WASHINGTON CENTER FOR
Equitable Growth
Evidence for a Stronger Economy

Issue Brief: Inequality & Mobility

Navigating the research on AI's impact on work, workers, and the labor market

May 21, 2026 By Chiara Chanoi

www.equitablegrowth.org

601 13th St NW, 12th Floor
Washington, D.C. 20005
(202) 545-6002

WASHINGTON CENTER FOR

Equitable Growth

Evidence for a Stronger Economy

The Washington Center for Equitable Growth is a non-profit research and grantmaking organization dedicated to advancing evidence-backed ideas and policies that promote strong, stable, and broad-based economic growth.

Our fundamental questions have been whether and how economic inequality—in all its forms—affects economic growth and stability, and what policymakers can do about it.

We work to build a strong bridge between academics and policymakers to ensure that research on equitable growth and inequality is relevant, accessible, and informative to the policymaking process. And we have the support and counsel of a steering committee that comprises leading scholars and former government officials. Members have included Melody Barnes, Alan Blinder, Raj Chetty, Janet Currie, Jason Furman, John Podesta, Emmanuel Saez, and Robert Solow.

Since our founding in 2013, we have funded the work of more than 150 scholars and built a broader network through our working papers series, events, and convenings. By supporting research and bringing these scholars together to exchange ideas, we have learned a great deal and advanced a broad range of evidence-based policy approaches to addressing economic inequality and delivering broad-based economic growth to communities and families.

Contents

Key takeaways

- The impacts of AI are difficult to measure. Work, workers, and the labor market are not a monolith, and AI is developing rapidly, which may change its impact over time.
- Across the research, the range of current estimates imply that the overall impact of AI on the labor market will come from the balance of potential benefits and harms felt across tasks, skills, occupations, people, places, and time.
- Research findings on AI and the labor market offer valuable information about AI within specific contexts. Equitable Growth’s new database of AI research allows for an easier synthesis of relevant literature and makes reading and understanding AI research more accessible. Those using our database can sort by several topics and subtopics.
- **What this means for economic growth:** Technological change through AI could have substantial economic impacts, but without deliberate policy interventions, workers might not share in those gains.

Overview

Artificial intelligence will have a largely unpredictable but potentially profound effect on the labor market. Individuals have speculated how AI will affect the workforce since its earliest developments. But public attention has skyrocketed since late 2022, when the proliferation of chatbots made forms of AI highly accessible to workers and firms. As AI capabilities are further advanced by developers and increasingly adopted by firms and growing segments of the public, speculation on AI's ultimate impact on work, workers, and the labor market abound.

Headlines often paint extreme and dissonant pictures of AI's future impact. In 2023, a report by Goldman Sachs predicted that AI would have widespread business applications that would complement jobs and industries, leading to 7 percent growth in Gross Domestic Product over a 10-year period. Dario Amodei, CEO of the AI company Anthropic, suggested in October 2024 that AI may generate significant economic growth in developing nations but later warned of "unusually painful" market disruptions in January 2026.

Workers themselves echo the conflicting feelings of optimism and dread expressed by industry leaders. In 2025, a survey on worker sentiment by the Pew Research Center found that while 52 percent of workers answered "yes" when asked whether they feel worried about AI's future in the workplace, only 36 percent answered "yes" to feeling hopeful about it.

There is little agreement among experts over how AI will impact the labor market. AI capabilities and applications are still emerging, so it is unsurprising that the research on their effects is also still in process. AI has myriad potential uses and thus there are myriad hypotheses about the potential effects AI may have. Academic researchers have hypothesized several ways the technology may influence the labor market. They largely anticipate changes that will be equally capable of helping the economy as hurting it.

One prevalent concern about AI is that it will perform jobs currently done by humans, causing large amounts of labor displacement. But researchers also anticipate that new AI-centric jobs will be created. Similarly, some researchers hypothesize that workers who lack the expertise needed for a job will be able to use AI to work in occupations previously inaccessible to them. Others worry that AI devalues expertise, lowering the returns to education for workers who have invested years into their training.

Table 1 below lists some of the potential effects researchers anticipate AI may have on work, workers, and the economy. This list of potential impacts is neither comprehensive nor certain, but it helps to illustrate the wide range of predicted impacts across the research. It may be possible that seemingly conflicting hypotheses come to pass—for example, the adoption of AI could reduce wages for some workers while increasing wages for others. (See Table 1.)

Table 1

Artificial intelligence has the potential to both harm and benefit the labor market

A sample of academic theories about how AI might affect the labor market

Potential harms of AI – AI might	Potential benefits of AI – AI might
Reduce employment	Increase worker or firm productivity
Reduce wages	Increase wages
Increase income inequality	Decrease income inequality
Reduce returns to education	Reduce barriers to entry
Make some jobs obsolete	Create jobs
Replace workers in decision-making	Empower workers in decision-making
Reduce workers' share of the economic pie	Reduce production costs
Increase workloads	
Introduce more and different kinds of mistakes into work	
Make it harder for workers to accumulate skills	

Source: Equitable Growth AI database.
Note: These effects are neither definitive nor comprehensive, but rather illustrate the range of academic theories about AI's potential impact on the labor market.

WASHINGTON CENTER FOR Equitable Growth

Given the myriad and often conflicting potential impacts of AI's distribution across the labor market, there are many open questions about how this emerging technology will change the nature of work or the economy. Countless questions still need to be studied. How are the impacts of AI concentrated, especially by industry or occupation? Are certain demographics of workers more at risk? Will there be enough work available? Who shares in the productivity gains from AI? How gradually or suddenly will these changes be felt?

Broadly, though, most questions pertaining to AI and the labor market can be broken down into three fundamental questions:

- Which workers are the most vulnerable to the potential negative impacts of AI?
- Who is likely to benefit the most from AI?
- What will be the aggregate effect of AI on work, workers, and the labor market?

This issue brief and the companion database of research papers seek to make reading and understanding AI research easier and more accessible. These resources are meant to help readers understand key terms and frameworks in AI research, discover relevant AI literature, read abbreviated summaries of insightful papers, and feel better-equipped to think critically about this emerging field.

Much like its subject matter, current research examining the influence of AI on work, workers, and the labor market is wide-ranging, uncertain, and conflicting. This is due in part to a wide range of definitions, data sources, frameworks, and motives for the research. Parsing through the literature to determine how to think about AI and the workforce—much less come to any conclusions—can feel daunting. The resources catalogued in this issue brief are not meant to tell readers what to think about AI or AI research but rather how to navigate the evidence.

Readers new to this topic should start here for a broad overview before browsing the database. Let's begin with some basic definitions.

What is AI?

The term “artificial intelligence” broadly refers to the use of computational systems to learn, reason, perceive, make decisions, and solve problems—functions otherwise associated with human intelligence. “Narrow” or “weak” AI operates solely within its training domain and is the only type of AI currently available. Narrow AI encompasses applications from recommendation engines to autonomous vehicles and from speech to image recognition.

At the time of writing, artificial general intelligence, or AGI, and superintelligence—or technologies that would learn, reason, and perform tasks as well as or better than humans—remain unrealized. When we refer to AI in these resources, we refer to narrow forms of AI.

Artificial intelligence can be further (and more usefully) classified based on the technology underpinning its function and capabilities. Machine learning is a subfield of AI and allows systems to learn from data without being explicitly programmed with rules. Via machine learning, AI models can recognize patterns, predict outcomes, and can be retrained to improve their accuracy. A deeper subfield is known as deep learning, which analyzes complex patterns in large datasets and powers generative AI and large language models.

Generative AI produces content, often in response to prompts. The generative AI service DALL-E, for example, creates images based on user requests and specifications. Large language models, or LLMs, are generative AI tools that process and generate human language. OpenAI’s ChatGPT and Anthropic’s Claude are examples of text-based generative AI, as is Gemini (from Alphabet Inc.’s Google unit) and Copilot (from Microsoft Corp.). They all utilize large language modeling. The interface for these services is sometimes called a chatbot, but the term chatbot also includes much simpler technologies.

Large language models represent a significant leap in AI systems' ability to process unstructured language at scale, facilitating natural language interaction between humans and machines. Anyone with an internet-enabled device can use large language models, creating wide access for use in the workplace and at home.

Because large language models are widely available, many headlines, public debates, and academic papers focus on this technology's effect on the labor market, particularly since the release of ChatGPT in November 2022. But generative AI is far from the only type of artificial intelligence that may change how work is performed.

Our approach to navigating the research on AI

For this study, we limit our scope to research that discusses how a type of AI technology may affect the nature, productivity, or distribution of work. At present, we do not include research into how AI may be used for worker surveillance or how AI might affect competition between firms. For a wider discussion on the potential impacts of AI, please see Equitable Growth's AI and Technology resources.

This project casts a wide net on the kinds of AI included in the database but stops short of robotics, computer vision, and other physical AI systems. We include research at both the microeconomic and macroeconomic levels. At the micro level, AI may impact specific components of certain occupations for a subset of workers. Additionally (or alternatively), AI might change how firms grow and structure their workforces. But these influences may also be felt at the national or global level if economies contract or expand with fluctuating employment and productivity levels.

How are experts studying the impact of AI on the labor market?

Frameworks for thinking about AI and labor

At the time of writing, one framework for studying AI has emerged as a popular theory: the Task-Biased Technological Change framework, hereafter referred to as the task-based framework. This theory begins with the idea that jobs can be broken down into the tasks people perform in the course of their jobs. Receptionists, for example, might be abstracted into the tasks they perform, such as greeting customers, transcribing, proofreading, and other administrative duties. According to the task-based framework, technology does not replace (or complement) workers themselves, but rather, it impacts the individual tasks they perform. Grammar software may entirely perform proofreading tasks, and voice-to-text software may improve speed and accuracy in transcription, for example, but neither type of AI helps or replaces the task of greeting and interacting with guests.

The task-based framework is popular in AI and labor research. Researchers use the framework to estimate which parts of a job AI could influence or transform, whether that occupation could experience consequential shifts in supply or demand, and how the labor market could change as a result. But there are some drawbacks to this framework. In particular, the task-based framework does not capture the more holistic elements of an occupation. There are likely spillover effects between tasks, for instance, that cannot be accounted for within the task-based framework. Another major disadvantage is that, so far, the task-based framework can only produce estimates. Estimates should be treated with caution; they offer informed projections about what might happen but cannot replace analyses of what is actually happening.

The task-based framework is not the only way to assess how labor is impacted by advances in AI technology. Other methods include (but are not limited to) measurements of how many firms have adopted AI technology, the use of an AI program as a treatment in a controlled experiment, or qualitative discussions with experts and impacted populations.

Academic papers in Equitable Growth's database can be filtered by types of framework and other characteristics. For a full list and explanation of filters, please see our AI research database.

Data sources for measuring the prevalence of AI in the labor market

Once researchers decide how to conceptually think about AI and labor, they need a way to measure its presence. Researchers can derive information about work, workers, and the labor market from publicly available government data. U.S. Census Bureau data, such as the American Community Survey, quantify workers and their characteristics, while other Census sources, such as the Annual Business Survey, provide insights into firm behaviors.

Below, we discuss other commonly used data sources.

The Occupational Information Network, or O*NET

One dataset from the U.S. Bureau of Labor Statistics is especially popular among studies that utilize a task-based framework is the Occupational Information Network, or O*NET. This is a public dataset that breaks down approximately 1,000 occupations to the abilities, skills, and actions needed to perform them. Crucially, the dataset links a series of unique tasks to each occupation and provides metrics on how frequently they are performed in the course of work.

These tasks can be aggregated to several levels of work activities, which are shared between occupations. For instance, a bartender's task to create drink recipes and a bioinformatic scientist's task to develop data models and databases both fall under the O*NET work activity "Thinking Creatively." Researchers often study how AI could be used during an O*NET task, which allows them to examine how the technology might interact with a particular occupation or activity. (See Figure 1.)

Figure 1

O*Net contains data on the abilities, skills, and actions needed to perform a job

Unique tasks making up an occupation can be aggregated to work activities, which are shared between occupations

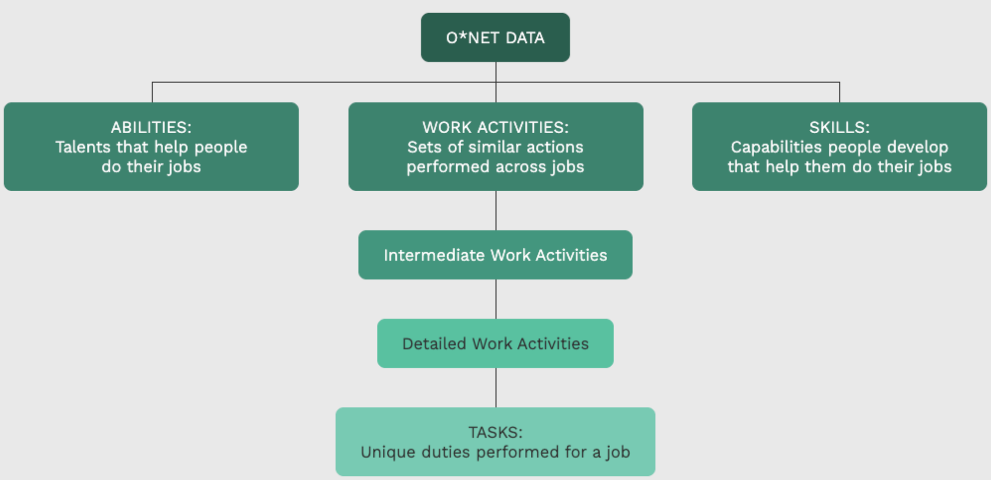

Source: Bureau of Labor Statistics, Occupational Information Network (O*NET).

Washington Center for Equitable Growth

O*NET has become a common dataset to study AI and labor due to its compatibility with the task-based framework, its public availability, and the ease with which it can be linked to other Census Bureau data. But some researchers point out that the dataset includes assumptions on how jobs are broken down into their tasks and is not fully updated every year, making it a less-than-ideal option, given AI's rapid development.

Other AI data sources

Researchers also use other creative data sources relevant to the field of labor and AI. Some research studies compile resumes or job opportunities posted by firms to study changes in job posting language. Others track how workers are highlighting AI-based experience in their resumes. Researchers can use these types of information to track when and how AI is described as a desired skill within the labor market, estimate rates of AI adoption by firms, and associate other worker characteristics with the technology.

Alternatively, some researchers construct proprietary datasets, conduct field surveys, or source qualitative input from focus groups and use other methods of data collection. The firms developing AI technologies also are a common source of private data, such as data from Microsoft's Copilot and Anthropic's Claude. Researchers may draw from these data sources when public alternatives do not exist.

Private data sources have limitations. In the case of the two sources above, for example, a major limitation is that the data only capture the use of a specific large language model by voluntary users. In other words, they capture the patterns of a small subset of people using one service for personal and/or professional use and therefore cannot be used to extrapolate larger trends about AI writ large. A study that employs ChatGPT usage data, for example, can show how users might ask ChatGPT accounting-related questions but cannot draw conclusions about how the work of accountants might be impacted by AI more generally.

Furthermore, when private datasets are not publicly available, study findings cannot be externally replicated. Importantly, private data sources typically describe fee-for-service products and thus may be imbued with different motives than public sources, such as U.S. Census Bureau data.

Academia, advocates, and policymakers are calling for the federal government to collect detailed, real-time, impartial data on AI and its economic impacts. At the time of writing, a few congressional bills have been introduced that would require government statistical agencies to begin scoping such data collection, but no concrete action has been taken yet. Until such data are available, students of AI must judiciously engage with the existing literature and limited empirical evidence of AI's overall impact.

How are experts estimating the effects of AI on the labor market?

AI has the potential to affect work and workers in different ways, depending on the worker's occupation, industry, and demographic characteristics. Worker interactions with AI may replace some occupations, augment others, or simply change how jobs are performed.

The potential for AI to automate, augment, or simply interact with a task, job, or industry is typically referred to as AI exposure. The definition of exposure varies and can change depending on a study's author, data source, or methods. But, in broad terms, AI exposure can be thought of as the potential for AI to interact with a task, worker, occupation, firm, industry, or region.

While many researchers tailor their interpretation of AI exposure to their specific study, a few common measures of exposure are beginning to emerge. Stanford University economist Michael Webb developed an early example of a popular approach. This method can be used to estimate exposure to tasks for several types of technologies, including AI. Webb constructs this dataset by matching patent language to similar O*NET tasks and estimating how new technology might overlap with current human tasks. The occupations with the highest number of tasks that relate to AI's capabilities (as described by its patent language) are predicted to be the most exposed to AI automation. Occupations are then ranked by their exposure percentiles. Webb, for instance, estimates that computer programmers have more AI-automatable tasks than 94 percent of sampled occupations and thus assigns them a score of 94. His paper only considers how exposed tasks might be automated, not augmented.

Economists Edward Felten at Princeton University and Manav Raj and Robert Seamans at NYU Stern School of Business provide another approach to estimating AI exposure, linking a list of AI capabilities to O*NET data. Where Webb's exposure measure focuses on tasks, Felten, Raj, and Seamans look at abilities. The authors match AI capabilities to O*NET abilities and then use this relatedness to estimate AI exposure by occupation. The authors' AI occupational exposure scores are relative. The exposure score for computer programmers, for example, is 1.272, while the exposure score for manicurists and pedicurists is 0.034. The authors interpret the difference in scores to mean that the human abilities needed to program computers overlap with AI capabilities far more than they do for nail care. Notably, these scores do not assess whether an occupation can be supported or replaced by AI.

Later approaches, such as one by Open AI's Tyna Eloundou, Sam Manning, and Pamela Mishkin, alongside Daniel Rock at the University of Pennsylvania's Wharton School, narrows the scope of estimated exposure to a specific type of AI—in this case, large language models. The authors consider either an O*NET task or detailed work activity as "exposed" if the time a human would take to complete it is reduced by 50 percent or more when performed by an large language model or LLM-powered system. They offer a small range of calculated scores per occupation, which represent the lower to upper bound of possible AI exposure. The four authors, for example, estimate at least 90 percent of computer programming tasks can be directly automated or augmented by large language models, giving computer programmers an estimated exposure score of 0.9.

AI company Anthropic's Kunal Handa and others provide their own measurements for use of Anthropic's chatbot, Claude, by linking queries made to Claude with O*NET tasks. In this case, task exposure is calculated as the proportion of user conversations with Claude that concern a certain task. Because each O*NET task is unique to an occupation, individual task-exposure scores are grouped by profession and added together, creating an aggregated occupational exposure score. (The authors caveat that the data include questions asked to Claude for personal use.) The authors calculate the exposure score for computer programmers to be 0.061, which estimates that 6.1 percent of all queries made to Claude can be associated with tasks performed by computer programmers. Estimates include both automation and augmentation in the aggregated occupational exposure score.

The list of exposure measurement methods in Table 2 below illustrates the variety of methods used to define and estimate AI's impact on different occupations. Existing exposure measures are estimates, rather than empirical evidence of where and how AI is changing work. Morgan Frank at the University of Pittsburgh and Yong-Yuel Ahn and Esteban Moro at the Massachusetts Institute of Technology, for example, find that no single exposure estimate is fully accurate, at least when it comes to predicting AI job losses. They conclude that while existing AI exposure scores still have limited predictive power individually, they may each identify a separate pathway through which AI could impact the labor market and that the predictive power of these metrics may come from their combined use.

Researchers are increasingly combining approaches to estimating AI exposure. Anthropic's latest effort to measure "observed exposure," for example, leverages their own usage data with those estimated by Eloundou and her co-authors. The strengths and weaknesses of any exposure estimate, including data sources, should be considered when evaluating how and where to apply a study's insights.
(See Table 2.)

Table 2

Worker exposure to artificial intelligence is estimated many ways

Examples of occupational AI exposure estimates

Exposure definition	Exposure score for computer programmers	Range of values	Score interpretation
Source: Webb (2020)			
How closely an occupation's tasks match AI patent language. Only measure automation exposure.	94	1 to 100	Computer programmers' tasks match AI capabilities more than 94% of other occupations.
Source: Felten, Raj, and Seamans (2021)			
How related AI applications are to the human abilities needed to do a job.	1.272	−2.670 to 1.528	Relative interpretation. The abilities required for computer programming overlap with AI capabilities more than most other occupations.
Source: Eloundou and others (2023)			
The degree to which occupational tasks can be performed faster using a large language model, whether augmented or automated.	90%	0% to 100%	An estimated 90% of tasks performed by computer programmers could be performed faster with LLM-powered tool at minimum.
Source: Handa and others (Anthropic 2025)			
The proportion of real user conversations with Anthropic chatbot, Claude, that relate to a given occupation.	6.1%	0% to 100%	6.1% of queries to Claude are associated with tasks performed by computer programmers.

Source: Michael Webb, "The Impact of Artificial Intelligence on the Labor Market." Working Paper (2019); Edward Felten, Manav Raj, and Robert Seamans, "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses," Strategic Management Journal 42 (12) (2021): 2195–2217; Tyna Eloundou and others, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." Working Paper (2023); Kunal Handa and others, "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations." Working Paper (Anthropic, 2025).

WASHINGTON CENTER FOR Equitable Growth

What are some common themes across the literature on AI and work, workers, and the labor market?

As mentioned earlier (and regularly in media headlines), predictions about AI's effect on the labor market range in both their level of certainty and the types of changes that could occur. At least for now, research mostly finds that AI's effect on the labor market is much more complicated than what both the "doomsayers" and the "optimists" assert. In fact, few scientific papers make predictions about how, or even if, AI will change labor.

The impacts of AI are difficult to measure

There are many reasons why definitive, scientifically backed statements about the impact of AI on the labor market are so far largely unavailable. First and most obviously, AI is a new technology. While researchers have been theorizing about how AI may change the labor market, the long-term effects of technology shocks are typically felt over the course of several decades.

Furthermore, AI is evolving quickly, and the research process can take significant time. OpenAI's release of ChatGPT in November 2022 was a significant step forward in AI commercial accessibility and use. Studies published as recently as 2023 may lack data on how the proliferation of large language models such as ChatGPT might be impacting work and workers.

Even studies that include generative AI have to navigate difficult macroeconomic distortions. Because all studies of generative AI have taken place after 2022, it is incredibly challenging to isolate its effects from contemporary phenomena. The rise of remote work, increased outsourcing and contracting, high interest rates, and other labor market disruptions have likely impacted labor at the same time as generative AI, making it difficult to separate the effects of one from the other.

A second complication is that AI might not just reshape or displace jobs but also is likely create them. A pillar of general purpose technologies, which many consider AI to be, is that they spur innovation, giving rise to new occupations that complement their new capacities. Just one case in point: Individuals studying how electricity affected the labor market shortly after its invention were not able to fully predict the rise of the energy sector and could only quantify the obsolescence of the lamp lighter. It is impossible to design forward-looking studies that wholly and accurately predict how future sectors and occupations might change work.

Work, workers, and the labor market are not a monolith

The effects of AI on work, workers, and the labor market are not likely to be uniform. Employment outcomes are heterogeneous by gender, race, age, and education—and so too may be the effects of AI. Inequalities already present in the economy could be exacerbated or mitigated, depending on where, when, and how AI disrupts or restructures the workforce.

AI is developing rapidly, which may change its impact

The complications inherent to measuring the current impact of AI on the labor market (much less predicting its future effects) may be one source of conflicting narratives and may be why, even in empirical research, findings are not consistent.

Take, for example, the findings of Tania Babina at the University of Maryland (formerly at Columbia Business School) and her co-authors. Their working paper finds that nontech firms that invest in AI technology hire more junior workers while reducing middle-management and senior roles. The authors credit such flattening of workforce hierarchies to the predictive qualities of AI, which empower junior workers to make more decisions and decrease their need for supervision.

In contrast, José Azar, Mireia Gine, and Javier Sanz-Espín at the University of Navarra find that for the firms they studied, as firm exposure to generative AI rose, the share of junior positions shrank, and mid-level positions increased by a comparable amount. The authors theorize that as AI replaces the tasks that junior workers traditionally perform, less-experienced employees are pushed up the organizational ladder. Consequently, the supply for mid-level positions swells with less-credentialed workers, and wages for mid-level roles decrease.

Similarly, in 2025, Harvard University's Seyed Mahdi Hosseini Maasoum and Guy Lichtinger find similar reductions in junior employment, relative to their senior counterparts, in firms that adopted generative AI. They theorize, however, that the restructuring was not the result of automated junior roles, but rather the anticipatory slowed hiring of firms that expected such work to be performed by generative AI in the near future.

Such inconsistent findings underscore the influence that timing and data have on this emerging field of research. Babina and her co-authors note, for example, that their data covers 2010–2018 and thus only examines firms' use of AI before the popularization of generative AI tools such as ChatGPT. Conversely, Azar, Gine, and Sanz-Espín and Hosseini and Lichtinger specifically examine firms' structures following ChatGPT's release and exclusively examine firms' adoption of generative AI. These conflicting results demonstrate the rapid pace at which AI technology develops. Its adoption and impact on the labor market could be equally as dynamic.

Another explanation for incongruous findings is how each study defines AI investments, use, and exposure. While Babina and her co-authors track firm-level AI investments by the composition of employees' AI-related skills, Azar, Gine, and Sanz-Espín use the data provided by Open AI's Eloundou and co-authors to create an occupational exposure index that describes the feasibility of substituting the generative AI large language model GPT-4 for specific work activities. These exposure metrics are not wholly comparable and may contribute to the diverging results. Hosseini and Lichtinger also draw on Eloundou and co-authors' exposure metrics, and though they use a different primary measure than Azar, Gine, and Sanz-Espin, they derive similar results.

The overall effect of AI on labor might be the sum of potential contradictions

At the time of writing, most academic papers, if they make predictions at all, propose that the overall effect of AI on work, workers, or the labor market will be a matter of balance. A few examples of such proposals include:

- The effect of automation on workers may rest partially on how many of their tasks can be automated and their individual skillsets. A moderate amount of automation might lead to higher wages, but two workers in the same job could have different outcomes.
- AI could facilitate a rebuilding of the U.S. middle class if deployed in ways that democratize expertise, but those same lowered barriers to entry could reduce wages even as they expand employment opportunities.
- Technological change through AI could have substantial economic impacts, but without deliberate policy intervention, workers might not share in those gains.

When viewed as a whole, the research currently estimates that the overall impact of AI on the labor market will be a combination of potential benefits and harms felt across tasks, skills, occupations, people, places, and time. How such impacts will be ultimately be felt across the labor market, and on what time scale, cannot be fully predicted. But existing research provides a valuable foundation and highlights where our understanding still needs to grow as AI evolves and integrates into the economy.

Equitable Growth's database of AI literature can help readers navigate the complexity of this nascent field and inform decisions as the evidence base grows.

Navigating the Equitable Growth AI database

The Equitable Growth AI database collects research on how AI technologies may affect work, workers, and the labor market. This resource is meant to help readers find and navigate relevant AI literature to understand key terms and frameworks, browse summaries by topic, or review papers by common data sources such as O*NET. A full list of filters can be found below.

Equitable Growth's database provides context and aims to optimize the research process for policymakers, researchers, and other students of AI. The database will be updated semi-regularly with new papers.

How to use the database: A sample use case

Papers can be sorted by topic of interest using the filters below. For instance, a common question individuals have about AI and the labor market is if the technology is causing job displacement. One could start with the filter "job displacement" to see all the papers that might discuss the topic, or one can narrow down with additional filters, such as by selecting "academia" to limit the search to academic papers, or "government data" to focus on papers that might use U.S. Census Bureau data or U.S. Bureau of Labor Statistics data.

When one reads the relevant papers on job displacement in the first edition of our database, a nuanced and conditional picture of AI-driven job displacement emerges. Across the papers, there is little evidence so far of large net job displacements. U.S. Census Bureau researcher Kathryn Bonney and her co-authors find that among firms using AI, a significant share report using it to replace worker tasks while very few report changes in overall employment levels.

They anticipate both increases and decreases in employment due to AI firm adoption but do not predict a net employment decline. Similarly, Menaka Hampole at the Yale School of Management and her co-authors suggest that AI substitutes for labor at the task level but is offset by productivity gains at the firm level, making its overall employment effect currently negligible.

When researchers do identify potential job displacements, the expected effects vary across workers. Sam J. Manning at the Foundation for American Innovation and Tomas Aguirre at the Center for the Governance of AI hypothesize in their recent paper that workers who might be highly exposed to AI may have protective factors that prepare them for smoother job transitions, such as liquid financial resources. Yueling Huang at the International Monetary Fund found in 2025 that workers may experience disparate outcomes due to AI by age or gender. Huang also suggested that geographically, job displacement may concentrate in regions with higher AI adoption. As discussed earlier, depending on the year, the data, and the model used, some papers may suggest junior-level workers gain workplace advantages, while other papers propose junior-level workers are at higher risk of job displacement.

As with any subject, the frameworks used offer insight into where and how to apply findings and where they may be limited. The findings of Morgan Frank at the University of Pittsburgh, Yong-Yeol Ahn at the University of Virginia, and Esteban Moro at Northeastern University emphasize this for exposure estimates. Different AI exposure estimates identify different channels of AI-driven labor adjustments. In other words, there are multiple ways in which AI might affect job displacement.

As new AI technology is released, and new exposure rates are estimated, new papers will be published, and a clearer understanding of how AI impacts the labor market will emerge.

Description of database filters

Equitable Growth's database of research on AI and work, workers, and the labor market is a tool to navigate the complex, fast-moving field. Updated quarterly, this resource collects the latest research on AI's effects on the labor market. We briefly summarize the research, provide context to help critically evaluate the findings, and allow users to filter across a variety of issues. Users can sort papers by the following topics and subtopics:

Sector

- Academia: Studies done by researchers at universities
  - Published papers have been peer reviewed and are available in an academic journal.
  - Working papers have not been published or peer reviewed and may be edited further before official publication.
- Industry: Studies done by researchers at nonacademic institutions or funded and published by nonacademic entities
  - AI firms often publish their own research.
  - Consulting firms often study AI when performing business-based research.
- Government: Studies published or funded by government entities
  - Government publishing includes studies published by federal or local governments.
  - Government funding may also support other papers.
- Other
  - Think tanks are research institutions, such as the Washington Center for Equitable Growth, that analyze policy-relevant issues.
  - AI-specific research groups are composed of researchers who specifically perform AI research.
  - Independent researchers submit AI research that is unaffiliated with any institution.

Type of study

- Empirical papers use statistical models to interpret data and test hypotheses. Empirical refers to the methods used and does not imply the findings are causal.
- Summary statistics summarize data, often visually, but do not test hypotheses or model relationships.
- Theory/discussions consider how AI might affect the labor market through reasoning, frameworks, and historical examples rather than data analysis.
- Literature reviews synthesize and analyze other research.

Type of data

- Government data: Data provided by government agencies, including:
  - U.S. Census Bureau
  - U.S. Bureau of Labor Statistics
- O*NET data
  - While also a government dataset, O*NET is used frequently enough to merit its own filter.
- Other data
  - Resumes and job postings data are derived from resumes and/or job postings.
  - Private data is exclusive to the researchers performing the study, either because they collected it themselves through experiments or have proprietary access.

Level of study

The level at which AI's impact is measured or estimated, ranging from its potential microeconomic effects to its potential macroeconomic effects.

- Tasks/skills/abilities
- Worker level
- Occupation level
- Firm level
- Industry level
- Regional level
- U.S. level
- Global level

Other topics

- Generative AI
- Large Language Models
- Worker demographics
- Growth/productivity
- Task-based framework
- Job transformation
- Job displacement
- Automation
- Augmentation
- AI adoption
- Worker input
- Discusses the impact of AI
- Uses AI in the study
- Specific policy implications
- Inequality
- Pre-November 2022 (ChatGPT's release)
- New jobs/skills

Glossary

- Artificial intelligence: the use of computational systems to learn, reason, perceive, make decisions, and solve problems.
- Deep learning: a subset of machine learning that analyzes complex patterns and makes predictions and decisions based on data.
- AI exposure: the specific definition of AI exposure can change paper to paper, but broadly can be thought of as the potential for AI to interact with a task, worker, occupation, firm, industry, or region. It does not differentiate between automating or augmenting AI impacts unless specified.
- Generative AI: produces content in response to prompts such as text, images, or code.
- Frameworks: theories, methods, and other tools that help to analyze a subject.
- Large language models: AI systems that use pattern recognition to process and generate unstructured text and speech. Can include chatbots such as ChatGPT.
- Machine learning: a core method used in artificial intelligence. Allows systems to learn, identify patterns, and sometimes make decisions from data without repeated programming.
- The Occupational Information Network, or O*NET: a public dataset provided by the U.S. Department of Labor. Breaks approximately 1,000 occupations down to the abilities, skills, and actions needed to perform them.
- Task-based framework: the idea that jobs can be broken down into the tasks people perform in the course of those jobs. According to the task-based framework, technology does not replace (or complement) workers but rather impacts the individual tasks they perform.

WASHINGTON CENTER FOR

Equitable Growth

Evidence for a Stronger Economy

The Washington Center for Equitable Growth is a non-profit research and grantmaking organization dedicated to advancing evidence-backed ideas and policies that promote strong, stable, and broad-based economic growth.

www.equitablegrowth.org

601 13th St NW, 12th Floor
Washington, D.C. 20005
(202) 545-6002

equitablegrowth.bsky.social

equitablegrowth

washington-center-for-equitable-growth

@equitablegrowth