Challenges of Measuring the "AI Revolution:"
Defining Data to Evaluate AI's Impact on the Labor Market

Erica L. Groshen, Demetra Nightingale, Mark Schweitzer,
Max Rokeach and Benyam Tesfaye[1]

Abstract

This paper reviews how the literature on the labor market impacts of artificial intelligence (AI) technologies can inform the development of a modern data infrastructure that could effectively track the labor market impact of AI adoption and other shocks to the economy. The paper considers the measurement challenges of AI's impacts on the labor market using an example of a previous technological shift. Our review of the current state of measurement of exposure to and impact of AI reveals results which range from negligible effects to quite substantial impacts on particular groups of workers. The uncertain state of the literature is not surprising given the early date of both the implementation and measurement of AI. We conclude that new data and information systems are necessary to resolve these discrepancies and better inform policy. In particular, standards-based employer administrative data could significantly improve the timely measurement of impacts and swifter adjustments of skill and occupational classifications could better identify which workers are most exposed to AI. Our approach could improve tracking, analyzing, and assessing AI's effects on the labor market over time to improve the policy responses.

This is a preliminary draft for discussion purposes. Please do not quote or cite without permission. Comments are invited; please send to Erica Groshen at erica.groshen@gmail.com.

The authors thank The Alfred P. Sloan Foundation and the US Chamber of Commerce Foundation for support for this work.

Conflict of Interest Statement: Beyond being recipients of grant funding from the Sloan Foundation, the authors have no financial or other interests with regard to this research.

The report and recommendations are the authors' alone and do not represent the Sloan Foundation, the Chamber of Commerce Foundation, or authors' affiliate organizations.

[1] Affiliations and emails: Groshen, Cornell University ILR School, erica.groshen@gmail.com; Nightingale, The Urban Institute, dnightingale@urban.org; Schweitzer, Weatherhead School of Management, Case Western Reserve University, mxs2012@case.edu, Rokeach, Department of Economics, Case Western Reserve University, max.rokeach@case.edu; Tesfaye, Department of Economics, Case Western Reserve University, bxt258@case.edu.

Introduction

This paper reviews how the literature on the labor market impacts of artificial intelligence (AI) technologies can inform the development of a modern data infrastructure that could effectively track the impact of AI adoption and other shocks to the economy.

By AI technologies, we mean computer systems designed to perform tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and decision-making. It encompasses various methods and tools, including machine learning and natural language processing, enabling machines to analyze data and make predictions or recommendations.

Experts on the development of AI have stressed that it is likely to have large impacts on the overall economy. For example, a recent report from OpenAI states: "As AI systems become more capable and more embedded across the economy, they may introduce new vulnerabilities alongside new abundance…[and]...place new strain on the institutions and norms that societies rely on to remain stable, secure, and free. (OpenAI (April 2026) p 9.)

Many, but not all, of the concerns about AI focus on the labor market. Yet, the economic estimates of impacts of AI technologies on the labor market remain highly uncertain, even though the topic is one the most professionally studied and popularly discussed topics in recent days. This literature is growing rapidly, in tandem with the development and adoption of AI technologies. The estimates from researchers and corporate leaders vary hugely in magnitude (from close to zero to cataclysmic), time period covered (past, present, or future), who is being affected (demographics, occupations, industries, geographies, etc.), and how they are being affected (employment, wages, layoffs, training, etc.).[2]

The range of estimates is unsurprising because researchers are still exploring how best to measure AI and its impacts. It is also the case that defining the exposure to AI is difficult when it is still evolving, much of our data predate the adoption of newer AI products, and pre-existing known shortcomings of data for measuring shocks to the labor market remain in place.

With this keen interest and divergent results, improving data collection sources and efforts with a focus on labor market adaptations to AI will be critical for the public and for

[2]For media coverage of the diversity of estimated impact, see this *New York Times* article: https://www.nytimes.com/2026/04/03/business/economists-once-dismissed-the-ai-job-threat-but-not-anymore.html?smid=url-share). Papers that discuss divergent adoption patterns and expectations of impact among firm CEOs include Baslandze, et al. (2026) and Yotzov et al. (2026).

policymakers striving to shape appropriate statistical and workforce development policies around AI.

The next section of this paper frames the challenges in measuring AI's labor market impact, and illustrates them with an example of a previous technological shift. Section 2 reviews the current state of measurement of the impact, focusing on highly-cited papers with transparent measures of workers' exposure to AI. Section 3 draws lessons about measurement needs from the current literature on the impact of AI. Section 4 discusses how new data and information systems could improve tracking, analyzing, and using data to assess AI's effects on the labor market over time, with a particular focus on the opportunities offered by standards-based employer administrative data. Section 5 concludes.

1. Challenges in Measuring How Shocks Affect the Labor Market

The imperative to measure AI's impact on the workforce comes from the widely shared perception that it has potential for great benefits for both workers and firms, but also is very likely to create high adjustment costs for both.[3] Business and policy decisions will be key to mitigating the costs without sacrificing the benefits. And those decisions will require accurate, timely and relevant data. As early as the late 1800's, Lord Kelvin championed the view now widely paraphrased as "You can't manage what you can't measure."[4]

A recent statement by OpenAI defines the data imperative as the need to:

> "[I]nvest in clear, real-time measurement of how AI is affecting work, wages, job quality, and sectoral dynamics, using public metrics such as unemployment rates and indicators of regional or industry-specific displacement. These systems should provide policymakers with timely visibility into where disruption is occurring and how severe it is." OpenAI (April 2026) p 7.

a. Why measuring AI impacts is challenging

An important element of the challenge can be seen when we clarify the difference between measuring "exposure to" AI versus the "impact of" the adoption of AI.

- "Exposure" is a current and prospective condition for workers, answering a question such as: "Are AI tools likely to be able to perform an important part of

[3] See https://www.brookings.edu/articles/building-pro-worker-ai/

[4]The exact quote from Lord Kelvin reads: "When you can measure what you are speaking about, and express it in numbers, you know something about it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts advanced to the stage of science." William Thomson, Lord Kelvin (1824-1907)

the tasks now performed by the workers?" It is usually tied directly to workplace uses of AI tools. Exposure metrics are usually expressed as a percentage–either the share of workers who are at risk or the percentage of their work that could be assumed by AI tools.

- Labor market "impact" is a retrospective measure and can extend beyond directly exposed workers, to those affected by others' use of AI or general equilibrium effects of AI adoption on the economy as a whole.[5] Impact metrics can include gains or losses in employment, layoffs, wages, working conditions, mobility, skill requirements, and broad measures of well-being.

Both are of interest to inform policy planning and research. Necessarily in the early stages of an economic shock, exposure metrics (some quite speculative) are more available than impact measures. These exposure measures are often used to infer or project direct effects of AI adoption while the response of firms and workers to the shock are still unknown.

However, the transmission from exposure to impact is far from simple. For one thing, the development process is ongoing. Early exposure metrics may not prove accurate in the long run; for example, they may be too inclusive or not inclusive enough. And, as noted some impacts will extend far beyond exposed workers. Impacts will evolve over time, potentially even changing from positive to negative or vice versa as the technology develops new applications and becomes more productive. Furthermore, impacts on exposed workers and the economy will reflect implementation opportunities, the demand for AI-affected products, and decisions, including those influenced by public policy, business strategies, and worker voice. Thus, the active debate about whether AI will predominantly replace workers (causing layoffs and unemployment) versus augmenting their skills (boosting productivity and economic growth) is largely about the transmission to impact, rather than about exposure to AI itself.[6]

As AI development and adoption proceed, the validity of current exposure and impact measures must be tested frequently and the measures adjusted as necessary. Furthermore, full assessments of the impact of AI must extend to indirect and general equilibrium effects, and ideally be delivered with short lags, if they are to accurately and reliably inform exposure estimates, policy, and research.

b. ATMs: An instructive tale for measurement of AI

Certain features of AI, as a large moving target, particularly complicate measurement. These challenges include the following:

[5] General equilibrium effects include changes in structure of the economy, such as labor and capital allocation, affecting total output (GDP), wages, and capital returns. In the case of AI, this includes the broadest and most optimistic and apocalyptic forecasts of an AI-based future.
[6] For more discussion, see Acemoglu et al. (February 2026).

1) AI is a general purpose technology, with widespread uses in possibly all industries.
2) Development and adoption of AI is uneven and ongoing.
3) AI inputs in production will continue to change in both quantity and quality.
4) AI use could restructure many features of the economy.

The challenges are substantial and subtle, but hardly unprecedented. A brief review of a relatively-straightforward, historical example of how a new technology affected work in one industry can illustrate these four measurement challenges.

In 1969, the introduction of automated teller machines (ATMs) started with a clear-cut case of occupational exposure. ATMs performed many of the tasks of human tellers. Thus, many observers in the 1970s expected employment in banking to fall because tellers' jobs were clearly strongly exposed to the introduction of ATMS and constituted a large share of bank employment. Instead, banking (depository institutions) jobs continued to increase long after the introduction of ATMs (see graph in Appendix A.)

What accounts for the unexpected strength of banking jobs after the introduction of labor-saving ATMS? Table 1 (first column) lists the challenges mentioned above. The second column summarizes how these challenges were manifested with ATMs, as described in studies decades after the introduction of ATMs (Fung, 2006; Frame and White, 2004 and January 2014). The reason is that installation of ATMs began in one of the largest, most sophisticated banks, and spread first to other large banks. These banks then grew their local branches rapidly. Now there are over 500,000 ATMs in the US and even the smallest credit unions have ATM access. Only in the last two decades do we see the leveling off and some decline in the number of banking jobs and the digitization (of which ATMs were a key part) changed the nature of the banking industry.

The history of ATMs shows that initial measures of exposure would have been very misleading about impacts in the short and medium run. Even now, in the longer run, when jobs for bankers and tellers are indeed shrinking, the pace of job loss is gradual enough to be absorbed by the labor market with little alarm.

Table 1: General Challenges in Measuring Labor Market Impacts of New Technology: Comparing ATMs/Banking to AI

General Challenge	How the Challenge Looked with ATMs	AI Measurement Challenge	AI Measurement Implications
AI, as a general technology, will have widespread applications.	At the outset, ATMs looked narrowly focused on teller services. Soon, ATMs were part of "digitization" that changed the nature of banking.	AI will alter the tasks of many/most workers so counter-factuals are hard to find.	Data must span occupations and, ideally, link to employer implementation of AI.
Adoption of AI is developing rapidly and unevenly.	1970s–2000s: Starting with sophisticated banks, ATMs raised productivity and reduced the number of tellers per branch, but expanded total teller employment via branch proliferation. 2010s–present: Digital banking and automation are reducing both branch density and teller jobs, with tellers performing more advisory and sales tasks.	Short run adoption patterns and impacts will be different from eventual long-run impacts.	Measurement of AI impacts must be designed to adjust with adoption patterns that will grow and evolve.
AI inputs in production will change in both quantity and quality.	ATMs become more versatile, capable of performing more sophisticated services, even as the price of ATMs fell.	Defining and measuring units when quality changes is harder than doing so for physical inputs.	New metrics will be needed to measure use of AI where pricing makes the value intangible.
AI use can potentially restructure many features of the economy.	Consumer expectations for access to their accounts anywhere helped eliminate restrictions on interstate banking. Finance is very different now.	General equilibrium Impacts can spread far beyond sectors touched by AI use or production.	Growth in tangentially affected sectors must be measured to fully gauge impacts.

For comparison with ATMs, the third column of Table 1 describes how each measurement challenge looks for the introduction of AI. The fourth column describes implications of those challenges for data needs. For each challenge, AI introduction will be far more complex than was the case for ATMs, affecting more tools and industries, as it is integrated into multiple processes and products throughout the economy.

In short, an ideal strategy to examine the impacts of AI on the labor market would look across well-defined groups of workers (across industries), track the tasks they perform

on the job and their career paths as they are hired, fired, or change jobs; note their compensation, and link all the worker and job data to information on their employers' AI strategies within given industries and regions. Researchers are currently far from being able to measure impacts at this level of disaggregation in a timely way and will not be able to move very far in this direction if constrained by the currently available employment and business data.

2. Current State of Measurement of AI in the Labor Market

With this framing in mind, we now compare how some key studies identify AI exposed workers, impacts, and their conclusions. The studies offer both exposure and impact metrics; most use some measure of "exposure to AI" and then apply their results in different ways. They also show evidence of the measurement challenges discussed above, although none specifically discuss the challenges.

a. Exposure: Where AI Investments Meet the Labor Market

Mirroring rising public concern, academic research on the impacts of AI on the labor market has grown swiftly. The number of articles indexed on Google scholar on the topic in 2020 numbered 51, but by 2025 the number of annual papers on Google scholar reached 496.[7] As Table 1 above showed, the most basic challenge is knowing which workers are being directly augmented or replaced by AI tools. This is no easy task as it relies both on measuring how AI is implemented and the employment of workers directly exposed to AI.

While there is a long literature focused on the effects of automation on labor markets (see Acemoglu and Restrepo (2019), Frey and Osborne (2013), and Autor et al. (2003)), the earliest paper we have found that provides a specific measurement of AI effects on the labor market is Webb (2019, updated 2020). Webb simplifies the tasks included in O*NET (Occupational Information Network) to verb-noun pairs and then searches for those pairs in AI-focused patents. The estimated exposures to AI at the task level are merged back into occupations, and then the effects are measured using occupation data. This general approach of measuring exposure at the occupation level and then contrasting economic outcomes between "more" and "less" affected workers is the measurement strategy (largely) followed to date. Webb concludes that AI is primarily directed at high-skilled tasks and likely to reduce the wage differential between 90th and 10th percentile workers.

Another early approach (Felten et al. 2021) used Electronic Freedom Foundation (EFF) assessments of AI model progress in computer tasks which are subsequently

[7] Based on a Google Scholar search on "AI Labor market Effects "Labor market effects" "artificial intelligence"" conducted on 4/14/2026.

connected to human tasks.[8] Unfortunately, the underlying EFF project was halted in 2018, so the Felten et al., measures should be viewed as based on 2018 AI capabilities.

Recently, measures of AI exposure have been developed using information from major Large Language Models (LLM). Anthropic publishes an economic Index which uses anonymized conversation data with Anthropic's Claude tool to infer which human skills (from O*NET) are being augmented or automated (Anthropic, 2025). This information is then aggregated into occupational exposure measures using the O*NET tasks list. Key steps were also performed by human experts for verification of the Claude results in the papers which developed the methodology (Handa et al., 2025). The Anthropic team concluded that "roughly 36% of jobs had some use of AI for at least 25% of their tasks."

Open AI's model was also applied to construct an AI exposure index by assessing whether time on a task can be reduced by using ChatGPT4 by at least 50% (Eloundou et al., 2023). This is again aggregated into occupations using the O*NET task lists. They also allow for other complementary software to be used to complete the task which probably yields a pattern more frequently applied by AI users. Overall, this estimate "indicates that up to 49% of workers could have half or more of their tasks exposed to LLMs."

Tomlinson et al. (2025) measures AI applicability (their term for exposure) by analyzing how users employ the Microsoft AI tool's Copilot in their general work. This is similar to the other work examining Anthropic and OpenAI tools, but Tomlinson et al. focus on broader skills (ONET Intermediate Work Skills [IWS]) that workers are using to supplement their activities. By focusing on IWS instead of occupation-specific tasks their measure aims at potential AI assistance regardless of the context (including the occupation) of the particular users.

A recent working paper, Hampole et al. (2025), analyzes the adoption of AI at the firm level. They are able to measure the use of "AI Integrators" in firms to show investment in AI and measure the impacts on occupational level associated with exposed tasks and the extent to which exposure is concentrated in narrow tasks allowing workers to substitute into other tasks within their occupation. Hampole et al. (2025) are also able to estimate productivity growth at the firm level. They conclude that "overall employment effects are modest, as reduced demand in exposed occupations is offset by productivity-driven increases in labor demand at AI-adopting firms." The ability to analyze change at the firm level clarifies important offsetting changes.[9]

[8] The connection from AI to human tasks was initially by expert assessments of computer science students and then later using Amazon's Mechanical Turk gigwork system for larger-scale assessments.
[9] Babina, Fedyk, He, and Hodson (2023) also analyze workforce composition changes following firms' AI investments, as determined by their hiring of AI-skill, to show substantial shifts in the general workforce composition, but they do not measure occupational exposure.

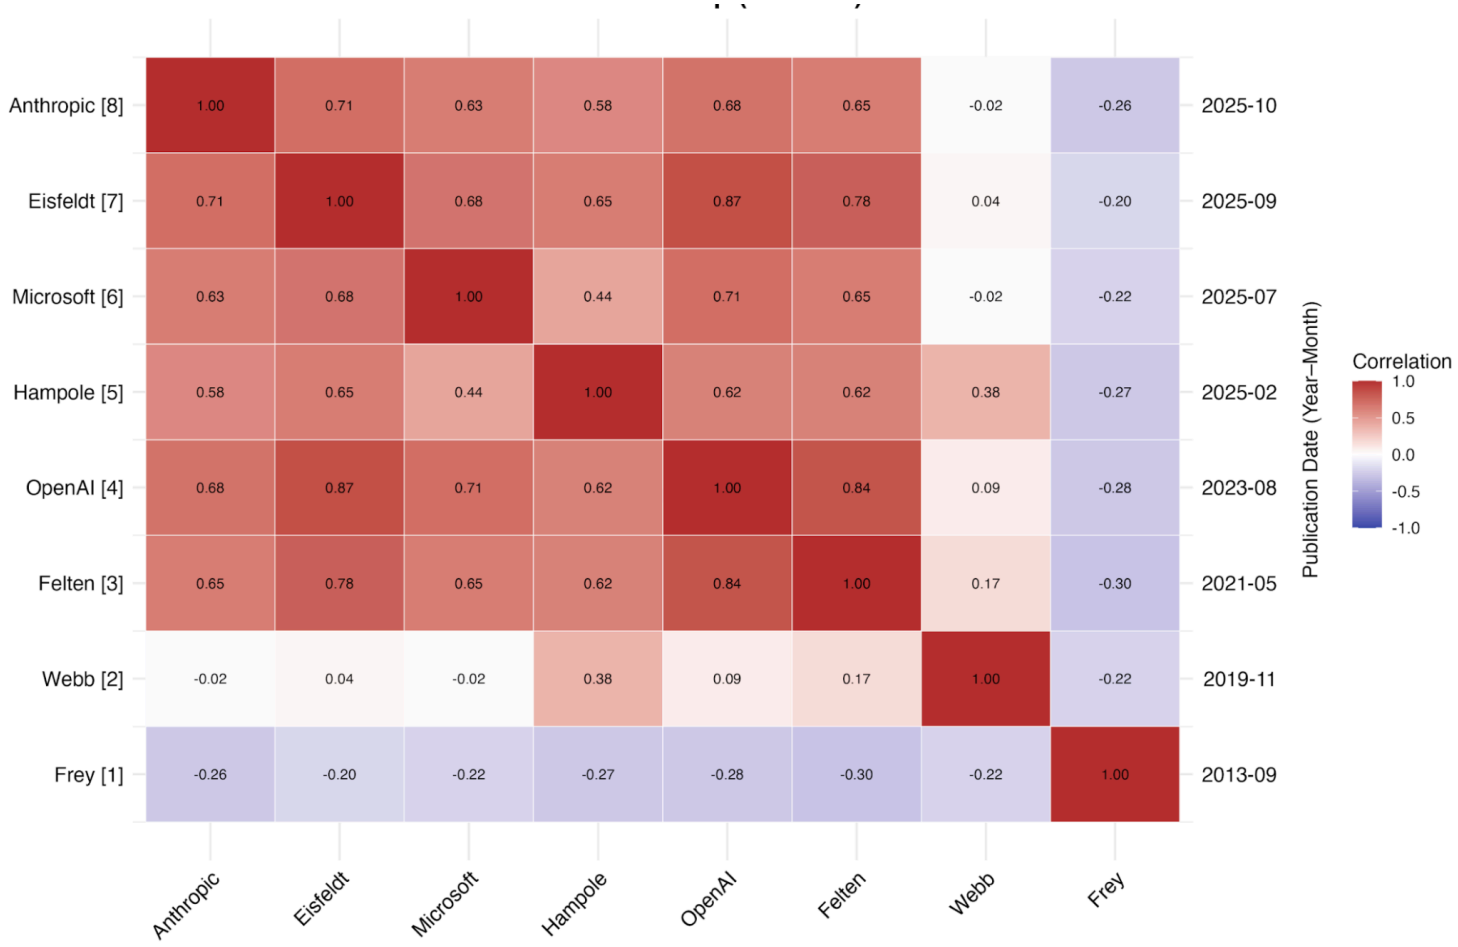
Given the range of approaches taken in these studies, we wanted to assess how similar the findings on exposure measures are.

Figure 1 shows the Pearson correlation for the standardized AI exposure measures used in key studies, whose various methods are described in Appendix E. Positive correlations are shown in red and negative in purple. We note the following:

- There is a clear break between the two oldest and the newer estimates. The oldest papers (Frey automation measure and Webb machine learning metric) cover earlier periods and do not focus specifically on AI tools that appear most relevant today. In most cases, these early exposure measures are negatively correlated with those developed later.
- Correlations among the most recent estimates (from Felten to Anthropic) correlations are all positive and generally above 0.5.
- Limiting comparisons to measures derived from **usage** of recent LLM tools (Anthropic, OpenAI, and Microsoft), the correlations range from 0.63 to 0.71. While the measures are correlated these values are low enough to imply that there are important deviations among the exposure measures.
- The two studies that use **company data** to infer usage (Eisfeldt and Hampole et al.) are imperfectly also correlated with each other at 0.64–indeed, no more correlated with each other than they are with the usage-based measures.
- Felten, which uses **EFF projections** that focus on machine learning (an older AI technology), is still similarly correlated (0.62 to 0.84) with the newer measurement approaches.

DRAFT

Figure 1. Pearson correlations of AI exposure measures



[1] Frey: Frey & Osborne (2017, originally a 2013 Oxford working paper) estimate automation risk by combining O*NET occupation task and skill data with expert labeling, then apply a machine learning model to assign each occupation a probability of computerization.

[2] Webb: Webb (2019, working paper) measures AI exposure by linking USPTO patent text data to O*NET task descriptions, capturing where AI innovation overlaps with what people do on the job.

[3] Felten: Felten, Raj & Seamans (2021, Strategic Management Journal) construct their AIOE measure by combining Electronic Frontier Foundation (EFF) AI progress benchmark data with O*NET occupational abilities, mapping AI capabilities to workplace skills.

[4] OpenAI (Eloundou): Eloundou et al. (2023, OpenAI working paper “GPTs are GPTs”) use human annotators and GPT-4 applied to O*NET task data to estimate how much generative AI can reduce the time needed to complete tasks across occupations.

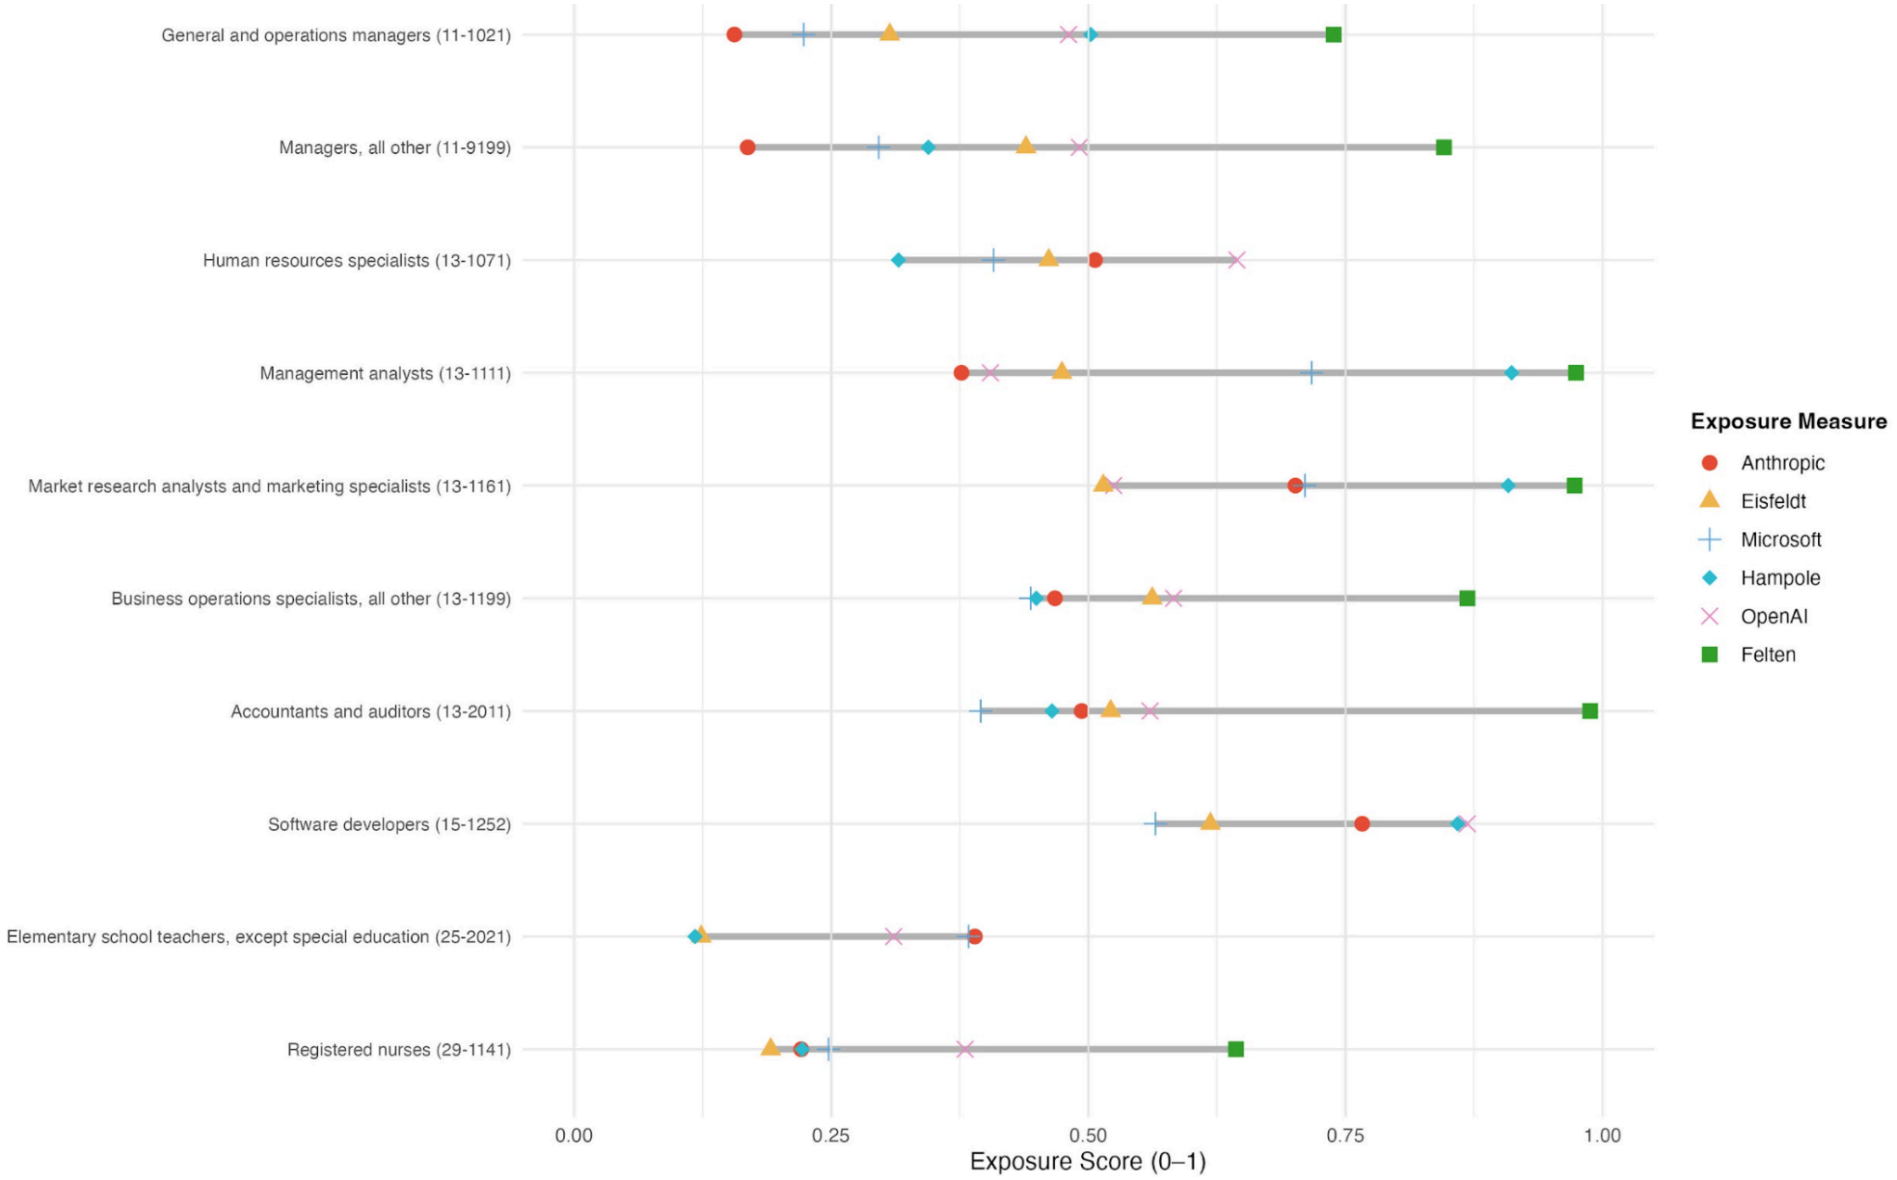
[5] Hampole: Hampole (working paper, ~2023–2024) builds exposure scores using O*NET-based task mappings combined with AI capability scoring, estimating how much of each occupation’s work AI can perform or assist with.

[6] Microsoft (Tomlinson): Tomlinson et al. (2024–2025, Microsoft Research “Working with AI”) use real-world Copilot usage data combined with task-level work activity mappings (O*NET-style) to measure how AI is actually used across occupations.

[7] Eisfeldt: Eisfeldt et al. (2023–2024, working paper) construct exposure using O*NET task data paired with GPT-based task classification, assigning exposure based on whether generative AI can complete or augment each task.

[8] Anthropic: Anthropic’s Economic Index (2024–2025), as used by the Yale Budget Lab (Gimbel, 2025), combines real-world Claude usage data with occupation task mappings (O*NET-style) to estimate observed AI usage across jobs.

Figure 2: AI Exposure range by occupation, for occupations typically requiring a college degree



But how much agreement is there actually? Even though the Pearson correlations show general agreement on occupational exposure among the more recent papers, examining measures for specific occupations reveals substantial differences.

Figure 2 shows the diversity of exposure estimates for ten occupations typically requiring a college degree and identified by BLS as likely to grow fastest from 2024 to 2034 (BLS 2026a). In the figure, we standardize estimates of AI exposure as follows: each measure's most affected occupation(s) receives a 1 and the least affected occupation receives a 0. This approach loses any sense of the scale difference of the different exposure measures but is appropriate for comparisons of the least versus most affected occupations that are used to estimate expected labor market effects. We limit the measures included in Figure 2 to those showing positive correlations in Table 1, which excludes Webb and Frey. More details on the construction of Figure 2 and results for other occupation groups (BLS 2026 b,c) are shown in Appendix D.

We find the following:

- The range of estimates is generally around half of the available range from 0 to 1 for this group expanding occupations.
- Felten, the earliest and most commonly-used exposure measure, tends to show higher values over this set of mostly white-collar, heavily administrative occupations, but it is relatively low for Software Developers.
- The three exposure measures that exploit current tool usage (Anthropic, OpenAI, and Microsoft) are relatively low in many of the occupations that rank highly by Felton's measure. This reflects relatively low usage figures for these occupations, while usage is relatively high for software developers.
- Even limiting the comparison to these three tool-related approaches reveals important differences. For example, for Software Developers, OpenAI has the highest exposure estimate of the alternatives, while Microsoft has the lowest.
- Finally, the two measures that rely on investment patterns (Eisfeldt and Hampole et al.) also don't agree much on exposure for the occupations shown here. Eisfeldt shows relatively little exposure for all these occupations while Hampole highlights relatively large exposure estimates for the two "analyst" occupations on the list and Software Developers.

Despite some broad patterns in the exposure estimates, our overall takeaway is that researchers' understanding of possible AI exposure is still in its early days. It is far from clear where we should expect large effects; we have no way to judge which particular approach to measuring current AI exposure will prove the most reliable. Until time and further research provide more guidance, we should continue to explore this measurement challenge actively and not rely on a single measurement.

b. Impact: Measuring AI-Induced Changes in the Labor Market

If identifying AI exposure is important and difficult, so, too, is identifying the broader labor market impacts of AI (or any large technological shock). And as with exposure analysis, the approach taken and data used in a study can alter the conclusions.

A key difference in the AI impact studies is how the analysts identify a counterfactual to which the affected population can be compared. This task is not easy because general technologies affect most groups and the general equilibrium impacts are likely to be large and widespread if AI effects are substantial. With large shocks to the labor market, economists often have little choice but to focus on the contrast between "more" and "less" affected workers, despite the recognition that general equilibrium effects will potentially alter all workers' employment, hours, and wages.

Table 2 reviews selected papers' summaries of effects and key measurement features of the labor market impacts of AI, including the exposure measure, time period of the analysis, and the characteristics of the focus populations. The papers are roughly sorted from the largest stated impacts to the smallest. Our review identified a few patterns:

- On the upper end of the estimates, Brynjolfsson et al. (2025) find that early-career workers (ages 22–25) in AI-exposed occupations experienced 16% relative employment declines, while employment for experienced workers remained stable.
- On the lower end, Gimbel et al. (2025) conclude "AI's impact on the labor market that emerges from our data is one that largely reflects stability, not major disruption at an economy-wide level." Both use very recent data and Anthropic and OpenAI exposure measures but distinct approaches to impacts on the labor market. Gimbel et al. (2025) look for change in the aggregate occupational shares in Current Population Survey data, while Brynjolfsson, [missing text]
- The largest reported effects are generally in small groups that arguably might be the leading edge of change. Chandar and Chen (2025) use ADP payroll data to examine exposed occupations by age category, a narrower slice of the labor market than is typically possible in official data. Younger workers are more active margin of adjustment, so the results are plausibly very different for younger workers while still showing little change in the occupational composition of the overall workforce. Hui, Reshef and Zhou (2024) also focus on a relatively narrow slice of the workforce, freelance writers and designers using an app to find work and again find more dramatic effects within this group.

Many data sources have been used to measure the labor market impacts in these papers.

- One set of studies rely on public data.

    - Early on, Webb (2020) simulates the effects based on prior technology adoptions on older government data (American Community Survey and decennial censuses), because the expected effects were presumably too small to measure at that point.
    - Johnston and Makridis (2025), Georgieff and Hyee (2021), and Gimbel et al. (2025) use relatively conventional government sources for the labor market data and find limited effects. By contrast, Johnston and Makridis (2025) find some sector level effects when occupational impacts are converted into sectoral impacts.
- Other studies draw on newer private sector data sources that can focus on the early impacts of a rapidly changing technology.
    - ADP data has the advantage of allowing age cuts in establishment data, but it is otherwise similar to official sources.
    - Burning Glass data reveal changes in job listings, which both can show more use of AI-specific skills (as in Autor et al. (2022)) and explore an active margin of companies' responses to a new technology.
    - The skills requested in active job listings provide relative to static skills specified in Revelio data includes both resumes and listings and is able to observe changes at the firm level, which enable Hampole et al. (2025) to identify detailed investments and labor market responses at the firm level over the period from 2014 to 2023. They find larger effects in firms that are investing in AI (based on hiring AI integrators) but these effects are more limited at the occupation level, because firms that are investing are also growing more quickly.
    - As each of these approaches gets closer to an adjustment margin there is more evidence of change.The time periods vary substantially. So, it is not surprising that studies based on earlier time periods tend to find more limited effects given the rapid development of AI tools particularly after 2023.

Overall, the results of AI impact studies thus far span from essentially zero to strong effects. We believe that this very large range is actually reasonable for *early* measures of impacts of a *rapidly* developing technology, especially with our current national data infrastructure. It is also clear that new data sources are able to provide information on important margins of adjustment to what is currently observable in official sources.

Just as with the exposure measures, there is room to learn from existing measures while also striving for a more targeted measurement approach.

DRAFT

Table 2: Literature on Estimated Labor Markets Impacts

Study	AI Exposure Measure	Labor Market Data	Time Period of Analysis	Authors' Characterization of Effects	Approximate Share of Focus Group
Brynjolfsson, Chandar & Chen (2025)	OpenAI, Anthropic	ADP microdata	2022-2025	Early-career workers in the most exposed occupations see about a 16% relative employment decline vs. least exposed	**Narrow** (Early career, exposed) Software Developers : 0.78% Customer Service : 1.96%
Hui, Reshef & Zhou (2024)	Just writers and designers	UpWork online app	Jan 2022 - April 2023	More exposed freelancers see 2% fewer monthly jobs and 5.2% lower monthly earnings after ChatGPT's release.	**Narrow** Writing- and design-focused freelancers (out of 4.6 million on app)
Johnston & Makridis (2025)	OpenAI	BLS QCEW	2021 - 2025	Sectors with more automatable tasks show declines in wage bills and headcount; sectors with more complementary tasks show gains.	**Broad** Exposed occupations imply industry employment and income
Hampole, Papanikolaou, Schmidt & Seegmiller (2025)	Hampole	Revelio Labs	2014-2023	Employment effects are modest: exposed occupations shrink within adopting firms, but this is offset by productivity-driven increases in overall labor demand.	**Broad** all occupations, but oversamples white-collar & limited to publicly traded firms
Acemoglu, Autor, Hazell & Restrepo (2022)	Felten, Webb, Brynjolfsson	Burning Glass (job listings)	2007 and 2010-2018	Significant decline in some of the skills and a significant emergence of new skills. No detected relationship between AI exposure and overall employment or wages.	**Medium** all occupations, but focused on the adjustment margin
Webb (2020)	Webb	Censuses + ACS	AI Patents to 2019 Labor to 2018	Simulated effects imply wage compression (high-skill tasks most exposed); employment effects are not yet directly observed in the data.	**Broad** but based on prior robot and software shocks
Georgieff & Hyee (2021)	Felten approach, international	Burning Glass, Int'l Labor Surveys	2012-2019	No clear relationship between AI exposure and employment growth across all occupation	**Broad** 36 Broad Occupations across 23 countries
Gimbel, Kinder, Kendall & Lee (2025)	OpenAI & Anthropic	CPS	2022 - 2025	AI exposure measures have no signs of being related to changes in employment	**Broad** looks for composition change in full labor market

DRAFT

3. How Well Do Current Data Meet the AI Measurement Challenges?

This section draws lessons about measurement from the current literature on the impact of AI. We consider exposure and impact measures separately and ask about the usefulness of available official public data and private data for meeting measurement challenges discussed in Section 1.

a. Exposure: Mapping AI-related labor activity to occupations

AI exposure analysis is particularly complicated by the first two measurement challenges, namely (1) AI is a general purpose, widespread technology and (2) adoption of the technology is rapid, with uneven exposure and adoption across industries, firms and occupations. This combination (widespread, rapid, and uneven adoption) clearly complicates mapping AI-related labor activity to occupations, so that the evolution can be tracked.

Existing federal statistics were not designed to overcome these challenges, even though the federal statistical system provides the most comprehensive, consistent and reliable labor market data available. One reason is that none were designed to track occupational evolution. A related reason is that the most widely used occupational classification system is updated infrequently. Finally, most statistical programs lack the capacity to add new questions or data fields to collect information on new topics.

Generalizing broadly, the existing public labor market datasets update a useful but typically static set of variables. They offer either comprehensive employment data with little or no occupational detail (QCEW, for example) or highly detailed occupations but lack broad enough coverage for analysis at that level (CPS) or not updated in a timely way (OEWS). The details behind this generalization are shown in Appendix C, which lists eight major sources of data and their key features.

Despite their limitations, the virtues of public data are strong enough that all the impact studies cited in Table 2 use public data in some way, usually combined with other information in various ways, including weighting, matching, modeling, and validation.

AI exposure metrics in studies that use public data must be imported from another source and tied either to AI investments by firms or industries or AI task capabilities. For AI task capabilities, exposure metrics are most often based on duties or activities or tasks associated with occupations.

Particularly for the latter (occupation) approach, most AI exposure studies rely heavily on Standard Occupational Classification (SOC) codes developed by a BLS-chaired federal statistical system committee and often on additional information built into O*NET (Occupational Information Network), developed under sponsorship of the US Department of Labor. SOC documentation designates occupational groups, providing

descriptions, job titles and duties. O*NET provides additional detail including tasks, a range of other occupational characteristics and some additional classifications.[10]

While SOCs have the advantages of statistical rigor, wide use, and continuity, their updating process is designed to move slowly. SOCs were last updated in 2018 and the next update is scheduled for 2028.[11] These lags can pose problems for tracking fast moving occupational changes anticipated as a consequence of AI adoption. AI is poised to change both tasks and how they are organized or associated with one or more occupations, as well as affecting the labor demand for various occupations or new occupations. The major classification scheme (SOC and O*NET) was designed without any expectation of AI tools, so their applicability to occupations being substantially augmented by AI tools is gradually being incorporated to consider the role of the tasks in the occupation, while some newer jobs, such as “AI Integrator," are challenging to place in the dated coding system.

O*NET updates a portion of occupational detail components each year to reflect any BLS SOC updates and other information it gathers. For example, some content (e.g., tools, technologies) in the “interest profile” component of O*NET (originally included for career guidance purposes) is updated annually, using machine learning and AI tools to merge information from job listing databases and other private sources. Beginning with the updates in 2025, profiles for about a dozen O*NET occupations indicate “AI Expert” under the workstyles component (O*NET Resource Center, 2026).

For this key correspondence analysis, while the reliance on O*NET to map AI activity to the approximately 950 occupational categories in O*NET should make study results more similar, there are three possible levels of work tasks in O*NET that may explain variation in findings discussed in the previous section:

- Detailed Work Activities (DWAs), with about 20,000 task statements;
- Intermediate Work Activities (IWAs), with about 2,000 activities generally tied to one job family; and
- Generalized Work Activities (GWAs), with 41 broad activities that span many occupations.

[10] For more information on O*NET, see https://www.onetonline.org/.

[11] Previous updates occurred in 2000 and 2010. The 2018 Standard Occupational Classification (SOC) system is a federal statistical standard used by federal agencies to classify workers into occupational categories for the purpose of collecting, calculating, or disseminating data. All workers are classified into one of 867 detailed occupations according to their occupational definition. To facilitate classification, detailed occupations are combined to form 459 broad occupations, 98 minor groups, and 23 major groups. Detailed occupations in the SOC with similar job duties, and in some cases skills, education, and/or training, are grouped together. BLS chairs the federal statistical system committee that updates SOC codes. See https://www.bls.gov/soc/ for more information about the SOC process.

Researchers have chosen different levels of work activities to identify affected occupations. At this point it is not clear which level of aggregation is most appropriate. While it is attractive to identify activities that define an occupation in the DWAs, that might neglect efficiencies that could alter employment levels based on less time spent on secondary activities in many occupations that might be better picked up by O*NET's GWAs.

Furthermore, an ideal task classification scheme is consistent, meaningful, and up-to-date. In practice, these are challenging goals. Regularly updating classifications of occupation, tasks and skills will be important to identifying where AI impacts are occurring even if the retrospective data compatibility is disturbed. Indeed, the classification updates themselves can provide important information for stakeholders and decision makers.

Private-sector data offers a critical additional source of information about AI exposure. Compared to official data, private sector labor market data sources have higher frequency and timeliness. Many online job listing and labor exchange systems have been developed, some providing almost real-time information about open jobs, often also including job seeker profiles or resumes and online links for submitting job applications all very useful for studying the evolution of work.

Many researchers in the US and other countries use data scraping and machine learning methods to analyze data from one or more job listing systems, usually to address particular questions, for instance, about a specific occupation or type of worker or job requirements or experience. In the examples listed in Table 2, Hampole et al. (2022) and Georgieff and Hyee (2021) use job listing data from Burning Glass and Papanikolaou, while Schmidt and Seegmiller (2025) use web-scraped data from Revelio.

Private sector data has its own set of limitations. First is representativeness. Each platform has special features that attract a subset of the labor market. Thus, information on the non-represented sectors may be omitted or less reliable. This is of particular concern because, currently, most online platforms (including Burning Glass, used by Hampole et al. (2022) and Georgieff and Hyee (2021)) disproportionately cover job listings for occupations requiring a college education, even though most US workers do not have a college degree.

Additional challenges in nonpublic data include lack of consistent history and structure, lack of information specifically targeted to address AI-related questions, and inaccessibility to researchers. There are also technical variations, that may not be transparent, in how analysts aggregate, link, scrape or merge data from different sources.

DRAFT

In sum, these private sector data are also highly relevant for employment research on particular topics (e.g., postings with AI terminology) and for job seekers, but they do not replace official government statistical labor market systems subject to stricter standardization and methodological rigor.

b. Labor Market Impact: Shifting AI usage

The second two AI measurement challenges (i.e., that technology inputs can lead to both quantitative and qualitative changes and technology could restructure some features of the economy as a whole) are particularly relevant for measuring AI impact on the labor market and economy, in part because the use of AI is shifting across industries throughout the economy, and shifting rapidly. That necessarily implies that AI impacts continue to shift. Academic analysis tends, thus far, to use static measures of expected impact for causal determination. In a general technology shock, we should expect wide-spread employment and wage adjustments in exposed occupations, along with related adjustments in occupations that seem distant from the shock. Of the approaches, Hampole et al. (2025) is the closest to measuring broader responses to AI investment, but they have a narrow measure of AI investment that could be expanded and a limited view of dynamics within the firms.

Today's estimates of AI impacts for specific categories of workers span the range from negligible to large effects. However, all that may change, because the scale of current and planned AI investments suggest the potential for broader and far larger changes in the future. While data limitations certainly contribute to today's divergent results, so does the fact that AI impacts will likely vary over time and in different realms. Thus, an inclusive and flexible approach to analysis will be needed.

Existing approaches are tied to static lists of occupations and information on worker tasks that are only updated slowly. This problem is fundamental and argues for an approach that also evolves with the developments of AI. In terms of direct exposure, authors have taken various approaches to this, including measuring AI tasks and linking to human tasks and by using information on employers' investment decisions and subsequent job demands. Both approaches should be developed, but the latter approach has some important advantages. A key advantage to working from employers' job demands is that it can identify indirect and general equilibrium effects more readily. Some of these may be difficult for policymakers and researchers to anticipate a priori.

c. Lessons gleaned

Despite the data limitations, the number of papers on AI impacts, as discussed in the previous section, is growing rapidly, as are the various methods, particularly blending data sources. The papers and challenges reviewed here strongly suggest the need for

more timely and detailed data on employment, occupations and work tasks. Ideally, the data should be standardized and accessible for ongoing, even real-time, analysis.

The following observations about data needs and limitations can help inform what data are needed to adequately address the challenges, as discussed in the following section:

- Public official data will continue to be an essential component of analysis of AI impacts.
    - The lack of a public data source designed to track occupational changes in a comprehensive, timely, granular and rigorous manner contributes to the uncertainty about the impact of AI.
    - Going forward, the aptness of occupational classification schemes and the frequency, timeliness, and rigor of updates will have a strong bearing on how well AI adoption and exposure can currently be tracked from AI developments and investments to use by and impacts on workers.
    - To track the full breadth of AI impacts, the public data infrastructure will need to be more nimble.
- Private sector data offers a rich source of timely information on the cutting edge of occupational changes, with its own set of limitations.
- Interoperability of data to allow better linking and matching between public and private data is necessary for blending data to track the impact of AI.

Opportunities to improve data to allow better analysis are discussed in the next section.

4. Data Developments that Could Help to Move Analysis Closer to the Desired Frontier

This section discusses a holistic approach to creating the data (particularly government statistical products) necessary to equip public and private decision makers with the information and research they need to manage the labor market disruptions associated with AI.

a. Standards-based, enhanced interoperable employer data infrastructure

Interoperable digital employer administrative records may offer the key to modernizing our labor market information system. Indeed, it is hard to envision another solution.

These records, which all employers maintain in some form, include payroll data (with timely information on dates, hours, components of pay, and work location), job postings (with dates, job titles and descriptions), and human resources files (with job

titles and descriptions, demographics, and work and training history). In short, every US employer manages a wealth of digital information on workers' jobs, characteristics, and compensation. They need these data for internal operations, to comply with regulations, and to respond to federal and state data requests. However, much of their information is not yet designed to be interoperable (that is, to work together) across purposes, software systems, employers, or government agencies.

One promising initiative that illustrates the opportunity offered by interoperable standards for employer administrative data is the Jobs and Data Exchange (JEDx) initiative being led by the US Chamber of Commerce Foundation.[12] JEDx is developing a vision and roadmap for a public-private data collaboration to improve the collection and use of standards-based jobs and employment data.

The business community and policymakers increasingly recognize that investment in interoperable employer records holds great promise for improving the timeliness and granularity of public and private evidence.[13] Furthermore, content standards designed to meet multiple needs reduce duplication of effort while technology standards can protect confidentiality consistently even as they simplify transmission. Thus, designing records for interoperability can lower employers' and public sector operating costs as they enhance individuals' privacy protections.[14]

Interoperability, particularly in terms of standardized definitions for occupations and data items, privacy protections, and secure access to data across all sources and the ways they are updated is critical to ensure timeliness and allow the body of analysis to remain useful and relevant as new AI technological developments occur.

More standardized, interoperable and timely data would benefit empirical research on AI by enabling more nuanced measures of its exposure, adoption and impact and also provide the government, employers, workers, researchers and other data users more timely and nuanced information as AI technology improvements continue.

DRAFT

[12] The JEDx initiative is part of the US Chamber of Commerce Foundation's The New Data Paradigm program. The JEDx platform has been developed over nearly a decade and is now being piloted with selected state partners, employers, and data technology firms. The New Data Paradigm represents the next step by improving the enhanced and standardized data in ways that can be linked to information about workers skills and credentials that can provide beneficial data to workers, training institutions and workers.

[13] See the JEDx Value Proposition Research Report 12 20 2022 and Megan Leonhardt, "The U.S. Needs Better Data. Why We Fell Behind and How We Can Catch Up." Barron's magazine, Nov 22, 2025.

[14] Note that state governments' requirements on employers for these records are set to increase substantially with the addition of federal work requirements for Medicaid and the Supplement Nutritional Assistance Program.

b. Labor market performance improvements from enhanced, standardized, and interoperable employment data

Standards-based employment data systems would have important benefits for the functioning of the labor market as it is buffeted by any shock, whether associated with Ai or not. Five potential features and benefits from creating this infrastructure are noted below.

1. **Improvements to key government statistical series** would be enabled by updated data collection processes UI/QCEW data collection processes that include more information on location, workers, and impacts. More timely and granular measures of employment dynamics and wages would help to refine impact analysis and target policies across many purposes.This could directly improve the measurement approach used in Johnston & Makridis (2025) by allowing narrower industry specifications and variation across geographies. Collection costs would also be lower.

2. **Valuable enhancements to employment and earnings records** could support swift updating of occupational classifications and characteristics (SOC and O*NET). The wage records collected by states' Unemployment Insurance systems are currently very sparse and typically include no information on occupations. The system JEDx envisions would allow states and the BLS to collect particular additional information with little or no additional burden to employers. High-value fields already held by employers include occupation or job title, hours, work and residence location and employment start and stop dates (Groshen, et al 2023). These enhancements would yield occupational updates that are uniquely valuable and critical to AI. In addition, such enhancements would yield other important advantages, such as greatly facilitating program and policy evaluations and better informing local workforce development efforts.

3. **Improved access to earnings records of workers** would allow workers, job seekers, recruiters, and policy evaluators with access to lower cost and better information about work history and combine it with information on skills and education. Workers could use certified Learning and Employment records (LERs) to aid in job search. Employers could use LERs to better assess their workforce and screen job applicants. Researchers could produce more precise evaluations and analysis of public education and training programs that prepare future workers for occupations and skills in demand and retrain workers displaced by potential disruptions to the labor market. Administrative earnings records are more reliable and detailed than household survey responses to questions about compensation.

DRAFT

4. **Incorporation of standardized electronic job listings from employers directly** is also possible. These listings paired with other employer records would reveal evolving hiring demand and job requirements. This information could extend and improve the intra-firm dynamics explored in Hampole, et al. (2025). This could be applied within existing occupation structures but the job descriptions that employers provide could also be used to more regularly create new occupations that reflect specific demands from employers that seek more AI specialists and AI-focused skills within existing occupations. Compared to Burning Glass or Revelio, these listings would be more universal, ensuring better coverage of occupations that do not require a college degree, similar to the National Labor Exchange or Kaggle.[15]

5. **Creating an infrastructure that allows for an easier path for introducing new questions of employers** would be valuable as the landscape changed, due to AI and other influences. For example, it would be desirable for the federal government to add specific AI-related questions to existing surveys or have regular periodic AI-specific surveys of employers and/or households.

In short, standardized and interoperable data on occupations, skills required, and workers could improve the analysis of AI impacts in many ways. Labor demand and supply patterns and trends could clarify shifts in the nature of jobs and how demand for workers and skills changes as the technology (undoubtedly) changes. Regular national statistical economic reports could benefit from more detailed, standardized and precise data on AI-related factors on a routine basis, interoperably combined with other administrative and survey data to gauge trends nationally, regionally and by state. Policymakers at the federal, state and local level could benefit from more specific data about labor needs of businesses, allowing more timely decisions about education and training programs, credentials appropriate for their labor markets, and learning and employment records (LERs).

Incorporating job listings information along with employer-provided training to incumbent workers into analysis of hiring demand and evolving job requirements could enhance the labor demand and supply analysis and help inform policy decisions about whether and what types of credential and training programs should be available. Employers would benefit from more timely, standardized and tailored information about their local labor market and industry trends. Program evaluators could access more timely and comprehensive employment and earnings data on graduates and

[15] The National Labor Exchange (NLx) is a public-private partnership of the nonprofit organization DirectEmployers Association and the National Association of State Workforce Agencies. NLx collects job openings from about 20,000 sources (e.g., state public job banks, government, corporate and career websites), and over 300,000 employers. NLx "syndicates" their data (e.g., eliminating duplicates and listings with "dead" links to the full posting, is "refreshed" daily, and is used by many state and local public employment agencies and job seekers. https://usnlx.com/

trainees to rigorously determine the effectiveness of training interventions and programs.

5. Conclusion

This review of the current state of measurement of exposure to and impact of AI reveals results which range from negligible effects to quite substantial impacts on particular groups of workers. The uncertain state of the literature is not surprising given the early date of both the implementation and measurement of AI. We conclude that new data and information systems are necessary to resolve these discrepancies and better inform policy. In particular, standards-based employer administrative data could significantly improve the timely measurement of impacts and swifter adjustments of skill and occupational classifications could better identify which workers are most exposed to AI. Our approach could improve tracking, analyzing, and assessing AI's effects on the labor market over time to improve the policy responses.

Our review also implies that we must improve our data quickly to prepare ourselves to act on public policy solutions. We have time, but it will take time to develop the needed infrastructure. Without an initial start to get good data, we will be behind the curve and have a hard time keeping up with changes in our economy. Although data requirements will likely evolve as we go along, we offer these steps as a useful start.

DRAFT

References

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6), 2188–2244. https://shapingwork.mit.edu/wp-content/uploads/2023/10/Robots-and-Jobs-Evidence-from-US-Labor-Markets.p.pdf

Acemoglu, Daron, David Autor, and Simon Johnson (February 2026) “Building pro-worker artificial intelligence.” Building pro-worker artificial intelligence | MIT Economics

Acemoglu, Daron, David Autor, Jonathan Hazell, and Pascual Restrepo (2022). “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” NBER Working Paper. Available at: https://www.nber.org/papers/w28257

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics, 118(4), 1279–1333. https://economics.mit.edu/sites/default/files/publications/the%20skill%20content%202003.pdf

Baslandze, Salomé, Zachary Edwards, John R. Graham, Ty McClure, Michael Sparks, Brent Meyer, Sonya Ravindranath Waddell, and Daniel Weitz (2026). Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives. Federal Reserve Bank of Atlanta Working Paper 2026-4 March 2026. https://doi.org/10.29338/wp2026-04

Bessen, J. E. (2016). How computer automation affects occupations: Technology, jobs, and skills (Law & Economics Research Paper No. 15-49). Boston University School of Law. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2847189_code355514.pdf?abstractid=2690435

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025a). "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence." Available at: https://digitaleconomy.stanford.edu/app/uploads/2025/12/CanariesintheCoalMine_Nov25.pdf

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond (2025b). “Generative AI at Work.” NBER Working Paper. Available at: https://www.nber.org/papers/w31161

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang (2023). "Generative AI and Firm Values" Available at SSRN: papers.ssrn.com/sol3/papers.cfm?abstract_id=4440717

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models (arXiv:2303.10130, version 5). arXiv. https://doi.org/10.48550/arXiv.2303.10130

DRAFT

Felten, Edward, Manav Raj, and Robert Seamans. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." Strategic Management Journal (2021). doi:10.1002/smj.3286

Frame, W., Scott, and Lawrence J. White. (2004). Empirical Studies of Financial Innovation: Lots of Talk, Little Action? Journal of Economic Literature 42 (1): 116–144. DOI: 10.1257/002205104773558065

Frame, W. Scott and Lawrence J. White, (January 2014). Technological Change, Financial Innovation, and Diffusion in Banking. NYU Working Paper No. 2451/33549, https://ssrn.com/abstract=2380060

Frey, Carl Benedikt, and Michael A. Osborne. “The Future of Employment: How Susceptible Are Jobs to Computerisation?” Oxford Martin School, September 17, 2013. Available at: https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment

Fung, M. K. (2006). Are labor-saving technologies lowering employment in the banking industry? Journal of Banking & Finance, 30(1), 179–198. https://doi.org/10.1016/j.jbankfin.2005.03.008

Georgieff, Alexander, and Jieun Hyee (2021). “Artificial Intelligence and Employment: New Cross-Country Evidence.” OECD Social, Employment and Migration Working Paper No. 255. Available at: https://www.oecd.org

Gimbel, Martha, Molly Kinder, Joshua Kendall, and Maddie Lee, Evaluating the Impact of AI on the Labor Market: Current State of Affairs (October 1, 2025). The Budget Lab at Yale. Available at: https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Graetz, G., & Michaels, G. (2018) — “Robots at Work” (Review of Economics and Statistics).100(5), 753–768.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli, "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations" (2025). Anthropic. Available at: https://arxiv.org/abs/2503.04761

DRAFT

Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024). “The Short-Term Effects of Generative AI on Employment: Evidence from Online Labor Markets.” Available at SSRN: https://ssrn.com/abstract=4573321

Johnston, Andrew, and Christos Makridis (2025). “Automation, Artificial Intelligence, and Labor Market Outcomes.” Available at SSRN: https://ssrn.com

Menaka Hampole, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," NBER Working Paper 33509 (2025), https://doi.org/10.3386/w33509

ONET Resource Center, "ONET Occupation Update Summary." https://www.onetcenter.org/dataUpdates.html (accessed 04/11/2026).

OpenAI. (April 2026). Industrial Policy for the Intelligence Age: Ideas to keep people first.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri, "Working with AI: Measuring the Applicability of Generative AI to Occupations" (July, 2025). arXiv preprint arXiv:2507.07935. Available at: https://arxiv.org/abs/2507.07935

U.S. Bureau of Labor Statistics (2026a). Projected job openings in occupations typically requiring a bachelor’s degree, 2024–34. https://www.bls.gov/opub/ted/2026/projected-job-openings-in-occupations-typically-requiring-a-bachelors-degree-2024-34.htm

U.S. Bureau of Labor Statistics (2026b). Projected job openings in occupations typically requiring some postsecondary education, 2024–34. https://www.bls.gov/opub/ted/2026/projected-job-openings-in-occupations-typically-requiring-some-postsecondary-education-2024-34.htm

U.S. Bureau of Labor Statistics (2026c). Employment projections: 2024–34 occupation data tables. https://www.bls.gov/emp/tables/occupations-most-job-growth.htm

Webb, Michael, The Impact of Artificial Intelligence on the Labor Market (November 6, 2019). Available at SSRN: https://ssrn.com/abstract=3482150 or http://dx.doi.org/10.2139/ssrn.3482150

DRAFT

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang (2026). Firm Data on AI. Federal Reserve Bank of Atlanta Working Paper 2026-3, March 2026. https://doi.org/10.29338/wp2026-03

DRAFT

Appendix A.

History of US Banking Jobs

Currently being revised

To be added later

DRAFT

Appendix B. Construction of Occupational Exposure Heatmap (Figure 1)

We start by aligning all exposure measures (named in the footnote to Figure 1 and methodologies described in Appendix E) within a common occupational framework using SOC 2018 codes. Since different papers report data at varying levels of detail, we consolidate everything to a consistent SOC level (usually 6-digit), assigning broader occupations (5-digit) to more detailed occupations (6-digit) when codes end in 0, then proceeding to remove the broader occupations. This approach allows us to compare datasets smoothly without overcomplicating the structure.

Next, we combine all measures into a single occupation-level dataset, where each column represents a measure and each row corresponds to an occupation. Pairwise Pearson correlations are calculated using only overlapping, non-missing observations for each pair of measures. This pairwise-complete method maximizes usable data while avoiding distortions caused by uneven coverage.

To incorporate timing, we add publication dates and sort measures by recency. This order is reflected in the visualization. The final output is a correlation heatmap with measures arranged consistently across both axes, along with each publication date shown along the right axis, indicating when each measure was published.

Overall, this approach maintains clear, comparable comparisons without introducing bias from varying occupational coverage across measures.

DRAFT

Appendix C. Characteristics of Official Data

Although the US has many federal data sources for studying the labor market, none is ready to meet the challenges listed above. Most are either comprehensive with little detail or highly detailed but lacking broad coverage.

Examples of particular relevance to studies of AI's impact are:

- The Current Population Survey (CPS) has a wealth of timely information on workers but does not have a large enough sample to monitor narrow occupations, industries and occupations. It can track one-year changes for workers in its sample, but not longer.
- The National Longitudinal Surveys (NLS) and Panel Study of Income Dynamics (PSID) have detailed information on workers and track workers for extended periods of time, but they are infrequent, less timely, and have sample sizes smaller than the CPS.
- The Quarterly Census of Employment and Wages (QCEW) is granular, fairly timely and has employer information and covers most employment, but does not track individual employees or their occupations.
- The National Directory of New Hires (NDNH) covers most workers, but does not routinely include occupation or skills or information about the employer. Furthermore, historical NDNH data is not kept and access to the data is very limited.
- The Longitudinal Employer Household Dynamics (LEHD) currently comes closest to filling these needs currently. Data on workers is linked to data on employers. However, states drop in and out of participation in the LEHD, there can be long lags in updating, and the terms of access through the Federal Statistical Research Data Centers do not allow continued monitoring. In addition, gaining access to the data can be time consuming and difficult.
- O*NET provides detailed information on occupational activities, but is updated irregularly and with information of unclear reliability.
- The Occupational Employment and Wage Survey (OEWS) collects job titles, duties, employment, and wages. However, OEWS data are not designed to be used longitudinally and have lags in publication.

DRAFT

Table C1: Comparison of official labor market data sets

Data source (producing agency)	Unit of observation	Coverage	Longitudinal	Frequency	Granularity	Occupa-tion	Employer-worker linkage
Current Population Survey - CPS (Census and BLS)	Individual	National sample	I-year changes	Monthly	Low, small sample	Yes	No
National Longitudinal Survey – NLS (BLS)	Individual	National sample	Yes	Biennial	Low, small sample	Yes	No
Panel Study of Income Dynamics – PSID	Individual	National sample	Yes	Biennial	Low, small sample	Yes	No
Quarterly Census of Employment and Wages – QCEW (BLS)	Establishment	National census	Yes	Quarterly	High, by geography & industry	No	No
National Directory of New Hires - NDNH (Social Security Admin.)	Individual	National census	Yes	Quarterly	High, by geography & industry	No	Yes
Longitudinal Employer Household Dynamics – LEHD (Census)	Individual	Census of most states	Yes	Quarterly	High by geography & industry	Modeled	Yes
Occupational Information Network–O*Net (US Dept of Labor)	Occupation	NA	No	NA	High, by occupation with tasks and skills	Yes	No
Occupational Employment and Wage Survey – OEWS (BLS)	Occupation in establishment	National sample	No	Annual	High, by occupation with duties	Yes	Yes

DRAFT

Appendix D. Construction of AI Exposure Across Occupation Groups with Additional Views

For the dumbbell charts beginning with Figure 2, we standardize exposure measures within each dataset before making comparisons. For each measure, we restrict the sample to only the occupations for which that measure reports values and apply a min–max normalization so that scores range from 0 to 1 within that measure. This ensures comparability across measures while preserving the relative ranking of occupations within each dataset.

When constructing the charts, occupations are grouped into cohorts based on BLS classifications, either by predefined occupation groupings or by Major Occupation Groups (2-digit SOC codes). Within each cohort, we include all relevant occupations available in the data.

Because some measures report exposure at less detailed levels, we expand these observations to a consistent level of detail. When a reported occupation corresponds to a broader category, we assign that value to all underlying detailed occupations within that subcategory and then proceed using the fully expanded set. This allows us to maintain consistency in granularity across measures.

The resulting figures display the range of exposure scores across measures for each occupation or occupation group, with minimum and maximum values connected by a line and individual measure estimates overlaid. This approach highlights variation across methodologies while maintaining a consistent and comparable structure across all charts.

DRAFT

Figure D1: AI Exposure range by occupation, for occupations typically requiring a minimum of Postsecondary Education

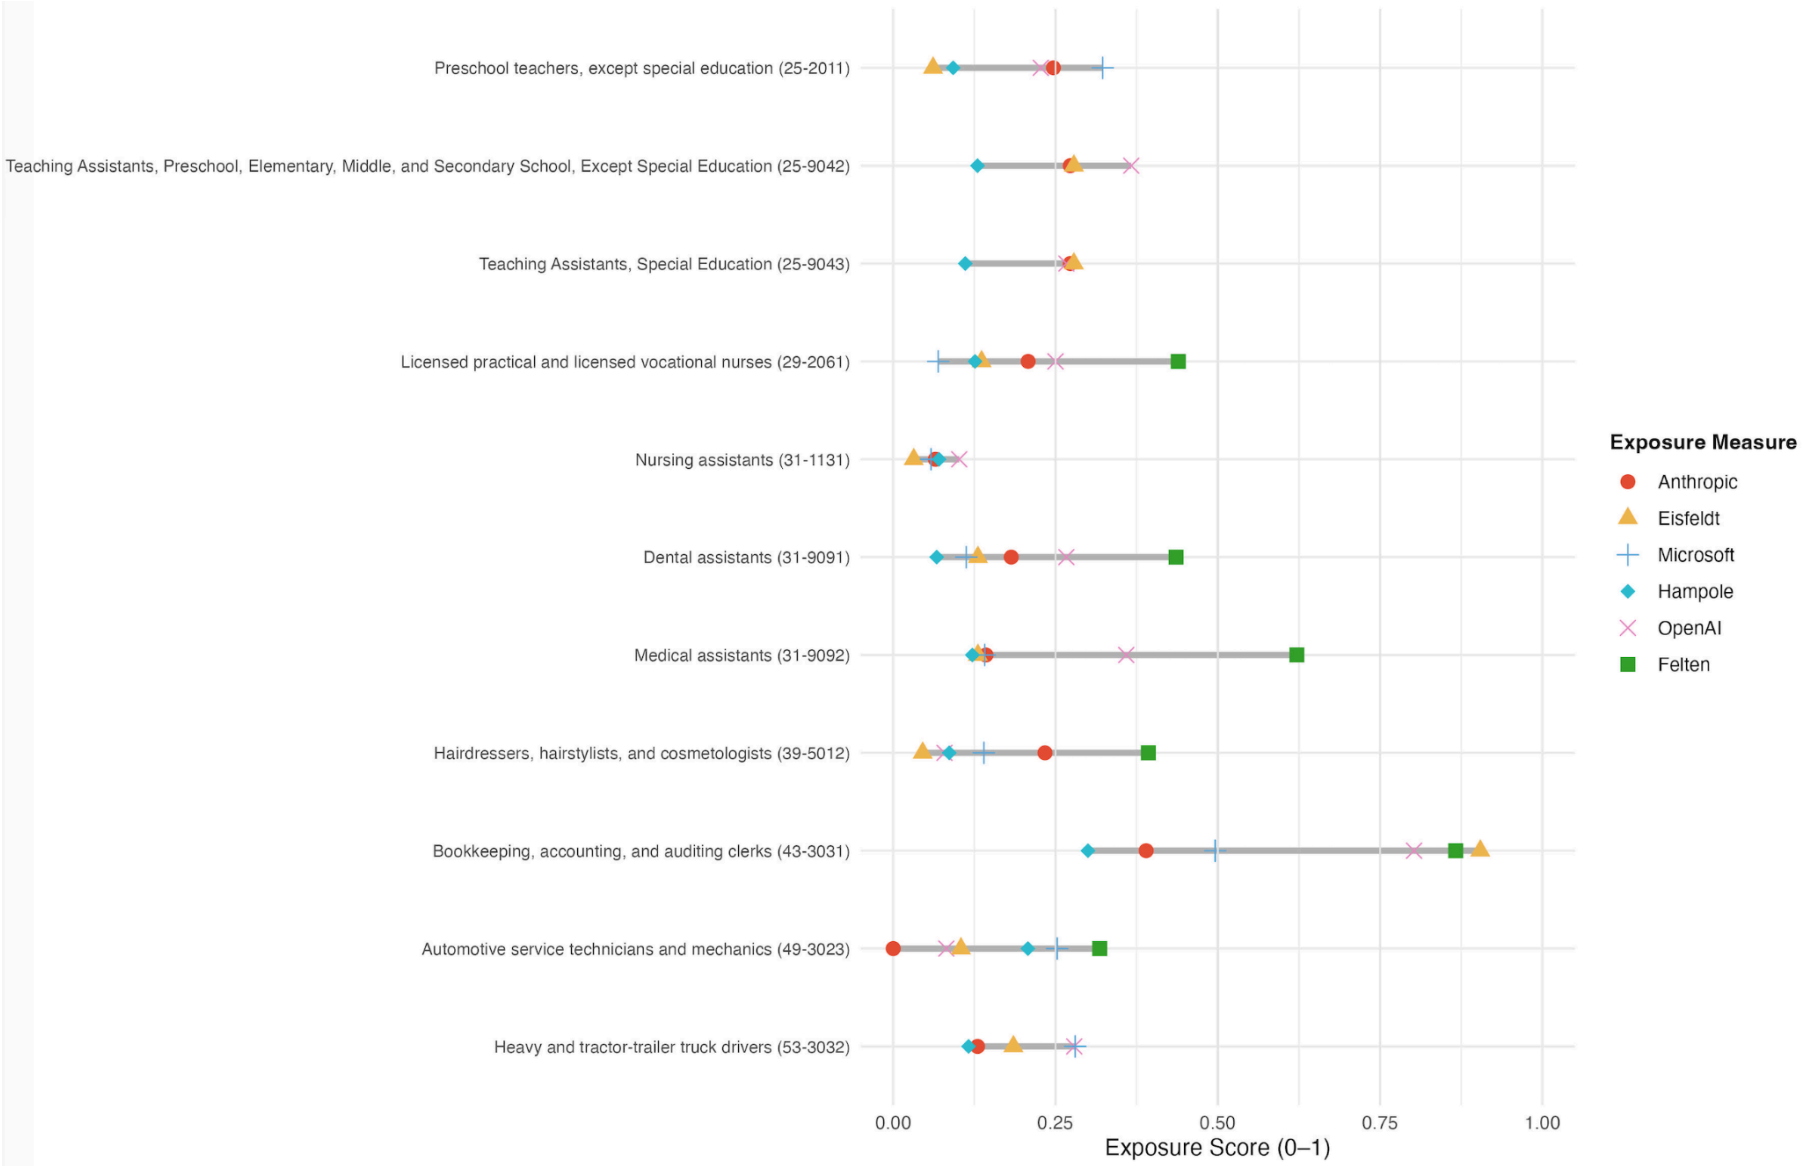
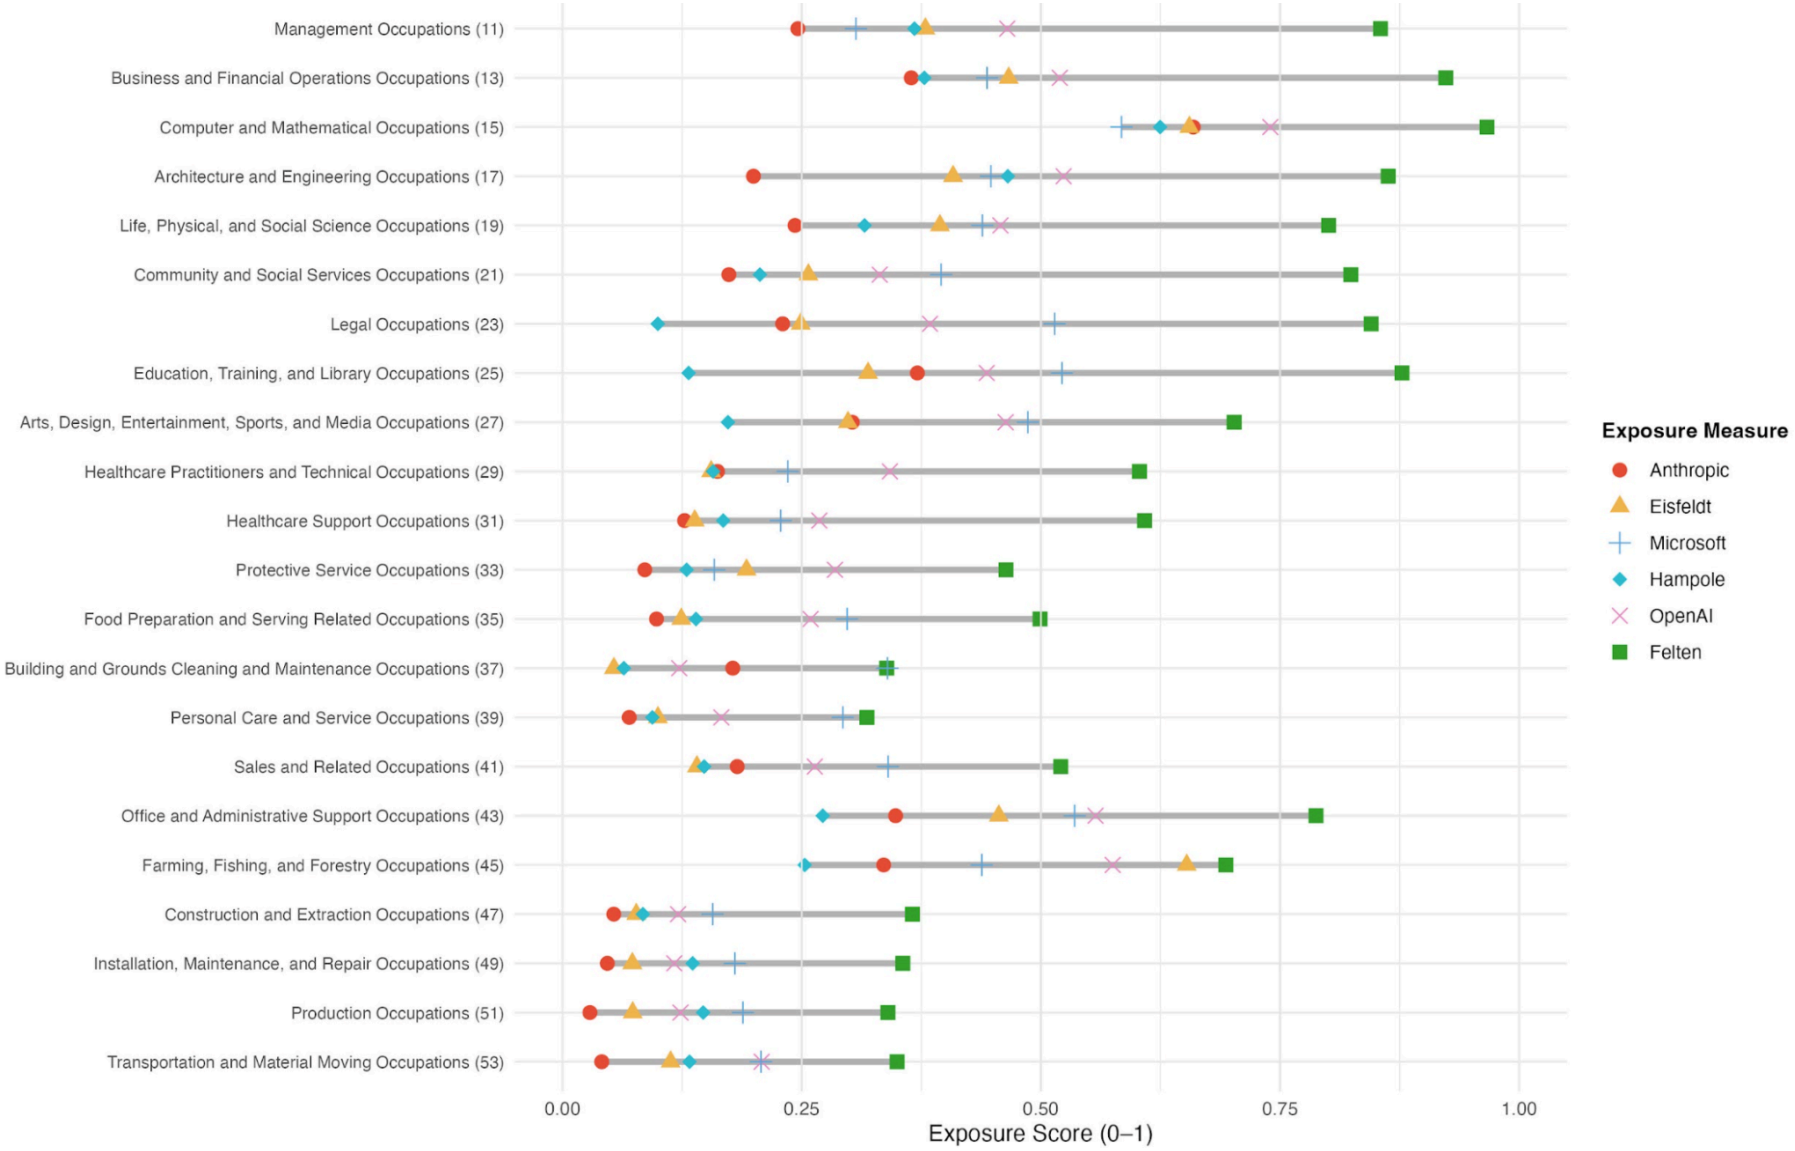


Figure D2: AI Exposure Range by Major Occupation Groups



Appendix E. Methodological Explanations for AI Exposure Measures

Table E1: Summary of AI Exposure Measures and Data Sources

Paper	Publication Date	Approximate Share of Focus Group
Frey	September 2013	Use O*NET bottleneck-related task/skill measures plus expert-labeled occupations to train a Gaussian Process model that assigns each occupation a probability of computerisation.
Webb	November 2019	Constructs occupation-level exposure scores by extracting task-like verb-noun pairs from AI-related patent text, comparing them with occupation task descriptions using NLP, and measuring how strongly each occupation's tasks overlap with patented technological capabilities.
Felten	May 2021	Construct AI exposure scores by having Amazon Mechanical Turk rate how closely AI application areas map to ONET human abilities, then aggregating those ability-level ratings to the occupation level using O*NET importance and level weights.
OpenAI	August 2023	Classify O*NET tasks by direct LLM exposure and software-enabled LLM exposure, then aggregate to occupations using weighted O*NET task data, with score = direct exposure + 0.5 × software-enabled exposure.
Hampole	February 2025	Uses LLM-extracted firm AI applications from worker resumes, matches them to O*NET tasks via semantic similarity, and aggregates task-level exposure into occupation scores using O*NET task-importance weights.
Microsoft	July 2025	Uses LLM-classified Microsoft Copilot conversations to map user goals and AI actions to O*NET work activities, then aggregates activity-level coverage, completion, and scope into occupation scores using O*NET importance and relevance weights.
Eisfeldt	September 2025	Uses O*NET task descriptions, classifies each task's generative AI exposure using GPT-based ratings of whether ChatGPT can reduce task completion time, and aggregates those task-level exposures to occupations and then firms using occupational employment weights.
Anthropic	October 2025	Maps Claude conversations to O*NET tasks with a hierarchical LLM classifier and aggregates task-level usage to occupations using O*NET task–occupation links.

DRAFT

References

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6), 2188–2244. https://shapingwork.mit.edu/wp-content/uploads/2023/10/Robots-and-Jobs-Evidence-from-US-Labor-Markets.p.pdf

Acemoglu, Daron, David Autor, and Simon Johnson (February 2026) “Building pro-worker artificial intelligence.” Building pro-worker artificial intelligence | MIT Economics

Acemoglu, Daron, David Autor, Jonathan Hazell, and Pascual Restrepo (2022). “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” NBER Working Paper. Available at: https://www.nber.org/papers/w28257

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics, 118(4), 1279–1333. https://economics.mit.edu/sites/default/files/publications/the%20skill%20content%202003.pdf

Baslandze, Salomé, Zachary Edwards, John R. Graham, Ty McClure, Michael Sparks, Brent Meyer, Sonya Ravindranath Waddell, and Daniel Weitz (2026). Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives. Federal Reserve Bank of Atlanta Working Paper 2026-4 March 2026. https://doi.org/10.29338/wp2026-04

Bessen, J. E. (2016). How computer automation affects occupations: Technology, jobs, and skills (Law & Economics Research Paper No. 15-49). Boston University School of Law. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2847189_code355514.pdf?abstractid=2690435

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025a). "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence." Available at: https://digitaleconomy.stanford.edu/app/uploads/2025/12/CanariesintheCoalMine_Nov25.pdf

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond (2025b). “Generative AI at Work.” NBER Working Paper. Available at: https://www.nber.org/papers/w31161

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang (2023). "Generative AI and Firm Values" Available at SSRN: papers.ssrn.com/sol3/papers.cfm?abstract_id=4440717

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models (arXiv:2303.10130, version 5). arXiv. https://doi.org/10.48550/arXiv.2303.10130

Felten, Edward, Manav Raj, and Robert Seamans. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." Strategic Management Journal (2021). doi:10.1002/smj.3286

Frame, W., Scott, and Lawrence J. White. (2004). Empirical Studies of Financial Innovation: Lots of Talk, Little Action? Journal of Economic Literature 42 (1): 116–144. DOI: 10.1257/002205104773558065

Frame, W. Scott and Lawrence J. White, (January 2014). Technological Change, Financial Innovation, and Diffusion in Banking. NYU Working Paper No. 2451/33549, https://ssrn.com/abstract=2380060

Frey, Carl Benedikt, and Michael A. Osborne. “The Future of Employment: How Susceptible Are Jobs to Computerisation?” Oxford Martin School, September 17, 2013. Available at: https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment

Fung, M. K. (2006). Are labor-saving technologies lowering employment in the banking industry? Journal of Banking & Finance, 30(1), 179–198. https://doi.org/10.1016/j.jbankfin.2005.03.008

Georgieff, Alexander, and Jieun Hyee (2021). “Artificial Intelligence and Employment: New Cross-Country Evidence.” OECD Social, Employment and Migration Working Paper No. 255. Available at: https://www.oecd.org

Gimbel, Martha, Molly Kinder, Joshua Kendall, and Maddie Lee, Evaluating the Impact of AI on the Labor Market: Current State of Affairs (October 1, 2025). The Budget Lab at Yale. Available at: https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Graetz, G., & Michaels, G. (2018) — “Robots at Work” (Review of Economics and Statistics).100(5), 753–768.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli, "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations" (2025). Anthropic. Available at: https://arxiv.org/abs/2503.04761

Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024). “The Short-Term Effects of Generative AI on Employment: Evidence from Online Labor Markets.” Available at SSRN: https://ssrn.com/abstract=4573321

Johnston, Andrew, and Christos Makridis (2025). “Automation, Artificial Intelligence, and Labor Market Outcomes.” Available at SSRN: https://ssrn.com

Menaka Hampole, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," NBER Working Paper 33509 (2025), https://doi.org/10.3386/w33509

ONET Resource Center, "ONET Occupation Update Summary." https://www.onetcenter.org/dataUpdates.html (accessed 04/11/2026).

OpenAI. (April 2026). Industrial Policy for the Intelligence Age: Ideas to keep people first.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri, "Working with AI: Measuring the Applicability of Generative AI to Occupations" (July, 2025). arXiv preprint arXiv:2507.07935. Available at: https://arxiv.org/abs/2507.07935

U.S. Bureau of Labor Statistics (2026a). Projected job openings in occupations typically requiring a bachelor’s degree, 2024–34. https://www.bls.gov/opub/ted/2026/projected-job-openings-in-occupations-typically-requiring-a-bachelors-degree-2024-34.htm

U.S. Bureau of Labor Statistics (2026b). Projected job openings in occupations typically requiring some postsecondary education, 2024–34. https://www.bls.gov/opub/ted/2026/projected-job-openings-in-occupations-typically-requiring-some-postsecondary-education-2024-34.htm

U.S. Bureau of Labor Statistics (2026c). Employment projections: 2024–34 occupation data tables. https://www.bls.gov/emp/tables/occupations-most-job-growth.htm

Webb, Michael, The Impact of Artificial Intelligence on the Labor Market (November 6, 2019). Available at SSRN: https://ssrn.com/abstract=3482150 or http://dx.doi.org/10.2139/ssrn.3482150

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka,

Gregory Thwaites, and Ben Zhe Wang (2026). Firm Data on AI. Federal Reserve Bank of Atlanta Working Paper 2026-3, March 2026. https://doi.org/10.29338/wp2026-03

DRAFT