The Potentially Large Effects of Artificial Intelligence on Economic Growth (Briggs/Kodnani)

Table of Contents

26 March 2023 | 9:05PM EDT

FOLLOW

- The recent emergence of generative artificial intelligence (AI) raises the question whether we are on the brink of a rapid acceleration in task automation that will drive labor cost savings and raise productivity. Despite significant uncertainty around the potential of generative AI, its ability to generate content that is indistinguishable from human-created output and to break down communication barriers between humans and machines reflects a major advancement with potentially large macroeconomic effects.
- If generative AI delivers on its promised capabilities, the labor market could face significant disruption. Using data on occupational tasks in both the US and Europe, we find that roughly two-thirds of current jobs are exposed to some degree of AI automation, and that generative AI could substitute up to one-fourth of current work. Extrapolating our estimates globally suggests that generative AI could expose the equivalent of 300mn full-time jobs to automation.
- The good news is that worker displacement from automation has historically been offset by creation of new jobs, and the emergence of new occupations following technological innovations accounts for the vast majority of long-run employment growth. The combination of significant labor cost savings, new job creation, and higher productivity for non-displaced workers raises the possibility of a productivity boom that raises economic growth substantially, although the timing of such a boom is hard to predict.
- We estimate that generative AI could raise annual US labor productivity growth by just under 1½pp over a 10-year period following widespread adoption, although the boost to labor productivity growth could be much smaller or larger depending on the difficulty level of tasks AI will be able to perform and how many jobs are ultimately automated.
- The boost to global labor productivity could also be economically significant, and we estimate that AI could eventually increase annual global GDP by 7%. Although the impact of AI will ultimately depend on its capability and adoption timeline, this estimate highlights the enormous economic potential of generative AI if it delivers on its promise.

The Potentially Large Effects of Artificial Intelligence on Economic Growth

The recent emergence of generative artificial intelligence (AI) has raised questions around whether we are at the brink of a rapid acceleration in task automation that will significantly save on labor costs, raise labor productivity, and increase the pace of economic growth. In this *Global Economics Analyst*, we provide an overview of AI's potential macroeconomic impacts, and argue that if AI delivers on its promised capabilities, it has the potential to significantly disrupt labor markets and spur global productivity growth over the coming decades.

Generative AI, Explained

We first discuss the current state of AI development and its key capabilities. Exhibit 1 provides an overview of generative AI, in comparison to its predecessor machine learning methods, sometimes referred to as narrow or analytical AI. In our assessment, the generative AI technologies currently in focus, such as ChatGPT, DALL-E, and LaMDA, are distinguished by three main characteristics: 1) their generalized rather than specialized use cases, 2) their ability to generate novel, human-like output rather than merely describe or interpret existing information, and 3) their approachable interfaces that both understand and respond with natural language, images, audio, and video.

The first two advances are key to expanding the set of tasks that AI can perform, while the third is key for determining its adoption timeline. Just as the migration from command line programming (e.g., MS-DOS) to graphical user interfaces (e.g., Windows) enabled the development of programs (e.g., Office) that brought the power of the personal computer to the masses, the intuitive interfaces of the current generation of AI technologies could significantly increase their speed of adoption. For example, ChatGPT surpassed 1mn users in just 5 days, the fastest that any company has ever reached this benchmark.

Exhibit 1: An Overview of Generative AI

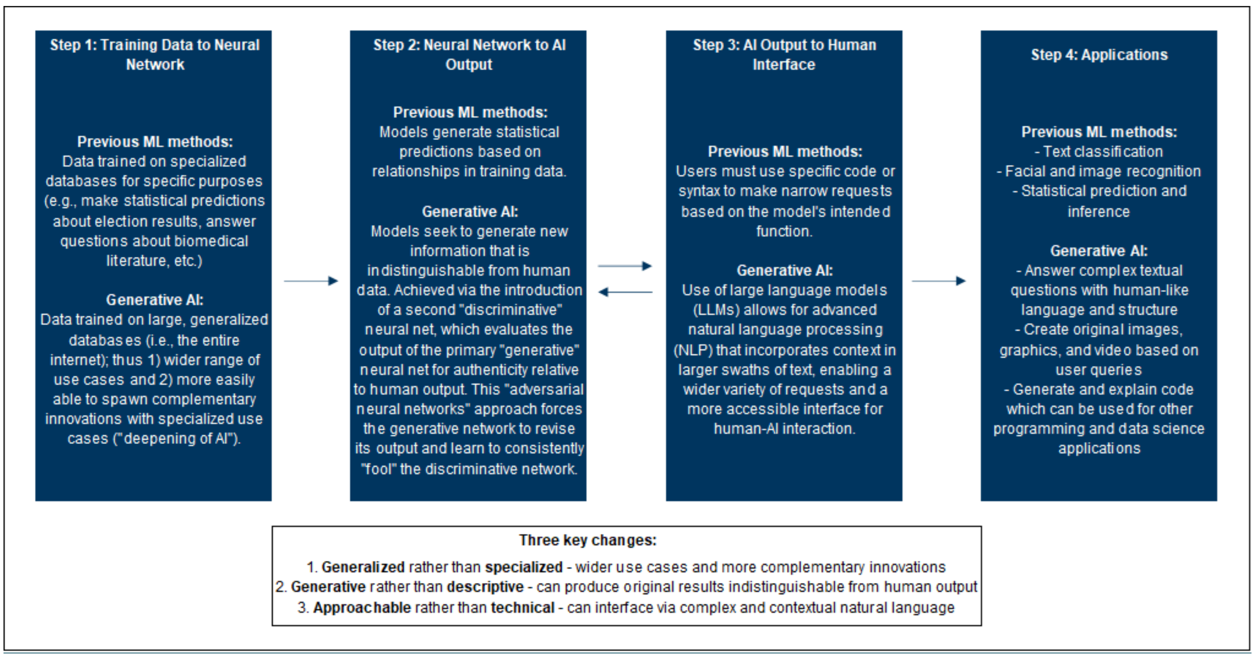

Source: Goldman Sachs Global Investment Research

Beyond these changes, exponential increases in the raw computing power available[1] has enabled rapid advances in the complexity of tasks AI can perform and the accuracy at which it can perform them. For example, the latest iteration of OpenAI's GPT model—GPT-4, released in March 2023, roughly one year after the GPT-3.5 model currently underlying ChatGPT finished training—scores 150 points higher on the SAT than its predecessor, is 40% more likely to produce accurate responses, and can now accept visual input (rather than just text).[2] As Exhibit 2 shows, the algorithms underlying generative AI had begun to surpass human benchmarks for tasks such as image classification and reading comprehension even before these recent advances.

Exhibit 2: AI Increasingly Outperforms Human Benchmarks

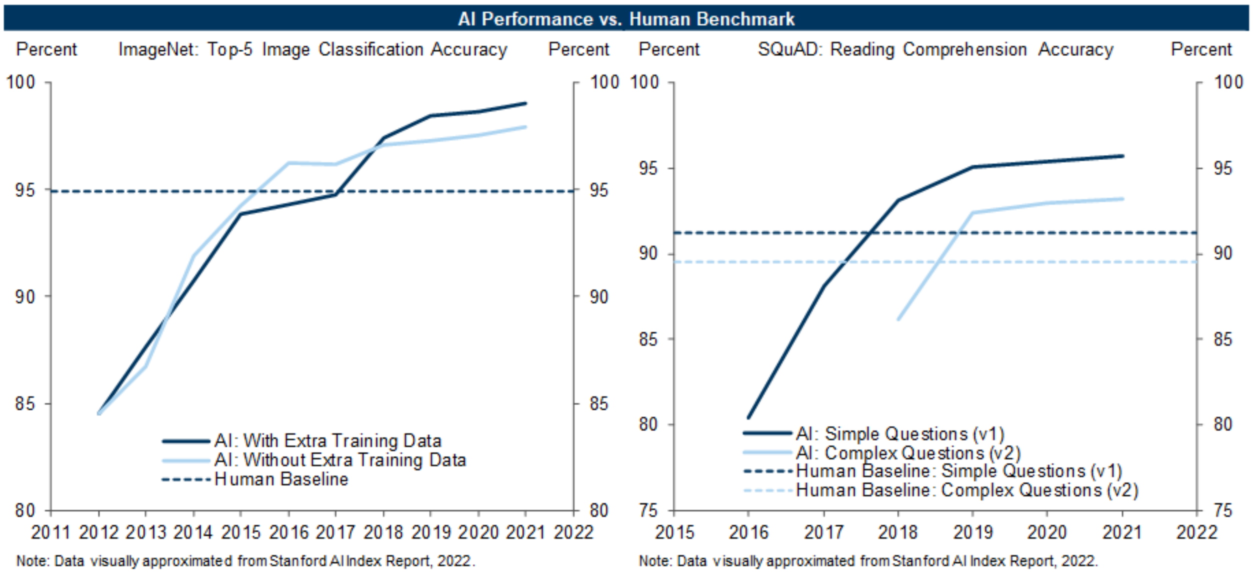

Source: Stanford Institute for Human-Centered Artificial Intelligence, Goldman Sachs Global Investment Research

As AI has become increasingly advanced and accessible, interest and investment have followed. Management teams of publicly-traded corporations increasingly cite AI in earnings calls—and at a rapidly increasing rate—and these indications of interest predict large increases in capital investment at the company level (Exhibit 3). As of 2021, US and global private investment in AI totaled $53bn and $94bn respectively—each up more than fivefold in real terms from five years prior[3]—and if investment continues to increase at the more modest pace that software investment grew at during the 1990s, US investment in AI alone could approach 1% of US GDP by 2030.

Exhibit 3: Management Teams Are Increasingly Focused on Opportunities from AI on Company Earning Calls, and More Mentions of AI Predict Higher Capex

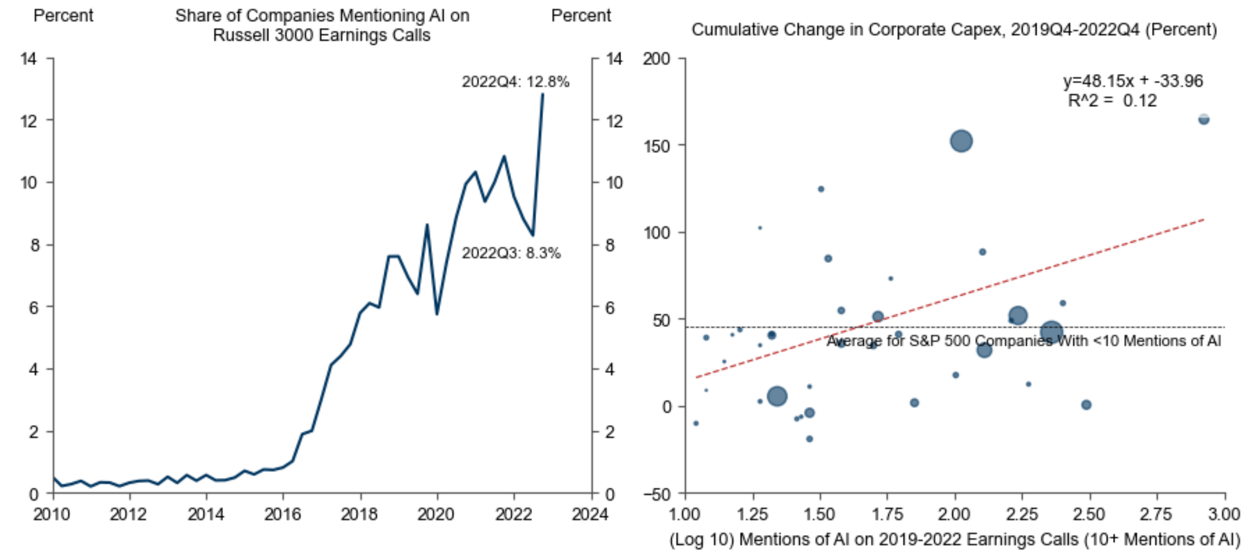

Source: GS Data Works, FactSet, Goldman Sachs Global Investment Research

While much uncertainty remains around both the capability and adoption timeline of generative AI, these developments suggest that AI is well-positioned to advance rapidly and grow in scale in the coming years.

The Future of Work: Substitute Sometimes, Complement Often

Generative AI's ability to 1) generate new content that is indistinguishable from human-created output and 2) break down communication barriers between humans and machines reflects a major advancement with potentially large macroeconomic effects.

To assess the size of these effects, we consider the likely impact generative AI will have on the labor market if it delivers on its promised capabilities. In particular, we use data from the O*NET database on the task content of over 900 occupations in the US (and later extend to over 2000 occupations in the European ESCO database) to estimate the share of total work exposed to labor-saving automation by AI by occupation and industry.

Based on our review of existing literature on the probable use cases of generative AI, we classify 13 work activities (out of 39 in the O*NET database) as exposed to AI automation, and in our base case assume that AI is capable of completing tasks up to a difficulty of 4 on the 7-point O*NET "level" scale (see Appendix for more details). We then take an importance- and complexity-weighted average of essential work tasks for each occupation and estimate the share of each occupation's total workload that AI has the potential to replace. We further assume that occupations for which a significant share of workers' time is spent outdoors or performing physical labor cannot be automated by AI.

In Exhibit 4, we report the occupation-level distribution of the share of tasks that AI can perform. We find that roughly two-thirds of US occupations are exposed to some degree of automation by AI, and that of those occupations which are exposed, most have a significant—but partial—share of their workload (25-50%) that can be replaced.

Exhibit 4: Two-Thirds of Current Occupations Could be Partially Automated by AI

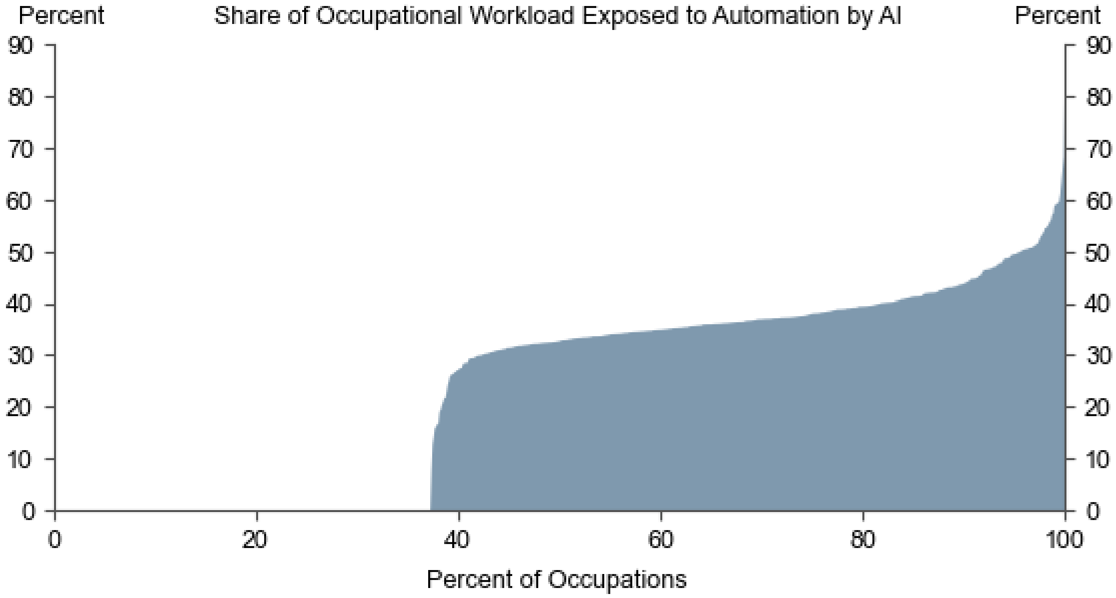

Source: Goldman Sachs Global Investment Research

Weighting our estimates by the employment share of each occupation in the US Occupational Employment and Wage Survey (OEWS) and aggregating to the industry level, we estimate that one-fourth of current work tasks could be automated by AI in the US (Exhibit 5, top panel), with particularly high exposures in administrative (46%) and legal (44%) professions and low exposures in physically-intensive professions such as construction (6%) and maintenance (4%). Matching our occupation-level estimates to the European ISCO occupation classification system and performing a similar exercise for the Euro Area using the Eurostat Labor Force Survey (LFS) database yields estimates of a similar magnitude, both in aggregate and across industries (Exhibit 5, bottom panel).

Exhibit 5: One-Fourth of Current Work Tasks Could Be Automated by AI in the US and Europe

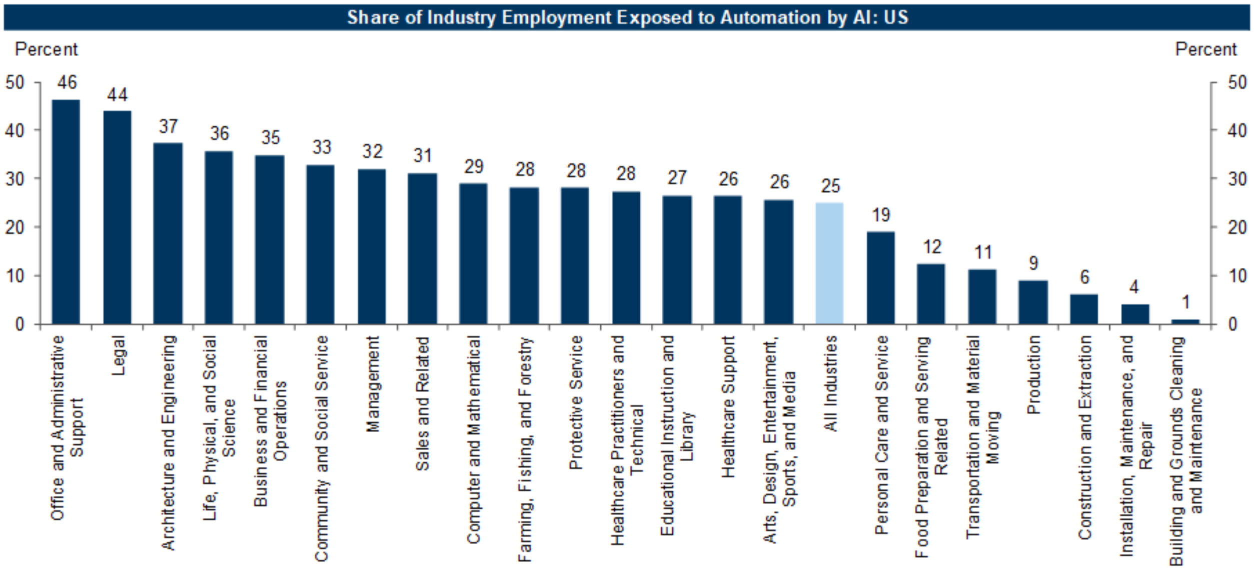

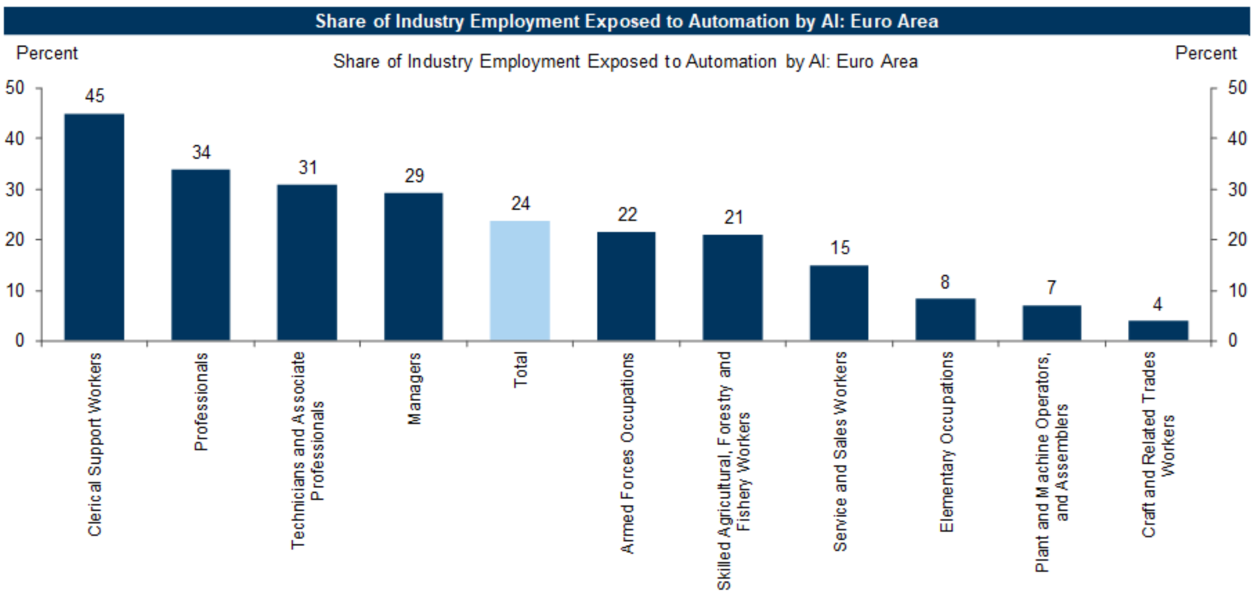

Source: Goldman Sachs Global Investment Research

We next extend our US and European estimates globally, adjusting for differences in industry composition across countries and further assuming that AI does not impact the agricultural sector in EM economies due to significant differences in the composition and production approaches in that industry between EM and DM economies.[4] Our estimates intuitively suggest that fewer jobs in EMs are exposed to automation than in DMs, but that 18% of work globally could be automated by AI on an employment-weighted basis (Exhibit 6).

Exhibit 6: Globally, 18% of Work Could be Automated by AI, with Larger Effects in DMs than EMs

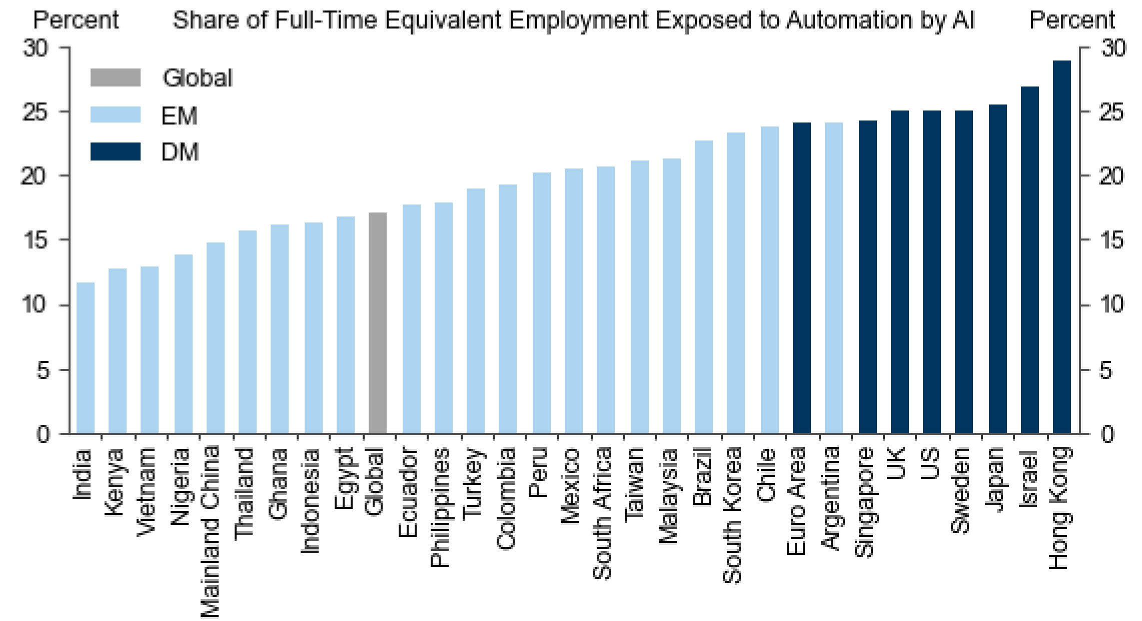

Source: Goldman Sachs Global Investment Research

Collectively, our estimates suggest that a large share of employment and work is at least partially exposed to automation by AI, raising the prospect of significant labor savings. To assess the robustness of our estimates, we compare our baseline US estimate to a wider range of scenarios, including those in which AI can perform more or less difficult tasks than we assume in our baseline, and in which we relax our assumption that AI cannot assist with jobs which are primarily outdoors or physical (i.e., a scenario in which AI is complementary with robotics and existing machinery). Our scenario analysis suggests that the ultimate share of work exposed to automation could range from 15-35% (Exhibit 7, left panel), a range which is consistent with—but on the conservative side of—existing estimates in the literature (Exhibit 7, right panel). Our relatively conservative baseline primarily reflects our narrower focus on the impact of generative AI, in contrast to other studies which sometimes consider a wider range of related technologies (including robotics) that increase the scope for automation.

Exhibit 7: Our Estimates Confirm that a Significant Share of Employment Is at Least Partially Exposed to Automation by AI, but Larger Effects Frequently Cited Include the Automation of Physical Tasks That Seem Less Likely in the Near-Term

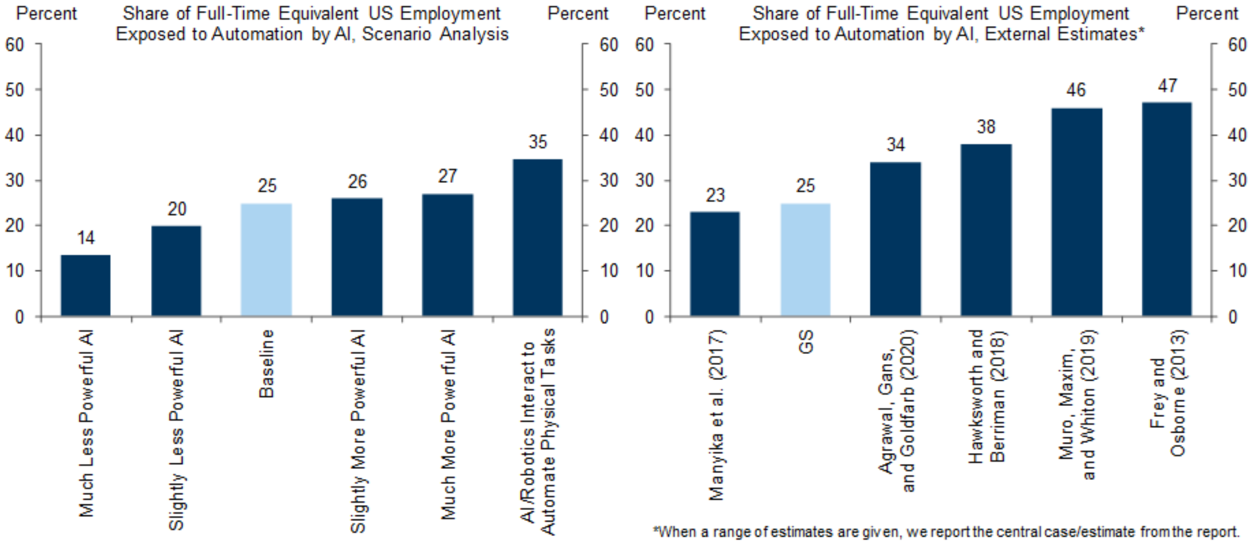

Source: Goldman Sachs Global Investment Research

Although the impact of AI on the labor market is likely to be significant, most jobs and industries are only partially exposed to automation and are thus more likely to be complemented rather than substituted by AI. In Exhibit 8, we assume for illustration that jobs for which at least 50% of importance- and complexity-weighted tasks are exposed to automation are likely to be substituted by AI, while jobs with an exposure of 10-49% are more likely to be complemented, and jobs with a 0-9% exposure are unlikely to be impacted. In our baseline, these assumptions would be consistent with 7% of current US employment being substituted by AI, 63% being complemented, and 30% being unaffected, though the ultimate effects will depend on how labor demand and occupational workloads evolve in response to partial labor savings in the majority of occupations.

Exhibit 8: Replacement in Legal and Administrative Fields, Little Effect in Manual and Outdoor Jobs, and Productivity-Enhancement Everywhere Else

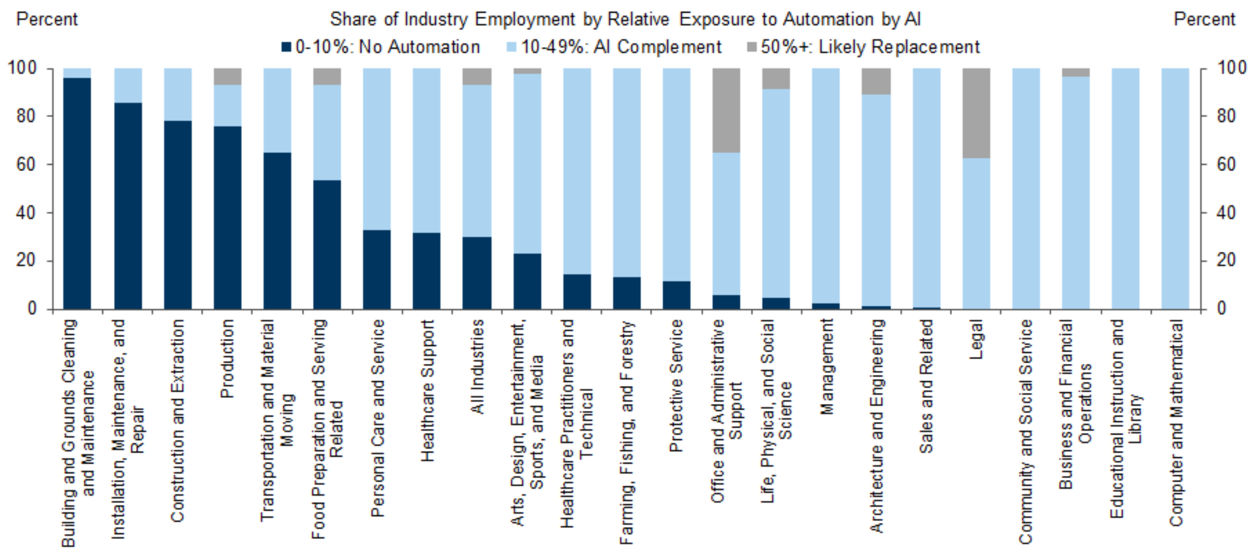

Source: Goldman Sachs Global Investment Research

Sizing the Boost to Productivity and Growth

The large share of employment exposed to automation from generative AI raises the potential for a boom in labor productivity that significantly increases global output. There are two main channels through which AI-driven automation could raise global GDP.

First, most workers are employed in occupations that are partially-exposed to AI automation and, following AI adoption, will likely apply at least some of their freed-up capacity toward productive activities that increase output.

Academic studies confirm that workers at early-adopting firms experience higher labor productivity growth following AI adoption, with estimates generally implying a 2-3pp/year boost (Exhibit 9). While differences in the capability of generative AI relative to earlier vintages make it hard to extrapolate these results forward, they clearly suggest that generative AI can drive an economically significant increase in productivity.

Exhibit 9: Academic Studies Generally Find That AI Adoption Increases Within-Firm Annual Worker Productivity Growth by 2-3pp

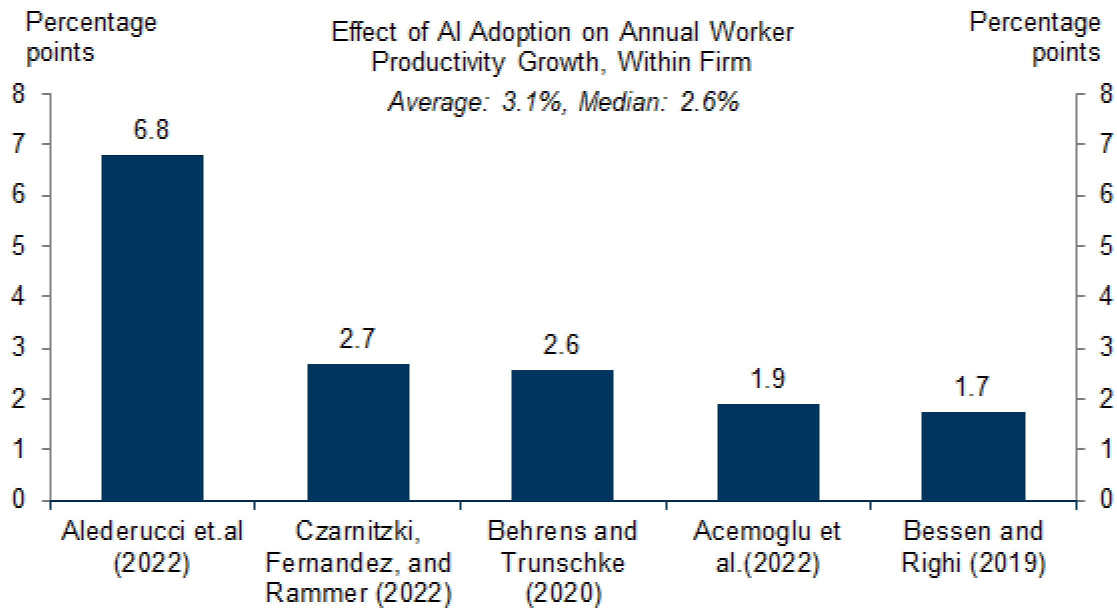

Source: Goldman Sachs Global Investment Research

Second, we anticipate that many workers that are displaced by AI automation will eventually become reemployed—and therefore boost total output—in new occupations that emerge either directly from AI adoption or in response to the higher level of aggregate and labor demand generated by the productivity boost from non-displaced workers. Both of these channels have plenty of historical precedent. For example, information technology innovations introduced new occupations like webpage designers, software developers, and digital marketing professionals, but also increased aggregate income and indirectly drove demand for service sector workers in industries like health care, education, and food services.

To demonstrate how technological innovation that initially displaces workers drives employment growth over a long horizon, in Exhibit 10 we show results from a recent study by economist David Autor and coauthors.[5] Using Census data, they find that 60% of workers today are employed in occupations that did not exist in 1940, implying that over 85% of employment growth over the last 80 years is explained by the technology-driven creation of new positions.

Exhibit 10: Technological Innovation Leads to the Creation of New Occupations That Account for the Bulk of Employment Growth

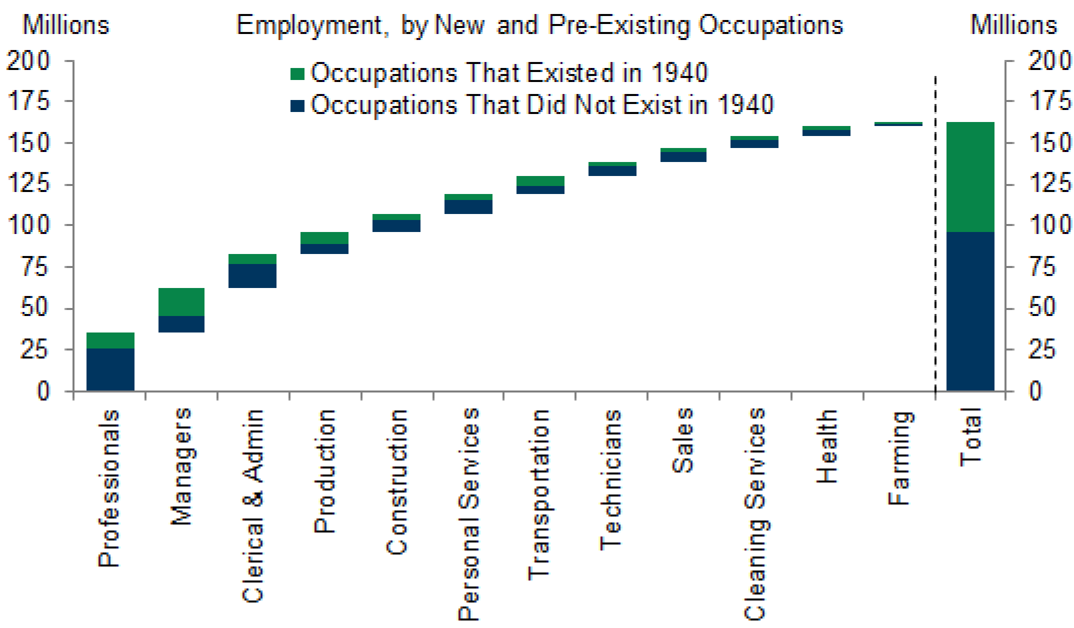

Source: Autor et al. (2022), Goldman Sachs Global Investment Research

Exhibit 11 leverages a different academic study by economists Daren Acemoglu and Pascual Restrepo that decomposes changes in labor demand into contributions from productivity growth and technology-driven worker displacement and reemployment (among other factors) to show how the drivers of labor demand change over time.[6]

Technological change displaced workers and created new employment opportunities at roughly the same rate for the first half of the post-war period, but has displaced workers at a faster pace than it has created new opportunities since the 1980s. These results suggest that the direct effects of generative AI on labor demand could be negative in the near-term if AI affects the labor market in a manner similar to earlier advances in information technology, although the effects on labor productivity growth would still be positive.

Exhibit 11: Historically, Worker Displacement from Automation Was Roughly Offset by Creation of New Roles/Tasks Prior to 1980, but Displacement Has Created a Net Drag on Labor Demand More Recently

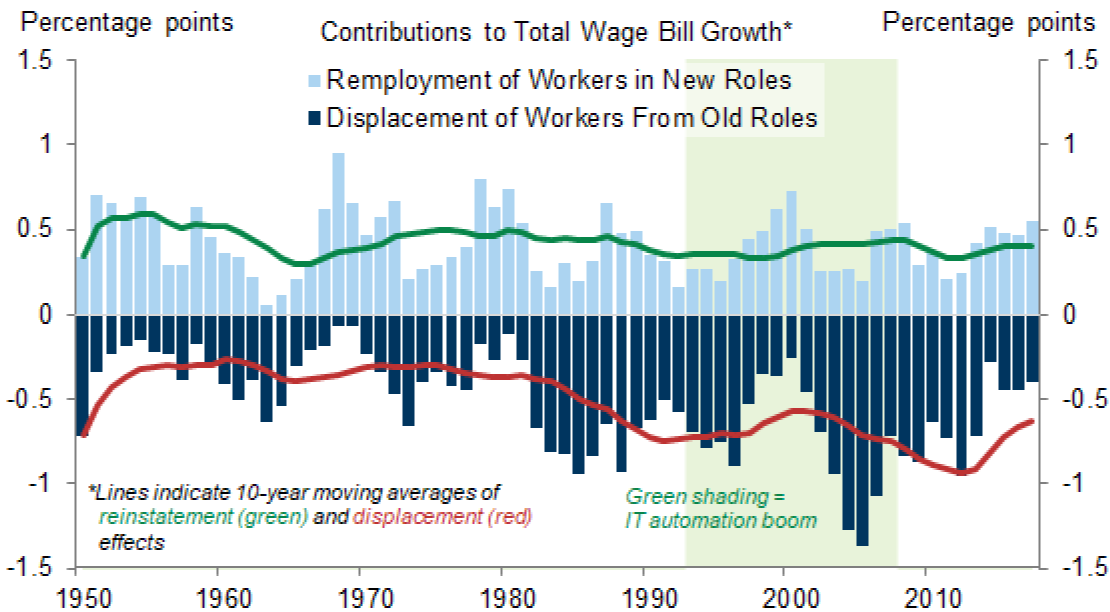

Source: Acemoglu and Restropo (2019), Goldman Sachs Global Investment Research

The combination of significant labor cost savings, new job creation, and a productivity boost for non-displaced workers raises the possibility of a labor productivity boom like those that followed the emergence of earlier general-purpose technologies like the electric motor and personal computer. These past experiences offer two key lessons.

First, the timing of a labor productivity boom is hard to predict, but in both cases started about 20 years after the technological breakthrough, at a point when roughly half of US businesses had adopted the technology (Exhibit 12, left panel). Second, in both of these instances, labor productivity growth rose by around 1.5pp/year in the 10 years after the productivity boom started, suggesting that the labor productivity gains can be quite substantial (Exhibit 12, right panel).

Exhibit 12: Previous Milestone Technologies Have Led to Labor Productivity Booms, but the Timing Is Hard to Predict

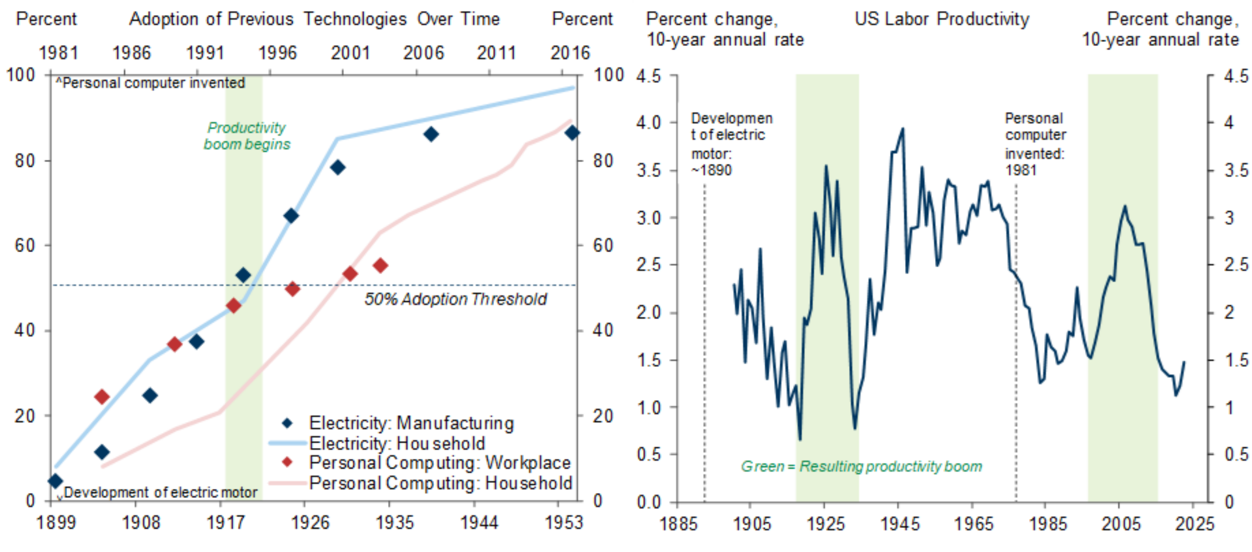

Source: US Bureau of Labor Statistics, Census Bureau, Our World in Data, Woolf (1987), Haver Analytics, Goldman Sachs Global Investment Research

To estimate the boost to labor productivity in the US from widespread adoption of generative AI, in Exhibit 13 we sum up the implied labor productivity effects from direct labor cost savings, a productivity boost for non-displaced workers, and a composition effect from the reemployment of displaced workers in new positions.

Our baseline analysis incorporates our key findings from above, including that about 7% of workers are fully displaced but that most are able to find new employment in only slightly less productive positions, that partially exposed workers experience a boost in productivity consistent with existing estimates (Exhibit 9), and that effects are realized over a 10-year period that starts around the time when roughly half of businesses have adopted generative AI. Under these assumptions we estimate that widespread adoption of generative AI could raise overall labor productivity growth by around 1.5pp/year (vs. a recent 1.5% average growth pace), roughly the same-sized boost that followed the emergence of prior transformative technologies like the electric motor and personal computer.

This estimated boost to labor productivity growth is both admittedly quite large and highly uncertain. Exhibit 13 therefore also considers other plausible scenarios and shows that the boost to US productivity growth could easily range from 0.3-3.0pp depending on the difficulty level of tasks generative AI can perform, how many jobs are ultimately automated, and the speed of adoption:

- First, we vary the O*NET difficulty level of the tasks that AI is capable of completing. In a much less powerful AI scenario where, for example, generative AI is only ultimately able to "skim a short article to gather the main point" (difficulty score 2) rather than "determine the interest cost to finance a new building" (difficulty score 4), the implied labor productivity growth boost would fall to 0.3pp/year. If AI is instead more powerful and is able to, again for example, "analyze the cost of medical care services for all US hospitals" (difficulty score 6), the implied labor productivity growth boost would rise to 2.9pp/year.

- Second, we vary the amount of labor that is fully displaced by generative AI. Assuming no labor displacement implies only a moderately smaller productivity growth boost of 1.2pp/year because non-displaced workers would still experience significant productivity gains, while assuming that a much larger share of workers are displaced would raise the boost to productivity growth to 2.4pp/year.
- Third, we vary the timeline of adoption. The productivity growth boost would only be roughly half as large if the gains are realized over a 20-year period and one-third as large if realized over a 30-year period.

Our key takeaway from these analyses is that the ultimate boost to labor productivity is uncertain, but in most scenarios would remain economically significant.

Exhibit 13: We Estimate That Generative AI Could Boost Aggregate Labor Productivity Growth by 1.5pp in the US, Although the Size of the Boost Will Depend on AI's Capability and Adoption Timeline

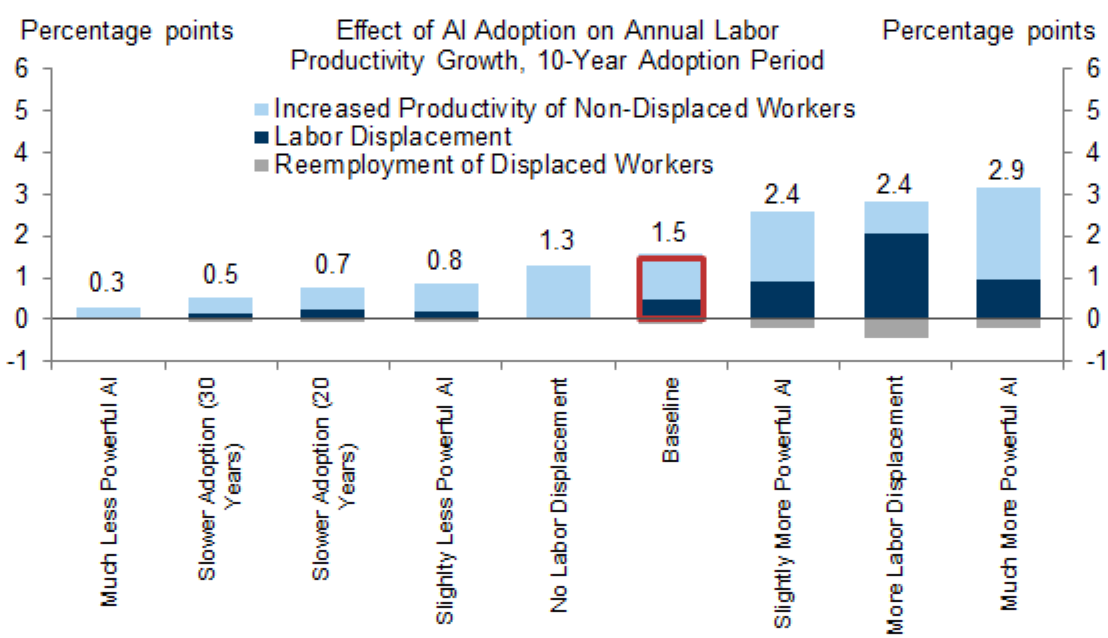

Source: Goldman Sachs Global Investment Research

In Exhibit 14 we extrapolate our analysis for the US to other countries under the assumption that differences in the industry-composition of labor can account for most of the differences in labor productivity growth. Our estimates imply that AI adoption could boost global annual productivity growth for countries in our coverage by 1.4pp (FX-weighted average) over a 10-year period, although we would likely expect a more delayed impact in EM economies.

Exhibit 14: Productivity Growth Boosts Could Be Sizable in Other Countries As Well; We Estimate Widespread AI Adoption Could Boost Global Annual Productivity Growth by 1.4pp Over a 10-Year Period

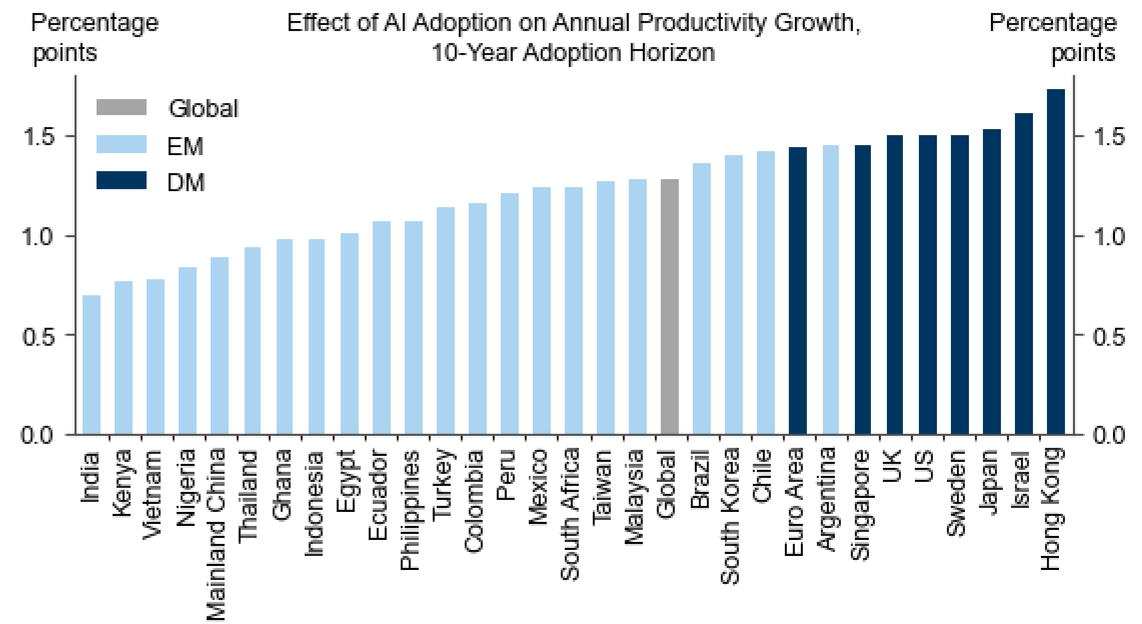

Source: Goldman Sachs Global Investment Research

Applying our estimated global labor productivity boost to countries in our coverage implies that widespread AI adoption could eventually drive a 7% or almost $7tn increase in annual global GDP over a 10-year period. Although the size of AI's impact will ultimately depend on its capability and adoption timeline—and uncertainty around both of these factors is sufficiently high that we are not incorporating our findings into our baseline economic forecasts at this time—our estimates highlight the enormous economic potential of generative AI if it delivers on its promise.

Joseph Briggs

Devesh Kodnani

Appendix

Our Baseline Assumes Tasks in 13 Categories Up to a Difficulty Level of 4 Could Be Automated

AI-Exposed Work Activity	Examples of Automation	Examples of Tasks by Difficulty (O*NET 1-7 Scale)
Getting Information	Web scrape data from online sources and consolidate into a clean dataset; conduct and summarize a review of prior research based on a textual query and answer follow-up questions	2: Follow a standard blueprint 4: Review a budget 6: Study international tax laws
Monitoring Processes, Materials, or Surroundings	Monitor sensor input and system logs for manufacturing and utilities system anomalies; monitor internet activity for changes in sentiment or trending themes	2: Check to see if baking bread is done 4: Test electrical circuits 6: Check the status of a patient in critical medical care
Identifying Objects, Actions, and Events	Identify objects, music, terminology, and people when provided with textual/visual/auditory input; provide context on identified subject	2: Test an automobile transmission 4: Judge the suitability of food products for an event 6: Determine the reaction of a virus to a new drug
Estimating the Quantifiable Characteristics of Products, Events, or Information	Produce market size estimates based on assumptions grounded in existing research; estimate parameters using statistical modeling on input data and select optimal model	2: Estimate the size of household furniture to be shipped 4: Estimate transportation delays from inclement weather 6: Estimate the size of resource deposits beneath the world's oceans
Processing Information	Process raw data from documents, sensors, and humans into clean datafiles that are easily subscriptable for analysis; provide summaries of data relevant to user needs	2: Calculate the costs for shipping packages 4: Calculate the adjustments for insurance claims 6: Compile data for a complex scientific report
Evaluating Information to Determine Compliance with Standards	Review documents and proposed actions for compliance with legal, regulatory, and corporate standards; provide arguments and scenarios for and against compliance in unclear cases	1: Review forms for completeness 4: Evaluate a complicated insurance claim for policy compliance 6: Make a ruling in court on a complicated motion
Analyzing Data or Information	Perform statistical analysis of and identify trends within large datasets; forecast future data based on optimal combination of variables and model with best out-of-sample predictive power	1: Skim a short article to gather the main point 4: Determine the interest cost to finance a new building 6: Analyze the cost of medical care services for all US hospitals
Updating and Using Relevant Knowledge	Draft and update reports in corporate knowledge base; update statistical and financial models based on new data which challenges prior scenarios/assumptions	2: Track price changes in a small retail store 4: Track changes in maintenance procedures for repairing SUVs 6: Learn information about a complex and rapidly-changing technology
Scheduling Work and Activities	Automatically schedule meetings and work activities using availabilities and emails; assign tasks and estimate time to completion based on past experience	2: Make appointments for patients using a predetermined schedule 4: Prepare the work schedule for salesclerks in a large retail store 6: Schedule a complex conference program with parallel sessions
Organizing, Planning, and Prioritizing Work	Delegate and prioritize tasks based on time to completion and importance; identify gaps or bottlenecks in work plans and target resources or managerial attention	2: Organize a work schedule that is repetitive and easy to plan 4: Plan and adjust a to-do list according to changing demands 6: Prioritize and plan multiple tasks several months ahead
Documenting/Recording Information	Transcribe and summarize the content of in-person meetings; write system reports based on sensor and human data input	2: Record the weight of a patient during a routine health exam 4: Document the results of a crime scene investigation 6: Maintain information about satellite use for industry communications
Interpreting the Meaning of Information for Others	Explain the structure and function of code or statistical results in easy-to-understand language; translate code and text between languages; summarize and contextualize text with technical jargon	1: Interpret a blood pressure reading 4: Interpret how foreign tax laws apply to U.S. exports 6: Interpret a complex experiment in physics for general audiences
Performing Administrative Activities	Draft automated email responses; schedule and manage meetings and work calendars; file and organize paperwork; book reservations	2: Complete routine paperwork 4: Complete tax forms for a small business 6: Serve as the benefits director for a large computer sales firm

Source: O*NET, Goldman Sachs Global Investment Research

1 ^ Sevilla et al. (2022) show that since the advent of deep learning approaches in the 2010s, the “training compute” used to train AI models (i.e., the number of computations used to train AI models) has doubled approximately every 6 months, less than one-third the doubling time implied by Moore’s law. See: Sevilla, James et al., “Compute Trends Across Three Eras of Machine Learning”, 2022 International Joint Conference on Neural Networks, 9 March 2022.

2 ^ OpenAI, “GPT-4 Technical Report”, 2023.

3 ^ Zhang, Daniel et al., “The AI Index 2022 Annual Report”, Stanford Institute for Human-Centered AI, March 2022.

4 ^ This assumption is conservative, and we could imagine a scenario in which AI improves the efficiency of agricultural logistics and production, particularly if AI and robotics interact to automate physical tasks.

5 ^ Autor, David, Caroline Chin, Anna M. Salomons, and Bryan Seegmiller. New Frontiers: The Origins and Content of New Work, 1940–2018. No. w30389. National Bureau of Economic Research, 2022.

6 ^ Acemoglu, Daron and Pascual Restrepo. "Automation and new tasks: How technology displaces and reinstates labor." Journal of Economic Perspectives 33, no. 2 (2019): 3-30.