What Work Does Generative AI Do? *

Alexander Bick
Federal Reserve Bank of St. Louis & CEPR

Adam Blandin
Vanderbilt University

David J. Deming
Harvard Kennedy School & NBER

Tyler Schumacher
Vanderbilt University

Preliminary

February 3, 2026

Abstract

The impact of genAI on worker productivity, wages, and employment hinges on what work genAI actually does in practice. We address this measurement challenge using novel data from a nationally representative labor market survey that links workers' reported genAI use to detailed O*NET tasks. This allows us to construct occupation- and task-level adoption indexes that capture realized genAI use. We document substantial heterogeneity in adoption across tasks and occupations, with notable deviations from existing exposure predictions. While exposure measures are positively correlated with adoption, they account for less than half of the variation, in part because individual-level factors such as learning and cross-domain spillovers also shape adoption. Our indexes capture these dynamics and can be updated as genAI capabilities and adoption patterns evolve, providing a continuously updated input for research on the labor market effects of this rapidly diffusing technology.

JEL Codes: J24, O33, C83
Keywords: Generative AI, Technology Adoption, Occupation and Task-Level Measurement, AI Exposure, Labor Markets

* *Contact*: alexander.bick@stls.frb.org; adam.blandin@vanderbilt.edu; david_deming@harvard.edu; tyler.schumacher@vanderbilt.edu. Kevin L. Bloodworth provided excellent research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. This work was supported by a joint grant from by the Washington Center for Equitable Growth and the Russell Sage Foundation.

1 Introduction

Generative artificial intelligence (genAI) tools, such as ChatGPT, are rapidly diffusing into workplaces (Bick, Blandin, and Deming, 2026). A central question is how this new technology will affect worker productivity, wages, and employment. Prior research on technological change emphasizes that these labor market effects depend critically on the type of work a new technology performs (Autor, Levy, and Murnane, 2003). Tasks performed by a new technology tend to become less valued in the labor market, while complementary tasks that the technology does not do tend to become more valued. Understanding genAI's labor market impact therefore requires understanding what work genAI does.

We address this measurement challenge using novel data from the Real-Time Population Survey (RPS), a repeated national online labor market survey that has run since 2020 (Bick and Blandin, 2023). Since August 2024, the survey has included a genAI adoption module that asks workers whether they use genAI for their job, how intensively they use it, and how much time the technology saves them.

A key challenge in trying to learn what work genAI does is that it is infeasible to ask survey respondents about hundreds or thousands of specific work activities. We address this challenge in two steps. First, for each worker we elicit their detailed occupation using an autocomplete text box linked to the 2018 Standard Occupational Classification (SOC) system.[1] Second, in the two most recent survey waves in August and November 2025 we use the worker's occupation to identify the ten most important work tasks for their occupation, as described in the Bureau of Labor Statistics' O*NET database. We present these ten tasks to respondents and ask them to select which tasks they performed in their job. For workers who report using genAI for their job, we then ask which of these tasks they used genAI to help with. This information allows us to construct a task-level adoption rate: among workers who do a given task, what share use genAI for it?

We begin by validating our measures of task performance and task-level AI adoption rates, which we collect in two survey waves in August and November 2025. The median worker reports performing five tasks (out of a maximum of ten), while only 1.2% of workers select zero of the presented tasks. The share of workers who perform a given task varies substantially across tasks and is stable across survey waves, suggesting that our measure of task performance contains meaningful information. Interestingly, workers who share a detailed occupation frequently report performing different tasks: to the extent that genAI assists some tasks more than others,

[1] Our occupation autocompletion approach follows Miano (2025). We extend his method to job titles that do not exactly match a unique SOC code or an autosuggestion by presenting respondents with a choice of probabilistic matches using the job title to SOC code matching algorithm developed by the National Institute for Occupational Safety and Health (Laughlin, Song, Wisniewski, and Xu, 2024).

this within-occupation variation in tasks may help explain within-occupation variation in genAI adoption. Conditional on performing a given task, the share of workers who report using genAI for that task also varies substantially across tasks and is stable across survey waves, again suggesting that our task-level AI adoption measure contains meaningful information.

We use the genAI adoption data to establish four key findings on what work genAI does. First, we characterize the occupations and tasks that genAI is most likely to assist. Occupations with the highest genAI adoption rates are management and professional occupations, specifically in business, finance, and computer-related occupations. Tasks with the highest genAI adoption rates require cognitive and information-intensive work involving data analysis, interpretation, and writing. These tasks often involve synthesizing facts and can be partially decomposed into steps well suited to language or analytical models (e.g., drafting, summarizing, evaluating alternatives, or generating procedures).

In contrast, the occupations with the lowest genAI adoption rates are generally service occupations that require manual or in-person interaction. Tasks with the lowest genAI adoption rates are concentrated in physically embodied or situational activities that require direct interaction with people or equipment. Several of these tasks also involve care or safety work, which may require quick actions and entail important responsibility or liability concerns.

Second, we find that existing task-level "genAI exposure" predictions are moderately correlated with actual genAI adoption. Recent work constructs predictions of genAI exposure by mapping predicted model capabilities onto the tasks and occupations in which genAI is likely to be applicable (Eloundou, Manning, Mishkin, and Rock, 2024; Felten, Raj, and Seamans, 2021). These predictions attempt to characterize the distribution of potential genAI assistance across tasks and occupations based on task content and assumed technological feasibility. Our approach complements this work by measuring realized genAI adoption at the task level and occupation level, capturing how and which workers report actually using the technology in practice. (Future drafts will also include comparisons to earlier measures of exposure to machine learning and automation more broadly, e.g. Brynjolfsson, Mitchell, and Rock (2018) and Webb (2020).)

At the task-level, the correlation between exposure and actual adoption ranges from 0.17 to 0.62, and at the occupation level from 0.19 to 0.71, depending on the measure of exposure used. Put differently, exposure measures alone can explain at most 40% of the variation in adoption rates across tasks and 50% across occupations, respectively. Education / Legal / Community occupations and Reasoning and Decision Making tasks exhibit particularly high adoption rates relative to predicted exposure. Office / Administrative Support occupations and Administering tasks exhibit particularly low adoption rates relative to predicted exposure.

Third, we plan to compare our survey-based task adoption rates to task-based measures constructed from large-scale genAI conversation logs (Chatterji et al., 2025; Handa et al., 2025). These papers classify prompts into O*NET tasks and describe how usage within a given product is distributed across tasks and occupations. Conversation-based measures have important strengths: they avoid measurement error from self reports and cover a broader set of activities than a pre-selected task list. However, this approach also faces limitations, including non-representative use cases and user bases for individual genAI products, limited information on worker occupation and labor market characteristics, misclassification of prompts into tasks, and difficulties classifying a prompt as work- or non-work related (as we document, non-work use accounts for a majority of total genAI use). In ongoing work, we will obtain the task-based indexes in Handa et al. (2025) and Chatterji et al. (2025) and construct analogous survey-based measures capturing a task's share of genAI use. Where the measures differ, we will investigate the potential drivers of disagreement, such as non-representative use or errors in prompt classification.

Fourth, while task characteristics strongly shape how genAI is used, we find that individual choices also play a key role in determining adoption. Within narrowly defined tasks, adoption is far from universal: even in the highest-adoption tasks a substantial share of workers do not use genAI. Consistent with this heterogeneity, adding worker fixed effects to task-level adoption regressions substantially increases explanatory power, indicating the importance of unobserved individual-level factors such as preferences, skills, beliefs, or prior experience. We provide evidence on one potential mechanism that can help explain individual variation in genAI use: the diffusion of genAI use within a given worker from one domain to another. Observationally equivalent workers who use genAI for high-adoption tasks are significantly more likely to also use it for tasks with low average adoption rates. Similarly, observationally equivalent workers whose employers encourage them to use genAI for their job are more likely to use genAI at home. Moreover, workers who have been using genAI for longer report using genAI for more tasks. These patterns are consistent with a learning or experimentation mechanism: workers may initially incur a fixed cost to adopt genAI for tasks where time savings are most salient, and then expand use to a broader set of activities as they become familiar with the technology. These results underscore the value of adoption data, which reflect not only a technology's capabilities but also the individual-level choices and barriers that shape its use.

Additional ongoing work will investigate the link between tasks and self-reported time savings from genAI. A primary goal is to learn which tasks are associated with greater time savings. We will also investigate whether time savings are primarily driven by which tasks genAI assists, the number of tasks genAI assists, or additional factors.

By measuring realized genAI use at the task level, our approach provides a simple and

portable input for research on the economic effects of genAI. The task- and occupation-level adoption indexes we construct can be directly incorporated into existing analyses of productivity, wages, task reallocation, and inequality, and can be updated as adoption patterns change. The resulting indexes will help researchers move beyond assumptions about technological feasibility toward an understanding of what work genAI actually does.

The remainder of the proposal is organized as follows. Section 2 describes our data source and how we measure genAI adoption. Sections 3 - 6 present an outline of preliminary work completed thus far.

2 Data & Measurement

2.1 The Real-Time Population Survey (RPS)

Our main data source is the Real-Time Population Survey (RPS), a national labor market survey of U.S. adults aged 18-64 (for a detailed discussion, see Bick and Blandin 2023). The RPS is fielded online by Qualtrics, a large commercial survey provider, and has collected multiple survey waves each year starting in 2020.

The RPS is designed to mirror the Current Population Survey (CPS) along key dimensions. First, the survey responses are collected within two weeks, with either the first or second week overlapping with the CPS reference week. Second, the RPS matches questions on demographics and labor market outcomes in the basic CPS and CPS Outgoing Rotation Group, using the same word-for-word phrasing when practical and replicating the intricate sequence of questions necessary to elicit labor market outcomes in a manner consistent with the CPS (US Census Bureau, 2015). Replicating key portions of an existing high-quality survey ensures that survey concepts are comparable, which allows researchers to validate RPS outcomes against a widely used benchmark with a larger sample size and to construct sample weights. Bick and Blandin (2023) and Bick, Blandin, Caplan, and Caplan (2025a,b) show that the RPS closely aligns with government household surveys on distributions of employment, hours worked, earnings, industry composition, employee tenure, and work from home.

Since August 2024, the RPS is conducted quarterly with a module measuring genAI use both at work and at home. The analysis in this paper is based on the August and November 2025 surveys, which include our task-level measures, and will be extended to incorporate the February 2026 survey.

2.1.1 RPS Sample

Qualtrics panel members are recruited to the panel online and, in our case, can participate in exchange for 30 to 50 percent of the $5.50 paid per completed survey. The Qualtrics panel includes about 15 million members and is not a random sample of the U.S. population. However, Qualtrics can target survey invitations to specific demographic groups. The RPS sample was designed to be nationally representative across key demographics. We also include measures to filter out low quality responses. Details on sampling and data quality procedures can be found in Appendix A.1.

Each survey round collected responses for about 5,000 individuals. Appendix Table A.1 compares the unweighted sample composition between the CPS and RPS along the demo-

graphics targeted in the sampling procedure for our main surveys (columns 1 and 2). The bottom panel also compares employment status in the CPS and RPS, which was not targeted in the sampling procedure. The two surveys yield very similar population shares.

2.1.2 Sample Weights and Validation

To address remaining discrepancies, we construct sample weights using the raking algorithm of Deming and Stephan (1940), ensuring that the weighted sample proportions align with the demographic characteristics targeted in the sampling procedure and that the employment shares in the RPS align with the CPS as well. Details on weighting can be found in Appendix A.2.

In Appendix A.3, we show that the distribution of usual weekly earnings, occupation, and industry shares in the RPS, which were not targeted during sampling, closely align with the CPS, in both the raw and weighted data.

2.2 Measurement of Occupation

We elicited respondents' job titles through a free text response:

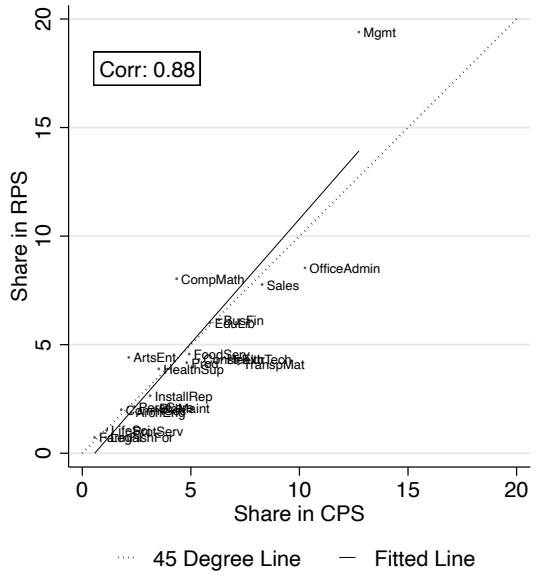
> *What is your main occupation? Please type your occupation in the box below and select one of the suggested options. Try to be specific. For instance, write preschool teacher or high school teacher rather than just teacher. If none of the options corresponds to your occupation, try adding more detail or rephrasing.*

Once they have starting typing, respondents are presented with autocomplete suggestions covering over 40,000 occupations from O*NET and the Occupational Outlook Handbook. For job titles that do not exactly match a unique SOC code, we present respondents with a choice of the five most probabilistic matches using the job title and industry, which we elicit in a prior step, to a SOC code matching algorithm developed by the National Institute for Occupational Safety and Health [NIOSH] (Laughlin, Song, Wisniewski, and Xu, 2024). We proceed in the same way for respondents whose text response does not match one of the autosuggestions exactly.

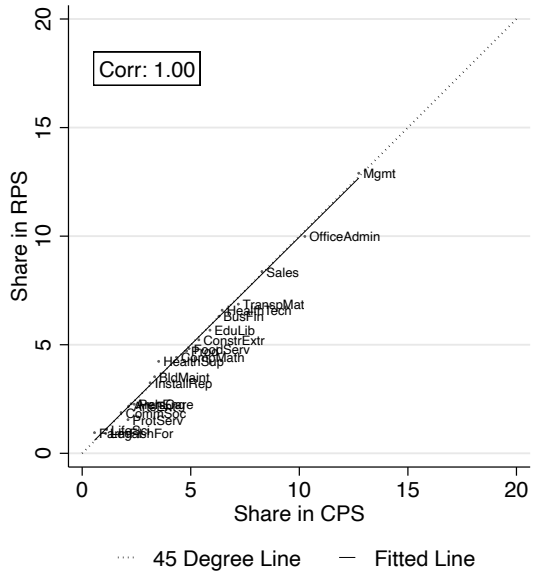
Figure 1 compares the distribution of occupation employment shares in the RPS to shares in the CPS. Occupation shares were not targeted by our sampling scheme. Figure 1a shows that occupation shares in the unweighted RPS data are overall very similar to those in the CPS. Figure 1b verifies that occupation shares in the weighted RPS data nearly perfectly match shares in the CPS.

Figure 1: Validation Check – Occupation Shares in the RPS and CPS

(a) Occupation Shares: Unweighted



(b) Occupation Shares: Weighted



Notes: The figure on the left use unweighted RPS data, the figure on the right use weighted RPS data. We use the same sample of RPS respondents in both figures. All figures use weighted CPS data. Data samples for the occupation comparison are employed respondents ages 18-64 in the pooled August and November 2025 RPS and CPS with sample sizes of 7181 and 80005, respectively.

2.3 Measurement of Generative AI Use

When designing our survey, a key goal was to enable comparisons to the historical adoption of other technologies. To facilitate these comparisons, we use existing technology adoption questions from the CPS Computer and Internet Use supplement (CIU) as a template for our own questions regarding genAI. In the CIU, the leading question about computer adoption at work for employed respondents was:

> *Do you [directly] use a computer for your job? (No/Yes)*

A second question asks about computer use at home:

> *Do you [directly] use a computer at home? (No/Yes)*

These questions were asked in 1984, 1989, 1993, 1997, 2001, and 2003, with the word "directly" being omitted in the last two years.[2]

[2]Starting in 2001, the CIU also asked about internet usage.

The RPS genAI module begins with a definition of genAI:

> *Generative AI is a type of artificial intelligence that creates text, images, audio, or video in response to prompts. Some examples of Generative AI include ChatGPT, Gemini, and Midjourney.*

We included examples of popular genAI products because we thought some respondents may be more familiar with those product names than with the broader concept of genAI. For employed respondents, the next question asks about genAI use at work:

> *Do you use Generative AI for your job? (No/Yes)*

This question is designed to mirror the analogous computer use question from the CIU, discussed above. We opted for the version asked from 2001 onwards that omitted the word "directly."

Another question later in the survey asks about genAI usage outside work:

> *Do you use Generative AI [outside your job]? (No/Yes)*

Non-employed respondents are not shown the term in brackets. While the CIU asks about computer use "at home", we use slightly different language to account for respondents who use genAI on mobile devices outside of their home.

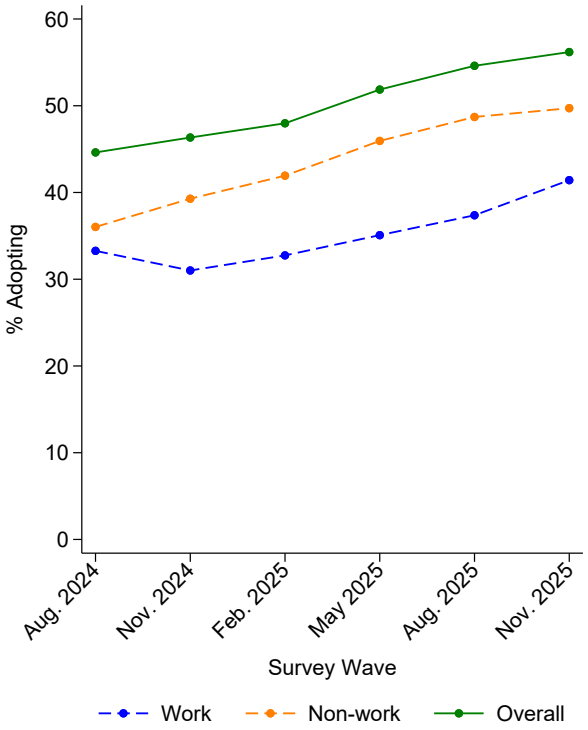
The survey asks several follow-up questions of genAI users regarding which products they used, how intensively they used it, and how much time they saved, always distinguishing between work and non-work use.

An important caveat regarding our measurement approach is that it will only capture genAI use that respondents are aware of. For example, while typing prompts into ChatGPT is an obvious case of using genAI, automatic AI-assisted writing suggestions are a more subtle case that respondents may not consider. Because our estimates may not detect some passive or embedded uses of the technology, we interpret our estimates as a lower bound on the extent of genAI adoption.

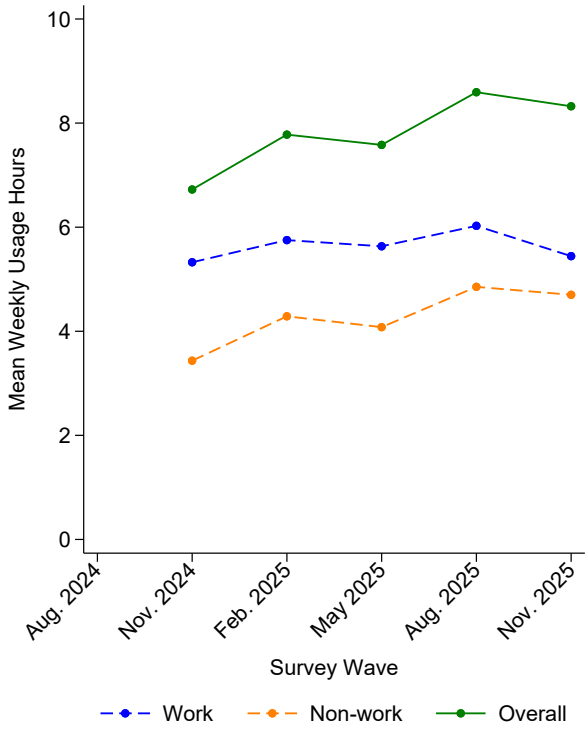
Figure 2a plots the share of individuals who report using genAI for work, non-work, and overall. Our analysis focuses on the two latest surveys (August and November 2025). As of November 2025, 50% of US adults aged 18-64 use genAI for non-work reasons, while 41% of workers use genAI for work. The overall adoption rate, i.e., the share of individuals who use genAI either for work or non-work purposes, is 56%. Along all three margins, adoption has increased substantially since our first reading in November 2024.

Figure 2: genAI Adoption Over Time

(a) Extensive Margin: AI Adoption Rate



(b) Intensive Margin: Weekly Usage Hours



Notes: In both panels, the work use lines limit the sample to respondents who work. Panel 2a shows the share of respondents using genAI for work, non-work, and overall by RPS survey wave. Panel 2b shows the mean usage hours per person by survey wave, conditional on using genAI (and in the case of the work line, conditional on at-work use of genAI). Information on time spent using genAI is not comparable to the other survey waves in the Aug. 2024 survey, so it is excluded from Panel 2b.

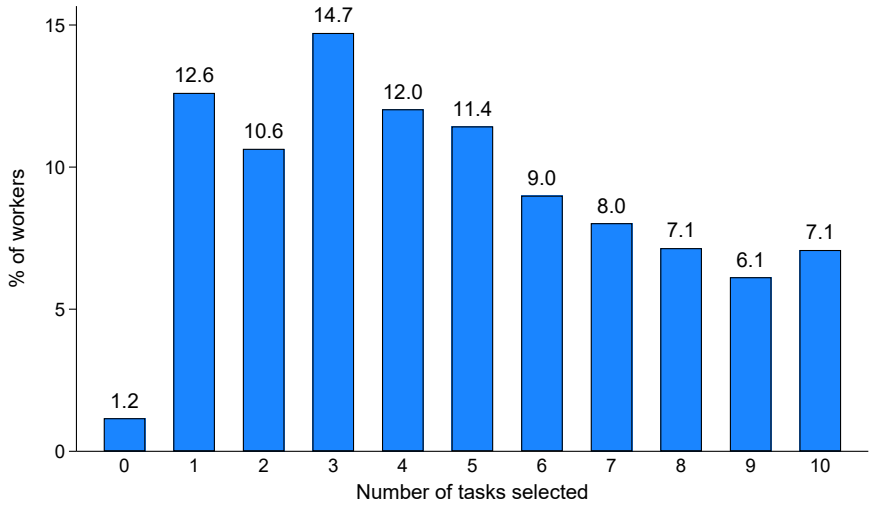
Figure 2b shows the mean hours that genAI users spent using the technology in the past week, providing an indication of the intensity of genAI use over time. (The series begin in November 2024 because our time use categories changed between the August and November 2024 waves.) Mean weekly hours among genAI users are flat for work users, but increase mildly for non-work use.

2.4 Measurement of O*NET Tasks

After identifying a respondent's occupation, we ask whether they perform each of the 10 most important O*NET detailed work activities (DWA) for their occupation. For genAI users, we then ask which of those activities they use genAI for.

We use ONET work activities rather than task statements because activities are more concise and intuitive for survey respondents. Using crosswalks from ONET, we link task-level

Figure 3: Number of O*NET Task that Workers Report Performing



Notes: Dataset is pooled from Aug. 2025 and Nov. 2025 RPS survey waves. Survey respondents are shown the top 10 O*NET activities in their detailed occupation and asked whether they actually perform each task at work. The figure shows the distribution of the number out of the top 10 task which respondents perform.

importance scores to 2,040 DWAs. We present respondents with the 10 most important DWAs for their occupation. DWAs can be aggregated into 333 intermediate work activities (IWAs), which we use for most of our analysis. Details on the construction of importance scores, the linking procedure and precise questions are provided in Appendix A.5.

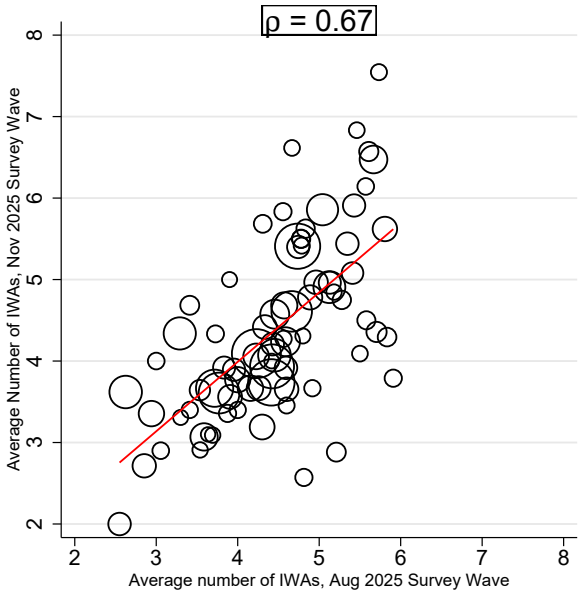
2.4.1 Work Tasks: Basic Facts and Validation

For the remainder of the paper we use the parlance of the literature and refer to O*NET work activities as "tasks." Figure 3 shows the distribution of the number of work tasks that people report performing. The median number of tasks selected is five; 1.2% of workers select zero tasks and 7.1% of workers report performing all ten.

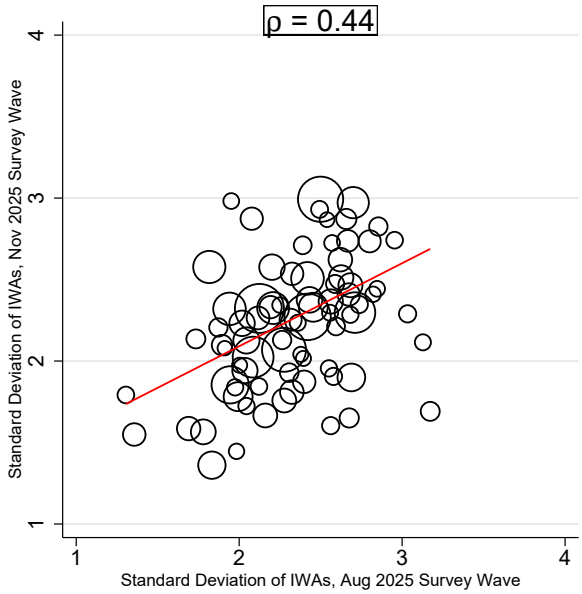
Figures 4a and 4b plot the mean and standard deviation of intermediate work activities performed in a given occupation in November 2025 against the corresponding values in August 2025. Figure 4c compares the share of intermediate work activities among all activities performed for the two surveys. These comparisons show that workers in a given occupation select tasks at similar rates across survey waves, indicating that our estimates of tasks performed within a given occupation are internally consistent.

Figure 4: Average and SD Tasks by Occupation, and Tasks Shares by Survey Waves

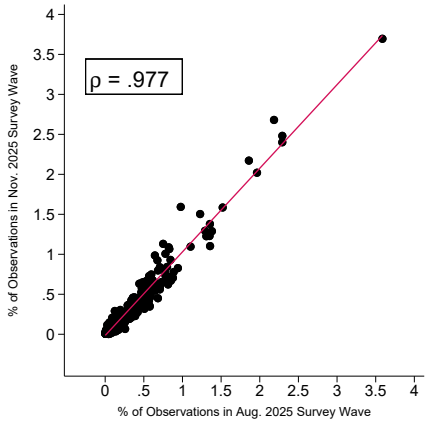
(a) Average Tasks by Occupation



(b) SD of Tasks by Occupation



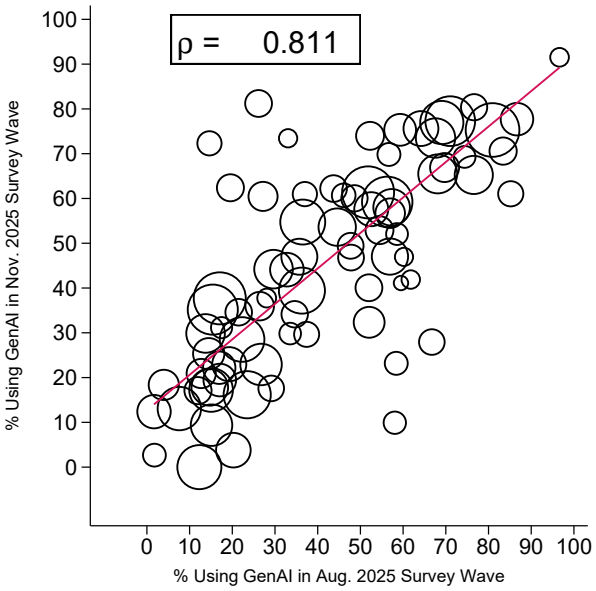
(c) Task Shares



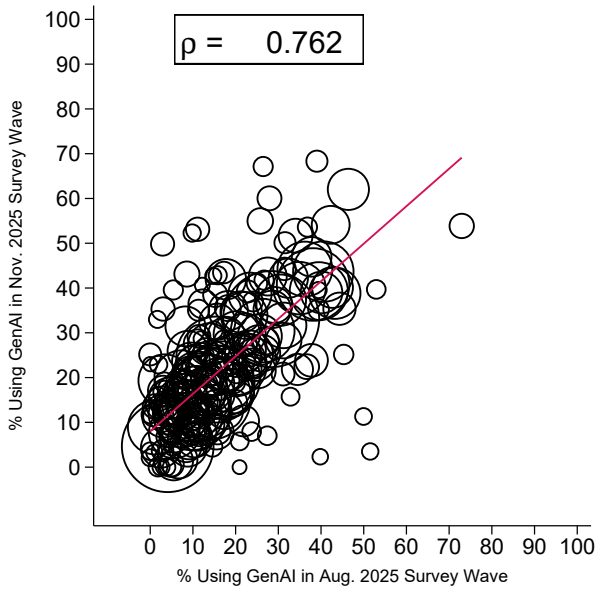
Notes: Panel (a) and (b) plots the mean number and standard deviation of intermediate work activities associated with each occupation in the November 2025 survey wave against the corresponding mean in the August 2025 wave. Each observation is an ONET SOC title with at least 10 observations. Panel (c) plots the share of each task among all tasks performed. Each observation is an intermediate work activity. The values inside the boxes, ρ, report the correlation coefficient between August 2025 and November 2025s.

Figure 5: Consistency of Occupation- and Task-Level genAI Adoption Rates

(a) Occupation-Level Adoption Rates



(b) Task-Level Adoption Rates



Notes: In Panel (a), an observation is an occupation at the CPS detailed (3-digit) level. In Panel (b), an observation is an O*NET Intermediate Work Activity. Circle sizes correspond to observation weights. We include only observations with at least 10 observations in each survey wave. The number inside the box, ρ, is the weighted correlation coefficient between adoption in the August 2025 and November 2025 waves.

2.5 Validating Occupation- and Task-Level genAI Adoption Rates

We construct an occupation-level measure of genAI adoption, which indicates the share of people in a CPS detailed (3-digit) occupation who report using genAI for work. Figure 5a shows that occupation-level genAI adoption rates are strongly correlated across two waves of the RPS, in August 2025 and November 2025. Overall, the correlation coefficient is 0.811.

We construct a task-level measure of genAI adoption, which is the share of workers who use genAI for a given task (among all workers who perform that task). Figure 5b shows that task-level genAI adoption rates are strongly correlated across two waves of the RPS, in August 2025 and November 2025. Overall, the correlation coefficient is 0.762.

The stability of occupation- and task-level genAI adoption rates across survey waves provides evidence that our measures of genAI adoption are internally consistent across adjacent RPS survey waves and is not plagued by excessive measurement error.

Note on remaining sections

In the remainder of the paper, we present a set of novel facts and brief interpretations of our findings. This analysis is preliminary and will be expanded in future drafts.

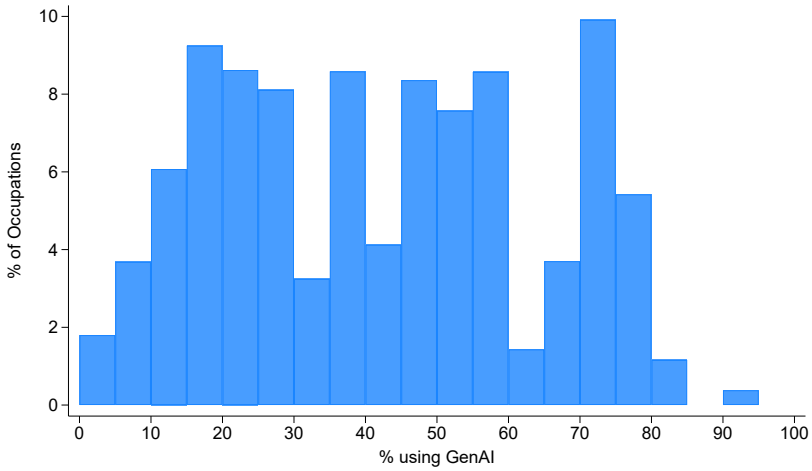
3 What Work Does genAI Do?

This section characterizes how genAI is actually being used across occupations and tasks. We document substantial heterogeneity in adoption rates: high-adoption occupations are concentrated in management, finance, and computing, while high-adoption tasks involve analysis, writing, and evaluation. In contrast, occupations and tasks involving physical, manual, or in-person work exhibit low adoption rates. These patterns indicate that genAI is primarily assisting cognitive and information-intensive work.

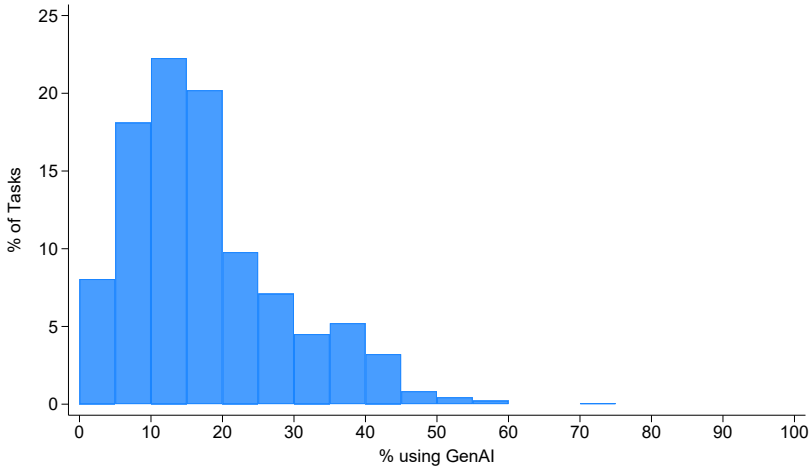
GenAI Adoption Rates Vary Substantially Across Occupations and Tasks

Figure 6: Distribution of genAI Adoption Rates by Occupation and Task

(a) genAI Adoption Rates by Occupation



(b) genAI Adoption Rates by Task



Notes: Tasks are at the Intermediate Work Activity (IWA) level. Figure includes only occupations (tasks) with at least 20 observations in the pooled Aug. 2025 and Nov. 2025 RPS survey waves. Tasks are weighted by the number of observations in the pooled Aug. 2025 and Nov. 2025 RPS survey waves, and occupations are weighted by the number of people. Occupation is at the CPS detailed (3-digit) occupation level.

Figure 6 shows that genAI adoption rates vary strongly by occupation and by task. Figure 6a shows that 12% of three-digit occupations have adoption rates below 10%, while 22% have adoption rates exceeding 60%. Figure 6b shows that 27% of tasks have adoption rates below 10%, while 8% of tasks have adoption rates exceeding 40%.

GenAI adoption is highest for cognitive information-intensive work and lowest for physical and in-person work

Table 1: Occupations and Tasks With Highest and Lowest genAI Adoption Rates

Top 10 occupations		**Top 10 tasks**	
Occupation	**Rate (%)**	**Task**	**Rate (%)**
Network and computer systems administrators	93.6	Develop models of systems, processes, or products.	71.0
Financial managers	82.7	Analyze business or financial risks.	57.7
Producers and directors	78.4	Analyze scientific or applied data using mathematical principles.	56.7
Computer and information systems managers	77.7	Manage human resources activities.	51.9
Personal financial advisors	76.1	Analyze market or industry conditions.	49.7
Lawyers	75.6	Edit written materials or documents.	47.7
Chief executives	73.8	Direct scientific or technical activities.	47.5
Software developers	72.8	Develop operational or technical procedures or standards.	45.5
Computer support specialists	72.6	Analyze performance of systems or equipment.	43.7
Management analysts	71.4	Design computer or information systems or applications.	43.4
Bottom 10 occupations		**Bottom 10 tasks**	
Occupation	**Rate (%)**	**Task**	**Rate (%)**
Cooks	16.3	Clean medical equipment or facilities.	3.3
Miscellaneous agricultural workers	13.8	Operate agricultural or forestry equipment.	3.0
Stockers and order fillers	12.6	Operate office equipment.	2.8
Maids and housekeeping cleaners	11.9	Operate lifting or moving equipment.	2.7
Janitors and building cleaners	10.5	Apply protective solutions or coatings.	2.6
Landscaping and groundskeeping workers	9.5	Prepare medical equipment or work areas for use.	2.1
Driver/sales workers and truck drivers	6.8	Perform agricultural activities.	1.9
Tellers	5.9	Escort others.	1.9
Fast food and counter workers	4.9	Repair electrical or electronic equipment.	1.7
Receptionists and information clerks	2.3	Prescribe medical treatments or devices.	0.8

Notes: Occupations are defined using CPS detailed (3-digit) occupation codes. Tasks are Intermediate Work Activities (IWA). "Rate" for occupations is the share of workers in that occupation who report using genAI at work (pooled Aug. 2025 and Nov. 2025 RPS survey waves). "Rate" for tasks is the share of worker-task observations in which the worker reports using genAI for that task. The occupation (task) panel excludes occupations (tasks) with fewer than 20 observations in the pooled waves.

Table 1 shows the top 10 and bottom 10 occupations and tasks by adoption rate, among occupations and tasks with at least 20 observations. The top occupations are primarily computing, finance, and professional occupations, while the bottom occupations are service, agricultural, and manual occupations. The top tasks involve analytical, modeling, and information-processing activities, while the bottom tasks involve physical, operational, or medical procedures with limited scope for automation.

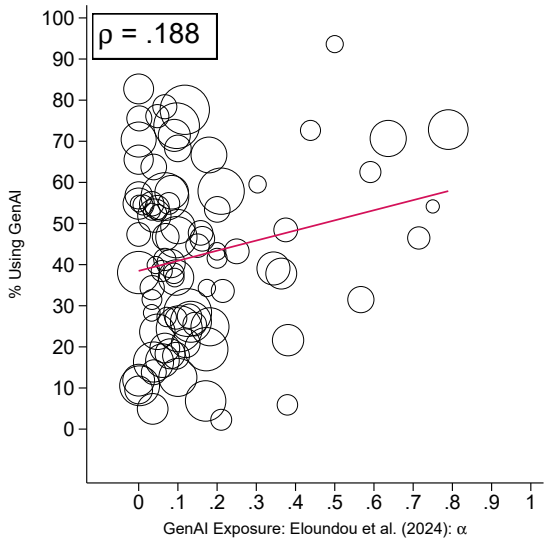
4 GenAI Exposure Predictions versus Actual genAI Adoption

This section compares our realized adoption measures to widely-used genAI exposure predictions from Eloundou, Manning, Mishkin, and Rock (2024) and Felten, Raj, and Seamans (2021). We find that exposure measures are positively but imperfectly correlated with actual adoption at both the occupation and task level. Occupation-level correlations are stronger than task-level correlations. This could reflect variation in tasks within occupation or that occupations vary along dimensions besides task content (*e.g.*, demographic or other individual characteristics) that predict adoption. We document which broad occupation and task categories deviate from exposure predictions, identifying domains where adoption exceeds or falls short of what existing measures would suggest. Finally, we show that our occupation-based adoption index substantially outperforms occupational exposure measures in predicting individual workers' genAI use.

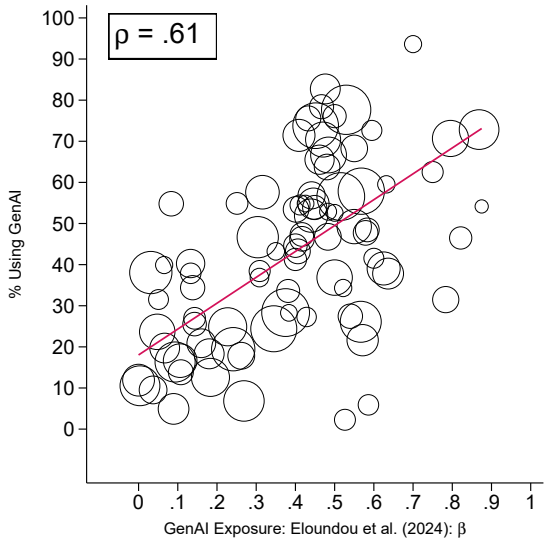
Exposure Measures are Positively but Imperfectly Correlated with Occupation-Level Adoption

Figure 7: GenAI Adoption versus Exposure Predictions by Occupation

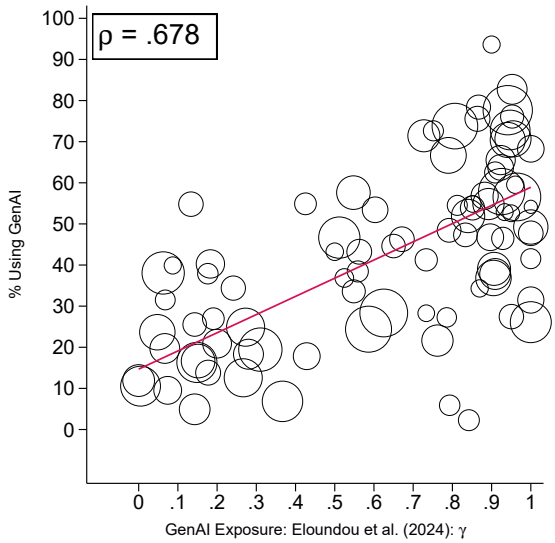
(a) Eloundou, Manning, Mishkin, and Rock (2024): α Prediction



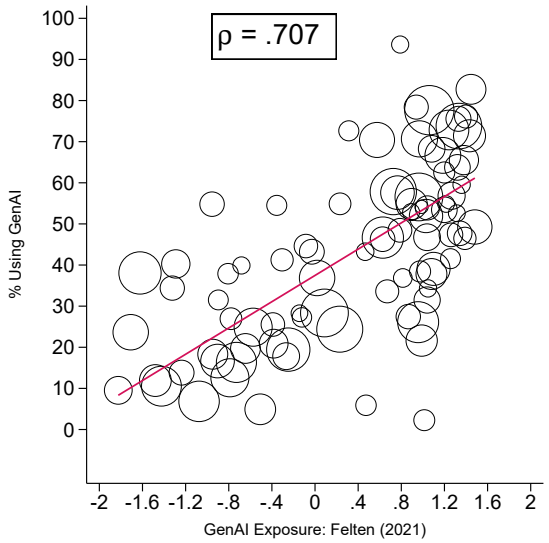
(b) Eloundou, Manning, Mishkin, and Rock (2024): β Prediction



(c) Eloundou, Manning, Mishkin, and Rock (2024): γ Prediction



(d) Felten, Raj, and Seamans (2021) Exposure



Notes: Each observation in this regression is an observation, with tasks at the CPS detailed occupation level. Circle sizes correspond to regression weights (size of the occupation in people in the pooled Aug. 2025 and Nov. 2025 RPS survey waves). Figure includes only occupations with at least 20 observations in the pooled data. The number inside the box, ρ, is the weighted correlation coefficient between occupation-level adoption and the given Eloundou, Manning, Mishkin, and Rock (2024) task measure.

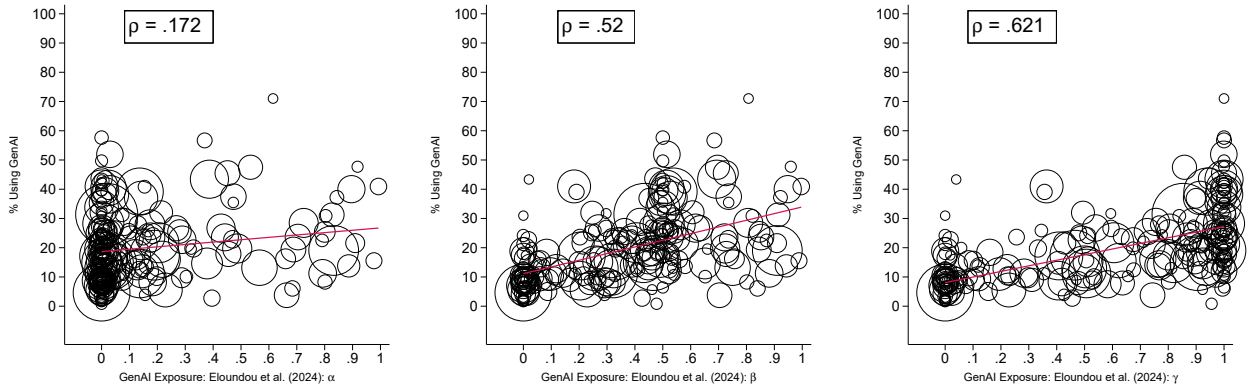
Figure 7 compares occupation-level genAI adoption rates from the RPS to four predictions of genAI exposure. Figures 7a - 7c compare genAI adoption to three measures of exposure from Eloundou, Manning, Mishkin, and Rock (2024): α, β, and γ, which progressively assume greater genAI capabilities. Figure 7d compares genAI adoption to an exposure measure from Felten, Raj, and Seamans (2021). The α measure is weakly correlated with actual adoption. The β and γ measures from Eloundou, Manning, Mishkin, and Rock (2024) and the measure from Felten, Raj, and Seamans (2021) are more strongly correlated with adoption. These results help validate that these widely-used exposure predictions are predictive of which occupations are actually assisted by genAI.

At the same time, all measures have correlation coefficients substantially below 1, indicating substantial gaps between predicted exposure and actual adoption thus far. In particular, the exposure measures alone can at most predict 50% of the variation in adoption rates across occupations.

Exposure Measures are Positively but Imperfectly Correlated with Task-Level Adoption

Figure 8: Eloundou, Manning, Mishkin, and Rock (2024) Exposure Predictions versus Adoption by Work Task

(a) α Prediction (b) β Prediction (c) γ Prediction



Notes: Each observation in this regression is a task, with tasks at the Intermediate Work Activity (IWA) level. Circle sizes correspond to regression weights, (number of appearances of the task in the pooled Aug. 2025 and Nov. 2025 RPS survey waves). Figure includes only tasks with at least 20 observations in the pooled data. The number inside the box, ρ, is the weighted correlation coefficient between task-level adoption and the given Eloundou, Manning, Mishkin, and Rock (2024) task measure.

The occupation-level exposure predictions in Eloundou, Manning, Mishkin, and Rock (2024) are derived from task-level exposure predictions. We can therefore compare their exposure predictions to actual genAI adoption rates in the RPS. (The exposure prediction in Felten, Raj, and Seamans (2021) is not conducted at the task-level.) Figure 8 displays the results. The α exposure is weakly correlated with adoption. The β and γ measures are more strongly correlated with adoption. The correlation between adoption and the β and γ exposure predictions is lower at the task-level than at the occupation level. As with the occupation-level results, we emphasize that all the correlation coefficients are substantially below 1, indicating important gaps between predicted exposure and actual adoption thus far. In particular, the exposure measures alone can at most predict 40% of the variation in adoption rates across occupations.

In the following pages, we will shed more light on which types of occupations and tasks had higher or lower adoption rates than predicted.

Largest Discrepancies between Predicted Exposure and Actual Adoption for Occupations

Table 2: Occupations Most Over- and Under-Predicted by Exposure-Based Measures

Most over-predicted occupations				Most under-predicted occupations			
Occupation	**Actual**	**Predicted**	**Gap**	**Occupation**	**Actual**	**Predicted**	**Gap**
Receptionists and information clerks	2.3	51.9	−49.7	Network and computer systems administrators	93.6	54.5	39.1
Tellers	5.9	49.8	−43.9	Automotive service technicians and mechanics	54.8	20.6	34.2
Customer service representatives	26.1	58.9	−32.8	Financial managers	82.7	56.8	25.9
Sales representatives, wholesale and manufacturing	27.5	56.7	−29.3	Producers and directors	78.4	53.0	25.4
Bookkeeping, accounting, and auditing clerks	31.5	58.9	−27.4	Computer support specialists	72.6	47.9	24.7
Office clerks, general	21.6	48.4	−26.8	Management analysts	71.4	46.9	24.6
Driver/sales workers and truck drivers	6.8	30.9	−24.1	Chief executives	73.8	50.4	23.4
Food service managers	27.3	49.4	−22.2	Lawyers	75.6	52.9	22.7
Farmers, ranchers, and other agricultural managers	28.3	47.1	−18.8	Preschool and kindergarten teachers	54.9	33.5	21.4
Business operations specialists, all other	34.3	53.1	−18.8	Hairdressers, hairstylists, and cosmetologists	39.9	18.5	21.4

Notes: Occupation is at the CPS detailed (3-digit) occupation level. This table reports occupations with the largest negative (over-predicted) and positive (under-predicted) residuals from a regression of genAI adoption rates on exposure scores constructed using Eloundou et al. (2024). "Actual" is the share of workers in the occupation reporting genAI use at work. "Predicted" is the fitted value based on exposure. "Gap" is the difference between actual and predicted adoption (actual minus predicted). Rows are ordered by the absolute value of the residual. All statistics are based on pooled Aug. 2025 and Nov. 2025 survey waves.

In Table 2 we regress genAI adoption on the occupation-level on the occupation-level γ from Eloundou, Manning, Mishkin, and Rock (2024). Occupations with the largest positive (negative) residuals have higher (lower) adoption rates than one would predict from exposure measure. "Over-predicted" occupations where adoption is high relative to predictions are largely clerical and administrative tasks. "Under-predicted" occupations are generally management and professional occupations, along with some service occupations involving information gathering.

Largest Discrepancies between Predicted Exposure and Actual Adoption for Tasks

Table 3: Tasks Most Over- and Under-Predicted by Exposure-Based Measures

Most over-predicted tasks				Most under-predicted tasks			
Task	**Actual**	**Predicted**	**Gap**	**Task**	**Actual**	**Predicted**	**Gap**
Prescribe medical treatments or devices.	0.8	26.4	−25.5	Develop models of systems, processes, or products.	71.0	27.2	43.8
Examine materials or documentation for accuracy or compliance.	6.1	27.1	−21.1	Position tools or equipment.	43.3	8.8	34.5
Develop organizational or program goals or objectives.	6.4	27.2	−20.8	Analyze business or financial risks.	57.7	27.2	30.4
Monitor environmental conditions.	5.7	25.8	−20.1	Analyze scientific or applied data using mathematical principles.	56.7	27.2	29.5
Notify others of emergencies or problems.	3.8	22.3	−18.6	Resolve personnel or operational problems.	41.0	15.0	25.9
Verify personal information.	10.1	27.2	−17.1	Manage human resources activities.	51.9	27.2	24.7
Examine financial activities, operations, or systems.	10.7	27.2	−16.5	Set up computer systems, networks, or other information systems.	39.1	14.9	24.2
Maintain sales or financial records.	11.2	26.7	−15.4	Direct scientific or technical activities.	47.5	24.5	23.0
Assist others to access additional services or resources.	12.1	27.2	−15.2	Manage agricultural or forestry operations.	31.0	8.1	22.9
Process shipments or mail.	8.3	23.4	−15.1	Analyze market or industry conditions.	49.7	27.2	22.5

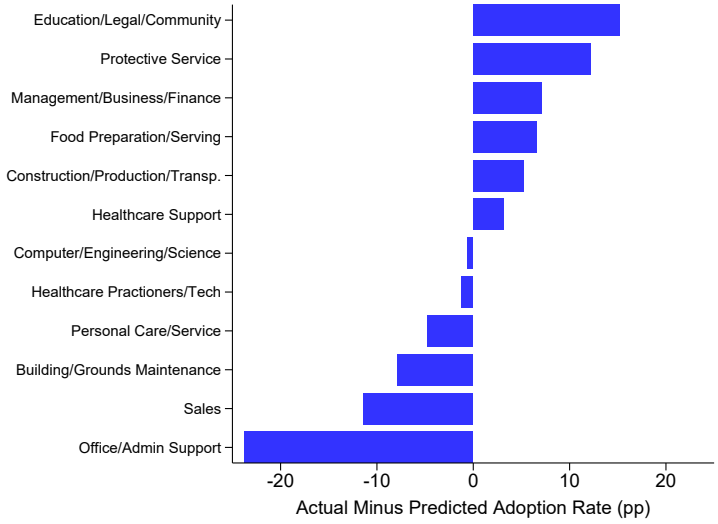
Notes: This table reports tasks with the largest negative (over-predicted) and positive (under-predicted) residuals from a regression of task-level genAI adoption rates on exposure scores constructed using Eloundou, Manning, Mishkin, and Rock (2024)'s γ measure. "Actual" is the share of worker-task observations in which the worker reports using genAI for that task. "Predicted" is the fitted value based on exposure. "Gap" is the difference between actual and predicted adoption (actual minus predicted). Rows are ordered by the absolute value of the residual. All statistics are based on pooled Aug. 2025 and Nov. 2025 survey waves.

In Table 3 we regress AI adoption on the task-level on the task-level γ from Eloundou, Manning, Mishkin, and Rock (2024). Tasks with the largest positive (negative) residuals have higher (lower) adoption rates than one would predict from exposure measure. "Over-predicted" tasks where adoption is high relative to predictions seem to be largely diagnostic or clerical tasks with clear privacy concerns, *e.g.* involving medical or legal records. "Under-predicted" tasks are generally the ones with the highest adoption rates and are primarily white collar tasks involving planning and coordinating.

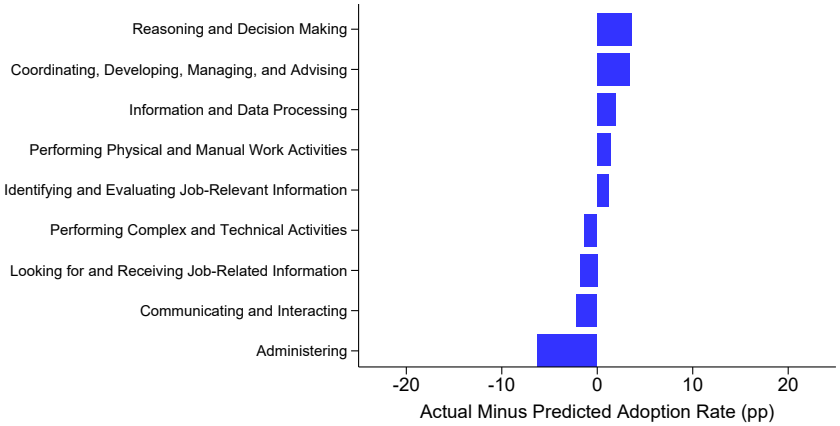
Discrepancies between Predicted Exposure and Actual Adoption for Broad Occupations & Tasks

Figure 9: Predicted genAI Exposure versus Actual Adoption: Broader Categories

(a) Broad Occupation



(b) Broad Task



Notes: Eloundou, Manning, Mishkin, and Rock (2024)'s measure at the occupation (task) level occupation (task)-level genAI adoption rate are aggregated to the CPS broad occupation level/broad work activity (BWA) level. Actual adoption rates are then regressed onto Eloundou, Manning, Mishkin, and Rock (2024)'s measure at this level. "Predicted Adoption" has the fitted values for adoption rates from this regression and "Actual Adoption" has the actual adoption rates. Occupations (tasks) where the adoption rate exceeds the prediction based on the exposure score exhibit higher adoption than would be expected based on Eloundou, Manning, Mishkin, and Rock (2024)'s measure.

Figures 7 and 8 showed that many occupation-level and task-level adoption rates differed from predicted exposure in Eloundou, Manning, Mishkin, and Rock (2024). In Tables 2 and 3 we zoomed in on the top 10 over- and under-predicted tasks. In Figure 9, we take a broader perspective. We collapse detailed tasks into nine broad work activity groups and aggregate occupations into CPS broad occupation categories. To facilitate comparison across measures, exposure predictions are converted to a predicted adoption rate based on the fitted values of a linear regression of adoption rates onto exposure. Figure 9 displays the results.

"Education / Legal / Community" are the most under-predicted occupations, while and "Office / Admin Support" are the most over-predicted occupations. At least for the latter, this goes hand in hand with "Administering" being the most-overpredicted task, while "Reasoning and Decision Making" is the most-underpredcited task.

Low adoption rates relative to exposure may reflect barriers to adoption not captured by exposure predictions. High adoption rates relative to exposure may reflect unanticipated capabilities of genAI or subtle aspects of ONET tasks that genAI can assist that were not evident when exposure predictions were created. We plan to further investigate these deviations, which may occur for two reasons. First, within a broad occupation group, workers may be more likely to perform high-exposure tasks than is implied by ONET. Second, the genAI adoption rate for a given task may exceed its exposure score. (This same logic holds in reverse for occupation groups with genAI adoption rates below their exposure score.)

Occupation Fixed Effects Substantially Outperform Existing Exposure Measures in Predicting Individual Workers' genAI Use.

Table 4: Predicting genAI Adoption at the Worker Level

Controls	Adjusted R^2
Eloundou, Manning, Mishkin, and Rock (2024) Exposure Measure	0.064
Felten, Raj, and Seamans (2021) Exposure Measure	0.084
Detailed Occupation Fixed Effects	0.203

Notes: Observations are at the worker level and the dependent variable in all regressions is an indicator for whether a worker uses genAI at work. Uses pooled Aug. 2025 and Nov. 2025 survey waves, and all regressions include a fixed effect for survey wave.

Table 4 evaluates how well different measures predict whether an individual worker actually uses genAI for their job. We compare three measures: (i) the occupation-level exposure measure from Eloundou, Manning, Mishkin, and Rock (2024), which aggregates predicted task exposure to the occupation level; (ii) the occupation-level exposure measure from Felten, Raj, and Seamans (2021); and (iii) detailed occupation fixed effects, which are equivalent to our adoption measure. Consistent with the evidence in Figures 7 and 8, all three measures are predictive of worker-level adoption. However, using detailed occupation fixed effects yields an adjusted R^2 of 0.202 versus 0.064 using Eloundou, Manning, Mishkin, and Rock (2024) exposure and 0.084 using Felten, Raj, and Seamans (2021).

5 Task-Level Adoption: Data from Surveys versus Chat Logs

This section discusses an alternative approach to measuring task-level genAI use based on large-scale conversation logs, and explains how our survey-based adoption measures complement this approach.

Two recent papers construct task-based indexes using conversation data: Handa et al. (2025) using chat data from Claude.ai and Chatterji et al. (2025) using chat data from ChatGPT. Both classify prompts into O*NET tasks based on conversational context and characterize how genAI usage is distributed across imputed tasks and occupations.

A clear advantage of the conversation-based approach is that it relies on revealed behavior rather than self-reports, and therefore does not suffer from recall error or subjective reporting biases. In addition, because it is not constrained by a pre-selected list of tasks, it can capture a wider range of activities than survey-based approaches.

At the same time, conversation-based measures face several important limitations. First, imputing an O*NET task from a single prompt inevitably introduces classification error. Second, inferring a user's occupation or whether a prompt pertains to work is inherently challenging, since many prompts reflect general-purpose or ambiguous tasks. The work versus non-work distinction is quantitatively important: non-work genAI use accounts for a majority of total use (see Figure 2, which shows that roughly 55% of genAI hours reflect non-work use). Third, usage patterns from a single product may reflect its specific capabilities or user base rather than nationally representative use.

Finally, for many economic questions, the relevant object is not the share of genAI prompts associated with a task, but rather the share of instances in which a given task is assisted by genAI. Conversation data observe only assisted instances and therefore cannot distinguish between tasks that generate many prompts because they are common and tasks that generate many prompts because they are frequently assisted by genAI.

For these reasons, we view conversation-based and survey-based measures as complementary. In ongoing work, we plan to obtain the task-based indexes in Handa et al. (2025) and Chatterji et al. (2025), construct analogous survey-based measures capturing a task's share of genAI use (rather than its adoption rate), and investigate sources of disagreement where the measures diverge.

6 Individual Determinants of genAI Use

This section shows that individual characteristics play a central role in genAI adoption, above and beyond task characteristics. Worker fixed effects dramatically increase explanatory power for task-level adoption, indicating that unobserved individual factors such as skills, beliefs, or prior experience are key. That is, our ability to predict whether a given worker uses genAI for a given task is improved if we take into account not only which task it is, but also which worker it is. We then provide three pieces of evidence consistent with a fixed-cost / learning mechanism: workers with high-adoption tasks in their portfolio are more likely to adopt genAI for any given task; experienced genAI users use it for more tasks; and plausibly exogenously induced workplace adoption spills over into home use.

Unobserved Worker-Level Factors are a Key Driver of Adoption

Table 5: Predicting Adoption: Individuals vs. Tasks

Controls	Adjusted R^2
Task Fixed Effects	0.090
Task Fixed Effects + Demographic Controls	0.112
Task Fixed Effects + Individual Fixed Effects	0.431

Notes: Tasks are at the Intermediate Work Activity (IWA) level. Observations are at the worker-task level, and a given observation is whether the worker who is doing that task reports using genAI for it. Includes only people who do at least two tasks. Each regression has 27,669 person-task observations. Demographic Controls consist of education, race, sex, and a quadratic for age.

Table 5 regresses task-level genAI adoption on task fixed effects and various individual-level controls. Task fixed effects alone generate an adjusted R^2 of 0.09, indicating that 9% of the variation in adoption occurs between tasks. This reflects the fact that adoption rates vary substantially across tasks, and yet in most tasks adoption is not uniformly zero or one.

What can explain the substantial variation in adoption decisions within tasks? One natural candidate is that, for some reason, individuals vary in their propensity to adopt genAI. Individual variation in the propensity to adopt could reflect differences in demographic characteristics, such as sex, age, race, and education. However, the second row of Table 5 shows that controlling for these characteristics only marginally increases the share of adoption variation explained.

The third row of Table 5 shows that including individual fixed effects dramatically increases the adjusted R^2 to 0.431. (Note that in this regression an observation is at the worker-task level, which is why we can include individual fixed effects.) This sharp increase implies that a large share of variation in task-level genAI adoption reflects persistent individual differences rather than task characteristics or observable demographics. In other words, for a given task, some workers are systematically more likely to adopt genAI than others. This suggests that individual characteristics, such as skills, beliefs, learning costs, willingness to experiment, or work environment, play a central role in adoption choices.

We now ask why some workers systematically adopt genAI more than others.

Workers who Perform High-Adoption Tasks are More Likely to Adopt genAI for Other Tasks, Consistent with Cross-Task Learning.

Table 6: Regression of Task Adoption onto Mean Other-Task Adoption Index

	Adopted genAI
Other-Task Adoption Index	0.0399***
	(0.00990)
Adjusted R^2	0.172
Observations	26,703

Notes: Tasks are at the Intermediate Work Activity (IWA) level. Observations are at the worker-task level, and a given observation is whether the worker who is doing that task reports using genAI for it. Dependent variable is an indicator for if a worker reports using genAI for that task. Other-Task Adoption Index is constructed as the leave-one-out mean adoption rate for a worker-task over all of the tasks a worker performs besides that task, then z-scored. Regression includes controls for education, detailed occupation, survey wave, sex, race, a quadratic for age, the number of tasks the worker does, and task fixed effects. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

One possibility for why some workers systematically adopt genAI for more tasks than others is that a worker's initial adoption choice may involve fixed costs. These fixed costs may include learning how to communicate with genAI effectively and understanding its capabilities and limitations. Workers who have paid these fixed costs because they perform a task in which genAI excels may then be more likely to extend use to other tasks or domains. We now provide three related but distinct pieces of evidence consistent with this mechanism.

First, Table 6 shows that workers are more likely to use genAI for a given task if their other tasks have high adoption rates. The key explanatory variable, the Other-Task Adoption Index, is the average genAI adoption rate across the worker's other tasks, calculated excluding the worker's own adoption decisions. The regression controls for detailed occupation, demographics, and task fixed effects, ensuring that we compare workers performing the same task but who differ in the genAI intensity of their broader task portfolios.

For intuition, consider two workers who both perform a moderately genAI-suited task such as writing. One worker also performs software coding tasks, where genAI is highly effective, while the other performs manual tasks, where genAI offers little benefit. The positive coefficient in the regression indicates that, on average, the worker who also codes is more likely to use genAI for writing (compared to the worker whose other task is manual). This pattern suggests that experience with high-return tasks lowers barriers to adoption in other contexts.

Workers who have Used genAI for At Least Six Months Use it for More Rasks

Table 7: Prior genAI Experience and Task-Level Adoption

	Number of genAI-Assisted Tasks
Have used genAI for at least 6 months	0.456***
	(0.168)
Adjusted R^2	0.476
Observations	3,253

Notes: Regression is at the worker level and includes only workers who report using genAI for at least one task. The dependent variable is the number of genAI-assisted tasks the worker performs. The key independent variable is an indicator for whether the worker reports having used genAI for at least six months. The regression controls for a quadratic in age, detailed occupation, survey wave fixed effects, sex, and the total number of tasks performed by the worker. Regression includes a constant term that is not in the table. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

A key prediction of the fixed-cost interpretation is that genAI adoption should expand across tasks as workers gain experience with the technology. If workers incur learning or setup costs when they first adopt genAI, but not when extending its use to other tasks, then we would expect to see workers gradually adopt genAI for more tasks as they gain experience.

Table 7 provides evidence consistent with this prediction. Among workers who report using genAI for at least one task, those who have used genAI for six months or more report using it for a significantly larger number of work tasks. The estimated effect is economically meaningful and remains robust after controlling for detailed occupation, demographics, survey wave fixed effects, and the total number of tasks the worker performs.

Plausibly Exogenous Variation in Work-Use of genAI Increases Home Use.

Table 8: 2SLS Regression of Non-Work Usage Onto Work Usage

	Uses genAI for Non-Work Purposes
Uses genAI for Work	0.522***
	(.070)
Observations	2,253
Adjusted R^2	0.216

Notes: Regression is at the worker level, and instruments for work usage with employer encouragement of work usage. Regression includes controls for detailed occupation, industry, sex, education, college major, and survey wave fixed effects. Regression includes a constant term that is not in the table. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Finally, Table 8 examines whether genAI adoption spills over across domains within workers by estimating the effect of workplace genAI adoption on home adoption. Because work adoption may be endogenous to individual preferences or unobserved traits, we instrument for genAI use at work using whether the respondent's employer encourages genAI use. The regression also includes detailed controls for occupation, industry, education, college major, demographics, and survey wave fixed effects.

The results show a strong positive relationship: workers induced to use genAI at work are significantly more likely to use genAI at home. This finding provides suggestive evidence of a within-person spillover across contexts. One interpretation is that plausibly exogenous exposure to genAI in the workplace induces the worker to pay fixed costs associated with adoption, leading workers to apply the technology more broadly in their non-work activities. This result reinforces the broader pattern in this section that genAI adoption in one domain increases the probability that it is adopted in a separate domain.

7 Conclusion

This paper measures what work generative AI actually does by linking survey-based adoption data to detailed O*NET tasks. We document substantial heterogeneity in genAI adoption across tasks and occupations: high-adoption work is concentrated in cognitive and information-intensive activities such as analysis, writing, and evaluation, while physical, manual, and in-person work exhibits low adoption. These patterns are correlated with existing exposure predictions, but exposure measures leave the majority of variation in adoption unexplained.

A key contribution of this paper is the construction of task- and occupation-level genAI adoption indexes that can be used by researchers to study the economic effects of genAI. Unlike exposure measures, which capture potential applicability based on task content, our indexes reflect how workers are actually using the technology in practice. We show that these adoption-based measures substantially outperform exposure measures in predicting individual workers' genAI use.

We also provide evidence that adoption patterns reflect not only task characteristics but also individual-level factors. Worker fixed effects explain far more variation in task-level adoption than task characteristics alone, and we document patterns consistent with a learning mechanism: workers who use genAI for high-adoption tasks are more likely to adopt it for other tasks, experienced users apply genAI to more tasks, and workplace adoption spills over into home use.

Our methodology can be updated as genAI capabilities and adoption patterns evolve, providing a continuously updated input for research on the labor market effects of this rapidly diffusing technology.

References

Autor, D. H., F. Levy, and R. J. Murnane (2003). "The skill content of recent technological change: An empirical exploration". In: *The Quarterly journal of economics* 118.4, pp. 1279–1333.

Bick, A. and A. Blandin (2023). "Employer reallocation during the COVID-19 pandemic: Validation and application of a do-it-yourself CPS". In: *Review of Economic Dynamics* 49, pp. 58–76.

Bick, A., A. Blandin, A. Caplan, and T. Caplan (2025a). "Heterogeneity in Work From Home: Evidence from Six U.S. Datasets". In: *Federal Reserve Bank of St. Louis Review* 107.14, pp. 1–23.

Bick, A., A. Blandin, A. Caplan, and T. Caplan (2025b). "Measuring Trends in Work From Home: Evidence from Six U.S. Datasets". In: *Federal Reserve Bank of St. Louis Review* 107.15, pp. 1–23.

Bick, A., A. Blandin, and D. J. Deming (2026). "The Rapid Adoption of Generative AI". In: *Management Science*. Published online January 20, 2026.

Brynjolfsson, E., T. Mitchell, and D. Rock (2018). "What can machines learn and what does it mean for occupations and the economy?" In: *AEA papers and proceedings*. Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 43–47.

Chatterji, A., T. Cunningham, D. J. Deming, Z. Hitzig, C. Ong, C. Y. Shan, and K. Wadman (2025). *How people use chatgpt*. Tech. rep. National Bureau of Economic Research.

Deming, W. E. and F. F. Stephan (1940). "On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known". In: *The Annals of Mathematical Statistics* 11.4, pp. 427–444.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). "GPTs are GPTs: Labor market impact potential of LLMs". In: *Science* 384.6702, pp. 1306–1308.

Felten, E., M. Raj, and R. Seamans (2021). "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses". In: *Strategic Management Journal* 42.12, pp. 2195–2217.

Handa, K., A. Tamkin, M. McCain, S. Huang, E. Durmus, S. Heck, J. Mueller, J. Hong, S. Ritchie, T. Belonax, K. K. Troy, D. Amodei, J. Kaplan, J. Clark, and D. Ganguli (2025). *Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations*. arXiv: 2503.04761 [cs.CY].

Laughlin, L., X. Song, M. Wisniewski, and J. Xu (2024). *From Job Descriptions to Occupations: Using Neural Language Models to Code Job Data*. Working Paper. U.S. Census Bureau, University of Pennsylvania, and Pennsylvania State University.

Miano, A. (Jan. 2025). *Search Costs, Outside Options, and On-the-Job Search*. Working Paper. University of Naples Federico II and CSEF.

US Census Bureau (2015). "Current Population Survey Interviewing Material". In: *https://www2.census.gov/programs-surveys/cps/methodology/intman/CPS_Manual_April2015.pdf*.

Webb, M. (2020). *The Impact of Artificial Intelligence on the Labor Market*. Working Paper. Stanford University.

ONLINE APPENDIX

A RPS: Measurement and Definitions

A.1 Sample Restrictions

The Qualtrics panel includes about 15 million members and is not a random sample of the U.S. population. However, researchers can instruct Qualtrics to target survey invitations to specific demographic groups. The RPS sample was designed to be nationally representative of the U.S. across gender, age, education, household income, and other demographic characteristics. Individuals who take the survey too fast, i.e. in less than 50% of the median time in a soft launch, or who do not state that they will provide their best answers, are automatically dropped from the sample.[1] Finally, once fielding is completed Qualtrics staff goes over all responses and filters out any that look suspicious. This is typically the case for around 10% of the responses. We also screened out responses based on various open text questions who appear to be bots.

Table A.1 compares the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure for our main surveys (columns 1 and 2). The most notable discrepancies are that individuals with household income of $50,000 or less are overrepresented, while those with more than $150,000 are underrepresented. The bottom panel of Table A.1 compares employment status in the CPS and RPS, statistics that have not been targeted in the sampling procedure. Individuals classified as employed at work and unemployed are overrepresented in the RPS.

A.2 Weighting

As described in the body of the paper, we asked Qualtrics to administer the survey to a sample of respondents who match the U.S. population along a few broad demographic characteristics: gender, five age bins (18-24, 25-34, 35-44, 45-54, 55-64), race and ethnicity (non-Hispanic White, non-Hispanic Black, Hispanic, other), education (high school or less, some college or associate degree, bachelor's degree or more), marital status (married or not), number of children in the household (0, 1, 2, 3 or more), three income bins for household income over the last 12 months (<$50k, $50k-100k, >$100k), and four Census regions. Table A.1 compare the sample composition between the pooled August and November RPS and the CPS along the demographics targeted in the sampling procedure for our main survey. The tables look very similarly for each survey wave separately, including the June pilot, and are available upon

[1] The exact phrasing of the screener question is: "We care about the quality of our survey data and hope to receive the most accurate measure of your opinions, so it is important to us that you thoughtfully provide your best answer to each question in the survey. Do you commit to providing your thoughtful and honest answers to the questions in this survey?" with the following three answer options (1) I will provide my best answers, (2) I will not provide my best answers, and (3) I can't promise either way. According to Qualtrics staff, this question provides the best results in terms of screening out respondents.

Table A.1: Sample Composition in the Pooled August and November 2025 CPS and RPS

	Everyone		*Employed*	
	CPS	RPS	CPS	RPS
	(1)	(2)	(3)	(4)
Gender: Women	50.3	51.3	47.1	47.9
Age				
18-24	15.0	13.9	11.9	14.5
25-34	22.4	20.7	24.0	22.3
35-44	22.4	25.2	24.8	26.4
45-54	19.9	22.6	21.6	22.1
55-64	20.3	17.5	17.7	14.7
Race/Ethnicity				
Non-hispanic White	55.5	57.8	56.7	58.1
Non-hispanic Black	12.8	15.5	11.9	15.5
Hispanic	21.2	19.7	21.0	19.4
Other	10.6	7.0	10.4	7.0
Education				
Highschool or less	36.3	33.0	32.3	29.5
Some college/Associate's degree	25.9	27.7	25.3	27.6
Bachelor's or Graduate degree	37.8	39.3	42.4	42.9
Marital Status: Married	50.1	48.9	53.0	50.9
Number of children				
0	61.0	57.5	60.0	55.5
1	17.6	18.7	17.9	19.8
2	13.6	14.9	14.4	15.8
3+	7.8	9.0	7.7	9.0
Household Income in Last 12 Months				
\$0-\$50,000	24.2	32.5	18.4	26.5
\$50,000-\$100,000	29.7	28.7	30.2	30.7
\$100,000-\$150,000	18.7	18.6	20.5	20.6
\$150,000+	27.5	20.2	30.9	22.3
Region				
Northeast	17.1	18.6	17.2	19.5
Midwest	20.1	17.7	20.9	17.5
South	38.7	40.1	37.9	39.5
West	24.1	23.6	24.1	23.5
Employment Status				
Employed, at work last week	71.8	76.3		
Employed, absent from work last week	2.5	3.1		
Unemployed	3.4	7.1		
Not in the labor force	22.3	13.5		
Observations	107817	9844	80005	7819

Notes: Column 1 reports the sample composition in the pooled August and November 2025 Current Population Survey (CPS) for the variables targeted by Qualtrics in the sampling procedure. The employment status was the only variable not targeted. Column 2 reports the sample composition in the pooled August and November 2025 Real-Time Population Survey (RPS). The sample in both data sets is restricted to the civilian population ages 18-64. Columns 3 and 4 report the same outcomes for the employed (at work and absent from work last week).

request. As discussed in the paper, we also compare at the bottom of the table employment status in the CPS and RPS, as well as the demographic composition for the employed. None of this latter sets of moments were targeted during the sampling of survey respondents.

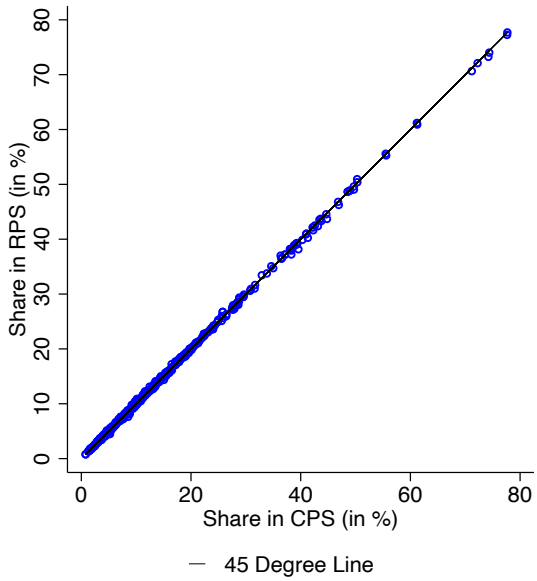
Using the iterative proportional fitting (raking) algorithm of Deming and Stephan (1940), we construct sampling weights to ensure the RPS matches the CPS sample proportions for the same set of demographic characteristics included in the Qualtrics sampling targets for the overall sample, i.e., independent of employment status. However, we use more disaggregated categories for education, marital status, and household income and we interact all categories with gender. In particular, for education, we distinguish between less than high school, high school graduate or equivalent, some college but no degree, associate degree, bachelor's degree, and graduate degree. For marital status, we distinguish between married + spouse present, divorced, never married, and "other." We also condition on relationship status (spouse living in the same household, partner living in the same household, other). For household income over the last 12 months we distinguish between \$0-\$25,000, \$25,000-\$50,000, \$50,000-\$75,000, \$100,000-\$150,000, and more than \$150,000.

In addition, our sampling weights replicate the employed-at-work rates, the employment rates, and the labor force participation rates in each of the subsequent months. We match these key labor market statistics not only in the aggregate but also conditional on demographic characteristics. More specifically, we match the employed-at-work rate, the employment rate, and the labor force participation rate for the current month by the same demographics described in the previous paragraph. We also include 2-digit occupation codes in our weighting scheme.

We also include occupation into our weighting scheme. Among the 22 broad occupations, we merge a) "Community and Social Service Occupations" and "Legal Occupations", b) "Healthcare Support Occupations" and "Protective Service Occupations", and c) "Farming, Fishing, and Forestry Occupations" and "Construction and Extraction Occupations." to ensure a sufficient sample size

Including all the interaction terms, we have a total of 48 statistics (e.g., gender is one, and gender x education is another) which we weight on, with a combined 295 categories (e.g., gender has two categories, and gender x education has 2 x 6 = 12 categories. To visualize the goodness of fit, we plot in Figure A.1 for all statistics used in the weighting scheme, the weighted fraction of individuals with the respective characteristics in the RPS (on the y-axis) and the CPS (on the x-axis) for August and November 2025. The lack of sizeable deviations from the 45-degree line demonstrates how well the weighting procedure works. The results are similar for each survey wave separately, including the June pilot, and are available upon request.

FIGURE A.1: Sample Composition in the Weighted RPS vs. CPS for August and November 2025



Notes: The figure shows for all statistics used in the weighting scheme, the weighted fraction of individuals with the respective characteristics in the RPS (on the y-axis) and the CPS (on the x-axis).

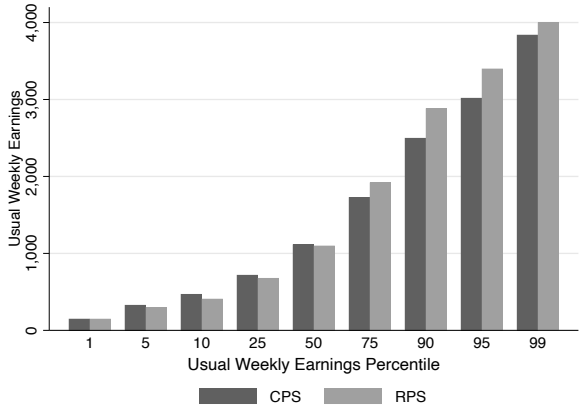
A.3 Validation Checks

To illustrate how the RPS and CPS compare for characteristics that are not part of the sampling scheme and the effect of weighting, panels (a) and (b) of Figure A.2 compare the usual weekly earnings distribution in the RPS and CPS in the pooled August and November 2025 waves, with and without weights, respectively. Both the unweighted and weighted distributions are very similar.

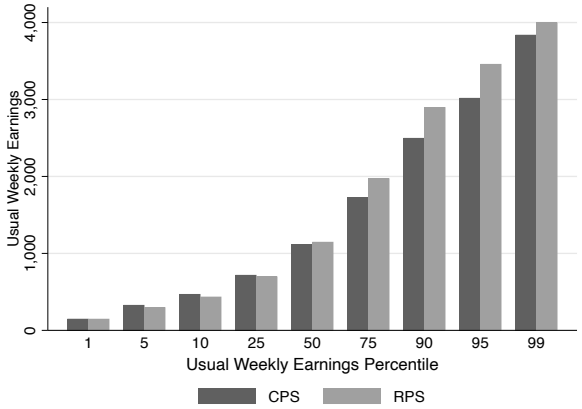
Figure A.3 compare the industry shares in the RPS and CPS in the pooled August and November 2025 waves, with and without weights, as well as college major shares in the RPS and 2022 American Community Survey (ACS). The unweighted correlations of the industry shares is 0.73, comparable to those for occupations. The correlation of college majors in the RPS and the ACS is 0.77. The major with the highest share relative to the ACS is Computers, Mathematics, & Statistics, while the major with the lowest share relative to the ACS is Science and Engineering Related Fields. Weighting the data marginally improves the fit between the RPS and the two other data sets.

FIGURE A.2: Validation Checks – Usual Weekly Earnings and Occupation Shares

(a) Weekly Earnings Percentiles: Unweighted



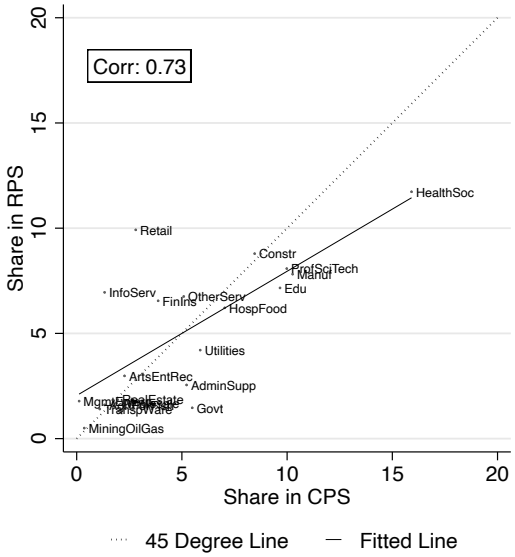
(b) Weekly Earnings Percentiles: Weighted



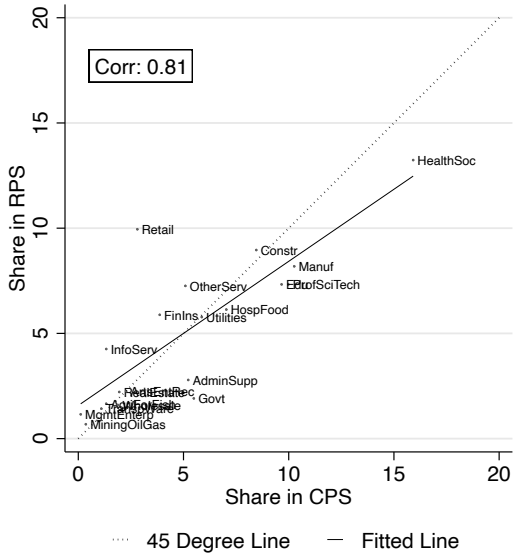
Notes: The figure on the left uses unweighted RPS data, the figure on the right uses weighted RPS data. We use the same sample of RPS respondents in both figures. All figures use weighted CPS data. Data samples for the weekly earnings figures are employees ages 18-64 in the pooled August and November 2025 RPS and CPS-ORG with weekly earnings below the CPS topcode of $4300 and an implied hourly wage of at least the federal minimum wage of $7.25. Since the CPS topcode varies across month, we apply the lower topcode in both months.

FIGURE A.3: Validation Checks – Industry and Major Shares

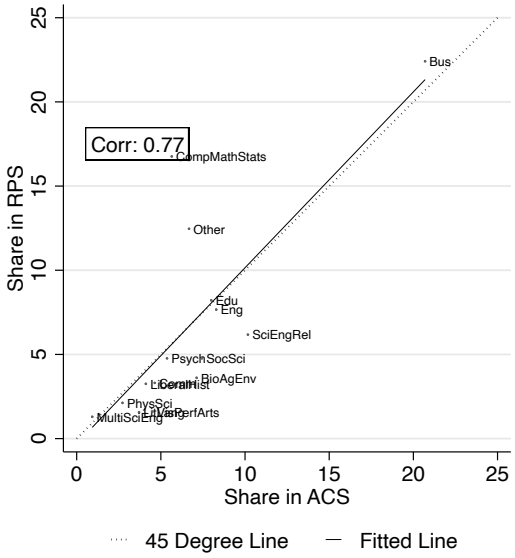
(a) Industry: Unweighted



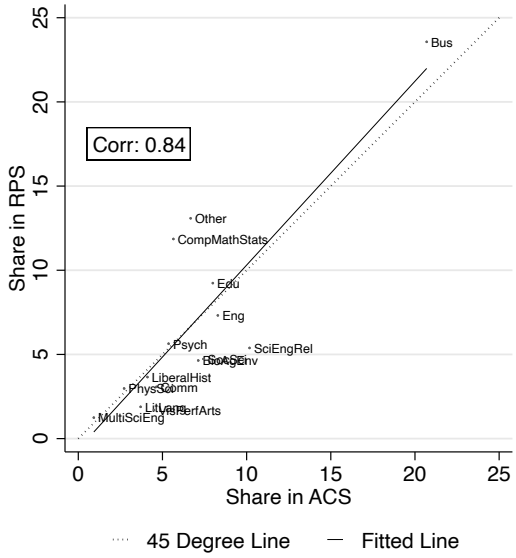
(b) Industry: Weighted



(c) Major Shares: Unweighted



(d) Major Shares: Weighted



Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the same sample of RPS respondents in both figures. All figures use weighted CPS and ACS data, respectively. Data samples for the industry comparison are employed respondents ages 18-64 in the pooled August and November 2025 RPS and CPS with sample sizes of 7181 and 80005 , respectively. Data samples for the college major comparison are respondents with a college degree ages 18-64 in the pooled August and November 2025 RPS with sample sizes of 3627 and 80005 in the 2022 ACS, respectively.

A.4 Definition of Industries and Occupations

Industries correspond to the 22 major industries in the NAICS: Agriculture, Forestry, Fishing, and Hunting (NAICS=11); Mining, Quarrying, and Oil and Gas Extraction (NAICS=21); Utilities (NAICS=22); Construction (NAICS=23); Manufacturing (NAICS=31-33); Wholesale Trade (NAICS=42); Retail Trade (NAICS=44-45); Transportation and Warehousing (NAICS=48-49); Information (NAICS=51); Finance and Insurance (NAICS=52); Real Estate and Rental and Leasing (NAICS=53); Professional, Scientific, and Technical Services (NAICS=54); Management of Companies and Enterprises (NAICS=55); Administrative and Support and Waste Management Services (NAICS=56); Educational Services (NAICS=61); Health Care and Social Assistance (NAICS=62); Arts, Entertainment, and Recreation (NAICS=71); Accommodation and Food Services (NAICS=72); Other Services (except Public Administration) (NAICS=81); Public Administration (NAICS=99).

The industry groupings used in the main text figures are: Agriculture / Extraction (sectors 11, 21), Construction (23), Manufacturing (31-33), Wholesale / Retail Trade (42, 44-45), Transportation / Utilities (22, 48-49), Info Services (51), Finance / Real Estate (52, 53), Professional / Business Services (54, 55, 56), Education / Health/ Social / Public Services (61, 62, 92), Leisure / Accommodation / Other (71, 72, 81).

Occupations correspond to the 22 major occupations in the SOC: Management Occupations (SOC=11); Business and Financial Operations Occupations (SOC=13); Computer and Mathematical Occupations (SOC=15); Architecture and Engineering Occupations (SOC=17); Life, Physical, and Social Science Occupations (SOC=19); Community and Social Service Occupations (SOC=21); Legal Occupations (SOC=23); Educational Instruction and Library Occupations (SOC=25); Arts, Design, Entertainment, Sports, and Media Occupations (SOC=27); Healthcare Practitioners and Technical Occupations (SOC=29); Healthcare Support Occupations (SOC=31); Protective Service Occupations (SOC=33); Food Preparation and Serving Related Occupations (SOC=35); Building and Grounds Cleaning and Maintenance Occupations (SOC=37); Personal Care and Service Occupations (SOC=39); Sales and Related Occupations (SOC=41); Office and Administrative Support Occupations (SOC=43); Farming, Fishing, and Forestry Occupations (SOC=45); Construction and Extraction Occupations (SOC=47); Installation, Maintenance, and Repair Occupations (SOC=49); Production Occupations (SOC=51); Transportation and Material Moving Occupations (SOC=53).

A.5 Measurement of O*NET Tasks and Activities

ONET task statements are detailed descriptions of the activities that workers in various occupations perform. ONET comprises 19,530 task statements, which are suggested by actual

workers in each occupation who complete surveys indicating which tasks they perform. Workers also rate each task's importance on a 1-5 scale; these ratings are averaged across all surveyed workers in that occupation to create an importance score on a scale from 0 to 100.

To access the importance scores, we download the tasks and importance scores for each 2018 SOC code. We use 2018 codes to stay consistent with the NIOSH Industry and Occupation Computerized Coding System that we use when eliciting respondents' occupations. Some occupations are slightly aggregated. For example, the 2019 SOC lists occupation 11-1011.00 (Chief Executives) and 11-1011.03 (Chief Sustainability Officers) separately; according to the ONET-SOC 2019 to 2018 crosswalk, both map to 2018 SOC code 11-1011 (Chief Executives) and are thus assigned the tasks and importance scores from that code. This procedure yields 15,355 ONET tasks, with importance scores unavailable for 727.

We prefer using O*NET work activities rather than task statements for several reasons. Task statements are long, wordy, and often technical, while detailed work activities are more intuitive. For example:

> ***Task Statement:*** *Typeset and measure dimensions, spacing, and positioning of page elements, such as copy and illustrations, to verify conformance to specifications, using printer's ruler or layout software.*
>
> ***Work Activity:*** *Proofread documents, records, or other files to ensure accuracy.*
>
> ***Task Statement:*** *Analyze job orders, drawings, blueprints, specifications, printed circuit board pattern films, and design data to calculate dimensions, tool selection, machine speeds, and feed rates.*
>
> ***Work Activity:*** *Study blueprints or other instructions to determine equipment setup requirements.*

Using crosswalks from O*NET, we link 13,803 of the 15,355 task statements with importance scores to detailed work activities; 10,659 have a unique match, and for the remaining 3,144, most link to two activities. When multiple task statements link to a single occupation-specific detailed work activity, we assign the mean importance score.

The 2,040 detailed work activities can be aggregated into 333 intermediate work activities, 37 broader work activities, 9 broad work activities, and 4 very broad work activities. For 25 occupations, we do not have importance scores for any designated tasks or work activities. Based on the respondent's occupation, we ask:

> *Below is a list of work activities commonly associated with your occupation. Please select all the activities that you personally perform as part of your job.*

We present respondents with the top 10 activities by importance, except for 57 occupations with fewer than 10 activities (affecting 48 respondents in August 2025 and 50 in November 2025). GenAI users are then asked:

> *Earlier in the survey you indicated that you personally perform the following activities in your job. For which of these activities do you regularly use Generative AI? Please select all that apply.*