Generative AI in the Labor Market: Evidence and Some Theory*

Alexander Bick
Federal Reserve Bank of St. Louis & CEPR

Adam Blandin
Vanderbilt University

David J. Deming
Harvard University & NBER

Tyler Schumacher
Vanderbilt University

Preliminary — Please do not circulate

March 20, 2026

Abstract

This paper develops a novel approach to measuring the labor market impact of a new technology. Rather than relying on assumptions about which tasks a new technology can perform, we infer its current capabilities from empirical adoption patterns through the lens of a structural model. In the model, a new technology can perform some portion of work, which differs by occupation. Heterogeneous workers adopt the new technology once its productivity in automatable work exceeds their own. Empirical adoption rates reveal the technology's productivity by occupation, while usage intensity among adopters reveals the share of automatable work. We apply this framework to generative AI (genAI) using nationally representative adoption data from the US. Our calibrated model implies genAI has increased aggregate labor productivity by 2.0% and wage inequality by 0.9%, with gains concentrated in management, computer, and math occupations. A counterfactual simulation using ex-ante predictions of genAI capabilities from the literature implies a 27% aggregate productivity gain, suggesting that genAI's impact thus far is modest relative to these predictions.

Keywords: Generative AI, Technology Adoption, Labor Productivity, Wage Inequality

* *Contact*: alexander.bick@stls.frb.org; adam.blandin@vanderbilt.edu; david_deming@harvard.edu; tyler.schumacher@vanderbilt.edu. 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.

1 Introduction

An immense literature seeks to understand the labor market impact of technological change. Within this literature, the prevailing paradigm begins from the premise that new technologies perform some types of work but not others. A central question, then, is which work the technology performs. An influential approach to addressing this question is to partition the space of work activities into a set of tasks and then impose assumptions about which jobs or tasks the technology can perform based on the underlying nature of the technology. This approach is well suited to mature technologies whose capabilities are well understood, such as studies of information and communication technologies conducted more than a decade after their widespread diffusion (Autor, Levy, and Murnane, 2003; Krusell et al., 2000). More recently, researchers have adapted this approach to nascent technologies like generative AI (genAI) by developing ex-ante predictions about which tasks the technology can perform (Eloundou et al., 2024; Felten, Raj, and Seamans, 2021) and embedding these predictions in quantitative models to assess potential labor market effects, see e.g., Acemoglu (2024), Althoff and Reichardt (2026), and Freund and Mann (2025). However, these predictions face inherent limitations when the technology's capabilities are unclear and rapidly evolving.

In this paper, we offer an alternative approach. Rather than relying on ex-ante predictions, we infer a technology's capabilities from revealed preferences through the lens of a structural model of technology adoption. We show that we can parametrize the model using data on technology adoption without requiring prior assumptions about technological capabilities. The result is a framework that provides contemporaneous estimates of the technology's labor market impact, which can be updated as the data evolves.

In our model, workers differ in their occupation and allocate time between two types of work activities: "generic" labor, which the new technology can perform, and "nuanced" labor, which it cannot. Workers are heterogeneous in their productivity at each activity. Occupations differ in both the relative importance of generic work and the technology's productivity at performing it. A worker adopts the technology when its productivity for generic work exceeds their own.

This structure maps directly to observable adoption patterns. For a given occupation, the share of workers who adopt the technology identifies the technology's productivity: higher adoption rates reveal occupations where the technology is more productive. Similarly, the share of work time spent using the technology identifies an occupation's generic work share: more time spent using the technology indicates a larger scope for automation. Together with data on wages and employment shares, these empirical adoption patterns allow us to fully parametrize the model and quantify the technology's labor market impact.

We apply this method to study the labor market impact of genAI. This requires data on both

the extensive and intensive margins of genAI adoption by occupation: what share of workers use the technology, and what share of their work time do they spend using it? 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 (Bick, Blandin, and Deming, 2026). For this paper we pool data from multiple surveys in 2025, covering roughly 15,000 workers.

We quantify genAI's labor market impact by comparing the calibrated model to a counterfactual with genAI productivity set to zero. Our model delivers three key results. First, as of late 2025 genAI has increased aggregate labor productivity by 2.0%. While genAI increases productivity in all occupations, roughly one third of this aggregate gain is driven by management, computer, and math occupations.

Second, genAI has increased the variance of log wages by 0.9%. This is driven by similar increases in both between-occupation inequality and within-occupation inequality. Wage inequality increases between occupations because occupations with higher wages ex-ante tend to have higher genAI productivities. Wage inequality increases within occupations because higher wage workers are more likely to adopt the technology.

Third, we compare our estimates to a counterfactual simulation by mapping genAI exposure predictions from Eloundou et al. (2024) into our model. That is, rather than using adoption data to infer genAI's contemporaneous capabilities, we feed in one set of ex-ante predictions about genAI's potential capabilities in the future. Under this counterfactual, we find that genAI increases aggregate productivity by 27%, versus 2% in our primary calibration. We conclude that genAI's current impact is modest relative to recent predictions of its future potential.

2 Data Sources and 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. This section summarizes the description of the RPS in Bick, Blandin, and Deming (2026), which presented initial estimates of genAI adoption from 2024.

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 et al. (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. This paper analyzes results from four surveys conducted in 2025.

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. Table 1 compares the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure pooling February, May, August, and November 2025 CPS (column 1) and RPS (column 2). The most notable discrepancies are that individuals with no more than a high school degree or with household income in the last 12 months exceeding $150,000 degree are underrepresented in the RPS relative to the CPS.

The bottom panel of Table 1 compares employment status in the CPS and RPS and columns 3 and 4 the demographic composition of the employed. None of these statistics were targeted in the sampling procedure. The RPS features a higher share of employed and unemployed individuals, while the differences in the demographic composition of the employed mirrors the differences for the whole population aged 18-64. This is not surprising given that the vast majority of the CPS and RPS samples are employed.

Table 1: Sample Composition in the 2025 CPS and RPS

	Everyone		*Employed*	
	CPS	RPS	CPS	RPS
	(1)	(2)	(3)	(4)
Gender: Women	50.4	50.2	47.2	46.3
Age				
18-24	15.0	13.5	11.9	14.1
25-34	22.4	20.0	24.1	22.0
35-44	22.4	25.8	24.7	27.7
45-54	19.9	22.0	21.5	21.7
55-64	20.3	18.7	17.8	14.5
Race/Ethnicity				
Non-hispanic White	55.6	59.8	56.9	59.9
Non-hispanic Black	12.8	14.3	11.9	14.2
Hispanic	21.1	18.2	20.8	18.2
Other	10.5	7.7	10.3	7.7
Education				
Highschool or less	36.5	31.1	32.4	27.1
Some college/Associate's degree	25.9	27.5	25.3	27.0
Bachelor's or Graduate degree	37.6	41.4	42.3	46.0
Marital Status: Married	50.0	49.5	52.8	52.1
Number of children				
0	60.4	57.0	59.4	53.8
1	17.7	18.9	17.9	20.2
2	13.9	16.4	14.7	18.0
3+	8.0	7.8	7.9	8.0
Household Income in Last 12 Months				
\$0-\$50,000	24.6	29.8	18.7	23.2
\$50,000-\$100,000	29.5	29.2	30.2	30.9
\$100,000-\$150,000	18.7	22.8	20.5	25.4
\$150,000+	27.2	18.3	30.5	20.5
Region				
Northeast	17.0	19.2	17.2	19.9
Midwest	20.2	18.9	20.8	18.6
South	38.8	39.2	38.2	38.6
West	24.0	22.6	23.9	22.8
Employment Status				
Employed, at work last week	71.9	75.1		
Employed, absent from work last week	2.3	3.0		
Unemployed	3.3	6.8		
Not in the labor force	22.5	15.1		
Observations	218186	19774	161694	15455

Notes: Column 1 reports the sample composition in the pooled 2025 Current Population Survey (CPS) for the variables targeted by Qualtrics in the sampling procedure. Employment status was the only variable not targeted. Column 2 reports the sample composition in the pooled August and November 2024 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).

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. We weight by disaggregated categories for education, employment, and marital status and interact all categories with gender. We match these key labor market statistics both in the aggregate and conditional on demographic characteristics. Finally, occupational composition is also included in our weighting scheme; however, due to the relatively small number of observations for some occupations, we do not interact occupation categories with demographics. Appendices A.1 and A.2 detail the sample restrictions and the construction of the sample weights.

Figure 1 shows the distribution of usual weekly earnings and occupational shares in the RPS and CPS samples (Appendix A.5 explains how we elicit occupation). Occupation shares were included as targets in the weighting procedure, but earnings were not, and neither earnings nor occupation were targeted during sampling. Left-hand-side figures show unweighted distributions and right-hand-side figures show weighted distributions. The distributions are already similar without weights, and applying weights further improves the alignment (by construction for occupation shares).

Appendix A.3 presents analogous plots for industry and college major, neither of which were targeted during sample collection. The correlation between unweighted industry shares in CPS and college major shares in the ACS and the RPS are 0.73 and 0.79, respectively. Applying sample weights increases these correlations slightly, to 0.81 and 0.85, respectfully.

2.2 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:

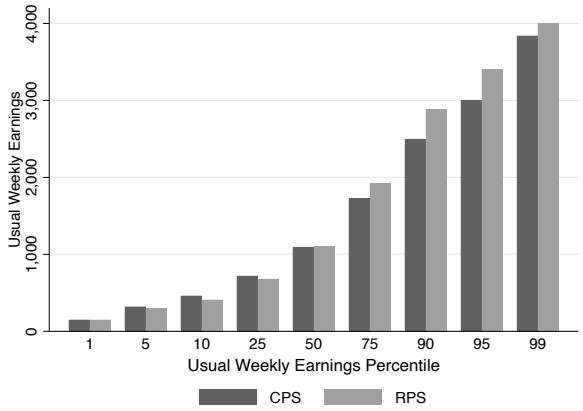
> *Do you [directly] use a computer for your job? (No/Yes)*

A second question asks about computer use at home:

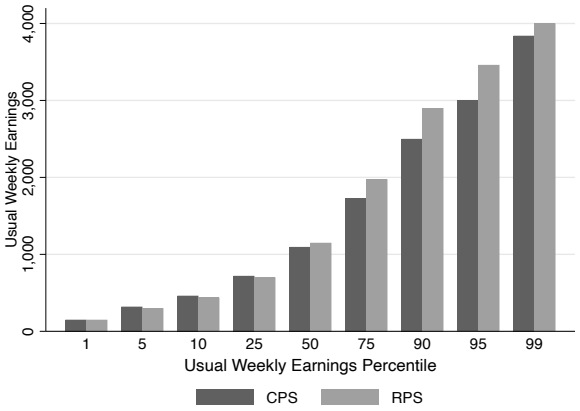
> *Do you [directly] use a computer at home? (No/Yes)*

Figure 1: Validation Checks – Usual Weekly Earnings and Occupation Shares

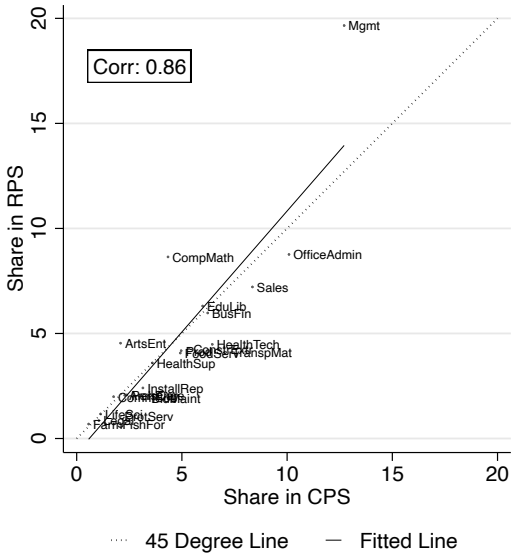
(a) Weekly Earnings Percentiles: Unweighted



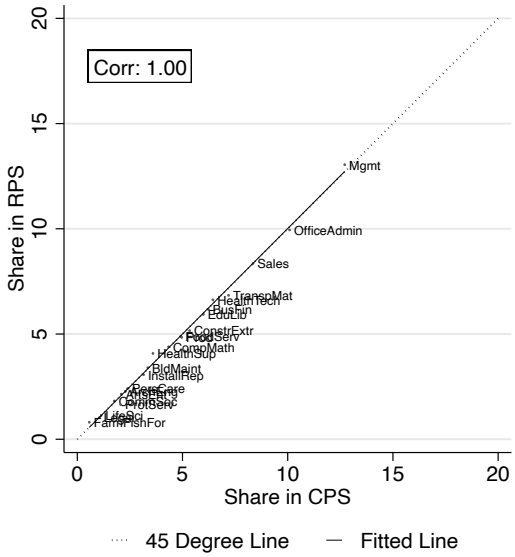
(b) Weekly Earnings Percentiles: Weighted



(c) Occupation Shares: Unweighted



(d) Occupation 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 data. Data samples for the weekly earnings figures are employees ages 18-64 in the pooled 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. Sample sizes for the RPS and CPS are 5042 and 23416, respectively. Data samples for the occupation comparison are employed respondents ages 18-64 in the pooled 2025 RPS and CPS with sample sizes of 161705 and 14207, respectively.

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

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. After defining genAI, the module asks respondents we ask employed respondents:

> *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."

The survey asks several follow-up questions of genAI work users regarding which products they used, how intensively they used it, which tasks they used it for, and how much time they saved. It also asks similar questions about non-work use.

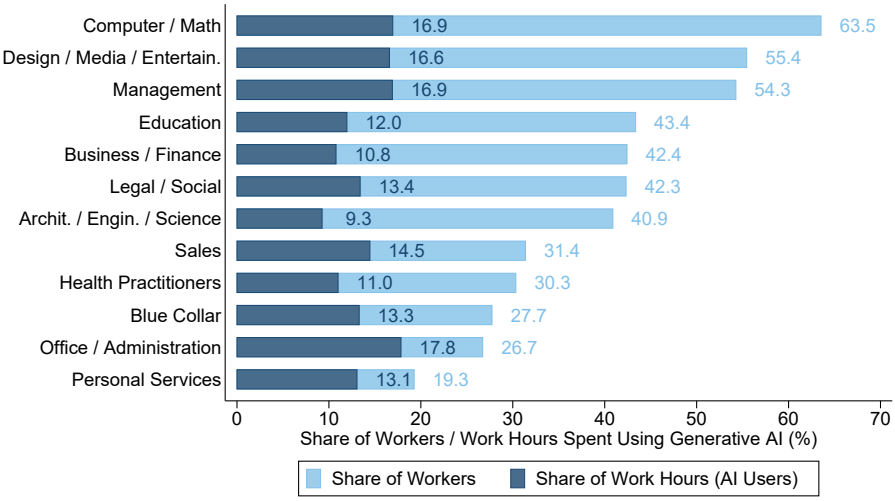
3 Evidence on Worker Adoption of Generative AI

This section uses RPS data to establish three key empirical patterns that guide our modeling choices. First, the share of workers who adopt genAI varies substantially across occupations. Second, most workers who adopt genAI only report using the technology for a modest portion of their work week. Third, workers who use the technology for a larger share of their workweek report saving more time from the technology.

3.1 The Extensive and Intensive Margin of Worker Adoption

Share of Workers who Adopt GenAI Figure 2 displays genAI adoption rates for 12 major occupation groups. The light blue bars indicate the share of workers in each occupation group who report using genAI for their job. On the high end, 63% of Computer / Math workers use genAI, as well as over 50% of Design / Media / Entertainment and Management workers. On the low end, 19% of Personal Service workers use genAI, and fewer than 30% of Blue Collar and Office / Administration workers do so.

Figure 2: Generative AI Adoption Across Occupation Groups



Notes: Dataset is the pooled Feb. 2025, May 2025, and Aug. 2025 RPS survey waves. Occupations are aggregated into broad groups based on the CPS major occupation group. The light bars report the share of workers in each occupation group who report using Generative AI. The dark bars report the mean share of work hours spent using Generative AI among users within each occupation group.

Adoption shares vary to an even greater degree at more granular measures of occupation. Table 2 summarizes regressions of genAI adoption on detailed occupation dummies, excluding occupations with fewer than 10 observations. With only occupation dummies, we find an adjusted R^2 of 0.144. Including controls for worker demographics and industry increase the adjusted R^2 to 0.169. For perspective, an analogous regression of computer adoption on detailed occupation dummies in the 1984 CIU yields an adjusted R^2 of 0.267 without demographic controls and 0.276 with demographics.

Share of Work Hours Spent Using GenAI The RPS also allows us measure the share of work time that workers spend using genAI. First, the RPS ask workers who use genAI for their job:

> *Did you use Generative AI for your job LAST WEEK?*
>
> – *No, I did not use Generative AI for my job last week* **(12.9%)**
> – *Yes, 1 workday last week* **(20.8%)**
> – *Yes, more than 1 workday last week* **(34.2%)**
> – *Yes, every workday last week* **(31.8%)**

The bold text following each answer option indicates the overall response shares. Roughly one third of users used the technology every day last week, one third used it on multiple days but not every day last week, and one third used it on one or fewer days last week.

Table 2: Predictors of GenAI Adoption and Computer Adoption

	GenAI adoption	Computer adoption
Controls	Adjusted R^2	Adjusted R^2
Detailed Occupation FE	0.144	0.267
Detailed Occupation FE + Demographic Controls	0.169	0.276

Notes: Data are from the pooled Feb. 2025, May 2025, and Aug. 2025 RPS surey waves for genAI adoption, and the 1984 CPS Computer and Internet Use Supplement (CIU) for computer adoption. Each row reports the adjusted R^2 from a separate regression with the dependent variable indicated by the column header. The dependent variable is an indicator for whether the employed worker reports using the technology. "Demographic controls" include sex, race, education, and a quadratic for age. "Detailed occupation FE" are at the CPS detailed (3-digit) occupation level in both cases.

Second, the RPS asks people who reported using genAI the previous week:

> *Please think back to the days LAST WEEK on which you used Generative AI at work. On average, how much time did you spend actively using Generative AI at work?*
>
> - *Less than 15 minutes per day* **(15.0%)**
> - *Between 15 and 59 minutes per day* **(43.0%)**
> - *Between 1 hour and 3 hours per day* **(31.6%)**
> - *More than 4 hours per day* **(10.4%)**

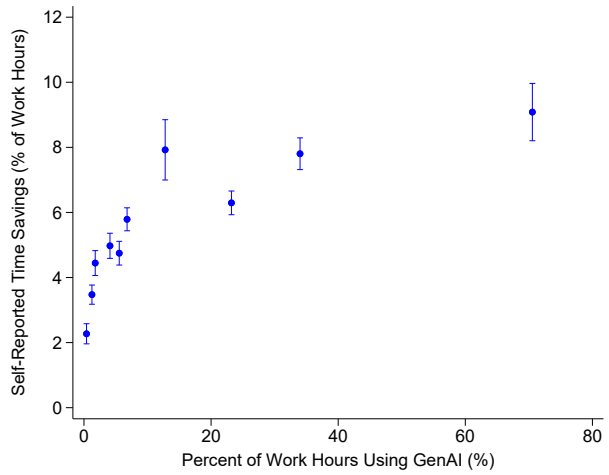
The bold text following each answer option indicates the overall response shares. Roughly 60% used the technology for one hour per day or less last week, 30% used it for between one and three hours per day, and only 10% used it for more than four hours per day.

We estimate bounds on genAI time share by combining usage intensity with days and hours worked in the previous week. For example, consider a worker who reports (i) working 40 hours last week over 5 days, (ii) using genAI on multiple but not all workdays last week, and (iii) using genAI for between 15 and 59 minutes per day on days that they used it. The lower bound implies the respondent used genAI on two days for 15 minutes each day ($15/59 \cdot 2/40 = 1.3\%$ of their weekly hours); the upper bound implies using genAI on four workdays for 1 hour each day ($1 \cdot 4/40 = 10\%$ of their weekly hours). Going forward, we use the mean of the lower and upper bound on genAI time share for each worker.

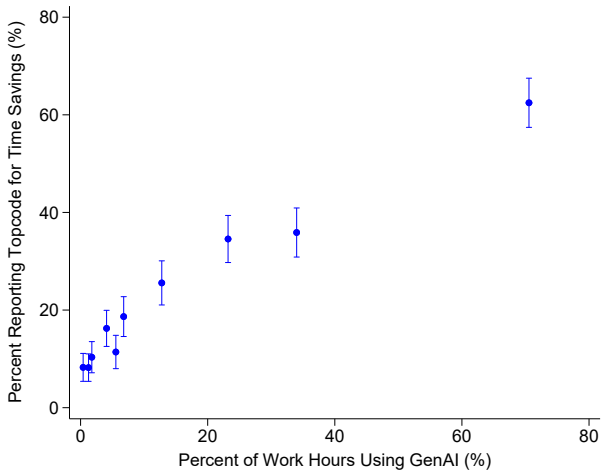
The dark blue bars in Figure 2 indicate the mean share of work time that genAI users report using genAI for their job by occupation group. The mean genAI time share among users ranges from 9.3% to 17.8%. Interestingly, genAI users in occupations with higher adoption rates do not systematically use the technology for a larger share of their workweek. We conclude that in all occupation groups the mean time share of work time spent using genAI among users is modest.

Figure 3: GenAI Time and Time Savings

(a) Mean Time Savings by Usage



(b) Share Reporting Topcoded Time Savings



Notes: Dataset is the pooled Feb. 2025, May 2025, and Aug. 2025 RPS survey waves. The x-axis (time using GenAI as a share of work hours) is split into deciles on both figures. In Panel 3a self-reported time savings (share of work hours) are averaged within each decile. In Panel 3b the shares of people reporting the topcoded time savings value of 4 hours are reported.

3.2 Time Savings Due to Generative AI

The RPS asked workers who used genAI in the previous week to estimate how much time genAI saved them:

> *You indicated that LAST WEEK you worked _X_ hours and that you used Generative AI for your job. Now, imagine that LAST WEEK you did not have access to Generative AI. How many additional hours of work would you have needed to complete the same amount of work?*
>
> - *Less than 1 hour* **(15.7%)**
> - *1 hour* **(18.3%)**
> - *2 hours* **(25.0%)**
> - *3 hours* **(18.0%)**
> - *4 hours* **(11.7%)**
> - *More than 4 hours* **(11.3%)**

In the survey, X is replaced with the respondent's reported hours worked last week. Overall, 34% of users reported saving one hour or less, 55% reported saving between two and four hours, and 11% reported saving more than four hours.

Table 3: GenAI Time and Time Savings

Controls	Adjusted R^2
% of Work Hours Using genAI	0.149
% of Work Hours Using genAI + Demographic Controls	0.168
% of Work Hours Using genAI + Demographic Controls + Detailed Occupation FE	0.250
Demographic Controls + Detailed Occupation FE	0.121

Notes: Data is from the Aug. 2025 wave of the RPS. Dependent variable in all regressions is self-reported time savings from using genAI as a percentage of work hours. Control "% of Work Hours Using genAI" is the worker's reported time spent using genAI at work last week as a percentage of the worker's work hours last week. Demographic controls include a quadratic for age, sex, race, and education level. Detailed occupation fixed effects use the CPS detailed (3-digit) occupation.

Panel (a) of Figure 3 shows that workers who spend more time using genAI report greater time savings. For instance, workers in the first decile of genAI usage spend 0.4% of work hours using genAI and report 2.3% time savings, while sixth-decile workers spend 6.8% of work hours using genAI and report 5.8% time savings. Time savings plateau at higher deciles, however. Panel (b) shows why: the share of respondents selecting the top-coded (maximum) time savings response increases sharply with usage intensity, suggesting we underestimate true time savings among intensive users.

Table 3 summarizes results from a regression of time savings on time spent using genAI among workers who use the technology. The dependent variable is the percent of work hours saved due to genAI. The independent variable of interest is the percent of work hours spent using genAI. The first row shows that genAI time alone yields an adjusted R^2 of 0.149. The second and third rows show that worker demographics and fixed effects for occupation contribute less explanatory power for time saved, increasing the adjusted R^2 to 0.250. The fourth row shows that excluding genAI time but keeping all other regressors reduces the adjusted R^2 to 0.121. We conclude that, among genAI users, a central predictor of genAI time savings is time spent using the technology.

External Validation of Time-Saving Estimates Data based on self-reported time savings is subject to measurement error and is not derived from exogenous variation in access to AI. One way to assess the plausibility of these estimates is to compare them to existing experimental micro estimates of genAI's productivity impact. Bick et al. (2026) review this literature and find productivity gains ranging from 5% to 96%, with a mean of 31% across studies.[1]

[1]Specifically, the reported time savings / productivity gains are: Kanazawa et al. (2025), 5%; Dell'Acqua et al. (2023), 12%; Brynjolfsson, Li, and Raymond (2025), 15%; Noy and Zhang (2023), 18%; Choi, Monahan, and Schwarcz (2024), 22%; Cui et al. (2024), 26%; Peng et al. (2023), 55%; Goldsmith-Pinkham, Tan, and Zentefis (2026), 96%. We do not incorporate quality improvements into these calculations, since it is not obvious how

Among genAI users in our sample, we find a mean time savings among AI users of 4.9%. (For example, an AI user who works 40 hours per week would be expected to save roughly 2 hours per week; i.e., they would complete 42 hours worth of "pre-genAI" work in 40 hours.) Alternatively, we can ask how much time workers save per hour that they spend using AI. In our sample, we find a mean time savings of 0.38 hours per hour spent using AI. (For example, a worker who used AI for 5 hours in a week would be expected to save 1.9 hours that week.) We conclude that our mean time savings estimates of 4.9% per user, or 0.38 hours per hour spent using AI, are broadly in line with existing experimental estimates.

4 A Theory of Technology Adoption by Workers

This section lays out a theoretical model of technology adoption by workers. For a given set of parameters, the model delivers quantitative predictions of how a new technology will impact productivity and wages in the aggregate, across occupations, and within occupations.

Like most of the literature on technology adoption in the labor market, we assume that a new technology can assist some, but not all, types of work (Autor, Levy, and Murnane, 2003; Krusell et al., 2000). Within this paradigm, a central question is to understand *which work* the new technology performs. An ideal dataset to address this question would be a rich, minute by minute time log documenting workers' activities and use of the technology. In the absence of such a dataset, one common method is to (i) categorize all work activities into a set of tasks and all jobs into a set of occupations, (ii) estimate the relative importance of each task for each occupation, (iii) make ex-ante predictions about which tasks the technology will perform (Eloundou et al., 2024; Felten, Raj, and Seamans, 2021), and (iv) feed these task-importance and task-capability estimates into a structural model (Acemoglu, 2024; Althoff and Reichardt, 2026; Freund and Mann, 2025).

This approach has two potential limitations. First, even a very detailed task may involve some activities that a new technology can assist and others that it cannot assist. Second, ex-ante predictions of technology capabilities may not align with realized productivity gains or adoption patterns. This concern is less salient for mature technologies whose capabilities are fairly well understood: for example, studies of ICT written more than a decade after personal computers had been introduced. This concern becomes more acute for nascent technologies like genAI, whose capabilities are unclear and rapidly evolving.

This paper offers an alternative approach: rather than relying on ex-ante predictions, we infer technology current capabilities from revealed preferences. In our model, the share of

to weight these gains relative to time savings. In most studies, AI improved quality on average, which would tend to increase the quality-adjusted productivity gains.

workers who adopt a new technology and the share of work time spent using it are informative about which work the technology performs and the productivity gain it provides. This revealed-preference approach allows us to parametrize the model using RPS adoption data without requiring prior assumptions about technological capabilities.

4.1 Environment Before the New Technology

We begin with the environment before the new technology's arrival. The economy is static and consists of a unit mass of workers distributed across O occupations. Workers are exogenously assigned to occupations and do not switch. Each occupation o contains a mass $I_o > 0$ of workers indexed by $i \in [0, I_o]$, where $\sum_o I_o = 1$.

Production Technologies A single consumption good is produced by a representative firm that aggregates output across occupations using a CES production function:

$$Y = \left(\sum_o \beta_o Y_o^{\frac{\omega-1}{\omega}} \right)^{\frac{\omega}{\omega-1}}, \tag{1}$$

Here, β_o is the weight on occupation o and $\omega > 0$ is the elasticity of substitution across occupations. Aggregate output in occupation o is given by

$$Y_o = \int_{i=0}^{I_o} y_{i,o}\, di, \tag{2}$$

where $y_{i,o}$ denotes the output of worker i in occupation o.

Workers Workers value the consumption good, which they purchase using wages paid by their occupation's firm. Worker wages are $w_o y_{i,o}$, the product of an occupation-specific piece rate w_o and their individual output $y_{i,o}$.

Each worker is endowed with a unit of time, which they can split between two activities: generic labor, g, and nuanced labor, n. The new technology, when it arrives, will only be able to perform generic labor. Workers differ in their generic and nuanced productivities. Worker i in occupation o draws a pair of productivity shocks (γ_i, ν_i) from a bivariate log-Normal distribution. Given a pair of productivities and a time allocation choice (g_i, n_i), worker i produces $y_{i,o}$ units of output:

$$y_{i,o}(g, n) = \left(\theta_o^{\frac{1}{\phi_o}} (\gamma_i g_i)^{\frac{\phi_o - 1}{\phi_o}} + (1 - \theta_o)^{\frac{1}{\phi_o}} (\nu_i n_i)^{\frac{\phi_o - 1}{\phi_o}} \right)^{\frac{\phi_o}{\phi_o - 1}} \tag{3}$$

The worker's output function is CES in generic and nuanced labor with elasticity of substitution

$\phi_o > 0$. The weights on generic versus nuanced labor are given by θ_o and $1 - \theta_o$, respectively. These weights can differ by occupation, reflecting the potential of a given technology to perform a greater share of work in some occupations than others.

Optimization Problems The firm rents labor services from each occupation to maximize profits, taking the occupational piece rates as given:

$$\max_{\{Y_o\}_o} \quad \left(\sum_o \beta_o Y_o^{\frac{\omega-1}{\omega}} \right)^{\frac{\omega}{\omega-1}} - \sum_o w_o Y_o \tag{4}$$

The price of the consumption good is normalized to one.

Workers split their time endowment between generic and nuanced work to maximize their earnings, taking their occupational piece rate as given:

$$\max_{g,n \geq 0} \quad w_o y_{o,i}(g, n) \tag{5}$$

$$s.t. \quad g + n = 1 \tag{6}$$

Alternatively, we could assume that the firm tells each worker how to allocate their time; however, because the firm maximizes profits, it would choose the same time allocation as workers.

4.2 Environment with the New Technology

A new technology arrives which workers can use to perform generic labor. Denote by p_i the share of time that worker i spends using the technology and by π_o the productivity of the technology in occupation o. Since our application is about the adoption of genAI, we will refer to p as "prompting". The key idea is that workers must spend time engaging with genAI, similar to how workers who used to perform calculations by pencil and paper switched to actively using a calculator or PC to perform those calculations as these technologies became available. The worker's output is now

$$y_{i,o}(g, p, n) = \left(\theta_o^{\frac{1}{\phi_o}} (\gamma_i g_i + \pi_o p_i)^{\frac{\phi_o - 1}{\phi_o}} + (1 - \theta_o)^{\frac{1}{\phi_o}} (\nu_i n_i)^{\frac{\phi_o - 1}{\phi_o}} \right)^{\frac{\phi_o}{\phi_o - 1}}. \tag{7}$$

Generic labor assisted by the technology, p_i, and unassisted by the technology, g_i, are perfect substitutes, differing only in their productivity levels. The worker's problem becomes

$$\max_{g,p,n\geq 0} \quad w_o y_{i,o}(g,p,n) \tag{8}$$

$$s.t. \quad g+p+n=1. \tag{9}$$

The firm's problem remains unchanged.

4.3 Definition of Equilibrium

A competitive equilibrium is a collection of: time choices for each worker, (g_i, p_i, n_i); occupational output, Y_o; and piece rates for each occupation, w_o; subject to

1. Given (γ_i, ν_i), π_o, and w_o, (g_i, p_i, n_i) solves the worker's problem (8)-(9)
2. $Y_o = \int_0^{I_o} y_{i,o}(g_i, p_i, n_i)di$
3. $w_o = \frac{\partial Y}{\partial Y_o}$

4.4 Characterizing Equilibrium

We say that a worker i adopts the technology if they use it, i.e. $p_i > 0$. Because g_i and p_i are perfect substitutes in production, worker i adopts if and only if $\pi_o > \gamma_i$. Conditional on not adopting the technology, $p_i = 0$, the worker splits time between generic and nuanced labor as follows:

$$\frac{g_i}{n_i} = \left(\frac{\theta_o}{1-\theta_o}\right)^{\phi_o} \left(\frac{\gamma_i}{\nu_i}\right)^{\phi_o - 1} \tag{10}$$

The ratio of generic to nuanced labor is increasing in θ_o, the CES weight on generic labor. The dependence on worker productivities is governed by the elasticity of substitution ϕ_o. When generic and nuanced labor are substitutes ($\phi_o > 1$), g_i/n_i is increasing in γ_i/ν_i: those with relatively high generic productivity devote relatively more time to generic labor. When generic and nuanced labor are complements, $\phi_o \in (0,1)$, g_i/n_i is decreasing in γ_i. In the Cobb-Douglas case ($\phi_o \to 1$), the ratio of generic to nuanced labor depends only on the occupation-level parameter θ_o and is independent of individual productivities.

Conditional on adopting the technology, workers set $g_i = 0$ and split their time between

prompting and nuanced labor as follows:

$$\frac{p_i}{n_i} = \left(\frac{\theta_o}{1-\theta_o}\right)^{\phi_o} \left(\frac{\pi_o}{\nu_i}\right)^{\phi_o - 1} \tag{11}$$

The ratio is identical to (10), except that prompting p_i replaces generic labor g_i and technology productivity π_o replaces the worker's generic productivity γ_i.

In competitive equilibrium, the piece rate in occupation o is determined by the marginal product of occupational output:

$$w_o = \beta_o \left(\frac{Y}{Y_o}\right)^{\frac{1}{\omega}} \tag{12}$$

A key implication of this equation is that, all else equal, an increase in occupational output Y_o reduces the occupation's piece rate: $\frac{\partial w_o}{\partial Y_o} < 0$.

4.5 The Impact of an Increase in Technology Productivity

What is the labor market impact of an increase in the technology's productivity in occupation o, from 0 to $\pi_o > 0$?

To begin, wages increase for all workers in all other occupations $k \neq o$. We can see this in three steps. First, from equations (10) and (11) we see that the time allocation of workers in other occupations do not change. This implies that Y_k is unchanged as well. Second, Y_o increases: individual output is unchanged for all non-adopters, and strictly increases for all adopters. Third, the piece rate of all other occupations increase: $\frac{\partial w_k}{\partial w_o} > 0$.

Second, wages decrease for all non-adopters in occupation o. This is because non-adopters produce the same output as before, but the piece rate in occupation o declines as implied by Equation (12).

Third, the sign of wage changes for adopters in occupation o varies across workers. Adopters benefit from higher individual output due to the increase in π_o, but face a lower piece rate w_o due to the increase in aggregate occupational output Y_o. The net effect on wages $w_o y_{i,o}$ depends on the magnitude of these two opposing forces. Workers who rely more heavily on the technology—those with low γ_i relative to π_o—experience larger gains in individual output and are more likely to see wage increases. Conversely, adopters with γ_i just below π_o gain little from the technology but still suffer the full decline in the piece rate, and thus are more likely to see wage decreases.

More formally, whether a given adopter experiences a wage gain or loss depends on the

pair (γ_i, ν_i). Holding ν_i fixed, there exists a threshold $\bar{\gamma}(\nu_i)$ such that adopters with $\gamma_i < \bar{\gamma}(\nu_i)$ experience wage gains while those with $\gamma_i \in (\bar{\gamma}(\nu_i), \pi_o)$ experience wage losses. The dependence of $\bar{\gamma}$ on ν_i is governed by the elasticity of substitution ϕ_o. When generic and nuanced labor are substitutes ($\phi_o > 1$), $\bar{\gamma}(\nu_i)$ is decreasing in ν_i: higher nuanced productivity workers optimally devote more time to nuanced labor and less to prompting, benefit less from the increase in π_o, and therefore need to have had lower pre-existing generic productivity for the output gain to outweigh the piece rate decline. When generic and nuanced labor are complements ($\phi_o < 1$), the opposite holds: higher ν_i workers devote more time to prompting and so benefit more from the increase in π_o, implying that $\bar{\gamma}(\nu_i)$ is increasing. In the Cobb-Douglas case ($\phi_o \to 1$), the threshold is independent of ν_i.

Taken together, these results imply that the technology generates both winners and losers within an occupation. Non-adopters unambiguously lose. Among adopters, those who were least productive at generic labor prior to the technology's arrival are most likely to benefit.

What are the implications for within-occupation wage inequality? The technology raises wages for adopters with low γ_i and lowers wages for non-adopters and for adopters near the adoption margin. Whether this compresses or widens the wage distribution within the occupation depends on the joint distribution of (γ_i, ν_i). When the correlation between γ_i and ν_i is positive or close to zero, workers with low generic productivity, who benefit most from the technology, also tend to have low nuanced productivity, and are therefore low-wage workers prior to the technology's arrival. In this case, the technology compresses the within-occupation wage distribution by raising wages at the bottom. However, when the correlation is sufficiently negative, workers with low γ_i tend to have high ν_i, making them already high-wage workers prior to the arrival of the technology. The technology then widens within-occupation inequality by disproportionately benefiting high wage workers.

The technology also affects between-occupation inequality. An increase in π_o raises output Y_o in occupation o, which lowers the piece rate w_o while raising piece rates w_k in all other occupations. Whether the total wage bill $w_o Y_o$ in occupation o rises or falls depends on the elasticity of substitution across occupations, ω. When $\omega = 1$, the decline in the piece rate exactly offsets the increase in output, leaving the total wage bill (and mean wages within occupation o) unchanged. When occupations are complements ($\omega < 1$), the piece rate effect dominates: the wage bill in occupation o falls, reducing mean wages in the occupation and raising them elsewhere. Conversely, when occupations are substitutes ($\omega > 1$), the output effect dominates and the wage bill rises. The implications for between-occupation inequality thus depend on whether the technology disproportionately improves productivity in initially high- or low-wage occupations, and on the degree of substitutability across occupations.

4.6 Future Extensions

Future extensions will include occupational choice, alternative production functions in which productivity gains from genAI may correlate with generic skill, and adoption costs (both fixed and variable).

5 Application: The Labor Market Impact of Generative AI

5.1 Calibration

We calibrate the model to RPS data from 2025, after the introduction of genAI. At the aggregate level, we set the elasticity of substitution between occupational outputs to $\omega = 1.4$, implying that occupations are gross substitutes. In calculating aggregate output, we weight each occupation's output by its employment share in the data.

For each occupation, we must calibrate 9 parameters: five parameters governing the bivariate log-normal distribution of skills $(\mu_{\gamma,o}, \sigma_{\gamma,o}, \mu_{\nu,o}, \sigma_{\nu,o}, \rho_{\gamma_o,\nu_o})$, three parameters for the worker production function (genAI productivity π_o, the generic labor share θ_o, and the elasticity of substitution ϕ_o), and the occupational weight in aggregate production (β_o).

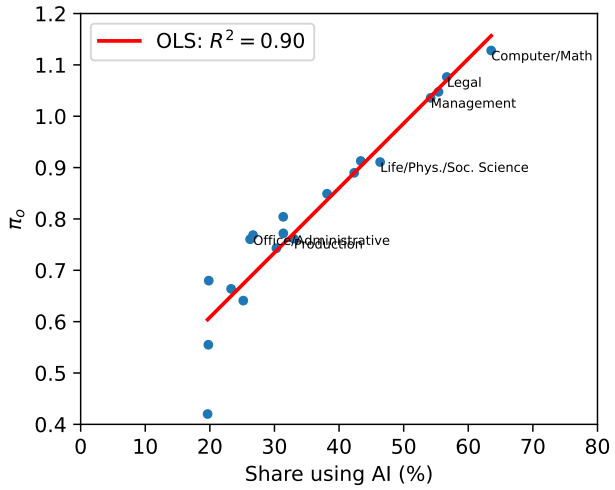
We set $\phi_o = 1$ for all occupations, implying a Cobb-Douglas production function in which each worker's time allocation depends only on θ_o and is independent of individual productivity draws. In future work, we will explore deviations from this assumption using recently added RPS questions designed to calibrate occupation-level elasticities. Without loss of generality, we normalize the mean of both skills to zero: $\mu_{\gamma,o} = \mu_{\nu,o} = 0$.

For the remaining parameters for each of the 20 broad occupations, we use simulated method of moments, selecting values so that the model matches fourteen target moments per occupation. Figure 4 visualizes the relationship between calibrated parameters and their most closely related target moments. While all parameters influence all moments, we can provide intuition about the primary linkages by examining these parameter-moment pairs. Red lines show OLS fits, and Appendix Tables C.1–C.6 provide detailed calibration moments and parameter values and show that we match all targets perfectly.

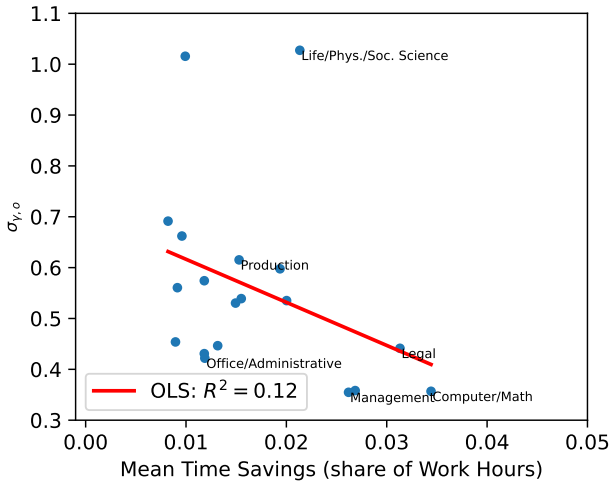
Panel (a) shows that genAI productivity π_o exhibits a tight relationship with adoption rates. For a given $\sigma_{\gamma,o}$, the technology productivity π_o determines the adoption threshold: workers adopt if and only if $\pi_o > \gamma_i$, creating a mechanical link between π_o and the share of workers adopting. The relationship is not perfect because productivity also affects time savings, which in turn informs $\sigma_{\gamma,o}$. This is shown in panel (b). Greater dispersion in generic productivity implies larger gaps between the adoption threshold and the mean productivity

Figure 4: Calibration: Parameters and Target Moments

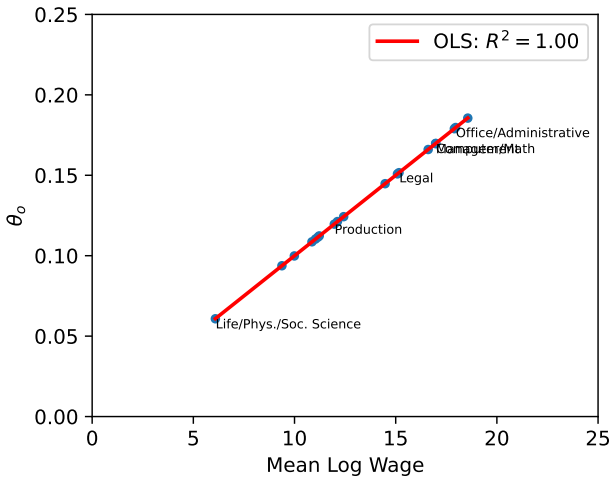
(a) GenAI Productivity, π_o



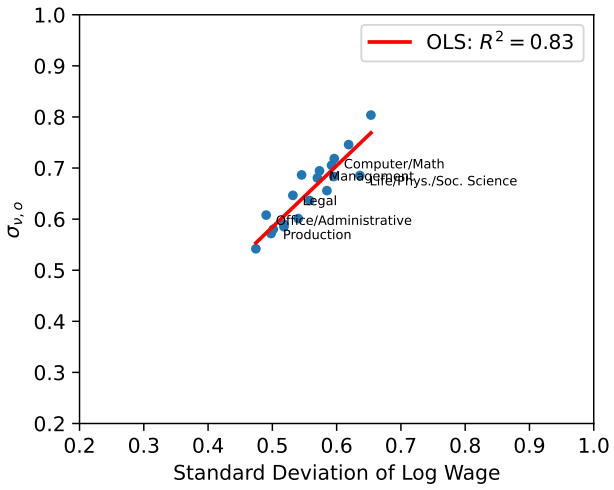
(b) Generic Skill Dispersion, $\sigma_{\gamma,o}$



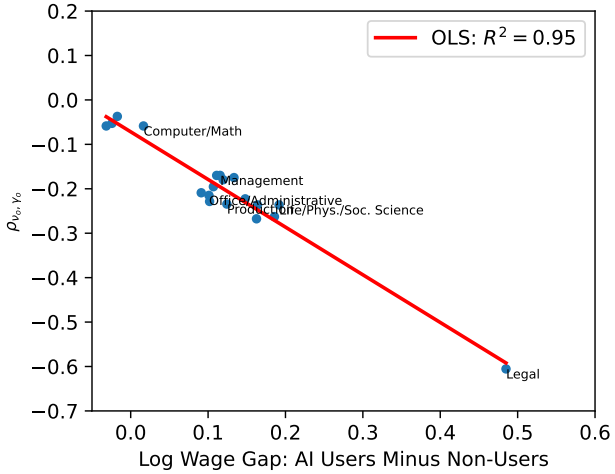
(c) Generic Labor Shares, θ_o



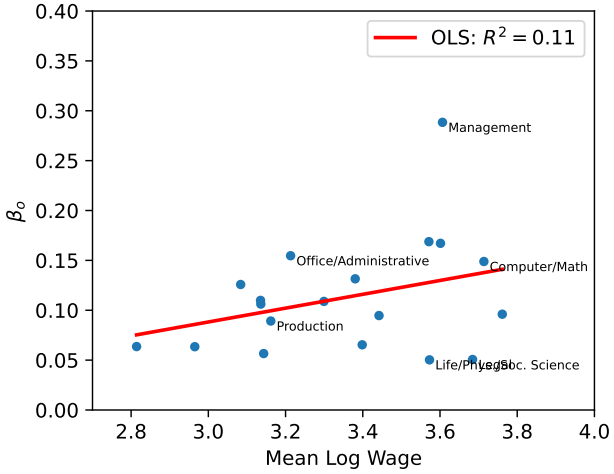
(d) Nuanced Skill Dispersion, $\sigma_{\nu,o}$



(e) Skill Correlation, ρ_{ν_o,γ_o}



(f) Occupational Weight, β_o



Notes: Model parameters are calibrated to match target moments from the pooled Feb., May, and Aug. 2025 RPS waves. Each panel plots a calibrated parameter against its most closely related data moment. Red lines show OLS fits. Appendix Tables C.1–C.6 provide detailed calibration moments and parameter values.

of adopters, generating larger time savings. The weaker relationship reflects interaction with other parameters, especially productivity π_o. Outliers like "Life/Physical/Social Science" and "Food Preparation and Service" show high time savings despite low adoption because the model reconciles this pattern through high $\sigma_{\gamma,o}$. This creates a fat tail of workers with very low baseline generic productivity who gain substantially from the technology. Raising π_o would increase adoption (since more workers would satisfy $\pi_o > \gamma_i$), so the model instead adjusts $\sigma_{\gamma,o}$ to create sufficient heterogeneity to match both adoption patterns and productivity gains simultaneously.

Panel (c) shows that the generic labor share θ_o is perfectly identified by time use among adopters. This direct mapping occurs because workers using genAI allocate all generic labor time to the technology in our Cobb-Douglas specification.

Since genAI time represents at most 20% of total work time on average, within-occupation wage inequality is primarily driven by variation in nuanced productivity. Panel (d) demonstrates this clearly: nuanced skill dispersion $\sigma_{\nu,o}$ strongly correlates within-occupation wage inequality. We therefore set $\sigma_{\nu,o}$ to match the variance of log wages in each occupation.

Panel (e) shows that skill correlation ρ_{ν_o,γ_o} tightly correlates with the wage gap between genAI users and non-users. This parameter directly determines the selection pattern of who adopts: if adopters earn higher wages than non-adopters, they must have low generic productivity but high nuanced productivity, implying negative correlation.

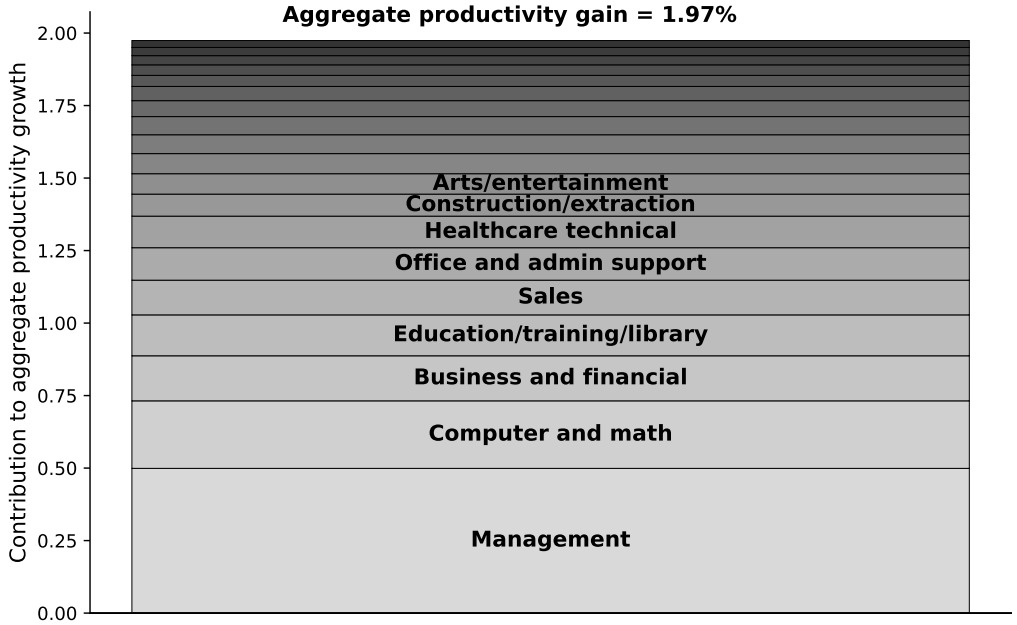
Finally, panel (f) shows occupational weights β_o versus mean wages. The occupation-specific weight β_o determines relative occupational output and thus mean wages across occupations conditional on the exogenously imposed employment shares. The weaker relationship arises because in calculating aggregate output, we weight each occupation's output by its employment share, so β_o must adjust to match mean wages while accounting for differences in occupational size.

5.2 Model-Implied Effects of Generative AI

We quantify genAI's labor market impact by comparing our calibrated economy to a counterfactual without genAI (i.e., setting $\pi_o = 0$ for all occupations, corresponding to the pre-ChatGPT era). Since labor is supplied inelastically, all output changes represent labor productivity gains.

Aggregate Productivity: GenAI increases aggregate output by 1.97% relative to the no-genAI counterfactual. Figure 5 decomposes this gain into occupation-level contributions using value-share weights from our CES aggregator. An occupation contributes to aggregate growth through two channels: high productivity gains per worker, or large employment share.

Figure 5: Estimated Aggregate Productivity Gains from GenAI



Notes: Figure decomposes the aggregate productivity gain into occupation-level contributions using value-share weights from a CES aggregator with elasticity of substitution across occupations of $\omega = 1.4$. Baseline and counterfactual occupation outputs are aggregated using value-share weights (see Sato (1976); Vartia (1976)) derived from model-implied prices and output levels, so that each segment reflects an occupation's contribution to the change in log aggregate productivity. By construction, the stacked bars sum to the total aggregate gain.

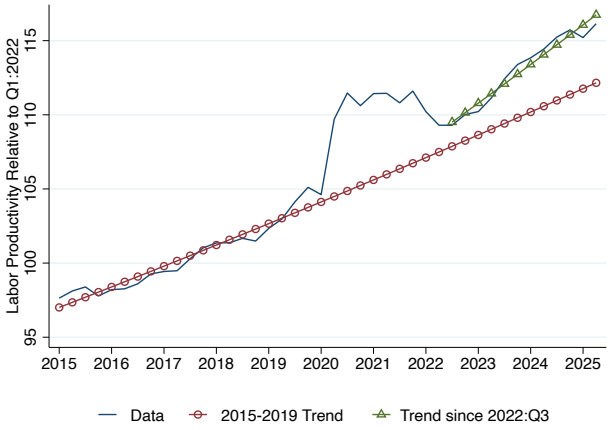
Management and Computer/Math occupations alone account for over one-third of the total gain.

Wage Inequality: GenAI increases the variance of log wages by 0.9% overall, with similar contributions from between-occupation and within-occupation inequality. Between-occupation variance increases by 1.1% because genAI productivity is higher in occupations with higher pre-genAI wages. Within-occupation variance rises by 0.9% because higher-wage workers within each occupation are more likely to adopt.

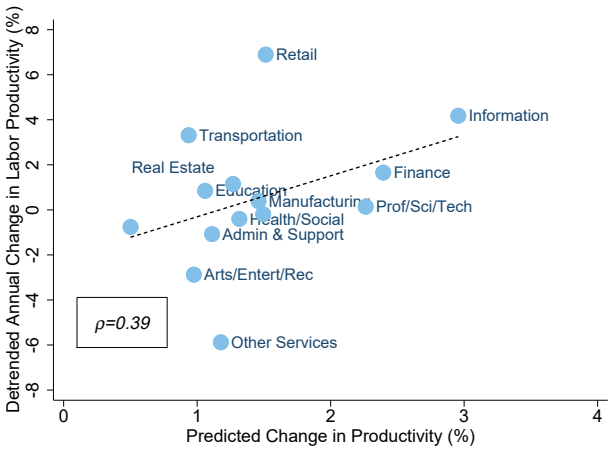
Distributional Effects: The model's distributional implications arise from worker heterogeneity and comparative advantage. Because adoptions decisions depend on the relative productivity of the technology and the worker's baseline skill, technological change generates systematic selection into adoption and heterogeneous wage effects. This mechanism is closely related to models which explore task-specific technical change and comparative advantage (*e.g.,* Althoff and Reichardt, 2026). In most occupations, all workers experience wage gains. For occupations with low genAI adoption, the piece rate w_o rises due to increased complementarity with high-adopting occupations, benefiting all workers. However, in the six highest-adoption occupations, the piece rate declines due to increased occupational output. In these occupations, non-adopters and marginal adopters (those with γ_i close to π_o) experience wage losses

Figure 6: Estimated Aggregate Productivity Gains Per Worker in U.S.

(a) Aggregate Trends



(b) By Industry



Notes: Data from BLS output per hour worked in nonfarm business sector. Panel (a): The line "Data" shows the actual value, the line "2015-2019 Trend" shows a linear fit based on 2015-2019 data, and the line "Trend Since 2022: Q3" shows a linear fit from Q3 2022 to Q1 2025. Panel (b): Each observation is a broad industry from the BLS. The vertical axis shows the change in labor productivity from BLS estimates, and the horizontal axis shows our prediction of change in productivity using data from the RPS. The number inside the box, ρ, shows the correlation between the two.

averaging 1%, while strong adopters experience gains averaging 5%. Despite these losses for some workers, mean wages rise in all occupations because winners gain more than losers lose.

5.3 Model Validation

We validate our model's predictions against three independent sources of evidence: aggregate and industry-level productivity data and the historical experience of computer adoption.

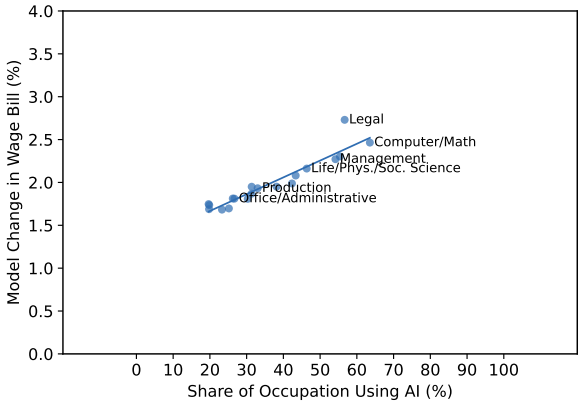
5.3.1 Aggregate and Industry Productivity Growth

Figure 6a compares our model-implied productivity gain to aggregate U.S. labor productivity growth since genAI's introduction. Using BLS data on output per hour in the nonfarm business sector, we fit trend lines for 2015-2019 (pre-pandemic baseline) and Q3 2022-Q1 2025 (post-genAI). Productivity growth in the post-genAI period exceeds the pre-pandemic trend by 2.1%. Adjusting our model-implied 1.97% gain by labor's share of GDP (approximately 60%) yields a predicted productivity increase of 1.2%, suggesting that worker adoption of genAI accounts for more than half of the observed acceleration in productivity growth.

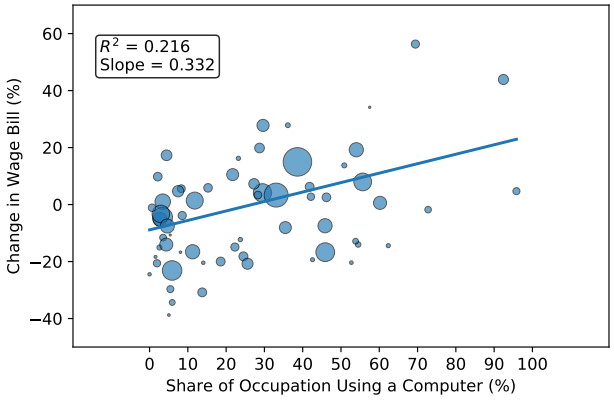
We emphasize that this comparison is suggestive rather than causal as other factors could contribute to the acceleration. We therefore examine industry-level variation for a more strin-

Figure 7: Change in Occupation-Level Wage Bill vs. Technology Adoption

(a) Model Prediction: GenAI Adoption



(b) Data (1990-1982): Computer Adoption



Notes: Occupation is at the CPS detailed occupation level. Panel (a): Share adopting genAI by occupation is calculated using the RPS. Change in the the occupation-level Panel (b): Share adopting computers by occupation is calculated using the 1984 CPS Computer Use Supplement. Change in the occupation-level wage bill is calculated as the change in total payments to workers within the occupation based on the CPS-ASEC, including changes in employment, hours, and hourly wages, between 1990-1982.

gent test. Figure 6b compares BLS industry-level productivity growth with the time-savings on the industry-level (this is equivalent to calculating for each industry the average of the occupation-specific productivity gains using the industry-specific occupation shares as weights). The correlation between predicted and observed productivity growth across industries is 0.39, indicating meaningful predictive power of our model.

5.3.2 The Historical Experience of Computer Adoption

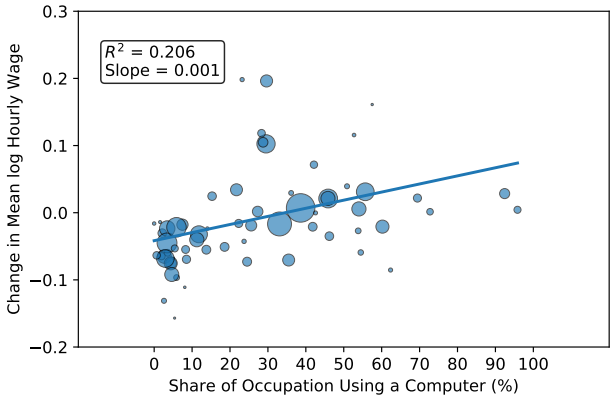
Finally, we compare our model's predictions to the historical episode of personal computer adoption in the 1980s. We cannot calibrate our model to the PC era because the 1984 CPS Computer and Internet Use supplement measured only the extensive margin of adoption. We can however examine whether the cross-sectional relationship between early technology adoption and subsequent wage growth aligns with our framework.

Figure 7 presents this comparison. Panel (a) shows the model-implied relationship between genAI adoption rates and occupation-level wage bill changes, demonstrating a strong positive relationship. Panel (b) shows the historical analog: computer adoption in 1984 strongly predicts changes in occupation-level wage bills from 1982-1990. This pattern supports our modeling approach: early adoption patterns are indeed informative about subsequent labor market impacts and in the same direction as predicted by our model.

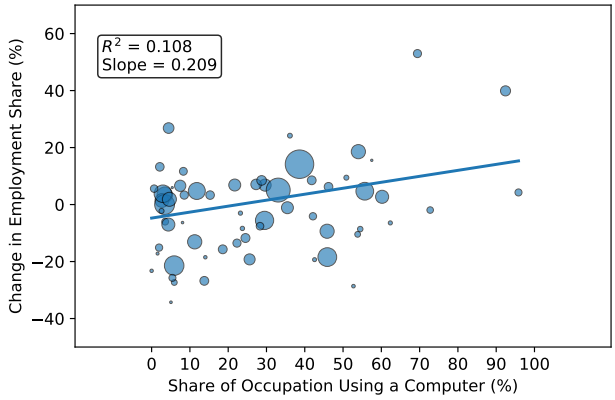
We acknowledge two limitations of this comparison. First, wage bill changes reflect both

Figure 8: Decomposition of Wage Bill Change (1990-1982) and Computerization

(a) Change in Wage Bill



(b) Change in Employment Shares



Notes: Figures use the CPS-ASEC for data on employment shares and wages, and the 1984 CPS Computer and Internet Use Supplment (CIU) for share using a computer by occupation. Panel 8a shows the change in mean log hourly wage by occupation from 1982-1990. Panel 8b shows the change in share employed by occupation from 1982-1990.

employment reallocation across occupations and within-occupation wage changes. Our current model abstracts from occupational choice but future extensions will incorporate this margin. Second, the appropriate time horizon for measuring effects is unclear. Adjustment takes time, and our eight-year window (1982-1990) was chosen because occupational codes were not changed throughout this period, enabling consistent measurement (shorter time windows within that period deliver qualitatively similar results).[2]

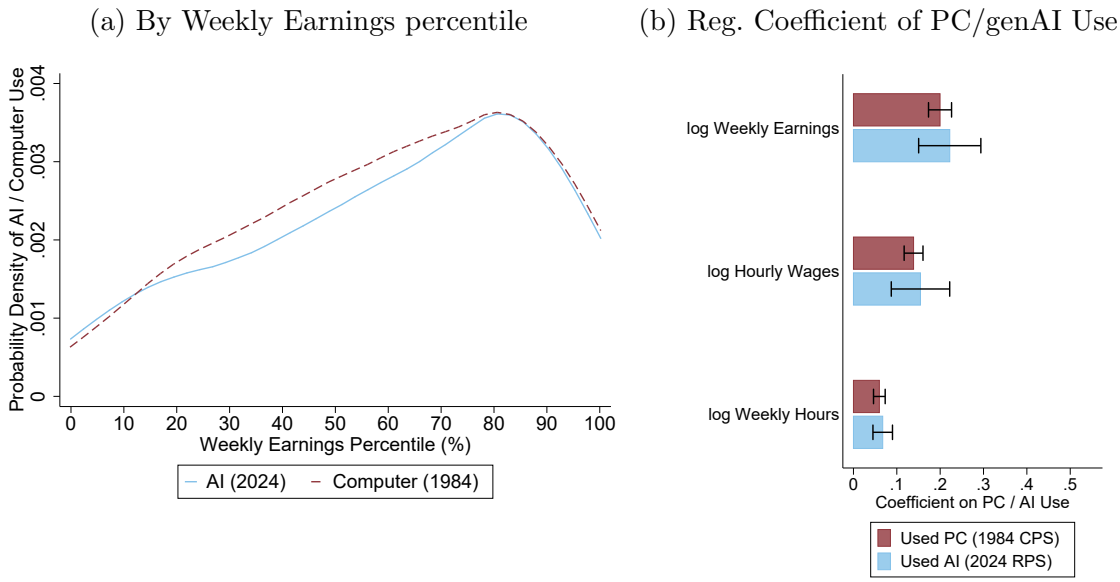
Figure 8 decomposes wage bill changes into wage growth (panel a) and employment share changes (panel b), showing that both margins contribute to the positive relationship with adoption.

Despite these limitations, the PC comparison provides valuable external validation. The similarity in adoption patterns is striking: Figure 9a shows that both genAI and PC adoption rates increase with earnings through the 80th percentile before declining. Regression analysis (Figure 9b) reveals nearly that both technologies are associated with 21 log point earnings premia, 15 log point wage premia, and 5 log point hours premia, conditional on occupation, industry, education, and demographics.[3]

[2]When we used the "occ1990dd" occupation mapping from Autor and Dorn (2013) to best facilitate comparisons pre 1982 or post 1990, we found for many occupations discontinuities in shares coinciding with switches in the occupation classification scheme that the CPS uses.

[3]Controls include indicators for female, age bins (18-29, 40-49, 50-64), bachelor's degree, graduate degree, Black, Hispanic, occupation group, and industry.

Figure 9: Earnings and Generative AI / PC Adoption



Notes: Panel (a) shows adoption rates at work for genAI and PCs by weekly earnings percentile (kernel density). Panel (b) shows the coefficient on PC adoption (red bars) and genAI adoption (blue bars) in a regression of (i) log weekly earnings, (ii) log hourly wages, and (iii) log weekly hours worked. For all regressions, the additional control variable used are indicators for: female, age 18-29, age 40-49, age 50-64, a bachelor's degree, a graduate degree, Black, Hispanic, occupation group, and industry. The data source for genAI is the August and November 2024 waves of the RPS. The data source for computers is the Outgoing Rotation Group (ORG) subset of the 1984 Computer and Internet Use Supplement of the CPS. The sample for each dataset is employed individuals ages 18-64 whose implied hourly wage is at least the federal minimum wage in the relevant year..

5.4 Comparing Realized and Predicted GenAI Impact

How does genAI's current impact compare to predictions of its future potential? Table 4 contrasts our estimates based on observed 2025 adoption with a counterfactual calibration based on Eloundou et al. (2024)'s predictions for genAI capabilities in 10 years.

Eloundou et al. (2024) construct occupation-level "exposure" measures predicting the share of tasks that genAI will be able to perform in half the time within 10 years. These estimates are used in several recent papers as inputs indicating where genAI is useful, , see e.g., Acemoglu (2024), Althoff and Reichardt (2026), and Freund and Mann (2025). We implement their predictions by setting each occupation's generic labor share θ_o equal to their exposure measure, and setting $\pi_o = 2$ for all occupations (implying 50% time savings for the median worker, since γ_i is log-normally distributed with mean zero).

Under this counterfactual, we predict that 90.8% of workers would adopt genAI and spend 34.0% of work time using it, generating a 27.2% increase in aggregate productivity. Achieving 27.2% growth over 10 years requires a compound annual growth rate of 2.4%. Comparing this to our baseline estimate of 2.0% cumulative productivity gain through 2025, we conclude that

Table 4: Model-Estimated Impact of GenAI Using Usage Data vs. Eloundou et al. (2024) Measures

	RPS (2025)	Eloundou et al. (2024)
GenAI adoption share (%)	34.9	90.8
GenAI work time share (%)	5.0	34.0
Δ Aggregate productivity (%)	+2.0	+27.2

Notes: Left column shows aggregate productivity predictions of genAI using the model-estimated parameters from the RPS data. Right column uses estimates of θ_o and π_o based on numbers from Eloundou et al. (2024)'s estimates of genAI's capabilities projected 10 years into the future. We make θ_o equal to the exposure score from Eloundou et al. (2024), and set $\pi = 2$ so that the average worker experiences 50% time savings from genAI.

genAI's current impact is modest relative to predictions of its potential.

This gap could narrow through several channels: continued improvements in genAI capabilities, learning-by-doing as workers develop complementary skills, organizational adaptation, or expanded use cases. Alternatively, the predictions may prove optimistic if genAI encounters fundamental limitations in performing nuanced tasks. Tracking genAI adoption and productivity effects over time will help discriminate between these possibilities.

6 Conclusion

This paper develops a revealed-preference approach to measuring the labor market impact of a new technology. Rather than imposing assumptions about which tasks a technology can perform, we infer its effective capabilities from observed adoption and usage behavior through the lens of a structural model. We apply this framework to generative AI using novel nationally representative survey data that jointly measure the extensive and intensive margins of use.

Our estimates imply that, as of late 2025, generative AI has had economically meaningful but modest aggregate effects. The calibrated model indicates a roughly 2 percent increase in aggregate labor productivity and a small increase in wage inequality. Productivity gains occur across all occupations in the model but are especially concentrated in a subset of high-adoption occupations, while changes in wage dispersion arise naturally from selection into adoption and equilibrium price adjustments. These findings contrast sharply with counterfactual simulations based on ex-ante predictions of genAI's capabilities, which imply substantially larger effects.

More broadly, our approach underscores the value of using observed economic behavior to discipline inference about new technologies. Assessments of technological impact often require

researchers to take a stand on what a technology can do, despite substantial uncertainty about its capabilities. By contrast, we show that adoption and usage decisions themselves provide direct information about the technology's productivity and scope of application. This perspective allows us to quantify realized effects without imposing strong ex-ante assumptions about a technology's capabilities. We validate this approach by looking back to the mass adoption of computers, showing that in hindsight computer adoption was substantially predictive of labor market changes around the mass adoption of computers.

Several extensions remain important for future work. In future work we intend to incorporate occupational choice into the model, and also to use newly designed RPS questions to measure whether generic and nuanced labor are complements or substitutes within-occupation. This will help to sharpen identification of the model and therefore of genAI's labor market effects. As genAI continues to develop, combining contemporaneous adoption data with disciplined structural modeling offers a promising path for tracking its economic consequences in real time.

References

Acemoglu, D. (2024). *The Simple Macroeconomics of AI*. Working Paper. National Bureau of Economic Research.

Althoff, L. and H. Reichardt (Jan. 2026). *Task-Specific Technical Change and Comparative Advantage*. CESifo Working Paper. Available at https://hugoreichardt.com/pdf/tstc_compadvantage.pdf. CESifo.

Autor, D. H. and D. Dorn (2013). "The growth of low-skill service jobs and the polarization of the US labor market". In: *American Economic Review* 103.5, pp. 1553–1597.

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.

Bick, A., A. Blandin, D. J. Deming, N. Fuchs-Schundeln, and J. Jessen (2026). "Mind the Gap: AI Adoption in Europe and the US". In: *Working Paper*.

Brynjolfsson, E., D. Li, and L. Raymond (2025). "Generative AI at work". In: *The Quarterly Journal of Economics* 140.2, pp. 889–942.

Choi, J. H., A. B. Monahan, and D. Schwarcz (2024). "Lawyering in the age of artificial intelligence". In: *Minnesota Law Review* 109, p. 147.

Cui, Z. K., M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz (2024). "The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers". In: *Available at SSRN 4945566*.

Dell'Acqua, F., E. McFowland, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2023). "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality". In: *Harvard Business School Technology & Operations Mgt. Unit Working Paper* 24-013.

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.

Freund, L. B. and L. F. Mann (2025). "Job Transformation, Specialization, and the Labor Market Effects of AI". In.

Goldsmith-Pinkham, P., C. Tan, and A. K. Zentefis (2026). "Human-AI Collaboration in Radiology: The Case of Pulmonary Embolism". In: *Working Paper: arXiv preprint arXiv:2601.13379*.

Kanazawa, K., D. Kawaguchi, H. Shigeoka, and Y. Watanabe (2025). "AI, skill, and productivity: The case of taxi drivers". In: *Management Science*.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). "Capital-skill complementarity and inequality: A macroeconomic analysis". In: *Econometrica* 68.5, pp. 1029–1053.

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.

Noy, S. and W. Zhang (2023). "Experimental evidence on the productivity effects of generative artificial intelligence". In: *Science* 381.6654, pp. 187–192.

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023). "The impact of ai on developer productivity: Evidence from github copilot". In: *arXiv preprint arXiv:2302.06590*.

Sato, K. (1976). "The ideal log-change index number". In: *The Review of Economics and Statistics*, pp. 223–228.

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

Vartia, Y. O. (1976). "Ideal log-change index numbers". In: *scandinavian Journal of statistics*, pp. 121–126.

Generative AI in the Labor Market: Evidence and Some Theory

Alexander Bick Adam Blandin David Deming Tyler Schumacher

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 who appear to be bots based on various open text questions.

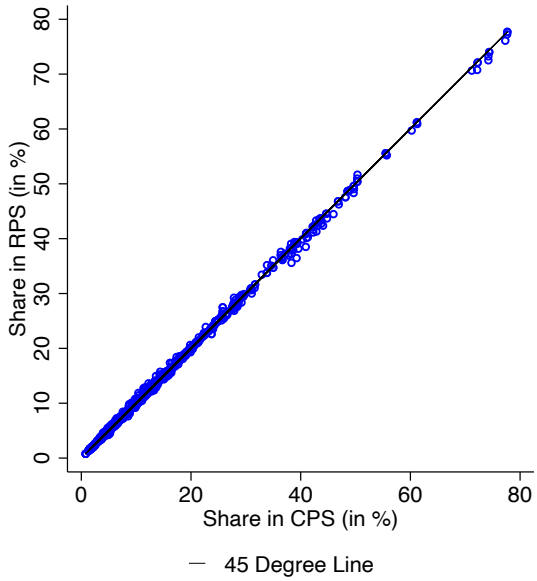
A.2 Weighting

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.

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

[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.

FIGURE A.1: Sample Composition in the Weighted RPS vs. CPS for 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).

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.

A.3 Validation Checks

Figure A.2 compare the industry shares in the RPS and CPS in the pooled 2025 waves, with and without weights, as well as college major shares in the RPS and 2024 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.

A.4 Measurement and Definition of Industries

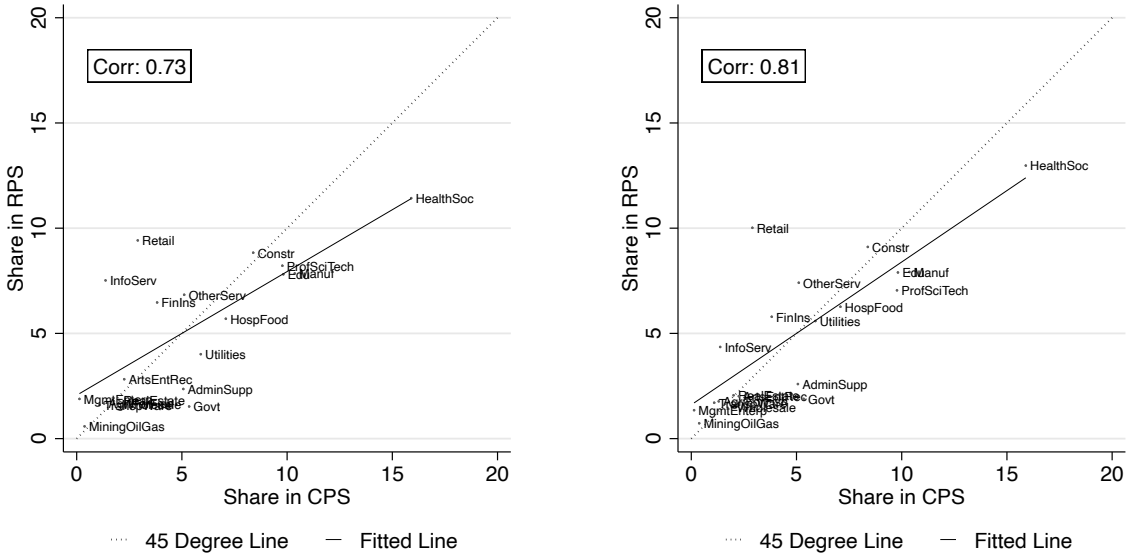
We first ask employed respondents what kind of industry or job their job is in choose their industry with a list of the broad major industries in the NAICS: Agriculture, Forestry, Fishing, and Hunting (NAICS=11); Mining, Quarrying, and Oil and Gas Extraction (NAICS=21); Utilities (NAICS=22); Construction (NAICS=23); Manufacturing (NAICS=31-33); Wholesale Trade (NAICS=42); Retail Trade (NAICS=44-45); Transportation and Warehousing (NAICS=48-49); Information Services (including Publishing and Media) (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 Services (including 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).

For each industry, except construction, we provide them with a follow-up list on more detailed industries within that broad group. Anyone who selected "Other Services" for the broad industry and "Other Personal Services" as the more detailed industry was asked the following question: "We would like to know some more details about what kind of business or industry this job is. Please include the main activity, product, or service provided at the location where you are employed. (For example: elementary school, residential construction)." We provided an open text field for respondents to type their answers. These responses were routed to an industry and occupational coder at the National Institute for Occupational Safety and Health (NIOSH) (https://csams.cdc.gov/nioccs/SingleCoding.aspx). Based on the provided responses, the system suggested the top five 6-digit NAICS codes plus an "Other" option. We assigned industry based on their choice, with "Other" being classified as "Other

FIGURE A.2: Validation Checks – Industry and Major Shares

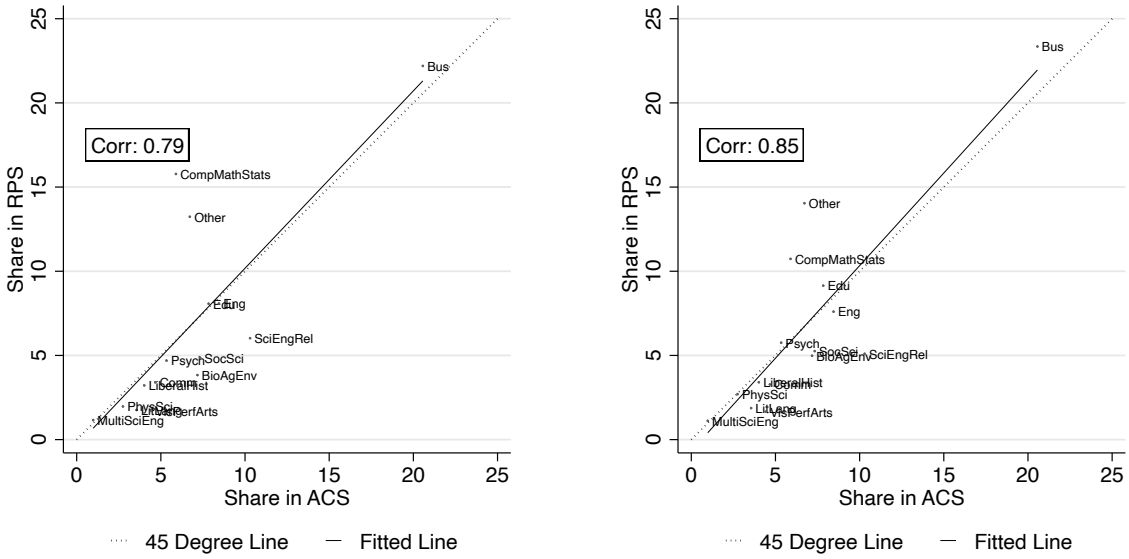
(a) Industry: Unweighted

(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 2025 RPS and CPS with sample sizes of 153279 and 153279, respectively. Data samples for the college major comparison are respondents with a college degree ages 18-64 in the pooled 2025 RPS with sample sizes of 7668 and 161705 in the 2024 ACS, respectively.

Services - Other Personal Services." We implemented this last step since in RPS waves prior to August 2024, "Other Services" was the most overrepresented industry relative the CPS. We attributed this to the difficulty for respondents in the service sector to select the right category. Our updated procedure resolved this issue.

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).

A.5 Measurement and Definition 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.*

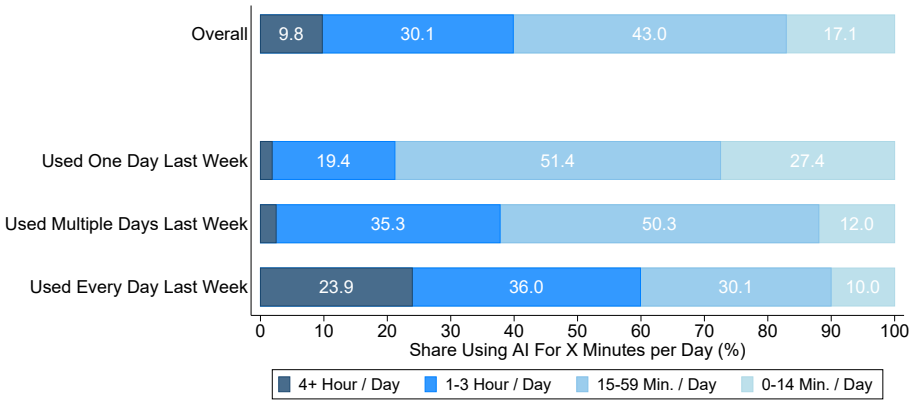
Once they have starting typing, respondents are presented with autocomplete suggestions covering over 40,000 occupations from O*NET and the Occupational Outlook Handbook. For job titles that do not exactly match a unique SOC code, we present respondents with a choice of the five most probabilistic matches 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 et al., 2024). We proceed in the same way for respondents whose text response does not match one of the autosuggestions exactly.

Throughout the draft, our occupation group often corresponds 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).

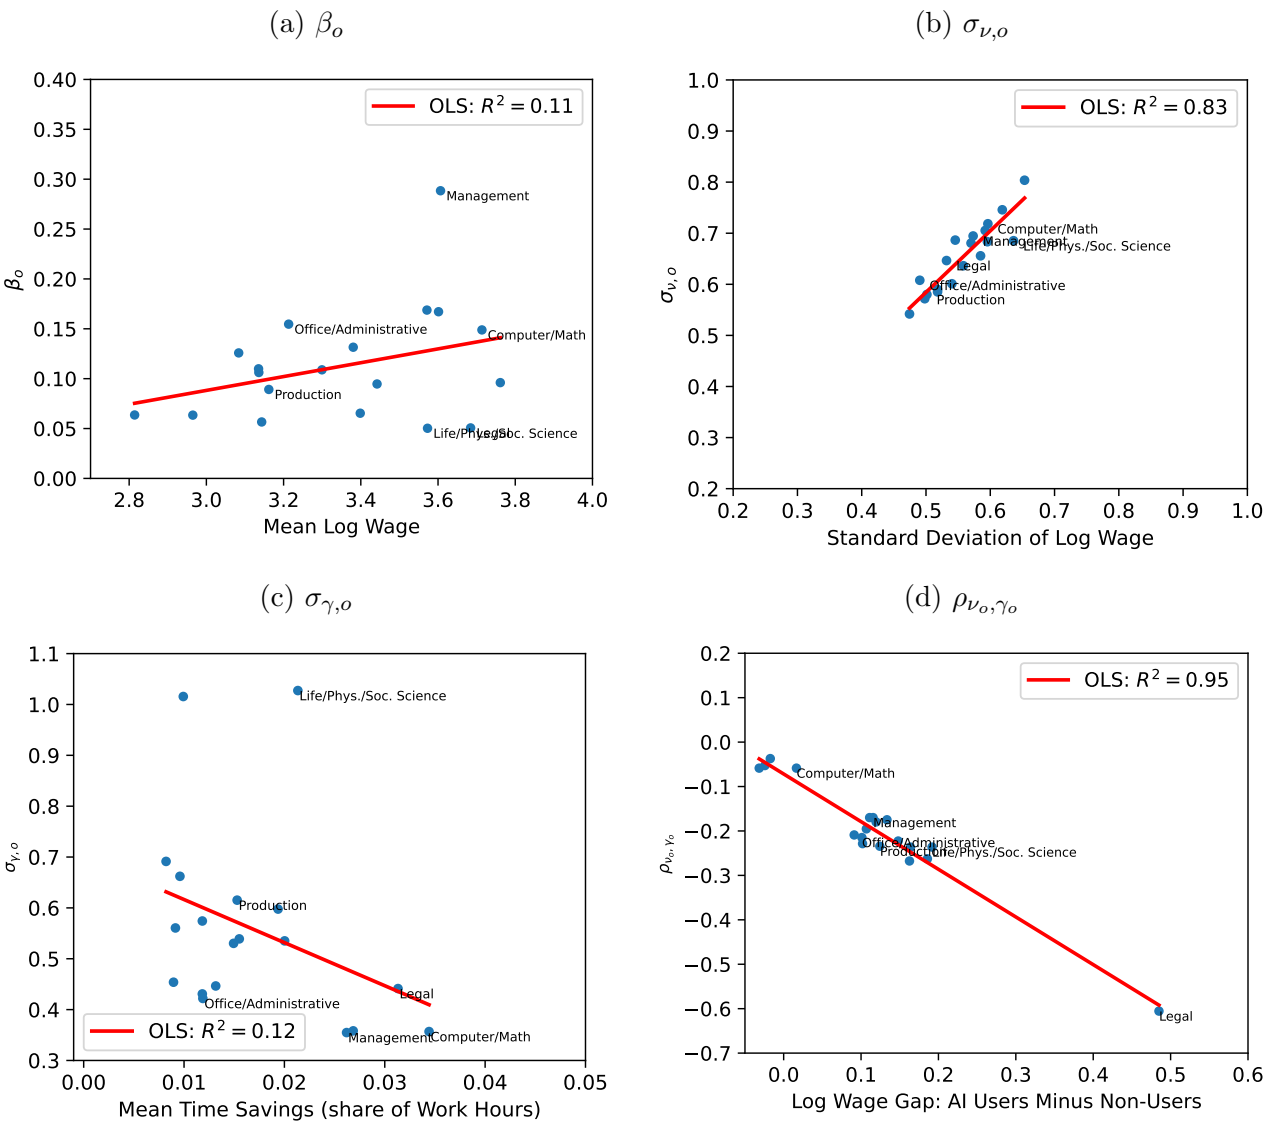
B Figures

FIGURE B.3: Time Spent Using Generative AI



Notes:

FIGURE B.4: Calibrated Parameters Versus Related Moments



Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Parameters are calibrated using data moments as described in Section 5.1. Each figure plots a calibrated model parameter by occupation against its most closely related target moment. The red line is an OLS fit of the y axis onto the x axis, and the box displays the R^2 from running this univariate OLS regression at the occupation level.

C Tables

Table C.1: Calibration: β_o

Occupation	Mean Log Wage Data	Mean Log Wage Model	β_o
Management	3.606	3.607	0.288
Business and financial	3.601	3.601	0.167
Computer and math	3.714	3.714	0.149
Architecture/engineering	3.761	3.761	0.096
Life/physical/social Science	3.573	3.573	0.050
Community/social service	3.442	3.442	0.095
Legal	3.684	3.684	0.051
Education/training/library	3.380	3.380	0.132
Arts/entertainment	3.399	3.399	0.065
Healthcare technical	3.571	3.571	0.169
Healthcare support	2.965	2.965	0.063
Food prep and service	2.814	2.814	0.064
Installation/maintenance/grounds	3.136	3.136	0.106
Personal care/service	3.143	3.143	0.057
Sales	3.084	3.084	0.126
Office and admin support	3.213	3.213	0.155
Construction/extraction	3.299	3.299	0.109
Production	3.162	3.162	0.089
Transportation/material moving	3.135	3.135	0.110

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Parameters are calibrated using data moments as described in Section 5.1. Column 1 shows the name of the occupation. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.

Table C.2: Calibration: $\sigma_{\nu,o}$

Occupation	SD of Log Wage		$\sigma_{\nu,o}$
	Data	Model	
Management	0.573	0.573	0.695
Business and financial	0.585	0.585	0.656
Computer and math	0.596	0.596	0.719
Architecture/engineering	0.558	0.558	0.636
Life/physical/social Science	0.636	0.636	0.685
Community/social service	0.540	0.540	0.601
Legal	0.532	0.532	0.646
Education/training/library	0.596	0.596	0.683
Arts/entertainment	0.619	0.619	0.746
Healthcare technical	0.518	0.518	0.590
Healthcare support	0.474	0.474	0.542
Food prep and service	0.518	0.518	0.585
Installation/maintenance/grounds	0.546	0.546	0.687
Personal care/service	0.653	0.653	0.804
Sales	0.592	0.592	0.705
Office and admin support	0.490	0.490	0.608
Construction/extraction	0.570	0.570	0.681
Production	0.501	0.501	0.580
Transportation/material moving	0.498	0.498	0.572

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Column 1 shows the name of the occupation. Parameters are calibrated using data moments as described in Section 5.1. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.

Table C.3: Calibration: π_o

Occupation	Share Using genAI Data	Share Using genAI Model	π_o
Management	0.542	0.542	1.036
Business and financial	0.423	0.423	0.890
Computer and math	0.635	0.636	1.128
Architecture/engineering	0.381	0.381	0.849
Life/physical/social Science	0.464	0.464	0.911
Community/social service	0.252	0.252	0.641
Legal	0.567	0.567	1.076
Education/training/library	0.434	0.434	0.913
Arts/entertainment	0.554	0.554	1.048
Healthcare technical	0.303	0.303	0.743
Healthcare support	0.233	0.233	0.664
Food prep and service	0.197	0.196	0.420
Installation/maintenance/grounds	0.198	0.198	0.680
Personal care/service	0.262	0.262	0.760
Sales	0.314	0.314	0.772
Office and admin support	0.267	0.267	0.769
Construction/extraction	0.314	0.314	0.804
Production	0.330	0.330	0.762
Transportation/material moving	0.198	0.198	0.555

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Column 1 shows the name of the occupation. Parameters are calibrated using data moments as described in Section 5.1. Column 1 shows the name of the occupation. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.

Table C.4: Calibration: θ_o

Occupation	Mean genAI Time Among Users		
	Data	Model	θ_o
Management	0.170	0.170	0.170
Business and financial	0.109	0.109	0.109
Computer and math	0.170	0.170	0.170
Architecture/engineering	0.112	0.112	0.112
Life/physical/social Science	0.061	0.061	0.061
Community/social service	0.100	0.100	0.100
Legal	0.152	0.152	0.152
Education/training/library	0.121	0.121	0.121
Arts/entertainment	0.166	0.166	0.166
Healthcare technical	0.110	0.110	0.110
Healthcare support	0.124	0.124	0.124
Food prep and service	0.094	0.094	0.094
Installation/maintenance/grounds	0.186	0.186	0.186
Personal care/service	0.179	0.179	0.179
Sales	0.145	0.145	0.145
Office and admin support	0.180	0.180	0.180
Construction/extraction	0.151	0.151	0.151
Production	0.120	0.120	0.120
Transportation/material moving	0.112	0.112	0.112

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Parameters are calibrated using data moments as described in Section 5.1. Column 1 shows the name of the occupation. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.

Table C.5: Calibration: $\sigma_{\gamma,o}$

Occupation	Self-Reported Time Savings		$\sigma_{\gamma,o}$
	Data	Model	
Management	0.026	0.026	0.355
Business and financial	0.019	0.019	0.598
Computer and math	0.034	0.034	0.357
Architecture/engineering	0.016	0.015	0.539
Life/physical/social Science	0.021	0.021	1.027
Community/social service	0.010	0.010	0.662
Legal	0.031	0.031	0.441
Education/training/library	0.020	0.020	0.535
Arts/entertainment	0.027	0.027	0.358
Healthcare technical	0.012	0.012	0.574
Healthcare support	0.009	0.009	0.561
Food prep and service	0.010	0.010	1.016
Installation/maintenance/grounds	0.009	0.009	0.454
Personal care/service	0.012	0.012	0.431
Sales	0.015	0.015	0.530
Office and admin support	0.012	0.012	0.422
Construction/extraction	0.013	0.013	0.446
Production	0.015	0.015	0.615
Transportation/material moving	0.008	0.008	0.691

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Parameters are calibrated using data moments as described in Section 5.1. Column 1 shows the name of the occupation. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.

Table C.6: Calibration: ρ_{γ_o,ν_o}

Occupation	Wage Premium (Users Minus Non) Data	Model	ρ_{γ_o,ν_o}
Management	0.116	0.115	-0.170
Business and financial	-0.017	-0.017	-0.037
Computer and math	0.017	0.016	-0.059
Architecture/engineering	0.164	0.163	-0.236
Life/physical/social Science	0.192	0.191	-0.236
Community/social service	-0.024	-0.024	-0.053
Legal	0.485	0.485	-0.605
Education/training/library	0.119	0.120	-0.180
Arts/entertainment	0.133	0.134	-0.175
Healthcare technical	0.107	0.106	-0.195
Healthcare support	-0.032	-0.031	-0.059
Food prep and service	0.102	0.102	-0.229
Installation/maintenance/grounds	0.163	0.163	-0.268
Personal care/service	0.111	0.111	-0.170
Sales	0.186	0.186	-0.263
Office and admin support	0.101	0.101	-0.215
Construction/extraction	0.148	0.149	-0.223
Production	0.124	0.125	-0.234
Transportation/material moving	0.091	0.091	-0.209

Notes: Target moments are calculated from the pooled RPS Feb. 2025, May 2025, and Aug. 2025 survey waves. Parameters are calibrated using data moments as described in Section 5.1. Column 1 shows the name of the occupation. Column 2 shows the data target in calibration. Column 3 shows the corresponding data moment generated by the model. Colum 4 shows the calibrated value of the most closely related parameter to the data moment being targeted.