

BROOKINGS

RESEARCH

Generative AI, the American worker, and the future of work

Molly Kinder, Xavier de Souza Briggs, Mark Muro, and Sifan Liu

October 10, 2024

- Existing generative AI technology already has the potential to significantly disrupt a wide range of jobs. We find that more than 30% of all workers could see at least 50% of their occupation's tasks disrupted by generative AI.
- Unlike previous automation technologies that primarily affected routine, blue collar work, generative AI is likely to disrupt a different array of "cognitive" and "nonroutine" tasks, especially in middle- to higher-paid professions.
- Despite the high stakes for workers, we are not prepared for the potential risks and opportunities that generative AI is poised to bring.
- The report emphasizes the importance of developing strategies to proactively shape AI's impact on work and workers. This includes fostering worker engagement in AI design and implementation, enhancing worker voice through unions or other means, and developing public policies that ensure workers benefit from AI while mitigating harms such as job loss and inequality.

The launch of ChatGPT-3.5 at the end of 2022 captured the world's attention and illustrated the uncanny ability of generative artificial intelligence (AI) to produce a range of seemingly human-generated content, including text, video, audio, images, and code. The release, and the many eye-catching breakthroughs that quickly followed, have raised questions about what these fast-moving generative AI technologies might mean for work, workers, and livelihoods—now and in the future, as

new models are released that are potentially much more powerful. Many U.S. workers are worried: According to a [Pew Research Center poll ↗](#), most Americans believe that generative AI will have a major impact on jobs—mainly negative—in the next two decades.

Despite these widely shared concerns, however, there is little consensus on the nature and scale of generative AI's potential impacts and how—or even whether—to respond. Fundamental questions remain unanswered: How do we ensure workers can proactively shape generative AI's design and deployment? What will it take to make sure workers benefit meaningfully from its gains? And what guardrails are needed for workers to avoid harms as much as possible?

These animating questions are the heart of this report and a [new multiyear effort](#) (<https://www.brookings.edu/collection/workers-and-ai/>) we have launched at Brookings with a wide range of external collaborators. Through research, worker-centered storytelling, and cross-sector convenings, we aim to enhance public understanding, inform policymakers and employers, and shape our societal response toward a future where workers benefit meaningfully from AI's gains and, as much as possible, avoid its harms.

In this report, we frame generative AI's stakes for work and workers and outline our concerns about the ways we are, collectively, underprepared to meet this moment. Next, we provide insights on the technology and its potential impact on jobs, drawing on our analysis of detailed data from OpenAI (described [here ↗](#)) that explores task-level exposure for over a thousand occupations in the labor market. Finally, we discuss three priority areas for a proactive response—employer practices, worker voice and influence, and public policy levers—and highlight immediate opportunities as well as gaps that need to be addressed. Throughout the report, we draw on insights from a recent Brookings workshop we convened with more than 30 experts from different disciplines—policy, business innovation and investment, labor, academic and think tank research, civil society, and philanthropy—to grapple with those fundamental questions about AI, work, and workers.

The scope of this report is more limited than the full suite of [concerns about AI's impact on workers ↗](#). Conscious that our effort builds on an already robust body of

academic work, dedicated expertise, and policy momentum on key aspects of job quality and harms from AI (including privacy, surveillance, algorithmic management, ethics, and bias), our primary focus is addressing some of generative AI's emerging risks for which society's response is far less developed, especially risks to livelihoods.

What is at stake?

Despite high stakes for workers, we are not prepared for the potential risks and opportunities that generative AI is poised to bring. So far, the U.S. and other nations lack the urgency, mental models, worker power, policy solutions, and business practices needed for workers to benefit from AI and avoid its harms.

To date, most of the discussion around ChatGPT and similar technologies has stayed away from work and workers. Other serious concerns are dominating the debates, with far more focus on national security, disinformation, privacy and surveillance, intellectual property, electricity consumption, and deception (epitomized by "deep fakes" as an instrument of financial and political fraud). Attention is spreading widely and quickly, but also diffusely.

What is not receiving nearly enough attention are the workers and the content and terms of their work, which are so crucial for delivering the value of AI for society. Attention to AI's impacts on the world of work and livelihoods has been secondary at best, and mostly conjecture.

To the extent work is discussed, conversations about AI's implications have been stuck at the extremes. On one end, [techno-optimists champion a world ↗](#) of abundance and unlimited possibility, of [drudgery-busting ↗](#) AI assistants in our pockets, [AI-powered scientists ↗](#) curing cancer, and turbocharged productivity creating prosperity for all. On the other extreme lie sweeping predictions of doom, mass job loss, and the [end of human employment ↗](#)—or [even existence ↗](#)—as we know it.

It is impossible to predict the future trajectory of technological advancements. Indeed, the range of possible AI futures is exceedingly broad, from a near-term plateau in useful capabilities to exponential improvements resulting in capabilities at the level of long-hypothesized artificial general intelligence (AGI), with sweeping economic and social consequences.

Though confident prediction is not possible, what is clear is that the design and deployment of generative AI technologies are moving far faster than our collective response to understand and shape them.

We are underprepared to meet this complex and growing challenge. Take public policy: While today, there are few partisan battlelines in policy responses to AI's threats to work, there is also relatively little urgency, momentum, or [concrete examples of legislation and regulations ↗](#) at the state or federal level that address automation risks or generative AI's workplace threats—or, conversely, that directly encourage responsible engagement of workers in making the most of AI's capabilities.

Second, worker organization and power (or lack thereof) remain [critical to shaping how AI is deployed in the economy ↗](#), yet they appear spotty and limited. While there have been a few high-profile examples of workers actively shaping AI safeguards through collective bargaining, such as [last year's landmark agreement between Hollywood writers and major studios \(https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/\)](#), we find that there is a stark mismatch between the industries and occupations most exposed to generative AI and the sectors where workers have substantial union strength or other access to worker organizations, voice, and influence.

Third, there's a widely reported "gold rush" mentality and hype driving AI deployment, with many companies rushing to adopt the technology despite fresh questions about [cost and expected profitability ↗](#). While large tech companies such as Google, Meta, and Microsoft are making big investments in developing AI, most other organizations—whether in business, government, or the nonprofit sector—will focus on using these AI tools rather than developing them. These "deployers" of AI technology are employers as well, with employees expected to adapt in some way to growing AI deployment. We

refer to such organizations in this report as “employer-deployers”—a key decisionmaking group that will influence how AI technologies are adopted and managed. Currently, there are few guidelines or codes of conduct for how companies should ethically implement AI with respect to their workforce. At the same time, many companies, especially those publicly traded or aiming to go public, feel intense pressure from competitors and investors to adopt AI to save on labor costs and increase efficiency.

“Currently, there are few guidelines or codes of conduct for how companies should ethically implement AI with respect to their workforce. At the same time, many companies, especially those publicly traded or aiming to go public, feel intense pressure from competitors and investors to adopt AI to save on labor costs and increase efficiency.”



Bumpy product deployment and broader uncertainty notwithstanding, the stakes for workers are unquestionably high. Even on its current trajectory, without any dramatic acceleration of capability gains, generative AI technology is poised to impact a broad range of workers in fields as diverse as law, marketing, finance, health care, computer programming, customer service, the creative arts, administrative support work, education, and media. For some industries and occupations, the first waves of that disruption are only months away, or are even quietly underway right now. Interacting with an AI-powered customer service agent or bot—something that is already commonplace—is just the tip of this iceberg.

These changes bring both opportunity and risk, as many observers have underlined. On one hand, generative AI has the potential to complement millions of workers’ skills, enabling them to be more productive, creative, informed, efficient, and accurate. On the other hand, employers may choose to automate some, or even all, of their employees’ work, leading to possible job losses and weakened demand for previously sought-after skills. For still other workers, especially those such as writers, journalists, and creatives who generate original content, generative AI presents troubling, existential questions around copyright and consent. AI also raises the specter of

potent new tools for employers to monitor and surveil employees, [undermining worker autonomy ↗](#), agency, and [power ↗](#).

Thus, even as generative AI has the potential to boost incomes, enhance productivity, and open up new possibilities, it also risks degrading jobs and rights, devaluing skills, and rendering livelihoods insecure.

Yet the future is [not preordained ↗](#). Ultimately, whether workers benefit from AI-driven productivity gains or suffer harm and precarity depends in part on the ability of workers and other stakeholders to shape the technology's deployment, as well as the specific choices that employers, technology companies, policymakers, consumers, and civil society make. We know from a long and mixed [economic history ↗](#) going back centuries that unrestrained technological advancement [can lead to greater inequality ↗](#) and lasting pain for workers and their communities. Technology is not destiny, but inaction is.

Not your grandparents' automation: Understanding generative AI's potential impact on work and workers

What are generative AI's likely impacts on work and workers? In this section, we briefly summarize several defining features of this new technology and glean insights from the data provided by OpenAI. We include summary findings of new Brookings research analyzing OpenAI data, which looks at task exposure to existing ChatGPT-4 technology across more than 1,000 occupations. The data is best interpreted as directionally useful in identifying the types of occupations that might see more (or less) disruption from current generative AI technology. But the analysis does not and cannot offer definitive predictions or precise accounting of specific impacts. For more information on our methodology and some of its limitations, please see the appendix.

The focus of our analysis: Generative AI and what makes it different

Popularized by ChatGPT-3.5's release at the end of 2022, "generative" AI is a compelling technological breakthrough with traits and sophisticated capabilities fundamentally different from past forms of computerization and automation. The tool's particular combination of traits makes it different: its capacity to generate new content, its relative ease of diffusion, and the fact that, for now, it is a mostly "disembodied" technology rather than a physical tool for work like an industrial robot (though that could change soon with advances in machine vision and other AI technologies).

Generative AI tools are, in some key respects, novel among information technologies because of their ability to create entirely new content from the data the AI models were trained on. That's what makes them "generative." As a type of machine learning, generative AI works as an algorithm that can produce a wide range of new content, including images, music, text, audio, video, and code. The technology is enabled by large language models (LLMs) that train on vast data sets, detecting statistical patterns and structures that the model then uses to generate new content.

Especially critical is generative AI's ability to predict and generate new "natural language" content useful to a user's momentary intent and need, not unlike an auto-prompt feature on a smartphone—whether it be to write correspondence, answer questions, produce computer code, develop business plans, or scrape the internet and then generate ideas for action. Advanced generative AI models such as Dall-E 3, Midjourney, and Stable Diffusion can create high-quality visual content from text input, while programs such as Sora have made striking advances in text-to-video content. Now, systems are coming that can combine different data types such as text, images, audio, and video for both input prompts and generated outputs.

"ChatGPT became the fastest-spreading tech platform in history by reaching 1 billion monthly visits just four months after its 2022 launch."

Generative AI is also distinct for its relative ease of diffusion, mostly through existing web browsers and apps on computer devices of all kinds. In other words, the rails to

carry generative AI are mostly already in place.

As one benchmark, the Census Bureau [reports ↗](#) that it took around two decades for the personal computer to become ubiquitous after its introduction in the late 1970s. More recently, it took six to seven years for the smartphone to become ubiquitous in the U.S. after the first iPhone's launch in 2007. By contrast, ChatGPT became the fastest-spreading tech platform in history, reaching [1 billion monthly visits ↗](#) (a rough proxy for users) just four months after its November 2022 launch.

While estimates vary widely, workplace adoption of generative AI is still modest today, as employers experiment with early use cases and face lingering concerns around privacy, security, and accuracy. (Workers experiment too, sometimes in clandestine fashion, regardless of their employer's rules.) And while more widespread diffusion may only happen over a longer time horizon, workplace adoption of AI may face lower barriers than previous forms of technology given: 1) its accessible and user-friendly interface that does not require machine learning expertise; and 2) its modest infrastructure requirements.

Finally, AI tools remain—for now—disembodied, unlike physical robots assembling goods in a factory or vacuuming your floor. Digital in nature, AI tools remain oriented toward information-based tasks. But that too could change, as LLMs are designed to communicate with material objects and their sensors.

Generative AI's capabilities portend a stark break from previous 'skill-biased' technologies

Generative AI's capabilities represent a departure from previous workplace technologies. For decades, as a large body of research shows, technology has been "[skill-biased](#)" ↗: It substituted for routine skills common in middle- and some low-wage jobs (such as manual accounting, production, and food preparation), while it complemented non-routine skills typical of higher-paid jobs (such as managerial decisionmaking, complex analysis, and the use of human creativity).

Technologies such as ChatGPT upend this paradigm. In fact, generative AI is not likely

to disrupt physical, routine, blue collar work much at all, barring technological breakthroughs in robotics. Instead, generative AI excels at mimicking the kinds of non-routine skills and interactive traits that just a few years ago experts considered impossible for computers to perform, including programming, prediction, writing, creativity, projecting empathy, communication and persuasion, and analysis. Most of the industries that face the greatest exposure to generative AI today are those that just a few years ago were ranked at the bottom of automation risk.

Already, generative AI technologies are capable of performing a wide range of tasks, often quite sophisticated, and at times even without human oversight. Some of the capabilities that the technology can perform autonomously without human oversight include those in the box below:

Box 1. A sample of ChatGPT-4's autonomous capabilities

Coding

Writing, editing, and transforming text and code

Debugging code or software

Programming in computer languages such as Python and C++

Assisting with data analysis

Writing and reading

Summarizing documents

Reading text from PDFs

Writing questions for an interview or assessment

Writing and responding to emails

Writing lessons plans

Preparing training materials

Information sharing, retrieval, and synthesis

Translating between languages; transcribing

Answering questions about a document

Searching an organization's existing knowledge, data, or documents, and retrieving information

Informing anyone of any information via any written or spoken medium

Conducting analysis and research

Making recommendations given data or written input

Analyzing written information to inform decisions

Performing legal research and counsel

Source: OpenAI and University of Pennsylvania [working paper](#)

Given this extraordinary set of capabilities (which human workers can tweak, guide, and complement to enhance AI's perceived responsiveness and traits such as empathy) and the huge interest in AI deployment, it's time to get a handle on which workers in which industries are most likely to be affected—and also how equipped and likely they are to be able to shape the deployment of AI in their fields.

Looking ahead: Potential widespread effects, with the greatest impacts on middle- to higher-paid occupations, clerical roles, and women

The exposure data from OpenAI suggests that generative AI technology may impact

broad swaths of the nation's workers. We find that more than 30% of all workers could see at least 50% of their occupation's tasks disrupted by generative AI, while some 85% of workers could see at least 10% of their work tasks impacted.

The sectors that face the greatest exposure are dominated by higher-paying fields with advanced degree requirements, such as STEM pursuits, business and finance, architecture and engineering, and law, in addition to lower-paying, "middle-skill" office and administrative support occupations. Manually intensive, blue collar sectors face the least exposure, while lower-paid service sector jobs will also likely see more modest effects.

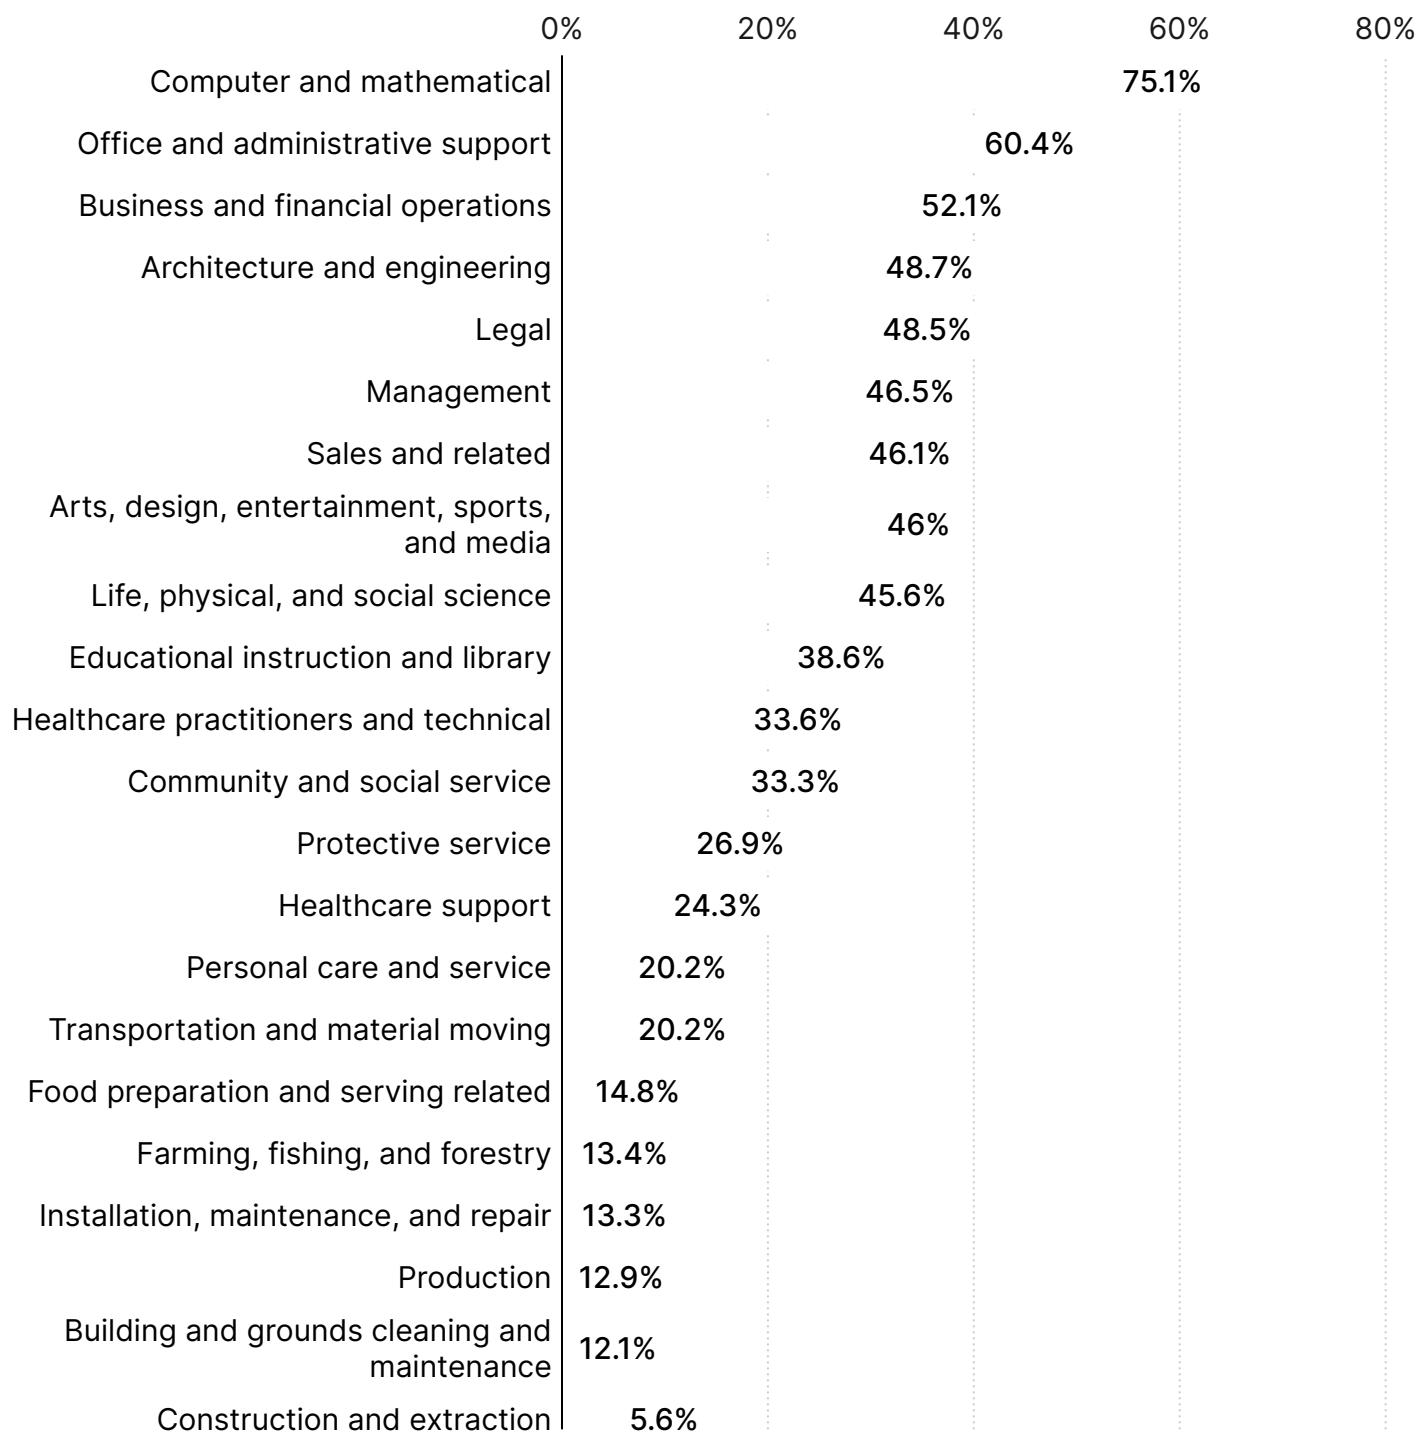
Education, health care, and community and social services have medium exposure according to our analysis. For instance, elementary school teachers and registered nurses could see substantial time savings in about one-third of tasks. A teacher might save time on tasks such as grading, planning activities, administering tests, maintaining records, and preparing reports. For a registered nurse, many manually intensive tasks that require in-person administration—such as performing a physical exam, conducting a lab test, or administering an IV—will see minimal impact, but generative AI could save time on other tasks such as evaluating diagnostic tests, recording patient information, modifying treatment plans, maintaining records, recommending treatments, and performing administrative and managerial functions.

Zooming out, Figure 1 shows how this looks across sectors, with exposure levels plotted for occupational groups as bars depicting major groups' exposure levels. In the chart, the bars' lengths reflect the share of the major occupational group's tasks that LLMs can reduce the time to complete by 50% or more. At a glance, the figure helps us spot that some fields—such as computer work, office and administrative support, business and financial operations, and engineering—stand out as having relatively high levels of exposure.

FIGURE 1

LLM task 'exposure' levels for occupation groups

2022



Source: Authors' analysis of task exposure data from OpenAI.

Note: This data estimates the share of each occupational group's tasks for which ChatGPT-4 can reduce the time for humans to complete by 50% or more.

[Download image](#)

Looking at generative AI exposure impacts more closely, it's possible to see how LLM exposure varies by occupations' pay levels. Figure 2 shows that for the most part, higher-paying occupational groups such as computer work, management, engineering, and business-financial roles stand out for being forecasted to encounter high exposure to ChatGPT-4 and other LLMs.

Further, the bubbles representing various occupational groups are sized according to the current number of workers in those jobs. That means that several very large occupational groups—such as business, management, and health care work—stand to undergo significant exposure to generative AI. This alone forecasts the technology's broad implications for the labor market.

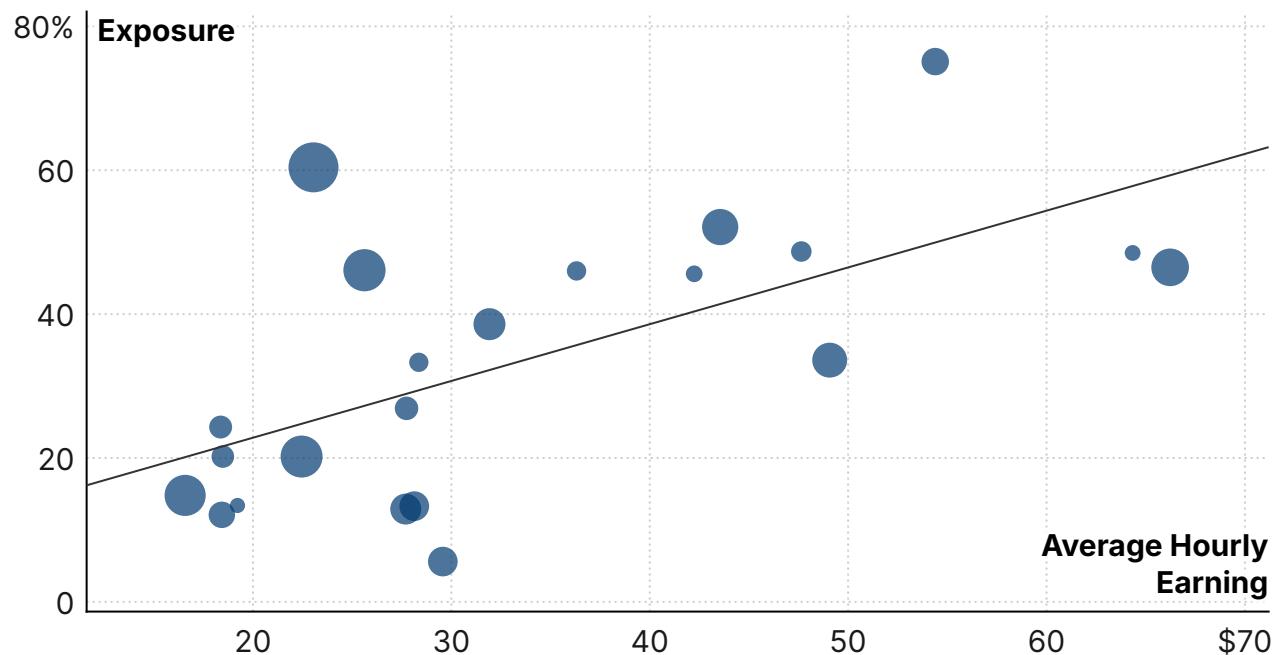
FIGURE 2

LLM task ‘exposure’ levels for occupation groups, with reference to average pay

2022

2022 Jobs

○ 2M ○ 8M ○ 20M



Source: Authors’ analysis of task exposure data from OpenAI.

Note: “Exposure” readings estimate the share of each occupational group’s tasks for which ChatGPT-4 can reduce the time for humans to complete by 50% or more.

[Download image](#)

B | Brookings Metro

In exploring this data, it is important to recall that these exposure rates in themselves do not predict—let alone determine—the nature of effects on workers. Rather, they reflect generative AI’s potential involvement with jobs or occupational groups, without distinguishing between labor-augmenting or labor-displacing effects.

However, it is obviously important to probe the specific potential for LLM-driven automation (or work replacement), given the technology’s widely feared potential to disrupt human work. To assess the technical feasibility of LLMs automating specific tasks within occupations, we analyzed data from OpenAI estimating the likelihood of

generative AI completing tasks with no human oversight, per the autonomous capabilities list in Box 1. Tasks with a high exposure and high likelihood of being completed without human oversight were categorized as “more likely to automate.” Five sectors emerge as having relatively high exposure and high automation potential, as detailed in the chart below alongside illustrative occupations.

TABLE 1

Occupations with high exposure and high automation potential

Occupations	% of tasks more likely automated	% women	Number of Jobs, 2023
Office and administrative support occupations	71%	71%	18,553,450
Bookkeeping, accounting, and auditing clerks	100%	86%	1,501,910
Customer service representatives	86%	65%	2,858,710
Office clerks, general	84%	81%	2,496,370
Secretaries and administrative assistants, except legal, medical, and executive	78%	92%	1,785,430
Business and financial operations occupations	44%	54%	10,087,830
Insurance underwriters	100%	57%	101,310
Insurance claims and policy processing clerks	100%	80%	241,650
Tax preparers	83%	66%	81,650
Sales and related occupations	46%	49%	13,380,660
Travel agents	88%	80%	58,250
Securities, commodities, and financial services sales agents	80%	25%	479,630
Sales representatives of services, except advertising, insurance, financial services, and travel	67%	32%	1,142,020
Legal occupations	49%	52%	1,240,630
Legal secretaries and administrative assistants	88%	96%	152,790
Lawyers	32%	40%	731,340
Paralegals and legal assistants	58%	83%	354,890
Computer and mathematical occupations	45%	27%	5,177,400
Database architects	67%		59,920
Data scientists	66%		192,710
Computer programmers	66%	22%	120,370

Source: Brookings analysis of OpenAI and Bureau of Labor Statistics data

B | Brookings Metro

[Download image](#)

Office and administrative support occupations stand out for the sector's high exposure, high automation potential, and large number of workers. What's more, women comprise the overwhelming majority of the nearly 19 million Americans employed in the sector, which has provided the lion's share of decent-paying, stable jobs with upward mobility potential for women without a college degree in jobs such as bookkeepers, legal secretaries, HR assistants, bank tellers, and payroll clerks. For decades, technology has contributed to the hollowing out of these jobs; generative AI could accelerate these trends.

The stakes are especially high for this racially and ethnically diverse group of lower-middle-class women, many of whom may risk falling into more precarious, lower-paid work if this work is displaced. In this regard, much more analysis is needed of the likely distribution of AI's employment effects by race, disability, and other statuses and identities. For example, some of the most exposed and possibly vulnerable jobs are currently disproportionately held by white workers, while others—the ranks of bank tellers and HR assistants, say—roughly mirror the racial makeup of the labor force as a whole. And regardless of race, not all workers in a given occupation will necessarily be affected equally.

The plight of clerical workers illustrates a broader trend: It is women, not men, who face both the highest exposure to generative AI and the highest automation risk, due to their overrepresentation in white collar work that requires a college degree and in administrative support roles. Altogether, 36% of female workers are in occupations in which generative AI could save 50% of the time on tasks, compared to 25% of male workers, according to Brookings' analysis of OpenAI's GPT-4 ratings of task susceptibility.

This reality runs counter to popular conceptions of technology and work: The dominant stereotype of a worker at high risk of automation is often that of a blue collar, male worker in manufacturing, warehousing, or truck driving, or perhaps a computer programmer. Yet generative AI is likely to have only a minimal impact on male-dominated blue collar industries, barring further advances in robotics technologies.

TABLE 2

Occupations with low exposure and low automation potential

Occupations	% of tasks more likely automated	% women	Number of Jobs, 2023
Construction and extraction occupations	2%	4	6,225,630
Construction laborers	1.8%	5	1,019,090
Sheet metal workers	1.5%	7	116,190
Plumbers, pipefitters, and steamfitters	3.9%	2	436,160
Production occupations	9%	30	8,770,170
Printing press operators	27%	25	151,450
Cutting, punching, and press machine setters, operators and tenders, metal and plastic	4.7%	7	179,230
Machinists	9%	6	290,720
Transportation and material moving occupations	20%	21	13,752,760
Hand laborers and freight, stock, and material movers	3%	24	3,008,300
Stockers and order fillers	21%	37	2,872,680
Bus drivers, transit and intercity	25.5%	37	184,990

Source: Brookings analysis of OpenAI and Bureau of Labor Statistics data



[Download image](#)

In sum, generative AI is not just the latest update of the various digital and automation technologies that have been reshaping wide-ranging segments of the labor market for decades. It's something new and distinct.

Open questions: What we still don't know

While the data here suggests some contours of how generative AI could impact a

range of workers and types of work, what we know about the likely impacts and how best to shape them remains radically incomplete. The technology is still in its early stages, and insights into its workings and potential impacts remain sketchy beyond an appreciation of great potential and necessary caution. Overall, we don't know a lot yet about how "exposure" to generative AI will translate into real-world impacts on workers.

"We don't know a lot yet about how "exposure" to generative AI will translate into real-world impacts on workers."



Several key questions loom large and suggest how we might at least locate the unknowns and generate more tangible experience and learning:

How much and how rapidly will generative AI augment—as opposed to automate—human labor? We don't know the extent to which generative AI will affect overall demand for human labor (types and numbers of jobs) or, when it comes to the content of work, how much AI will in fact augment (enhance capabilities and/or improve efficiency, productivity, and performance) versus automate jobs—and how soon these changes will unfold.

For example, there are multiple ways AI might augment the role of a computer programmer: enhancing productivity, debugging work, checking for errors, and teaching new skills. On the other hand, AI might also automate some or even much of the work, taking on routine tasks and even generating code. We need to learn when and how AI can complement or displace workers, and whether what looks initially like "augmentation" might ultimately lead to displacement. Relatedly, we need to clarify which exposed workers are vulnerable to displacement, and who will not be very vulnerable and instead will "roll with the changes." As we analyze and track these effects and potential others, we should very carefully consider gender, race, disability, and other differences—not only where workers of different backgrounds are concentrated in terms of industries and occupations, but also how well positioned and supported they are to respond.

Exactly which workers are most likely to benefit or suffer harmful dislocation?

Related to the previous point, we don't yet know which workers are most likely to benefit—or lose out—from generative AI within occupations and sectors. It is possible that the technology may be more beneficial or harmful to workers based on experience and skill, for instance. Recent academic experiments in industries ranging from [customer service support ↗](#) to [consulting ↗](#) to [computer programming ↗](#) documented a “leveling up” dynamic in which less skilled or less experienced workers experienced the biggest gains from using AI. But it is possible that the opposite may also be true: Some jobs may be “de-skilled.” For example, generative AI might “upskill” the ability of a novice grant writer to prepare higher-quality grant applications, potentially even rivaling one written by a seasoned and high-performing grant writer. But another possibility is that the job of grant writer may be de-skilled, with specialized (and to some extent, rarer) skill substituted by copying and pasting from generative AI. It is also possible that as the technology improves, more senior employees will experience productivity boosts while demand for lower-level employees erodes. All of these scenarios imply a massive need for workers to adapt and a wide spectrum of forms that could materialize.

The [case of Hollywood writers](#) (<https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/>) is instructive here, because their landmark agreement with major studios aimed to build in AI adaptation along with guardrails. The Writers Guild opted to support the use of AI technology in principle and assume its ongoing evolution, while also establishing a role in co-determining the technology's use (e.g., stipulating what AI cannot replace) and safeguarding intellectual property, employment levels, and key features of compensation.

“Tied to the question of how gains and losses will appear is the question of how both, in turn, will impact inequality across multiple dimensions: income, wealth, gender, race, educational attainment, and geography.”



How could AI-triggered changes affect inequality, and how can that best be

shaped? While a few [signals](https://www.brookings.edu/articles/ais-impact-on-income-inequality-in-the-us/) (<https://www.brookings.edu/articles/ais-impact-on-income-inequality-in-the-us/>) have been coming in, we don't yet know the likely overall impact of generative AI on inequality. Nor do we know what efforts to mitigate even greater inequality might prove helpful, or what steps might be delivered at scale to ensure LLMs deliver on "leveling up" or other gap-reducing changes. Tied to the question of how gains and losses will appear is the question of how both, in turn, will impact inequality across multiple dimensions: income, wealth, gender, race, educational attainment, and geography.

Can generative AI really level up lower performers, narrow gaps between them and "star" workers, and thereby [lower inequality and boost the middle class?](#) ↗ To what extent will workers benefit from AI-supported worker productivity gains, and via what mechanisms (stock ownership or other equity, performance bonuses, etc.)? What will happen to the value of a college degree? How difficult will it be for displaced workers to transition to new roles, especially if their education and training are highly specialized? Who will be positioned to benefit from the new jobs that will be created? Who—and in what places—will reap the biggest financial gains?

Beyond affecting the demand and rewards for human labor, how could generative AI add to harms for workers and their workplaces? A growing body of evidence documents multiple ways that employers' use of AI [can harm workers](#) ↗ beyond livelihood risks, including by undermining their [power](#) ↗, contributing to workplace injuries, violating copyright, monitoring and surveilling with little or no consent, introducing bias, collecting their data, and exacerbating scheduling and other pressures through algorithmic management and decisionmaking. For now, we know too little about how generative AI specifically might contribute to or exacerbate those harms, or perhaps introduce new ones.

Three priority areas for a proactive response

In this final section, we outline three priority challenges that we explored at the Brookings workshop earlier this summer. This short list is by no means exhaustive, but

it does encompass a range of leverage points for positive change. Again, our aim is to focus on emerging concerns—especially involving risks to livelihoods—that have received much less attention to date than well-documented harms such as bias and surveillance.

1. **What is 'good' in the employer-deployer arena, and what makes it good?**

The first key priority is establishing what good, responsible business and organizational practice looks like for employers and deployers of generative AI, and what could support it.

Generative AI has been all the corporate buzz over the past year and a half, with investor interest peaking and companies racing to demonstrate their embrace of the innovation and new ways to create value—or at least not be left too far behind.

In a highly unequal economic system—one centered on maximizing shareholder value and short-term returns while persistently concentrating huge market power in a small number of leading companies—a predominant focus of industry discussions has been on the potential labor cost savings and efficiency gains from deploying AI. There has been little public discussion or focus on worker impacts or worker engagement in shaping AI's use at work.

Yet there is a powerful business case for having that discussion. A [growing body of research ↗](#) is documenting the benefits of incorporating workers into the design and roll out of new technologies, compared to top-down implementation that does not incorporate workers' unique knowledge and insights. Who has not lived through an enterprise software or hardware rollout that went badly, with costly results? We have seen at least some of this movie before: Decades ago, comparisons of U.S. and Japanese manufacturers [revealed the powerful competitive advantages ↗](#) of engaging workers as active problem solvers rather than passive followers of formulaic rules.

In addition, there is a long-established set of norms and gold-standard examples for what it looks like to be a “high-road” employer in terms of worker pay and benefits, dimensions of job quality such as creativity and purpose, and investments in upskilling and job security. In some market economies, “co-investments” is more accurate:

Employers, worker organizations, and government invest together, so each has skin in the game. Beyond employment per se, newer standards labeled “sustainability” address—and also rate and rank—companies’ data privacy, environmental sustainability performance, and sometimes other factors. Investors, CEOs, and other corporate leaders can and do respond to those.

But here’s the urgent challenge: There is not yet any standard for acting as a high-road employer-deployer in the context of generative AI. These standards might include actions such as assessing risks and opportunities; setting goals that put the fortunes and capabilities of workers in the center and not the margins; engaging workers in designing and implementing AI deployment and its rewards (and in the process, redefining work and sharing gains from higher productivity); and responsibly supporting transitions for workers who need them (e.g., as demand for certain human skills and tasks declines).

Given that we are in a “pre-regulatory” moment for AI, voluntary action by employers is one available lever that can field test and socialize standards that inform regulation. Examples include several sets of principles developed by the Partnership on AI, including voluntary standards and collaboration among employers with regards to [responsible deployment of synthetic media ↗](#). The Partnership also launched a task force in 2023 that generated a structured set of questions for risk assessment and solution-building by employers and workers: [Guidelines for AI and Shared Prosperity ↗](#). Two other nonprofit, business-facing organizations, Chief Executives for Corporate Purpose and JUST Capital, are likewise exploring what “good” can mean—and how to realize it in corporate practice.

Leading companies and worker-led organizations are starting to launch promising collaborations. For example, in 2023, Microsoft and the AFL-CIO announced a first-of-its-kind [tech-labor partnership on AI and the future of the workforce ↗](#), which aims to educate workers, bring workers’ voice into AI development, and shape pro-worker policies.

Meanwhile, in addition to stronger leadership by employers, worker organizations, and others, more research on the ground—especially in the workplace—is required to answer key questions. How is AI being deployed in different settings and industries?

How are employers making decisions? How, if at all, are affected workers engaged as co-designers and/or user-deployers?

2. **How can we enhance worker voice and power in an economy with low unionization, especially in the industries and occupations most exposed to AI?**

The second key priority is adapting and scaling models for amplifying worker voice in the generative AI moment. This requires that we recognize and tackle a “great mismatch.”

As we documented in a [recent multimedia case study](https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/) (<https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/>), last year’s Hollywood writers’ strike illustrated the power of organized workers using their collective voice to protect their livelihoods. It even showed the potential of “sectoral bargaining”—rare for a U.S. labor dispute. The contract the Writers Guild secured with all the major studios includes far-reaching guardrails on generative AI—the first of their kind for any collective bargaining agreement.

By exercising their power, the writers pushed back on unchecked risks and succeeded in setting their own terms for the use of AI—not a ban on the technology, but instead regulation of its use in ways that could benefit writers and studios alike while reducing clear harms. Crucially, the writers emphasized both the income at stake as well as the creative purpose and meaning at stake in their technology-affected work lives and changing career paths.

“A great mismatch: The industries most exposed to generative AI have some of the lowest union representation in the economy.”

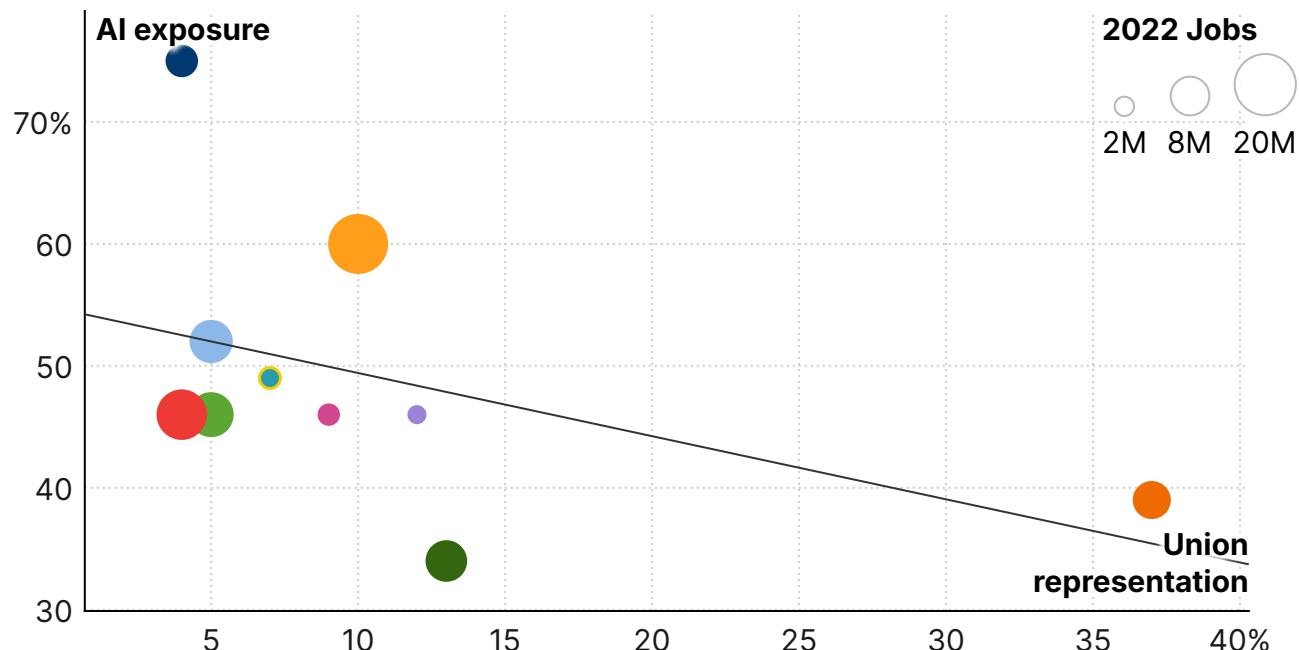
But the replicability of the writers’ success is limited by a great mismatch: The industries most exposed to generative AI have some of the lowest union representation in the economy.

Nationally, 10% of all workers and just 6% of private sector workers [were members of a union in 2023 ↗](#). In most of the industries with the highest AI exposure, unions represent an even smaller percentage of workers. For instance, unions represent only 1% of workers in finance—a highly exposed sector where corporate leaders [are reportedly ↗](#) exploring job and pay cuts to, for example, the entry-level analyst jobs that have traditionally offered the foundation for moving up. Education stands out as the lone exception: a medium-exposure sector with substantially higher union representation.

FIGURE 3

The 'great mismatch': Worker power in AI-exposed occupational groups

Legal
 Management
 Sales and related
 Computer and mathematical
 Architecture and engineering
 Business and financial operations
 Office and administrative support
 Life, physical, and social science
 Educational instruction and library
 Healthcare practitioners and technical
 Arts, design, entertainment, sports, and media



Source: OpenAI and Bureau of Labor Statistics data.

Note: Bubble sizes reflect employment levels of occupational groups.

[Download image](#)

B | Brookings Metro

Beyond formal power and bargaining rights through unions, workers in heavily exposed industries also lack voice and visibility in other forms of countervailing power, from worker justice organizations to sustained campaigns. For legal secretaries and HR assistants, there is no equivalent to "Fight for 15," the landmark campaign that shifted the goal posts and momentum on the minimum wage. Likewise, there is no

pro-worker alliance—equivalent to the National Domestic Workers Alliance or United for Respect—for bookkeepers or sales reps. Yet all of these large, ubiquitous occupations seem to be marketing targets for those offering AI training and tips, starting with popular and professionally oriented social media platforms such as LinkedIn and YouTube.

Compounding these twin challenges of low union density and limited countervailing worker power are the political obstacles to enacting broader labor law reforms, as well as the limitations to federal regulatory power stemming from recent [Supreme Court decisions](https://www.brookings.edu/articles/around-the-halls-after-chevron-whats-next/) (<https://www.brookings.edu/articles/around-the-halls-after-chevron-whats-next/>).

On the positive side, as workshop participants pointed out, there is a unique group of potential influencers: Sought-after technologists—including senior coders and other employees of Google, OpenAI, and other leading AI developers—are [speaking out](#) as whistleblowers and, in some cases, voting with their feet by seeking out new employers and expressing serious concerns about unregulated development and deployment of generative AI. Similarly, generative AI's disruption of new industries—including higher-paying, higher-status jobs that previously were not perceived at risk of technological change—may bring fresh opportunities to organize new and broader classes of workers.

“There is no “Fight for 15”-equivalent campaign for legal secretaries and HR assistants and no pro-worker alliance for bookkeepers or sales reps.”

Some innovative models at the local, state, and international levels may offer lessons on how to give workers more say in the use of a significant technological advancement in their field. For example, California is home to several promising approaches that could inform efforts to pilot AI-specific sectoral bargaining and other structural ways to give workers greater voice. During the pandemic, Los Angeles created [an innovative program of public health councils](#) comprised of frontline workers in sectors with high COVID-19 transmission rates; these workers were empowered to meet with

management, ensure compliance, mitigate transmission, and report concerns directly to public health officials. And in 2023, California enacted legislation creating the [Fast Food Council](#), a statewide council comprised of representatives from industry and labor that will set industry working conditions and standards. In [testimony to state lawmakers](#), Annette Bernhardt, a labor and technology policy researcher at the University of California, Berkeley, highlighted the important role of public sector workers and their unions (among other drivers) in shaping responsible AI deployment. Similarly, European [works councils](#) offer a well-documented model for incorporating worker voice.

To adapt models of worker voice to the challenges arising from generative AI, more experimentation is needed to enable positive AI use cases and models to be documented, replicated, and scaled. Other opportunities include exploring opportunities to require or support worker voice, such as through procurement standards and conditions attached to government grants or other public money. The growing number of occupations and industries perceived to be at risk from the technology may present fresh opportunities to organize concerned workers. Pressure from well-organized campaigns helps motivate leadership—both insisting on accountability and recognizing when employers do the right thing and serve as role models for peers and/or competitors. Labor history shows many ways to connect worker and work-centered campaigns to consumer power and “conscious consumerism” as well, both to discourage harmful practices and reward positive ones.

3. **Public policy solutions: What can policy solve for, and how?**

The third priority is developing public policy responses and using the public sector to pilot and model how to help workers succeed in an AI-affected economy.

On public policy around AI, there is a striking gap between the concerns of voters and the momentum on legislation or meaningful regulation. Attention in Congress, for example, is largely on AI risks such as disinformation, safety, and democracy. Congressional hearings always offer some attention-focusing and attention-expanding opportunity in principle, even during election season, but movement on legislation is very unlikely in the current session. Also unlikely is near-term action on key aspirations in a section on [“AI and the Workforce”](#) included in a recent Bipartisan Senate AI

Working Group document. Meanwhile, polling shows that among voters, [worries about livelihoods and work](#) are among the top, if not the top, concerns about AI.

The Biden administration's [Executive Order](#) on the safe development and use of AI, issued October 2023, was a promising catalyst for proactive policy responses to AI, with a clear focus on worker voice and power included alongside other risks and opportunities. Yet it mostly calls for more data and knowledge-building—both critical, to be sure—and does not offer broader solutions for shaping a positive future of work.

To be fair, it is not yet clear what those solutions should be, coming from government. But the potential mechanisms are many: from worker-protecting standards built into government procurement to investment in worker-centered innovation and creativity as part of value-creating AI use. As discussed at our workshop, a government response should include use of AI by the public sector itself, potentially to enhance a near-endless range of service delivery, enforcement, investment, research, and other government functions.

Helpfully, there are not yet entrenched partisan positions on policy responses—but neither are there models of state or federal legislation and regulation ready to scale, at least not well-understood models or those ready to be championed. Like the larger debate about AI risks, most of the [promising examples of state-level policies](#) focus on reducing harms from AI bias, algorithmic management, surveillance, productivity quotas, and privacy concerns, with far less policy attention to automation risk or potential fallout, such as rapid job dislocation and income loss. Yet some new ideas are emerging; for example, in keeping workers informed, as well as the right to not be forced to train AI to replace one's own job. These broader concerns defy easy regulatory or “silver bullet” legislative fixes, as it is challenging to parse out “good” versus “bad” automation.

Zooming out, the scope of issues that require proactive policy and regulatory responses range from copyright and fair use (particularly in the creative arts) to core labor rights, workforce protections, workforce development, and taxation. In this respect, the current policy inaction—and impasse—over the rapid rise of social media offers a cautionary tale: Lawmakers and other policymakers must act, and soon, if they are to have any hope of shaping the impact of this transformative technology on

workers.

States are moving faster than the federal government on enacting AI policies, but still face challenges. Given the pace of technological development and deployment, the next two years appear to be especially important for establishing worker rights and employer responsibilities in policy. Some experts at our workshop argued strongly that the focus should be broader than generative AI per se, and instead looking at technology systems and the full suite of technology into which generative AI embeds and through which it distributes (e.g., established enterprise software already in wide use).

As previewed above, one clear opportunity in public policy is for the public sector to act as a model employer and deployer of generative AI. Government can productively deploy AI in so many ways, and will clearly be a major buyer of AI-driven software and related tools. But the die is far from cast, and public sector use cases are especially early stage. Also, the large public sector workforce—about 24 million people, 80% of whom work at the state or local level—is much more unionized than the private sector workforce. Nearly one-third of government workers (32.5%) belong to a union, versus just 6% in the private sector, according to the [Bureau of Labor Statistics](#). And about one in seven workers in the U.S. economy work for government. Building out the model for public sector deployment would benefit from more policy experimentation, as well as coordination across states and between the state and federal levels.

Conclusion

As part of a larger, multiyear effort focused on understanding and shaping a positive future of work in an AI-affected world, this report has outlined some of the major stakes and questions that should guide and accelerate much-needed attention toward the issue. In collaboration with leaders across a range of sectors and many parts of the country, we will be tackling each of the major challenges—and conversely, the opportunities—we have outlined, from helping workers and employers tell their stories to supporting policymakers as they experiment with different approaches and respond

to a range of stakes and stakeholders.

Generative AI is poised to rewire how many of us work and earn a living. As the technology advances, however, the future of work will not be determined by technological capacity alone. Whether generative AI lives up to its potential to unlock new possibilities for workers and spread shared prosperity or realizes fears of exacerbating inequality and harm depends on the choices that employers, policymakers, technologists, and civil society make.

Appendix: Methodology and limitations

Methodology

To assess the “exposure” of work to generative AI in the form of specific occupations, we utilized estimates shared by OpenAI relating the predicted ChatGPT-4 exposure level of thousands of the tasks that make up the hundreds of occupations defined by the Department of Labor’s O*NET database.

To generate its task-exposure statistics, OpenAI employed a combination of human annotators and GPT-4 itself to assess the overall exposure of tasks to GPTs, following in the tradition of earlier work quantifying work’s exposure to machine learning. After that, the implied impacts of the thousands of task impacts were aggregated onto 1,016 occupations, following OpenAI’s use of a midrange “exposure” rating statistic that assumes a middling amount of future innovation in complementary applications that use GPT-4 technology.

Finally, to assess the technical feasibility of current or near-term generative AI automating away specific tasks and occupations, we analyzed data from OpenAI estimating the likelihood of ChatGPT-4 completing tasks with no human oversight, as per the autonomous capabilities list in Box 1. Tasks with a high exposure and high likelihood of being completed without human oversight were categorized as “more likely to automate.”

As before, task assessments developed using O*NET were aggregated up and considered at the occupational level to create job-level automation estimates.

Limitations

We believe the data, methods, and statistics employed here yield a plausible way of generating a rough-draft speculation on how generative AI might impact work in coming years. With that said, we note here a few caveats about the data and its limitations.

First, this data does not attempt to project future capability enhancements from next-generation AI models likely to be released (e.g., ChatGPT-5 or 6).

Second, these exposure analyses—based on technical feasibility exercises to gauge the performance requirements of specific tasks—often overstate the potential for job impacts by not accounting for the many practical constraints to real-world adoption of technology in workplaces, from legal and business risks to ethical and privacy concerns and consumer preferences.

Third, task-based analysis in some cases understates the potential impact by missing major technological disruptions that are not easily captured in a task-based approach. For instance, a fashion model appears to have a low exposure to generative AI when considering the job's key tasks such as "apply makeup," "wear costumes," and "pose as directed." But that approach misses the potential for major disruptions to the fashion industry and job risks from retailers using their own AI fashion models.

Fourth, this exercise tells only part of the story: It neither captures the impact of generative AI on important aspects of job quality (as opposed to the quantity of human labor required), nor is it able to capture the emergence of new tasks and occupations that could result from generative AI—undeniably an important impact of earlier waves of automation (e.g., mass-produced motor vehicles creating the need for lots of mechanics).

Finally, a job's "exposure" to generative AI doesn't necessarily mean the job will be lost. Actual demand for skills and occupations will be determined not only by technical feasibility, but also by the specific choices of employers, the market response, and the existence—or lack thereof—of guardrails to protect workers. In some cases, that set of forces will very likely encourage changing work to include using AI as a tool—like a crane, which requires a human operator and allows humans to do much more, says investor [Roy Bahat](#) of Bloomberg Beta, or what Ethan Mollick, an innovation expert at the Wharton School, refers to as "AI as a co-worker" in his recent book, ["Co-Intelligence."](#)

In spite of these limitations, then, there is significant value in the basic exposure statistics presented and discussed here and utilized in the narrative—precisely because our understanding of the risk to livelihoods is so limited and our national conversation about solutions lags behind the technology. The data suggests the overall distribution of potential impacts on a wide range of occupations and workers, not just the handful discussed in popular news coverage or social media to date: coders and writers, customer service agents, and tax preparers. Such broad estimates can and should inform how society understands and responds to the deployment of increasingly powerful AI.

AUTHORS



Molly Kinder Fellow - Brookings Metro  @MollyKinder



Xavier de Souza Briggs Senior Fellow - Brookings Metro
 @xavbriggs



Mark Muro Senior Fellow - Brookings Metro  @markmuro1



Sifan Liu Former Senior Research Analyst - Brookings Metro
 @Sifan_Liu

Acknowledgements and disclosures

Brookings Metro would like to thank the following partners for their generous support of this analysis and our research on AI and work more broadly: Omidyar Network, Microsoft, Google, and NVIDIA.

The authors would like to thank the following colleagues inside and outside of Brookings for important insights and helpful feedback: Bharat Ramamurti, Tom Kochan, Michelle Miller, and Alan Berube. Pamela Mishkin of OpenAI provided data and valuable advice at several stages of our research. We are grateful to the participants in our June convening, who provided tremendously useful perspectives and insights, and to Omidyar Network for providing support for the convening. Special thanks to Glencora Haskins for research support and to Mayu Takeuchi for fact checking.

In addition, the authors wish to thank Leigh Balon, Erin Raftery, Michael Gaynor, and Edward Paisley for their editorial and communications expertise. Thanks to Carie Muscatello for layout and graphic design.
