NBER WORKING PAPER SERIES

THE RAPID ADOPTION OF GENERATIVE AI

Alexander Bick
Adam Blandin
David J. Deming

Working Paper 32966
http://www.nber.org/papers/w32966

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We thank Kevin Bloodworth II and Eli Tsung for excellent research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the National Bureau of Economic Research. The survey was funded in part by the Walmart Foundation, grant 252882.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Alexander Bick, Adam Blandin, and David J. Deming. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Rapid Adoption of Generative AI
Alexander Bick, Adam Blandin, and David J. Deming
NBER Working Paper No. 32966
September 2024
JEL No. J24, O33

ABSTRACT

Generative Artificial Intelligence (AI) is a potentially important new technology, but its impact on the economy depends on the speed and intensity of adoption. This paper reports results from the first nationally representative U.S. survey of generative AI adoption at work and at home. In August 2024, 39 percent of the U.S. population age 18-64 used generative AI. More than 24 percent of workers used it at least once in the week prior to being surveyed, and nearly one in nine used it every workday. Historical data on usage and mass-market product launches suggest that U.S. adoption of generative AI has been faster than adoption of the personal computer and the internet. Generative AI is a general purpose technology, in the sense that it is used in a wide range of occupations and job tasks at work and at home.

Alexander Bick
Federal Reserve Bank of St. Louis
One Federal Reserve Bank Plaza
St. Louis, MO 63102
alexander.bick@stls.frb.org

Adam Blandin
Department of Economics
Vanderbilt University
VU Station B #351819
2301 Vanderbilt Place
Nashville, TN 37235-1819
adam.blandin@vanderbilt.edu

David J. Deming
Harvard Kennedy School
Malcolm Wiener Center for Social Policy
79 JFK St
Cambridge, MA 02138
and NBER
david_deming@harvard.edu

A data appendix is available at http://www.nber.org/data-appendix/w32966

1 Introduction

Generative Artificial intelligence (AI) has rapidly emerged as a potentially transformative workplace technology. The large language model (LLM) ChatGPT debuted in November 2022, and by March 2024 the most common generative AI tools had been accessed more than three billion times by hundreds of millions of users each month (Liu and Wang, 2024). Several recent studies have found that generative AI improves worker productivity (Brynjolfsson, Li, and Raymond, 2023; Cui et al., 2024; Dell'Acqua et al., 2023; Noy and Zhang, 2023; Peng, Kalliamvakou, Cihon, and Demirer, 2023). Yet other studies expect only modest impacts of AI on work, depending on how well AI substitutes for complex job tasks (Acemoglu, Autor, Hazell, and Restrepo, 2022; Bloom, Prettner, Saadaoui, and Veruete, 2024).

The ultimate impact of generative AI on the economy depends on how quickly and intensively the technology is adopted. Yet there is little systematic evidence of the extent to which generative AI is used at work and at home. Who uses generative AI, how much do they use it, and what do they use it for?

This paper presents results from the first nationally representative U.S. survey of generative AI adoption at work and at home. Our data come from the Real-Time Population Survey (RPS), a nationwide survey that asks the same core questions and follows the same timing and structure of the Current Population Survey (CPS), the monthly labor force survey conducted by the U.S. Census Bureau for the Bureau of Labor Statistics (BLS). We benchmark our survey to national estimates of employment and earnings, ensuring representativeness. Prior research has used the RPS methodology to study work from home during the COVID-19 pandemic, among other topics (Bick and Blandin, 2023; Bick, Blandin, and Mertens, 2023). The survey structure allows us to easily add and modify questions and to track generative AI usage over time within a large representative sample of the U.S. workforce.

We find that in August 2024, 39.4 percent of the U.S. population age 18-64 used generative AI, with 32.0 percent using it at least once during the week they were surveyed; 28.0 percent of employed respondents used generative AI at work, with most (24.2 percent) using it at least weekly; and 10.6 percent of the employed reporting daily usage at work. Generative AI use is more common outside of work, but less intensive. One in three respondents (32.7 percent) said that they used generative AI outside of work, but only 6.4 percent used it outside of work every day. ChatGPT is by far the most commonly used generative AI program, although many others are reported, including tools that embed AI inside standard office software packages (e.g., Microsoft Copilot).

How does the speed and intensity of the adoption of generative AI compare with other technologies? Prior research shows that better technologies are adopted faster, and the speed and intensity of technology adoption across countries is highly correlated with economic growth (Beaudry, Doms, and Lewis, 2010; Comin and Hobijn, 2010; Comin and Mestieri, 2018). We compare the speed of adoption of generative AI with two other technologies - the personal computer (PC) and the internet - using data from the CPS Computer and Internet Use Supplement and the International Telecommunication Union (ITU).[1]

Generative AI has been adopted at a faster pace than PCs or the internet. Generative AI has a 39.5 percent adoption rate after two years, compared with 20 percent for the internet after two years and 20 percent for PCs after three years (the earliest we can measure it). This is driven by faster adoption of generative AI at home compared with the PC, likely because of differences in portability and cost. We find similar adoption rates at work for PCs and for generative AI. (Note that we cannot separate internet usage between home and work.)

Some scholars argue that generative AI could reduce workplace inequality (e.g., Autor 2024). However, similar to PC adoption, generative AI usage is more common among younger, more educated, and higher-income workers. This is notable because the PC revolution was followed by rising labor market inequality, with computers substituting for routine "middle-skill" tasks while complementing high-skilled labor (Autor, Levy, and Murnane, 2003). The one exception is gender. We find that men are 9 percentage points more likely to use generative AI at work and 7 percent more likely to use it at home. In contrast, PC adoption at work was more common for women, possibly because of the transition between typewriters and word processors and the high female share of secretaries and other administrative occupations.

Generative AI is used by workers in a broad range of occupations to perform many different workplace tasks. Generative AI adoption is most common in management, business, and computer occupations, with usage rates exceeding 40 percent. Still, one in five "blue collar" workers and one in five workers without a college degree use generative AI regularly on the job as well. This is consistent with Eloundou, Manning, Mishkin, and Rock (2024), who compare generative AI capabilities with the task content of work and find that many occupations will be affected. We asked workers whether they used generative AI to help them perform ten different job tasks, including writing, searching for information, interpreting data or text, coding, data analysis, and others. Among generative AI users at work, all ten tasks in our list had usage rates of at least 25 percent, with writing, interpreting, and administrative help ranked as the most helpful.

[1] The CPS asked respondents about their PC and Internet usage at home and at work, and we closely replicate the wording of their questions for generative AI usage to facilitate comparison. For each technology we measure the speed of adoption after the first mass market product was released. For generative AI, it is the November 2022 release of ChatGPT. For PCs it is the release of the IBM PC in August 1981. For the internet it is April 1995, when the National Science Foundation (NSF) decommissioned NSFNet and allowed the internet to carry commercial traffic. This was also the year of Netscape's initial public offering (IPO).

Using responses to questions about both the frequency and the intensity of work usage, we estimate that between 0.5 and 3.5 percent of all work hours in the U.S. are currently being assisted by generative AI. If we assume that generative AI increases task productivity by 25 percent - the median estimate across five randomized studies - this would translate to increase in labor productivity of between 0.125 and 0.875 percentage points at current levels of usage. However, this calculation assumes that small-scale studies are externally valid and should be treated with caution.

Our results line up broadly with other published surveys of generative AI usage. The most similar study to our is Humlum and Vestergaard (2024), who survey a representative sample of workers in eleven occupations in Denmark about their usage of ChatGPT at work. We find broadly similar usage rates in the occupations covered by both surveys, although the lack of clean correspondence between job codes across countries makes an exact comparison difficult.[2] A Pew Research Center survey conducted in February 2024 found that 23 percent of adults in their survey reported ever having used ChatGPT, with higher rates of adoption for younger and more educated respondents (McClain, 2024).[3] A Reuters online survey conducted in six countries in April 2024 found that 18 percent of U.S. respondents used ChatGPT at least weekly, compared to less than 10 percent in Argentina, Denmark, France, Japan, and the United Kingdom (Fletcher and Nielsen, 2024).

Our study shows that the generative AI is being adopted much faster than previous waves of AI technology. McElheran et al. (2024) find that less than 6 percent of firms had used frontier AI technologies such as machine learning, computer vision, and natural language processing in 2017. Similarly, Acemoglu, Autor, Hazell, and Restrepo (2022) find that only about 3 percent of U.S. firms had adopted predictive AI tools between 2016 and 2018 and Humlum and Meyer (2022) found similarly low adoption rates in Denmark in 2017.

Generative AI may be adopted more rapidly because it targets consumers rather than firms. Bonney et al. (2024) report firm-level AI adoption using the Business Trends and Outlook Survey (BTOS), a Census Bureau study that asked firms about AI usage between December 2023 and February 2024. They found that AI adoption rose over the survey period from 3.7 percent in December to 5.4 percent in February, which is a rapid rise but still far below our estimates. Like Bonney et al. (2024), we also find that generative AI usage is higher in large firms. Still, gaps by firm size are far too small to explain the discrepancy between firm and worker usage, suggesting that workers are using generative AI even in firms that haven't officially adopted it.[4]

[2] Humlum and Vestergaard (2024) also finds similar demographic differences in AI usage, although they find much larger gender gaps (about 20 percentage points) in Denmark than we do in the U.S.

[3] The Pew survey included adults over age 65, who had very low adoption rates, implying a higher estimate of around 27 percent ages 18-64.

[4] In our data, 27 percent of workers report that their employer encourages them to use generative AI. Employer

The paper proceeds as follows. Section 2 describes the survey methodology and data. Section 3 presents our main results. Section 4 concludes.

2 Data Source and Measurement

2.1 The Real-Time Population Survey (RPS)

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

The RPS is designed to mirror the Current Population Survey (CPS) along key dimensions. The RPS survey 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, where necessary, to construct sample weights.

However, the RPS also collects information not contained in the CPS. Novel questions in the RPS have been previously used to study trends in employee reallocation across firms, work from home, and interstate migration as well as the relationship between inflation and job search (Bick and Blandin, 2023; Bick, Blandin, and Mertens, 2023; Bick, Blandin, Mertens, and Rubinton, 2024; Pilossoph and Ryngaert, 2023). In June and August 2024, the RPS introduced a module designed to measure Generative AI use both at work and at home.

The RPS produced very similar statistics for employment, hours worked, earnings, industry composition, and employee tenure during the pandemic (Bick and Blandin, 2023). This is in part due to careful, iterative fielding of the survey to match question wording and other details of the CPS.

encouragement is highly correlated with AI use: 82.9 percent of workers who report encouragement also report using generative AI, compared with only 7.1 percent of workers who report no encouragement.

2.2 Sample

Qualtrics panel respondents are recruited online and can participate in exchange for 30 to 50 percent of the fee charged by Qualtrics (we paid $6.90 per completed survey).[5] 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 several broad demographic characteristics: gender, age, race and ethnicity, education, marital status, number of children in the household, Census region, and household income over the past 12 months.

We fielded a pilot survey in June 2024 and received 2,551 responses. We then launched our full survey in August 2024 and received 5,014 responses. Both surveys were fielded during the same weeks that the CPS conducted its corresponding surveys. We dropped 14 and 33 respondents, respectively, from the June and August surveys because they reported their industry and/or occupation as military. An additional 9 respondents were dropped from each survey because they reported being employed but also reported being homemakers, retired, or unemployed as their occupation.

The August 2024 survey included several improvements on the pilot survey in June. It was a larger sample, and also included questions about intensity of generative AI usage and about the breadth of usage across job tasks. We also improved the coding of industry and occupations and made several other small improvements. We found very similar results for the questions that overlapped between surveys, with slightly higher work usage in August, although the increase was not statistically different from zero. Thus, for simplicity, we report only the August results in the main paper. Appendix B replicates the main figures of the paper using data from the June 2024 survey.

The first two columns of Table 1 compare the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure for our main survey in August 2024. The most notable discrepancies are that individuals aged 18 to 24 and with no more than a high school degree are under-represented in the RPS relative to the CPS, while individuals with household income of $50,000 or less are over-represented. The bottom panel of Table 1 compares employment status in the CPS and RPS, statistics that have not been targeted in the sampling procedure. Individuals classified as unemployed according to the CPS definition are over-represented in the RPS.

In columns three and four of Table 1, we compare the demographic composition between

[5] The median time to complete the survey is 10 and a half minutes, implying an hourly pay rate for respondents of roughly $11.80 to $19.70.

Table 1: Sample Composition in the August 2024 CPS and RPS

	Everyone		*Employed*	
	CPS	RPS	CPS	RPS
	(1)	(2)	(3)	(4)
Gender: Women	50.4	52.2	47.2	46.7
Age				
18-24	14.9	8.1	12.0	9.2
25-34	22.2	23.0	23.9	24.4
35-44	22.2	24.7	24.5	26.1
45-54	20.1	22.0	21.8	22.5
55-64	20.6	22.1	17.8	17.8
Race/Ethnicity				
Non-hispanic White	56.6	56.2	57.9	59.0
Non-hispanic Black	12.9	13.0	12.1	10.9
Hispanic	20.5	20.6	20.2	19.5
Other	10.0	10.3	9.9	10.5
Education				
Highschool or less	37.5	32.4	33.2	25.5
Some college/Associate's degree	25.6	26.7	25.1	27.1
Bachelor's or Graduate degree	36.9	40.8	41.7	47.4
Marital Status: Married	50.2	48.6	52.9	52.7
Number of children				
0	59.4	57.6	58.5	54.2
1	17.3	19.3	17.5	20.8
2	14.5	15.7	15.2	17.9
3+	8.8	7.4	8.7	7.1
Household Income in Last 12 Months				
$0-$50,000	25.8	32.7	19.8	22.7
$50,000-$100,000	30.3	28.0	31.0	31.1
$100,000+	43.9	39.3	49.2	46.2
Region				
Midwest	16.9	18.6	16.8	19.3
Northeast	20.3	18.1	20.9	17.4
South	39.0	37.8	38.3	37.7
West	23.8	25.5	24.0	25.6
Employment Status				
Employed, at work last week	71.0	67.4		
Employed, absent from work last week	3.2	3.1		
Unemployed	3.4	8.7		
Not in the labor force	22.5	20.8		
Observations	58068	4972	42987	3506

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

the CPS and RPS for employed respondents. This improves balance overall, although there are still some discrepancies.

2.3 Sample Weights and Validation

To address these discrepancies, we construct sample weights using the iterative proportional fitting (raking) algorithm of Deming and Stephan (1940). Our application of the raking algorithm ensures that the weighted sample proportions across key demographic characteristics match those in the CPS. We also use more disaggregated categories for education and marital status than those included in the Qualtrics sampling targets and interact all of these categories with gender. In addition, our sampling weights replicate the breakdown of employment status, both in the aggregate and conditional on our set of targeted characteristics for which use more aggregated groups to ensure sufficiently large cell sizes for some variables. We also include occupation in our weighting scheme. This requires us to drop another 108 and 112 observations due to missing occupation codes for the June and August surveys, respectively.[6] Appendix C.1 provides details on the categories targeted by our weighting scheme.

Bick and Blandin (2023) shows that the RPS replicates the CPS well along many dimensions neither targeted by the sampling procedure nor included in the weighting scheme, including usual and actual weekly hours worked, the share of workers who are paid hourly, the weekly earnings distribution, industry composition, and job tenure. Panels (a) and (b) of Figure 1 compare the usual weekly earnings distribution in the RPS to the CPS, without and with weights, respectively.[7] The unweighted distributions are already similar, and the weights improve the fit further.

Panels (c) and (d) of Figure 1 compare occupation shares in the RPS and the CPS, unweighted and weighted respectively. The two samples line up fairly well without weighting. The correlation between samples is 0.87, with management occupations particularly overrepresented in the RPS. For this reason, we apply occupation weights to all further analyses. Panel (d)

[6] We drop another 72 and 157 employed respondents after constructing weights because we lack information on their Generative AI usage at work last week. Almost all of these dropped respondents were classified as "employed, absent from work last week" as they by construction cannot have used Generative AI last week. Accounting for all individuals dropped from our analysis, this leaves us with a final sample size of 94.9% and 96.3% of the initially collected responses for the June and August 2024 survey, respectively.

[7] To ensure comparability and minimize concerns over measurement error, we restrict both the RPS and CPS samples to individuals with (i) weekly earnings below the CPS topcode of $3,960.00, (ii) an implied hourly wage of at least the federal minimum wage of $7.25. In the CPS, we subsequently drop 31.6% of individuals because they do not report all components required to calculate weekly earnings. In the remaining sample, 3.8% lack the necessary information on usual hours to calculate an hourly wage; among the remaining individuals, 1.1% earn less than $7.25 per hour and 3.3% earn more than $3,960.00. The respective numbers in the RPS are 1.0%, 1.7%, 8.0%, and 11.3%. Hence, while the RPS features a substantially larger share of individuals below the minimum wage and above the earnings topcode, the total number of observations dropped in the RPS is also substantially lower than in the CPS.

Figure 1: Validation Checks

(a) Weekly Earnings Percentiles: Unweighted

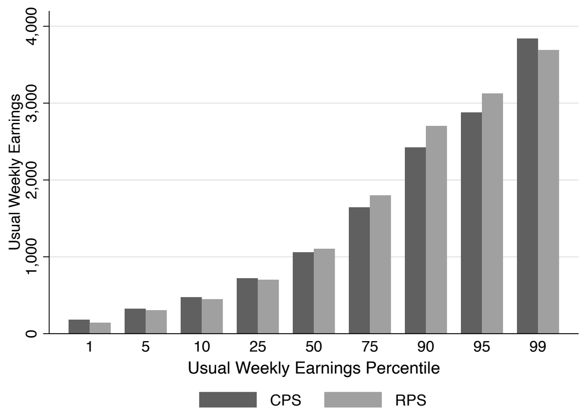

(b) Weekly Earnings Percentiles: Weighted

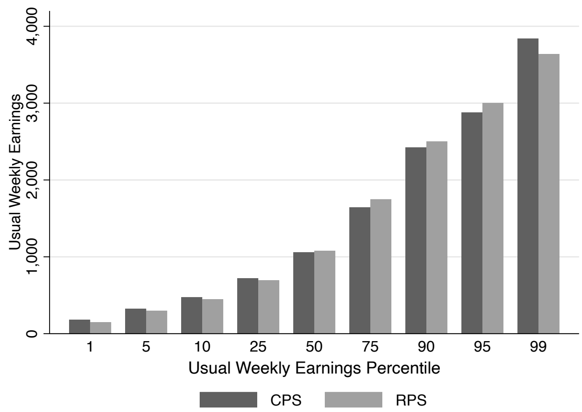

(c) Occupation Shares: Unweighted

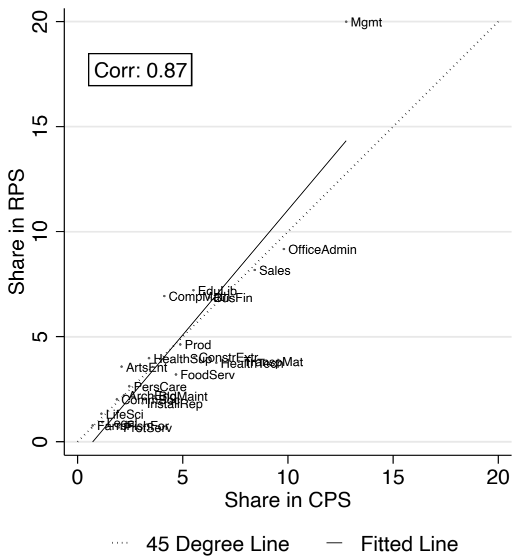

(d) Occupation Shares: Weighted

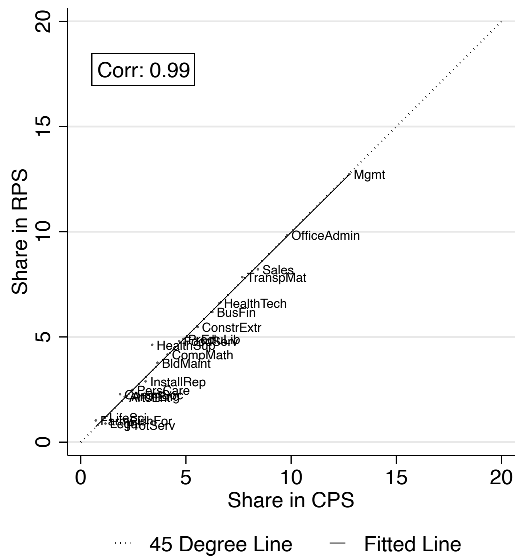

Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the sample of RPS respondents in both figures. All figures use weighted CPS data. Data samples for the weekly earnings figures are employees ages 18-64 in the August 2024 RPS and CPS-ORG with weekly earnings below the CPS topcode of $3,960.00 and an implied hourly wage of at least the federal minimum wage of $7.25. Sample sizes for the RPS and CPS are 2184 and 6078, respectively. Data samples for occupation are employed respondents ages 18-64 in the August 2024 RPS and CPS with sample sizes of 3216 and 42987, respectively.

shows that this adjustment mechanically balances the sample on occupation.[8]

2.4 Measurement of Generative AI Use

Definition and opening question. The Generative AI module begins with a definition of Generative AI:

> *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.*

Because Generative AI is a relatively new technology, we believed it was important to provide both a definition of the concept and some specific examples. We avoided mentioning specific generative AI methods such as "Large Language Models" because they seemed too technical for a general audience and because we wanted our definition to include a broad array of methods. At the same time, we mention examples of some popular generative AI products because we thought some respondents may be more familiar with those product names than with the broader concept of generative AI.

After defining generative AI, the module asks respondents whether they had heard of the concept prior to the survey. Respondents who answer "No" skip the remainder of the module, while those who answer "Yes" advance to the next question in the AI module.

Generative AI use at work. For employed respondents, the next question asks about Generative AI use at work:[9]

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

This question is designed to mirror an analogous question from the CPS Computer and Internet Use Supplement, which we discuss in Section 2.5. Respondents who answer "No" skip the remainder of the work-related generative AI questions. Respondents who answer "Yes" are asked additional questions about generative AI use at work, which fall into two broad categories. The first category relates to the intensity of generative AI use. We ask them within the last week, how many days they used generative AI, and on those days, how much time per day on

[8] Appendix Figure C.2 presents analogous plots for industry and college major. The unweighted correlation between occupation, industry and college major shares in the RPS and the CPS is 0.87, 0.88, and 0.80, respectively. Our weighting procedure merges the three smallest occupations with other close comparisons, which explains the very small discrepancies in Panel (d). In Appendix Section C.3, we show the same validation figures for the June 2024 survey. The results are similar to the August 2024 survey.

[9] Respondents with multiple jobs are told to refer to their "main job," which is defined, consistent with the CPS, as the job in which they normally work the most hours.

average. The second category asks which specific products the respondent used, which specific tasks generative AI helped with, and some broader questions about the uses and benefits of Generative AI.

Generative AI use outside of work. The final portion of the Generative AI module asks about use outside of work:[10]

Do you use Generative AI outside your job? (No/Yes)

Respondents who answer "No" to this question skip to the end of the generative AI module. Respondents who answer "Yes" are asked a set of additional questions analogous to those in the "at work" portion of the module.

2.5 Measurement of Computer and Internet Use

Beginning in 1984 the CPS fielded an occasional survey supplement with questions on computer and internet use, the Computer and Internet Use supplement (CIU). The questions relevant for our study were fielded by the CIU in 1984, 1989, 1993, 1997, 2001, 2003, 2007, and 2009. All CPS respondents who received the basic CPS questions also received the CIU questions.

We focus on two sets of questions from the CIU supplement that refer to computer use. The first question asks about computer use at work:[11]

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

The second question asks about computer use at home:

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

The CIU asks about computer use "at home," whereas we ask about Generative AI use "outside of work." While this means that our phrasing is not exactly the same as the CPS question, our broader phrasing has two advantages over the CIU phrasing. First, work from home is not uncommon today and asking about Generative AI use at home would not allow us to cleanly separate work from non-work uses. Second, many people access Generative AI using mobile devices, and asking only about use at home may not capture non-work use outside of the home.

We also use a question from the CIU about internet use:

[10] Non-employed respondents are simply asked "Do you use Generative AI?"

[11] This question and the one below were asked in the 1984, 1989, 1993, 1997, 2001, and 2003 waves of the CIU. The 2007 and 2009 waves do not contain the computer and internet use questions that we rely on in this paper. Prior to the 2001 wave, the question is phrased "Do you directly use ..."; since 2001, the question omits the word "directly." Our question about Generative AI use at work also omits the word "directly."

Do you use the internet at any location? (No/Yes)

This question was asked in the 2001, 2003, 2007, and 2009 waves of the CIU, and unlike the computer-related questions it does not condition on location.

To calculate internet adoption before 2001, we use data from the International Telecommunication Union (ITU). The ITU, in collaboration with the World Bank, has collected internet usage data in the U.S. and other countries since 1995. They combine internet usage data from national regulatory authorities and service providers on the number of subscribers, which they use to estimate the proportion of the population with internet access (Peña-López et al., 2009).

3 Results

Figure 2 presents our main results. The first bar shows that 39.4 percent of all August 2024 RPS respondents say that they used generative AI, either at work or at home. About 32 percent of respondents reported using generative AI at least once in the week prior to the survey, while 10.6 percent reported using it every day last week. About 28 percent of employed respondents used generative AI at work in August 2024, with the vast majority (24.1 percent) using it at least once in the last week and 10.9 percent using it daily. Usage outside of work was more common (32.7 percent), but slightly less intensive, with 25.9 percent using it at least once in the last week and 6.4 percent using it every day. Appendix Figure A.1 presents the share of respondents using specific generative AI products. ChatGPT is used most often (28.5 percent), followed by Google Gemini (16.3 percent).

Figure 3 presents adoption rates at work by gender, age, education, and college major. The survey shows 32 percent of men use generative AI at work, compared with 23 percent of women. Generative AI use declines with age, from about 34 percent for workers under age 40 to 17 percent for workers age 50 or above. About 40 percent of workers with a bachelor's degree or more use generative AI at work, compared with about 20 percent for those without a college degree. 46 percent of workers who majored in science, technology, engineering, or mathematics (STEM) use generative AI at work, compared with 40 percent for workers who majored in business, economics, or communication and 22 percent for all other majors, including liberal arts and humanities. Appendix Figure A.2 presents generative AI use outside of work by demographic characteristics, which are generally similar although the differences are less pronounced.

Appendix Table A.1 presents coefficients from a multivariate regression of generative AI adoption at work on demographic characteristics and occupations. We find that the patterns described here generally hold up to a multivariate specification, so we focus on simple bivariate

Figure 2: Share of Working Age Adults Using Generative AI

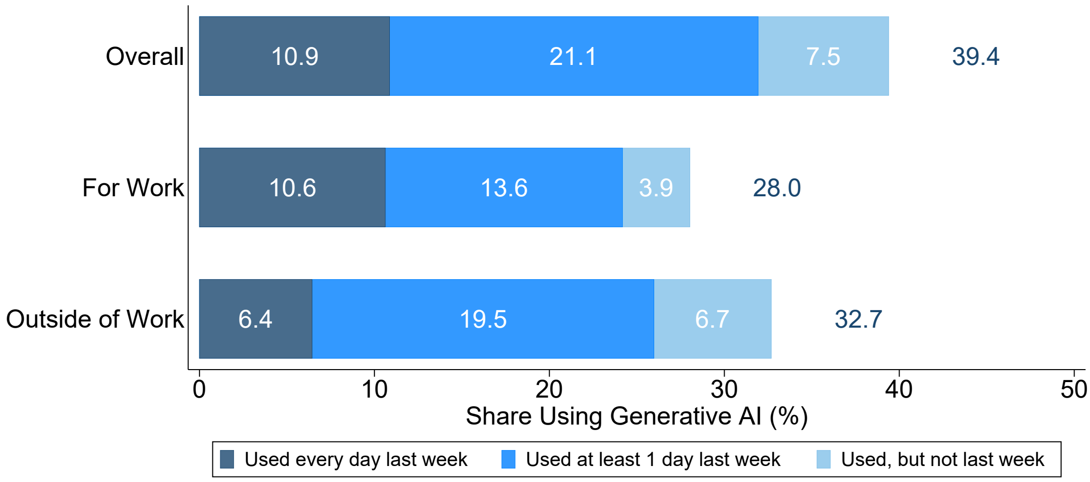

Notes: The figure shows the share of respondents who use AI for work, outside of work, and overall (either for work or outside of work). Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The "For Work" sample is employed individuals ($N = 3216$); the other bars include all respondents ($N = 4682$).

comparisons for simplicity.

3.1 Comparing Generative AI Adoption to PCs and the Internet

Figure 4 compares the speed of adoption of generative AI with two other technologies: PCs and the internet. The horizontal axis measures adoption relative to the first mass-market product. The first mass-market computer was the IBM PC, which was released in August 1981 and sold more than a million units (Abbate, 1999). This implies that our first data point from the CIU is three years after mass adoption. The first Generative AI model to eventually sell at least one million subscriptions was ChatGPT, which was released in November 2022, two years before our survey. Finally, we date mass-market availability of the internet to April 1995, when the National Science Foundation (NSF) decommissioned NSFNet and allowed the internet to carry commercial traffic (Leiner et al., 2009). This was also the year of Netscape's initial public offering (IPO).

We measure combined usage for all three technologies to facilitate comparison, since we cannot separate work and non-work internet usage. The blue dot in Figure 4 repeats the 39.4 percent adoption rate for generative AI reported in Figure 2. The red squares plot personal computer adoption from year 3 to year 22, spanning CPS supplements 1984 to 2003. PC adoption rose steadily from 20 percent in year 3 to nearly 70 percent in year 22. The dark

Figure 3: Demographic Differences in AI Use At Work

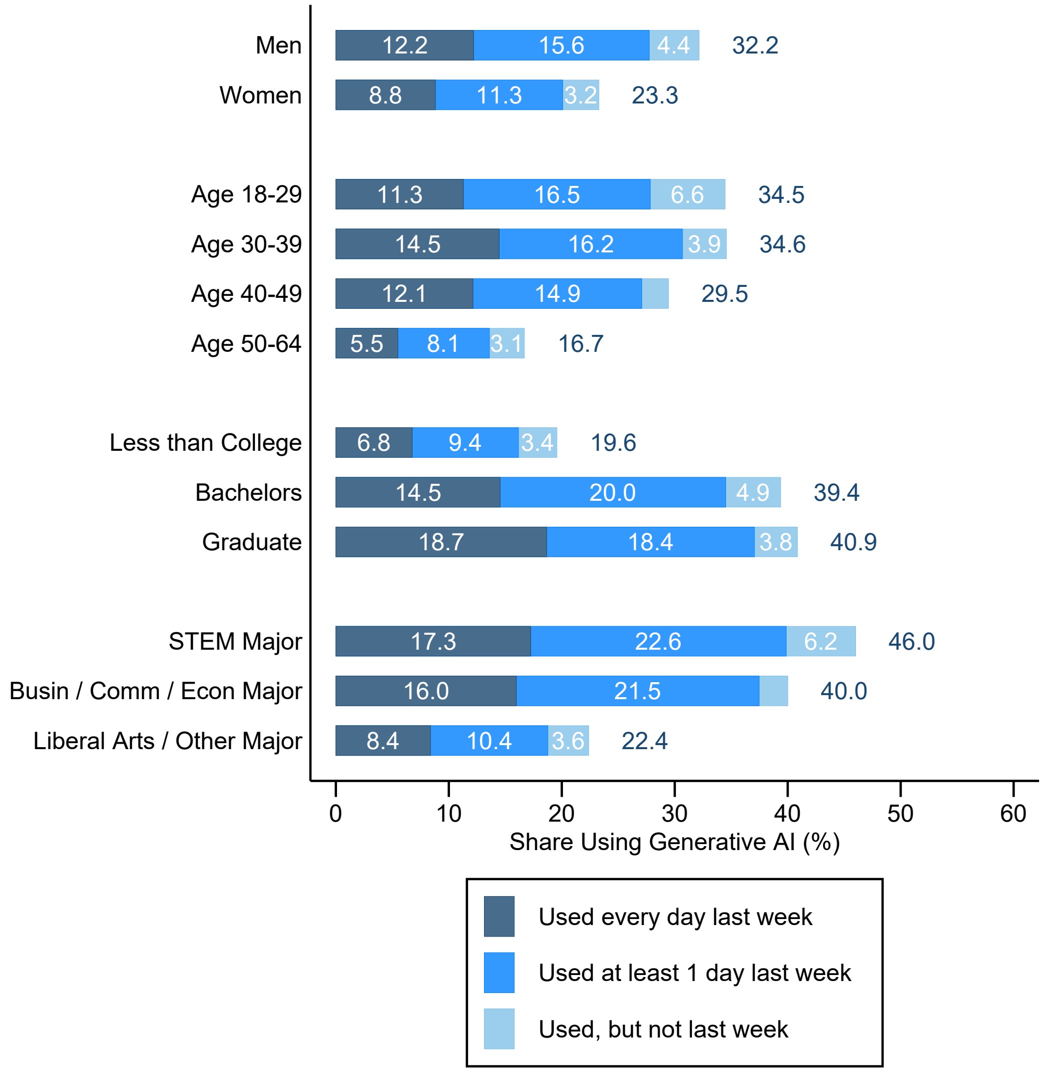

Notes: The figure shows the share of respondents who use AI for work, broken down by gender, age, education, and college major. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 3216$). The sample for college majors is employed individuals with at a bachelor's degree or more. STEM majors include biological, agricultural, environmental, physical, and related sciences; computers, mathematics, and statistics; and engineering. "Busin / Comm / Econ" includes business, communications, and economics majors. "Liberal Arts / Other" includes all other majors.

Figure 4: The Trajectory of Computer, Internet, and AI Adoption

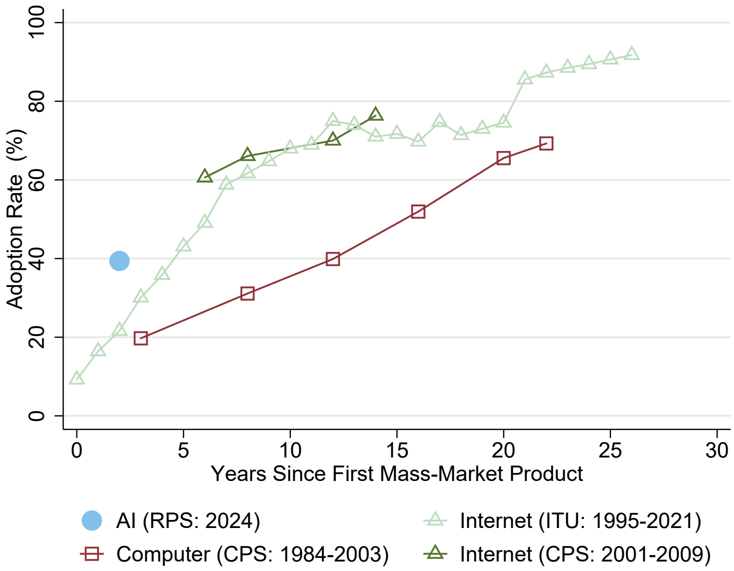

Notes: The figure shows usage rates at work for three technologies: AI, computers, and the internet. The horizontal axis represents years since the introduction of the first mass-market product for each technology. We use 2022 as the introduction year for AI, which was the year ChatGPT was released. We use 1995 as the introduction year for the internet, which was the year that the NSF decommissioned NSFNet and allowed the internet to carry commercial traffic. We use 1981 as the introduction year for computers, which was the year the IBM PC was released. The data source for AI is the August 2024 wave of the RPS (solid blue circle). The data source for computers is the 1984-2003 Computer and Internet Use Supplement of the CPS (hollow red squares). We plot two estimates of internet use: the 2001-2009 Computer and Internet Use Supplement of the CPS (dark green triangles) and the ITU (teal triangles). The sample for the RPS and CPS is all individuals ages 18-64. The RPS sample size is $N = 4682$. The sample for the ITU is individuals of all ages.

green triangles report internet usage from years 6 to 14, spanning the 2001 to 2009 CPS when questions about internet usage were added to the computer supplement. By year 6, internet adoption was already at about 60 percent. The light green triangles plot U.S. internet adoption from the ITU from years 0 to 26 (e.g., 1995 to 2021). Internet adoption increased rapidly from 20 percent in year 2 to 60 percent in year 7, and then increased steadily from 60 percent to 90 percent over the next two decades. The two datasets align closely for the years that overlap.

Figure 4 shows that so far, generative AI has been adopted at a faster pace than PCs or the internet. Faster adoption of generative AI compared with PCs is driven by much greater use outside of work, probably due to differences in portability and cost. Appendix Figure A.3 compares generative AI and PC adoption at work only, using CIU data (we cannot separate internet usage between home and work). We find an adoption rate of 28 percent in year two for generative AI, compared with a 25 percent adoption rate in year three for PCs.

Generative AI and PCs have similar early adoption patterns by education and income. Appendix Figure A.4 shows that about 42 percent of workers with a bachelor's degree or more had used a PC at work three years after mass-market adoption, compared with only 20 percent of workers with less than a college degree. We find nearly identical adoption rates by education for generative AI. Appendix Figure A.5 plots usage of generative AI and PCs by percentile of weekly earnings. We find a very similar pattern for both technologies, with usage increasing in income until about the 85th percentile and then declining slightly.

While adoption patterns are very similar overall for the two technologies, the one exception is gender. Appendix Figure A.6 shows that 32 percent of men use generative AI at work, compared with only 23 percent of women. In contrast, by 1984 only 22 percent of men used a PC at work, compared with 30 percent of women. One possible explanation is the high share of women in office and administrative support occupations, where PC adoption was nearly 50 percent.

3.2 Generative AI Usage Across Jobs and Tasks

Figure 5a presents generative AI adoption by occupation. We elicited respondents' job titles through a free text response and then match them to Standard Occupation Classification (SOC) codes using a parsing algorithm that identifies occupations in 97 percent of cases.[12]

[12]For job titles that do not exactly match a unique SOC code, we use the job title to SOC code matching algorithm developed by the National Institute for Occupational Safety and Health. Laughlin, Song, Wisniewski, and Xu (2024) show that this algorithm works as well as the Occupational Information Network (O*NET) autocoder, which is used widely. Both programs generate probabilistic matches based on the free text response and the respondent's industry, which we elicit in a prior step. When the match is inexact, the survey presents the respondent with the five most likely occupation codes for the text they entered, as well as a "none of the above" option. Respondents who choose "none of the above" are coded as missing; 76 percent of job titles

Figure 5: AI Use At Work by Occupation and Industry Groups

(a) Occupation Groups

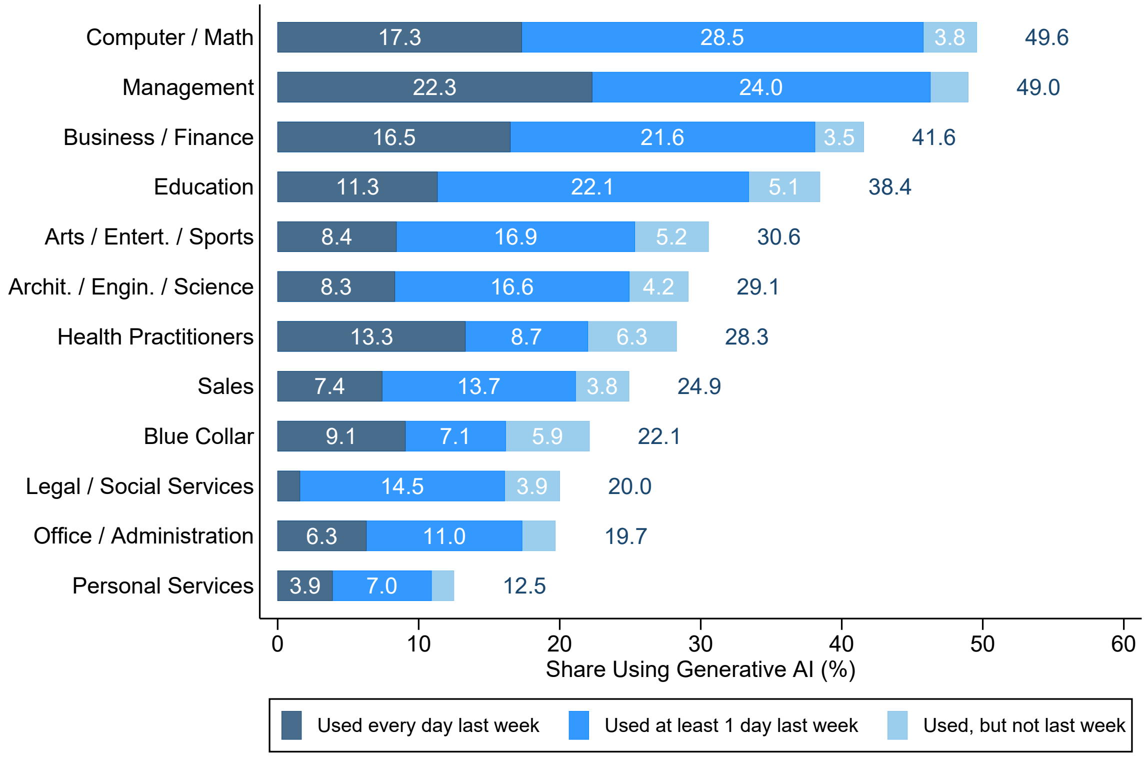

(b) Industry Groups

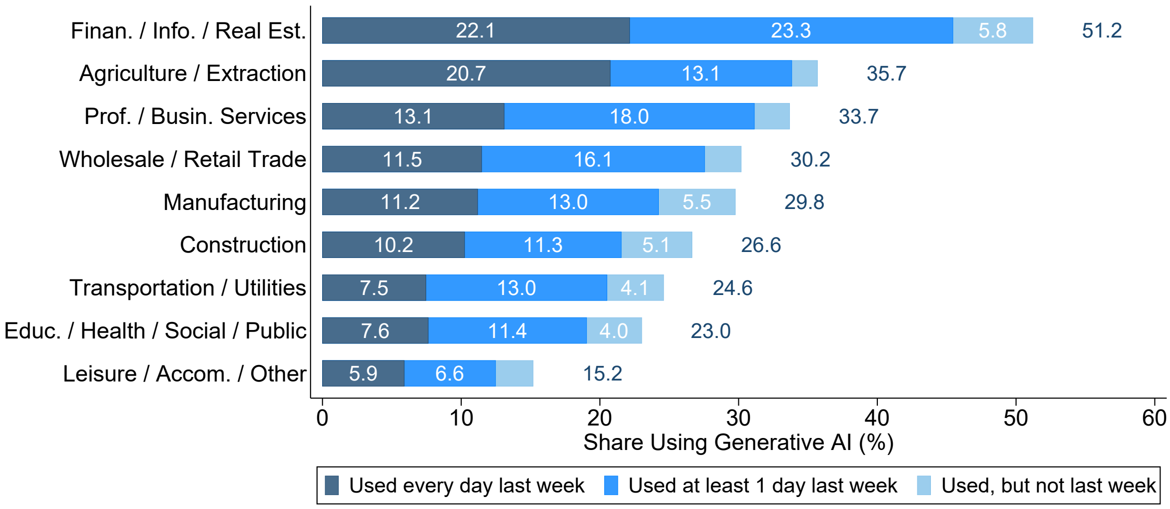

Notes: Figure 5a shows the share of respondents who use AI for work, broken down by occupation. Personal Services occupations combine SOC codes 31-39: Healthcare support, Protective services, Food preparation and serving, Cleaning and maintenance, and Personal care. Blue Collar occupations combine SOC codes 47-53: Construction, Extraction, Installation, Maintenance and Repair, Production, Transportation, and Moving. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 3191$). Figure 5b shows the share of respondents who use AI for work, broken down by industry. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 3216$).

Generative AI adoption at work is highest for computer/mathematical and management occupations, at about 49 percent. Usage at work is also high for business and finance and education occupations (42 and 38 percent, respectively). However, generative AI adoption is relatively common across a range of jobs. With the exception of personal services, at least 20 percent of workers from all major occupations groups use generative AI at work. Interestingly, 22 percent of workers in "blue collar" jobs - construction and extraction, installation and repair, skilled production, and transportation and moving occupations - use generative AI at work.

Appendix Figure A.7 presents generative AI adoption outside of work, by respondent occupation. The ordering is generally similar, although usage rates are higher overall, ranging from 27 percent among workers in personal services to 47 percent for managers. Figure 5b presents results by industry. Generative AI usage is highest for workers in Finance, Insurance, and Real Estate (51 percent) and lowest in Leisure and Accommodation (15 percent). Appendix Figure A.8 compares generative AI and PC adoption rates by occupation. PC usage was much more concentrated in a few occupations, ranging from more than 90 percent in computer and mathematical occupations to less than 10 percent for most blue collar and personal services occupations.

We also ask RPS respondents about the job tasks in which generative AI is most useful. Among respondents who indicated that they had used generative AI in the last week, we present them with the list of tasks in Figure 6 and ask them to select any that they had used generative AI for help with in the last week. They were also allowed to write in other tasks. Respondents were then asked to rank the tasks they selected in order of how helpful generative AI was in completing the task.

Figure 6 reports the share of respondents who ranked each task in the top two in terms of importance. The highest ranked tasks at work were writing (38 percent), administrative tasks (27 percent), and interpreting/translating/summarizing text or data (23 percent). However, the rankings were evenly distributed overall, with eight of the ten tasks in our list ranked in the top two by at least 10 percent of respondents. Outside of work, the highest ranked tasks were writing (27 percent), interpreting/translating/summarizing (23 percent), and personal assistance (21 percent). As with work usage, eight of the eleven tasks were ranked in the top two by at least 10 percent of respondents.

Appendix Figure A.9 reports the share of respondents who reported using generative AI for each task. The ordering is generally similar to Figure 6. Usage rates at work were at least 25 percent for all ten tasks in our list, with the most frequent tasks being writing (57 percent), searching for information (49 percent), and obtaining detailed instructions (48 percent). Outside of work, the most common tasks were personal assistance (lists, schedules, etc.), ideas and

match exactly, 21 percent are chosen from the list, and 3 percent are missing.

Figure 6: In Which Specific Tasks Is AI Most Useful?

(a) At Work

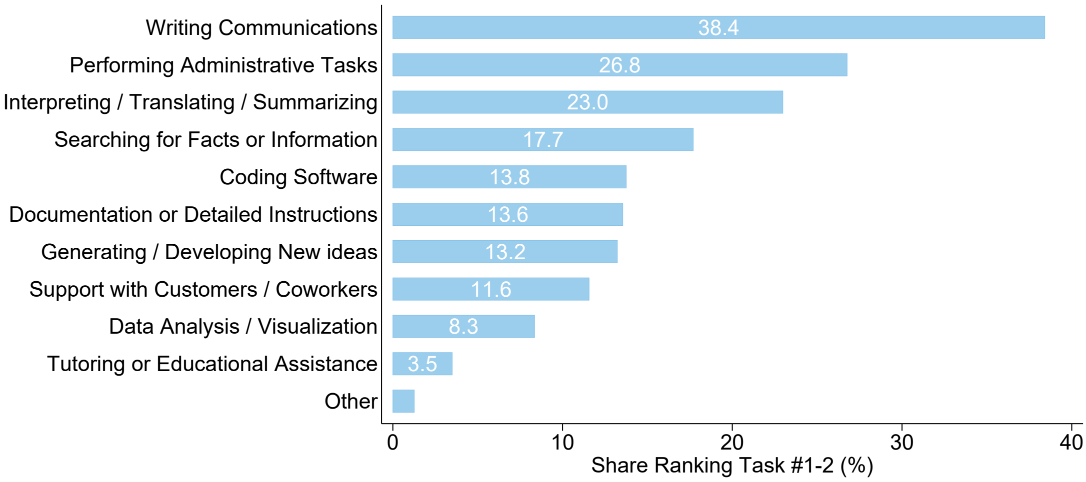

(b) Outside of Work

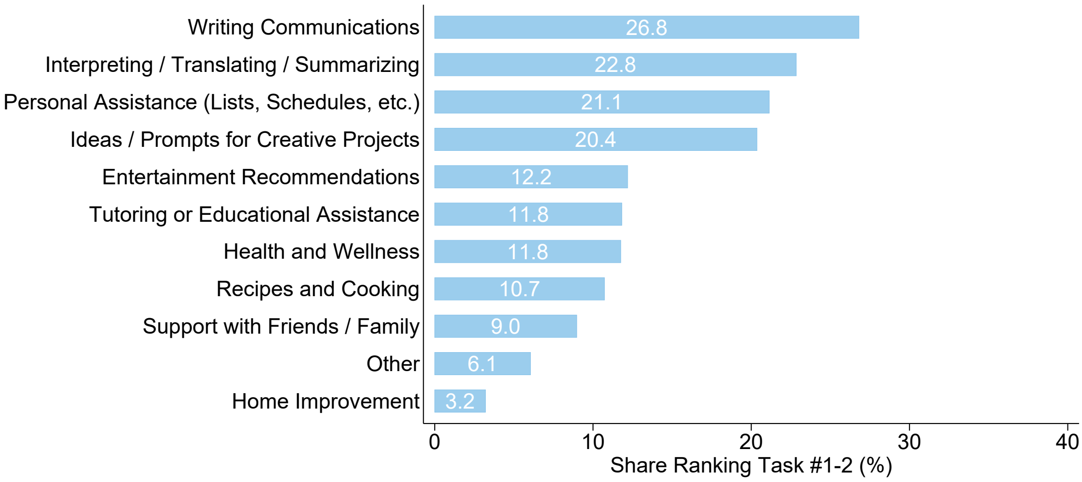

Notes: The figure shows which tasks AI users report that AI is most helpful in completing. Panel (a) refers to tasks at work; panel (b) refers to tasks outside of work. Separately for work and outside of work, respondents were first provided with a list of tasks and asked to select those that they had used AI to help with last week. Respondents were then asked to rank these selected tasks according to how helpful AI was in completing the task. The figure reports the share of AI users who ranked a particular task either #1 (AI was most helpful in this task) or #2. The bars do not have a natural sum because some respondents selected fewer than two tasks. Data source is the August 2024 wave of the RPS, ages 18-64. The samples for panels (a) and (b) are employed individuals ($N = 3216$) and all individuals ($N = 4682$), respectively.

prompts for creative projects, and writing. Overall, generative AI use is strikingly broad across occupations and job tasks.

3.3 How Much Could Generative AI Increase Labor Productivity?

In addition to asking respondents about the frequency of generative AI usage, we also ask about intensity of use within a day. Specifically, we asked respondents whether they used generative AI for 15 minutes or less, between 15 minutes and an hour, or more than an hour on the days that they used it at work and at home. Appendix Figure A.10 reports the intensity of daily usage of generative AI both at work and at home. Overall, 25 percent of generative AI users reported using it for an hour or more at work, and 52 percent used it for between 15 and 60 minutes on the days that they used it. Intensity and frequency of usage are positively correlated, with 42 percent of daily users reporting an hour or more of usage per day. The patterns are similar for usage outside of work.

We use these measures of generative AI intensity and frequency to estimate the share of all work hours assisted by generative AI. For each individual we multiply their reported hours worked by the lower and upper bounds of their self-reported generative AI usage frequency and intensity. For example, consider a worker who reports using generative AI on some but not all workdays last week, and for between 15 and 60 minutes per day on days that they used it. For a lower bound we assume the respondent used generative AI on exactly one day for 15 minutes. The upper bound in these cases would be using generative AI on all but one of the days they worked and 59 minutes each day. We assume zero hours for respondents who reported using generative AI but not in the last week.

Summing these figures together for all respondents yields lower and upper bounds of 0.5 and 3.5 percent of all work hours per week assisted by generative AI. The small overall share is mostly driven by the extensive margin, with 76 percent of workers reporting zero hours of generative AI usage.

Given these figures, how much could generative AI plausibly increase labor productivity? Five recent studies have estimated the impact of generative AI on task productivity using experimental or quasi-experimental designs. Noy and Zhang (2023) pay college-educated professionals to perform writing tasks and find that ChatGPT improved productivity by 40 percent. Brynjolfsson, Li, and Raymond (2023) find that a generative AI-based conversational assistant increases productivity by 14 percent. Dell'Acqua et al. (2023) find that providing generative AI access to strategy consultants increases productivity by about 25 percent in a pre-registered set of tasks they deemed fit for AI assistance. Cui et al. (2024) find that randomly providing GitHub copilot, an AI-based coding assistant, to software developers at a large electronics manufacturing company increases productivity by 26 percent. Peng, Kalliamvakou, Cihon, and Demirer (2023) also study the impact of GitHub copilot on a single coding task and find that it increases productivity by 56 percent.

If we (somewhat arbitrarily) multiply the median estimate of about 25 percent by the share of work hours assisted by generative AI, we estimate that it could plausibly increase labor productivity by between 0.125 and 0.875 percent at current levels of usage. If generative AI adoption continues at the rate of past technologies such as PCs and the internet, and there were no change in intensity or task composition, these numbers would double within the next decade. However, we caution that this calculation is extremely speculative. Workers and firms are probably using generative AI first in its most-productive applications, suggesting diminishing returns as usage expands. On the other hand, the technology may become better and more widely applicable over time.

4 Conclusion

Generative AI has rapidly emerged as an important new technology, yet its impact on the economy depends critically on the speed and intensity of adoption. This paper reports results from the first nationally representative U.S. survey of generative AI use at work and at home. Our data come from the Real-Time Population Survey (RPS), a survey that is constructed and weighted to be nationally representative and follows the same survey design as the CPS, a widely used national data source. We find that 39.4 percent of the U.S. population age 18-64 reported using generative AI during August of 2024, with 28.0 percent using at work and nearly one in nine workers reporting daily usage.

We compare the speed and intensity of generative AI adoption with two transformative technologies - PCs and the internet. We find that generative AI has been adopted more rapidly than both technologies. We also find that generative AI is used by workers in a broad range of occupations and job tasks. Nearly half of workers in computer and mathematical and management occupations use generative AI, but so do nearly one in four blue-collar workers. We asked respondents about whether generative AI was useful in eleven different job tasks such as writing, administrative support, interpreting and summarizing text or data, and coding. Usage rates at work exceeded 25 percent for all ten tasks in our list. Overall, we find strong support for the idea that generative AI truly is a general-purpose technology Eloundou, Manning, Mishkin, and Rock (2024).

We estimate that between 0.5 and 3.5 percent of all work hours in the U.S. are currently assisted by generative AI. Assuming that the productivity gains from recent experimental studies are externally valid, this suggests that generative AI could plausibly increase labor productivity by between 0.125 and 0.875 percentage points at current levels of usage, although we caution that this calculation should be considered highly speculative given the assumptions it requires.

Our findings suggest many directions for future work. In particular, it will be important to track the adoption of generative AI as the technology matures and to monitor whether its usage expands broadly across workers, firms, and occupations. Future RPS surveys will include more detailed questions about the frequency and intensity of generative AI adoption so that we can track its evolving impact on the U.S. economy.

References

Abbate, J. (1999). "Getting small: a short history of the personal computer". In: *Proceedings of the IEEE* 87.9, pp. 1695–1698.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

Autor, D. (2024). *Applying AI to rebuild middle class jobs*. Tech. rep. National Bureau of Economic Research.

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.

Beaudry, P., M. Doms, and E. Lewis (2010). "Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas". In: *Journal of political Economy* 118.5, pp. 988–1036.

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, and K. Mertens (2023). "Work from home before and after the COVID-19 outbreak". In: *American Economic Journal: Macroeconomics* 15.4, pp. 1–39.

Bick, A., A. Blandin, K. Mertens, and H. Rubinton (2024). "Work from Home and Interstate Migration". In: *Bick, A., Blandin, A., Mertens, K., Rubinton, H.*

Bloom, D. E., K. Prettner, J. Saadaoui, and M. Veruete (2024). *Artificial intelligence and the skill premium*. Tech. rep. National Bureau of Economic Research.

Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, and K. Savage (2024). *Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey*. Tech. rep. National Bureau of Economic Research.

Brynjolfsson, E., D. Li, and L. R. Raymond (2023). *Generative AI at work*. Tech. rep. National Bureau of Economic Research.

Comin, D. and B. Hobijn (2010). "An exploration of technology diffusion". In: *American economic review* 100.5, pp. 2031–2059.

Comin, D. and M. Mestieri (2018). "If technology has arrived everywhere, why has income diverged?" In: *American Economic Journal: Macroeconomics* 10.3, pp. 137–178.

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.

Fletcher, R. and R. Nielsen (2024). "What does the public in six countries think of generative AI in news?" In.

Humlum, A. and B. B. Meyer (2022). *Artificial intelligence and college majors*. Rockwool Foundation Research Unit.

Humlum, A. and E. Vestergaard (2024). "The Adoption of ChatGPT". In: *University of Chicago, Becker Friedman Institute for Economics Working Paper* 2024-50.

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.

Leiner, B. M., V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff (2009). "A brief history of the Internet". In: *ACM SIGCOMM computer communication review* 39.5, pp. 22–31.

Liu, Y. and H. Wang (2024). "Who on Earth Is Using Generative AI?" In.

McClain, C. (2024). "Americans' use of ChatGPT is ticking up, but few trust its election information". In.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas (2024). "AI adoption in America: Who, what, and where". In: *Journal of Economics & Management Strategy* 33.2, pp. 375–415.

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

Peña-López, I. et al. (2009). "Manual for measuring ICT access and use by households and individuals". In.

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*.

Pilossoph, L. and J. Ryngaert (2023). *Job Search, Wages, and Inflation*. Working Paper. Duke University and University of Notre Dame.

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

The Rapid Adoption of Generative AI

Alexander Bick Adam Blandin David Deming

ONLINE APPENDIX

A Additional Results on Generative AI Use: August 2024

Table A.1: Predictors of Generative AI Use at Work

Constant	0.2074***
	(0.0300)
Female	−0.0845***
	(0.0183)
Age: 18-29	0.0152
	(0.0289)
Age: 40-49	−0.0622**
	(0.0271)
Age: 50-64	−0.1880***
	(0.0237)
Educ: Bachelor's Degree	0.1621***
	(0.0233)
Educ: Graduate Degree	0.1701***
	(0.0288)
Management	0.2724***
	(0.0341)
Business / Finance	0.2103***
	(0.0441)
Computer / Math	0.2577***
	(0.0451)
Archit. / Engin. / Science	0.0591
	(0.0515)
Legal / Social Services	0.0410
	(0.0489)
Education	0.1808***
	(0.0425)
Arts / Entert. / Sports	0.1087*
	(0.0556)
Health Practitioners	0.1172**
	(0.0497)
Sales	0.0871**
	(0.0368)
Office / Administration	0.0770**
	(0.0338)
Blue Collar	0.0606*
	(0.0321)
R^2-adj	0.13
N	3191

Notes: Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this table is employed individuals.

Figure A.1: Share Using Specific AI Products

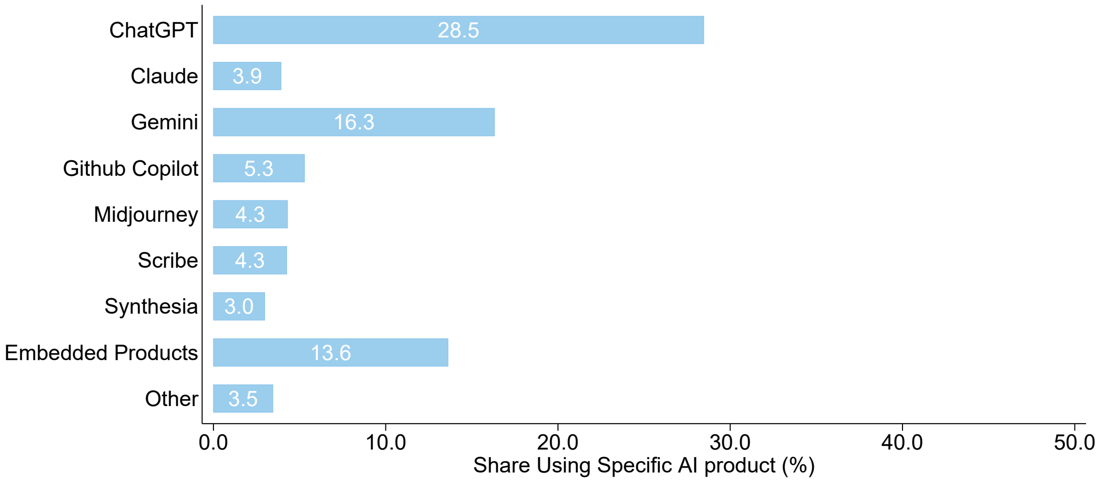

Notes: The figure shows the share of respondents who report using particular AI products. "Embedded products" are AI features embedded within existing software, such as Microsoft Copilot. Data source is the August 2024 wave of the RPS, ages 18-64 ($N = 4682$). Individuals who report using multiple AI products are reflected in multiple bars.

FIGURE A.2: Demographic Differences in AI Use Outside of Work

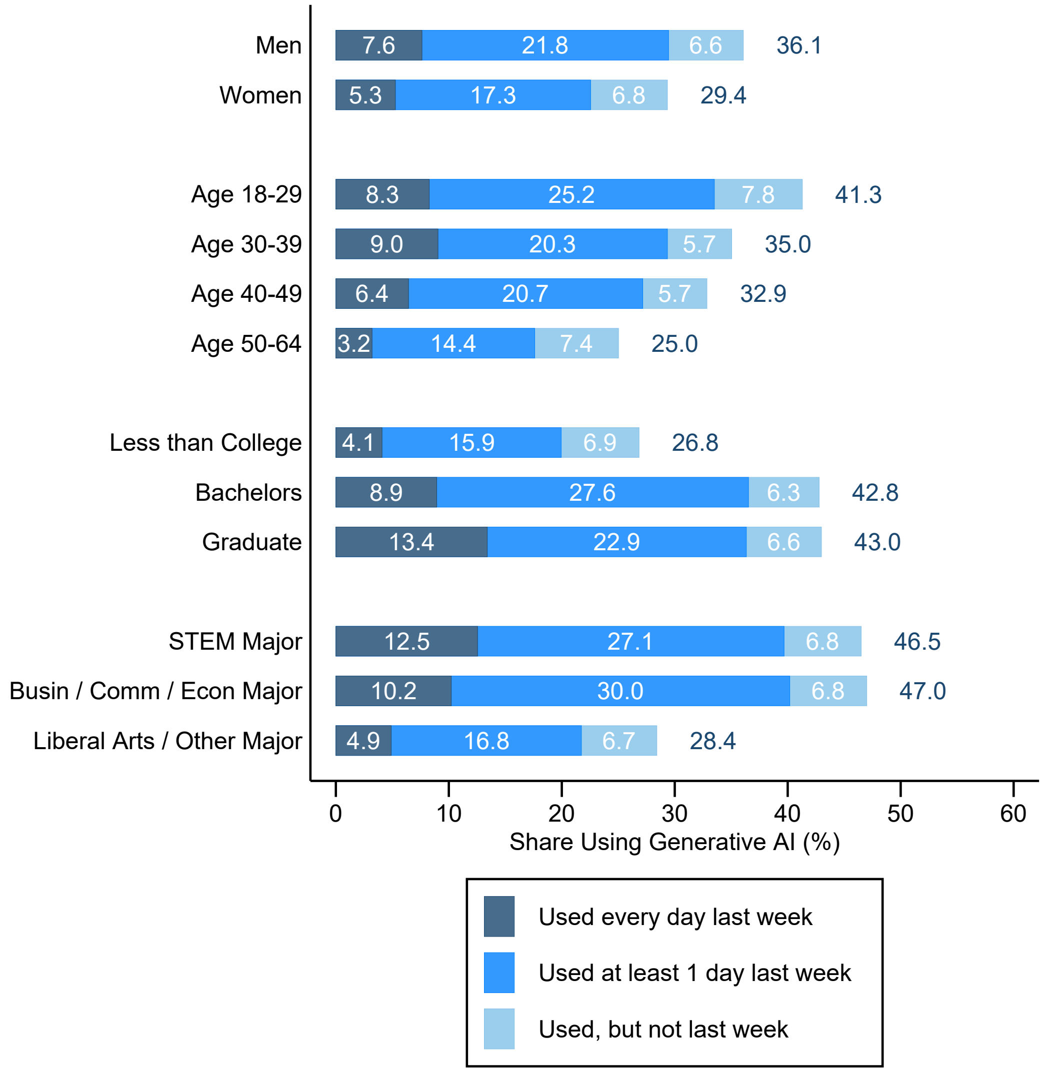

Notes: The figure shows the share of respondents who use AI outside of work, broken down by gender, age, education, and college major. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this figure is all individuals ($N = 4682$). The sample for college majors is individuals with at a bachelor's degree or more. STEM majors include biological, agricultural, environmental, physical, and related sciences; computers, mathematics, and statistics; and engineering. "Busin / Comm / Econ" includes business, communications, and economics majors. "Liberal Arts / Other" includes all other majors.

Figure A.3: The Trajectory of AI and Computer Adoption At Work

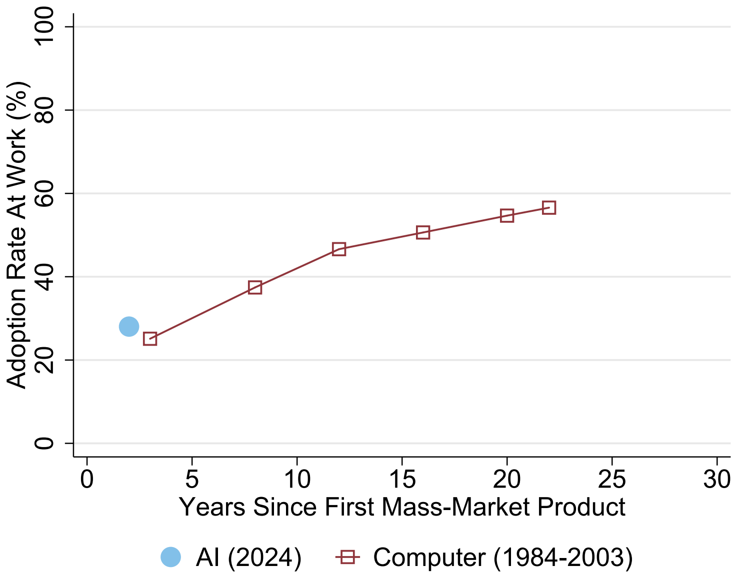

Notes: The figure shows usage rates at work for two technologies: AI and computers. The horizontal axis represents years since the introduction of the first mass-market product for each technology. We use 2022 as the introduction year for AI, which was the year ChatGPT was released. We use 1981 as the introduction year for computers, which was the year the IBM PC was released. The data source for AI is the August 2024 wave of the RPS (solid blue circle). The data source for computers is the 1984-2003 Computer and Internet Use Supplement of the CPS (hollow red squares). The sample for each dataset is employed individuals ages 18-64. The RPS sample size is $N = 3216$.

Figure A.4: AI and Computer Use At Work By Education

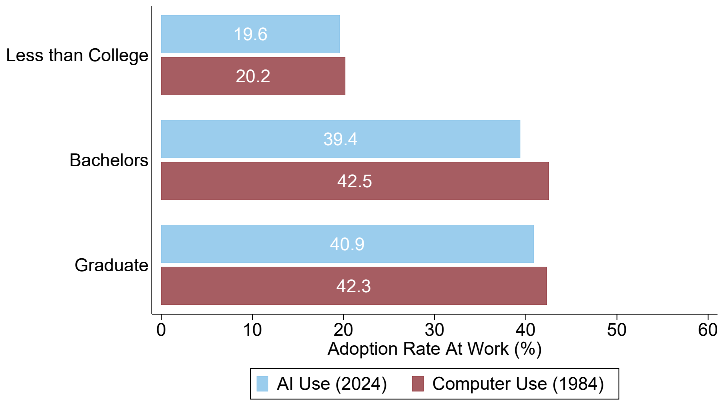

Notes: The figure shows usage rates at work by education for two technologies: AI and computers. The data source for AI is the August 2024 wave of the RPS (blue bars). The data source for computers is the 1984 Computer and Internet Use Supplement of the CPS (red bars). The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 3216$; CPS, $N = 61708$).

Figure A.5: AI and Computer Use At Work Across the Earnings Distribution

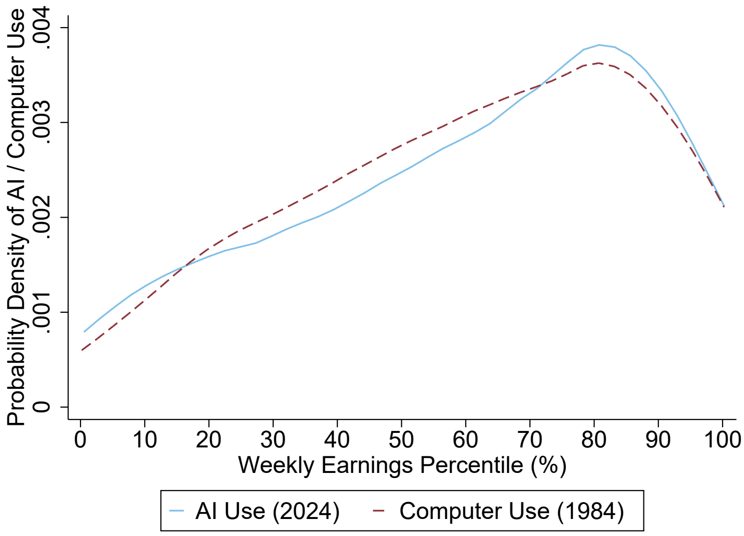

Notes: The figure shows usage rates at work by earnings for two technologies: AI and computers. The horizontal axis corresponds to the percentile of the weekly earnings distribution. The data source for AI is the August 2024 wave of the RPS (blue lines). The data source for computers is respondents in the 1984 Computer and Internet Use Supplement of the CPS who also were a part of the Outgoing Rotation Group (red lines). The sample for each dataset is employed individuals ages 18-64 whose implied hourly wage is above the minimum wage of that year and, to avoid top-coding, whose earnings are below the 99^{th} percentile. The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 2809$; CPS, $N = 13284$).

Figure A.6: AI and Computer Use At Work By Gender

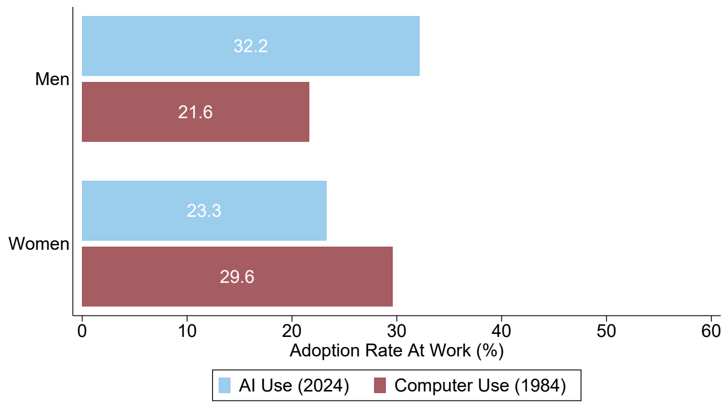

Notes: The figure shows usage rates at work by gender for two technologies: AI and computers. The data source for AI is the August 2024 wave of the RPS (blue bars). The data source for computers is the 1984 Computer and Internet Use Supplement of the CPS (red bars). The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 3216$; CPS, $N = 61708$).

Figure A.7: AI Use Outside of Work by Occupation Groups

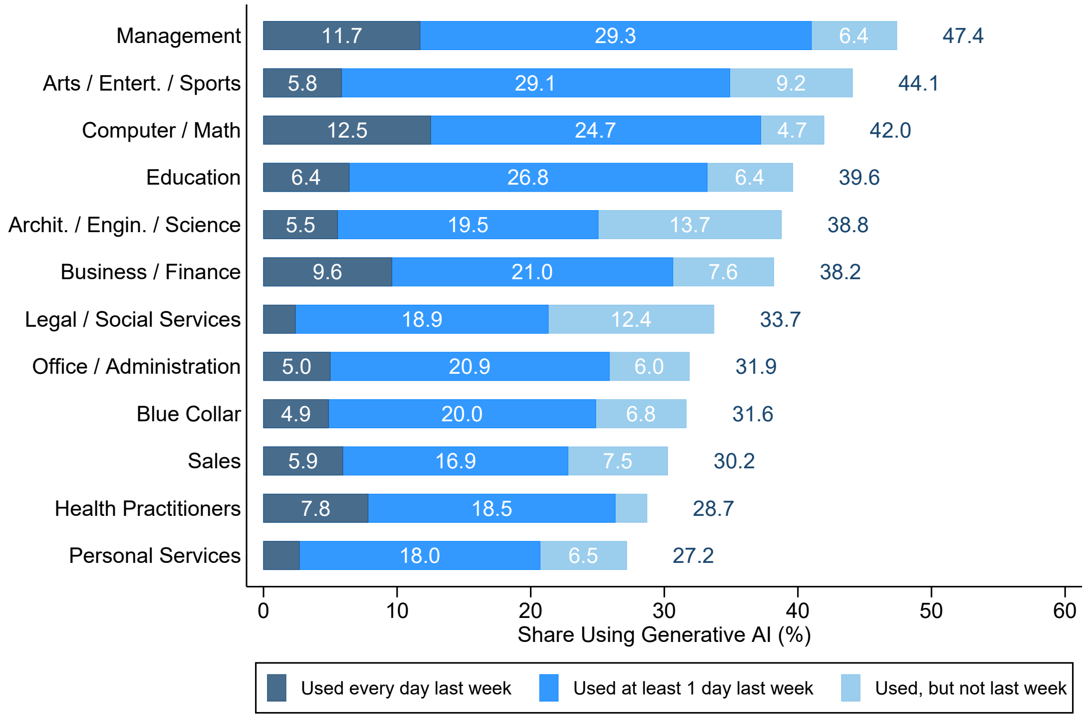

Notes: The figure shows the share of respondents who use AI outside of work, broken down by occupation. Personal Services occupations combine SOC codes 31-39: Healthcare support, Protective services, Food preparation and serving, Cleaning and maintenance, and Personal care. Blue Collar occupations combine SOC codes 47-53: Construction, Extraction, Installation, Maintenance and Repair, Production, Transportation, and Moving. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the August 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 3191$).

Figure A.8: AI and Computer Use At Work By Occupation

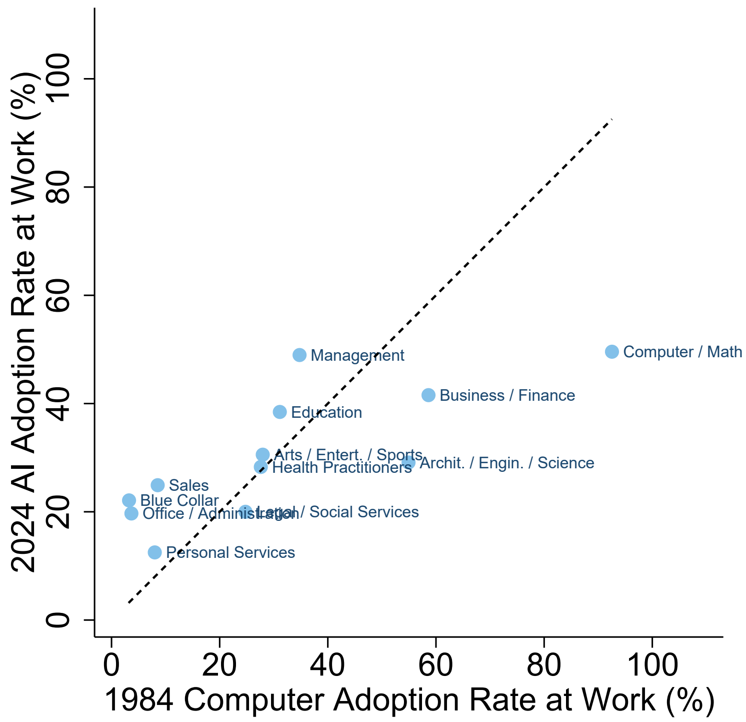

Notes: The figure shows usage rates at work by occupation for two technologies: AI and computers. The data source for AI is the August 2024 wave of the RPS (vertical axis). The data source for computers is the 1984 Computer and Internet Use Supplement of the CPS (horizontal axis). The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 3191$; CPS, $N = 60929$).

Figure A.9: Which Tasks Do People Use AI For?

(a) At Work

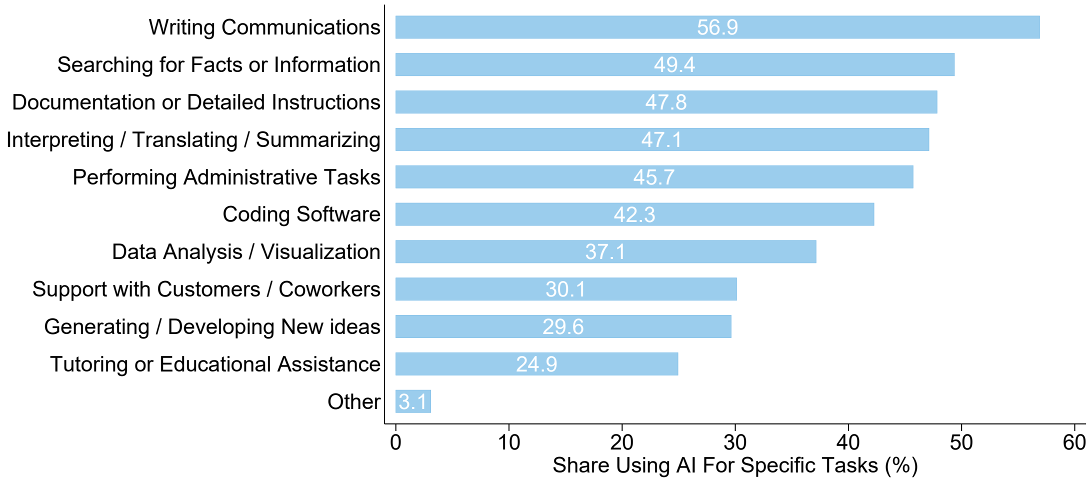

(b) Outside of Work

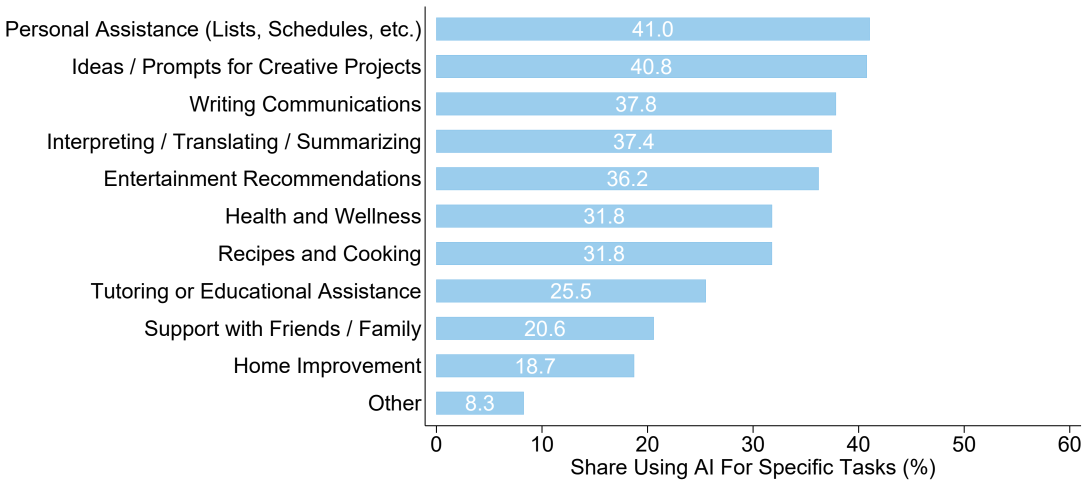

Notes: The figure shows the share of AI users that used AI for specific tasks. Panel (a) refers to tasks at work; panel (b) refers to tasks outside of work. The bars do not have a natural sum because respondents could select multiple tasks. Data source is the August 2024 wave of the RPS, ages 18-64. The samples for panels (a) and (b) are employed individuals ($N = 3216$) and all individuals ($N = 4682$), respectively.

Figure A.10: Daily Time Spent Using AI

(a) At Work

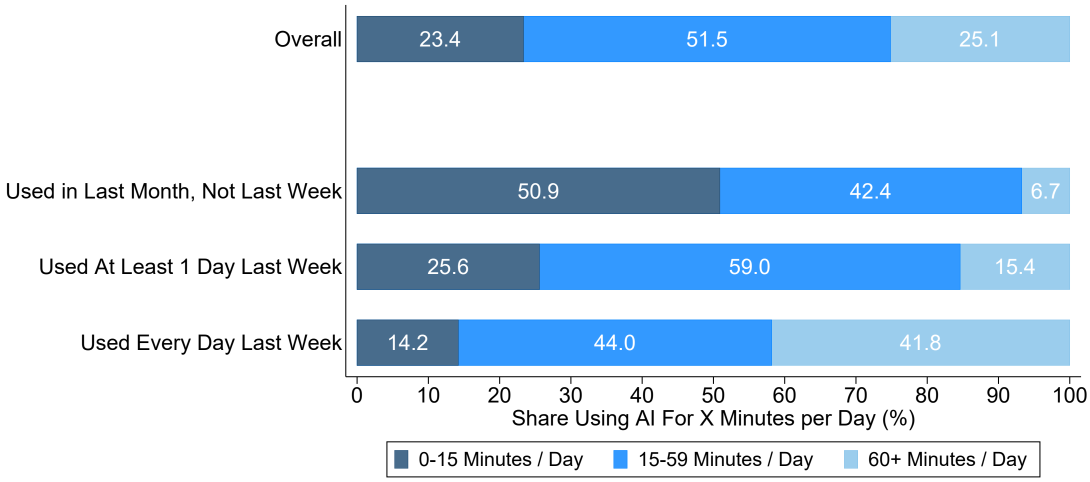

(b) Outside of Work

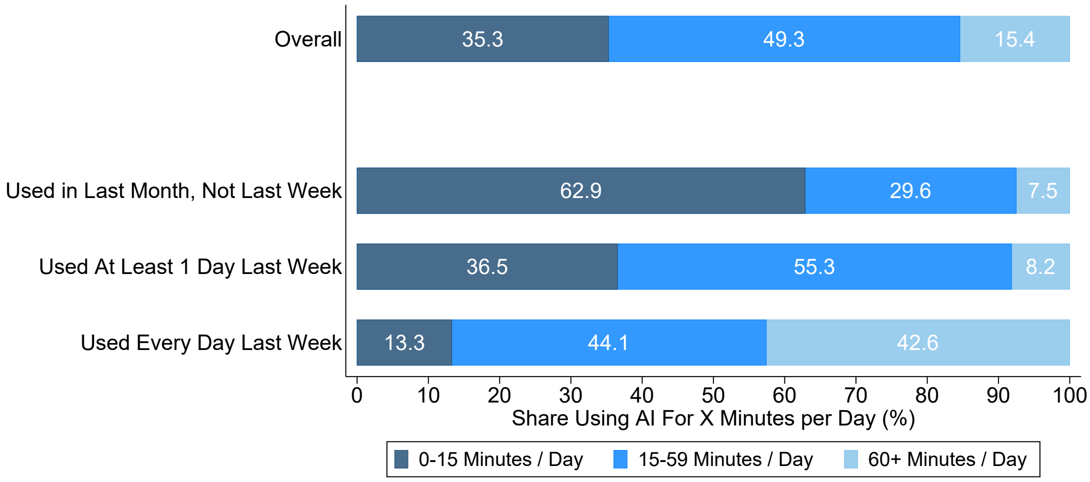

Notes: The figure shows the distribution of daily time spent actively using AI, among AI users. Panel (a) refers to AI use at work; panel (b) refers to AI use outside of work. Usage time is broken down into 0-14 minutes per day (dark blue), 15-59 minutes per day (medium blue), and 60 or more minutes per day (light blue). The "Overall" bar reflects the distribution among all AI users. The "Used in Last Month, Not Last Week" bar reflects users who did not use AI at work last week but did use it within the last four weeks. The "Used at least 1 day last week" bar reflects users who used AI at least one day last week but not every workday. The "used every day last week" bar reflects users who used AI for work every workday last week. Data source is the August 2024 wave of the RPS, ages 18-64. The samples for panels (a) and (b) are employed individuals who use AI ($N = 984$) and all individuals who use AI ($N = 1444$), respectively.

B Results on Generative AI Use: June 2024

Figure B.1: Share of Working Age Adults Using Generative AI

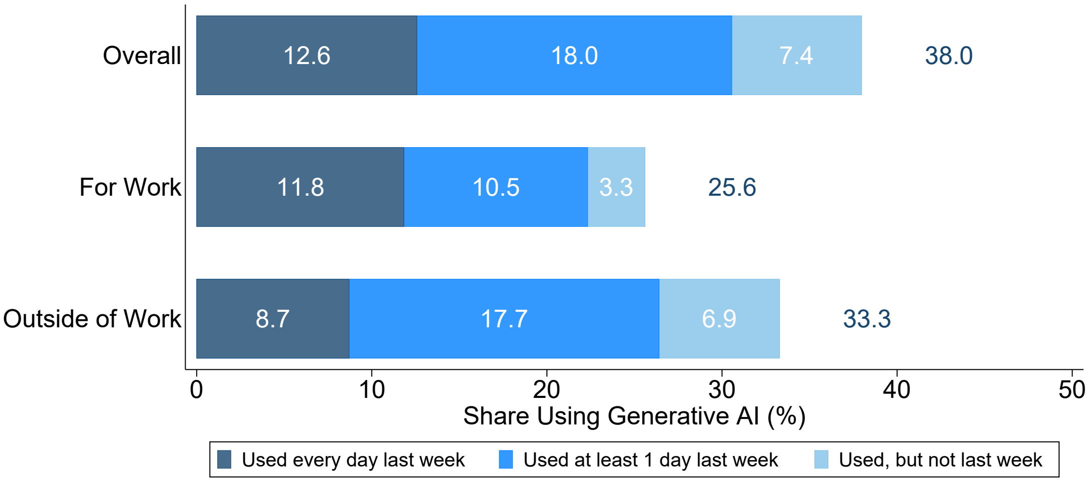

Notes: The figure shows the share of respondents who use AI for work, outside of work, and overall (either for work or outside of work). Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the June 2024 wave of the RPS, ages 18-64. The "For Work" sample is employed individuals ($N = 1576$); the other bars include all respondents ($N = 2354$).

Figure B.2: Demographic Differences in AI Use At Work

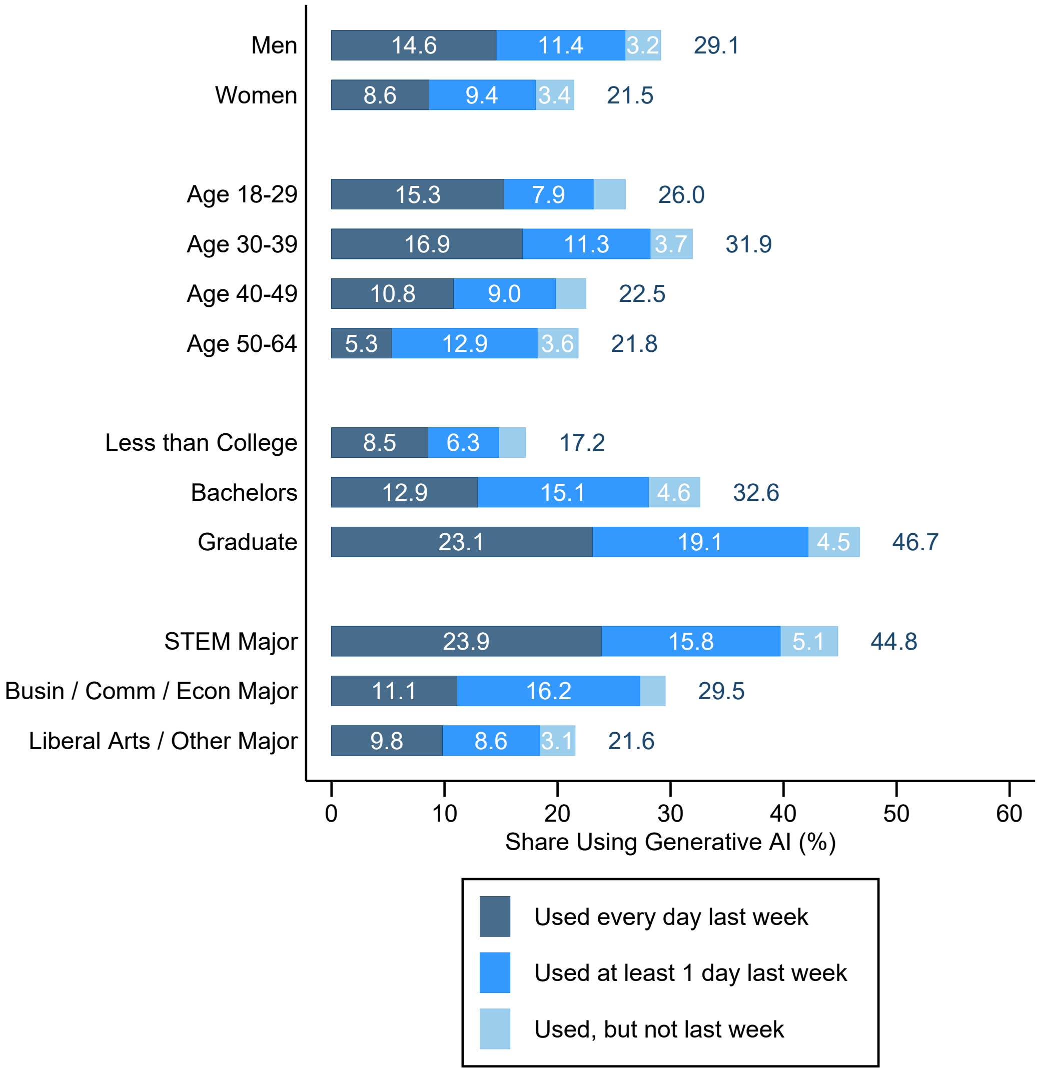

Notes: The figure shows the share of respondents who use AI for work, broken down by gender, age, education, and college major. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the June 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 1576$). The sample for college majors is employed individuals with at a bachelor's degree or more. STEM majors include biological, agricultural, environmental, physical, and related sciences; computers, mathematics, and statistics; and engineering. "Busin / Comm / Econ" includes business, communications, and economics majors. "Liberal Arts / Other" includes all other majors.

Figure B.3: AI Use At Work by Occupation Groups

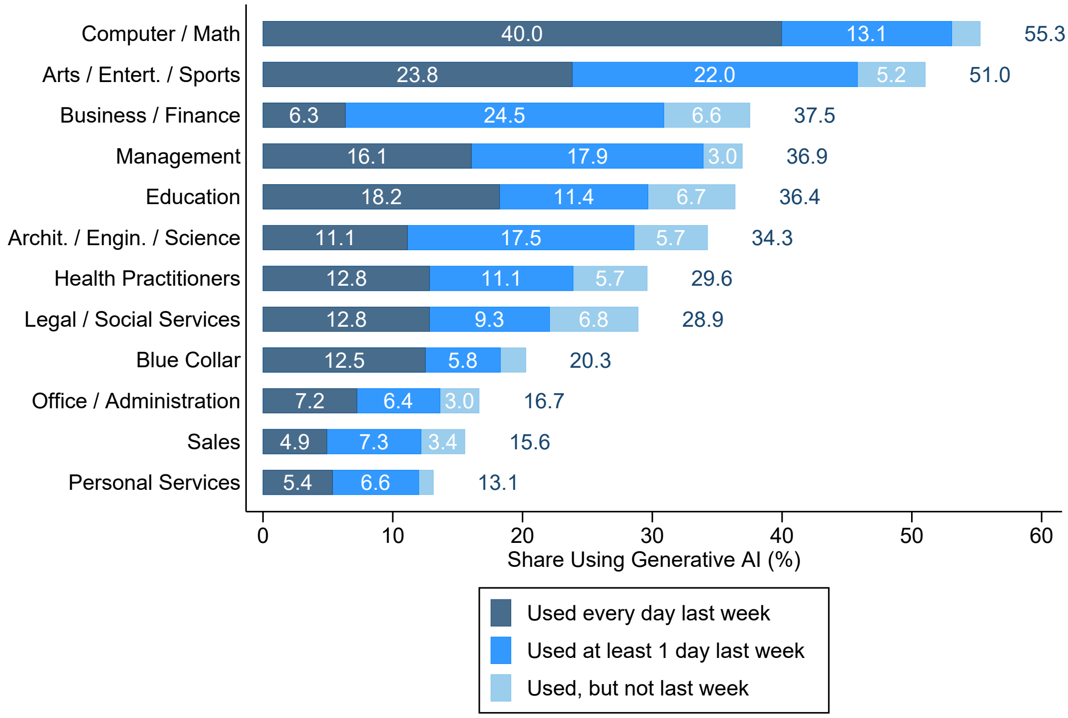

Notes: The figure shows the share of respondents who use AI for work, broken down by occupation. Personal Services occupations combine SOC codes 31-39: Healthcare support, Protective services, Food preparation and serving, Cleaning and maintenance, and Personal care. Blue Collar occupations combine SOC codes 47-53: Construction, Extraction, Installation, Maintenance and Repair, Production, Transportation, and Moving. Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source is the June 2024 wave of the RPS, ages 18-64. The sample for this figure is employed individuals ($N = 1570$).

C RPS: Measurement and Definitions

occupation in our weighting scheme. This requires us to drop another 108 and 112 observations due to missing occupation codes for the June and August surveys, respectively.

C.1 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. Columns 1 and 2 of Table 1 compare the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure for our main survey in August 2024. As discussed in the paper, we also compare at the bottom of the table employment status in the CPS and RPS, as well as the decompgraphic composition for the employed. None of this latter sets of moments were targeted during the sampling of survey respondents.

Using the iterative proportional fitting (raking) algorithm of Deming and Stephan (1940), we construct sampling weights to ensure the RPS matches the CPS sample proportions for the same set of demographic characteristics included in the Qualtrics sampling targets for the overall sample, i.e., independent of employment status. However, we use more disaggregated categories for education and marital status, 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).

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 gender, age (18-24, 25-34, 35-44, 45-54, 55-64), race and ethnicity (non-Hispanic White, non-Hispanic Black, Hispanic, all other racial and ethnic groups), education (high school or less, some college or associate degree, bachelor's degree or more), marital status (married + spouse present, never married, other), relationship status (spouse living in the same household, partner living in the same household, other), presence of children in the household (yes or no), household income over the last 12 months (<\$50k, \$50k-100k, >\$100k), and region (Midwest, Northeast, South, and West using the Census definition). These groupings were chosen to ensure that each cell size is at least 30. We also include 2-digit occupation codes in our weighting scheme. Among the 22 occupations, we merge several occupations to ensure a sufficient sample size. In particular, for the June survey, we merge a) "Architecture and Engineering Occupations" and "Life, Physical, and Social Science Occupations", b) "Community and Social Service Occupations" and "Legal Occupations", c) "Healthcare Support Occupations" and "Protective Service Occupations", and d) "Farming, Fishing, and Forestry Occupations" and "Construction and Extraction Occupations." For the August survey, we proceeded similarly but, due to the larger sample size, did not need to merge "Architecture and Engineering Occupations" and "Life, Physical, and Social Science Occupations."

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 C.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). The lack of noticeable deviations from the 45-degree line demonstrates how well the weighting procedure works.

C.2 Validation Checks for the August 2024

Figure C.2a shows that the correlation in industry shares between the unweighted RPS and the CPS is 0.88. Information Services has the highest share relative to the CPS, while Health Care and Social Assistance has the lowest share relative to the CPS. In the weighted data, the RPS and CPS are, on average, even more closely aligned, as shown in Figure C.2a.

Figure C.2c shows that the correlation in college major shares among college graduates between the unweighted RPS and the ACS is 0.8. 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 agreement between the RPS and ACS.

Figure C.1: Sample Composition in the Weighted August 2024 RPS vs. CPS

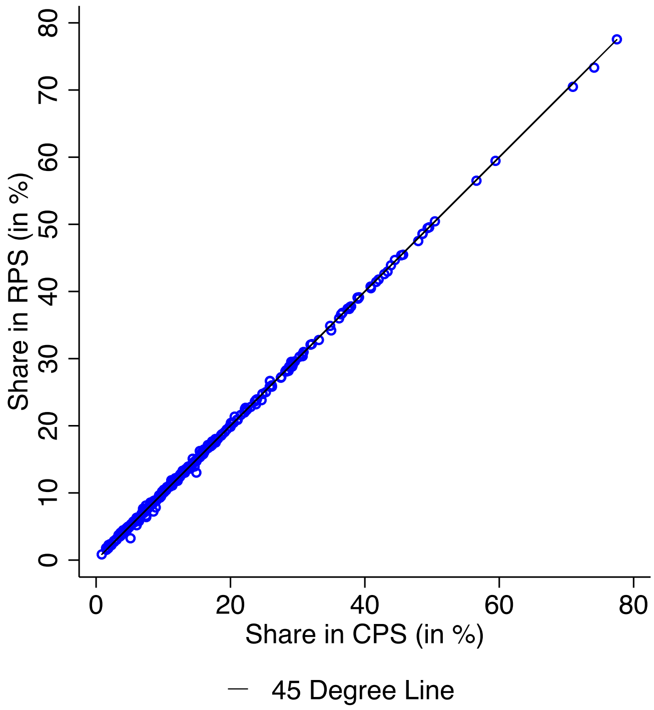

Notes: The figure shows for all statistics used in the weighting scheme, the weighted fraction of individuals with the respective characteristics in the RPS (on the y-axis) and the CPS (on the x-axis).

Figure C.2: Validation Check: Industry and College Major for August 2024

(a) Industry Shares: Unweighted

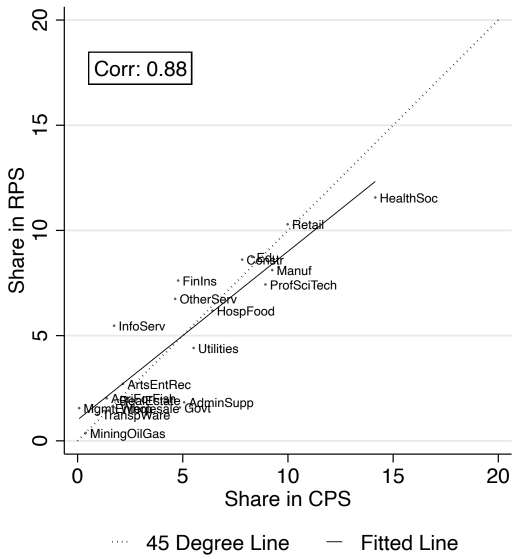

(b) Industry Shares: Weighted

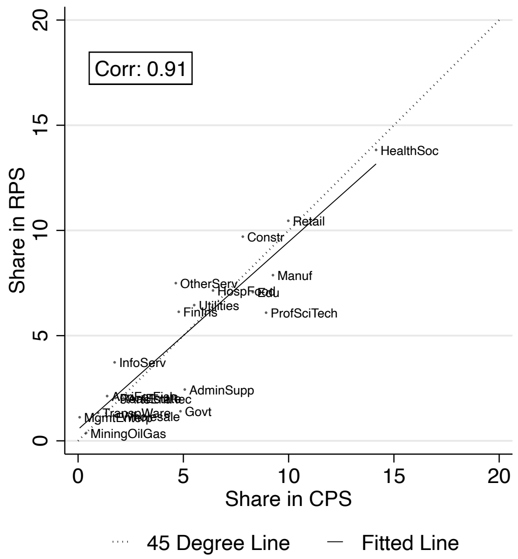

(c) College Major Shares: Unweighted

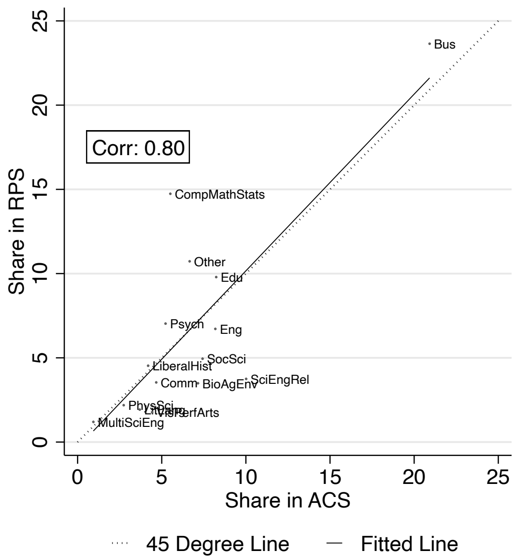

(d) College Major Shares: Weighted

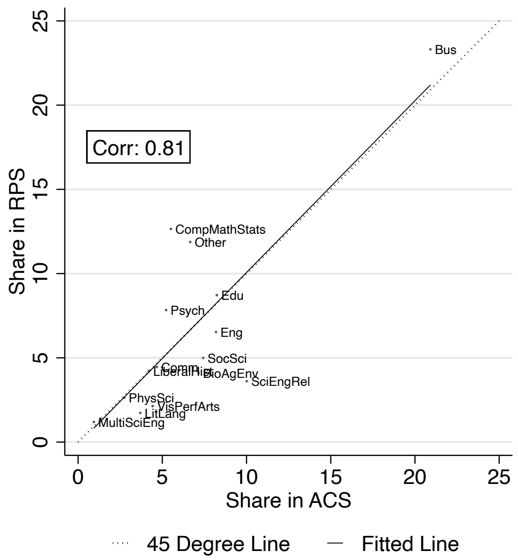

Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the sample of RPS respondents in both figures. All figures use weighted CPS and ACS data, respectively. For the industry comparison, sample sizes for the August 2024 RPS and CPS are 3216 and 45028, respectively. For the major comparison, sample sizes for the August 2024 RPS and the 2022 ACS are 1920 and 682280, respectively.

C.3 Weighting and Validation Checks for the June 2024

Table C.1 compares the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure for our pilot survey in June 2024 for the overall sample (columns 1 and 2), and also provides this comparison for the employed (non-targeted). The discrepancies are qualitatively and quantitatively very similar to those in the August 2024 survey.

Fig. C.3 and C.4 show the validation checks for June 2024. The comparisons for the usual weekly earnings distribution, occupation, and industry shares are very similar for both the June 2024 and August 2024 surveys. For industry shares, in the June 2024 survey, "Other Services" has a much higher share in the RPS than in the CPS. This pattern has been prevalent in all previous versions of the RPS (see Bick and Blandin, 2023). In the June 2024 survey, we asked employed respondents to first choose from a list of 22 2-digit NAICS industries. Based on that choice, we then presented them with a list of more detailed (sub)industries. The over-representation of "Other Services" in the RPS stems from the over-representation of "Other Personal Services" in the follow-up question. This suggests that many individuals in the service sector are uncertain about which industry to select. This issue is not relevant for the CPS because, in that survey, respondents report a business or industry orally, and professional coders make the assignment after the fact.

In the August 2024 survey, we addressed this issue by asking individuals from this group: "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 Services - Other Personal Services."

The comparison between Fig. C.2a and C.4a suggests that this change resolved the issue.

Table C.1: Sample Composition in the June 2024 CPS and RPS

	Everyone		*Employed*	
	CPS	RPS	CPS	RPS
	(1)	(2)	(3)	(4)
Gender: Women	50.4	53.2	47.0	44.8
Age				
18-24	14.9	9.4	12.5	10.0
25-34	22.3	23.0	23.9	24.9
35-44	22.1	24.4	24.2	25.4
45-54	20.1	21.9	21.7	22.2
55-64	20.7	21.4	17.7	17.5
Race/Ethnicity				
Non-hispanic White	56.8	54.5	58.2	58.9
Non-hispanic Black	12.9	14.0	12.0	11.9
Hispanic	20.4	21.1	20.1	19.5
Other	9.9	10.4	9.7	9.8
Education				
Highschool or less	37.0	32.7	32.9	24.5
Some college/Associate's degree	25.8	28.5	25.5	29.1
Bachelor's or Graduate degree	37.2	38.8	41.6	46.5
Marital Status: Married	50.3	52.2	52.6	56.7
Number of children				
0	59.0	56.5	58.2	53.3
1	17.8	19.8	18.1	21.9
2	14.7	15.5	15.5	17.1
3+	8.5	8.2	8.3	7.8
Household Income in Last 12 Months				
$0-$50,000	25.6	31.0	19.6	19.8
$50,000-$100,000	29.9	28.4	30.6	31.0
$100,000+	44.5	40.7	49.8	49.1
Region				
Midwest	20.2	19.8	20.8	19.7
Northeast	17.0	18.9	17.4	18.7
South	38.7	38.8	38.0	38.3
West	24.0	22.5	23.8	23.4
Employment Status				
Employed, at work last week	71.0	66.4		
Employed, absent from work last week	3.5	2.8		
Unemployed	3.3	8.4		
Not in the labor force	22.3	22.3		
Observations	56969	2528	42365	1750

Notes: Column 1 reports the sample composition in the June 2024 Current Population Survey (CPS) for the variables targeted by Qualtrics in the sampling procedure. The employment status was the only variable not targeted. Column 2 reports the sample composition in the June 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).

Figure C.3: Validation Checks: Usual Weekly Earnings and Occupation for June 2024

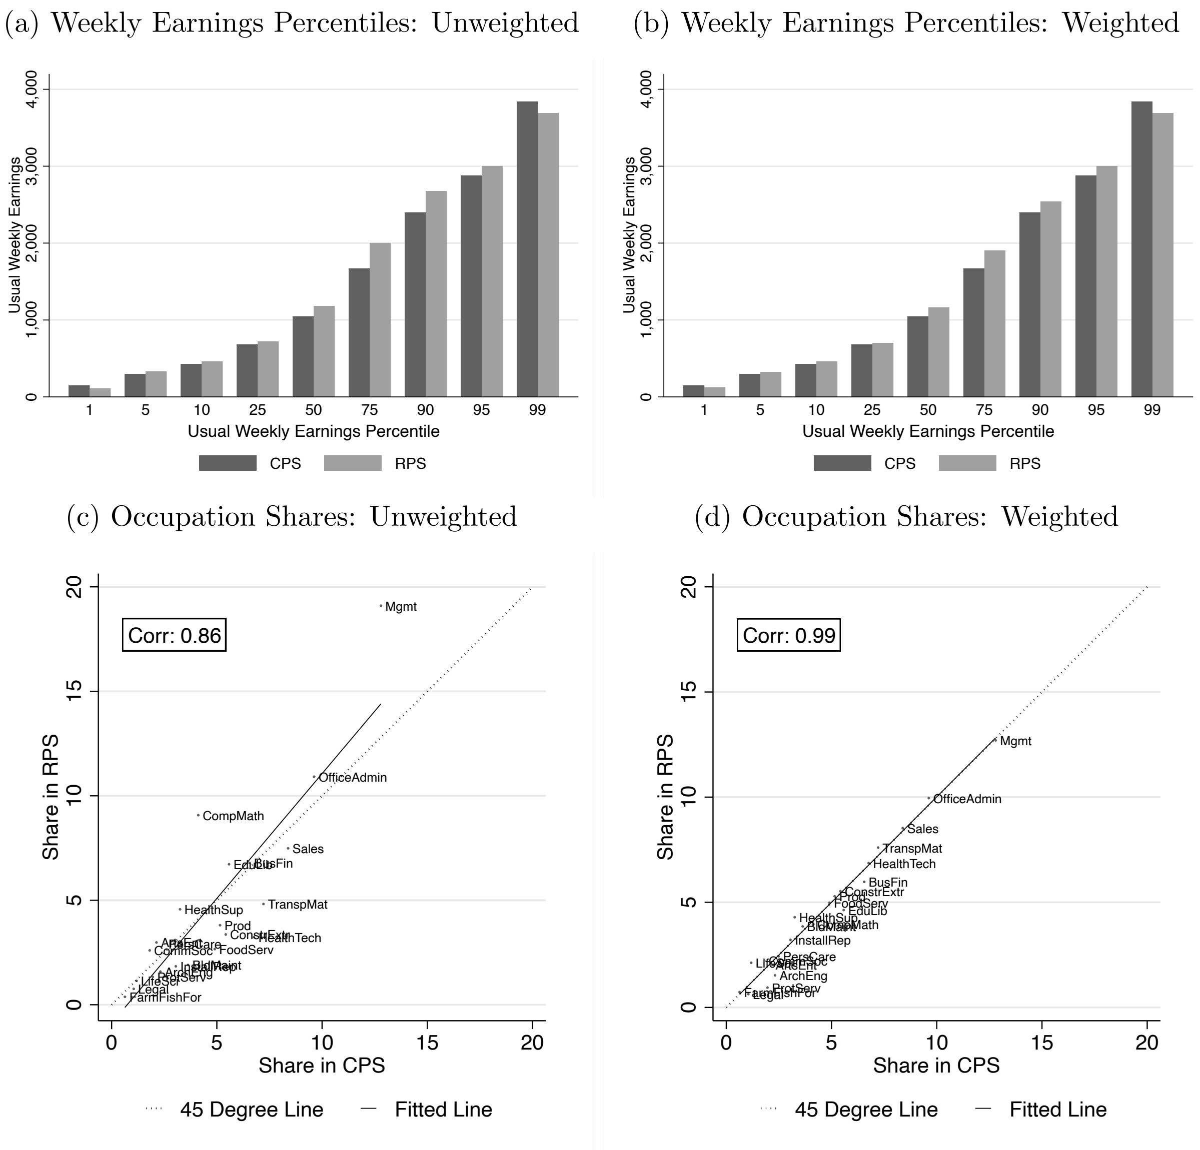

Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the sample of RPS respondents in both figures. All figures use weighted CPS data. Data samples for the weekly earnings figures are employees ages 18-64 in the June 2024 RPS and CPS-ORG with weekly earnings below the CPS topcode of $3,960.00 and an implied hourly wage of at least the federal minimum wage of $7.25. Sample sizes for the RPS and CPS are 1118 and 5899, respectively. Data samples for occupation are employed respondents ages 18-64 in the June 2024 RPS and CPS with sample sizes of 1576 and 42365, respectively.

Figure C.4: Validation Check: Industry and College Major for June 2024

(a) Industry Shares: Unweighted

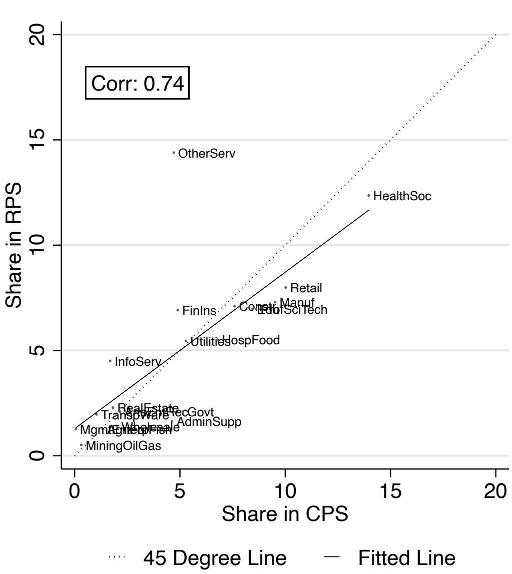

(b) Industry Shares: Weighted

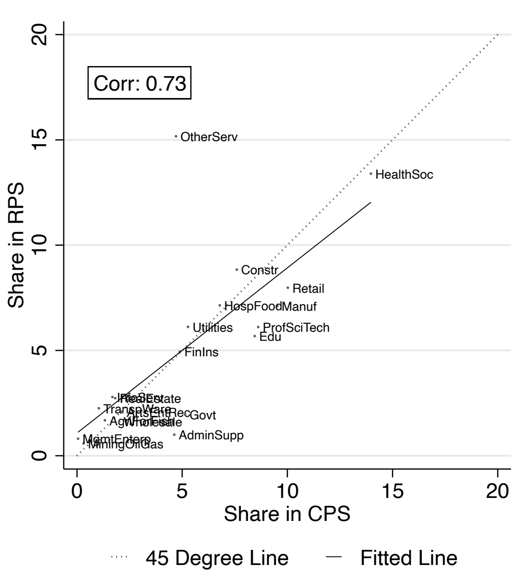

(c) College Major Shares: Unweighted

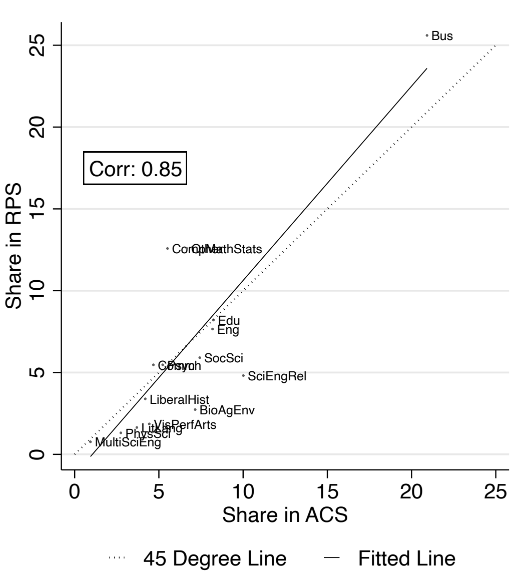

(d) College Major Shares: Weighted

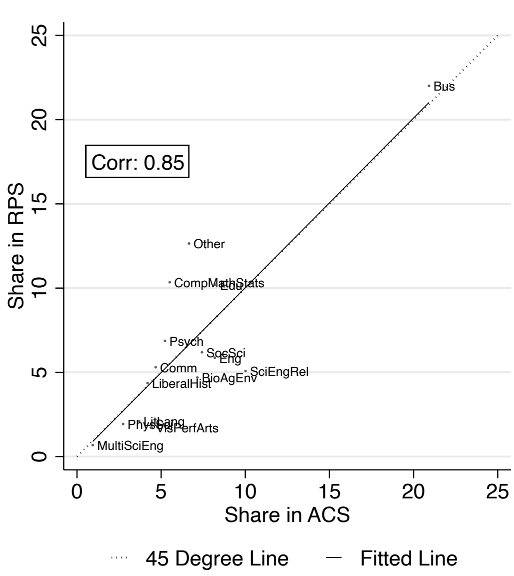

Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the sample of RPS respondents in both figures. All figures use weighted CPS and ACS data, respectively. For the industry comparison, sample sizes for the June 2024 RPS and CPS are 3216 and 45028, respectively. For the major comparison, sample sizes for the June 2024 RPS and the 2022 ACS are 1920 and 682280, respectively.