The transformative power of generative AI

Supercharged productivity or mass joblessness?

In brief

- Artificial intelligence (AI) technologies, particularly generative AI, appear likely to revolutionize the way we work, innovate and create. Generative AI can create novel, human-like output across various domains, making it highly versatile and intuitive – it even helped us draft parts of this paper – and as such, it has the potential to become a **"general-purpose technology"** like the steam engine and computer, transforming the global economy.
- The most positive economic implication of AI disruption will likely be **accelerating labor productivity** after many years of near stagnation; estimates of the potential impact span a wide range, though most analyses posit 1.5%-3.0% per year globally over the next decade. A boost to labor productivity should result in a similar boost to real GDP.
- We expect a large share of AI's productivity impact to come from **automating many tasks humans currently do**, helping to offset increased retirements, but the potential to accelerate innovation could make productivity gains even more significant.
- AI will also have **significant implications for labor.** Automating some tasks means needing fewer human workers to produce the same output, which could result in transitional job displacement, put downward pressure on wages and increase income inequality. However, if AI technologies stimulate demand, the creation of new jobs and higher overall economic growth should offset job displacement.
- AI may fundamentally change the way we, as humans, drive value in the workplace, requiring us to **focus on the skills where we have a comparative advantage.** These changes may be rapid and unpredictable, increasing the importance of career flexibility, re-training and effective action from governments.
- For markets, AI-driven productivity gains are likely to be **positive for corporate earnings and equity returns**; implications for bonds are more ambiguous, though we think the most likely impact is **modestly higher yields.**
- We remain humble in our projections of the economic and market implications of AI technologies, given **tremendous uncertainty** over how powerful and capable they can become, what kinds of unforeseen innovations and industry transformations they'll cause and, ultimately, how governments and society will respond.

J.P.Morgan
ASSET MANAGEMENT

The transformative potential of generative AI

Artificial intelligence (AI) – or the process of making machines smart – has existed in some form since the 1950s, making occasional headlines, like when the chess-playing AI *DeepBlue* beat Gerry Kasparov in 1997. In recent years, AI has quietly become more prevalent in our day-to-day lives, predicting arrival times of our online delivery orders, populating our social media feeds with personalized ads and filtering spam from our email inboxes. Such applications of **"traditional" AI** (also known as "narrow" or "weak" AI) can be very advanced, even exceeding human expert levels, but they are trained to perform only in specific domains.

Generative AI is the latest stride in AI development, and by contrast, its key ability is to generate novel, open-ended content. The recent launch of several generative AI applications (**Exhibit 1**) has brought this technology to the fingertips of the masses and captured global attention. The most popular of these applications take the form of chatbots, like ChatGPT, powered by **large language models (LLMs)** that string together words based on patterns in vast troves of text data, such as significant slices of the internet. By training related models on other forms of unlabeled, unstructured data,[1] ranging from photographs to the entire bodies of work

Exhibit 1: Generative AI tools can revolutionize the way we create and interpret diverse forms of data

	WRITER	CODER	VISUAL ARTIST	MATHEMATICIAN
Applications	• OpenAI **GPT-3*** (2020) and **GPT-4*** (2023) which power **ChatGPT** • Google **PaLM*** (2022) and **PaLM 2*** (2023) • Anthropic **Claud 2*** (2023)		• OpenAI **DALL-E** (2021) • **Midjourney** (2022) • **Stable Diffusion** (2022)	• Google **Minerva*** (2022)
Capabilities	Writes novel, high-quality responses to prompts that are often indistinguishable from human writing. It can write a college essay, explain a joke, summarize a book, or help draft an email.	Writes functional code in various languages from a specification. For existing code, it can explain, debug, and analyze (e.g., calculating time complexity).	Creates high-quality images—photorealistic or artistic—based on written descriptions. Early progress has also been made on generating video.	Solves complex numerical problems at the college level in subjects including algebra, physics, number theory and machine learning.
	*Describe what generative AI is in four bullets that rhyme.*	*Write a Python function that takes as input a file path to an image, loads the image into memory as a numpy array, then crops the rows and columns around the perimeter if they are darker than a threshold value...*	*Give me an image of a Pomeranian sitting on a throne, wearing a crown, with two tiger soldiers by his side.*	*Assume that variance of the first n natural numbers is 10 and the variance of first m even natural numbers is 16. Compute m+n.*

	• *Generating content, a skill so prime,*
	• *Imitating humans, it's quite sublime,*
	• *From art to text, it can create,*
	• *Machine learning models, it does relate*
		`import numpy as np`
		`import cv2`
		`def crop_dark_perimeter (image_path, threshold=100):`
		`# load the image using OpenCV`
		`image = cv2.imread(image_path)`
		`if image is None:`
		`...`
				*Variance of first* n *natural numbers is* $\frac{n^2-1}{12}$
				Variance of first m *even natural numbers is* $\frac{2^2(m^2-1)}{12}$
				So, $\frac{n^2-1}{12} = 10$ *and* $\frac{2^2(m^2-1)}{12} = 16$
				Solving for n *and* m*, we get* $n = 11$ *and* $m = 7$.
				So, $m + n = 18$.

Source: Google, OpenAI, J.P. Morgan Asset Management. *GPT-3, GPT-4, PaLM, PaLM 2, Claud 2, and Minerva are all large language models (LLMs). Note that all the responses here – including the image – are totally original AI creations.

[1] Structured data is highly organized and made up mostly of tables with rows and columns that define their meaning, such as Excel spreadsheets. Unstructured data is everything else, such as the substantial contents of email messages, books, customer service recordings, images, memes and PowerPoint presentations.

of famous artists, generative AI technologies can create high-quality content spanning a wide range of domains: image, video, audio, text, computer code and even entirely synthetic datasets.

Whereas traditional AI technologies might be able to identify photos with "bumble bees" in them, generative AI could produce a photorealistic image of a bumble bee wearing a hat or write a children's story about a bumble bee learning to fly. As whimsical as these examples sound, this difference matters because the open-ended nature of these tasks represents so much of what we humans do for work (and tend to think ourselves uniquely capable of doing).

Generative AI may still be in its infancy, but the technology has advanced to an extent that we can begin to imagine its transformative implications across the global economy. If a program can author a fictional story about a bumble bee, then it could write a movie script – or at least help automate a big part of the process – a real concern of screenwriters in the United States today. To be sure, screenwriters account for a vanishingly small share of U.S. jobs, but generative AI can also help software engineers write and debug computer code, lawyers research legal opinions and draft contracts or scientists read and summarize dense research papers. Visual and auditory generative AI technologies might likewise automate tasks for jobs ranging from graphic designers to video editors. While generative AI is named for its functional differentiation from traditional AI, we think the more economically significant distinction is how *general* it is.[2] When we start to tally what generative AI could do across the whole economy (**Exhibit 2**), the potential impact seems massive.

To be sure, **the output of generative AI applications is imperfect**, with chatbots like ChatGPT occasionally even including "hallucinations" of false information. Many applications are therefore likely to require a layer of human supervision, especially where the costs of mistakes are high, such as in medicine. However, further progress in developing generative AI might reduce some of these existing imperfections, and even supervised AI could still significantly boost human workers' output.

There is also **plenty that generative AI technology – or any form of AI – cannot yet do.** Interacting with the physical world is still one large obstacle. While robotics has made many impressive advances in recent years (with the notable exception of driverless cars), these systems are designed to perform specific tasks and typically require higher investment and maintenance costs in proportion to their potential output. The robot prototype designed to make guacamole for Chipotle, "Autocado," may quicken the food assembly line, but it cannot also fill a customer's cup or wipe down tables. Even so, in an information age with lots of desk jobs, non-physical problems are a big part of what we do.

Before generative AI, no other technology has arguably had as much potential to automate so much of our work. Such potential has sparked both public excitement and fear – the excitement of ridding ourselves of mundane and time-consuming tasks through automation, but also the fear of losing our jobs and livelihood. In this publication, we seek to answer a few key questions: Is generative AI truly transformational? Could generative AI become the next "general-purpose technology," like the steam engine and computer? What does broad-scale automation mean for labor markets? And, if AI can make us all a lot more productive, what impact will that have on the economy, inflation and financial markets?

More like a steam engine than a smartphone, economically speaking

Generative AI will have broad implications for the economy, but its most significant may be accelerating labor productivity after many years of near stagnation. Labor productivity – total output per unit of labor input – has been the main driver of U.S. economic growth over the last century or more. However, productivity growth stagnated in the last decade, registering just 1.2% per year on average. The last 10 years have seen plenty of technological advancements that have improved many aspects of our lives, yet productivity statistics have told a different story.

[2] Generative AI has the potential to handle a broader, more general range of tasks compared to conventional AI, but it is still not artificial general intelligence (AGI). AGI, an unattained concept, refers to machine intelligence capable of performing any intellectual task humans can do.

Exhibit 2: Generative AI has the potential to accelerate efficiency, quality improvement and innovation across industries

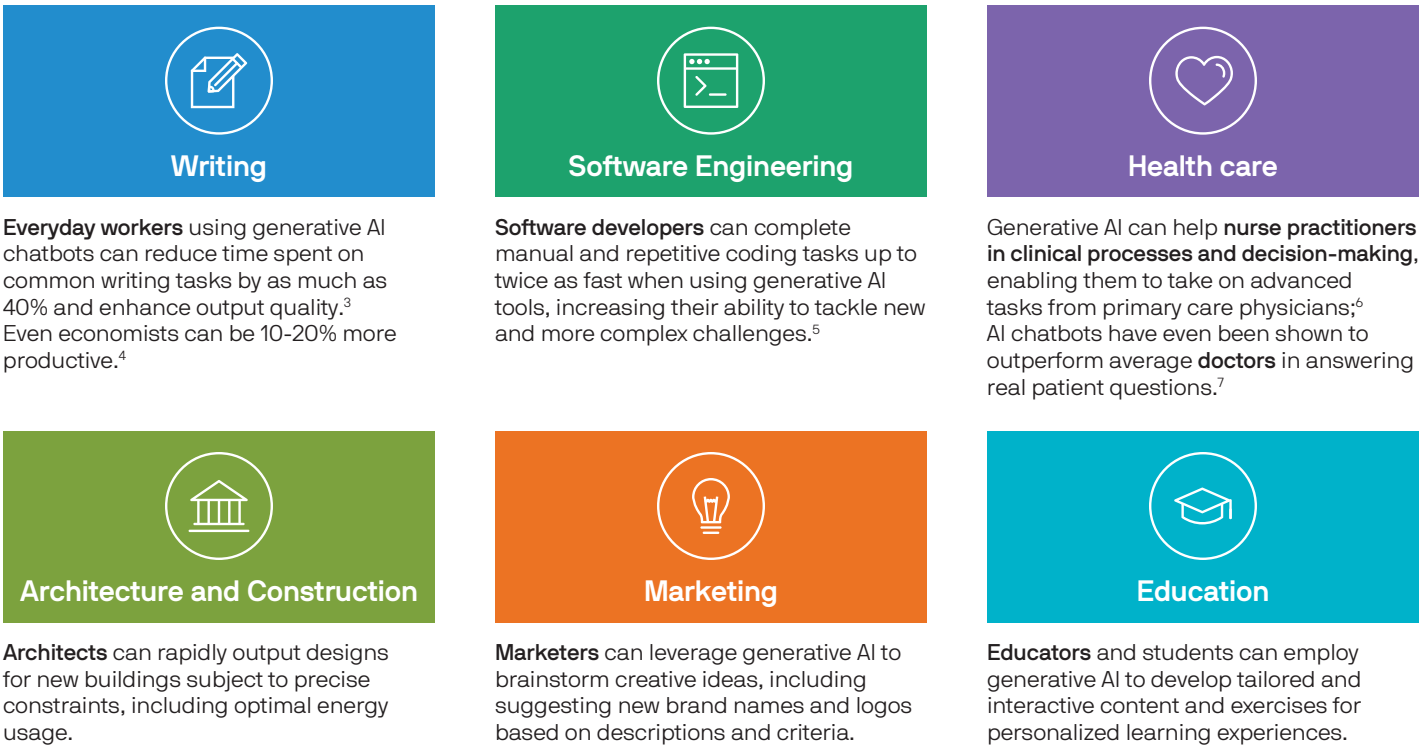

Source: J.P. Morgan Asset Management.

This disconnect may be due to a few factors. Foremost, although advances like smartphones and online media have made huge impacts on our daily lives, since they do so at a relatively low cost to consumers, they have limited impact on the market economy (various studies show a willingness among many consumers to pay for such services they receive for almost free).[8] Additionally, by distracting workers and delivering "information overload," they may detract from productivity in other activities. Some growth disappointment may also simply be due to mismeasurement in government statistics of the real value of new forms of software and human and organizational capital.[9]

Generative AI, by contrast, may be the advancement that finally ushers in a large, sustained boost to productivity. First, the broad-scale automation of existing activity – producing similar outputs with less labor input – should, essentially by definition, result in a more directly measurable productivity impact.

[3] Using ChatGPT for mid-level professional writing tasks allowed workers to reduce time spent by 37% and improve output quality by 0.4 standard deviations. See Noy, Shakked, and Whitney Zhang. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," Science 381, no. 6654, July 2023.

[4] Korinek estimates, based on 25 use cases for language models, that economists can be 10-20% more productive using large language models. See Korinek, Anton. "Language Models and Cognitive Automation for Economic Research," February 1, 2023.

[5] A McKinsey study found that developers using generative AI can increase task completion speed by 35-50% for lower complexity tasks and were 25-30% more likely to complete higher complexity tasks with time savings. Similarly, research by GitHub found that 88% of surveyed developers felt more productive, 73% felt more "in the flow" and 87% spent less mental effort on repetitive tasks when using the AI-powered GitHub Copilot. See Deniz et al. "Unleashing Developer Productivity with Generative AI," McKinsey & Company, June 27, 2023; and Kalliamvakou, Eirini. "Research: Quantifying GitHub Copilot's Impact on Developer Productivity and Happiness" The GitHub Blog, March 17, 2023.

[6] Louis Raymond et al., "Nurse Practitioners' Involvement and Experience with AI-Based Health Technologies: A Systematic Review," Applied Nursing Research 66, August 2022.

[7] In one study on patient questions randomly drawn from a social media forum, chatbot responses were preferred over physician responses and rated significantly higher for both quality and empathy. See John W. Ayers et al., "Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum," JAMA Internal Medicine 183, no. 6, April 2023.

[8] See, for instance, Brynjolfsson, Erik, Avinash Collis, and Felix Eggers, "Using Massive Online Choice Experiments to Measure Changes in Well-Being," Proceedings of the National Academy of Sciences of the United States of America 116, no. 15, March 26, 2019.

[9] Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," November 1, 2017.

Even more profound implications for productivity, output and welfare gains could come if generative AI is the tipping point that enables AI to become the 'general-purpose technology' of the 21st century.[10] Like earlier general-purpose technologies, such as electricity, the steam engine and the internet, generative AI could fundamentally change how a wide range of goods and services are produced, transform industries and create entirely new jobs, owing to its potential to:

- **Be pervasive.** Its *general* capabilities mean generative AI can be integrated in many different contexts to supplement or replace many activities currently done by humans.
- **Spawn complementary technologies and infrastructure.** Companies across industries are rushing to adopt AI in their fields, and the development of ancillary business applications is necessary to fully leverage AI's benefits. As we discuss later in this paper, generative AI may also enhance the performance of existing "traditional AI" technologies and vice versa.
- **Experience exponential growth and economies of scale.** AI's computing workload has been doubling every three to four months since 2012 and is likely to accelerate even further. OpenAI's GPT-3 and GPT-4[11] were released just two years apart, and the latter is significantly more complex, can interpret images received as inputs, is 40% more accurate in its responses and scores significantly higher percentiles on many standardized tests.[12]
- **Reshape industries.** Broad-scale automation will reshape the nature of jobs and business models, with transformative implications across industries.
- **Accelerate innovation.** AI has the potential to accelerate research and development and unlock new insights that inform and inspire innovation efforts. Many leaders in the field think this may be AI's paramount application.

Innovating innovation itself

Expanding on that last point, generative AI's greatest potential might not be in merely automating what humans do, but in enhancing human efforts to create novel solutions to all sorts of real-world problems. Such efforts could lead to considerable productivity and welfare gains beyond automation.

Simply making workers more efficient could perpetually accelerate technological progress.[13] For generative AI specifically, further upside likely lies in its ability to:

- **Quickly sift through vast datasets.** In a world of "information overload," generative AI can be a potent filtering tool, automating many of the time-consuming tasks in research and development.
- **Unlock new ideas and insights that inform researchers of where to concentrate their efforts.** Generative AI can analyze vast troves of unstructured data, something that is virtually impossible for humans to do, and in doing so can identify new patterns, reveal insights and discover better ways of doing things.
- **Conduct comprehensive predictive and evaluative analyses on new ideas.** AI can improve the accuracy of our predictions and models, or even provide a sounding board for new ideas. Try prompting ChatGPT to list the pros and cons of your latent business idea, for instance.

Joining forces with "traditional" AI: greater than the sum of its parts

While generative AI technologies are currently in vogue, recent years have seen the proliferation and refinement of many "traditional" AI technologies that have been trained to perform specific tasks very well. These tasks tend to be the kind where generative systems still fall short, particularly in the performance of accurate predictive modeling, numerical calculations and optimization. McKinsey estimates that these applications will account for a majority of the overall potential economic value added from AI.[14]

[10] Bresnahan and Trajtenberg first coined the term "General Purpose Technology" (GPT) to describe technologies that drove new eras of technological progress and growth. GPTs are characterized by pervasiveness, inherent potential for technical improvements and "innovational complementarities," giving rise to increasing returns-to-scale. See Bresnahan, Timothy F., and Manuel Trajtenberg, "General Purpose Technologies 'Engines of Growth'?," Journal of Econometrics 65, no. 1, January 1, 1995.

[11] Open AI's GPT-3 and GPT-4 refer to generative pre-trained transformers (not to be confused with "general-purpose technology"), which are the family of neural network models that power generative AI applications like ChatGPT.

[12] GPT-4 has 170 trillion parameters compared to GPT-3's 175 billion parameters, enabling many of the improvements in GPT-4 on processing and generating text with greater accuracy and fluency. See OpenAI, "GPT-4" https://openai.com/research/gpt-4.

[13] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work", NBER, April 2023.

[14] McKinsey & Company. "The Economic Potential of Generative AI: The Next Productivity Frontier," June 14, 2023.

Beyond generative AI, traditional AI is still delivering major solutions: some examples

Predicting the complex folding structure of proteins is one of the most exciting use cases of non-generative AI. In the last 60 years, scientists have determined the structure of 180,000 proteins, a small number in proportion to the millions yet undiscovered. This arduous task is an important part of drug discovery, but it can take years to execute. DeepMind's *AlphaFold* is now carrying out the same task in minutes with unprecedented accuracy,[15] a milestone in the application of AI to scientific research with immediate potential to advance drug development, biological research and our understanding of diseases at a molecular level.

Environmental sustainability is another notable application of AI systems. AI systems are increasingly helping optimize energy production, storage, distribution and use. In 2016, Google's DeepMind developed an AI framework which reduced energy usage for data center cooling by 40%.[16] More recently, AI systems are aiding clean energy transitions. Whereas traditional weather models fare poorly at predicting clouds, AI systems trained on satellite and weather data could help solar grid and wind turbine operators optimize power generation and reduce fossil fuel energy held as reserve. In the United Kingdom, *Open Climate Fix* is currently working with the country's electric grid operator to better forecast cloudy British weather.[17]

Source: J.P. Morgan Asset Management.

Integrating generative and traditional AI systems could yield value far beyond what each alone could deliver, since each has its own strengths and weaknesses.

No generative AI systems could achieve the accuracy of AlphaFold's predictions or estimate the exact hours of sunshine tomorrow. Both abilities required specialized training on structured datasets. A generative AI chatbot like ChatGPT even struggles with some simple quantitative reasoning.[18] Ask it to multiply two large numbers and it is likely to produce a close but incorrect answer. However, ChatGPT is fully capable of writing computer code to perform the very same calculation. Simply granting such chatbots access to code interpreters might be one way to supply the correct answer – not unlike calculators help humans solve math problems that most of us couldn't solve in our own heads.

But why stop there? **Generative AI chatbots could draw on the vast library of specialized traditional AI tools,** from mathematics engines to commute time-forecasting models, that have already been quite capable for over a decade – one by one, expanding their capabilities. Already, OpenAI is privately testing several such additions to ChatGPT.[19]

Some of these applications might be highly specialized. For instance, Bloomberg's approach to integrating generative AI into its terminal allows users to prompt a system that is especially fluent in matters of finance, tapping into decades of financial data collection and development of specifically trained models that tackle matters of financial complexity.[20] Indeed, we often hear now that **"English will be the coding language of the future,"** and it seems likely to be in many cases.

[15] "AlphaFold," DeepMind, accessed August 15, 2023, https://www.deepmind.com/research/highlighted-research/alphafold.

[16] "DeepMind AI Reduces Google Data Centre Cooling Bill by 40%," Google DeepMind, accessed August 17, 2023, https://www.deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40.

[17] "Sun in Their AIs: Nonprofit Forecasts Solar Energy for UK Grid," NVIDIA, https://resources.nvidia.com/en-us-energy-utilities/ai-forecasts-solar-e?xs=363482.

[18] Even Google's Minerva, a language model that achieves improved performance by gathering training data primarily from scientific papers, still makes simple mistakes with high frequency. See "Minerva: Solving Quantitative Reasoning Problems with Language Models," Google Research (blog), June 30, 2022.

[19] "ChatGPT Plugins," OpenAI, March 23, 2023, https://openai.com/blog/chatgpt-plugins.

[20] Shaohua Wu et al., "BloombergGPT: A Large Language Model for Finance," ArXiv (Cornell University), March 30, 2023.

Beyond written English, all these capabilities could be made more accessible by incorporating **speech recognition and synthesis,** areas where traditional AI already excels (along with handwriting and image recognition; see **Exhibit 3**). We can imagine, for instance, verbally requesting a conference room smart assistant to draw a new logo idea and display it on a screen, without the need for typing. In comparison to the current generation of "smart assistants" that rely on users' remembering pre-trained command phrases, generative AI could make interacting with all these modules **truly conversational experiences.** Considering that we already spend an estimated 25% of our total work time communicating with one another, being able to communicate just as seamlessly with machines opens the door to working alongside them.

A digitalized speed of adoption

AI's implications for economic growth and societal change can be profound, but the other factor to consider is timing. Although generative AI has suddenly become dinner-table conversation, its ultimate power and impact will not be seen for some time, though this may take place faster than with earlier transformative technologies.

Technological breakthroughs can take considerable time to raise productivity, with the peak impact of many industrial and post-industrial era technological breakthroughs, including general-purpose technologies, often only coming after 20-30 years. In 1987, Robert Solow famously quipped, "You can see the computer age everywhere but in the productivity statistics;" in that case, a significant impact did ultimately show up in productivity statistics, albeit over a decade later (**Exhibit 4**).

Exhibit 3: Traditional and generative AI capabilities are increasingly comparable to those of humans

Test scores of AI relative to human performance
Initial performance for each AI capability set to -100

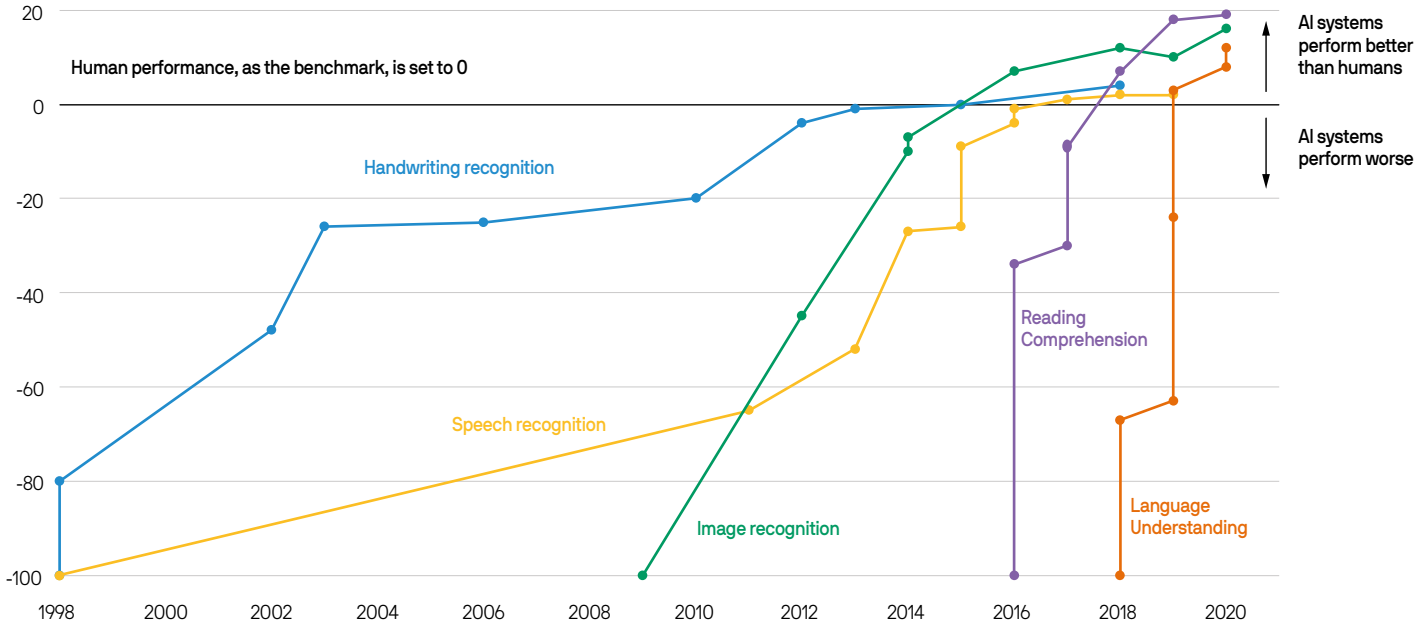

Source: Douwe Kiela, Max Bartolo, Yixin Nie et al., "Dynabench: Rethinking Benchmarking in NLP," Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, June 2021; J.P. Morgan Asset Management. Published online at OurWorldInData.org. Licensed under CC-BY under the author Max Roser.

This lag owes to the considerable time it can take to: (1) establish wide-scale familiarity and access to the technology, (2) reshape business models to integrate the technology, (3) achieve a sufficiently large capital stock of it and (4) develop complementary innovations and infrastructure that allow for full benefits of the technology to be harnessed.

Exhibit 4: Past general-purpose technologies have taken considerable time to deliver gains in labor productivity, which has been the main engine of U.S. GDP over the last century

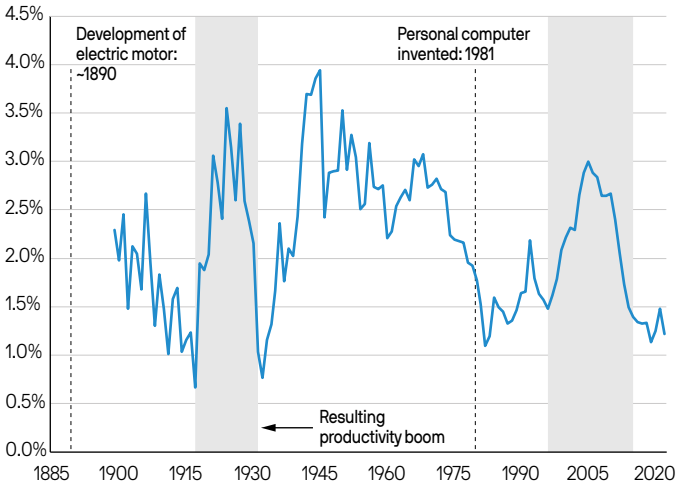

Source: BLS, NBER, J.P. Morgan Asset Management. Data from 1888-1957 reflect productivity data for the total private economy from John Kendrick, "Productivity Trends in the United States," NBER. Data from 1958-2022 reflect non-farm productivity data from the BLS. Data are as of August 31, 2023.

We think that AI adoption could be faster. Over time, an increasingly digitized world has helped accelerate the pace of technological adoption (**Exhibit 5**), and there are some reasons to believe generative AI could be adopted faster still:

- **Generative AI is very accessible and easy to use for the average person,** and its rapid accession to the mainstream is a testament to this fact. ChatGPT shattered records by amassing 100 million monthly users in just three months, compared to the time it took TikTok (nine months) and Instagram (two and a half years) to reach the same milestone.[21]
- **Generative AI is decades in the making, with considerable progress already made.** While much of the public hadn't heard about generative AI until this year, its most notable underlying innovations were developed in 2014.[22] Meanwhile, decades of advancement in cloud infrastructure and an explosion of data and computing power[23] have helped train these systems.
- **Massive business investment has already been made...** In the five years ending in 2021, global business investment in all types of AI grew more than sixfold in real terms, with the United States leading the pack at $73bn invested.[24] From 2017 to 2022, the share of businesses that have adopted AI, and the number of AI capabilities used, more than doubled.[25] Moreover, compared to some earlier technologies, generative AI infrastructure and service providers are bearing a larger share of the necessary capital investments potentially increasing adoption rates by lowering the financial barrier for end users.
- **...and more is underway to integrate AI or develop applications for business use.** Company management teams are increasingly focused on AI, with 40% of S&P 500 management teams mentioning AI in their 2Q 2023 earnings calls, up from 19% a year earlier (**Exhibit 6**). These businesses are rushing to develop AI infrastructure and applications across a wide range of domains, with many launched this year or in the development pipeline.[26] The surge in investor interest has propelled hefty gains for stocks; the global AI market is currently valued at $150bn and projected to grow to $1.3tn by 2030.[27]

[21] Statista, Reuters.

[22] The introduction of generative adversarial networks (GANs) in 2014 marked a breakthrough in generative AI. GANs quickly became one of the most influential generative AI models, allowing for the high-quality generation of images, audio, text and other types of content.

[23] The amount of computing power used to train AI systems has been doubling every six months over the past decade. See Jaime Sevilla et al., "Compute Trends across Three Eras of Machine Learning," ArXiv (Cornell University), February 11, 2022.

[24] Netbase Quid Companies dataset, 2022, as cited in "The 2023 AI Index Report," Stanford Institute for Human-Centered Artificial Intelligence, 2023.

[25] McKinsey & Company. "The State of AI in 2022—and a Half Decade in Review," December 6, 2022.

[26] Sequoia Capital's "AI 50 2023" identifies the emerging trends in privately held AI companies, stemming across large-language models, infrastructure for model training, generative AI applications and predictive AI applications.

[27] "Artificial Intelligence Market Size & Trends, Growth Analysis, Forecast," MarketsandMarkets, June 2023.

Exhibit 5: Technological adoption has accelerated over time

U.S. Technology, Rate of Adoption*

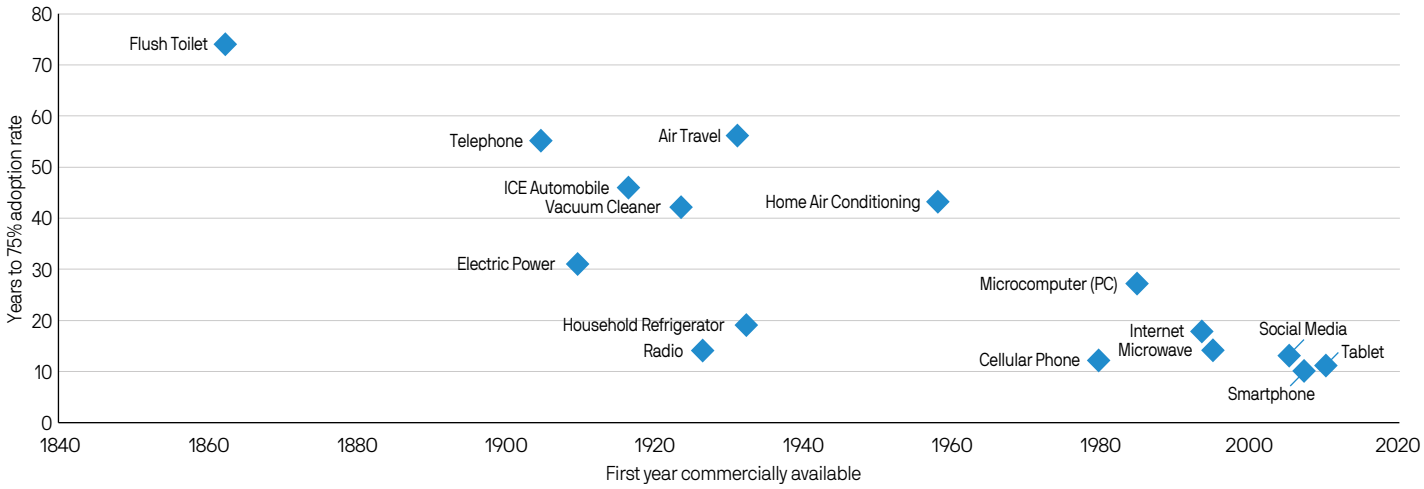

Source: Asymco, compiled from various sources with support of the Clayton Christensen Institute, J.P. Morgan Asset Management.
*Estimated from current adoption trends.

Exhibit 6: 40% of S&P 500 companies mentioned AI in 2Q 2023 earnings calls

Share of S&P 500 companies mentioning AI in earnings calls

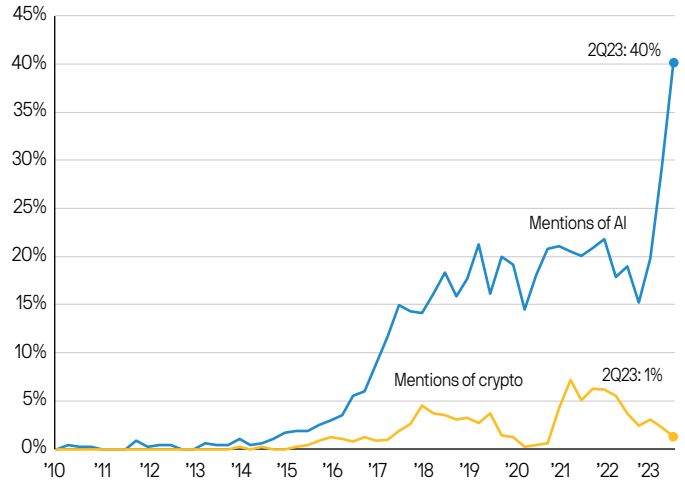

Source: J.P. Morgan Asset Management. Mentions of AI include the keywords: artificial intelligence (AI), deep learning, machine learning, chatbots and natural language processing. Mentions of crypto include the keywords: cryptocurrency, Bitcoin, Ethereum, blockchain, stablecoin and altcoin. Data are quarter-to-date for 2Q23 as of August 25, 2023.

Sizing the potential AI productivity gain

AI appears well positioned to significantly boost labor productivity,[28] but by how much? In our own analysis, we estimate **annual productivity gains between 1.4% and 2.7% per year across developed markets** over 10 years. This estimate, if realized, would be comparable to past periods of technologically driven surges in productivity (as shown earlier in **Exhibit 4**).

Importantly, our estimates quantify the **impact of automation alone.** Such productivity enhancement would be in addition to any other productivity growth, such as the acceleration of innovation, which we believe presents significant upside potential. On the other hand, external factors could partly limit productivity gains, for instance, if an ineffective or overly restrictive regulatory response impedes AI development, or social resistance stymies adoption.

[28] In this publication, we focus on labor productivity, which is real economic output divided by the total number of hours worked. However, total factor productivity (TFP) is often considered a more direct proxy for technological progress since it measures the efficiency with which all economic inputs, including capital investments, are used to produce output. TFP growth represents the portion of output growth that is achieved above the accumulation of these inputs. Since the 1970s, TFP has also seen relatively modest gains – just 0.6% per year in the U.S., according to commonly cited estimates by the Penn World Tables. However, TFP is more difficult to measure and, as a result, historical estimates are less available; additionally, labor productivity may be the more relevant measure from the perspective of individual human workers.

In line with several recent studies that come to similar conclusions, our approximation of the potential productivity impact of AI-related automation takes a task-level approach. We estimate the aggregate time spent on many types of tasks across the whole economy and judge the share of each of these tasks that might be automated (for more details, see *"Appendix: Sizing AI productivity gains"*). Ultimately, we find that **traditional and generative AI applications could potentially automate 14% to 27% of current work activities** in the United States over the next 10 years (we would expect similar results across other developed markets). The wide range of these estimates owes to considerable uncertainty around our assumptions outlined below.

Automation can materialize in productivity gains through three channels (**Exhibit 7**). The first and most straightforwardly positive channel is direct labor cost savings from **fewer workers** being needed to produce the same amount of output. Alternatively, instead of reducing headcounts, companies can produce even **more output** by retaining their more productive workers. The combined size of these two productivity impacts should equal the total productivity-weighted share of tasks automated. Finally, there is a **composition effect** that accounts for changes in the productivity of workers displaced from automation. Our projections assume that this effect is zero, or that, on average, displaced workers are reemployed in new jobs where they are equally as productive as they were in their former ones. We believe this is a conservative assumption, given the potential for AI to lower the barrier to entry of many jobs and because AI seems likely to generate entirely new jobs with high productivity. However, if displaced workers are employed in jobs where they are less productive, overall productivity benefits would be lower; if many are not reemployed or work less, then productivity benefits would not fully accrue to real GDP.

Exhibit 7: Sizing the potential AI productivity boost

1. Direct labor cost savings
2. Increased output from more productive workers
3. Composition effects from re-employment of displaced workers

= Total productivity impact

Source: J.P. Morgan Asset Management.

Ultimately, the greatest uncertainties in estimating the labor productivity boost involve the capabilities of AI itself. Our more conservative estimates of AI capabilities, assuming a broader range of tasks are immune to automation, would see productivity gains of just 0.4%-0.9% per year over 10 years; on the other hand, our most optimistic assumptions would see gains of as much as 7.7% per year. To be sure, we think this upside is extremely unlikely, but its possibility does illustrate the significance automation could have dependent on how powerful and pervasive AI ultimately becomes. Similarly, if our baseline scenario for AI automation were to take 20 rather than 10 years to take hold, then annual productivity gains would be proportionally smaller, i.e., 0.7%-1.4% per year.

Our 2024 forthcoming **Long-Term Capital Market Assumptions** (LTCMAs) will incorporate a small first step in accounting for the role of AI technologies, particularly generative AI, in automating current economic activity. Although we feel it is too early to fully embed the plausible productivity upsides estimated in this paper into our base-case long-run assumptions, we plan to re-evaluate this impact in the coming years and size the productivity gain depending on the progress made on automation.

Work in an age of AI automation

AI may bring on considerable productivity gains, but by doing so through automation, the idea that robots will "take our jobs" is becoming a popular concern. Are we in for mass unemployment?

We don't think so – at least not in the foreseeable future – but the future of work will likely look quite different.

Automating tasks, not jobs

AI seems unlikely to automate many *entire jobs,* but it does have significant potential to automate many of the *tasks* involved in those jobs, with most estimates of aggregate task exposure to automation ranging from 20% to 30%.[29] Such exposure will be broad-based across industries; OpenAI estimates that LLMs could affect 80% of the U.S. workforce in some form.[30]

The degree of exposure, however, varies considerably by job type. Highly exposed jobs include those responsible for documentation and review in legal professions, providing administrative support in businesses and customer service representatives. Given generative AI's advanced abilities to understand language and draw upon vast bodies of information, exposed tasks include those that involve a degree of knowledge or expertise. As such, higher-skilled jobs, such as STEM professionals and health care providers, are also exposed. At the opposite end of the spectrum, where tasks seem the least exposed to automation, are jobs involving physical work or where the human component is invaluable (i.e., construction workers and daycare providers).

Transforming existing jobs

In the vast majority of cases, **AI will *augment* but not entirely replace human capabilities.** While AI's abilities are impressive, there are many domains where AI technologies still fall short compared to humans or benefit from a layer of human supervision and feedback. As repetitive and time-consuming "grunt work" becomes automated, **workers can spend more of their time on higher complexity tasks, meaningful critical-thinking or creative endeavors.** As such, automation can also provide humans with the opportunity to deepen their skills, thereby expanding their overall potential and even increasing happiness at work.[31]

Consider a couple of instances where we already see this dynamic playing out.

- Financial "robo-advisers" can provide customized investment advice and algorithmic portfolio management. A human financial adviser, though, is still needed to provide appropriate advice on financial matters involving complexity, counsel against impulsive trading behavior during market crashes or bubbles and offer empathy in times of crisis. Combine a human financial adviser with the tools of a robo-adviser, and this new "team" has both the benefit of advanced technology, as well as a layer of emotional intelligence for when contextual understanding is necessary.
- Similarly, for software engineers, generative AI can significantly reduce time spent on research and trial and error, especially when working with a new programming language or software framework. With these time savings, one study found that developers were 25-30% more likely to complete higher complexity tasks within the same time limit.[32]

[29] McKinsey estimates that generative AI and other technologies have the potential to automate 30% of current hours worked today in the United States by 2030. Goldman Sachs estimates that roughly two-thirds of the U.S. workforce may be exposed to some degree of automation from AI, with up to one-fourth fully substituted. The OECD estimates that 27% of jobs in major countries are at high risk of automation.

[30] OpenAI, OpenResearch and the University of Pennsylvania estimate that 80% of the U.S. workforce would have at least 10% of their work tasks affected by LLMs, with 19% of workers seeing at least 50% of tasks impacted. See Tyna Eloundou et al., "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," ArXiv (Cornell University), March 17, 2023.

[31] Survey results from the OECD's 2023 Employment Outlook found that 63% of workers in finance and manufacturing said using AI in the workplace improved their enjoyment in the job. Similarly, Noy and Zhang (2023) found that exposure to Chat-GPT among participants using the tool for writing tasks was associated with a substantial increase in job satisfaction of 0.40 standard deviations.

[32] Deniz et al., "Unleashing Developer Productivity with Generative AI."

This pivot will require workers to focus on skills where humans have a comparative advantage (**Exhibit 8**). Human intelligence can still better understand context and matters of complexity, apply intuition and employ emotional intelligence in social settings or where cultural norms are relevant. Place a human professional in a room full of potential new clients, and within minutes she can adapt and respond to social cues across all five senses, adjusting everything from her handshake to the level of detail in her presentation.

Humans also still have some edge in conceptual thinking. Each human mind is intrinsically capable of developing or working with an infinite number of abstract representations and models of the world – from solving physics problems to deciding how to structure an organization. No form of AI boasts such general abilities. Traditional AI applications can work with some such abstract representations, albeit only for the relatively narrow set of instances for which they are designed. Generative AI chatbots, despite producing novel content across a wide variety of domains, do so by essentially echoing patterns in the troves of text on which they are trained. Since this process involves no underlying conceptualization, it falls short in purely conceptual domains like mathematics, as we noted earlier.

Perhaps also owing to this edge in conceptual understanding, humans should maintain an edge in artistic creativity. While AI might soon be able to create a lot of music and visual art that sounds or looks like what has already been produced, it seems less likely to generate entirely new genres on its own.

Where AI's capabilities fall short, humans will be needed to fill in the gaps, and the jobs of tomorrow will be increasingly focused in these areas. Naturally, these are also the sorts of skills we will want to emphasize in education and job training, but many tough questions remain. Can everyone be a great problem-solver? How do we evaluate and train such skills? Do we need to foster creativity in classes, instead of teaching students how to write code?

Artificial intelligence has some clear advantages, including in its ability to process information at a speed impossible for humans to match. In these cases, leaning on artificial intelligence can expand our overall potential.

Exhibit 8: Humans and AI should still have distinct comparative advantages

AI now outperforms average humans on:

- ✔ Handwriting, speech and imagine recognition*
- ✔ Reading comprehension, language understanding*
- ✔ Breadth of knowledge
- ✔ Computational power
- ✔ Speed

Humans still have the advantage on:

- ✔ Conceptual reasoning
- ✔ Depth of understanding
- ✔ Complex problem-solving
- ✔ Emotional intelligence
- ✔ Morality and ethics
- ✔ Creativity
- ✔ Intuition

Source: J.P. Morgan Asset Management. *See **Exhibit 3**.

Exhibit 9: Technological innovation has coincided with the creation of new jobs that account for the bulk of employment growth

Employment, by new and pre-existing occupations
Millions, 2018 employment compared to 1940 occupations

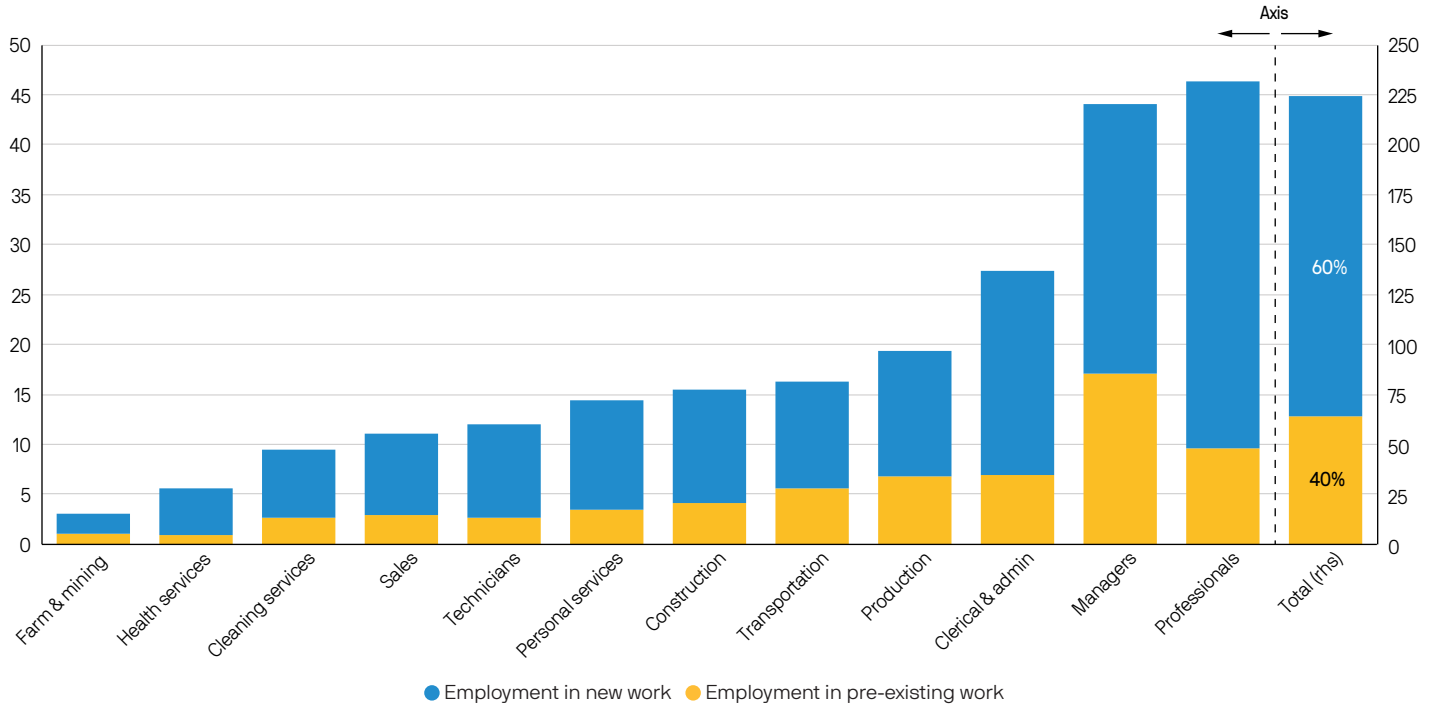

Source: Autor, David, Caroline Chin, Anna M. Salomons, and Bryan Seegmiller. "The New Frontiers: The Origins and Content of New Work, 1940-2018." No.w30389. National Bureau of Economic Research, 2022 (and subsequent calculations by the authors as of August 2023); J.P. Morgan Asset Management.

Creating new jobs

Although automation means that some of us will level up our daily tasks, increasing productivity also means businesses need less workers to create the same output, which could, in turn, result in job cuts.

However, if AI ushers in an economic boom, there should be all sorts of new jobs to which we can pivot. The long history of technological advancement has been associated with the continued process of job displacement and re-instatement that has supported an economy at full employment. For instance, in the early 1800s, over 80% of the U.S. workforce was in agriculture,[33] yet even after improvements in everything from mechanization to crop rotation dramatically reduced the need for farm labor, a subsequent economic boom re-employed many of those formerly working in the fields. Displaced farmers moved to the cities to find jobs in the industrial sector, where they then contributed to an explosion in associated commerce. With less time spent in the fields, people also had more time to spend on art and science. Today, in a similar vein, an estimated 72% of workers are employed in occupations that did not exist in 1940, implying that **over 87% of employment growth over the last 80 years has come from the tech-driven creation of new jobs** (Exhibit 9).

This proliferation of new jobs is ultimately because *productivity stimulates consumer demand* (Exhibit 10)**.** The ability to produce more output with fewer inputs inherently reduces production costs which tends to drive down consumer prices, enriching consumer wallets and enabling the consumption of all sorts of new goods and services. Higher consumer demand then stimulates business demand for workers in new jobs. Indeed, the rise in computerization has been associated with broad employment growth and the birth of many jobs in computer science, software engineering, graphic design, social media marketing and more.

[33] "The Story of U.S. Agricultural Estimates." U.S. Department of Agriculture, April 1969.

Exhibit 10: Higher productivity ultimately drives demand for new jobs

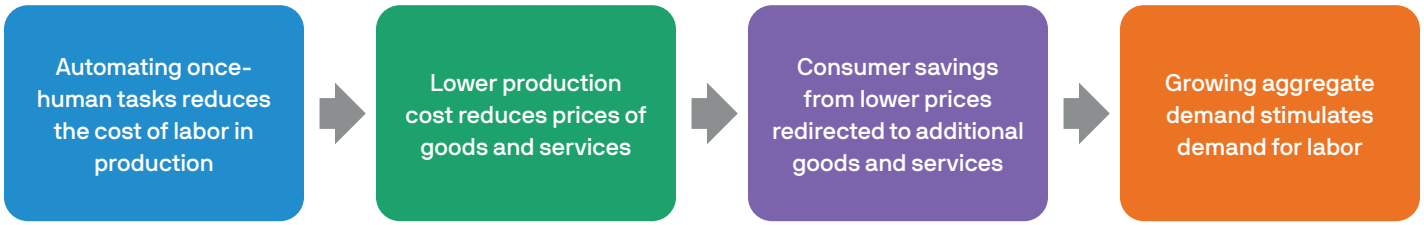

Source: J.P. Morgan Asset Management.

Generative AI ought to similarly increase our purchasing power, ultimately stimulating growth in new occupations – the web designers and app coders of tomorrow. AI, in some instances, may perform the work of a thousand humans for the cost of one. Such economies of scale seem likely to generate new business models we can only begin to imagine, along with demands for humans to complement and manage them. Despite the many headlines about job cuts at the hands of AI, the World Economic Forum found that 50% of global employers expect AI to create job growth versus just 25% who expect it to create job losses.[34]

We may even see labor demand increase in the same sectors undergoing automation. One telling example is the impact that the rise of automated teller machines (ATMs) had on the employment of bank tellers (**Exhibit 11**).[35] Although ATMs automated the cash handling tasks bank tellers had been doing, the number of bank tellers ***grew*** concurrently with the rise in ATMs for about a decade. How did this happen? ATMs allowed banks to operate branches more efficiently, lowering operating costs and prompting banks to open many more branches. Increased accessibility of banking services also spurred greater consumer demand for them – more people wanted to use ATMs, and with it, opted for additional banking services. Ultimately, growth in consumer banking drove greater demand for bank tellers, even if those new branches were staffed with fewer bank tellers per branch. The type of work bank tellers did also changed. With cash handling tasks mostly automated, bank tellers could focus more on customer service and sales, with more of them receiving skills training and college education than they did in the past.

Exhibit 11: The expansion of bank automated teller machines (ATMs) coincided with growth in human bank teller employment

Total employment of bank tellers and ATMs, thousands

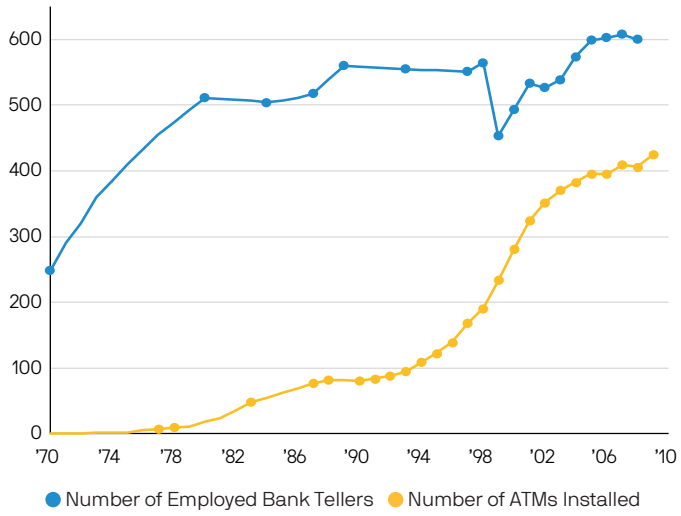

Source: Bank for International Settlements, Bessen (2016), Bureau of Labor Statistics, Occupational Employment Survey, J.P. Morgan Asset Management.

[34] World Economic Forum, "The Future of Jobs Report 2023," May 2023.

[35] James Bessen, "How Computer Automation Affects Occupations: Technology, Jobs, and Skills," Boston Univ. School of Law, Law and Economics, no. 15–49 (October 3, 2016).

Easing labor shortages...

Generative AI may be coming at an opportune time for the global economy when aging populations in most developed regions stand to meaningfully slow growth in the decades ahead. Globally, the ratio of young and elderly people who are economically "dependent" on those of working age will gradually increase. This dependency ratio will rise especially quickly in developed markets, where it has been climbing from a near-bottom of 48% in 2010 and is projected to rise to 82% by 2080.[36] In other words, by 2080 there will be nearly two consumers for every working-age person, and each year until then, these workers will need to support about 0.5% more people. All else constant, workers would face pressure to work longer hours or postpone retirements in order to support the same per capita economic output.[37]

Against this backdrop, AI presents a major opportunity to counterbalance building labor shortages. Every job automated by AI is also one more person who can retire without reducing overall economic output. This dynamic could be especially helpful in the United States, where the most acute labor shortages today include many occupations where AI technologies are likely to have a significant impact, such as healthcare providers and skilled software engineers.

... And even allowing us to work less

With freed up time from automation, and machines doing more of the heavily lifting on driving the economy, workers might also enjoy some more time for family, rest and fun.

Indeed, **over time and across countries today, rising productivity has coincided with fewer hours worked.** 150 years ago, workers in today's richest countries used to work *a lot*, but average working hours declined significantly in the wake of the Second Industrial Revolution (**Exhibit 12a**). In 1930, after a period of particularly large productivity gains, John Maynard Keynes suggested that further advances in technology and productivity might lead to a 15-hour workweek. In the decades since that prediction, though, declines

Exhibit 12a: Over time, people have worked less as productivity has risen

Average weekly working hours per worker

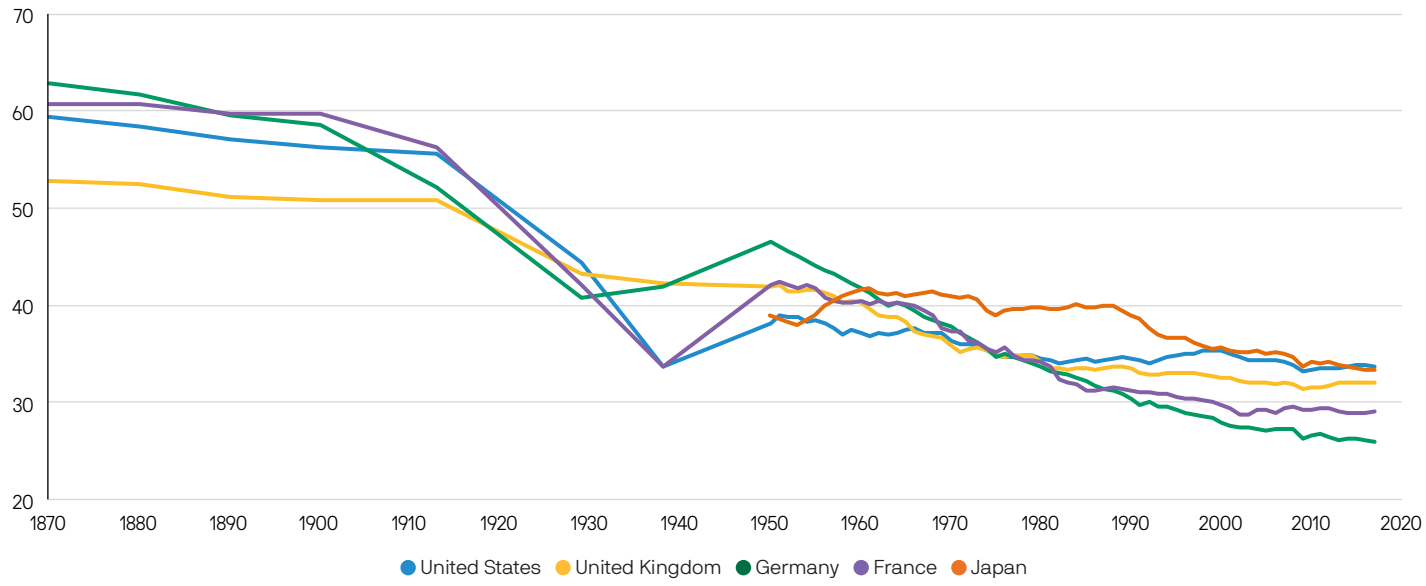

Source: Huberman & Minns (2007), Our World in Data, PWT 9.1 (2019). Published online at OurWorldInData.org. Retrieved from: 'https://ourworldindata.org/time-use'. Licensed under CC-BY by Esteban Ortiz-Ospina, Charlie Giattino and Max Roser.

[36] The United Nations defines the "dependency ratio" as the ratio of young population (under age 15) and elderly population (aged 65 and over) to the working-age population (aged 15 to 64). See "World Population Prospects" dataset, United Nations Department of Economic and Social Affairs Population Division, 2022.

[37] These developed regions could also offset their demographic headwind by increasing net immigration from, or expanding trade deficits with younger, less developed regions.

Exhibit 12b: Today, people work less when they work in more productive economies

Weekly working hours vs. labor productivity

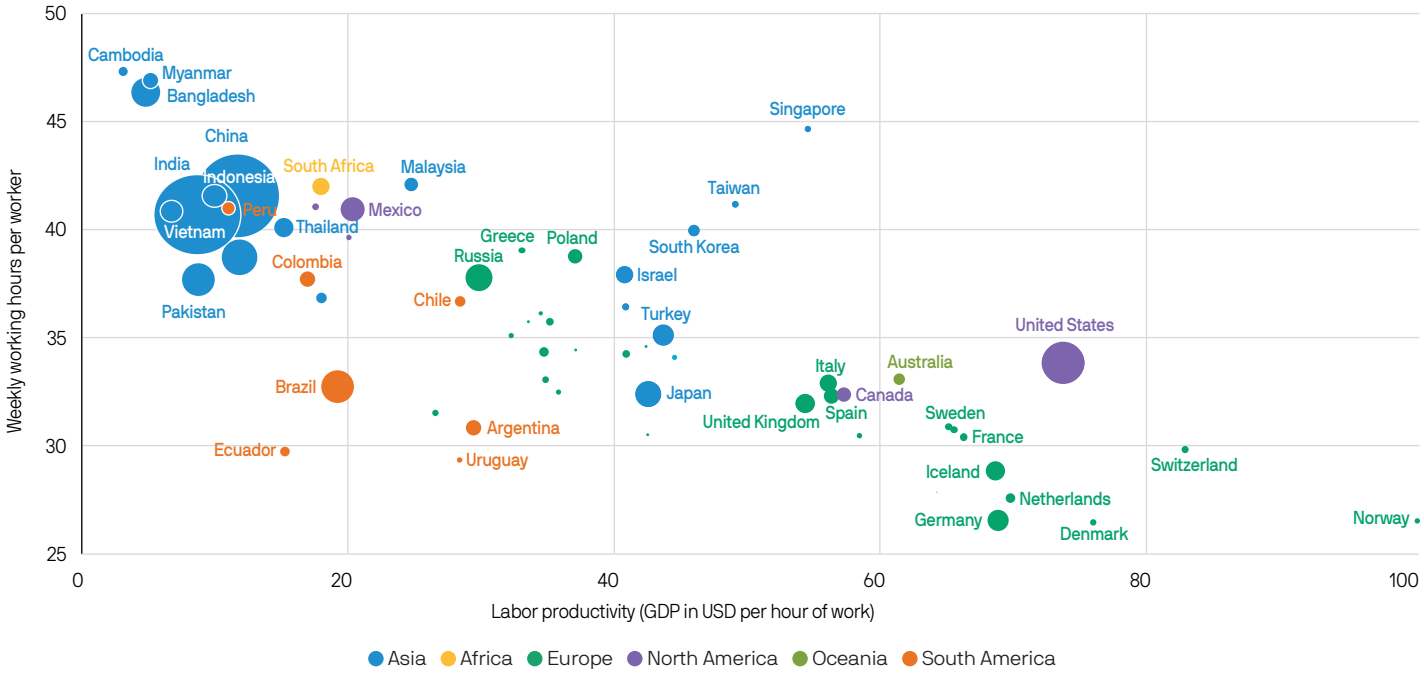

Source: Feenstra et al. (2015), Our World in Data, Penn World Table (2021), J.P. Morgan Asset Management. GDP is adjusted for inflation and for differences in the cost of living between countries. Feenstra, R. C., Inklaar, R. and Timmer, M.P. (2015), "The Next Generation of the Penn World Table". American Economic Review, 105(10), 3150-3182. Published online at OurWorldInData.org. Retrieved from: 'https://ourworldindata.org/time-use' [Online Resource] Licensed under CC-BY.

in working hours have leveled off for major developed economies despite further technological progress. It may be that some cultures intrinsically value hard work, while consumerism may keep us from ever feeling like we have enough. However, across countries today, the relationship between labor productivity and hours worked is consistently negative, suggesting we could still get closer to Keynes' vision (**Exhibit 12b**).

For developed regions, **we estimate that a hypothetical AI-driven 30% increase in labor productivity over the coming decade could drive a 5%-10% reduction in average hours worked.** However, in order to realize this outcome, individual workers will also need to earn enough income that they are able to give up potential working hours in exchange for leisure. That, in turn, requires mitigating further pressure on income inequality.

Considerations for income inequality

The automation of routine and manual work, or the potential to work less, are exciting prospects. For many, the cost could be increasing income inequality, which for many reasons has been rising across developed markets in recent decades. In the United States, the share of pre-tax national income accounted for by the top 10% of earners has grown from 34% to 57% since 1951, leaving the bottom 50% with only 10% of national income.[38] A similar trend is seen in wealth inequality, with the impressive growth in financial assets[39] concentrating economic gains among those with the means to invest. Some argue that technological advancement has played a significant role in these trends, with one study estimating that automation explains 50 to 70% of the increase in wage inequality from 1980 to 2016.[40]

[38] Pre-tax national income represents total labor and capital income before taxes and excludes government transfers. Prior to 1976, income is defined as market income and excludes government transfers but includes capital gains and is sourced from "Income Inequality in the United States, 1913-1998" by Thomas Piketty and Emmanuel Saez, updated to 2021. Top decile includes all families with annual income above $135,000. Data for 2022 are J.P. Morgan Asset Management estimates utilizing data sourced from realtimeinequality.org.

[39] The ratio of U.S. financial assets relative to nominal GDP has grown from 2.3x in the late 1970s to about 4.3x today.

[40] Acemoglu and Restrepo argue that a significant portion of the rise in US wage inequality over the last four decades has been driven by automation (and to a lesser extent offshoring) displacing certain workgroups from employment opportunities for which they had comparative advantage. See Acemoglu, Daron, and Pascual Restrepo. "Tasks, Automation, and the Rise in US Wage Inequality," NBER, June 1, 2021.

A key pitfall of automation is that it can lead to the concentration of gains in the holders of capital, at least initially. If some part of what a worker does is replaced with an AI program, then the owner of that AI capital will receive the "wages" the worker used to earn. While this may propel gains for technology companies, and their investors, this dynamic doesn't bode well for labor's share of income, particularly in an economy where worker bargaining power has already dwindled.[41] Over time, these inequalities can fade, as cost savings from automation pass through to consumer prices and as new jobs emerge that reemploy displaced labor. However, while the U.S. economy has maintained full employment, there is some evidence that automation has been outpacing the creation of new tasks and jobs in recent decades.[42] If so, generative AI could complicate this challenge by further narrowing the set of skills that are uniquely human, increasingly including those of higher-skilled white-collar professionals who largely escaped the effects of prior waves of automation.

Mixed effects among workers

Among workers, the greatest beneficiaries are likely to be those whose skills are *complementary* to AI, rather than replaced by it. Those who work in such complementary roles already tend to earn relatively more, and increased demand for their skills could add to inequality among workers. Consider a hypothetical customer service center that is made significantly more efficient by generative AI. Customer-facing workers may experience a direct productivity enhancement, but efficiency gains mean needing fewer of them to produce the same output, leaving them exposed to replacement. By contrast, a manager of this center who effectively develops and maintains systems that integrate the work of humans and AI would be performing work that is more complementary to AI.

At the same time, generative AI technologies **can level the playing field between lower-skilled and higher-skilled workers,** by "lending" expertise to those who lack it, without the need for formal training and investment. These skill-leveling effects might slightly offset inequality among workers. Recent studies on the impact of ChatGPT on customer service workers[43] and on college-educated professional performing writing tasks[44] found that the greatest productivity gains came from novice and low-skilled workers. Higher performers saw less benefit, perhaps because they were already delivering results closer to their peak potential, while lower-skilled workers were not only able to complete tasks faster but also perform tasks with greater complexity, "leveling up" in their responsibilities. Indeed, advanced technologies are already enabling nurse practitioners to take on more tasks usually performed only by primary care physicians.[45]

The net impact of AI of inequality – the inequality-increasing effects of benefitting those with complementary skills and the inequality-reducing effects of leveling the playing field among workers – will likely vary considerably by industry, and ultimately depend on how it is developed and deployed.

Effective policy management

Importantly, the speed at which adoption is taking place suggests governments, businesses and workers will need to act swiftly to reshape education and skills training and implement fiscal policies to smooth the transition for labor. Skills mismatches might be offset by investing in education and reskilling programs to ensure workers are keeping pace with the new skills demanded in an AI economy, while proper public safeguards will be needed in cases of job displacement. A greater concentration of wealth may call for further redistribution of economic income; reducing inequality, moreover, should help drive the demand boost that creates new jobs and new incomes, promoting a virtuous cycle that helps reduce inequality sustainably. If accompanied with the right policy approach, an economic boom from AI automation should be a "win-win" that ensures all income levels stand to benefit for the foreseeable future.

[41] According to BLS data, the share of the U.S. workforce represented by unions has fallen by more than half since the early 1980s, amounting to just 11.3% in 2022.

[42] Acemoglu and Restrepo find that automation corresponded to greater displacement effects and weaker reinstatement effects, or the acceleration of automation compared to the creation of new tasks, over the last three decades than the preceding decades. See K. Daron Acemoglu and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," Journal of Economic Perspectives 33, no. 2, May 2019.

[43] Brynjolfsson, Li, and Raymond, "Generative AI at Work."

[44] Noy and Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence."

[45] Raymond et al., "Nurse Practitioners' Involvement and Experience with AI-Based Health Technologies: A Systematic Review."

With considerable promise comes considerable risk

While the potential economic fruits of AI are bountiful, the technology brings with it several sociological and ethical concerns that we will need to navigate and confront in the coming years. Challenges include both those intrinsic to the underlying technology and its economic implications, as well as the potential for manipulation and misuse of powerful AI technologies by bad actors:

- **Social instability** is one potential consequence of rapid development and deployment of AI technologies that could manifest in a few ways:
  - **AI-generated disinformation** like falsified photos and videos intended to deceive, otherwise known as "deepfakes," might soon flood the internet. Such a development could make it even harder for members of the public to discern and agree on facts, ultimately amplifying ideological and affective political polarization.
  - **The concentration of AI ownership** among a few large corporations or countries could likewise concentrate power among them. At what point do AI corporations become "too big" for the greater good? Competition over such power could also lead to an unmitigated "arms race" between competing AI superpowers that might have second-order consequences, including for the careful assurance of safety of more powerful AI technologies.
  - **Economic hardship** from transitional unemployment and increased inequality could encourage further political extremism.
- **AI bias** is one key ethical concern, in that AI can perpetuate and amplify existing biases present in the data on which it is trained. As one old computer science adage says, *garbage in, garbage out*. The incidence of racial discrimination of facial recognition technology has been studied extensively and images produced by generative AI evidently amplify existing stereotypes.[46]
- **Data privacy** is also a major concern. AI systems often require vast amounts of personal information to function effectively, raising concerns about the collection and storage of sensitive data without proper consent or security measures. AI technologies can also inadvertently facilitate cross-border data transfers, resulting in potential violations of international data privacy laws. Few regulations exist so far to ensure responsible use of both national and cross-border data sharing.

However, AI is perhaps also uniquely positioned to address some of the same societal challenges it could potentially worsen. Properly designed and trained, AI may prove more objective and less biased than human counterparts. We can imagine generative AI-powered tools that present information in ways that help ideologically opposed individuals understand and relate to one another, and perhaps even serve as a real-time mediator or fact-checker for online discourse, encouraging objectivity and even social stability.

> Many fears around AI focus on the potential for a "doomsday" scenario that puts all of humanity at risk. As AI systems become more autonomous and capable of making decisions, many experts forecast a significant risk that humans fail to contain a powerful AI system that is not aligned with our values. This sort of risk, however, is a longer-term consideration that depends on considerable further advancement in AI technology that could be many decades away and is thus beyond the scope of this paper.

[46] Alexandra Sasha Luccioni et al., "Stable Bias: Analyzing Societal Representations in Diffusion Models," ArXiv, March 20, 2023.

Key asset class implications

The broad implications of AI for global economies also leave much for investors to consider.

Higher equity prices

For equities, if AI delivers on its promises, the implications should be straightforwardly positive. Our work on *Long-Term Capital Market Assumptions* suggests that an acceleration in potential GDP growth, all else the same, is likely to drive an acceleration in earnings by a similar degree. Moreover, a greater share of national income flowing to capital – the owners of AI technologies – could give an added boost to equity returns.

While these impacts could take several years to materialize, markets – and equities in particular – are likely to price in AI optimism long before then. Indeed, strong global equity performance in 2023 so far, particularly in U.S. large caps, has been largely influenced by excitement around generative AI technologies. Such performance has inspired comparisons to the early 2000s dot-com "bubble," where enthusiasm about the internet propelled blind exuberance that drove stock prices significantly above intrinsic values.

In our view, performance so far does not nearly resemble a "bubble." Price multiples are not yet significantly stretched as enthusiasm has also been accompanied by strong upward revisions to earnings outlooks for stocks with the most AI exposure, and that multiple expansion has been relatively broad based, in contrast to the narrow leadership of the dot-com bubble (**Exhibit 13**).

The future potential, however, for AI to drive something more like a bubble presents some upside risk to equities. Historically, bubbles have commonly involved some new technology with no direct historical comparisons, making the impact hard to precisely quantify. When a consensus emerges that this technology is the "next big thing," investors' imaginations tend to run wild. AI certainly seems capable of satisfying this criterion. On the other hand, most bubbles have also developed in periods of highly available credit, whereas today's environment is one of restrictive monetary policy and tightening bank lending standards, which should keep investors' optimism at least somewhat in check.

Exhibit 13: Rise in valuations has so far been broad based, unlike in the early 2000s dot-com bubble

12-month foward price-to-earnings (P/E) mutiple, S&P 500

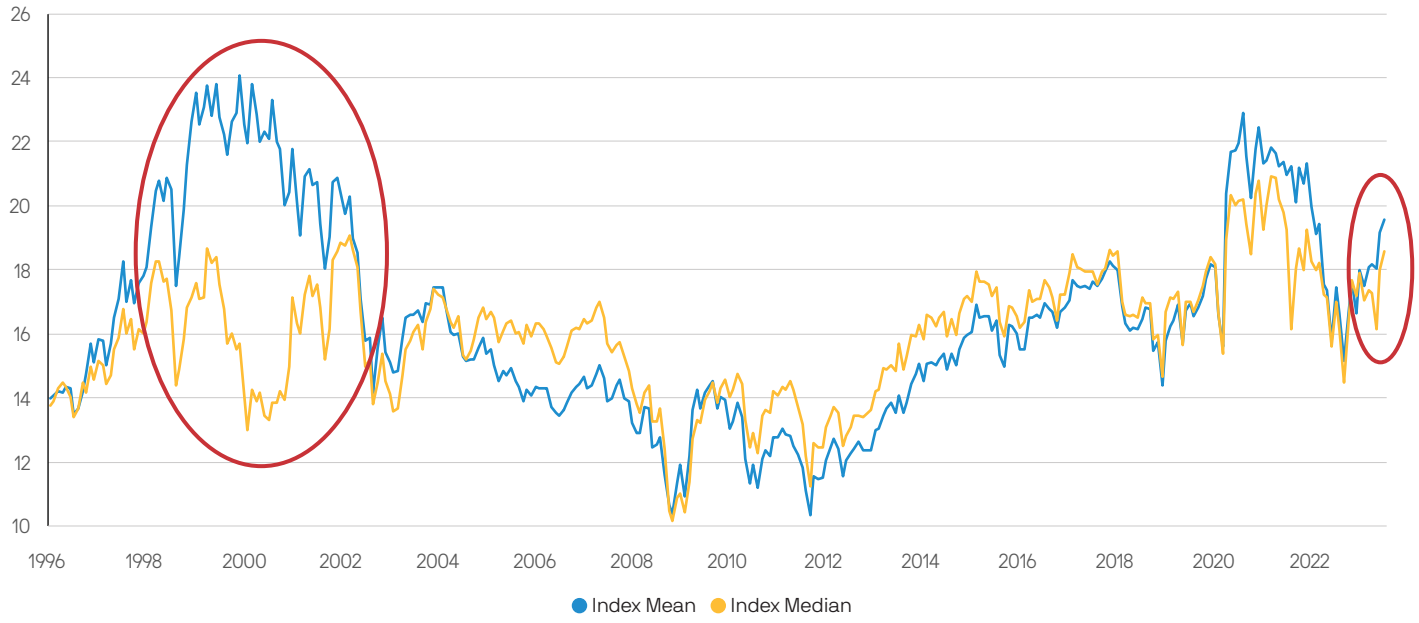

Source: FactSet, J.P. Morgan Asset Management. Data as of August 2023.

AI excitement has already led to considerable gains for technology companies in the S&P 500. Among individual equities, Big Tech – now including chipmaker Nvidia – has obvious exposure to AI and especially generative AI, but as in past technology cycles, tomorrow's winners from AI may include relatively new players that are not in vogue today. For our U.S. Equity Group's thinking on navigating AI investment opportunities, see "Artificial intelligence: Powering the next wave of technological innovation."

Potentially higher yields

For government bonds, an AI productivity shock will likely contribute to modestly higher yields across developed markets. Faster growth in productivity and thus real economic activity is likely to push cycle-neutral real yields higher by roughly the same degree, as has been the case over the long term (especially in the United States since the middle of the 20th century). However, greater income inequality and downward pressure on wages could reduce inflationary pressures and thus breakeven inflation rates, partially offsetting upward pressure on nominal yields. One key point of reference is the late 1990s, the most recent period of strong productivity growth; over this period, nominal yields did rise modestly – with the U.S. 10-year up by over 2% between late 1998 and early 2000, a period that also roughly corresponded to the strongest NASDAQ appreciation.

A similar dynamic may play out this time, but we are mindful of the possibility that greater inequality is met with increased pressure to fund policies like universal basic income (UBI) that are not met with greater tax revenues. Such policies would result in greater government debt issuance, which tends to result in higher long-term yields.

One consequence for monetary policy is that structurally higher yields could reduce the need for unconventional policies like quantitative easing, which have placed downward pressure on longer-dated yields.

Conclusions

When we discuss AI with our clients, many of them are more concerned than excited. Fear can be a good thing, but it's easy to simply fear what we don't understand, so part of our goal with this publication is to help identify what we do know and where the real challenges are with this emerging technology. AI certainly does present many challenges, especially for labor, but an era of mass unemployment seems highly implausible. After all, we wouldn't bet against our ability as humans to always find new ways to challenge ourselves.

Managed properly, we do think AI has the potential to make us all more productive, lower the real costs of many goods and services, reignite economic growth and offset aging demographics. Beyond economic growth, AI could also help solve some of our hardest societal challenges, such as in medicine and energy sustainability, and even accelerate the pace of innovation itself. With all of this potential, AI may prove to be the major transformative technology of the 21st century, a rare occurrence that has historically preceded significant industry and societal change. Generative AI, and its rapid accession to the mainstream, may be the tipping point.

For investors, all of these outcomes provide significant multi-asset investment opportunities, but with generative AI still in its early innings, we would also emphasize the importance of humility and discretion. It is in such uncertain environments, however, where we believe active management ultimately excels at identifying the winning companies of tomorrow, and the paths various asset classes may take along the way. As the importance of and attention to AI continues to grow, we are working hard at J.P. Morgan Asset Management to develop and share our insights on the many questions surrounding this technological wave. As always, we welcome your feedback.

An audio accompaniment of this paper is featured on the "Insights Now" podcast series, available on Spotify, Apple and Google Podcasts. The episode is entitled "The economic implications of generative AI."

Appendix: Sizing AI productivity gains

Our methodology

In line with a number of peer estimates, we took a task-based approach to approximating the impact of AI automation based on the task dataset from O*NET. This dataset provides many details on the task composition of different jobs, including task difficulty and relative importance, based on large-scale surveys of U.S. workers. We then multiply the O*NET task content for each job by U.S. Bureau of Labor Statistics (BLS) data on the share of this job in the economy, providing an estimate of the national aggregate share of each task, and from there, the total potential impact of automation on economic activity.

We identify 33 out of the 41 catalogued occupational tasks as likely to be exposed to automation from AI technologies in some form over the next decade, leaving eight tasks immune (mainly due to physical-world requirements, as we exclude robotics from our assumptions). We then scored those tasks based on (1) the share of the task we expect to be potentially automated (i.e., AI automates up to 80% of the tasks bucketed under "interpreting the meaning of information for others"), (2) the skill intensity of the automation (i.e., interpreting Python code would have a higher difficulty than interpreting customer feedback) and (3) the impact of such automation scaled across the entire economy (i.e., some tasks, such as "documenting information," are more prevalent in the economy than others, such as "repairing mechanical equipment").

Key assumptions

Our framework involves a few key assumptions:

1. AI automation is widespread across tasks and industries with most impacts taking hold in the next 10 years.
2. AI can automate up to a difficulty level of 3 on the 0-7 O*NET difficulty scale (see **Exhibit 15** for examples); varying this figure produces quite a wide range of alternative outcomes.
3. Displaced workers are re-employed in jobs where they are equally as productive as they were before, translating to a "composition effect" of zero, as we reviewed in the prior section "Sizing the productivity impact."

Our assumptions on automation exposure include both generative AI and "traditional" AI technologies but exclude robotics. Our analysis only accounts for the potential effects of AI automation; further sources of productivity gains could skew these results higher.

Analytical limitations

The O*NET dataset on task content presents some inherent limitations and challenges. Most notably, reported task difficulty corresponds to difficulty for humans rather than for AI, requiring our own estimates on the latter, and the reported figures for relative "importance" reflect a subjective measure of the task's respective importance to the overall job rather than time spent. Presumably, these two metrics are related, but there are certainly cases where the tasks that take up the most time in a worker's day aren't the most value-added (i.e., so-called "grunt work"). Our sample captures 94% of the labor force as measured by BLS, as some roles in the BLS dataset did not exist in the O*NET database and were therefore excluded from our analysis.

We also note limitations in the extrapolation of automation estimates of specific O*NET tasks across all professions, since the same "task" can vary between occupations. For example, our analysis found that fashion models, a profession that should arguably be immune to automation, had over 30% exposure to automation (in theory, image generating AI could probably automate some fashion models, if societal norms permitted it). This example underscores the difficulty in translating the "task difficulty" measure provided by O*NET to the tasks AI could realistically automate.

Sensitivities

The above limitations result in estimation error in our projections, and indeed, minor adjustments in the parameters of our analysis can lead to varied results. If we assumed AI technologies could only automate up to a level 2 on the seven-point scale employed by O*NET and took a more conservative cut of tasks subject to automation, with 15 tasks considered immune, estimated productivity gains would then be cut to 0.4%-0.9% per year. The most aggressive assumptions, assuming AI automates all tasks up to a difficulty of 5, would see

productivity gains of up to 7.7% per year. This wide range of outcomes underscores the high uncertainty behind any projections of the magnitude of productivity gains. Regardless, we do believe such exercises are worthwhile as we imagine the extent of potential automation across the economy, and the channels through which productivity gains could materialize.

Comparisons to other estimates

Many peer estimates on the potential productivity gain are within the range of what we find plausible (**Exhibit 14**). For instance, McKinsey came to similarly high estimates of 0.2%-3.3% per year through 2040 including a broad set of AI technologies. This kind of productivity gain would be a significant feat, translating to $2-$4tr annually added to the world economy from generative AI use cases alone. Brynjolfsson et al. found that generative AI alone could raise productivity by 18% over 10 years, or 1.8% per year, and acknowledge further potential upside from the acceleration of innovation.[47]

Potential limitations to productivity upside

While the potential boost to productivity is significant, several factors may constrain its effects. The composition effect hinges on the ability for workers to pivot to new tasks and occupations where they are similarly productive. An AI boom that does not coincide with the creation of new tasks and jobs would hinder overall productivity benefits. In addition, AI technologies are still very expensive to develop and integrate into business use cases; if AI is not met with monetizable demand drivers, its uptake could be more limited. While regulation is necessary to ensure a sustainable rollout of AI, overly restrictive or ineffective policies could impede the development and deployment of AI technologies. Moreover, a broader resistance to change and technological limitations may further constrain productivity gains.

Exhibit 14: Others' estimates on the productivity gain from AI

Source	Estimated annual AI productivity gain	Region	Period
Alderucci et al.	6.8%*	United States	next 5 years
Brynjolfsson et al.	1.8%**	United States	next 10 years
Goldman Sachs	1.5% (range: 0.3% to 2.9%)	United States	next 10 years
McKinsey	0.2% to 3.3%***	World	2023 to 2040
J.P. Morgan Asset Management	1.4% to 2.7%	Developed markets	next 10 years

Source: J.P. Morgan Asset Management; estimates compiled as of August 31, 2023.

* Alderucci et al. (2022) found that on a firm-level, manufacturing firms with AI-related inventions experienced a 7% increase in TFP, along with an 8.3% increase in total revenue per employee and an 8.9% increase in value-added per employee. Across the economy, the impact of AI-related invention is associated with a 6.8% increase in revenue per employee in the following 5 years.

** Brynjolfsson et al. (2023) estimate that generative AI will raise productivity by an added 18% over ten years (or 1.8% per year) above the current Congressional Budget Office (CBO) projection of 1.5% productivity growth.

*** McKinsey (2023) estimates that generative AI alone could enable labor productivity growth of 0.1% to 0.6% annually through 2040. Combined with a broad set of other AI technologies, work automation could add 0.2% to 3.3% points annually to productivity growth.

[47] Brookings. "Machines of Mind: The Case for an AI-Powered Productivity Boom," May 10, 2023. https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom/.

Exhibit 15: A task-based approach to sizing the productivity impact of AI technologies

Task difficulty (1-7 scale)	2	4	6
Getting information	Follow a standard blueprint	Review a budget	Study international tax laws
Monitoring processes, materials or surroundings	Check to see if baking bread is done	Test electrical circuits	Check the status of a patient in critical medical care
Identifying objects, actions and events	Test an automobile transmission	Judge the suitability of food products for an event	Determine the reaction of a virus to a new drug
Estimating the quantifiable characteristics of products, events or information	Estimate the size of household furniture to be shipped	Estimate the time required to evacuate a city in the event of a major disaster	Estimate the amount of natural resources that lie beneath the world's oceans
Processing information	Calculate the costs for shipping packages	Calculate the adjustments for insurance claims	Compile data for a complex scientific report
Evaluating information to determine compliance with standards	Review forms for completeness	Evaluate a complicated insurance claim for compliance with policy terms	Make a ruling in court on a complicated motion
Analyzing data or information	Skim a short article to gather the main point	Determine the interest cost to finance a new building	Analyze the cost of medical care services for all hospitals in the country
Updating and using relevant knowledge	Keep up with price changes in a small retail store	Keep current on changes in maintenance procedures for repairing sports cars	Learn information related to a complex and rapidly changing technology
Scheduling work and activities	Make appointments for patients using a predetermined schedule	Prepare the work schedule for salesclerks in a large retail store	Schedule a complex conference program with multiple, parallel sessions
Organizing, planning and prioritizing work	Organize a work schedule that is repetitive and easy to plan	Plan and adjust a personal to-do list according to changing demands	Prioritize and plan multiple tasks several months ahead
Documenting/recording information	Record the weight of a patient during a routine health exam	Document the results of a crime scene investigation	Maintain information about the use of satellites for industry communications
Interpreting the meaning of information for others	Interpret a blood pressure reading	Interpret how foreign tax law applies to U.S. exports	Interpret a complex experiment in physics for general audiences
Performing administrative activities	Complete routine paperwork	Complete tax forms for a small business	Serve as the benefits director for a large computer sales organization

Source: O*NET, J.P. Morgan Asset Management. Table is for illustrative purposes.

Authors

Michael Albrecht

Global Strategist, Executive Director,
Multi-Asset Solutions

Stephanie Aliaga

Research Analyst, Senior Associate,
Global Market Insights

Acknowledgements

We would like to express our gratitude to several individuals across J.P. Morgan Asset Management who generously contributed their time and expertise to the development of this paper – through their insights and constructive feedback, help with research and broad support for our efforts.

For his sponsorship, we would first like to thank **Dr. David Kelly,** Chief Global Strategist and Head of the Global Market Insights Strategy Team, who was the first to support our foray into the subject as part of our Long-Term Capital Market Assumptions, and for providing feedback on multiple iterations of this publication, especially on economic implications.

We are likewise grateful to **Jed Laskowitz**, Chief Investment Officer and Global Head of Asset Management Solutions, and **John Bilton**, Head of Global Multi-Asset Strategy, who both supported our efforts, provided constructive feedback and helpfully suggested others whose inputs proved invaluable.

Special thanks also go to **Michael Cembalest**, Chairman of Market and Investment Strategy, who offered particularly extensive and detailed feedback on our paper – given his remarkable depth of knowledge in AI, among many other fields. Michael also brought our attention to additional relevant research that we have since incorporated into ours.

We are also grateful for feedback from **Jeffrey Geller**, a Chief Investment Officer in Multi-Asset Solutions; **Jared Gross**, Head of Institutional Portfolio Strategy; **Evan Grace**, Head of Multi-Asset Portfolio Management for the International Private Bank; and **Anthony Werley**, Chief Investment Officer of the Outsourced CIO/Endowments and Foundations Team. We would like to thank several members of the U.S. Equity Group, particularly for their insights on AI advances beyond generative AI, corporate applications of AI and equity investment opportunities within AI: **Robert Bowman**, a semis and technology hardware analyst; **Timothy Parton**, a portfolio manager, and **James Connors** and **Doug Stewart**, both investment specialists. We are also grateful to **Dillon Edwards**, an AI strategist in our Intelligent Digital Solutions Team, for his partnership and feedback on technology specific issues, and to our Data Science team for supplying bespoke data on mentions of AI in earnings calls. Finally, from Market Insights, **Tai Hui**, Asia Chief Market Strategist, and **David Lebovitz**, Global Market Strategist, and from Multi-Asset Solutions, **Sylvia Sheng** and **Thushka Maharaj**, both Global Strategists. All of these colleagues gave useful feedback in their respective areas of expertise.

Last though certainly not least, we would like to thank direct contributors to this paper, foremost from **Michael Di Pentima**, a portfolio management associate in Multi-Asset Solutions, who helped develop calculations to quantify the economic impact of AI automation. Additionally, he, as well as **Brandon Hall** and **Mary Park Durham**, research analysts in Market Insights, and **Ayesha Khalid**, a Global Strategist in Multi-Asset Solutions, helped us collect and organize various exhibits and data used in this publication. This publication, and all of our published content, could not have been possible without the incredible support and assistance from our marketing partners, namely **Maria Bernice Halaguena**, a marketing strategy associate, and **Haley Baron**, marketing strategy lead.

This paper is a joint effort between Market Insights and Multi-Asset Solutions with input and support from across the firm. We hope our insights help inform our clients' long-term investment planning and asset allocations.

J.P. Morgan Asset Management

277 Park Avenue | New York, NY 10172

The Market Insights program provides comprehensive data and commentary on global markets without reference to products. Designed as a tool to help clients understand the markets and support investment decision making, the program explores the implications of current economic data and changing market conditions.

For the purposes of MiFID II, the JPM Market Insights and Portfolio Insights programs are marketing communications and are not in scope for any MiFID II/MiFIR requirements specifically related to investment research. Furthermore, the J.P. Morgan Asset Management Market Insights and Portfolio Insights programs, as non-independent research, have not been prepared in accordance with legal requirements designed to promote the independence of investment research, nor are they subject to any prohibition on dealing ahead of the dissemination of investment research.

This document is a general communication being provided for informational purposes only. It is educational in nature and not designed to be taken as advice or a recommendation for any specific investment product, strategy, plan feature or other purpose in any jurisdiction, nor is it a commitment from J.P. Morgan Asset Management or any of its subsidiaries to participate in any of the transactions mentioned herein. Any examples used are generic, hypothetical and for illustration purposes only. This material does not contain sufficient information to support an investment decision and it should not be relied upon by you in evaluating the merits of investing in any securities or products. In addition, users should make an independent assessment of the legal, regulatory, tax, credit and accounting implications and determine, together with their own professional advisers, if any investment mentioned herein is believed to be suitable to their personal goals. Investors should ensure that they obtain all available relevant information before making any investment. Any forecasts, figures, opinions or investment techniques and strategies set out are for information purposes only, based on certain assumptions and current market conditions and are subject to change without prior notice. All information presented herein is considered to be accurate at the time of production, but no warranty of accuracy is given and no liability in respect of any error or omission is accepted. It should be noted that investment involves risks, the value of investments and the income from them may fluctuate in accordance with market conditions and taxation agreements and investors may not get back the full amount invested. Both past performance and yields are not reliable indicators of current and future results.

J.P. Morgan Asset Management is the brand name for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at https://am.jpmorgan.com/global/privacy.

This communication is issued by the following entities:

In the United States, by J.P. Morgan Investment Management Inc. or J.P. Morgan Alternative Asset Management, Inc., both regulated by the Securities and Exchange Commission; in Latin America, for intended recipients' use only, by local J.P. Morgan entities, as the case may be.; in Canada, for institutional clients' use only, by JPMorgan Asset Management (Canada) Inc., which is a registered Portfolio Manager and Exempt Market Dealer in all Canadian provinces and territories except the Yukon and is also registered as an Investment Fund Manager in British Columbia, Ontario, Quebec and Newfoundland and Labrador. In the United Kingdom, by JPMorgan Asset Management (UK) Limited, which is authorized and regulated by the Financial Conduct Authority; in other European jurisdictions, by JPMorgan Asset Management (Europe) S.à r.l. In Asia Pacific ("APAC"), by the following issuing entities and in the respective jurisdictions in which they are primarily regulated: JPMorgan Asset Management (Asia Pacific) Limited, or JPMorgan Funds (Asia) Limited, or JPMorgan Asset Management Real Assets (Asia) Limited, each of which is regulated by the Securities and Futures Commission of Hong Kong; JPMorgan Asset Management (Singapore) Limited (Co. Reg. No. 197601586K), which this advertisement or publication has not been reviewed by the Monetary Authority of Singapore; JPMorgan Asset Management (Taiwan) Limited; JPMorgan Asset Management (Japan) Limited, which is a member of the Investment Trusts Association, Japan, the Japan Investment Advisers Association, Type II Financial Instruments Firms Association and the Japan Securities Dealers Association and is regulated by the Financial Services Agency (registration number "Kanto Local Finance Bureau (Financial Instruments Firm) No. 330"); in Australia, to wholesale clients only as defined in section 761A and 761G of the Corporations Act 2001 (Commonwealth), by JPMorgan Asset Management (Australia) Limited (ABN 55143832080) (AFSL 376919). For all other markets in APAC, to intended recipients only.

For U.S. only: If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

© 2023 JPMorgan Chase & Co. All rights reserved.

MI_MAS-AI bulletin

09nr231409193311

J.P.Morgan
ASSET MANAGEMENT