation to standard systems and processes. As new systems are designed and developed over time, additional insights and capabilities from LLMs and other AI systems will likely emerge.

Economic and investment implications

One of the largest challenges with an emerging technology is accurately predicting future implications. This is further complicated by the fact that the developments within generative AI are occurring at an astonishing pace. Regardless, below we seek to provide some insight into how generative AI may impact three areas: 1) macroeconomy, 2) capital markets, and 3) asset/portfolio management.

Macroeconomy

At the most basic level, generative AI provides relatively cheap and almost unlimited access to human-like capabilities as it relates to knowledge and creative capabilities. This implies that generative AI could improve economic productivity but potentially at the cost of replacing higher price labor. We will first focus on the latter implication.

There are many historical examples of workers fearing that their jobs would be displaced by technology. Beginning with British textile workers during the Industrial Revolution, there has often been a level of "automation anxiety" as emerging technologies take hold.

Although concerns about technology displacing jobs are recurrent, research indicates these fears are often exaggerated, particularly when new technologies are first introduced.[37] This is not to say that all jobs will be safe – any form of technological advance will likely displace some jobs. However, a critical difference between historical job augmentation or replacement due to automation and the potential impacts of generative AI pertain to the types of jobs that may be affected. Historically speaking, the jobs that have primarily been displaced by technology/automation have been those in "blue collar" industries such as basic manufacturing. Generative AI represents the first technology that may not only impact what would be considered "white collar" jobs but could also impact an area that was formerly thought to be largely immune to technology – the arts and other creative industries.

> *"Generative AI is likely to have the biggest impact on knowledge work, particularly activities involving decision making and collaboration, which previously has the lowest potential for automation."*
>
> McKinsey & Company, The economic potential of generative AI (June 2023)

A significant challenge when examining the potential economic implications of generative AI is the wide spectrum of projections from reputable sources. For

[37] For example, a 2016 paper by a Boston University economist examined 271 occupations that were listed in the 1950 census. Only one was eliminated by 2010. "How Computer Automation Affects Occupations: Technology, Jobs, and Skills," James Bessen, Boston University School of Law, Revised Working Paper, October 2016.

example, McKinsey & Company projects that "half of today's work activities could be automated between 2030 and 2060, with a midpoint in 2045."[38] Research from OpenAI is generally in the same direction, but with a different framing: "Our findings indicate that approximately 80% of the US workforce could have at least 10% of their work tasks affected by the introduction of GPTs, while around 19% of workers may see at least 50% of their tasks impacted. The influence spans all wage levels, with higher-income jobs potentially facing greater exposure."[39] The conclusion of most research can generally be summarized as follows: generative AI has the potential to augment and/or replace a segment of industries/roles, and as a result, improve labor productivity. For example, McKinsey states that generative AI could enable labor productivity growth of 0.1% to 0.6% annually.[40] This would represent a meaningful increase in labor productivity compared to the long-term average (1947-2023) of roughly 2.1% and the medium-term average (2007-2023) of approximately 1.4%.

[38] Source: *The economic potential of generative AI – the next productivity frontier*, McKinsey & Company, June 2023.

[39] Source: *GPTs are GPTs: An early look at the labor market impact potential of large language models*, Open AI, March 2023, https://arxiv.org/abs/2303.10130

[40] Source: *The economic potential of generative AI – the next productivity frontier*, McKinsey & Company, June 2023.

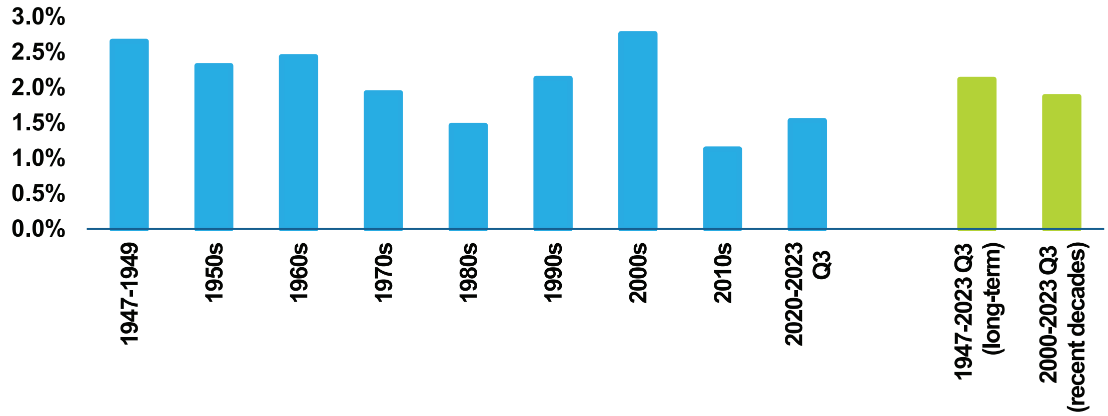

FIGURE 6

Productivity Change in the Nonfarm Business Sector, 1947 Q1 - 2023 Q2

Source: US Bureau of Labor Statistics. Data as of September 7, 2023.

Generative AI also has the potential to assist developed nations fill workforce gaps. It is well documented that most advanced nations are facing material demographic headwinds. Specifically, their populations are getting older and fertility rates are declining to below the estimated replacement rate of 2.1 births per woman.

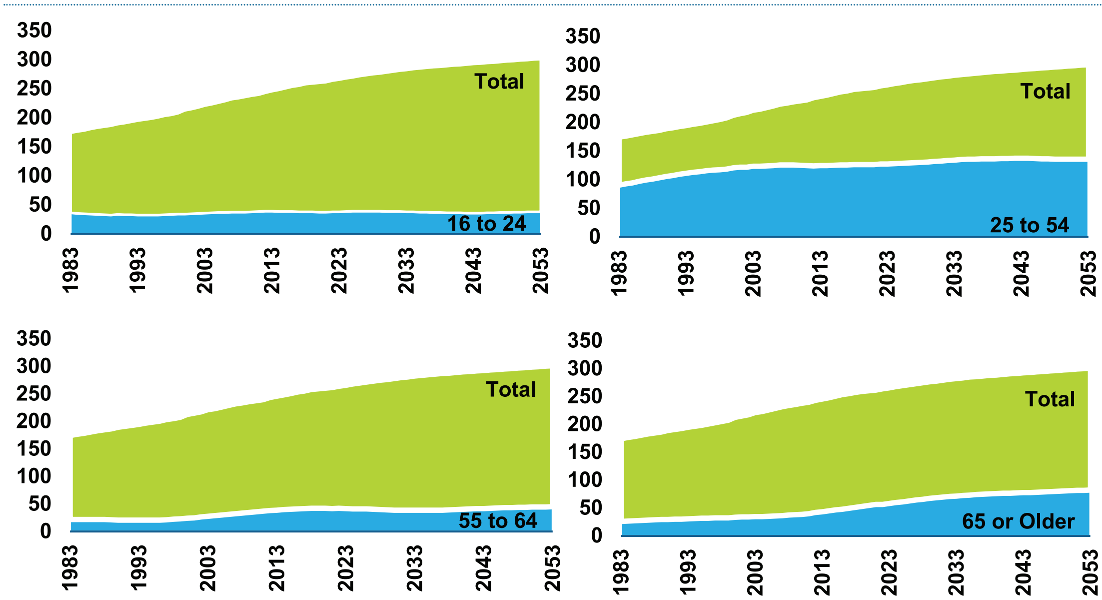

FIGURE 7

Demographic Information: United State

Source: CBO's January 2023 report *The Demographic Outlook: 2023 to 2053.* www.cbo.gov/publication/58612

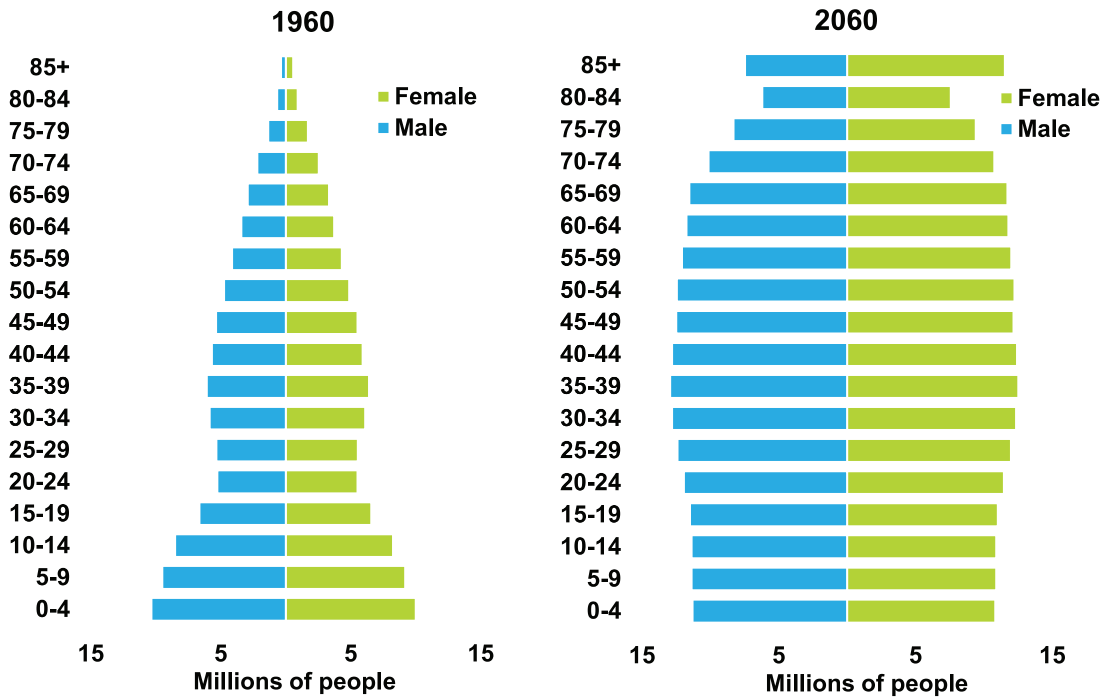

FIGURE 8

From Pyramid to Pillar: A Century of Change

Source: National Population Projections, 2017. www.census.gov/programs-surveys/popproj.html

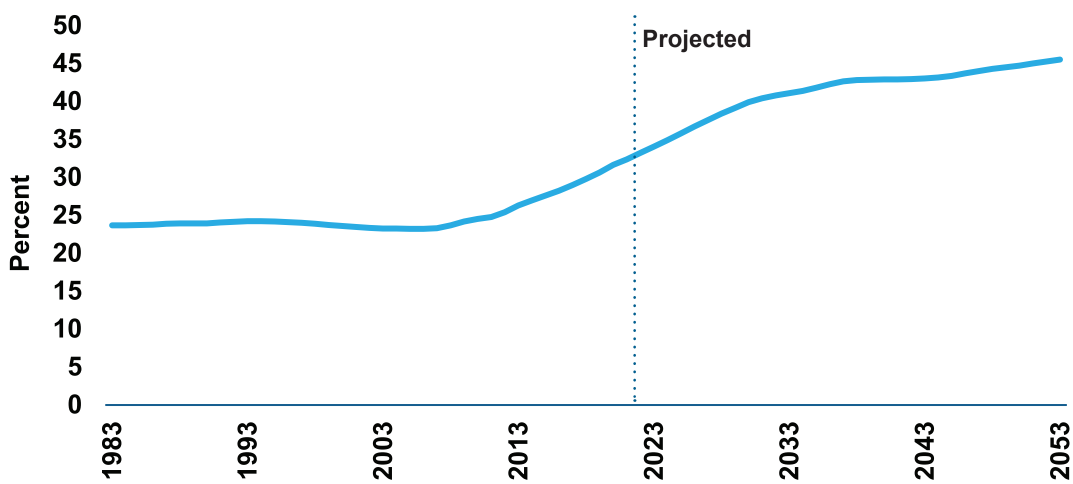

FIGURE 9

Population Age 65 or Older as a Share of the Population Ages 25 to 64

Source: CBO's January 2023 report *The Demographic Outlook: 2023 to 2053.* www.cbo.gov/publication/58612

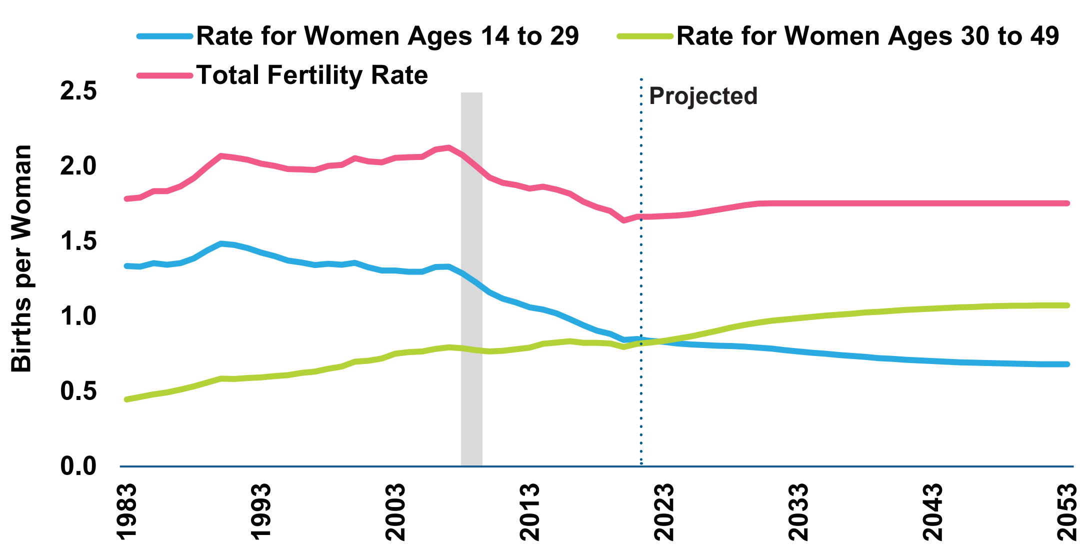

FIGURE 10

Fertility Rates

Source: CBO's January 2023 report *The Demographic Outlook: 2023 to 2053.* www.cbo.gov/publication/58612

With the demographic challenges that advanced nations are facing, the potential for generative AI to augment and/or replace certain tasks may turn out to be more of a necessity rather than a risk – we simply may not have enough individuals of working-age to fill the roles/jobs needed to maintain economic growth. This is one of our main (but preliminary) conclusions – generative AI may be a vital resource in helping advanced nations address one of their many impending challenges (e.g., debt levels, income inequality, climate change, etc.).

Capital Markets

With respect to capital markets, we are primarily concerned with the fundamental drivers of the expected returns across traditional asset classes such as stocks and bonds. To get a sense of the potential impact, we examine the building blocks we use in developing our capital markets expectations. For example, for public equity markets, we project expected returns based on three components: 1) current income (e.g., dividend yield), 2) growth (e.g., earnings growth), and 3) valuation changes.

Under this framework, only the latter two factors have the potential to be impacted by generative AI. As discussed previously, current projections indicate that generative AI can improve productivity, and in-turn, economic growth. This implies that all-else-equal, generative AI has the potential to improve equity market returns. However, economic growth is a highly complex variable that is impacted by a myriad of factors. When examining history, for example, technological advancements have been occurring on a continual basis and thus it is difficult to properly attribute their role on economic growth and equity returns. Moreover, equity returns can be heavily influenced by exogenous factors as well (e.g., wars), despite the consistency of technological progress that occurs during those events. Finally, there is always an element of what is "priced in" to market prices. In other words, current prices (which serve as the foundation that future performance is calculated on) may already reflect the potential for higher economic growth in the future.

As it relates to valuation changes, much of this is determined by the cost of capital and interest rates more broadly, which is heavily influenced by inflation. Regarding inflation, research suggests demographic challenges may lead to increased inflation rates. The thesis is that older populations have large amounts of savings that they can spend yet they contribute less to labor markets as they fade from the work force.[41] Under the assumption that generative AI can help negate some of these demographic challenges, one may conclude that generative AI could be disinflationary. Furthermore, if generative AI lowers the cost of production/services, this would also be disinflationary. Our general conclusion would then be that generative AI may be disinflationary in aggregate. On the topic of real interest rates, research has generally indicated that productivity growth has not generally driven changes in real interest rates.[42]

It is always difficult to project the potential impacts of a given technology on the capital markets, and generative AI is no different. This challenge is further

[41] Source: *The Great Demographic Reversal: Ageing Societies, Waning Inequality, and an Inflation Revival*, Goodhart, Charles, and Pradham, Manoj.

[42] Source: *Productivity Growth and Real Interest Rates in the Long Run*, Kurt Lunsford, Economic Commentary (November 2017), Federal Reserve Bank of Cleveland.

compounded by the technology's relative infancy and pace of evolution. At the highest level, we are comfortable operating under the assumption that generative AI will be a tailwind for economic growth as well as a buffer against inflationary pressures. Exactly how this impacts the projected and experienced returns across asset classes will likely vary in degree and timing. At this point in time, we do not believe that investors should fundamentally change how they develop expected returns for the broad capital markets.

Asset/Portfolio Management

Investors across the globe are exploring the potential usages of generative AI in their workflows. From data analysis and code generation to task automation and content generation, individuals and firms responsible for managing money are in the early stages of deploying generative AI. When it comes to examining the potential implications of these efforts, we may revert back to three mechanisms by which investors can "outperform":[43]

1. Be faster;
2. Be "smarter" at interpreting information; and
3. Use different/better information/data.

When we examine these three mechanisms and their relationship to generative AI, one can make the argument that generative AI has the potential to enable all three. However, because generative AI is becoming increasingly ubiquitous and accessible, this leads to the conclusion that generative AI will likely make the broad capital markets more efficient (i.e., less mispriced assets and thus alpha will become harder to generate) in the long term. Yet, it may enable different sources of alpha/outperformance during its nascent adoption period. We expect generative AI to increase the potential for alpha in the short term, as some early adopters find innovative ways to make use of AI. However, we expect that broad adoption and copycat strategies will eventually make markets more efficient and thus alpha harder to generate, at least in the long term. By extension, this implies that the relative importance of strategic asset allocation will only increase.

Further, it is certainly plausible that generative AI may enable firms to develop alternative datasets (a theme that has been occurring for decades now) that may lead to more insight. However, it should also be expected that the pace of alpha decay[44] will increase due to the declining cost of these technologies.

One of the main takeaways regarding currently available generative AI is that it effectively provides users with the equivalent of an unlimited supply of college student/graduate-like labor for certain tasks. As generative AI technologies improve over time, it seems plausible that the quality of this unlimited labor will increase.

[43] We put "outperform" in quotations because there are many measures/metrics by which this can be calculated, not all of which are used by all market participants, nor would they all be accepted by others.

[44] Alpha decay refers to the time by which an investment edge/advantage is eventually discovered by more market participants, leading to its demise as a source of outperformance over time.

Conclusion

Generative AI is an evolving technology that can create new, multimodal content based on its training data. While it has a long history, recent advancements are groundbreaking, with potential applications across various domains. Across industries, business leaders are exploring as many different applications as possible. There will be many attempts with variable levels of success/failure as AI is incorporated into corporate workflows.

From a macroeconomic and capital markets perspective, generative AI is expected to enhance productivity and should be viewed through a similar lens as prior technological leaps. It is likely that its biggest applications – and disruptions – have yet to be envisioned. One of the largest benefits of generative AI could be its potential to address demographic challenges in advanced nations, filling in the gaps in work forces via improved labor productivity and task replacement. From an industry standpoint, routine white collar and creative jobs are likely to be the main areas disrupted in the near term. These are jobs where “correctness” is either easy to verify and/or subjective.

Across the industry in which Meketa and our clients operate, generative AI may improve speed and access to data/knowledge. Generative AI could also lead to greater alpha potential in the short term, but that alpha will likely decay over time, leading to long-term levels of market efficiency that reinforce the importance of strategic allocation for investors.

While the future impact of generative AI may be unclear, we feel confident that it is here to stay.

Disclaimers

This document is for general information and educational purposes only, and must not be considered investment advice or a recommendation that the reader is to engage in, or refrain from taking, a particular investment-related course of action. Any such advice or recommendation must be tailored to your situation and objectives. You should consult all available information, investment, legal, tax and accounting professionals, before making or executing any investment strategy. You must exercise your own independent judgment when making any investment decision.

All information contained in this document is provided “as is,” without any representations or warranties of any kind. We disclaim all express and implied warranties including those with respect to accuracy, completeness, timeliness, or fitness for a particular purpose. We assume no responsibility for any losses, whether direct, indirect, special or consequential, which arise out of the use of this presentation.

All investments involve risk. There can be no guarantee that the strategies, tactics, and methods discussed in this document will be successful.

Data contained in this document may be obtained from a variety of sources and may be subject to change. We disclaim any and all liability for such data, including without limitation, any express or implied representations or warranties for information or errors contained in, or omissions from, the information. We shall not be liable for any loss or liability suffered by you resulting from the provision to you of such data or your use or reliance in any way thereon.

Nothing in this document should be interpreted to state or imply that past results are an indication of future performance. Investing involves substantial risk. It is highly unlikely that the past will repeat itself. Selecting an advisor, fund, or strategy based solely on past returns is a poor investment strategy. Past performance does not guarantee future results.