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, Revised February 2025

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.

The Rapid Adoption of Generative AI
Alexander Bick, Adam Blandin, and David J. Deming
NBER Working Paper No. 32966
September 2024, Revised February 2025
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 a series of nationally representative U.S. surveys of generative AI use at work and at home. As of late 2024, nearly 40 percent of the U.S. population age 18-64 uses generative AI. 23 percent of employed respondents had used generative AI for work at least once in the previous week, and 9 percent used it every work day. Relative to each technology's first mass-market product launch, work adoption of generative AI has been as fast as the personal computer (PC), and overall adoption has been faster than either PCs or the internet. Generative AI and PCs have very similar early adoption patterns by education, occupation, and other characteristics. Between 1 and 5 percent of all work hours are currently assisted by generative AI, and respondents report time savings equivalent to 1.4 percent of total work hours. This suggests that substantial productivity gains from generative AI are possible.

Alexander Bick
Federal Reserve Bank of St. Louis
One Federal Reserve Bank Plaza
St. Louis, MO 63102
alexander.bick@mailbox.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 (genAI) has rapidly emerged as a potentially important workplace technology. The large language model ChatGPT debuted in November 2022, and by December 2024 it reported over 300 million global weekly users. Several recent studies find that genAI 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; Toner-Rodgers, 2024) and some economists project that genAI will produce large aggregate economic effects (Chui et al., 2023; Filippucci, Gal, and Schief, 2024). Yet other studies predict only modest impacts of genAI on employment and productivity (Acemoglu, 2024). Uncertainty over the impact of genAI on the economy in part reflects a lack of systematic evidence on the frequency and intensity of genAI adoption, and most studies pre-date the release of widely accessible genAI tools (Acemoglu, Autor, Hazell, and Restrepo, 2022; Albanesi et al., 2024; Gathmann, Grimm, and Winkler, 2024).[1]

This paper studies the adoption of genAI using data from the first nationally representative U.S. survey of genAI usage at work and at home. Our data come from the Real-Time Population Survey (RPS), a national online labor market survey of working age adults aged 18-64 that has run since 2020 (Bick and Blandin, 2023). We report the combined results of two surveys fielded in August and November 2024, which collectively included more than 10,000 respondents.[2]

The RPS asks the same core questions and follows the same timing and structure of the Current Population Survey (CPS). This parallel structure allows us to benchmark our survey estimates against the CPS and to construct weights that ensure a nationally representative sample. Benchmarking to the CPS also allows us to compare the adoption of genAI to the personal computer (PC) and the internet. Beginning in 1984 the CPS fielded the Computer and Internet Use (CIU) supplement, which fueled an influential literature studying the impact of computerization on the labor market (Autor, Katz, and Krueger, 1998; Card and DiNardo, 2002; Krueger, 1993). We closely follow the CIU question wording and structure, which allows us to compare adoption of different technologies relative to their first mass market product release. But we also introduce novel questions to measure how intensively individuals use the technology and how much time it saves them.

We report six main results. First, a substantial share of respondents already use genAI at work and at home. In August and November of 2024, 26 percent of employed respondents reported using genAI for work: 9 percent used it every work day, 14 percent on some but not

[1] One exception is Eisfeldt, Schubert, Taska, and Zhang (2024) who show that occupations with core tasks that are more exposed to genAI have had lower job vacancies and wage growth relative to other occupations.

[2] Results from the August and November waves separately, as well as a pilot survey from June 2024, are included in Online Appendix B. All three waves yield similar estimates.

all work days in the previous week, and 3 percent did not use it in the last week. One third of respondents report using genAI outside of work, with 27 percent using it at least once in the previous week. Overall, 39 percent of respondents report using genAI either for work or outside of work. The most commonly used products are ChatGPT (28 percent of all respondents), Gemini (17 percent) and embedded products such as Microsoft Copilot (14 percent).

To validate our results we also fielded our genAI module in the Survey of Working Arrangements and Attitudes (SWAA) in December 2024. The SWAA is an online survey that has run monthly since May 2020 and uses a different online panel than the RPS (Barrero, Bloom, and Davis, 2021). The RPS and SWAA waves yield very similar adoption rates, both overall and by frequency of use, as well as the type of products being used. Our findings are also consistent with other studies of ChatGPT use in the US, including surveys that do not require internet access (Fletcher and Nielsen, 2024; McClain, 2024). Relative to these previous studies, our data provides several advantages: we have information about all genAI use rather than only ChatGPT; we distinguish between work and non-work use; and we collect more detailed data such as which workers use genAI, how frequently they use it, and how much time the technology saves them.

The widespread adoption of computers and related information technologies beginning in the 1980s contributed to rising aggregate productivity and growing income inequality (Autor, Levy, and Murnane, 2003). Comparing early adoption patterns of genAI to computers can help us understand whether genAI will have a similar impact to these earlier technologies (Autor, 2024). Our second main finding is that the adoption of genAI has so far been at least as fast as adoption of computers and the internet. In 1984, 25 percent of workers reported using computers for their job, compared with 27 percent of workers who report using genAI for their job in 2024. (The year 1984 was three years after the release of the IBM PC, the first mass-market computer; from this perspective, workplace adoption of genAI in 2024, two years after the release of ChatGPT, can be viewed as being on a similar trajectory.)

In contrast, genAI has been adopted at a faster rate than computers outside of work. While we cannot distinguish between work and non-work internet use, overall genAI adoption (including work and non-work usage) has been faster than both the personal computer and the internet. This could be due to differences in cost and user-friendliness. Computer adoption required the purchase of expensive hardware, and the internet required the purchase of a modem and an internet service provider contract, whereas ChatGPT is available for free or with a low-cost subscription, and using it does not require any technical expertise. Low adoption costs and the consumer focus of genAI products may also explain why individual adoption is faster than official firm-wide adoption (Babina, Fedyk, He, and Hodson, 2024; Bonney et al., 2024).[3]

[3] According to the Census Bureau's Business Trends and Outlook Survey, genAI adoption rose from 3.7 percent in December 2023 to 5.4 percent in February 2024, a rapid rise but still far below our estimates. This

Third, we find substantial variation in genAI adoption across demographic and labor market characteristics that closely matches early patterns of computer adoption. Both computer and genAI adoption have been faster for younger, more educated, and higher-wage workers. We also find similar adoption patterns by occupation and industry, although computer adoption was somewhat more concentrated in a few occupations whereas genAI adoption is more evenly dispersed. The one exception is gender: early genAI usage is higher for men, while computer usage was initially higher for women, reflecting rapid computer adoption among secretaries and other administrative occupations, which had large female shares.

Fourth, recent estimates of genAI task-based exposure are highly predictive of actual adoption. Eloundou, Manning, Mishkin, and Rock (2024) assign each occupation a predicted exposure score based on what fraction of that occupation's tasks they believe will be substantially affected by genAI. Their measure of "predicted genAI exposure" is highly correlated with adoption, which supports the utility of these estimates for researchers. Almost all occupation groups have at least a modest amount of genAI adoption, which is consistent with the breadth of exposure found in Eloundou, Manning, Mishkin, and Rock (2024). At the same time, some occupations have high adoption relative to predicted exposure (especially managers) and some have low adoption relative to predicted exposure (especially office and administrative support occupations).

Fifth, we find wide variation in the intensity of genAI usage. Among those who use genAI for work, 34 percent used it every workday in the previous week, 52 percent used it some but not all days in the previous week, and 14 percent did not use it at all the previous week. On days that workers used genAI, 32 percent used it for an hour or more per day, 47 percent used it between 15-59 minutes per day, and 21 percent used it for less than 15 minutes per day. Workers who use genAI on more days tend to use it more time per day. Combining these figures with information on days and hours worked in the previous week, we estimate that between 1 and 5 percent of all U.S. work hours involve the direct use of genAI.

Sixth, we provide a rough estimate of the early impact of genAI on aggregate productivity using data from self-reported time savings. We ask users how many additional hours they would have needed to complete the same amount of work in the previous week if they had not had access to genAI. We estimate a mean time savings of 5.4 percent among all genAI work users, which implies a mean time savings of 1.4 percent among all workers (including non-users).

suggests that firm-based measures of may not capture an important share of genAI adoption. Aaron Levie, the founder and CEO of Box, a cloud storage provider, suggests that this pattern has played out with past technologies: "It's almost a one-to-one parallel with the early wave of consumerization of IT. We saw that exact same dynamic in the early 2010s, which was employees of enterprises were pretty frustrated with their traditional document management and collaboration technology, and so they would go out and just go onto a web browser, sign up for a Box account, and then bring it into their organization. And for the first maybe couple of years, this was sort of happening really maybe not within the purview of IT, it was just employees and teams bringing in technology." (Levine, 2024)

Time savings are highly correlated with the intensity of genAI use and thus vary widely across occupations and industries. Using a standard model of aggregate production, we estimate a potential aggregate productivity gain of 1.1 percent at current levels of worker genAI usage. This is similar to Acemoglu (2024), who estimates a potential productivity gain of 0.7 percent using estimates of task-based exposure to genAI rather than actual adoption data.

The paper proceeds as follows. Section 2 describes the survey methodology and data. Section 3 presents our estimates of genAI adoption. Section 4 documents heterogeneity in adoption, and Section 5 explores genAI usage, time savings, and potential productivity gains. Section 6 concludes.

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.

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. Bick and Blandin (2023) show that the RPS produced very similar statistics for employment, hours worked, earnings, industry composition, and employee tenure during the pandemic.[4]

The RPS also collects information not contained in the CPS, including 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). Beginning in June and continuing through November 2024, the RPS introduced a module designed to measure

[4]As an additional point of comparison, Bick, Blandin, Caplan, and Caplan (2024a,b) show that estimates of work from home rates in the RPS align with estimates in the American Community Survey, the American Time Use Survey, the CPS, and the Survey of Income and Program Participation.

genAI use both at work and at home.

2.1.1 RPS Sample

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.[5] Finally, once fielding is completed Qualtrics staff goes over all responses and filters out any that look suspicious. This is typically the case for 1% to 3% of the responses.

We fielded the RPS with a pilot genAI module in June 2024 and received 2,551 responses. We then launched the full genAI module with the August and November 2024 waves and received 5,014 and 5,329 responses respectively. All surveys started fielding during the same weeks that the CPS conducted its corresponding surveys. For simplicity, we report the combined August and November results in the main paper. Online Appendix B replicates the main figures of the paper separately by survey waves.

Table 1 compares the sample composition between the CPS and RPS along the demographics targeted in the sampling procedure for our main surveys (columns 1 and 2). The most notable discrepancies are that individuals aged 18 to 24 and with no more than a high school degree are underrepresented in the RPS relative to the CPS, while individuals with household income of $50,000 or less are overrepresented. The bottom panel of Table 1 compares employment status in the CPS and RPS, statistics that have not been targeted in the sampling procedure. Employment rates are very similar across the two surveys, although individuals classified as unemployed according to the CPS definition are somewhat overrepresented in the RPS.

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

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

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

	Everyone		*Employed*	
	CPS	RPS	CPS	RPS
	(1)	(2)	(3)	(4)
Gender: Women	50.4	52.1	47.3	47.4
Age				
18-24	15.0	11.7	12.0	12.0
25-34	22.2	22.0	23.9	23.8
35-44	22.2	23.9	24.4	25.0
45-54	20.1	21.2	21.7	21.7
55-64	20.6	21.2	17.9	17.4
Race/Ethnicity				
Non-hispanic White	56.5	56.5	57.8	57.9
Non-hispanic Black	12.9	13.1	12.1	12.3
Hispanic	20.5	20.3	20.2	19.8
Other	10.1	10.0	9.9	10.1
Education				
Highschool or less	37.1	32.9	32.8	27.1
Some college/Associate's degree	25.8	27.5	25.2	27.6
Bachelor's or Graduate degree	37.1	39.5	42.0	45.4
Marital Status: Married	50.1	49.1	52.8	52.4
Number of children				
0	58.1	55.7	57.2	52.7
1	17.8	19.7	18.1	21.2
2	14.7	16.5	15.5	18.5
3+	9.3	8.1	9.2	7.6
Household Income in Last 12 Months				
$0-$50,000	26.1	31.3	19.9	22.6
$50,000-$100,000	29.9	29.5	30.7	32.0
$100,000+	44.0	39.2	49.5	45.4
Region				
Northeast	17.0	18.3	16.9	18.8
Midwest	20.3	19.4	21.1	19.2
South	38.8	38.1	38.1	37.8
West	23.8	24.2	23.9	24.2
Employment Status				
Employed, at work last week	71.7	69.9		
Employed, absent from work last week	2.5	2.9		
Unemployed	3.2	8.0		
Not in the labor force	22.6	19.2		
Observations	115477	10264	85546	7473

Notes: Column 1 reports the sample composition in the pooled August and November 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 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).

categories for education, employment, and marital status and interact all categories with gender. We match these key labor market statistics both in 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. The weighting scheme necessitated dropping some observations due to missing occupation and other discrepancies, resulting in a final sample size of 96.3% of initially collected responses. Appendices D.1 and D.2 detail the sample restrictions and the construction of the sample weights.

To illustrate how the RPS and CPS compare for characteristics that are not part of the sampling scheme and the effect of weighting, panels (a) and (b) of Figure 1 compare the usual weekly earnings distribution in the RPS and CPS, with and without weights, respectively.[6] 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 unweighted correlation is 0.88, with Management as the only major outlier. Applying the weights corrects this imbalance and mechanically increases the correlation to 1.

Appendix D.3 presents the same figures for each survey wave and analogous plots for industry and college major, which show qualitatively and quantitatively similar patterns to the occupation shares.

2.2 Measurement of Generative AI Use

A key goal when designing our survey was to provide a comparison to the historical adoption of other technologies. To facilitate this comparison, we use existing technology adoption questions from the CPS as a template for our own questions regarding genAI. In October 1984, the CPS launched a supplemental module to study the adoption of a new technology: the computer. (It was eventually named the Computer and Internet Use supplement, 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)*

[6] To minimize 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.

Figure 1: Validation Checks

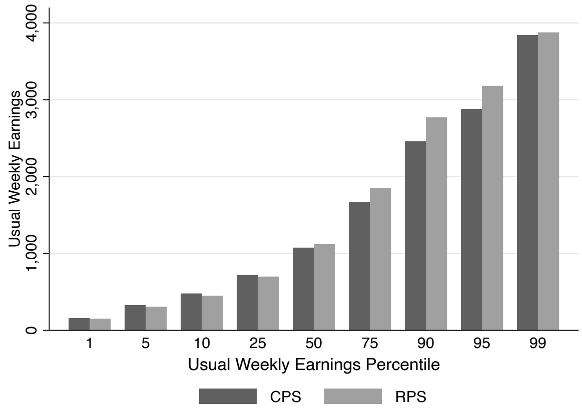

(a) Weekly Earnings Percentiles: Unweighted

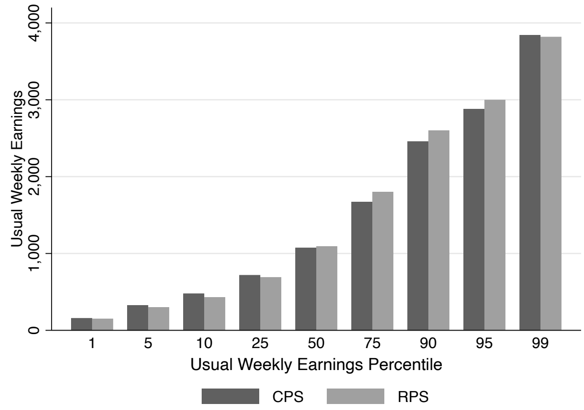

(b) Weekly Earnings Percentiles: Weighted

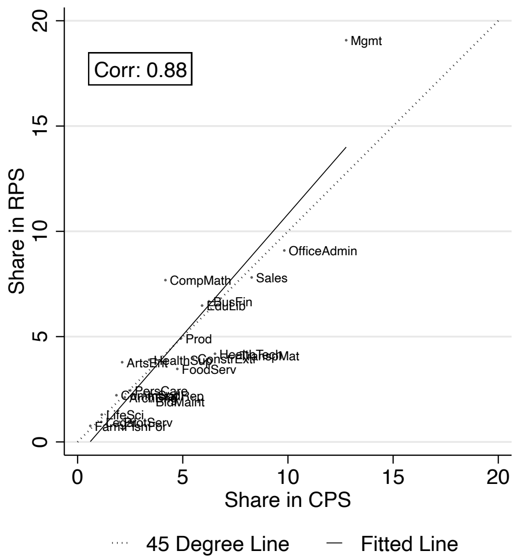

(c) Occupation Shares: Unweighted

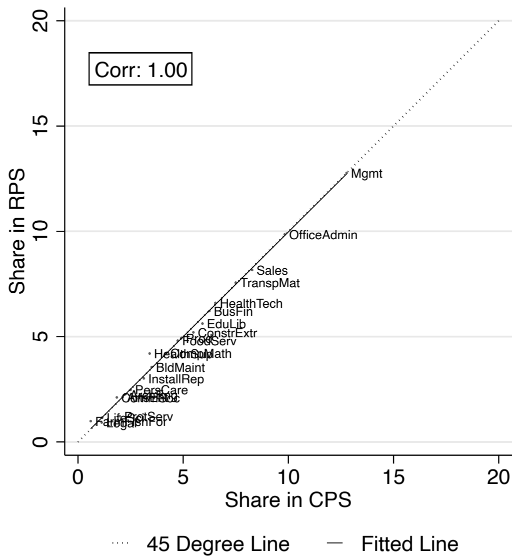

(d) Occupation Shares: Weighted

Notes: Figures on the left use unweighted RPS data, figures on the right use weighted RPS data. We use the same sample of RPS respondents in both figures. All figures use weighted CPS data. Data samples for the weekly earnings figures are employees ages 18-64 in the pooled August and November 2024 RPS and CPS-ORG with weekly earnings below the CPS topcode of $4140 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 4786 and 12133, respectively. Data samples for occupation are employed respondents ages 18-64 in the pooled August and November 2024 RPS and CPS with sample sizes of 6951 and 85546, respectively.

These questions were asked again in 1989, 1993, 1997, 2001, and 2003, with the word "directly" being omitted in the last two years. In 2001 a question about internet usage was added:

> *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.

Because genAI is a relatively new technology, we believe it is important to provide both a definition of the concept and some specific examples to our survey respondents. We therefore start the genAI survey module 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 whether they had heard of the concept prior to the survey. Overall, 74.5% report having heard of genAI; the remaining 25.5% who answered "No" skip the remainder of the module. Respondents who answer "Yes" continue with the survey, which then later asks them to select from a list of products or to write in others that they have used.

For employed respondents, the next question asks about genAI use at work:

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

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

A second question—though not immediately followed by this question—asks about genAI usage outside work.

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

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

All genAI users are asked about which product they used and can select from a list of products (e.g. ChatGPT, Google Gemini) or to write in others that they have used.

Among employed respondents who say they use genAI for their job, we ask how many days they used genAI in the last week. If they did not use it last week, we asked whether they used it in the last four weeks. We also asked how much time per day they spent using genAI on the days they used it. Second, we ask which specific products the respondent used, which specific tasks genAI helped with, and some broader questions about the uses and benefits of genAI.

Non-employed respondents who use genAI receive the same set of follow-up questions as do the employed who use genAI outside their job, in the latter case with an explicit reference to usage outside their job.

3 How Prevalent is Generative AI Use?

Figure 2a displays the share of August and November 2024 RPS respondents who report using genAI. The samples for the "Overall" and "Outside of Work" bars are all respondents, while the sample for the "For Work" bar is employed respondents. The topline estimate reflects the share who answer "Yes" to either the question "*Do you use Generative AI for your job?*" or the question "*Do you use Generative AI [outside your job]*"? We then further categorize genAI users according to their response to the follow up question "*Did you use genAI [for your job / outside your job] LAST WEEK?*"

The first bar shows that 39.6 percent of respondents report using genAI either at work or at home: 9.2 percent used it every day, 23.1 percent used it some but not all days, and 7.3 percent did not use it in the previous week. The second bar shows that 26.5 percent of workers used genAI at work: 9.0 percent used it every work day and 13.9 percent on some but not all workdays, bringing total weekly usage to 22.9 percent.[7] Outside of work, usage is higher overall (35.5 percent) but less intensive, with only 5.2 percent of all respondents using it every day last week.

3.1 How Do Our Results Compare to Other Estimates?

The RPS is administered by the commercial survey provider Qualtrics. To provide a second source of data, in December 2024 we ran our genAI module through a different online survey, the Survey of Working Arrangements and Attitudes (SWAA) (Barrero, Bloom, and Davis, 2021). The SWAA has been running monthly since May 2020 and is conducted using a different survey provider (IncQuery), but targets a similar sample that is intended to be representative of the U.S. ages 20-64 by age, sex, income, race, and region. To maximize comparability between the

[7] We included more detailed response categories in the November 2024 wave of the RPS, see Figure B.1. Among respondents who report using genAI at least once in the last week but not every day, about 55 percent reported using it on multiple days and 45 percent reported using it exactly one day.

Figure 2: Share of Working Age Adults Using Generative AI

(a) RPS: August & November 2024

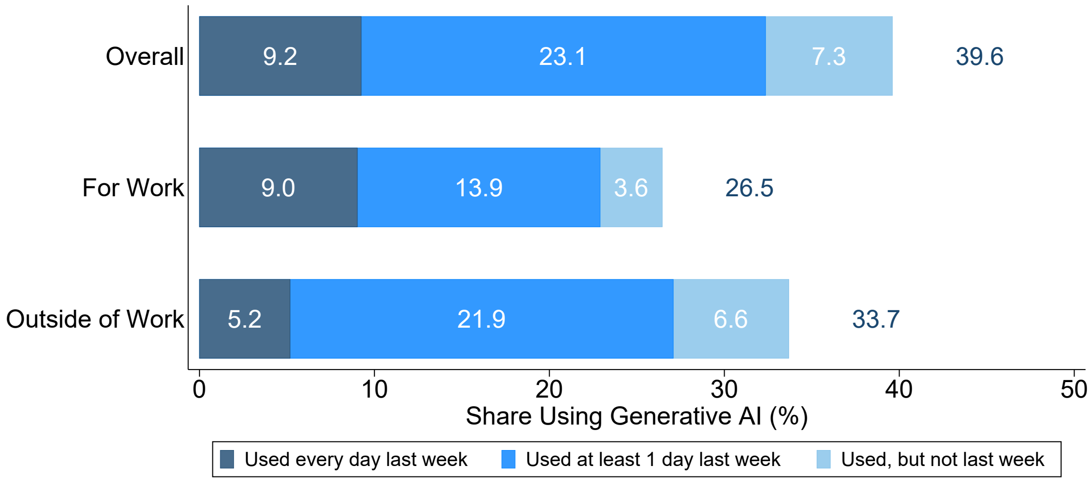

(b) SWAA: December 2024

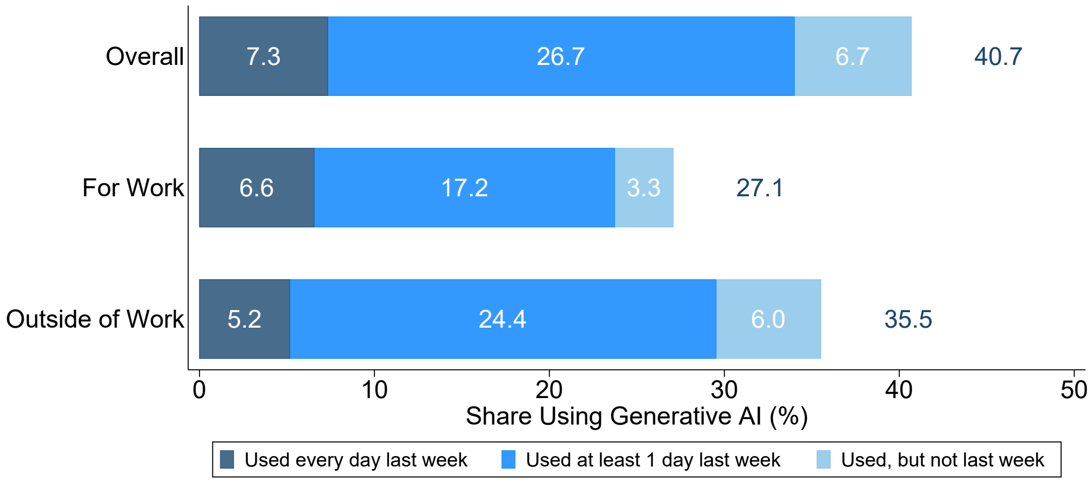

Notes: The figure shows the share of respondents who use genAI for work, outside of work, and overall (either for work or outside of work). Intensity of use is broken down into every day last week (dark blue), at least one day but not every day last week (medium blue), and not last week (light blue). Data source for panel (a) is the August and November 2024 waves of the RPS, ages 18-64. The "For Work" sample is employed individuals ($N = 6951$); the other bars include all respondents ($N = 9742$). Data source for panel (b) is the December 2024 wave of the SWAA, ages 20-64. The "For Work" sample is employed individuals ($N = 3516$); the other bars include all respondents ($N = 4698$).

two surveys, we construct sample weights according to the RPS procedure rather than using the weights provided by the SWAA.[8]

[8]Our weighting scheme incorporates a broader set of variables, including employment status. The employment-to-population ratio calculated using the SWAA's provided weights is substantially lower than the

Figure 2b displays parallel results from the Survey of Working Arrangements and Attitudes (SWAA). The December 2024 wave ($N = 4,698$) of the SWAA yields very similar rates of genAI usage to the RPS. 40.7 percent of SWAA respondents reported using genAI either for work or outside of work, compared with 39.6 percent of RPS respondents. We also find similar estimates of use for work (27.1 percent in the SWAA versus 26.5 percent in the RPS) and outside of work (35.5 percent in the SWAA versus 33.7 percent in the RPS). Online Appendix C includes additional comparisons between the SWAA and RPS. Because the two surveys yield similar estimates but the SWAA only includes a subset of the genAI questions in the RPS, the remainder of the main text reports results only from the RPS.

While our survey asks about all forms of genAI use, we also collect usage data for specific genAI products. This enables us to compare ChatGPT use reported in the RPS to other existing surveys. We provided genAI users with a list of common products and asked them to select which products they used (respondents could select multiple products and could write in an "Other" option if it was not on the list we provided). Appendix Figure A.1 presents the share of all respondents who reported using each genAI product. The three most common were ChatGPT (28.1 percent of all respondents), Gemini (16.6 percent) and genAI products embedded within existing software, such as Microsoft Copilot (14.1 percent). In April 2024, Fletcher and Nielsen (2024) conducted an online survey of ChatGPT adoption using a different survey provider (YouGov). They found that 18 percent of U.S. adults used ChatGPT at least weekly, compared to 19 percent who used ChatGPT in the previous week in our study.[9]

A potential concern with online-based surveys is that they may suffer from selection based on unobservable characteristics that are correlated with genAI use. For example, it may be that online survey respondents are more comfortable using technology and are therefore more likely to use genAI than the overall U.S. population. From this perspective, a valuable point of comparison is a survey conducted by the Pew Research Center (McClain, 2024). That survey sampled respondents using a nationally representative sample of U.S. Postal Services addresses. Importantly, the survey included respondents without internet access.[10] They find that in February 2024, 27 percent of U.S. adults age 18-64 reported ever having used ChatGPT, compared to 29 percent of individuals who report using ChatGPT in our survey in August and November 2024.

CPS ratio for the 20–64 age group. Additional details on the construction of our weights and a comparison of the raw data with the December 2024 CPS are provided in Appendix E.

[9] We also find patterns of ChatGPT adoption by demographic and labor market characteristics that align with other studies. Humlum and Vestergaard (2024) survey a representative sample of workers in eleven occupations in Denmark about their usage of ChatGPT at work. We find similar usage rates by age, education and gender, and in the occupations covered by both surveys, although the lack of clean correspondence between job codes across countries makes an exact comparison difficult.

[10] Internet access per se is unlikely to substantially impact our results: in the 2022 American Community Survey, 96.8% of individuals aged 18-64 report living in a household with access to the internet.

Figure 3: The Trajectory of Computer, Internet, and Generative AI Adoption

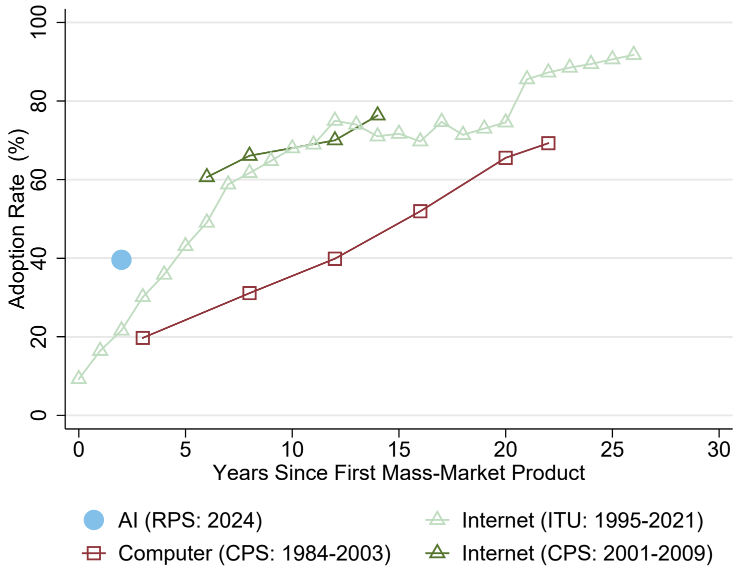

Notes: The figure shows usage rates at work for three technologies: genAI, computers, and the internet. The horizontal axis represents years since the introduction of the first mass-market product for each technology. We use 1981 as the introduction year for computers, which was the year the IBM PC was 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 2022 as the introduction year for genAI, which was the year ChatGPT was released. The data source for genAI is the August and November 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 $N = 9742$. The sample for the ITU is individuals of all ages.

3.2 How Fast is Generative AI Being Adopted Compared to Other Technologies?

Figure 3 compares the pace of genAI adoption with two other important technologies: PCs and the internet. (As described in Section 2, adoption was measured using very similar questions for all three technologies.) The horizontal axis represents years since the release of the first mass-market product for a given technology. The first mass-market computer was the IBM PC, which was released in August 1981. 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). The first mass-market genAI product was ChatGPT, which was released in November 2022.[11]

[11] Three popular computer products had already been released in 1977 (the Apple II, the Commodore PET, and the TRS-80); choosing 1977 as the year of the first mass-market PCs would imply a slower pace of computer adoption by 1984. The 1981 IBM PC was the first computer to sell over one million units (Abbate, 1999). With regards to the internet, 1995 was also the year of Netscape's initial public offering and the year that AOL 3.0 was released. We chose these dates based in part on ChatGPT's answers to questions about the year of the first mass market product for each technology - see Appendix F for the exact text of the prompt and the answers given.

The blue dot in Figure 3 repeats the 39 percent adoption rate for genAI reported in Figure 2a, which corresponds to 2 years since the first mass-market product. The red squares plot PC adoption in the CPS between years 3 to 22 (1984 to 2003). The dark green triangles plot internet adoption in the CPS between years 6 to 14 (2001 to 2009). The light green triangles are based on data from the International Telecommunication Union (ITU), which has gathered U.S. and global internet usage data since 1995, in partnership with the World Bank. The ITU combines subscriber data from national regulatory bodies and service providers to estimate the population's internet access rate (Peña-López et al., 2009). We plot internet adoption in the ITU between years 0 to 26 (1995 to 2021). The two series for the internet align closely for the years that overlap.

PC adoption rose steadily from 20 percent in year three to 70 percent in year 22. Internet adoption increased rapidly from 20 percent in year two to 60 percent in year seven, and then more gradually to 90 percent over the ensuing two decades. We conclude that, relative to the introduction of the first mass-market product, genAI has been adopted at a faster pace than both PCs and the internet.

Because all three technologies can be used by both consumers and workers, an obvious question is to ask whether the faster pace of genAI adoption was driven by work or non-work use. Unfortunately, neither dataset on internet adoption distinguishes between work and non-work use. However, we can distinguish between these cases for PCs and genAI.

The faster overall pace of genAI adoption compared with PCs is driven by non-work adoption. Figure 4a plots work adoption of PCs and genAI. We find very similar work adoption rates for genAI after two years (27 percent) and PCs after three years (25 percent). By contrast, Figure 4b reveals much higher non-work adoption rates for genAI (34 percent) compared with PCs (5 percent).

The choice of whether to adopt a given technology is a function of both its benefits and its cost. PC adoption required the purchase of relatively expensive hardware that was initially not very mobile, implying a substantial cost of adoption. Internet adoption required the purchase of a modem and an Internet Service Provider subscription. By contrast, many genAI products are currently free or inexpensive and user-friendly. This suggests that one likely factor for the relatively rapid adoption of genAI outside of work is its lower adoption cost relative to PCs and the internet. On the one hand, this means that similar adoption rates for PCs and genAI at work do not necessarily imply equivalent benefits from the two technologies. On the other hand, lower adoption costs may also allow genAI to diffuse more rapidly.

Figure 4: The Trajectory of Generative AI and PC Adoption For Work and Outside of Work

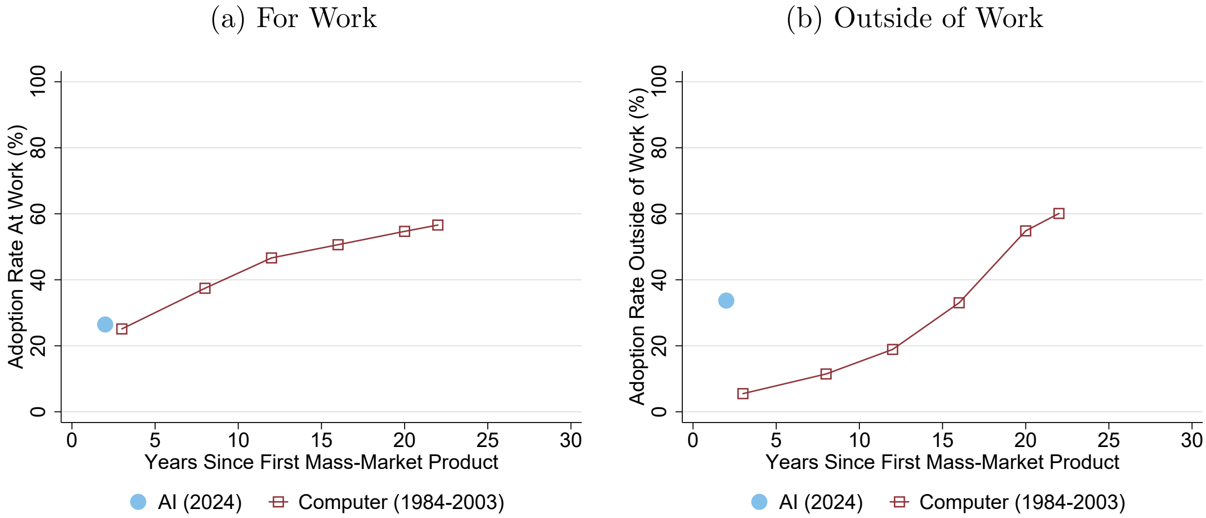

Notes: The figure shows usage rates for work (Panel a) and at home (Panel b) for two technologies: genAI and computers. The horizontal axis represents years since the introduction of the first mass-market product for each technology. We use 2022 as the introduction year for genAI, which was the year 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 genAI is the August and November 2024 waves 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 Panel a is employed individuals ages 18-64 (RPS, $N = 6951$). The sample for Panel b is all individuals ages 18-64 (RPS, $N = 9742$).

4 Which Workers Use Generative AI?

Figure 4a documented a very similar work adoption rate for PCs in 1984 compared with generative AI in 2024. Did the same groups of workers that drove PC adoption in 1984 also drive genAI adoption 40 years later? Answering this question is a preliminary step toward understanding the heterogeneous labor market impact of genAI (Acemoglu, 2024; Autor, 2024; Autor, Levy, and Murnane, 2003).

Figure 5 presents adoption by demographic group for PC's in 1984 and genAI in 2024. Figure 5a shows that genAI adoption is about twice as high (40 percent vs. 20 percent) for workers with Bachelors and/or Graduate degrees, compared to workers with less education. This pattern is strikingly similar for PC's. Figure 5b presents results by age, and again we find very similar patterns for genAI and PCs. In 2024, workers under the age of 50 were much more likely to use genAI than those over 50, and the same was true for PCs in 1984. However, we find different adoption patterns by sex. Figure 5c shows that genAI adoption is 7.5 percentage points higher for men, while early PC adoption was 6 percentage points higher for women. The high rate of PC adoption among women is driven by frequent and early use of PCs among secretaries and other administrative support occupations, which was a highly female occupation in 1984.

Figure 5: Generative AI and PC Adoption For Work by Demographic Group

(a) By Education

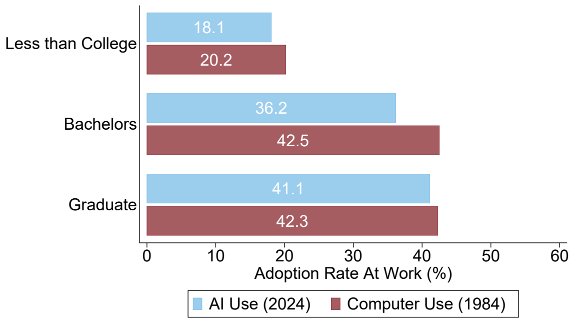

(b) By Age

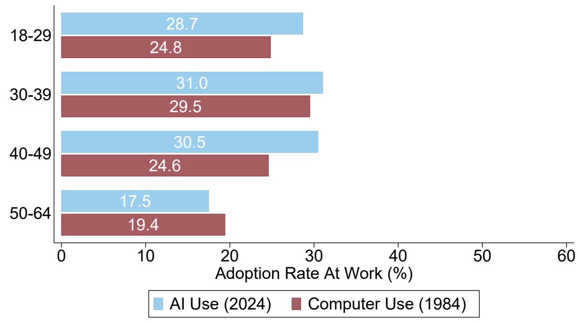

(c) By Sex

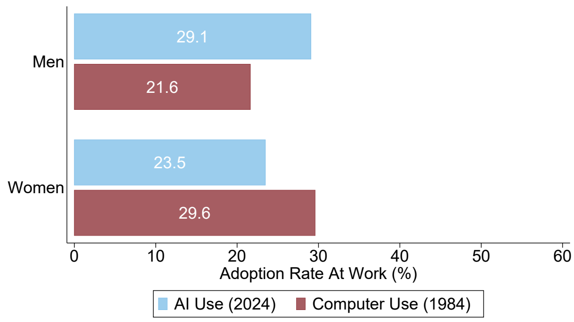

Notes: The figure shows adoption rates at work for genAI and PCs by education (Panel a), age (Panel b), and sex (Panel c). The data source for genAI is the August and November 2024 waves of the RPS (blue bars). The data source for computers is the 1984 Computer and Internet Use Supplement of the CPS (red bars). The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 6951$).

Figure 6: Earnings and Generative AI / PC Adoption

(a) By Weekly Earnings Percentile

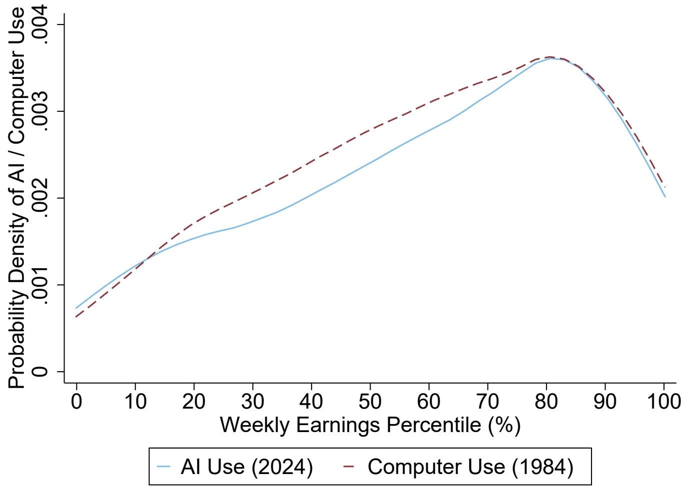

(b) Reg. Coefficient of PC / genAI Use

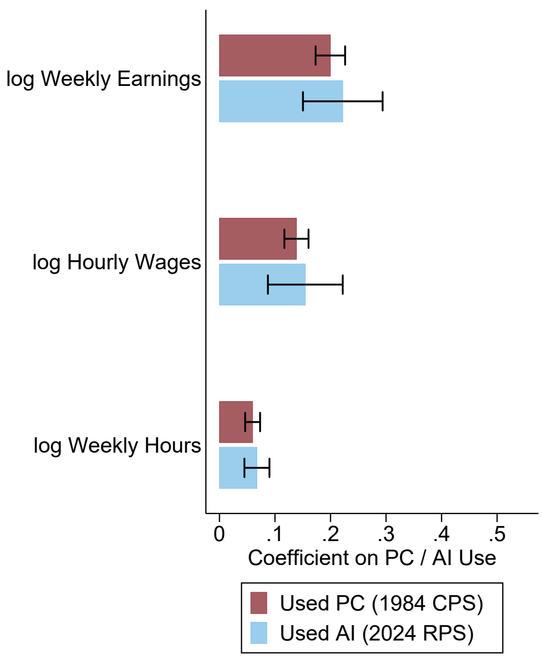

(c) Reg. Coefficient of Intensity of genAI Use

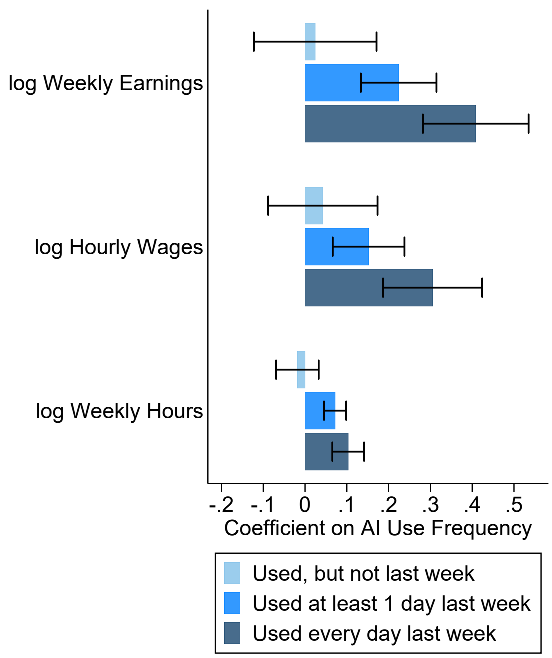

Notes: Panel (a) shows the kernel density of genAI adoption (blue solid line) and PC adoption (red dashed line) by weekly earnings percentile. Panel (b) shows the coefficient on PC adoption (red bars) and genAI adoption (blue bars) in a regression of (i) log weekly earnings, (ii) log hourly wages, and (iii) log weekly hours worked. Panel (c) shows the coefficients on genAI intensity-of-use (zero days last week, some but not all workdays last week, and every workday last week) for the same variables (i-iii). 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 1984 Computer and Internet Use Supplement of the CPS. The sample for each dataset is employed individuals ages 18-64 (RPS, $N = 6951$).

Figure 7: Generative AI and PC Adoption by Occupation and Industry

(a) By Occupation (b) By Industry

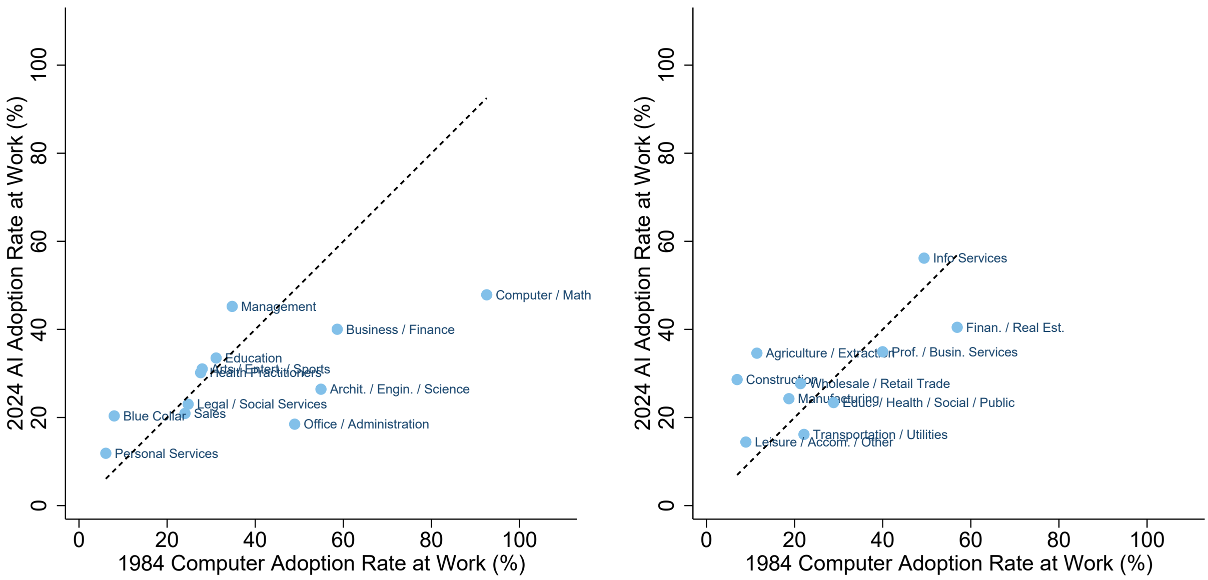

Notes: The figure shows adoption rates at work for PCs (horizontal axis) and genAI (vertical axis) by occupation group (Panel a), and industry group (Panel b). For Panel a, Personal Services occupations combine SOC codes 31-39: Healthcare support, Protective services, Food preparation and serving, Cleaning and maintenance, and Personal care. Blue Collar occupations combine SOC codes 47-53: Construction, Extraction, Installation, Maintenance and Repair, Production, Transportation, and Moving. For panel (b), the industry groupings 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). The data source for genAI is the August and November 2024 waves of the RPS. The data source for PCs is the 1984 Computer and Internet Use Supplement of the CPS. The sample for each dataset is employed individuals ages 18-64 with valid occupations and industries (RPS, $N = 6898$).

Figure 6a plots the kernel density of PC and genAI adoption by weekly earnings percentile in 1984 and 2024 respectively. We again find strikingly similar patterns for the two technologies. Adoption rates increase steadily with income until roughly the 80th percentile, and then decline.

To better understand the relationship between technology use and labor market outcomes, we regress indicators for PC and genAI use on earnings, wages, and hours while also controlling for occupation, industry, education, and demographic characteristics.[12] The results are displayed in Figure 6b. We find again similar results for PCs and genAI. Use of both technologies is associated with a weekly earnings premium of 21 log points, a wage premium of about 15 log pints, and 5 log points weekly hours.

Figure 6c estimates a version of our regression with separate indicators for infrequent, weekly, and daily users of genAI. We cannot run an analogous regression for PCs since the CPS does not ask questions about usage intensity. Controlling for demographics, occupation, and industry, workers who used genAI every workday in the previous week earn 40 percent more than non-users, and workers who used it some but not all workdays earn 22 percent more than non-users. By contrast, there is no statistical difference in earnings for workers who report using genAI but did not use it within the previous week. Most of these earnings differences are attributed to differences in hourly wages: daily users receive 30 percent higher wages and weekly users receive 15 percent higher wages.

Higher wages by more intensive users could reflect several factors. First, there is substantial variation in job tasks even within occupation and industry, and tasks that benefit more from genAI assistance may be associated with higher wages (Deming and Kahn, 2018). Second, productive workers may adopt new technologies more rapidly (Nelson and Phelps, 1966). Third, genAI may directly increase workers' productivity and wages. Our data do not allow us to distinguish between these explanations.

Figure 7 compares early adoption patterns for the two technologies by occupation and industry. Figure 7a plots genAI adoption against PC adoption by occupation group.[13] genAI adoption at work is highest in computer and mathematical, management, and business and finance jobs (43, 42, and 41 percent respectively). However, genAI use is broadly distributed, with every occupation group between 15 and 50 percent. In contrast, PC use is more concentrated, with three occupation groups above 50 percent and two occupation groups below 10

[12] The control variables used in these regressions are indicators for: female, age 18-29, age 40-49, age 50-54, a bachelor's degree, a graduate degree, Black, Hispanic, occupation group, and industry.

[13] We elicited respondents' job titles through a free text response with autocomplete suggestions covering over 40,000 occupations from O*NET and the Occupational Outlook Handbook and then match them to Standard Occupation Classification (SOC) codes using a parsing algorithm that identifies occupations in 97 percent of cases. For job titles that do not exactly match a unique SOC code, we present respondents with a choice of probabilistic matches using the job title to SOC code matching algorithm developed by the National Institute for Occupational Safety and Health (Laughlin, Song, Wisniewski, and Xu, 2024).

Figure 8: Predicted Generative AI Exposure vs. Adoption By Occupation

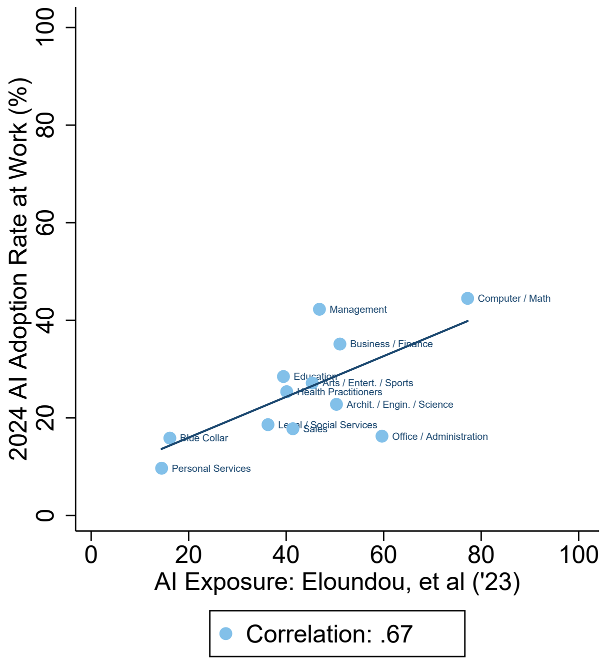

Notes: The figure compares predicted genAI exposure with genAI use by occupation group. The data source for predicted genAI exposure is ChatGPT-β predicted exposure from Eloundou, Manning, Mishkin, and Rock (2024). The data source for genAI use is the August and November 2024 RPS. The sample for the RPS is employed individuals age 18-64 with a valid occupation ($N = 6898$).

percent. Still, early adoption rates are highly correlated across the two technologies ($\rho = 0.66$).

Figure 7b plots genAI adoption against PC adoption by industry. genAI adoption at work is highest in information services (58 percent). GenAI adoption is lowest in leisure, accommodation, and other services and transportation and utilities (15 percent and 17 percent, respectively). Overall, adoption rates by industry are highly correlated across the two technologies ($\rho = 0.71$). It is interesting to note that early differences in PC adoption by occupation and industry were highly persistent over the next several decades, see Figure A.4.

4.1 Tasks, Predicted Exposure, and Generative AI Adoption

Why do some groups of workers adopt genAI at higher rates than others? One source of variation in genAI adoption could be that the technology is simply more helpful in some jobs than in others. To investigate this possibility, we use the predicted-task-exposure analysis of Eloundou, Manning, Mishkin, and Rock (2024).[14] They construct a rubric that assigns each

[14] Eisfeldt, Schubert, Taska, and Zhang (2024) construct a similar task-based-exposure measure and show that firms whose employees were more exposed to genAI experienced abnormally large increases in stock prices in the days following the release of ChatGPT.

Figure 9: In Which Specific Work Tasks Is Generative AI Most Useful?

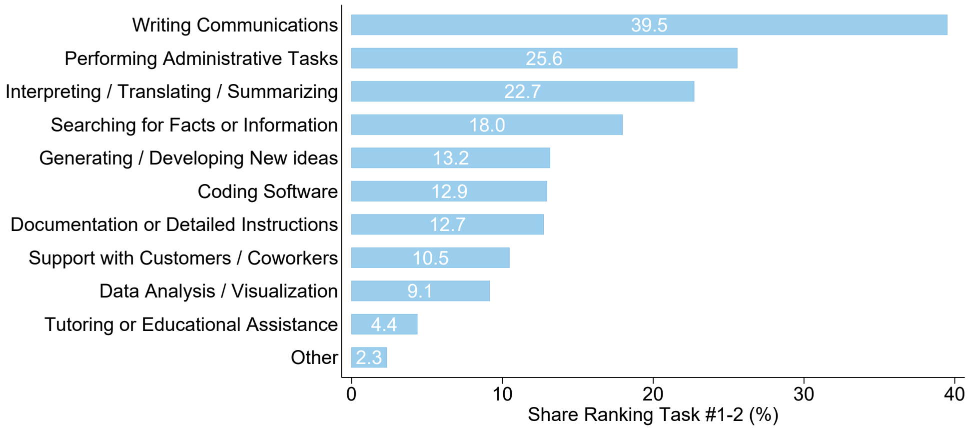

Notes: The figure shows which work tasks genAI users report that genAI is most helpful in completing. Employed respondents who used genAI were first provided with a list of tasks and asked to select those that they had used genAI to help with last week. Respondents were then asked to rank these selected tasks according to how helpful genAI was in completing the task. The figure reports the share of genAI 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 and November 2024 waves of the RPS, ages 18-64. The sample is employed individuals who used genAI and selected at least one task ($N = 4698$).

task in the O*NET database a zero or one corresponding to whether LLMs (a popular class of genAI) could feasibly produce large productivity gains for that task. They then aggregate these task exposure scores by occupation to assign each occupation a predicted exposure score based on what fraction of tasks could be impacted by LLMs.

Figure 8 plots an occupation group's predicted genAI exposure from Eloundou, Manning, Mishkin, and Rock (2024) against actual genAI adoption in the 2024 RPS.[15] We find that predicted genAI exposure is highly correlated with actual genAI adoption in that occupation ($\rho = 0.67$). Personal service and blue-collar occupations have relatively low predicted exposure and also relatively low adoption, while computer and mathematical occupations have both high predicted exposure and high adoption.

Our results offer partial validation of task-based exposure predictions, supporting the utility of these measures for researchers. At the same time, managers have relatively high adoption rates relative to predicted exposure, while legal and social services and office and administration occupations have relatively low adoption rates. These disparities could reflect some mismeasurement in which tasks can benefit from genAI. Alternatively, they could reflect variation in adoption costs or regulatory barriers.

The close correspondence between predicted exposure based on occupational tasks and actual adoption suggests that the task content of a worker's job may be an important determinant

[15] Note that while both variables are defined over the interval 0 to 100%, the measure of an occupation's predicted task exposure is different from the adoption rate.

of their adoption choice. A natural follow up question is to ask which tasks workers say they use genAI to help them complete. For respondents who indicated that they had used genAI in the last week, we present the list of tasks in Figure 9 and ask them to select any for which they had used genAI 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 genAI was in completing the task.

Figure 9 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 (39.5 percent), administrative tasks (25.6 percent), interpreting/translating/summarizing text or data (22.7 percent), and searching for facts or information (18.0 percent). While genAI seems most useful for creating, absorbing, and organizing written information or communications, eight of the ten tasks in our list were ranked in the top two by at least 10 percent of users, indicating a variety of use cases for genAI.

5 Frequency of Generative AI Use and Time Savings

5.1 How Much Do Workers Use Generative AI?

Figure 10 reports how much time genAI users spent using the technology in a given day. Specifically, we ask "*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?*" Respondents could select from three options: 15 minutes or less per day, between 15 minutes and one hour per day, or more than an hour per day.[16]

Figure 10 shows that 32 percent of genAI users report using genAI for an hour or more per day at work, 47 percent used it for between 15 and 60 minutes per day, and 21 percent used it for less than 15 minutes per day. We also show these estimates separately by workers' daily frequency of use. Frequency and intensity of genAI use are positively correlated. 52 percent of daily genAI users also report using it an hour or more each day, compared to only 7 percent of those who used it in the last month but not the last week. Overall, our results indicate wide variation in the frequency and intensity of genAI use at work.

We estimate bounds on the share of total work hours assisted by genAI by combining data on usage intensity with data on days and hours worked in the previous week. For this calculation we restrict attention to our November 2024 wave, which contains more detailed intensity categories.[17] For example, consider a worker who reports (i) working 40 hours last week

[16] Respondents who did not use genAI last week but did use it in the last four weeks were asked about their average usage on the days they used genAI over the last four weeks.

[17] In the August wave, workers could report using genAI (i) every workday last week, (ii) some but not all workdays last week, or (iii) zero days last week. In the November wave, we disaggregated (ii) into two options:

Figure 10: Intensity of Generative AI Use For Work

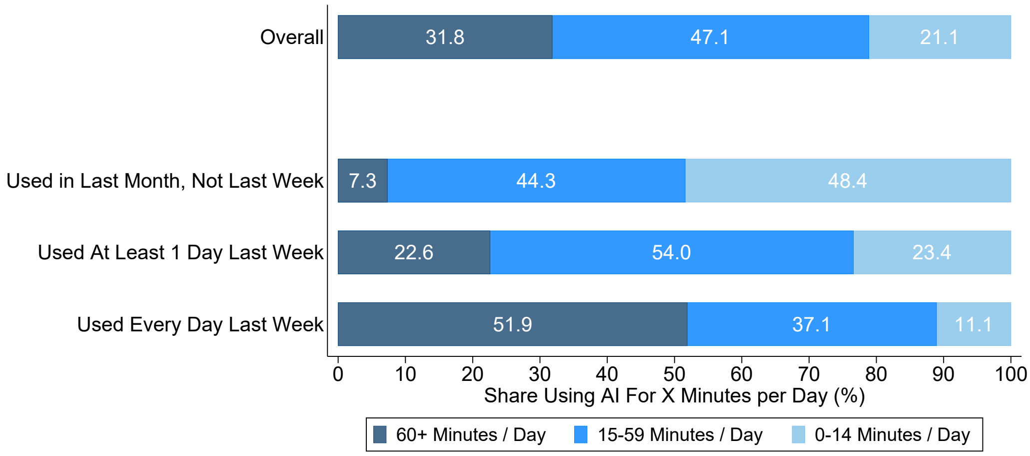

Notes: The figure shows the distribution of daily time spent actively using genAI for work, among genAI users. Usage time is broken down into 0-14 minutes per day (light blue), 15-59 minutes per day (medium blue), and 60 or more minutes per day (dark blue). The "Overall" bar reflects the distribution among all workers who use genAI. The "Used in Last Month, Not Last Week" bar reflects users who did not use genAI for 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 genAI at least one day last week but not every workday. The "Used every day last week" bar reflects users who used genAI for work every workday last week. Data source is the August and November 2024 waves of the RPS, ages 18-64. The sample is employed respondents who use genAI for work ($N = 1918$).

over 5 days, (ii) using genAI on multiple but not all workdays last week, and (iii) using genAI for between 15 and 60 minutes per day on days that they used it. For a lower bound we assume the respondent used genAI on exactly two days for 15 minutes, implying $15/60 \cdot 2/40 = 1.3$ percent of that worker's hours last week directly used genAI. The upper bound in these cases would be using genAI on four workdays for 1 hour each day, implying $1 \cdot 4/40 = 10$ percent of that worker's hours last week directly used genAI.

We first aggregate these estimates across the 21.7 percent of workers in the November 2024 wave who report using genAI in the past week. Among workers who use genAI, between 6.0 and 24.9 percent of all work hours were assisted by genAI. Next, we include the 78.3 percent of non-genAI-workers in our calculation, who spend zero percent of their time using genAI by construction. Among all workers, between 1.3 and 5.4 percent of total work hours were assisted by genAI.

Figure 11: Reported Time Savings Due to Generative AI

(a) Time Savings by Frequency of genAI Use

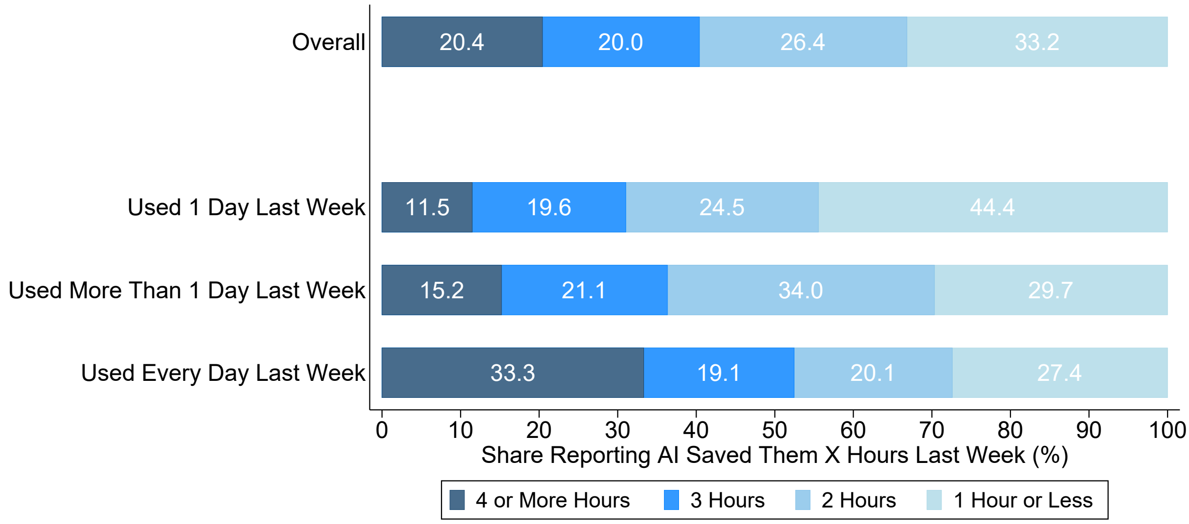

(b) genAI Time Savings and Use by Occupation

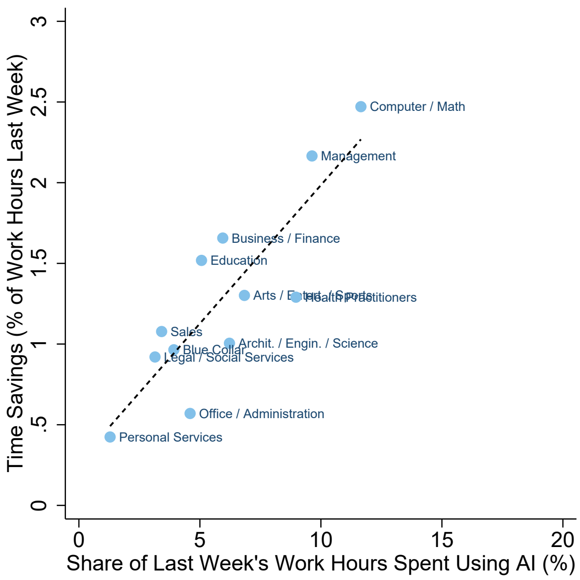

(c) genAI Time Savings and Use by Industry

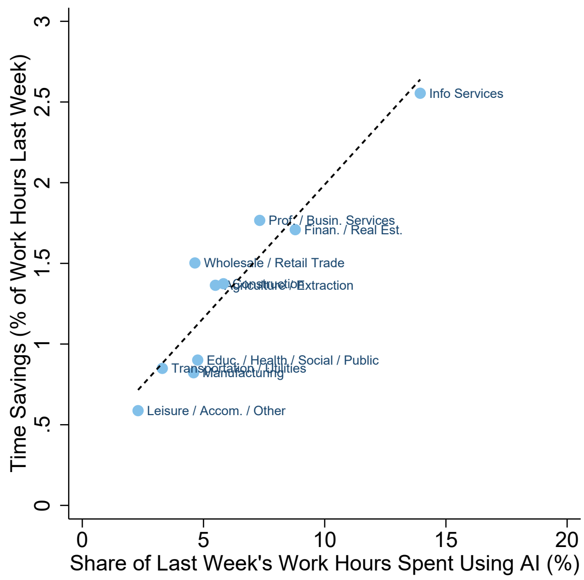

Notes: Panel a shows the distribution of work time saved in the previous week due to genAI, among individuals who use genAI for work. Time savings is broken down into 1 hour or less (lightest blue), 2 hours (light blue), 3 hours (medium blue), and 4 hours or more (dark blue). The "Overall" bar reflects the distribution among all workers who use genAI. The "Used 1 Day Last Week" bar reflects users who used genAI on exactly one workday last week. The "Used More Than 1 Day Last Week" bar reflects users who used genAI on at least two workdays but not all workdays last week. The "Used Every Day Last Week" bar reflects users who used genAI for work every workday last week. Data source is the November 2024 wave of the RPS, ages 18-64 (the August 2024 wave did not ask about time savings). The bottom two panels plot the mean share of last week's work hours spent using genAI against the mean time savings due to genAI by occupation group (Panel b) and industry group (Panel c). The sample is employed respondents who use genAI for work ($N = 933$).

5.2 Reported Time Savings Due to Generative AI

In the November 2024 wave, we asked genAI users to estimate how many additional hours they would have needed to complete the same amount of work they did last week if they had *not* had access to the technology.[18] Figure 11 reports the results from this question. Among workers who used genAI in the previous week, 33.2 percent said that genAI saved them an hour or less of time last week, 26.4 percent reported saving two hours, 20.0 percent reported three hours, and 20.4 percent reported four hours or more. Not surprisingly, more frequent users also reported greater time savings. Among workers who used genAI every day last week, 33.3 percent said it saved them four hours or more, compared to only 11.5 percent of those who used it only one day last week.

For each genAI user, we compute the percent of working hours saved as the ratio of time saved last week to hours worked last week. We find an average time savings of 5.4 percent of work hours in the November 2024 wave. For an individual working 40 hours per week, saving 5.4 percent of work hours implies a time savings of 2.2 hours per week. When we include all workers, including non-users, workers saved 1.4 percent of total hours due to genAI.

There is substantial variation in time savings across occupations and industries. Figures 11b and 11c plot average time savings by industry (including non-users) against hours spent using genAI in the previous week. Time savings and overall usage are highly correlated. Workers in computer, mathematical, and management occupations use genAI between 9 and 12 percent of their work hours, and it saves them between 2.1 and 2.5 percent of their work hours. By contrast, workers in personal service occupations use genAI in only 1.3 percent of their work hours, and it saves them only 0.4 percent of work time. The slope of the dashed OLS line is 0.170, indicating that a 10 percentage point increase in time spent using genAI is associated with an increase in time savings of 1.7 percent.

Across industries, information services has both the largest share of work hours spent using genAI (14.0 percent) and the highest time savings (2.6 percent). Leisure, accommodation, and other services has both the lowest share of work hours spent using genAI (2.3 percent) and the lowest time savings (0.6 percent).

exactly one day last week (6.3 percent of workers) and more than one day last week (7.8 percent of workers). With regards to time per day, in the August wave users could report using genAI (i) more than an hour per day last week, (ii) between 15-59 minutes per day, or (iii) less than 15 minutes per day. In the November wave, we disaggregated (i) into two options: 1-4 hours per day (30.6 percent of users), and more than 4 hours per day (8.1 percent of users).

[18] Specifically, the question is: *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?* For *X* we fill in the respondent's reported actual hours worked last week.

5.3 The Aggregate Productivity Gain from Generative AI

We now present a standard model of aggregate production to illustrate how our data can be used to estimate the increase in aggregate labor productivity due to genAI. We model the economy's aggregate output using a Cobb-Douglas aggregate production function where labor supply is perfectly substitutable across workers. The economy contains N workers. Let Y denote the aggregate output, and let L denote the aggregate supply of labor measured in efficiency units. The aggregate production function is given by:

$$Y = AK^{\alpha}L^{1-\alpha}, \tag{1}$$

where A is TFP, K is the aggregate capital stock, and α is the Cobb-Douglas share on capital. The aggregate supply of labor L is defined as the weighted sum of individual labor supplies:

$$L = \sum_{i=1}^{N} \ell_i e_i \tag{2}$$

where ℓ_i is the number of hours worked per week by worker i, and e_i is the efficiency units of labor supplied by worker i.

In a competitive labor market, each worker is paid their marginal product of labor:

$$w_i = (1-\alpha)A(K/L)^{\alpha}e_i \tag{3}$$

where w_i is the worker's hourly wage. Normalizing the mean efficiency unit to $\overline{e} = 1$, we have:

$$\tilde{w}_i \equiv \frac{w_i}{\overline{w}} = \frac{w_i}{\left(\sum_i w_i\right)/N} = \frac{(1-\alpha)A(K/L)^{\alpha}e_i}{\left(\sum_i (1-\alpha)A(K/L)^{\alpha}e_i\right)/N} = \frac{(1-\alpha)A(K/L)^{\alpha}e_i}{(1-\alpha)A(K/L)^{\alpha}\overline{e}} = e_i \tag{4}$$

where $\tilde{w}_i$ is worker i's wage relative to mean wages $\overline{w}$.

Suppose that worker i saves s_i hours per week due to genAI and spends this time on additional production within their job. The effective weekly labor supply for worker i is then $\ell_i + s_i$. Substituting this into the expression for aggregate labor supply, we have:

$$L' = \sum_i (\ell_i + s_i)e_i = \sum_i \ell_i e_i + \sum_i s_i e_i. \tag{5}$$

The change in aggregate labor supply attributable to genAI-induced time savings is:

$$\Delta L \equiv L' - L = \sum_i s_i e_i = \sum_i s_i \tilde{w}_i \tag{6}$$

Assuming no change in TFP, capital, and hours worked, the percent change in aggregate productivity from genAI equals the percent change in aggregate output. The approximate percent change in aggregate output from genAI is given by:

$$\frac{\Delta Y}{Y} \approx (1-\alpha)\frac{\Delta L}{L} = (1-\alpha)\frac{\sum_i s_i \tilde{w}_i}{\sum_i \ell_i \tilde{w}_i}. \tag{7}$$

This equation states that the percent change in output due to genAI is the ratio of mean time savings to mean hours worked, weighted by worker's wages, and scaled by labor's share of production costs.

Using November 2024 RPS data on hourly wages w_i, weekly hours worked ℓ_i, and weekly genAI time savings s_i, we use (7) to estimate that genAI currently increases aggregate labor productivity by 1.9%.[19] If we assume an AI-exposure-adjusted labor share of 0.57, following Acemoglu (2024), the implied potential productivity gain is 1.1%.

However, labor productivity gains will only translate into aggregate productivity gains if firms adjust their expectations. If workers are now able to complete the same tasks in less time without their employer's knowledge, they may take their time savings as on-the-job leisure, which would increase welfare but not productivity. Bonney et al. (2024) find that only 5.4 percent of firms had formally adopted genAI as of February 2024. While firm adoption may increase over time, it still lags far behind, suggesting that worker adoption is still mostly informal. Thus we emphasize that these *potential* productivity gains from genAI may not show up in productivity statistics, at least for now.[20]

5.4 Comparisons to Micro and Macro Estimates of Productivity Gains from Generative AI

How does our estimate of the aggregate productivity gain from genAI compare with experimental estimates from the literature? To answer this question, we reformulate the model to express time savings as a linear function of time spent using the technology:

$$L' = \sum_i (l_i + \gamma u_i) e_i = \sum_i l_i e_i + \gamma \sum_i u_i e_i \tag{8}$$

In this expression, u_i is the weekly hours spent using genAI by worker i, and γ is the productivity gain associated with one hour of genAI use.

[19] This estimate is slightly higher than the average time savings due to genAI from the previous section (1.4 percent). Intuitively, this is because the average genAI users tend to have above-average wages (see Figure 6).

[20] We also assume above that labor supply is perfectly substitutable across workers. However, imperfect substitution across worker types with differing usage patterns would reduce the impact of genAI on aggregate productivity.

We can then write the percent increase in aggregate labor productivity due to genAI as

$$\frac{\Delta L}{L} = \frac{\gamma \sum_i \tilde{w}_i u_i}{\sum_i \tilde{w}_i l_i} \tag{9}$$

Next, we solve for the value of γ that would generate a 1.9 percent increase in aggregate labor productivity given our estimates of genAI use:

$$\gamma = 1.9\% \cdot \left(\frac{\sum_i \tilde{w}_i l_i}{\sum_i \tilde{w}_i u_i} \right) \tag{10}$$

Recall that for each worker in the RPS data we do not have a single estimate for l_i; instead, we have a lower and upper bound, $\underline{l}_i$ and $\bar{l}_i$, respectively. Evaluating (10) using the midpoint of these bounds, $l_i \equiv (\underline{l}_i + \bar{l}_i)/2$ yields a wage-weighted share of total work hours spent using genAI of 5.7%, which implies a productivity parameter value of $\gamma = 0.33$. This implies that each hour spent using genAI increases the worker's productivity for that hour by 33%. This is similar in magnitude to the average value of 27% from several randomized experiments of genAI usage (Cui et al., 2024; Dell'Acqua et al., 2023; Noy and Zhang, 2023; Peng, Kalliamvakou, Cihon, and Demirer, 2023).

Our estimate of the potential aggregate productivity gain from genAI of 1.1% is slightly higher than the estimate of 0.7% found by Acemoglu (2024). Our estimate is based on reported genAI use and time savings, while Acemoglu (2024) uses experimental evidence on productivity gains from genAI usage and task-based predictions of genAI exposure. If we adjusted our estimate of labor cost savings from 33% down to the 27% from the experimental studies discussed in Acemoglu (2024), we would arrive at an aggregate productivity estimate of 0.9%. The remaining 0.2 percentage point gap is explained by different estimates of genAI adoption. Acemoglu (2024) uses expert assessments to estimate a wage-weighted share of all tasks for which genAI can be adopted of 4.6%. We use survey data on actual adoption to estimate a wage-weighted usage share of 5.7%, which yields slightly higher estimates. A final difference is that Acemoglu (2024) assumes that these productivity gains would evolve gradually over the next ten years, whereas our estimates are based on current usage.

6 Conclusion

Generative AI is a potentially important new technology. Its economic impact will ultimately be determined by how many people use it, how intensively they use it, and the productivity gains it provides. This paper documents results from the first nationally representative U.S. survey of genAI 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.

As of August and November of 2024, 39 percent of working-age adults reported using genAI, either for work or personal use. 9 percent of employed respondents used it every work day and 14 percent on some but not all workdays, bringing total weekly usage to 23 percent. ChatGPT is used most frequently, followed by Gemini, and our estimates of ChatGPT usage are similar to other recent surveys. We also validate our findings in a different survey panel, the Survey of Working Arrangements and Attitudes (SWAA).

We compare the speed of genAI adoption with PCs and the internet relative to the release date of the first mass market product for each technology. We find that overall adoption of genAI has been more rapid than these previous technologies. While work adoption rates were similar for genAI and PCs, adoption rates outside of work were much faster for genAI. Low adoption costs and the consumer focus of genAI products are a likely reason for relatively fast adoption rates, and the longer-run impacts of genAI remain uncertain. Still, it is notable that the speed of genAI adoption has been similar to the adoption of personal computers in the 1980s, since the latter led to an accelerate of productivity growth and rising inequality (e.g. Autor, Levy, and Murnane 2003).

We also find substantial variation in genAI adoption by education, age, gender, and occupation. genAI and computer adoption have similar demographic patterns, with greater early usage among younger, more educated workers in professional occupations. However, women were earlier users of PCs, in part because of rapid adoption of computers in white-collar clerical occupations. Interestingly, there is a strong correlation between adoption patterns by occupation and expert assessments of task-based AI exposure, which supports the value of these estimates for researchers studying the labor market impacts of AI (e.g. Eloundou, Manning, Mishkin, and Rock 2024). Workers report that genAI helps in a broad range of tasks, the most useful being writing, administrative support, and interpreting and summarizing information.

Combining detailed data on the frequency and intensity of genAI usage, we estimate that between 1 and 5 percent of total work hours in the previous week were spent using genAI. We also ask users about how much time they save by using genAI and estimate a mean time savings of 5.4 percent of all hours among users and 1.4 percent including non-users. Combining these figures with a standard model of aggregate production suggests a potential productivity gain from genAI adoption of around 1.1%. However, we caution that this assumes all of the gains are internalized as higher worker output rather than increased leisure.

Our findings suggest many directions for future work. It will be important to continue tracking the scale, intensity, and heterogeneity of genAI adoption over time. Another important question is to investigate the detailed tasks for which workers use genAI.

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. (2024). *The Simple Macroeconomics of AI.* Tech. rep. National Bureau of Economic Research.

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.

Albanesi, S., A. Dias da Silva, J. F. Jimeno, A. Lamo, and A. Wabitsch (Nov. 2024). *New Technologies and Jobs in Europe.* Tech. rep. OIGI Working Paper.

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

Autor, D. H., L. F. Katz, and A. B. Krueger (1998). “Computing inequality: have computers changed the labor market?” In: *The Quarterly journal of economics* 113.4, pp. 1169–1213.

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.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). “Artificial intelligence, firm growth, and product innovation”. In: *Journal of Financial Economics* 151, p. 103745.

Barrero, J. M., N. Bloom, and S. J. Davis (2021). *Why Working From Home Will Stick.* Working Paper. National Bureau of Economic Research.

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 (2024a). *Heterogeneity in Work From Home: Evidence from Six U.S. Datasets.* Working Paper. Federal Reserve Bank of St. Louis and Vanderbilt University.

Bick, A., A. Blandin, A. Caplan, and T. Caplan (2024b). *Measuring Trends in Work From Home: Evidence from Six U.S. Datasets.* Working Paper. Federal Reserve Bank of St. Louis and Vanderbilt University.

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.* Working Paper. Federal Reserve Bank of St. Louis and Vanderbilt University.

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.

Card, D. and J. E. DiNardo (2002). “Skill-biased technological change and rising wage inequality: Some problems and puzzles”. In: *Journal of labor economics* 20.4, pp. 733–783.

Chui, M., E. Hazan, R. Roberts, A. Singla, and K. Smaje (2023). “The economic potential of generative AI”. In.

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, D. and L. B. Kahn (2018). "Skill requirements across firms and labor markets: Evidence from job postings for professionals". In: *Journal of Labor Economics* 36.S1, S337–S369.

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.

Eisfeldt, A. L., G. Schubert, B. Taska, and M. B. Zhang (Mar. 2024). *Generative AI and Firm Values*. Tech. rep. SSRN Working Paper.

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.

Filippucci, F., P. Gal, and M. Schief (2024). "Miracle or Myth? Assessing the macroeconomic productivity gains from Artificial Intelligence". In.

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

Gathmann, C., F. Grimm, and E. Winkler (2024). *AI, Task Changes in Jobs, and Worker Reallocation*. Tech. rep. 11585. CESifo Working Papers.

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

Krueger, A. B. (1993). "How computers have changed the wage structure: evidence from microdata, 1984–1989". In: *The Quarterly Journal of Economics* 108.1, pp. 33–60.

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.

Levine, A. (2024). "Consumers Are Diving Into AI, New Data Show. Companies Are Being More Cautious." In: *Barron's*. Accessed: 2025-01-24.

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

Nelson, R. R. and E. S. Phelps (1966). "Investment in humans, technological diffusion, and economic growth". In: *The American economic review* 56.1/2, pp. 69–75.

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.

Toner-Rodgers, A. (2024). "Artificial intelligence, scientific discovery, and product innovation". In: *arXiv preprint arXiv:2412.17866*.

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