Large Language Models, Small Labor Market Effects*

Anders Humlum† Emilie Vestergaard‡

July 15, 2025

Abstract

We examine the early labor market impacts of AI chatbots using two large-scale adoption surveys of 11 exposed occupations, linked to Danish administrative data. Chatbots are now widespread: most employers encourage use, many deploy enterprise solutions, and training is common. These initiatives boost adoption, narrow gender gaps, and create new tasks. Yet, impacts on earnings and recorded hours remain minimal. Using difference-in-differences, we estimate precise null effects at both the individual and workplace levels, with confidence intervals ruling out effects above 1%. While adoption is associated with occupational switching and task restructuring, these shifts do not result in net changes to hours or earnings. Modest productivity gains (3% time savings) and the emergence of integration and oversight tasks help explain these muted effects. Our findings challenge narratives of imminent disruption from Generative AI, highlighting instead the slow burn of productivity J-curves.

JEL Codes: J23, J24, J31, O33

*We thank David Autor, Marianne Bertrand, Jacob Conway, Matthew Notowidigdo, Chad Syverson, and seminar participants at the Becker Friedman Institute, Chicago Booth, Chicago Fed, INSEAD, Microsoft Research, MIT Sloan, and the SED 2025 Meeting for helpful comments and suggestions. This project received financial support from the Center for Applied Artificial Intelligence and the Polsky Center for Entrepreneurship and Innovation. Caspar Ringhof provided excellent research assistance.

†University of Chicago, Booth School of Business, and NBER, anders.humlum@chicagobooth.edu

‡University of Copenhagen, Department of Economics, emilievestergaard@econ.ku.dk

The emergence of AI chatbots marks the rise of Generative Artificial Intelligence (AI). By some measures, these technologies are already living up to the immense hype: AI chatbots have seen the fastest worker take-up of any new technology (Bick, Blandin and Deming, 2025; Humlum and Vestergaard, 2025), randomized controlled trials (RCTs) demonstrate substantial productivity gains for users (Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023), and case studies indicate notable effects on online labor market platforms (Teutloff et al., 2025). These effects on productivity and platform demand are remarkable in both magnitude—ranging from 15% to 50%—and speed, materializing within a few months.

However, the broader labor market implications of Generative AI remain unclear for at least three reasons. First, while workers have embraced AI chatbots for their low costs and ease of use, it is unclear whether employer initiatives mainly aim to enhance the tools' productivity or prevent damage from misusage (Humlum and Vestergaard, 2025). Second, RCTs show that the effects of AI chatbots can turn negative if applied to the wrong tasks (Dell'Acqua et al., 2023; Otis et al., 2024*b*), raising caution in extrapolating effects from controlled settings to the broader economy. Third, while studies have documented effects on productivity, it remains unclear how these translate into earnings and hours, as high-quality microdata on such outcomes are rarely available.

This paper addresses these gaps by conducting two large-scale surveys on AI chatbot adoption and linking the responses to administrative labor market data in Denmark. Our dataset includes two survey rounds (late 2023 and 2024), each covering about 25,000 workers from 7,000 workplaces across 11 occupations exposed to AI chatbots, linked to monthly panel data on earnings, hours, and occupations.[1] Our analysis proceeds in four parts.

[1]Our list of occupations is accountants, customer support specialists, financial advisors, HR professionals, IT support specialists, journalists, legal professionals, marketing professionals, office clerks, software developers, and teachers.

We begin by documenting that employers are now heavily invested in AI chatbots: most encourage their use, 38% deploy enterprise versions, and 30% of employees have received training. Notably, enterprise chatbots and training are present even in workplaces that do not encourage usage, indicating that some initiatives may be intended more to mitigate misuse than to enhance productivity. More broadly, many otherwise similar firms have adopted markedly different strategies toward AI chatbots.

Second, we examine how these employer chatbot strategies influence workers' adoption of the tools. Encouraged-use policies, enterprise chatbots, and training each substantially increase take-up, nearly doubling adoption rates, with encouraged use particularly effective at promoting intensive usage. Notably, these initiatives also help narrow demographic gaps in adoption, with training playing a central role: the gender gap in take-up shrinks from 12 to 1 percentage points when employers implement the full suite of initiatives—encouraging use, deploying enterprise chatbots, and providing training.

Next, we examine how AI chatbots affect work processes, revealing a stark contrast across employer initiatives: while encouraged-use policies are consistently associated with greater reported productivity gains and a rise in productive AI tasks—such as content drafting and ideation—training and enterprise chatbots, when implemented in isolation, are linked to lower reported time savings and an increase in compliance and oversight tasks. This contrast reinforces the notion that some employer initiatives remain more focused on preventing misuse than on enhancing the "out-of-the-box" productivity of AI chatbots. Notably, the highest reported benefits arise when all three initiatives—encouraged use, enterprise chatbots, and training—are implemented together. The fact that initiatives can be counterproductive in isolation but complementary in combination echoes the idea of "Productivity J-curves," in which complementary firm-level investments are critical to unlocking the gains from new technologies (Brynjolfsson, Rock and Syverson, 2021; David, 1990).

Taking stock of our survey evidence, we estimate that the overall productivity gains from AI chatbots are modest: combining users' reported benefits with their usage frequency, average time savings amount to 2.8% of work hours. This contrasts with the significant productivity gains—often exceeding 15%—documented by RCTs in our study occupations. Two factors help explain this discrepancy. First, effect heterogeneity: while existing RCTs concentrate on occupations with the largest productivity gains from AI chatbots, our broader survey reveals that several exposed occupations see more modest gains. Second, the lack of employer encouragement and complementary investments: many real-world workers do not operate under the same favorable conditions as those in experimental settings, limiting realized productivity gains.

In the final and central part of the paper, we examine the effects of AI chatbots on labor market outcomes. To do so, we link our survey data to administrative records on monthly earnings, hours worked, and occupations through June 2024—one and a half years after the launch of ChatGPT. Using a difference-in-differences framework, we compare adopters and non-adopters before and after the arrival of AI chatbots.

We begin by analyzing outcomes for workers who use AI chatbots. If these tools enhance individual productivity, a natural question is whether such gains translate into higher earnings. Our first key finding is that AI chatbots have had minimal impact on workers who use them. Difference-in-differences estimates for earnings, hours, and wages are all precisely estimated zeros, with confidence intervals ruling out average effects larger than 1%. At the occupation level, estimates are similarly close to zero, generally excluding changes greater than 5%. Moreover, we find no evidence of differential trends over time, nor larger effects among early adopters, frequent users, those reporting greater benefits, those supported by employer initiatives, or in occupations with flexible pay systems. Together, these findings suggest that the limited earnings effects of chatbot use may not be merely a short-run phenomenon.

Next, we examine the job mobility of adopters. A consistent finding from experimental studies is that AI chatbots offer greater productivity benefits to workers with less prior experience in their occupations (Brynjolfsson, Li and Raymond, 2025; Dell'Acqua et al., 2023; Noy and Zhang, 2023). In line with this, we find that chatbot adoption is closely linked to occupational switching. Our results support the thesis of Autor (2024), who argues that Generative AI may help individuals extend their work into new occupations. Just one and a half years after the arrival of these tools, our administrative data show that such dynamics are already unfolding in the labor market. However, the effects on occupation switching are modest—amounting to 2% of full-time equivalent employment—and not large enough to generate net gains in hours or earnings for adopters.

Finally, we zoom out to the workplace level and examine whether firms that encourage and invest in AI chatbot use—the "AI front-runners"—have experienced differential changes in overall labor demand. If chatbots reduce the need for labor, they could lead to slower hiring or even layoffs. Conversely, if they enhance productivity, firms may expand hiring. Despite these possibilities, we find no differential changes in employment or wage bills among workplaces that encourage chatbot use. These workplaces also show no shifts in job creation or destruction, nor changes in the composition of hires and separations toward workers more prone to using chatbots. These null results on workplace outcomes hold across occupations and persist even in workplaces that combine encouragement with enterprise chatbot solutions and training. Taken together, our evidence suggests that while internal reorganization is underway—reflected in emerging new tasks—these changes have not yet translated into observable differences in labor market outcomes.

As a complement to our difference-in-differences analysis of labor market impacts, we directly asked workers: *"Have AI chatbots affected your labor earnings?"* Tellingly, nearly all respondents—97.7% of adopters and 99.5% of non-adopters—reported that AI chatbots have had *no* impact on their earnings. These worker perceptions further reinforce

our conclusion that AI chatbots have had limited impacts on labor markets to date.

Overall, our findings challenge narratives of imminent labor market disruption due to Generative AI. Although worker adoption has been rapid, many organizations remain in the trough of the productivity J-curve, devoting substantial resources to integration, training, compliance, and oversight. Partly as a result, productivity gains remain modest, and labor market effects limited. Yet, even among AI front-runners—workplaces that have adopted the full suite of proactive chatbot initiatives or workers who use the tools daily and report large benefits—chatbots have not translated into measurable labor market outcomes. This suggests the limited impacts we estimate may persist.

1 Data and Institutional Setting

Denmark offers an ideal setting for examining the labor market impacts of Generative AI.

First, Danish workers have been at the forefront of Generative AI adoption, with take-up rates comparable to those in the United States (Bick, Blandin and Deming, 2025; Humlum and Vestergaard, 2025; RISJ, 2024).

Second, Denmark's labor market is highly flexible, with low hiring and firing costs and decentralized wage bargaining—similar to that of the U.S.—which allows firms and workers to adjust hours and earnings in response to technological change (Botero et al., 2004; Dahl, Le Maire and Munch, 2013). Most workers in our sample engage in annual negotiations with their employers, providing regular opportunities to adjust earnings and hours in response to AI chatbot adoption during the study period. Appendix B.1 documents the wage-setting systems in the occupations we study, revealing substantial variation in exposure to centralized bargaining. Section 4.2.4 leverages this heterogeneity to examine how wage-setting institutions shape our labor market impact estimates.

Third, Denmark has exceptional infrastructure for tracking the adoption of new technologies. In particular, every Dane has a digital mailbox that Statistics Denmark can

use to distribute survey invitations. We use this infrastructure to conduct two large-scale, representative surveys on AI chatbot adoption, which we detail below.

Finally, our partnership with Statistics Denmark allows us to link these surveys to administrative matched employer-employee data, providing a unique opportunity to analyze labor market effects such as changes in earnings, working hours, and job mobility.

Taken together, these factors make Denmark a prime setting for observing early labor market effects of Generative AI, with insights that may extend to other advanced economies—including the US—and an unparalleled data infrastructure to assess these impacts rigorously.

1.1 Data Sources

This paper builds on two large-scale surveys on AI chatbot adoption, conducted in November–December of 2023 and 2024. The first survey provided the dataset for Humlum and Vestergaard (2025), which documented adoption patterns of ChatGPT, the dominant AI chatbot. This paper extends that dataset in two key ways. First, we link survey responses to administrative labor market data after the introduction of ChatGPT, enabling us to assess impacts on earnings, work hours, and job mobility. Second, we introduce a second survey round in 2024 that (*i*) broadens the scope to include all AI chatbots, including custom versions, (*ii*) provides extensive data on firm-led adoption initiatives, and (*iii*) offers deeper insights into workers' actual usage and perceived benefits of these tools. Section 1.3 outlines the 2024 survey, which is the primary focus of this paper.

1.2 Occupations

Our surveys focus on 11 occupations that are highly exposed to AI chatbots: *accountants, customer support specialists, financial advisors, HR professionals, IT support specialists, journalists, legal professionals, marketing professionals, office clerks, software developers, and teachers*. These occupations were selected based on three criteria: (*i*) they have

at least one O*NET job task where AI chatbots can save time, as measured by the "Direct Exposure (E1)" metric from Eloundou et al. (2024); (*ii*) they are captured by a well-defined set of ISCO codes; and (*iii*) they contain a sufficient number of workers for statistical analysis. Humlum and Vestergaard (2025) details the selection and empirical measurement of these occupations. Together, these 11 occupations represent the labor markets where the impacts of AI chatbots are likely to emerge first and most strongly. Our analysis centers on this set.

1.3 Survey Outline

Our 2024 survey is organized into the four blocks summarized below. The 2023 round followed a similar structure. The full questionnaires are in Appendix I.

Block 1: Occupation and tasks. Workers first select their occupation and report the importance of six representative tasks in their occupations.

Block 2: Adoption. Workers report their experiences with various AI chatbots, including the domains, frequency, and duration of usage.

Block 3: Employer initiatives. Workers are asked about any employer initiatives related to AI chatbots, including usage policies, enterprise chatbots, and employee training.

Block 4: Impact on work. Workers are asked about their experienced benefits and estimated effects of AI chatbots.

1.4 Register Data

We use several administrative registers at Statistics Denmark. Our matched employer-employee data come from the *E-Income Register,* which records earnings, hours, occupation, and industry for all job spells in Denmark on a monthly basis from 2008 onward. This

register is compiled by the Danish tax authorities and subsequently harmonized by Statistics Denmark into the *Employment Statistics of Employees* (BFL) dataset. We complement this with demographics data on individuals from the *Population Register* (BEF), and wealth information for the *Personal Wealth Register* (FORMPERS). Finally, we draw firm financial data (e.g., annual revenue and employment) from the *Firm Statistics* (FIRM) Register. Because our survey was sent to workers identified in the register data by their (deidentified) social security numbers (*pnr*), all respondents can be matched to the register data. While labor market data in the BFL register are updated continuously, we only have firm financial data up to 2022. For that reason, we focus our effect analyses on labor market outcomes.

1.5 Survey Sample

We invited 115,000 workers to participate in each of our survey rounds in 2023 and 2024. The registers at Statistics Denmark enabled us to target these invitations by occupation and workplace. In particular, for each of our 11 occupations, we conducted a workplace-based sampling procedure, first drawing a random set of workplaces within the occupations and then sampling all relevant workers for these workplaces. This sampling procedure maximizes the statistical power of the workplace-level analyses whilst keeping our sample representative. Appendix B.2 details the sampling protocol. We sent three reminders per survey round, two by e-mail and one by text. The invitation letters are in Appendix H. We received about 25,000 valid and complete responses to each of the survey. Appendix B.3 details our survey response rates. While our main analysis focuses on responses from the 2024 round, we use the 2023 round to examine the dynamics of our estimated effects.

1.5.1 Representativeness and Response Quality

Appendix B.3.1 conducts several checks on the representativeness and quality of our survey responses. These analyses extend the checks in Humlum and Vestergaard (2025) to the 2024 survey round. First, we ensure that our sample represents the population based on observables, including age, gender, experience, earnings, and wealth. Second, following Dutz et al. (2025), we use randomized participation incentives to show that our findings are also balanced on workers' latent willingness to participate in the survey. Finally, we cross-check that the survey responses align with variables that are also recorded in the administrative registers.

2 Adoption

2.1 Employer Initiatives

Figure 1, Panel (a) shows the prevalence of employer policies related to AI chatbot usage across our 11 occupations. Firms are now heavily invested in AI chatbots: about 43% of workers are explicitly encouraged to use them, another 21% are allowed, while only about 6% are explicitly prohibited from doing so. This marks a shift from early responses to ChatGPT, when many employers restricted its use due to concerns over data confidentiality and output accuracy (Humlum and Vestergaard, 2025).[2]

Employers not only set usage policies—they also put their money where their mouth is. Panel (b) shows that 38% of firms have an enterprise AI chatbot solution—most often customized versions—and Panel (c) shows that 30% of employees have participated in training courses on AI chatbot usage, most often organized by their employer.[3] Notably, these

[2]Although outright bans are now rare, they persist in some occupations that involve sensitive data (e.g., financial advisors) or require high factual accuracy (e.g., legal professionals).

[3]Customized chatbots are typically versions of ChatGPT, tailored for workplace use through adjustments such as compliance with data security, access to internal data, and custom prompt creation. Often, these are thin wrappers around ChatGPT Enterprise or the OpenAI API. Appendix C.2.1 details the usage of AI chatbot products, showing that ChatGPT remains the dominant player. It also provides additional information on the characteristics of enterprise chatbots.

firm-led investments are prevalent in all 11 occupations but are particularly widespread in journalism and marketing and more limited in teaching. The widespread investments in AI chatbots at the workplace are consistent with their potential as a general-purpose technology (Eloundou et al., 2024).

Which firms have adopted different AI chatbot initiatives? Table 1 shows that firms encouraging AI chatbot use tend to be slightly younger (each additional 10 years of age is associated with a 1.4 percentage point decrease in encouragement), more productive (a doubling of productivity corresponds to a 4.7 percentage point increase), and more likely to be privately owned (associated with a 4.1 percentage point increase), but not systematically larger than others.[4,5] Younger and more productive firms are also somewhat more likely to purchase enterprise chatbots and offer training. While firm size is unrelated to encouraged-use policies and employer-provided training—perhaps due to the low fixed costs of off-the-shelf chatbot solutions—larger firms are more likely to purchase enterprise versions.

Overall, however, chatbot initiatives are largely unrelated to observable firm characteristics, which account for just 1.3% of the within-occupation variation in encouragement—far less than the 10.6% explained by our 11 occupation categories alone. In other words, many otherwise comparable firms have adopted markedly different strategies toward AI chatbots, even for similar workers within the same occupation.

Figure 2 illustrates the chatbot packages adopted by different employers. Among those that encourage chatbot use, 61% have an enterprise chatbot solution, 39% of employees have received training, and 29% have both. Notably, enterprise chatbots and training are present even in workplaces that do not encourage usage: 21% have an enterprise

[4]Appendix Table C.1 shows that these correlations with firms' characteristics are robust to controlling for worker characteristics.

[5]Younger and more productive firms are commonly found to adopt new technologies faster; see, e.g., Acemoglu et al. (2022), who document adoption patterns for a range of advanced technologies in 2019, including AI and robotics.

solution, 11% provide training, and 4% do both. The fact that employers invest in these tools without encouraging use suggests some initiatives may be aimed more at mitigating misuse than enhancing productivity.

2.1.1 Employer Initiatives and Worker Outcomes

Throughout the paper, we examine how workers' adoption of and benefits from AI chatbots vary with employer-led chatbot initiatives. This analysis is motivated by the literature on "Productivity J-Curves," which emphasizes that complementary organizational investments may be critical for realizing the gains from new technologies (Brynjolfsson and Hitt, 2000; Brynjolfsson, Rock and Syverson, 2021). Moreover, because early adoption of AI chatbots was largely worker-driven, this analysis sheds light on the extent to which employers can move the needle and influence outcomes by also investing in the technology.

Specifically, we compare workers in the same occupations with similar characteristics and estimate how their adoption behaviors and reported benefits vary with employer initiatives for AI chatbots:

$$Y_i = \gamma' X_i + \beta' \text{EmployerInitiatives}_i + \varepsilon_i, \tag{1}$$

where Y_i is a chatbot-related outcome for worker i, such as use or reported benefits of AI chatbots for work; X_i is a vector of worker pre-determined characteristics, including age, gender, experience, and occupation fixed effects;[6] and $\text{EmployerInitiative}_i$ is a vector of employer-driven initiatives, including encouraging employees to use AI chatbots, investing in enterprise chatbot solutions, providing employee training— and combinations thereof. We estimate Equation (1) by OLS and cluster the standard errors by workplace-occupation level. Appendix A.1 presents a Roy-style model of chatbot adoption to help interpret the estimates in Equation (1). We draw on this framework in the discussions that follow.

Although employer initiatives are not assigned through a randomized design, our rich

[6] All worker characteristics are measured in 2022, before AI chatbots became available.

data on firms and workers allow us to examine the robustness of our findings to several potential confounders. Appendix F.1 presents these robustness checks. First, we show that all results hold when controlling for firm characteristics (Table 1), ensuring that differences across employer initiatives are not driven by variation in firm age, size, or productivity. Second, we show that the results are robust to controlling for workers' detailed task mixes within occupations, demonstrating that the effects of employer initiatives are not driven by differences in the types of tasks more amenable to AI chatbot use.[7] Finally, in our panel analyses in Section 4, we verify that adopters of different chatbot initiatives followed similar labor market trajectories prior to the arrival of AI chatbots. Together, these robustness checks help rule out confounding factors and support the validity of our estimated relationships.

Finally, since our data are collected from workers rather than senior leadership, our estimates reflect employer initiatives *as perceived by workers*. Since employer policies may not be perfectly communicated, worker self-reports may not fully align with employer intentions. Nonetheless, Appendix F.2 shows that our findings are largely robust to using coworker reports to measure employer initiatives, demonstrating these may be interpreted as workplace-wide treatments.

2.2 Worker Adoption

Figure 3 shows how AI chatbot adoption varies with employer initiatives for a typical worker. The top estimate (*No Initiative*) highlights the bottom-up nature of adoption: even among workers whose employers neither encourage chatbot use, provide an enterprise chatbot, nor offer training, about 41% have used AI chatbots at work at some point, and 7% use them daily.

Despite this high baseline, employer initiatives are associated with substantially higher

[7]We do not include task controls in our main specification, as task mix may itself be endogenous to chatbot adoption. Indeed, Section 3.2 shows that AI chatbots have created new job tasks for many workers.

adoption. Implemented individually, *Encouraged use*, *Own chatbot*, and *Training* each raise extensive-margin take-up rates to 76%, 57%, and 69%, respectively. Encouraged-use policies are especially effective at increasing intensive usage, raising daily use from 7% to 18%. In contrast, providing training or an enterprise chatbot without encouragement leads to only modest increases in daily use, to 9% and 8%, respectively.

The bottom rows show adoption rates when employer initiatives are combined. Adoption peaks when all three—encouragement, enterprise chatbots, and training—are in place: 93% of workers in such settings report having used AI chatbots at work, and 28% use them daily.

Employer initiatives not only boost adoption—they also influence *who* adopts the tools. Figure 4 examines how the gender gap in AI chatbot adoption varies with employer initiatives, controlling for workers' occupation, age, experience, earnings, and wealth.

The *No Initiative* estimate shows that a substantial gender gap has emerged in chatbot adoption: without encouraged use, enterprise chatbots, or training, women are about 13 percentage points less likely than comparable men in the same occupation to have used AI chatbots.[8]

Notably, employer initiatives can significantly reduce this gap, with training proving especially effective. For instance, providing training alone narrows the gap to 7 percentage points, and combining training with other initiatives reduces it further, to just 2–3 percentage points.

Appendix Table C.3 analyzes how employer initiatives affect other demographic gaps in adoption, and Appendix Figure C.6 shows gender gaps in daily usage. Employer initiatives also help close gaps by age, experience, and earnings. Training, in particular, is

[8]Appendix Figure C.4 shows that the gender gap in chatbot adoption persists even in workplaces where the tools are prohibited. Although employer bans are associated with a reduction in overall adoption to roughly one-third of the baseline rate, about 80% of the baseline gender gap remains. This pattern mirrors findings by Carvajal, Franco and Isaksson (2024), who document a similar gender gap in ChatGPT use among Norwegian bachelor's students and show, using a vignette experiment, that the gap largely reflects male students continuing to use the tools even when explicitly banned.

consistently effective in narrowing gaps in daily use as well. By contrast, encourage-use policies alone *widen* the gender gap in daily usage—from 3 to 10 percentage points. However, when encouragement is paired with training, the gap reverses: women become *more* likely than men to use AI chatbots daily, with a gender gap of –2 percentage points.

Comparison to the Literature. The widespread adoption of AI chatbots—and the substantial gender gaps in usage—are well documented in the literature.[9] We contribute to this evidence by documenting that employer initiatives for AI chatbot adoption are now common and play an important role in shaping overall adoption rates and demographic disparities in usage.[10]

The importance of employer policies in driving adoption aligns with Dillon et al. (2025), who show that firm effects are a key determinant of workers' uptake of a randomized offer for Microsoft Copilot. The importance of training in closing gender gaps in adoption is consistent with our 2023 survey round, in which women were significantly more likely to cite lack of training as a barrier to using ChatGPT (Humlum and Vestergaard, 2025).

More broadly, the importance of employer initiatives for AI chatbot adoption echoes an extensive literature on the diffusion of general-purpose technologies, which stresses the need for complementary investments and coinvention (Bresnahan and Greenstein, 1996; David, 1990; Feigenbaum and Gross, 2024). In the sections that follow, we examine how employer investments in AI chatbots shape the tools' effects on work processes and labor market outcomes.

[9]For example, our 2023 survey found that 39.8% of workers had used ChatGPT for work, with adoption rates 15.9 percentage points lower among women within occupations (Humlum and Vestergaard, 2025) (see Appendix C.2.2 for a comparison of the 2023 and 2024 adoption patterns). Higher adoption rates among younger, less experienced, and higher-earning workers also replicate the patterns observed in our earlier round. Bick, Blandin and Deming (2025) report similar adoption patterns among U.S. workers, and Otis et al. (2024*a*) show that gender gaps in Generative AI adoption are pervasive across contexts.

[10]By contrast, earlier surveys found relatively low take-up rates of AI at the firm level. For example, Bonney et al. (2024) document that as of February 2024, only 5.4% of U.S. firms reported using AI, with adoption in the leading sector (Information) still below 20%.

3 Work

3.1 Benefits for Users

What benefits do users report from AI chatbots? Figure 5 starts with workers' reported time savings from using AI chatbots, showing how these vary by employer initiatives.

The *No Initiative* estimates show that without employer encouragement, enterprise chatbots, or training, about 69% of users report any time savings, 46% report saving more than 15 minutes per day of usage, and 14% report saving over an hour on days they use AI chatbots.

The next three rows reveal a stark contrast: while encouraged-use policies are associated with consistently greater time savings, enterprise chatbots and training—when implemented in isolation—are linked to *lower* average time savings among users. For example, the share of users reporting more than 15 minutes of daily time savings increases from 46% to 57% when employers encourage usage, but drops to around 42% when employers provide a custom chatbot or training without also encouraging use.

The Roy adoption model in Appendix A.1 offers two explanations for why enterprise chatbots and training, when used in isolation, may be associated with lower reported benefits. First, these initiatives might reduce the time saved for all users. For instance, if employers implement training or invest in proprietary chatbots primarily to impose guardrails or prevent misuse, such constraints may also limit workers' ability to fully exploit the tools' productivity potential. Second, the lower average benefit may reflect compositional effects: training and enterprise chatbots may expand adoption to workers who inherently gain less from the technology. By contrast, the fact that encouraged-use policies increase both adoption and average reported benefits suggests that the common productivity gains are large enough to offset any negative selection from broader

adoption.[11,12]

The bottom four rows of Figure 5 present reported time savings under combinations of employer initiatives. Notably, average reported benefits are lowest in workplaces that offer both an enterprise chatbot and training but do not encourage usage—further suggesting that these initiatives may not be aimed at enhancing productivity. In contrast, the highest reported benefits are observed when all three initiatives—encouragement, custom chatbots, and training—are implemented together.

The fact that initiatives can be counterproductive in isolation, yet complementary in combination, aligns with the concept of "Productivity J-curves," where firm-based complementary investments are needed to unlock the potential of new technologies (Brynjolfsson, Rock and Syverson, 2021; David, 1990).

While time savings is the most commonly cited benefit, it captures only one dimension of the benefits that AI chatbots offer users: Nearly half of users say the chatbots enhance their creativity or improve the quality of their output, while about 21% report increased job satisfaction.[13] Table 2 examines how these other benefits—creativity, quality, and job satisfaction—vary across employer initiatives. The estimates show that the pattern observed in Figure 5 holds across these dimensions: while encouraged-use policies are associated with higher reported benefits, training and custom chatbots—when implemented

[11]Disentangling the individual-level causal effects from the selection effects would require additional assumptions (e.g., on the functional form of the adoption model à la Heckman (1979)) or alternative sources of variation (e.g., an instrument that shifts adoption costs without affecting perceived benefits). We do not attempt to separate these effects, as we view the bundled impact of employer initiatives as informative in itself, highlighting the average benefits of AI chatbots that may arise under these different initiatives.

[12]A less rosy interpretation of the benefits associated with encouraged-use policies in isolation is that these workplaces have adopted a "move fast and break things" approach to chatbots, prioritizing rapid deployment even at the risk of errors. In contrast, workplaces that offer training or enterprise chatbots may pursue a slower but potentially more sustainable path. If, women are more sensitive to such risks, this interpretation may also help explain why encouraged-use policies alone widen the gender gap in daily use (see Table C.3).

[13]Table D.1 breaks down these benefits by occupation, revealing significant variation. For instance, marketing professionals are more than twice as likely as teachers to report that chatbots improve work quality (69.8% vs. 32.1%), while software developers are more than twice as likely as journalists to report increased job satisfaction from chatbot use (30.5% vs. 12.6%). Despite these differences, AI chatbots help save time across all exposed occupations, with time savings reported by 64%–90% of users.

in isolation—are consistently linked to lower reported benefits.[14]

3.2 Workloads and Task Creation

AI chatbots not only affect existing tasks—they also transform the nature of work itself. Figure 7 shows how the emergence of new workloads from AI chatbots vary across employer initiatives. These new workloads may stem from individual workers taking on more work,[15] from shifts in task allocation within teams, or from new demands generated by the tools themselves.

The *No Initiative* estimates reveal that even in the absence of encouraged-use policies, enterprise solutions, or training, AI chatbots have generated new workloads for 12% of users (Panel (a)). Of these, 3 percentage points report doing more of the same tasks, 7 percentage points have taken on entirely new tasks, and 1 percentage point report doing both.

Notably, the share of users performing new tasks increases to roughly 17% in workplaces with active chatbot initiatives. Training, in particular, is strongly associated with new job responsibilities, suggesting that these programs may be designed to help workers handle new responsibilities introduced by AI tools.

The importance of task creation for the impact of AI chatbots on workloads supports key theoretical predictions for how automation technologies may reinstate the demand for labor in the production process (Acemoglu and Restrepo, 2019; Autor et al., 2024). To further understand how AI chatbots affect the nature of work, we asked respondents to describe the new job tasks they perform as a result of AI chatbots. In Appendix D.2, we

[14]Training and enterprise chatbots display important positive interaction effects with other initiatives, however, particularly for quality and creativity. This suggests that even if these initiatives do not help workers save time on specific tasks, they may still support higher-quality output (e.g., ensuring alignment with company standards) and enhance creative capabilities (e.g., enabling users to accomplish tasks they otherwise couldn't), especially when combined with other initiatives.

[15]Table D.2 shows that the vast majority (80%) of chatbot users say redirect time savings from AI chatbots to other job tasks. By contrast, fewer than 10% report taking additional breaks or leisure, and about 25% spend more time on the same tasks they initially saved time on.

categorize the free-text responses into common new tasks associated with AI chatbots in each occupation. AI chatbots have generated new job tasks for workers across all 11 occupations, with 50% to 95% of these directly linked to AI use.

Figure 7 breaks down the composition of AI-related job tasks by occupation and task category. Appendix D.2 provides example tasks for each pair. Notably, most new work revolves around integrating and overseeing the use of AI chatbots. Specifically, only 41% of new tasks focus on "productive AI use" (*AI Ideation*, *AI Content Drafting*, and *AI Data Insights*), while the remaining 59% relate to "AI implementation and oversight" (*AI Quality Review*, *AI Integration*, and *AI Ethics & Compliance*).[16]

Appendix Figure D.3 shows that ethics and compliance tasks are especially common among workers who have received employer-provided training in AI chatbots. In contrast, workplaces that merely encourage chatbot use have a higher share of "productive AI tasks" focused on content generation and ideation.

Consistent with the finding that AI chatbot workloads extend beyond direct use, Figure 7.(b) shows that 5% of non-users report new workloads resulting from AI chatbots, with the effects most pronounced in workplaces with chatbot initiatives. In particular, Figure D.1(b) shows that 10% of teachers who have not used AI chatbots report new workloads from the technology—almost entirely due to newly assigned tasks.

In summary, the widespread creation of new tasks and spillovers to non-users highlight the broader workplace transformations caused by AI chatbots. Yet, the high prevalence of "non-productive AI tasks" related to integration, compliance, and oversight suggests that many organizations remain in the trough of a productivity J-curve. Notably, the strong association between training and the emergence of non-productive tasks helps

[16]The most common task relates to the integration of AI chatbots into the workplace, accounting for 15%–40% of all new tasks, with the highest shares observed in IT support and software development. *AI Integration* refers to embedding AI chatbots into workflows to automate or enhance tasks. For example, software developers report "fine-tuning AI coding assistants by providing feedback and project-specific examples," legal professionals describe "developing organizational AI usage policies and guidelines," and teachers mention "adapting exams and assignments to account for AI tool usage."

contextualize the finding in Section 3 that training—when implemented in isolation—is linked to consistently lower productivity gains from AI chatbots.

3.3 Total Time Savings

What is the economic significance of the reported benefits of using AI chatbots? In Figure 8, we combine workers' frequency of use (from Figure 3) with their reported time savings per day of usage (from Figure 5) to estimate time savings as a percent of total work hours.

While time savings capture only one dimension of chatbot benefits—arguably the least transformative—their clear measurability in minutes per day makes them easy to quantify and compare. Moreover, because prior RCTs have also used time savings as a key metric for productivity gains, our estimates can be benchmarked against those studies. Finally, time savings remain the most commonly reported benefit among workers in our survey.[17]

Figure 8 shows estimated time savings by workers' occupations and whether their employer encourages chatbot use. We focus on the encouraged-use policy, as Sections 2 and 3 identify it as the primary driver of increased adoption and reported benefits. The occupational breakdown facilitates comparison with prior studies focused on specific job roles.

Average time savings from AI chatbots vary from 6.8% among marketing professionals whose employers encourage their use to 0.6% among teachers in schools that do not encourage their use. Across our 11 occupations and employer policies, the average time savings amount to 2.8% of the total work hours of users.

Comparison to the Literature. These reported time savings by adopters may seem small compared to the substantial productivity gains—often exceeding 15%—documented

[17]Section 4 examines how various AI chatbot benefits—including time savings, enhanced creativity, improved quality, and other work-related dimensions—affect workers' earnings and hours.

by RCTs in our study occupations.[18,19,20,21] How can we reconcile workers' modest reported time savings with the substantial effects observed in RCTs?

First, as Figure 8 shows, our reported time savings are highest precisely in the occupations covered by RCTs, such as customer and IT support (Brynjolfsson, Li and Raymond, 2025), marketing and human resources (Noy and Zhang, 2023), and software development (Peng et al., 2023). In contrast, reported savings are about half as large among teachers, accountants, and financial advisors, whose tasks may be less suited to chatbot assistance. This occupational heterogeneity supports the concept of a "jagged frontier" of AI chatbot costs and benefits (Dell'Acqua et al., 2023), warranting caution against extrapolating productivity estimates across tasks or occupations.

Relatedly, while RCTs often focus on jobs where AI chatbots are continuously applicable, the typical adopter in our data uses chatbots only every third to fourth day, diluting their overall time savings accordingly. On days when they do use the tools, however, average time savings still amount to just 6–7%, substantially below experimental estimates. Limited applicability may again be an important explanation: even on usage days, adopters spend only 5-6% of their work hours actively using AI chatbots (see Table D.7).[22]

[18]Noy and Zhang (2023) report time savings of 37% from ChatGPT in text writing tasks representative of marketing and HR professionals. Brynjolfsson, Li and Raymond (2025) estimate productivity effects of 15% from a GPT-based chat assistant among customer and IT support agents, with more pronounced effects for less experienced workers. Cui et al. (2024) estimate a 26% productivity gain from GitHub Copilot among software developers, while Peng et al. (2023) report a 58% increase.

[19]Our estimated productivity gains are closer to those of Acemoglu (2025), though for different reasons. While Acemoglu (2025) combines the experimental productivity estimates with a calibrated adoption curve, we measure both adoption and (self-reported) time savings directly. Bick, Blandin and Deming (2025) provide a thoughtful comparison between their survey-based estimates of productivity gains and the predictions of Acemoglu (2025).

[20]Our time savings estimates align more closely with those of Bick, Blandin and Deming (2025), who, in a nationally representative U.S. survey, find an average reported time savings of 5.4% among users. However, their study does not split the time savings by occupation or employer policies, limiting direct comparability to our estimates.

[21]Edelman, Ngwe and Peng (2023) shows that self-reported time savings from AI chatbots—particularly Microsoft Copilot—may overstate actual productivity gains. This further reinforces the point that workers' productivity improvements from AI chatbots are modest, as our self-reported time savings estimates would then represent an upper bound on realized benefits. That said, overstated gains may not be the full story: Section 4.3.1 shows that workers' self-reported earnings effects closely align with our difference-in-differences estimates based on administrative data, reflecting that workers recognize that AI chatbots have not materially improved their earnings.

[22]The time use estimates should be interpreted with caution, as we might not expect a 100% active

Second, reported time savings are significantly higher when employers actively encourage chatbot use. For example, marketers and software developers whose employers promote AI chatbots report time savings of approximately 7%. In this sense, RCTs may capture the productivity effects of a well-coordinated and carefully managed adoption of these tools. While such estimates provide valuable insight into the workplace potential of AI chatbots, our survey indicates that many users do not operate under similarly favorable conditions. This highlights the need for caution when extrapolating productivity benefits from controlled settings onto the broader economy.[23]

4 Labor Market Outcomes

Our preceding analysis shows that AI chatbots are now widespread, with adopters reporting time savings and other work-related benefits, especially in workplaces that encourage their usage. Have these changes affected workers' labor market outcomes, such as earnings or recorded work hours? To answer this question, we link the survey responses to administrative data on labor market outcomes. We first outline our empirical strategy in Section 4.1 before presenting our results in the subsequent sections.

4.1 Empirical Strategy

Evaluating the labor market impacts of AI chatbots raises several conceptual questions: What are the relevant treatments to consider? If the treatments have an effect, in which

utilization rate even in tasks where AI chatbots are directly applicable. For example, in the writing task studied by Noy and Zhang (2023), treated workers spent on average 3 minutes before pasting in a large block of text—presumably generated by ChatGPT. The workers then spend an additional 14 minutes manually editing and refining the output. In our data, adopters spend an average of 2.2% of their total work hours using AI chatbots—less than the time savings they report. This pattern is also consistent with Noy and Zhang (2023), where treated workers saved 10 minutes on average despite only 3 minutes of initial tool use.

[23]Another source of discrepancy between our time savings estimates and those from RCTs lies in the parameter of interest: whereas experiments typically identify Average Treatment Effects (ATE) for the study population, our estimates reflect self-reported time savings among users—closer to an Average Treatment Effect on the Treated (ATT). However, if workers adopt chatbots based on their expected individual benefits—such that users tend to derive greater gains than non-users—the modest time savings reported in our survey appear even more stark when contrasted with the larger effects found in RCTs.

labor market outcomes would they manifest? And how can we identify these impacts in the data?

Appendix A.2 presents an economic framework to guide our analysis of how adoption translates into labor market outcomes. The framework embeds the Roy model of worker adoption from Sections 2 and 3 into a labor market equilibrium model where firms adopt chatbot initiatives and hire workers. The model motivates three types of analyses:

1. *Worker earnings:* If AI chatbots raise individual productivity, these gains may translate into higher wages for users. Notably, the wage effects depend on whether the productivity benefits are worker- or firm-driven, the amenity value of chatbot use, as well as the relative scarcity of adopting workers compared to encouraged firms. Section 4.2.1 examines these predictions by estimating earnings impacts by workers' adoption and encouragement status.

2. *Job mobility and sorting:* If chatbots are especially useful for newcomers to an occupation, adopters may be more likely to switch occupations. Furthermore, AI chatbots may lead to a resorting of workers who wish to use the tools moving into firms that encourage their use. Sections 4.2.2-4.2.3 investigate these implications at the worker and workplace levels.

3. *Workplace employment:* The effect of employer chatbot investments on workplace employment (measured in headcounts or hours) is theoretically ambiguous, depending on firms' investment costs, the productivity gains they generate, and the elasticity of product demand. Section 4.2.3 examines these effects.

The model also informs our empirical strategy for identifying the relevant effects. We use a series of difference-in-differences comparisons of adopters and non-adopters before and after the introduction of AI chatbots. Section 4.1.1 presents the regression specifications and assesses the validity of this design. These estimates identify the effects

of AI chatbots up to occupation-wide general equilibrium effects (in the model, the wage impact on non-encouraged non-adopters). We assess the magnitude of such spillovers in Section 4.3.2.

For the relevant treatments, we build on Sections 2 and 3, which showed that employers who explicitly encourage AI chatbot use observe substantially higher adoption and reported benefits. We use these "AI front-runner" workplaces—those with encouraged use—as our primary workplace treatment, as they offer a window into the effects of AI chatbots as employers embrace the tools. We also examine the effects of additional investments in employer-provided chatbots and training, as well as worker-level variation in adoption intensity and reported benefits.

A central question for the difference-in-differences analysis is timing: When should we date the arrival of AI chatbots, and when should we expect their effects on labor market outcomes to appear? We use the launch of ChatGPT in November 2022 as a before-and-after moment in an event-study design, marking a sharp rise in public awareness of AI chatbots.[24] We also separately examine effects for early adopters—workers observed using chatbots in our 2023 survey round. However, even if adoption occurs rapidly, labor market effects may take time to materialize. We examine these adjustment dynamics in two ways. First, we implement a dynamic difference-in-differences design to trace how effects unfold monthly following the introduction of AI chatbots. Second, we analyze heterogeneity across occupations that differ in the rigidity of wage-setting and employment contracts.

4.1.1 Regression Specifications

Our difference-in-differences specification compares the labor market outcomes of adopters and non-adopters before and after the introduction of AI chatbots. We examine both

[24]This timing is motivated by the rapid uptake of ChatGPT among workers: most adopters began using it within a year of its launch (Humlum and Vestergaard, 2025). Appendix C.2.1 shows that ChatGPT remains the dominant AI chatbot to date. Moreover, event studies centered on its launch in November 2022 have become common in the literature (e.g., Eisfeldt et al. (2024); Schubert (2025); Teutloff et al. (2025)), facilitating comparison across studies.

worker-level outcomes (e.g., worker earnings by chatbot usage) and workplace-level outcomes (e.g., workplace employment by chatbot initiatives).

A key concern in this analysis is that adopters may have fared differently in the labor market even without AI chatbots. We take two steps to address this issue. First, we control for workers' pre-determined characteristics—including gender, age, and labor market experience—to ensure these factors do not drive our estimates (e.g., adopters being younger and naturally on upward earnings trajectories). Second, we leverage our panel data to implement a difference-in-differences approach indexed to November 2022, the release date of ChatGPT. This allows us to control for time-invariant differences between workers (e.g., high-ability adopters who would have earned more regardless) and examine whether adopters experienced differential changes after this date. Additionally, we assess whether adopters were on distinct labor market trends even before AI chatbots became available, enabling us to control for potential pre-existing differences.

Let Y_{it} denote a labor market outcome (e.g., earnings) for worker i in month-year t, let X_i represent a vector of workers' pre-determined characteristics (age, gender, experience, and occupation FEs), and let A_i indicate whether the worker has adopted AI chatbots.

Dynamic Difference-in-Differences. To assess overall impacts, we employ a dynamic difference-in-differences specification that compares the monthly outcomes of adopters and non-adopters before and after the introduction of AI chatbots:

$$Y_{it} = \underbrace{\sum_{\tau} \lambda_{1\tau} X_i \mathbf{1}_{\{t=\tau\}}}_{\text{Pre-Determined Controls}} + \underbrace{\sum_{m} \lambda_{2m} \mathbf{1}_{\{m=m(t)\}}}_{\text{Seasonality}} + \underbrace{\sum_{\tau \neq 2022M11} \beta_\tau A_i \mathbf{1}_{\{t=\tau\}}}_{\text{Dynamic Diff-in-Diffs}} + \beta_0 A_i + \varepsilon_{it}, \quad (2)$$

where λ_{2m} are month fixed effects that control for seasonality in labor market outcomes. The parameters of interest, β_τ, capture the differential changes in labor market outcomes for adopters, indexed to November 2022, the release of ChatGPT. We estimate Equation (2) by OLS and cluster the standard errors by workplace-occupation level.

Pooled Difference-in-Differences. To examine heterogeneous impacts across multiple dimensions, we use a pooled specification that compares the average outcomes of adopters and non-adopters before and after the introduction of AI chatbots. To study dynamics, we estimate separate effects for each year after adoption.[25] Our estimating equation reads:

$$Y_{it} = \underbrace{\sum_{\tau} \lambda_{1\tau} X_i \mathbf{1}_{\{t=\tau\}}}_{\text{Pre-Determined Controls}} + \underbrace{\sum_{m} \lambda_{2m} \mathbf{1}_{\{m=m(t)\}}}_{\text{Seasonality}} + \underbrace{\lambda_3 t + \lambda_4 A_i t}_{\text{Time Trends}} + \underbrace{\sum_{y=2023}^{2024} \beta_y A_i \mathbf{1}_{\{y(t)=y\}}}_{\text{Pooled Diff-in-Diffs}} + \beta_0 A_i + \varepsilon_{it}, \tag{3}$$

where we add adopter-specific time trends to ensure our pooled difference-in-differences estimates are not driven by spurious trends in outcomes. The parameters of interest, β_y, measure how labor market outcomes for adopters changed in year y after the introduction of AI chatbots, compared to the period before their arrival. We estimate Equation (3) by OLS and cluster the standard errors by workplace-occupation level.

4.2 Results

4.2.1 Worker Earnings

Figure 9 examines earnings of chatbot adopters, distinguishing between those with and without employer encouragement. Each group is compared to non-encouraged non-adopters—workers who are most insulated from AI chatbots, either through their own non-use or that of their coworkers.

Panel (a) begins by presenting differences in means, controlling for workers' pre-determined characteristics—occupation, gender, age, and experience. Adopters appear to earn substantial premia of about 9% for those with employer-encouraged usage and 4% for those without. However, leveraging the panel dimension of our administrative data reveals that these earnings differentials entirely predate the arrival of AI chatbots. When indexing differences to November 2022, the difference-in-differences estimates in Panel (b)

[25] To align with our dynamic difference-in-differences specification (Equation (2)), which defines post-periods as months after 2022M11, we include 2022M12 in the 2023 estimate.

show no differential changes for adopters following the introduction of AI chatbots. The estimates remain flat even among adopters whose employers explicitly encourage chatbot use. Figure E.1 shows that these null results hold separately for workers' hourly wages, intensive-margin hours, and extensive-margin employment.

Figure 9 provides three key insights. First, one and a half years after their launch, AI chatbots have had minimal average impact on adopters' earnings. The confidence intervals of our dynamic difference-in-differences estimates rule out changes larger than 2%, and our pooled difference-in-differences estimates rule out effects larger than 1%.[26] Second, these effects show no differential trends after the introduction of AI chatbots, indicating that the minimal impacts are not merely a very short-term phenomenon.[27] Lastly, despite evidence in Section 3 that employer encouragement boosts work-related benefits of AI chatbots, we find no differential changes in earnings for adopters whose employers encourage their use.

Heterogeneity by Strength of Treatment. The null results on worker earnings in Figure 9 may reflect that adoption remains a relatively weak treatment. Even when use is encouraged, firms may provide limited guidance on how to use the tools, and workers may struggle to realize their full potential. If so, these estimates may understate the impact of AI chatbots as firms and workers learn to use them more effectively. Figure 10 investigates this possibility.

The top three rows consider the effects of additional investments in employer-provided chatbots and training. These estimates offer insight into the impacts of AI chatbots when firms pursue a more comprehensive strategy. Yet even when the "full package" of

[26]The dynamic difference-in-differences in Figure 9.(b) reveal that adopters are on slightly stronger labor market trends. However, because these trends entirely predate AI chatbots, the pooled difference-in-differences in Figure 15.(a) (which control for pre-trends; see Equation (3)) are precise zeros, with confidence bands ruling out effects larger than 1%.

[27]Section 4.3.1 shows that survey responses from November 2024 similarly report no earnings impacts of AI chatbots.

encouraged use, enterprise chatbots, and training substantially increases reported benefits (Section 3), we find no detectable impact on adopter earnings.

The next two rows focus on more intensive users—those who engage with AI chatbots daily or who adopted early (within the first year after ChatGPT's launch). These workers have not experienced improved labor market outcomes either. The fact that even daily users do not see earnings gains suggests the null results are not simply due to a lack of integration into everyday workflows. The null results for early adopters are consistent with our flat dynamic difference-in-differences estimates, suggesting that the zeros may not merely be short-run.

The remaining rows examine workers who report greater benefits from AI chatbot usage: time savings exceeding one hour per day, higher-quality output, enhanced creativity, and engagement with new AI-related tasks. While Section 3 emphasizes time savings due to its easily quantifiable nature, improvements in the quality or scope of work may be more consequential for earnings. Nonetheless, the estimates show that none of these perceived benefits are associated with better labor market outcomes.

Taken together, Figure 10 suggests that the null earnings effects are not simply the result of weak treatments—at either the firm or worker level. Workers in firms that adopt a more active, multifaceted approach to AI chatbots do not see differential earnings impacts. Nor do those who use the tools more frequently, adopt them earlier, or report large work-related benefits.

Finally, even if the limited earnings impacts do not reflect weak adoption treatments, rigid wage-setting systems may hinder the short-run translation of productivity gains into worker earnings. In Section 4.2.4, we explore this possibility by analyzing heterogeneity across occupations with varying exposure to collective bargaining agreements. To preview the results, even adopters in highly flexible pay settings—such as marketers and software developers—have not experienced differential earnings outcomes.

Comparison to the Literature. While a host of studies have documented productivity gains of chatbot use, it remains unclear how these translate into workers' earnings and hours, as high-quality microdata on such outcomes is rarely available. We fill this gap by leveraging the comprehensive administrative labor market data offered in the Danish setting.

Our analysis of adopter earnings follows a longstanding tradition, beginning with Krueger (1993)'s seminal work on wage premia associated with computer use. Our findings echo the critique by DiNardo and Pischke (1997): although chatbot usage is associated with sizable earnings premia in the cross-section, these differentials are entirely driven by positive selection.

Are the weak wage pass-through rates from chatbot-driven productivity gains theoretically surprising? Our economic framework in Section A.2 implies that workers should be compensated for productivity gains that are not driven by employer initiatives.[28] From this perspective, it is surprising that individual-specific benefits from AI chatbots—such as those among non-encouraged adopters in Figure 9 or among high-use, high-benefit adopters in Figure 10—do not translate into higher earnings.[29] On the other hand, if employers do not endorse or invest in the tools, they may be reluctant to reward their use. From that perspective, it is surprising that employer chatbot initiatives—such as encouraged-use policies, enterprise chatbot solutions, and employer-provided training—do not translate into measurable changes in worker earnings.

[28]This is a standard implication in labor markets where wages reflect outside options, dating back to Becker's distinction between general and firm-specific human capital (Acemoglu and Pischke, 1998; Becker, 1964).

[29]Of course, one possibility is that workers simply overestimate their actual benefits from the tools. Given that the reported time savings are already modest, this would reinforce the conclusion from Section 3.3 that realized gains are substantially lower than RCT estimates. Nevertheless, overstated gains may not be the full story: As discussed in Section 4.3.1, workers themselves recognize that only a small fraction of their reported time savings translates into higher earnings. Indeed, workers' self-reported earnings effects closely align with our difference-in-differences estimates based on administrative data.

4.2.2 Job Mobility

Figure 11 shows that adopters have seen greater occupational switching since the arrival of AI chatbots. By 2024, the additional hours worked in their latest occupations amount to about 2% of full-time equivalent employment. The figure provides three insights.

First, the flat pre-trend prior to the arrival of chatbots suggests that the association between occupation switching and chatbot adoption reflects a real effect rather than omitted confounders (e.g., more dynamic workers both adopting chatbots and switching jobs more generally). This contrasts with the earnings gaps in Figure 9.(a), which were entirely attributable to type-specific selection.

Second, the figure provides strong support for the notion that chatbots appeal to workers who are new in their occupations. This correlation may reflect that chatbots facilitate occupational switching or, perhaps more likely, that they are simply more attractive after a switch. While the difference-in-differences estimates do not identify the direction of causality, they clearly establish that chatbot adoption and occupational switching are intrinsically linked.

Finally, although the patterns in occupational mobility are clear and statistically significant, the estimated effects are modest—amounting to roughly a 2% increase in full-time equivalent employment—and not large enough to generate net gains in hours or earnings for adopters (Section 4.2.1). In other words, the patterns reflect a relatively modest reallocation of work hours across occupations, without corresponding changes in overall hours or earnings.

In addition to occupational switching, AI chatbots may also induce mobility across workplaces. For instance, if chatbots help workers move up firm-specific learning curves, adopters may be more likely to switch employers. Alternatively, if encouraged-use policies enhance the utility of chatbots (as shown in Section 3), adopters may seek to sort into such workplaces. This prediction is formalized in our theoretical framework in Appendix A.2.

Despite these possibilities, Appendix Figure E.2 shows no clear link between chatbot adoption and workplace mobility. Moreover, job mobility effects—across either workplaces or occupations—are not stronger among adopters whose use is encouraged by their endline employers. These findings yield two implications. First, chatbot adoption appears to be specifically associated with occupational switching, rather than general labor market mobility, suggesting that chatbots may facilitate the acquisition of occupation-specific rather than firm-specific skills. Second, the mobility effects seem to be driven by individual-level factors rather than firm-level ones, as adopters in encouraged and non-encouraged workplaces exhibit similar transition rates.

The next section explores further heterogeneity by the strength of the adoption treatment, considering both worker- and firm-level factors. Section 4.2.3 further examines how employer initiatives shape job churn and mobility at the workplace level.

Heterogeneity by Strength of Treatment. Figure 12 examines how the relationship between chatbot adoption and occupational mobility varies with the strength of the adoption treatment. The first five categories show that this relationship is fairly similar across employer chatbot initiatives, ranging in 2024 from just below 2% for non-encouraged adopters to about 3% for those with both encouraged-use, enterprise chatbots, and training. The next two categories show that occupational switching is markedly higher among individuals with larger reported gains from the tools, with the 2024 effects rising to 4% and 6.5% for those who use the tools daily or report large time savings, respectively. The final four categories show that adopters who report increased creativity and quality, new workloads, and the receipt of AI-related tasks experience similar increases in occupational switching of around 2.5%.

Overall, Figure 12 reinforces two points. First, the association between chatbot adoption and occupational switching appears to be driven primarily by individual-level factors rather than firm-level initiatives. Second, occupational switching is closely linked

to chatbot gains: adopters who use the tools most intensively and derive the greatest benefits are also the most likely to be new to their occupations.

Comparison to the Literature. A consistent finding from experimental studies is that AI chatbots provide greater productivity benefits to workers with less prior expertise in their occupations (Brynjolfsson, Li and Raymond, 2025; Dell'Acqua et al., 2023; Noy and Zhang, 2023). This section presents large-scale evidence supporting a key implication of those studies: chatbot adoption is associated with occupational switching. In this sense, our findings align with the thesis of Autor (2024), who argues that Generative AI may help individuals extend their work into new occupations. Our administrative data shows that, just one and a half years after the arrival of Generative AI chatbots, these dynamics are already unfolding in the labor market.

4.2.3 Workplace Employment

Even if AI chatbots have not affected the net earnings or hours of workers who use them, employers may still use the tools to adjust overall labor demand. For instance, if chatbots reduce the need for labor, adopting workplaces might slow hiring or lay off workers.[30] Conversely, if the tools boost productivity and spur demand, these workplaces may instead expand employment and raise wage bills.

Moreover, even if overall employment remains constant, AI chatbots may still alter the mix of workers at different workplaces. In particular, our economic framework in Appendix A.2 predicts that AI chatbots may lead to a resorting of workers who are prone to using AI chatbots into workplaces that encourage their use.

Figures 13 examines how employer chatbot initiatives affect workplace labor demand, job churn, and sorting.[31] We focus the dynamic difference-in-difference analysis on the

[30]For example, Anthropic CEO Dario Amodei recently predicted that AI could eliminate half of all entry-level white-collar jobs, driving unemployment to 10-20% within the next one to five years (Axios, 2025).

[31]We measure workplace impacts within each occupation by tracking how, for example, the demand for

encouraged-use policy, as Sections 2–3 identified it as the strongest predictor of adoption and reported benefits. Figure 14 considers additional investments in employer-provided chatbots and training.

Figure 13, Panels (a) and (b) show that workplaces that encourage chatbot usage have seen no differential change in employment or wage bills after the arrival of AI chatbots. That is, despite encouraged-use policies substantially boosting adoption and reported benefits among users, they have had no detectable impact on workplace labor demands.[32]

Figure 13, Panel (c) examines impacts on job churn. Even with constant employment levels, AI chatbots may require a changed mix of skills, potentially increasing the rate of job creation and destruction in the workplace. For instance, adopting and operating new tools may require hiring new workers who displace incumbents (Bessen et al., 2023), whereas assistance from the tools could instead help retain workers (Brynjolfsson, Li and Raymond, 2025). However, despite these possibilities, the estimates in Panel (c) show that rates of job creation and destruction have remained stable for workplaces that encourage chatbot usage.

Finally, even if overall rates of job churn are constant at the workplace, encouraging AI chatbots may attract workers who would like to use the tools. Figure 13, Panel (d) explores this possibility by estimating the "adopter propensity" of hires and separations at the workplace—that is, the predicted adoption rate (based on workers' pre-determined characteristics) of workers who join or leave the workplace.[33] However, despite the model prediction that AI chatbots should lead to workers resorting across firms, we see no differential changes in the adopter propensity of workers who join or leave workplaces that encourage chatbot usage. This finding—that employer initiatives have not led to job

a given occupation evolves in workplaces where workers in that occupation report a specific employer chatbot initiative.

[32]While Figure 13 focuses on employment within workers' respective occupations, Appendix Figure E.3 shows that overall workplace employment remains flat as well.

[33]Appendix G.1 details the estimation of these adopter propensities, which our economic framework in Section A.2 microfounds as workers' latent willingness to use AI chatbots.

churn or worker resorting—is consistent with our results in Section 4.2.2, showing that the increased occupational mobility among adopters is driven primarily by worker-level factors rather than firm-level ones.

Taken together, the results in Figure 13 show that workplaces that encourage AI chatbots have not fared differently in terms of overall labor demand, job churn, or worker sorting since the arrival of AI chatbots. These limited impacts on labor market outcomes contrast with the findings from Section 3 that AI chatbots have caused widespread creation of new tasks, especially at workplaces that invest in the technology: while workplaces have started to reorganize internally in response to AI chatbots, we remain to see these developments manifested in their labor market outcomes.

Heterogeneity by Strength of Treatment. The limited impact of encouraged-use policies on workplace outcomes may reflect that these are relatively weak treatments. In particular, without additional investments in tools and people, these initiatives may fail to deliver workplace labor market transformations. Indeed, Section 2-3 showed that adoption, reported benefits, and new job tasks are greatest when employers simultaneously encourage use, invest in enterprise chatbots, and provide training. However, despite these larger reported effects, Figure 14 shows that the additional initiatives have had no detectable impacts on labor demand at the workplace. Similarly, Appendix Figure E.4 shows that workplace job churn rates and worker sorting have not changed differently at workplaces with additional investments in employer-provided chatbots and training.

Comparison to the Literature. We have limited evidence on how AI chatbots affect workplace outcomes, largely because firm-level data on Generative AI adoption remain scarce, and even fewer datasets link adoption to workplace economic outcomes.[34] We

[34]The few existing surveys on Generative AI adoption all sample individual workers, whose responses cannot be linked at the workplace level (Bick, Blandin and Deming, 2025; Hartley et al., 2024; Pew, 2024). Bonney et al. (2024) survey AI adoption in U.S. firms as of February 2024 but do not specifically measure Generative AI or link adoption to subsequent changes in economic outcomes.

leverage our matched employer-employee links to relate workers' survey responses to the labor market outcomes of their workplaces.

The closest related evidence comes from online job postings. Eisfeldt et al. (2024) show that firms with occupational compositions more exposed to AI chatbots—based on pre-ChatGPT job postings—experienced stock price increases following the release of ChatGPT.[35] Schubert (2025) further shows that hiring for Generative AI skills, as listed in job postings, interacts with the demand for remote work. While these studies based on job postings suggest notable changes in firms' labor demand immediately following the release of ChatGPT, we find no corresponding effects in administrative labor market data.[36]

4.2.4 Occupational Heterogeneity

The previous sections show that the adoption of AI chatbots has had no significant impact on the labor market outcomes of workers or workplaces. While adoption is associated with occupation switching and task restructuring, these shifts have not led to net changes in hours or earnings. However, these overall zero effects could mask heterogeneous gains and losses across different occupations. For example, since teachers are covered by collective bargaining with centralized wage setting, it may be less surprising that worker-specific productivity gains do not translate into higher earnings. By contrast, occupations like marketing and software development feature decentralized wage setting,

[35]This rise in market capitalization is consistent with firms being at an early stage of a productivity J-curve, where current investments are expected to yield economic returns over the longer term (Brynjolfsson, Rock and Syverson, 2021).

[36]While job postings offer a rich source of data, it is less clear how well they capture employees' actual use of tools, particularly non-specialized ones like ChatGPT. First, although the emergence of Generative AI skills in job postings is notable, it remains unclear how such listings correspond to employees' actual use of AI chatbots (Schubert, 2025). Second, while Eisfeldt et al. (2024); Schubert (2025) rely on variation in firm-level *exposure* to AI chatbots *across* occupations (i.e., variation in hiring across jobs with different relevance to AI chatbots), our analysis instead examines variation in firm-level *adoption* of AI chatbots *within* occupations (i.e., variation in actual use of AI chatbots for the same relevant jobs).

allowing more flexibility to adjust pay based on individual productivity.[37,38] Comparing across occupations thus offers a way to assess whether the limited impacts AI chatbots reflect labor market rigidities. Furthermore, splitting our mobility estimates by occupation offers insight into where chatbots are especially useful for newcomers.

Figure 15, Panels (a) and (b) show the estimated effects of AI chatbots on worker earnings and workplace employment for each of the 11 occupations in our sample. We find no significant effects in any occupation: the estimates are all close to zero, generally ruling out effects larger than 5%. Even in our two most contrasting occupation groups—software developers and marketers (with flexible wages) versus teachers (with centralized wage-setting)—we see no systematic differences in impact. These results suggest that the limited labor market effects are not simply due to rigid wage-setting or limited occupation-wide benefits. The absence of effects even in occupations with the largest reported gains and most flexible wages indicates that overall labor market impacts from AI chatbots may remain limited for the foreseeable future.

Figure 15, Panel (c), splits the effects of chatbot adoption on mobility by occupation. The positive effects are mainly driven by IT supporters, and to some extent by office clerks. These are jobs where individual workers likely have more flexibility in choosing their toolset, allowing newcomers to take full advantage of chatbot benefits. In contrast, financial advisors and customer support specialists often work with established IT systems, making it harder for newcomers to compensate for missing skills by using chatbots. These patterns align with Dillon et al. (2025), who show that workers primarily adopt Microsoft Copilot for tasks where they can use the tools without coordinating with others. Finally, some occupations, such as teachers and accountants, require licensing through several years of prior education. Thus, the absence of switching effects in these occupations

[37]Appendix B.1 provides an overview of the wage-setting systems in the occupations we study.

[38]Furthermore, marketing professionals and software developers are exactly the occupations with the highest reported use and productivity gains from the tools (Figure 8).

reinforces that the mobility effects we see elsewhere reflect a real link between chatbot use and newcomers gaining occupation-specific skills.

Comparison to the Literature. The null results for overall labor market outcomes in each of our 11 exposed occupations stand in contrast to existing studies of online labor market platforms for freelance work, which document remarkable shifts in labor demand following the launch of ChatGPT. For example, Teutloff et al. (2025) report that demand for substitutable freelance services, such as writing and translation, declined by 20–50% after ChatGPT. In contrast, we find no significant impact on labor market outcomes across our diverse set of 11 occupations. What might explain this difference? Our evidence points to three factors.

First, while freelance proofreading may be highly substitutable by AI chatbots, most exposed occupations do not fall into this extreme category.[39] Indeed, workers in our survey report only modest time savings from AI chatbots. This reiterates our observations from Section 3.3 that existing studies tend to concentrate on the occupations where the productivity effects from the tools are the largest.

Second, freelance work offers far less job security and rigidity, allowing productivity changes to manifest more quickly in labor market outcomes. In contrast, we find only a weak relationship between workers' time savings and labor market effects from AI chatbots—even in the occupations with the most flexible pay systems, such as marketers and software developers. While the flexibility of freelance work makes it useful for identifying short-run changes in labor demand, our findings on the limited pass-through of productivity gains caution against extrapolating these effects to the broader labor market.

Finally, although existing studies document shifts in the demand for certain tasks on specific platforms, it remains unclear how these shifts affect workers' total earnings and

[39]Teutloff et al. (2025) show that the overall impacts of ChatGPT on freelance work are more mixed, with tasks that require AI complementary skills seeing increases.

hours, as workers may reallocate their time across tasks. In our survey, most workers report using time saved by AI chatbots to perform other tasks, including new responsibilities introduced by AI. A key advantage of our administrative data is that it allows us to measure workers' total earnings and hours, regardless of how they reallocate their time in the labor market. In particular, while chatbot adoption is associated with occupation switching, it does not lead to net changes in hours or earnings.

4.3 Supplementary Evidence

4.3.1 Perceived Impacts

To complement the difference-in-difference analysis of labor market impacts, we simply asked respondents directly: *"Have AI chatbots affected your labor earnings?"* If so, we followed up with, *"By how much?"* If AI chatbots were meaningfully changing labor market conditions in exposed occupations, we might expect "front-line workers" in those occupations to notice and report such changes.

While individuals may struggle to assess such counterfactuals precisely, we treat their responses as supplementary evidence based on a fundamentally different methodology that helps address some of the limitations of our difference-in-differences analysis.

First, as shown by Arcidiacono et al. (2020); Fuster, Kaplan and Zafar (2021); Giustinelli and Shapiro (2024), subjective assessments can reveal individual-specific treatment effects that may vary along unobservable dimensions not captured in survey or register data. In particular, our heterogeneity analyses above may fail to isolate the subset of workers most affected by AI chatbots.

Second, while difference-in-differences estimates identify only *differential* impacts among adopters, subjective assessments are less susceptible to this "missing intercept" problem. For example, if chatbot availability constitutes an economy-wide shock, it could affect earnings across entire occupations, regardless of individual or workplace adoption.

In particular, our estimates in Section 4.2.1 identify effects only up to the wage impact on non-encouraged non-adopters. In contrast, we may still ask these workers about their perceived labor market impacts from AI chatbots. If substantial occupation-wide general equilibrium effects are unfolding, they may perceive the impact on their earnings even without adopting the tools themselves.

Finally, eliciting perceived earnings effects offers a bridge between the self-reported work benefits in Section 3 and the observed earnings outcomes in Section 4. This helps assess whether the weak link between perceived work-related benefits and observed earnings outcomes reflects misperception on the part of workers—or instead an accurate understanding that the pass-through from work-related gains to earnings is limited.

Appendix E.2.1 reports workers' perceived earnings impacts from AI chatbots. Strikingly, virtually all respondents—97.7% of adopters and 99.5% of non-adopters—state that AI chatbots have had *no* impact on their labor earnings. This finding has several implications.

First, even workers themselves report small average earnings impacts, ranging from 0% to 0.5%, depending on their occupation and employer use policy. These estimates are an order of magnitude smaller than the time savings reported in Section 3.3 and fall within the confidence intervals of the observed earnings effects in Section 4.2.1.

Second, because these perceptions were elicited in November 2024, they help address concerns that the zero difference-in-differences estimates in Section 4.2.1 merely reflect a short post-treatment window ending in June 2024. The continued absence of meaningful effects by November aligns with the dynamic difference-in-differences estimates, which show no differential trend following the introduction of AI chatbots.

Third, the limited average earnings impacts primarily reflect the overwhelming share of workers reporting no change in earnings due to AI chatbots—rather than large positive and negative effects offsetting each other. This supports our conclusion from Figure 10

that AI chatbots have not had significant heterogeneous impacts on worker earnings.

Fourth, the fact that virtually 100% of non-adopters report no earnings impact suggests that occupation-wide spillover effects are minimal. Despite broad adoption and changes in work content, AI chatbots do not appear to have shifted equilibrium wages or hours.

Finally, the weak link between self-reported productivity benefits in Section 3 and labor market outcomes in Section 4.2.1 is not simply due to workers overestimating the tools' impact. In particular, Figure E.9 shows that only 3–7% of workers' reported time savings manifests in higher perceived earnings.

4.3.2 Spillovers to Non-Adopters

Our main difference-in-differences estimates identify the effects of AI chatbots only up to occupation-wide effects—the so-called "missing intercept" problem. For example, by comparing earnings relative to non-encouraged non-adopters in Section 4.2.1, any general equilibrium effects that also affect these workers will be differenced out. This section assesses the potential magnitude of such spillovers.

First, we assess the extent of spillovers across different segments of the labor market by varying the level of aggregation in the difference-in-differences analysis. For instance, the fact that encouraged workplaces do not experience differential outcomes in Section 4.2.3 suggests that spillovers between adopters and non-adopters within workplaces are limited. Consistent with this interpretation, Figure E.7 shows that non-adopters in encouraged workplaces have fared similarly to their counterparts in non-encouraged workplaces following the introduction of AI chatbots. Extending this approach, Appendix E.2.2 shifts the analysis to the occupation level, showing that our exposed study occupations have also evolved similarly to the rest of the economy, maintaining a stable 20% share of aggregate hours and earnings.

Second, in Appendix A.2.6, we utilize an equilibrium model to demonstrate that our existing difference-in-differences estimates are informative about the magnitude of the

"missing intercept": under reasonable product demand elasticities, observing null effects in our workplace- and worker-level difference-in-differences estimates implies that the occupation-wide equilibrium effect is also negligible.

Finally, as discussed in Section 4.3.1, workers' perceived earnings impacts offer an independent lens for evaluating general equilibrium spillovers. In particular, the fact that virtually *all* non-adopters report no effect on their earnings further reinforces the finding that occupation-wide impacts of AI chatbots are minimal.

Taken together, these different analyses—difference-in-differences at varying levels of aggregation, equilibrium calibration of existing estimates, and perceived impacts among workers—all point to the conclusion that AI chatbots have not had any meaningful impact on occupation-wide equilibrium outcomes thus far.

5 Conclusion

Generative AI is heralded as the engine of a new industrial revolution (World Economic Forum, 2024), yet we lack evidence on its economic impacts outside laboratory settings and case studies (Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023).

This paper provides large-scale evidence on the early labor market impacts of AI chatbots, the most widely adopted Generative AI tool to date. Our study is based on a series of extensive surveys on AI chatbot usage in 11 exposed occupations, linked to administrative labor market data in Denmark.

Despite rapid adoption and substantial investments by both workers and firms, our key finding is that AI chatbots have had minimal impact on labor market outcomes to date. Moreover, we find no evidence of differential trends over time—even among "AI front-runners"—suggesting that the limited effects may persist. In this sense, our results echo Robert Solow's famous observation about the IT revolution: *"You can see the computer age everywhere but in the productivity statistics"* (Solow, 1987).

However, our analysis sheds light on mechanisms through which Generative AI could become transformative over time.

First, consistent with Brynjolfsson, Rock and Syverson (2021), we find that firm-driven complementary investments and workplace reorganizations are critical to unlocking AI's potential. While some initiatives—such as training or enterprise chatbots—can be counterproductive in isolation, adoption and productivity gains are substantially larger when firms combine encouragement, enterprise chatbots, and training.

Second, aligning with theoretical predictions for how new technologies may reinstate labor demand (Acemoglu and Restrepo, 2018), we find that AI chatbots have created new job tasks—extending even to workers who do not use the tools directly—signaling broader workplace transformations. Crucially, the nature of this new work depends on firms' chatbot strategies: firms focused on preventing misuse tend to generate more compliance and oversight tasks, while those actively embracing chatbots see more tasks related to content creation and ideation.

Still, we believe that the key finding in this paper will remain central to understanding the labor market effects of Generative AI. Any account of transformational change must contend with a simple fact: two years after the fastest technology adoption ever, hard labor market outcomes—whether at the individual or firm level—remain untouched.

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy*, 40(121): 13–58.

Acemoglu, Daron, and Jörn-Steffen Pischke. 1998. "Why do firms train? Theory and evidence." *The Quarterly journal of economics*, 113(1): 79–119.

Acemoglu, Daron, and Pascual Restrepo. 2018. "Artificial intelligence, automation, and work." In *The economics of artificial intelligence: An agenda*. 197–236. University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of economic perspectives*, 33(2): 3–30.

Acemoglu, Daron, Gary W Anderson, David N Beede, Cathy Buffington, Eric E Childress, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas. 2022. "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey." National Bureau of Economic Research.

Altonji, Joseph G, Todd E Elder, and Christopher R Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of political economy*, 113(1): 151–184.

Arcidiacono, Peter, V Joseph Hotz, Arnaud Maurel, and Teresa Romano. 2020. "Ex ante returns and occupational choice." *Journal of Political Economy*, 128(12): 4475–4522.

Autor, David. 2024. "Applying AI to Rebuild Middle Class Jobs." National Bureau of Economic Research.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *The Quarterly Journal of Economics*, qjae008.

Axios. 2025. "Behind the Curtain: A White-Collar Bloodbath." Accessed: 2025-05-28.

Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.* Chicago:University of Chicago Press.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan Van den Berge. 2023. "Automatic Reaction-What Happens to Workers at Firms That Automate?" *The Review of Economics and Statistics*, , (Feb. 6, 2023).

Bick, Alexander, Adam Blandin, and David J Deming. 2025. "The Rapid Adoption of Generative AI." National Bureau of Economic Research.

Bonney, Kathryn, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, and Keith Savage. 2024. "Tracking Firm Use of AI in Real Time: A Snapshot From the Business Trends and Outlook Survey." National Bureau of Economic Research.

Botero, Juan C, Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. "The Regulation of Labor." *The Quarterly Journal of Economics*, 119(4): 1339–1382.

Bresnahan, Timothy F., and Shane Greenstein. 1996. "Technical Progress and Co-invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity: Microeconomics*, 1996: 1–83.

Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic perspectives*, 14(4): 23–48.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics*.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics*, 13(1): 333–372.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36(S1): S13–S70.

Carvajal, Daniel, Catalina Franco, and Siri Isaksson. 2024. "Will Artificial Intelligence Get in the Way of Achieving Gender Equality?" *NHH Dept. of Economics Discussion Paper*, , (03).

Cui, Zheyuan Kevin, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2024. "The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers." *Available at SSRN 4945566*.

Dahl, Christian M, Daniel Le Maire, and Jakob R Munch. 2013. "Wage Dispersion and Decentralization of Wage Bargaining." *Journal of Labor Economics*, 31(3): 501–533.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *The American Economic Review*, 80(2): 355–361.

Dell'Acqua, Fabrizio, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, , (24-013).

Dillon, Eleanor Wiske, Sonia Jaffe, Nicole Immorlica, and Christopher T Stanton. 2025. "Shifting Work Patterns with Generative AI." *arXiv preprint arXiv:2504.11436*.

DiNardo, John E, and Jörn-Steffen Pischke. 1997. "The returns to computer use revisited: Have pencils changed the wage structure too?" *The Quarterly journal of economics*, 112(1): 291–303.

Dutz, Deniz, Ingrid Huitfeldt, Santiago Lacouture, Magne Mogstad, Alexander Torgovitsky, and Winnie Van Dijk. 2025. "Selection in Surveys: Using Randomized Incentives to Detect and Account for Nonresponse Bias." Working Paper.

Edelman, Benjamin G, Donald Ngwe, and Sida Peng. 2023. "Measuring the Impact of AI on Information Worker Productivity." *Available at SSRN 4648686*.

Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska. 2024. "Generative AI and Firm Values." *Available at SSRN 4436627*.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384(6702): 1306–1308.

Feigenbaum, James, and Daniel P Gross. 2024. "Organizational and Economic Obstacles to Automation: A Cautionary Tale from at&T in the Twentieth Century." *Management Science*, 70(12): 8520–8540.

Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2021. "What would you do with 500?$Spendingresponsestogains, losses, news, andloans."TheReviewofEconomicStudies, 88(4) : 1760-$ $-1795.$

Giustinelli, Pamela, and Matthew D Shapiro. 2024. "Seate: Subjective ex ante treatment effect of health on retirement." *American Economic Journal: Applied Economics*, 16(2): 278–317.

Groes, Fane, Philipp Kircher, and Iourii Manovskii. 2015. "The U-Shapes of Occupational Mobility." *The Review of Economic Studies*, 82(2): 659–692.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore. 2024. "The Labor Market Effects of Generative Artificial Intelligence." *Available at SSRN*.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica: Journal of the econometric society*, 153–161.

Humlum, Anders, and Emilie Vestergaard. 2025. "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities Among Workers." *Proceedings of the National Academy of Sciences*, 122(1): e2414972121.

Hvidberg, Kristoffer B, Claus T Kreiner, and Stefanie Stantcheva. 2023. "Social Positions and Fairness Views on Inequality." *Review of Economic Studies*, 90(6): 3083–3118.

Jäger, Simon, Suresh Naidu, and Benjamin Schoefer. 2024. "Collective Bargaining, Unions, and the Wage Structure: An International Perspective." National Bureau of Economic Research.

Krueger, Alan B. 1993. "How computers have changed the wage structure: evidence from microdata, 1984–1989." *The Quarterly Journal of Economics*, 108(1): 33–60.

Munch, Jakob, and Will Olney. 2024. "Offshoring and the Decline of Unions."

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, 381(6654): 187–192.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics*, 37(2): 187–204.

Otis, Nicholas G, Solène Delecourt, Katelyn Cranney, and Rembrand Koning. 2024*a*. *Global Evidence on Gender Gaps and Generative AI*. Harvard Business School.

Otis, Nicholas, Rowan P Clarke, Solene Delecourt, David Holtz, and Rembrand Koning. 2024*b*. "The Uneven Impact of Generative AI on Entrepreneurial Performance." *Available at SSRN 4671369*.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from Github Copilot." *arXiv preprint arXiv:2302.06590*.

Pew. 2024. "Pew Research Center: Americans' Use of ChatGPT Is Ticking Up, but Few Trust Its Election Information." *https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information*, [Accessed 28-June-2024].

RISJ. 2024. "Reuters Institute for the Study of Journalism: What Does the Public in Six Countries Think about Generative AI in News?" *https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news#:~:text=Averaging%20across%20six%20countries%2C%20people,tried%20to%20use%20generative%20AI*, [Accessed 28-June-2024].

Roy, Andrew Donald. 1951. "Some thoughts on the distribution of earnings." *Oxford economic papers*, 3(2): 135–146.

Schubert, Gregor. 2025. "Organizational Technology Ladders: Remote Work and Generative AI Adoption." *Available at SSRN*.

Solow, Robert M. 1987. "We'd Better Watch Out." *New York Review of Books*.

Teutloff, Ole, Johanna Einsiedler, Otto Kässi, Fabian Braesemann, Pamela Mishkin, and R Maria del Rio-Chanona. 2025. "Winners and Losers of Generative AI: Early Evidence of Shifts in Freelancer Demand." *Journal of Economic Behavior & Organization*, 106845.

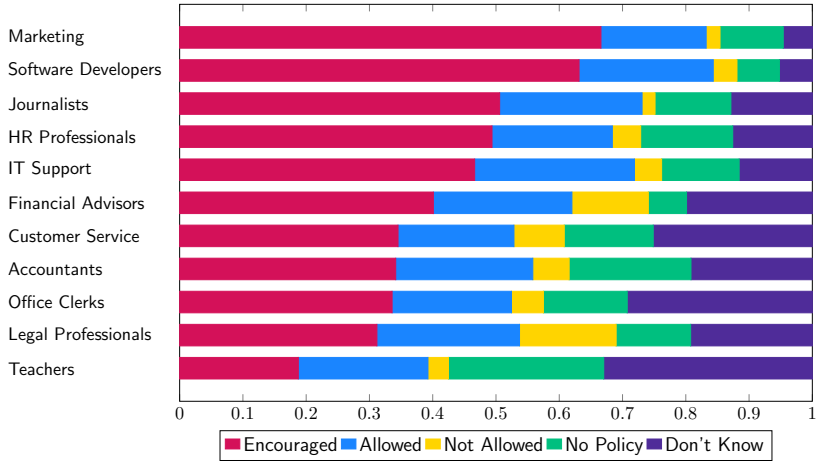
Walters, Christopher. 2024. "Empirical Bayes Methods in Labor Economics." In *Handbook of Labor Economics*. Vol. 5, 183–260. Elsevier.

World Economic Forum. 2024. "Generative AI: Steam Engine of the Fourth Industrial Revolution?" Accessed: 2024-12-11.

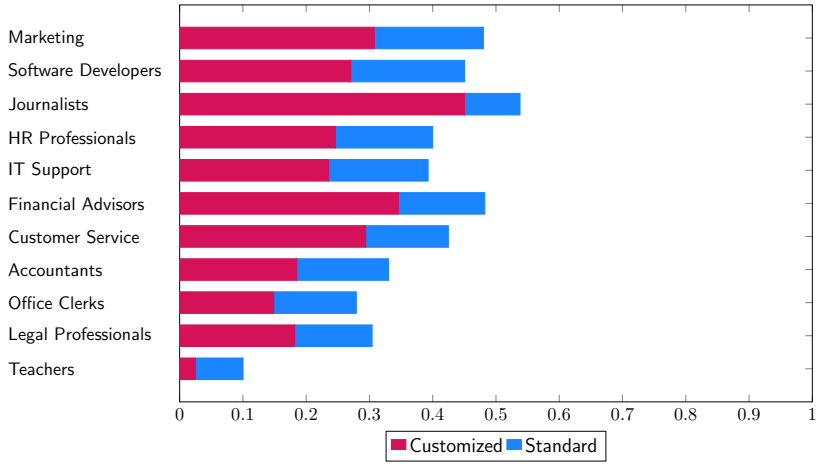
Figures and Tables

Figure 1: Prevalence of Employer Initiatives for AI Chatbot Adoption

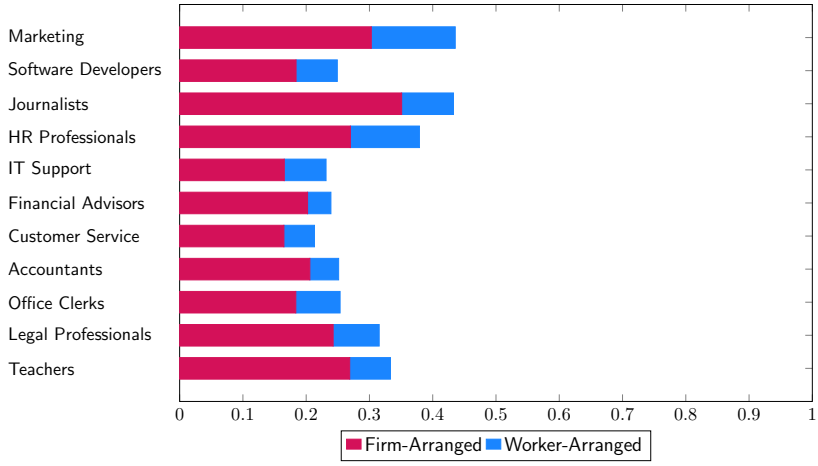
(a) Usage Policies



(b) Have Enterprise AI Chatbot

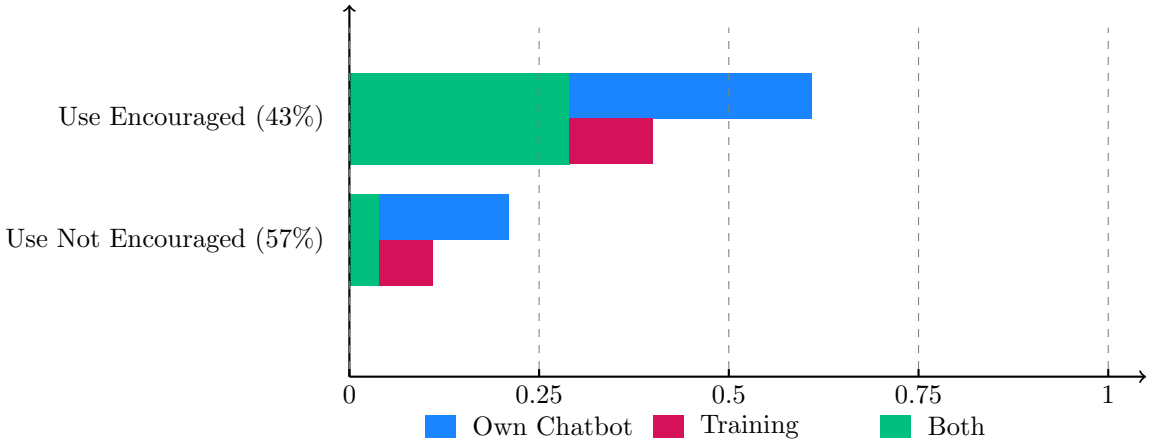


(c) Have Received Training in AI Chatbots



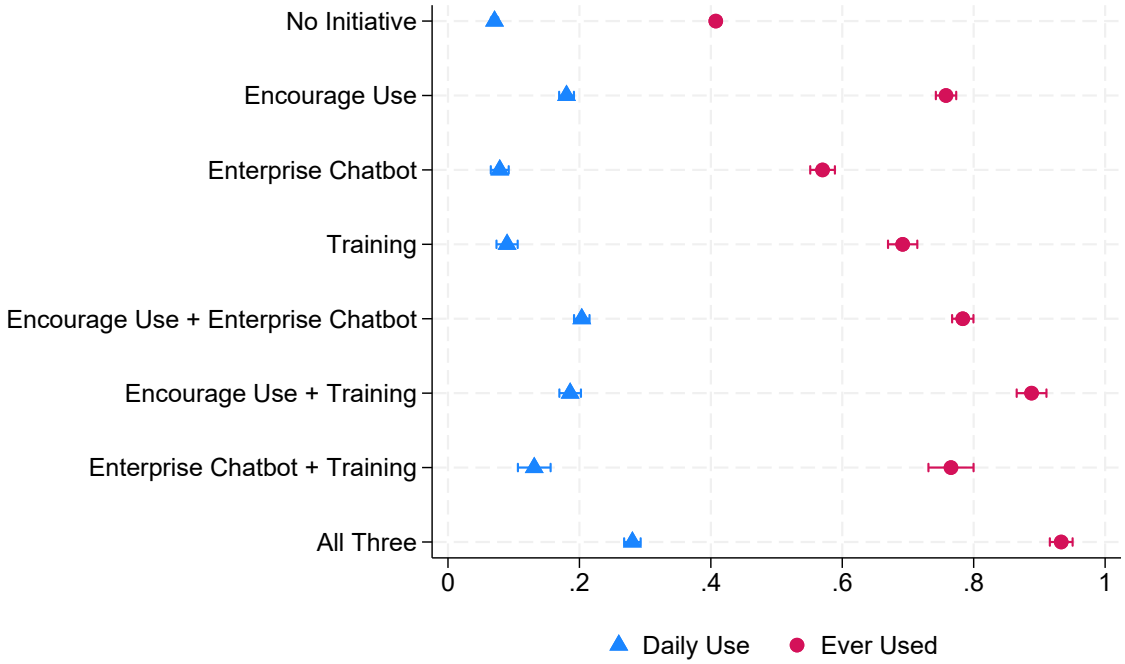
Notes: This figure shows the share of workers affected by various employer initiatives related to AI chatbot adoption. Panel (a) shows employers' policies on AI chatbot usage for work. Panel (b) indicates whether the employer has its own AI chatbot. Panel (c) reports whether workers have participated in AI chatbot training courses. Figure C.1 provides a workplace-level version of this graph, yielding similar results. *Sample*: All completed responses from the 2024 survey.

Figure 2: Packages of Employer Initiatives for AI Chatbot Adoption



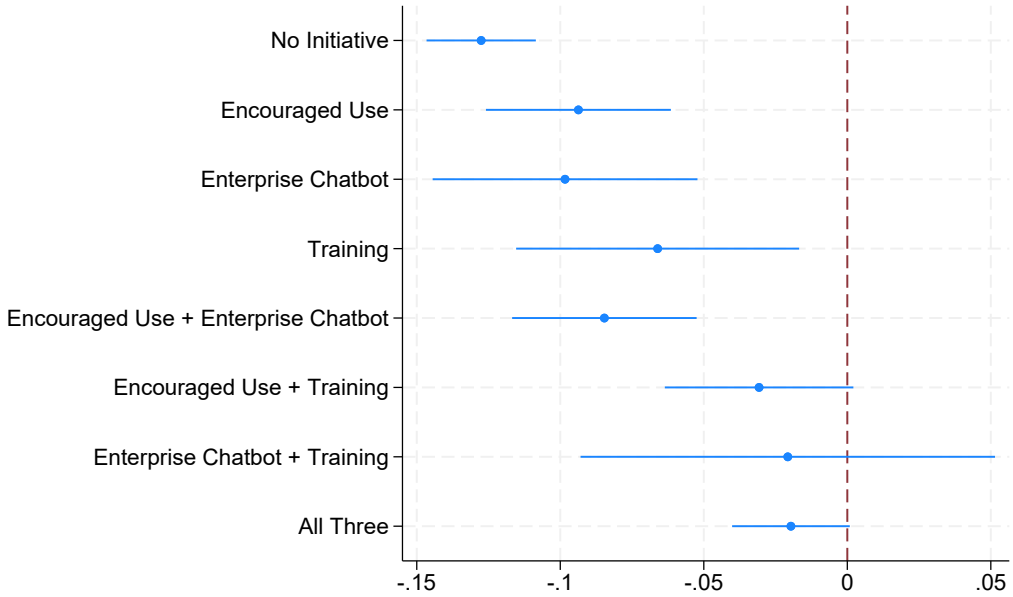
Notes: This figure shows the share of workers affected by different combinations of employer-led initiatives for chatbot adoption. Workers are grouped by whether their chatbot use is employer-encouraged or not. Within each group, the bars indicate the share of workers whose employer provides an enterprise AI chatbot, provide chatbot training, or both. *Sample*: All completed responses from the 2024 survey.

Figure 3: Employer Initiatives and Chatbot Adoption



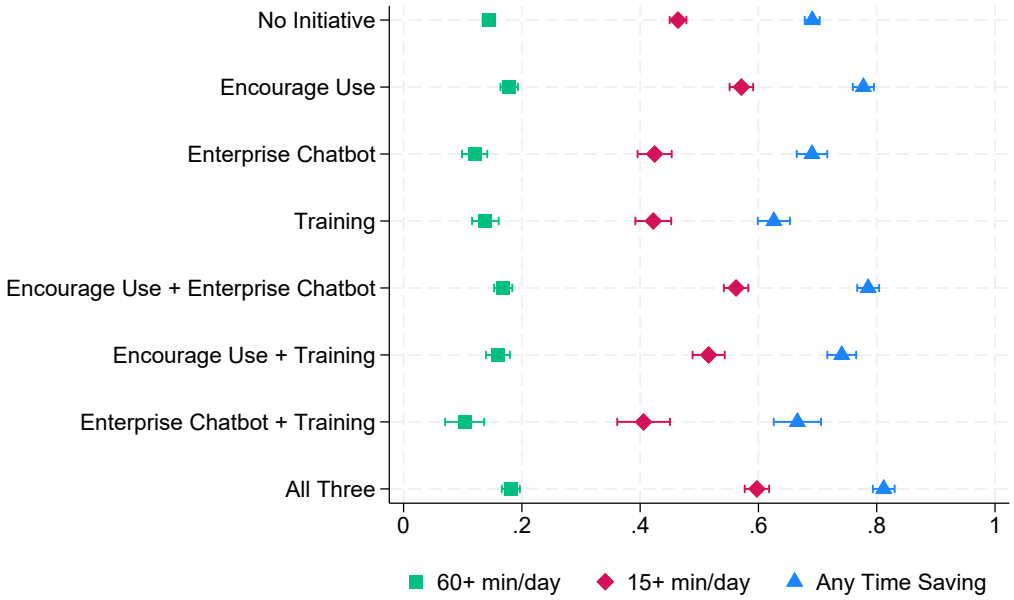
Notes: This figure illustrates how workers' use of AI chatbots vary with their employer initiatives. The estimates are based on predicted values from Equation (1), varying EmployerInitiatives while holding workers' characteristics X at their mean values. Whiskers represent 95% confidence bands of the predicted values. Some confidence intervals are too narrow to be distinguished from the point estimates at this scale. Table C.2 reports the full set of regression estimates, including specifications that use the share of workers with *Monthly or More* and *Weekly or More* usage as dependent variables. Figure C.2 visualizes these additional usage frequencies. *Sample*: All completed responses from the 2024 survey linked to registry data.

Figure 4: Employer Initiatives and Gender Gaps in Adoption



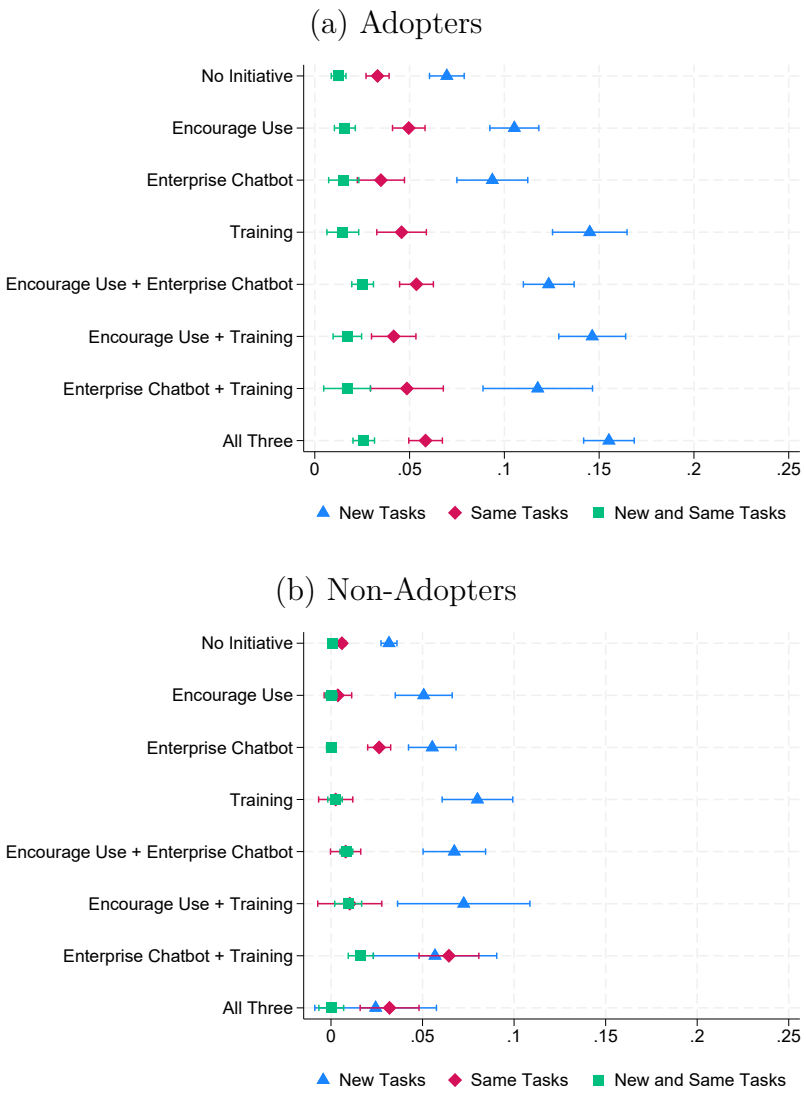
Notes: This figure illustrates how gender gaps in AI chatbot adoption varies with employer initiatives. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives. Worker characteristics include gender, age, experience, earnings, and net wealth; see Table C.3 for the full set of regression estimates. Whiskers represent 95% confidence intervals. *Sample*: All completed responses from the 2024 survey linked to registry data.

Figure 5: Employer Initiatives and Reported Time Savings



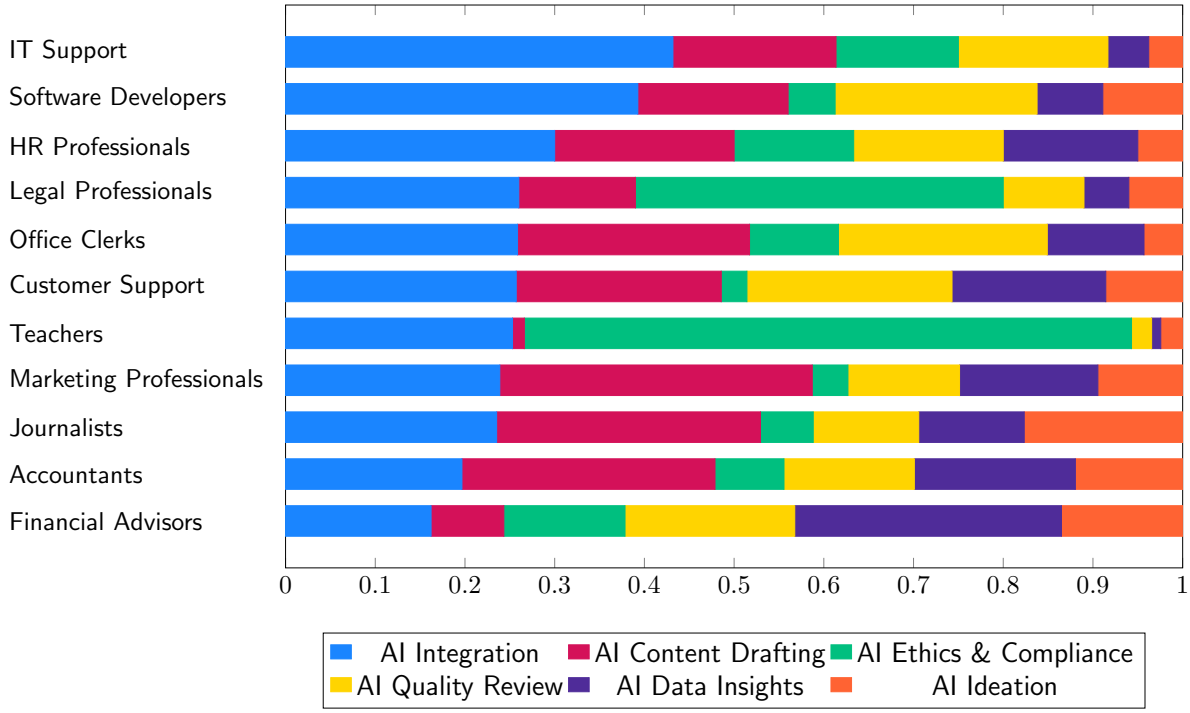
Notes: This figure illustrates show workers' time savings report from using AI chatbots vary with their employer initiatives. The estimates are based on predicted values from Equation (1), varying EmployerInitiatives while holding workers' characteristics X at their mean values; see Table D.2 for the full set of regression estimates. Whiskers represent 95% confidence bands of the predicted values. Some confidence intervals are too narrow to be distinguished from the point estimates at this scale. *Sample*: All completed responses from the 2024 survey linked to registry data.

Figure 6: Employer Initiatives and New Workloads From AI Chatbots



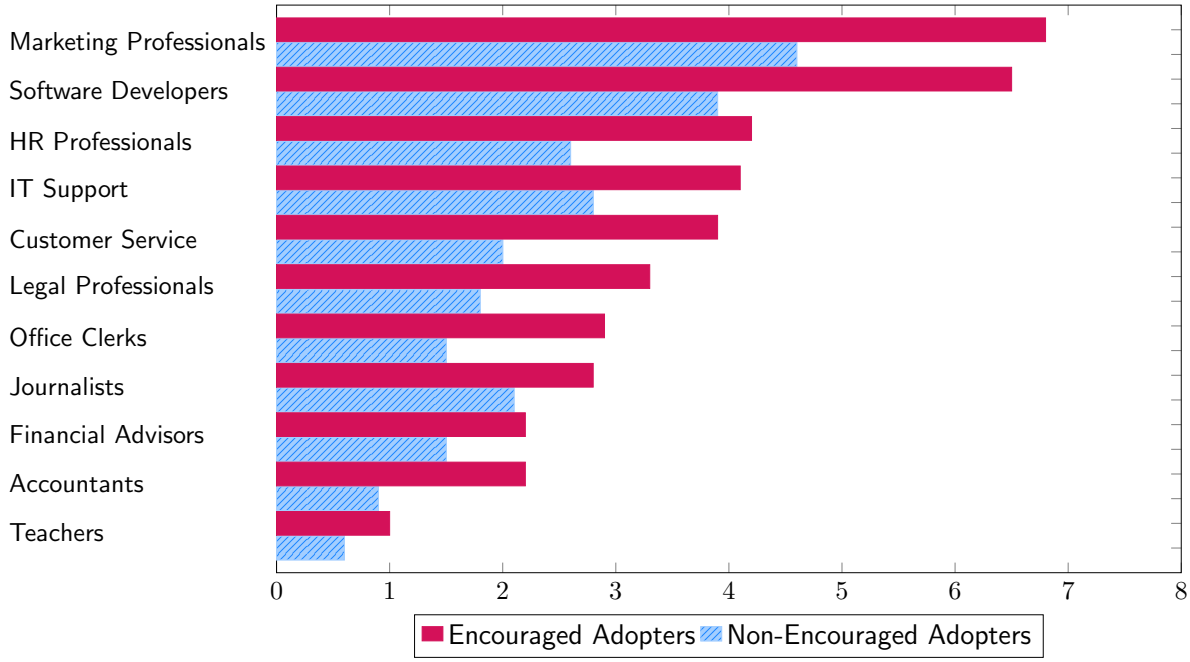
Notes: This figure illustrates how new workloads created by AI chatbots vary with employer initiatives. Panel (a) presents results for adopters, while Panel (b) focuses on non-adopters (workers who have never used AI chatbots for work). Estimates are predicted values from Equation (1), varying EmployerInitiatives while holding workers' characteristics X at their mean values; see Table D.4 for the full set of regression estimates. Whiskers represent 95% confidence bands of the predicted values. Some confidence intervals are too narrow to be distinguished from the point estimates at this scale. *Sample:* All completed responses from the 2024 survey linked to registry data.

Figure 7: Composition of AI Tasks



Notes: This figure shows the distribution of reported new job tasks across major task categories for each occupation. *AI Ideation* refers to "using AI to generate or expand creative ideas, such as concepts, strategies, or solutions," *AI Content Drafting* refers to "using AI tools to produce initial drafts of text or media," *AI Quality Review* refers to "reviewing and correcting AI-generated content for accuracy, clarity, and relevance," *AI Data Insights* refers to "using AI to analyze data or documents and extract key patterns or insights," *AI Integration* refers to "embedding AI into workflows to automate or enhance tasks," and *AI Ethics & Compliance* refers to "ensuring AI use follows ethical, legal, and organizational standards." See Appendix D.2 for details and occupation-specific examples. Tasks are ordered according to their average shares among the eleven occupations. Occupations are ordered according to their shares on *AI Integration*, the most frequent tasks across the occupations. *Sample*: All completed responses from the 2024 survey who reported new job tasks due to AI chatbots.

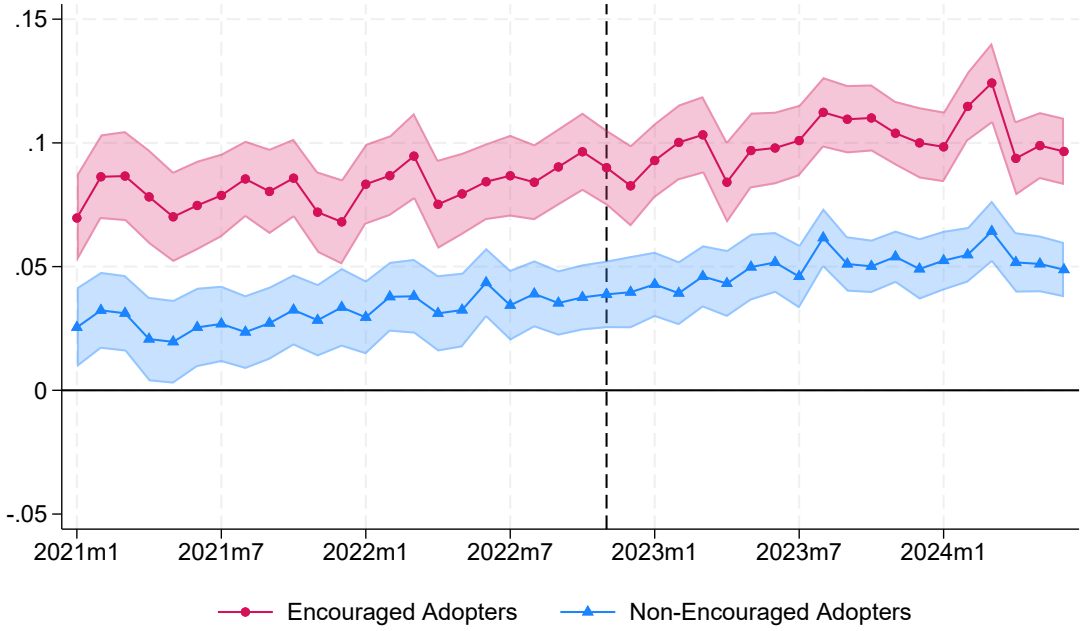
Figure 8: Total Time Savings from AI Chatbots in % of Work Hours



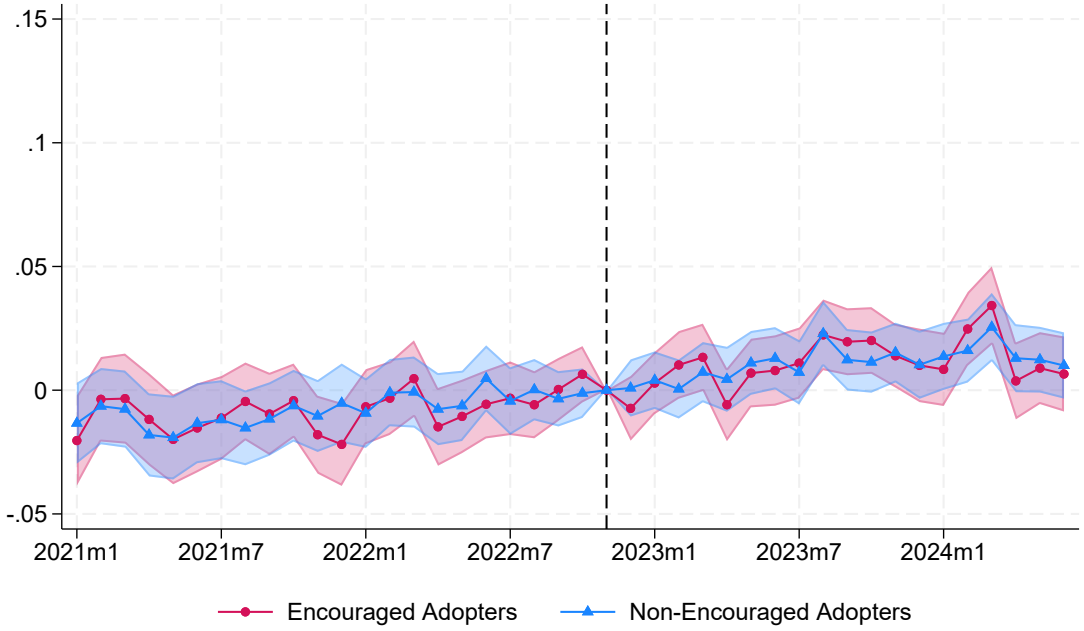
Notes: This figure presents workers' time savings from AI chatbots, categorized by workers' occupations and employer policies. The table focuses on workers who have ever used AI chatbots for work. Occupations are sorted by their average time savings among encouraged users. We estimate total time savings by combining workers' frequency of use (from Figure 3) with their reported time savings per day of usage (from Figure 5). Table D.6 decomposes the time savings into each of its components: frequency of use, time savings per day of usage, and the covariance between the two. The table notes provide additional details on how time savings and usage frequencies are coded. *Sample*: The table is based on all completed responses from the 2024 survey that can be linked to the registry data.

Figure 9: Have Adopting Workers Experienced Earnings Gains?
(Log Earnings Relative to Non-Encouraged Non-Adopters)

(a) Differences in Means

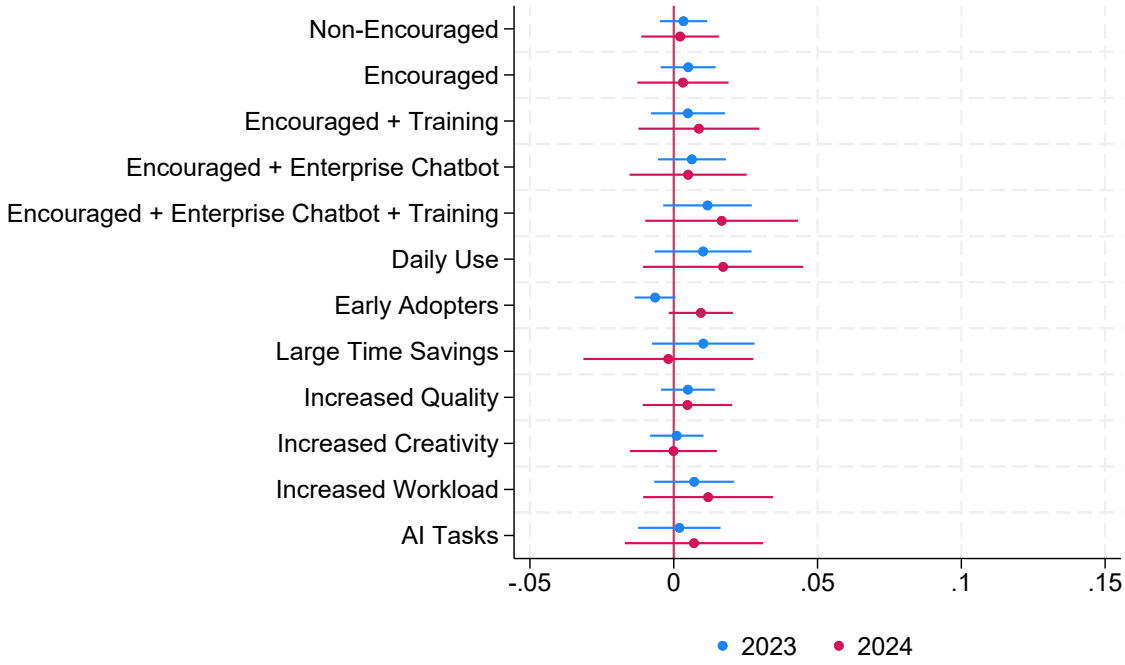


(b) Difference-in-Differences (DiD)



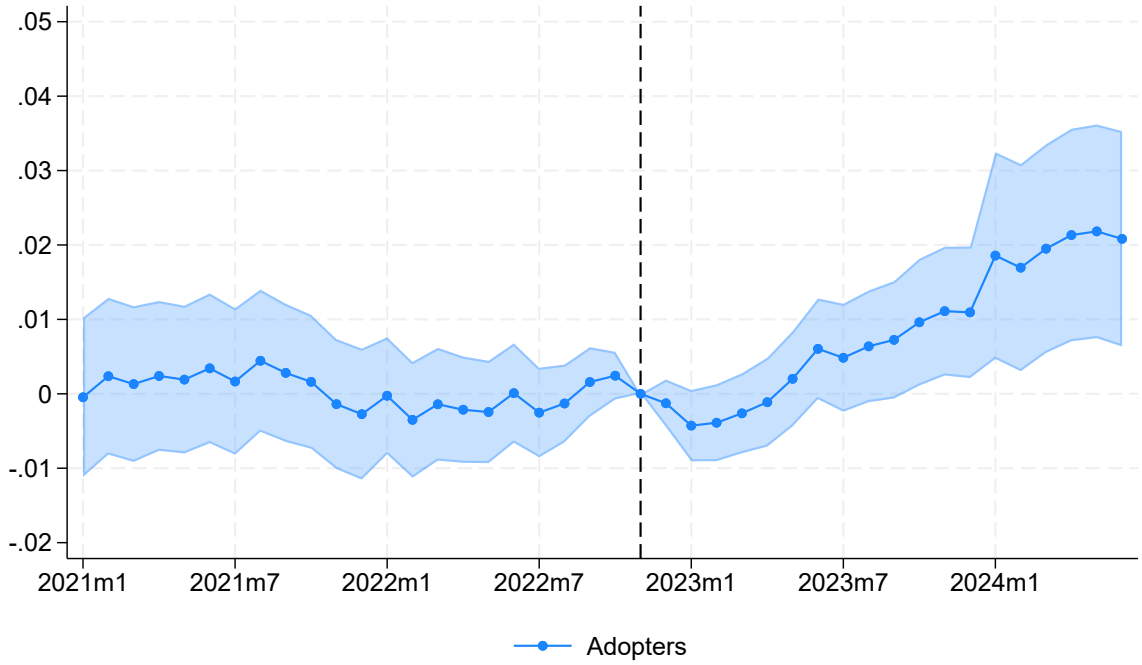
Notes: This figure shows the earnings gap of chatbot adopters, distinguishing between those with and without employer encouragement. Each group is compared to non-encouraged non-adopters. Panel (a) reports the difference in log earnings between AI chatbot adopters and non-adopters, controlling for occupation fixed effects, pre-determined worker characteristics, and seasonality. Panel (b) presents the difference-in-differences that correspond to Panel (b), indexed to November 2022, the launch date for ChatGPT. The difference-in-difference estimates are based on the specification in Equation (2). Shaded areas represent 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

Figure 10: Earnings Effects of Adoption: Heterogeneity by Strength of Treatment (Log Earnings; DiD to Non-Encouraged Non-Adopters)



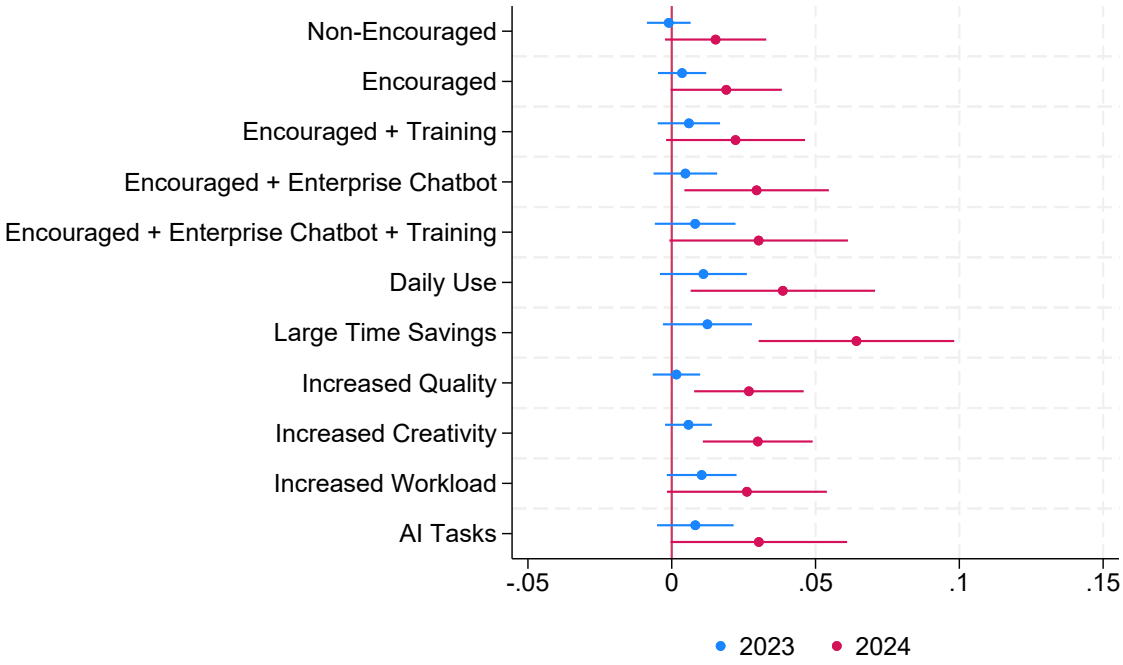
Notes: This figure shows difference-in-differences estimates of log earnings for encouraged adopters relative to non-encouraged non-adopters, indexed to the launch of ChatGPT in November 2022. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Encouraged + Training* restricts the treatment group to encouraged adopters who received employer-provided training; *Encouraged + Enterprise Chatbot*, to those whose employer deployed an enterprise chatbot; and *Encouraged + Enterprise Chatbot + Training*, to those with both. *Daily Use* refers to users reporting daily chatbot use at work. *Early Adopters* adopted ChatGPT by our 2023 survey. *Large Time Savings* indicates users reporting time savings of over 60 minutes per day. *Increased Quality / Creativity* refers to users reporting improved output quality or creativity. *Incr. Workload* refers to users reporting new tasks due to AI chatbots. *AI Task* restricts to workers assigned AI-related job tasks. *Sample:* All completed survey responses linked to registry data.

Figure 11: Have Adopters Experienced Greater Occupational Mobility?
(FTE Employment in Latest Occupation; DiD to Non-Encouraged Non-Adopters)



Notes: This figure shows difference-in-differences estimates of work hours in workers' latest occupations. The estimates compare changes for all adopters relative to non-encouraged non-adopters, indexed to the launch of ChatGPT in November 2022. Estimates are based on the dynamic difference-in-differences specification in Equation (2). Shaded areas represent 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

Figure 12: Mobility Effects of Adoption: Heterogeneity by Strength of Treatment (FTE Employment in Latest Occupation; DiD to Non-Encouraged Non-Adopters)



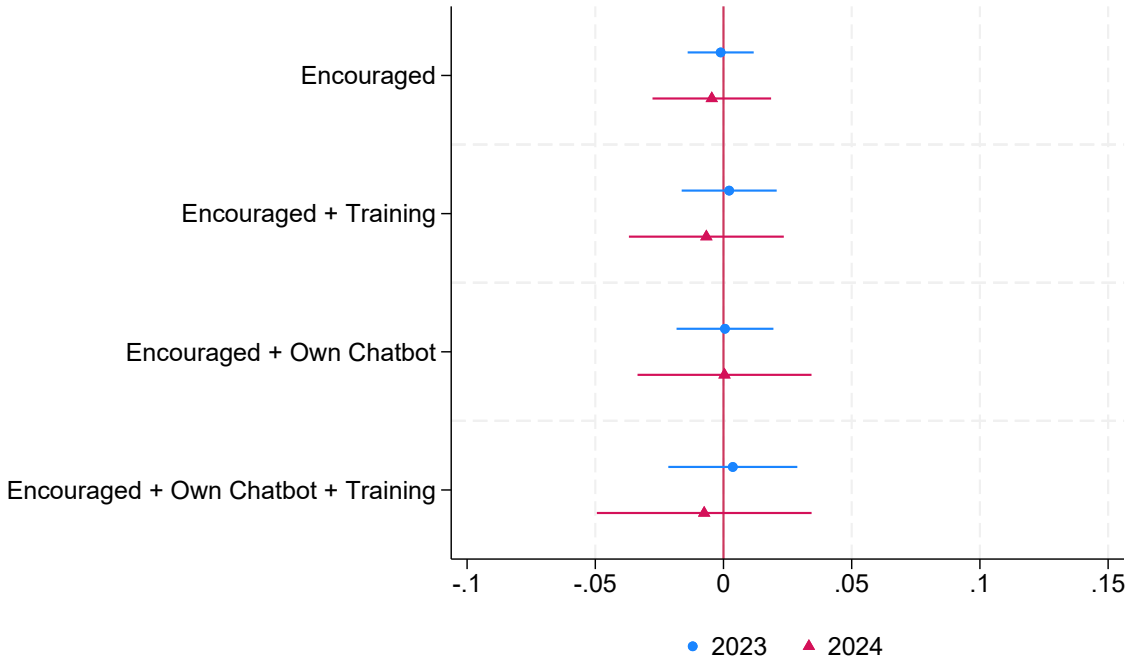
Notes: This figure shows difference-in-differences estimates of work hours in workers' latest occupations. The estimates compare changes for adopters relative to non-encouraged non-adopters, indexed to the launch of ChatGPT in November 2022. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Encouraged + Training* restricts the treatment group to encouraged adopters who received employer-provided training; *Encouraged + Enterprise Chatbot*, to those whose employer deployed an enterprise chatbot; and *Encouraged + Enterprise Chatbot + Training*, to those with both. *Daily Use* refers to users reporting daily chatbot use at work. *Early Adopters* adopted ChatGPT by our 2023 survey. *Large Time Savings* indicates users reporting time savings of over 60 minutes per day. *Increased Quality / Creativity* refers to users reporting improved output quality or creativity. *Incr. Workload* refers to users reporting new tasks due to AI chatbots. *AI Task* restricts to workers assigned AI-related job tasks. *Sample:* All completed survey responses linked to registry data.

Figure 13: Have Encouraged Workplaces Fared Differently? (Encouraged DiD)



Notes: This figure shows workplace outcomes for workers in encouraged-use workplaces relative to those without encouragement, indexed to November 2022. Panel (a) reports log wage bill; Panel (b), log employment; Panel (c), job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment); and Panel (d), the adopter propensity of hires and separations (see Appendix G.1 for details). Estimates are based on the dynamic difference-in-differences specification in Equation (2), with shaded areas representing 95% confidence intervals. *Sample:* All completed 2024 survey responses linked to registry data.

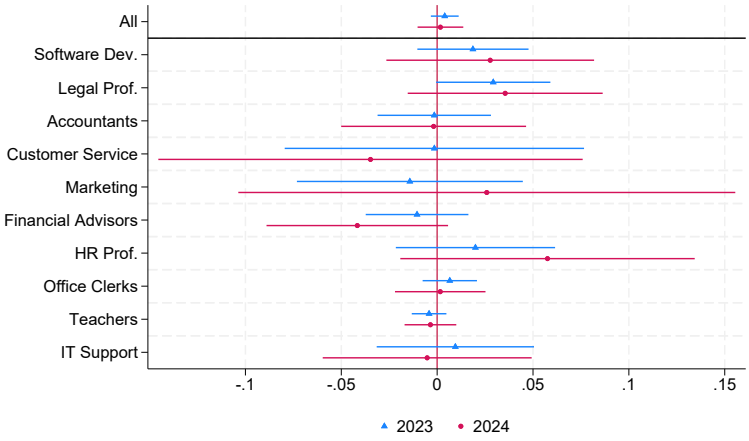
Figure 14: Workplace Effects of Chatbot Initiatives:
Heterogeneity by Strength of Treatment (DiD to Non-Encouraged)



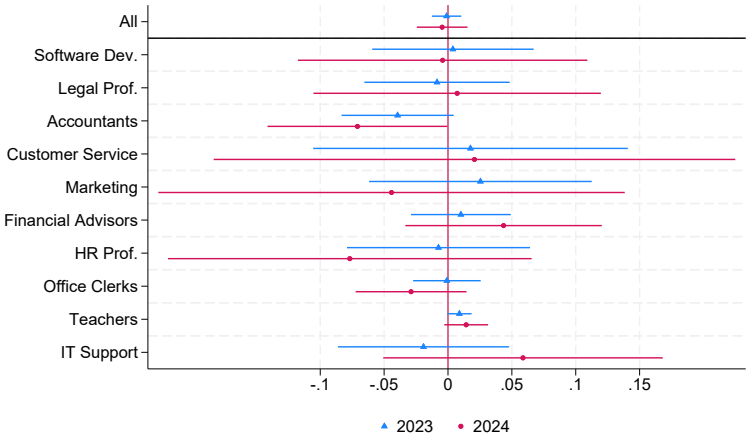
Notes: This figure shows workplace outcomes (log wage bill and log employment) for workers who report encouraged usage policies relative to those without encouragement, indexed to November 2022. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Encouraged + Training* restricts the treatment group to encouraged workers who received employer-provided training; *Encouraged + Enterprise Chatbot*, to those whose employer deployed an enterprise chatbot; and *Encouraged + Enterprise Chatbot + Training*, to those with both. *Sample:* All completed 2024 survey responses linked to registry data.

Figure 15: Labor Market Effects of AI Chatbots: Heterogeneity by Occupation

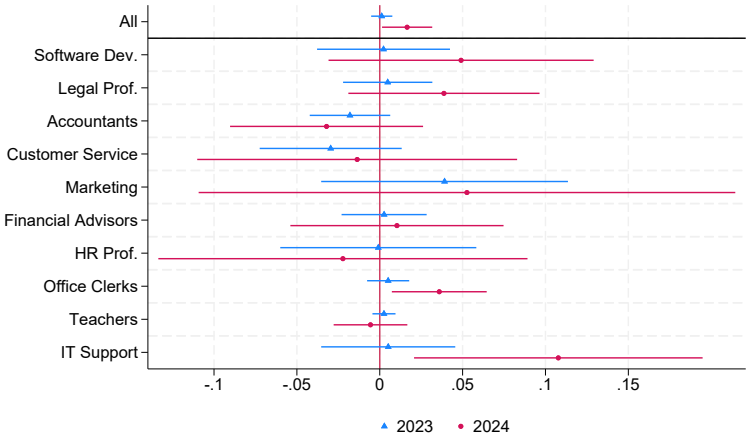
(a) Worker Earnings (Adopter DiD)



(b) Workplace Employment (Encouraged DiD)



(c) Occupational Mobility (Adopter DiD)



Notes: This figure shows occupation-specific effects of AI chatbot adoption. Panel (a) reports difference-in-differences estimates for adopters relative to non-encouraged adopters (see Figure 10); Panel (b) reports estimates for encouraged workplaces (see Figure 14). Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Sample:* All completed 2024 survey responses linked to registry data.

Table 1: Which Workplaces Have Adopted AI Chatbot Initiatives?

	Encouraged (1)	Allowed (2)	Not Allowed (3)	No Policy (4)	Don't Know (5)	Firm-Arranged Training (6)	Enterprise Chatbot (7)
Firm Age (10 Years)	-0.0140**	0.0030	0.0028	0.0038*	0.0044	-0.0122**	-0.0091*
	(0.0044)	(0.0021)	(0.0026)	(0.0016)	(0.0025)	(0.0039)	(0.0043)
log(Firm Employment)	0.0036	-0.0082***	0.0036	-0.0195***	0.0204***	-0.0048	0.0362***
	(0.0048)	(0.0023)	(0.0023)	(0.0022)	(0.0029)	(0.0043)	(0.0056)
log(Firm Labor Productivity)	0.0468**	0.0275***	0.0051	-0.0374***	-0.0420***	0.0765***	0.0697***
	(0.0153)	(0.0077)	(0.0051)	(0.0077)	(0.0084)	(0.0135)	(0.0176)
Private Firm	0.0414*	-0.0076	-0.0141*	-0.0104	-0.0092	0.0257	0.0554**
	(0.0168)	(0.0094)	(0.0072)	(0.0066)	(0.0096)	(0.0144)	(0.0180)
Occupation FE's	✓	✓	✓	✓	✓	✓	✓
Mean of Outcome	0.419	0.209	0.061	0.132	0.178	0.231	0.375
Within R^2	0.013	0.003	0.009	0.016	0.016	0.010	0.042
R^2	0.106	0.005	0.032	0.046	0.071	0.023	0.123
Observations	24184	24184	24184	24184	24184	24184	24184

Notes: This table examines which firm characteristics predict the adoption of various employer initiatives to promote AI chatbot use. Columns (1)–(5) report results for employer usage policies, while Columns (6)–(7) report results for enterprise chatbots and firm-provided training. Firm characteristics are measured in 2021, our latest data year available. Labor productivity is measured as value added per full-time equivalent worker. The regressions control for whether the firm reports value added. The regressions also control for occupation fixed effects, and Table C.1 shows that the estimates are robust to adding controls for worker characteristics. Standard errors, reported in parentheses, are clustered at the firm level. *Sample*: The table is based on all completed responses from the 2024 survey that can be linked to the registry data.

Table 2: Employer Initiatives and Reported Benefits by Adopters

	(1) Time Savings	(2) Quality	(3) Creativity	(4) Job Satisfaction	(5) No Benefits	(6) Don't Know
Encouraged	0.087***	0.076***	0.065***	0.065***	-0.038***	-0.027***
	(0.011)	(0.013)	(0.013)	(0.010)	(0.008)	(0.006)
Enterprise Chatbot	-0.000	-0.033*	-0.016	-0.003	0.027**	-0.016*
	(0.015)	(0.017)	(0.017)	(0.014)	(0.010)	(0.007)
Employer Training	-0.065***	-0.041*	-0.031	-0.008	0.086***	-0.010
	(0.015)	(0.017)	(0.017)	(0.014)	(0.010)	(0.007)
Encouraged × Enterprise Chatbot	0.009	0.045*	0.009	-0.007	-0.018	0.015
	(0.019)	(0.022)	(0.023)	(0.018)	(0.013)	(0.010)
Encouraged × Employer Training	0.028	0.066**	0.079**	0.014	-0.058***	0.005
	(0.021)	(0.024)	(0.025)	(0.020)	(0.014)	(0.011)
Enterprise Chatbot × Employer Training	0.040	0.095**	0.022	0.013	-0.076***	0.017
	(0.028)	(0.032)	(0.033)	(0.026)	(0.019)	(0.014)
Encouraged × Enterprise Chatbot × Employer Training	0.023	-0.047	-0.020	0.013	0.029	-0.021
	(0.035)	(0.039)	(0.040)	(0.032)	(0.023)	(0.017)
No Initiative, Level	0.702	0.460	0.448	0.165	0.102	0.056
Worker Controls	✓	✓	✓	✓	✓	✓
Occupation FEs	✓	✓	✓	✓	✓	✓

Notes: This table shows how self-reported benefits from AI chatbot use vary by employer initiative. The table focuses on workers who have ever used AI chatbots for work. Estimates are based on Equation (1), with standard errors in parentheses. *Sample*: All completed 2024 survey responses linked to registry data.

Online Appendix
Large Language Models, Small Labor Market Effects

Anders Humlum
University of Chicago

Emilie Vestergaard
University of Copenhagen

A Theoretical Framework **3**

A.1 Adoption and Work 3

A.2 Labor Market Outcomes 8

B Data and Institutional Setting **27**

B.1 Wage Setting Systems in Denmark 27

B.2 Sampling Protocol 29

B.3 Survey Sample 30

C Adoption **35**

C.1 Employer Initiatives 35

C.2 Worker Adoption 37

D Work **49**

D.1 Benefits for Users 49

D.2 Workloads and Task Creation 52

D.3 Total Time Savings 60

E Labor Market Outcomes **62**

E.1 Results 62

E.2 Supplementary Evidence 69

F Robustness Analysis **73**

F.1 Additional Controls 73

F.2 Coworker Encouragement 76

G Estimation Procedures **82**

G.1 Adopter Propensity 82

G.2 Empirical Bayes Shrinkage 83

H Invitation Letter **86**

H.1 English Translation 87

H.2 Original Danish Version 89

I Survey Questionnaire **91**

I.1 English Translation 92

I.2 Original Danish Version 97

A Theoretical Framework

In this section, we introduce a theoretical framework to guide our empirical analysis.

In Section A.1, we first introduce a simple partial-equilibrium model of chatbot adoption to interpret the empirical patterns in Sections 2 and 3. We posit a Roy (1951) style selection model of chatbot adoption, in which employer initiatives may shift both the costs and benefits of adoption. We use the model to analyze how these shifts affect workers' adoption decisions and the work-related benefits they report from using the tools.

In Section A.2, we embed the adoption model within a labor market equilibrium model, in which differentiated firms adopt chatbot initiatives and hire workers to maximize profits. We use this model to guide and interpret the empirical analysis of labor market outcomes in Section 4.

A.1 Adoption and Work

A.1.1 Setup

Worker i derives benefits B_i from using AI chatbots:

$$B_i = \alpha_b + \beta_b \times i, \tag{4}$$

where individuals $i \sim \mathcal{U}([0,1])$ are drawn uniformly from the unit interval and sorted by their individual benefits, such that $\beta_b \leq 0$. While we assume linear functional forms for analytical simplicity, our predictions also hold under more general monotonicity assumptions.

The worker also incurs a cost C_i of adopting chatbots:

$$C_i = \alpha_c + \beta_c \times i, \tag{5}$$

where $\beta_c \geq 0$ if workers who benefit more from AI also face lower adoption costs.[40]

[40] A common finding from RCTs is that AI chatbots tend to yield greater benefits for less experienced

We conceptualize employer initiatives $E \in \{0, 1\}$ as shifting workers' costs and benefits of using AI chatbots through the parameters α_c and α_b:

$$\alpha = \alpha_0 + \alpha_1 E. \tag{6}$$

While we focus on a binary initiative (e.g., encouraged use vs. not) for analytical simplicity, our propositions naturally extend to combinations of employer initiatives—such as use policies, enterprise chatbots, and training.

Workers adopt AI chatbots if their perceived benefits exceed the associated adoption costs ($B_i \geq C_i$):

$$A_{ij} = \mathbf{1}\left[B_{ij} \geq C_{ij}\right] = \mathbf{1}\left[\hat{\alpha}_0 + \hat{\alpha}_1 E_j + \hat{\beta} \times i \geq 0\right], \tag{7}$$

where $\hat{x} = x_b - x_c$ denotes net benefits.

A.1.2 Model Predictions

Figure A.1 illustrates how employer initiatives shift adoption rates and the average benefits reported by workers; formal derivations follow below.

Adoption rates. Figure 3 shows how chatbot adoption rates vary with employer initiatives. In the Roy model outlined above, the optimal adoption rate is given by:

$$i^*(E) = \mathbf{P}\left[B_{ij} \geq C_{ij}\right] = -\frac{\hat{\alpha}}{\hat{\beta}}, \tag{8}$$

where $0 < \hat{\alpha} < -\hat{\beta}$, corresponding to the empirically relevant case in which adoption rates lie between 0 and 1.

Equation (8) implies that employer initiatives which raise net benefits ($\hat{\alpha}$) increase adoption, with stronger effects when $\hat{\beta}$ is numerically small—i.e., when workers are

(and thus typically younger) workers (Brynjolfsson, Li and Raymond, 2025; Dell'Acqua et al., 2023; Noy and Zhang, 2023). If these workers also find it easier to adopt new tools, then $\beta_c \geq 0$.

relatively homogeneous in their net benefits of using AI chatbots.

In this light, the fact that employer-encouraged use in Figure 3 is associated with a *doubling* of adoption rates suggests that this initiative is powerful relative to the amount of individual heterogeneity in net benefits.

Benefits for users. Figure 5 and Table 2 show how reported benefits among chatbot users vary with employer initiatives. The average benefit reported by users is:

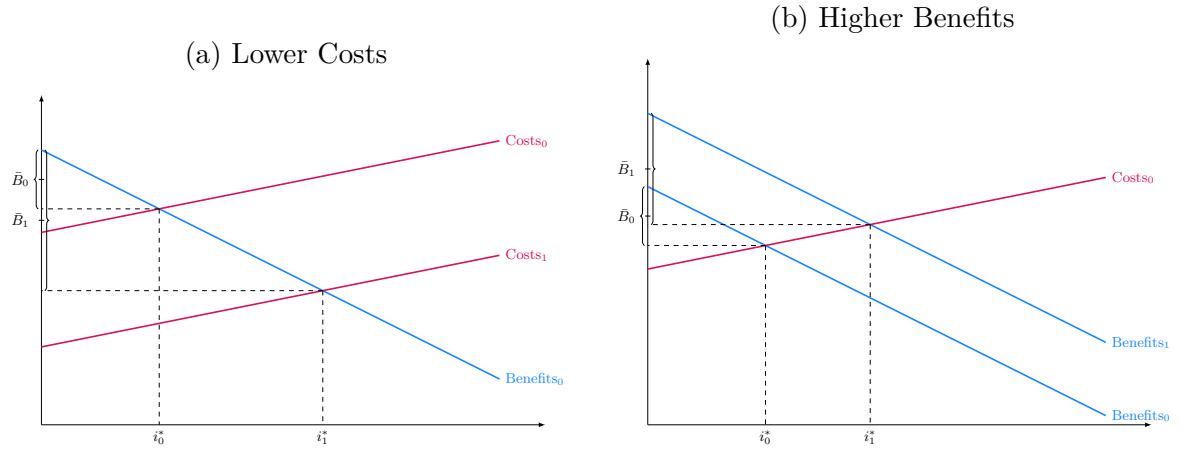
$$\bar{B}(E) = \mathbb{E}[B_i(\alpha) \mid i \leq i^*(\alpha)] = \gamma\alpha_c + (1-\gamma)\alpha_b, \tag{9}$$

where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$. Equation (9) provides two key insights about how employer initiatives relate to reported user benefits.

First, employer initiatives that reduce adoption costs (α_c) *lower* the average benefit among users due to a *selection effect*: lower costs induce adoption by workers with lower idiosyncratic benefits (i.e., higher i). This negative selection effect, γ, is stronger when workers are more homogeneous in their net benefits from using AI chatbots (i.e., when $\hat{\beta}$ is numerically small, so that employer initiatives generate larger increases in adoption; cf. Equation (8)) but more heterogeneous in their gross benefits from the tools (i.e., when β_b is large in absolute value, implying that marginal users have lower idiosyncratic benefits).

Second, employer initiatives that increase benefits (α_b) have two offsetting effects on average reported benefits. On the one hand, a higher α_b directly raises each user's individual benefit B_i. On the other hand, it expands adoption, drawing in users with lower idiosyncratic benefits (i.e., higher i), thereby triggering the same negative selection effect as a reduction in α_c. The net effect of an increase in α_b on the average reported benefits among users $\bar{B}$ is positive if and only if $\gamma < 1$. This condition is more likely to hold when workers are relatively similar in their gross benefits from the tools (i.e., when β_b is small) but differ more in their adoption costs (i.e., when β_c is large).

Figure A.1: The Impact of Employer Initiatives on Adoption Rates and Reported Benefits



Notes: This figure illustrates how employer initiatives affect adoption rates (i^*) and average reported benefits among users ($\bar{B}$). Panel (a) shows the impact of initiatives that reduce adoption costs (α_c), while Panel (b) depicts the impact of initiatives that enhance benefits (α_b). Lower adoption costs (α_c) unambiguously increase adoption rates and reduce average reported benefits due to a negative selection effect. In contrast, initiatives that raise benefits (α_b) also increase adoption, but their impact on average reported benefits is theoretically ambiguous and, at most, equal to their individual-level effect.

In summary, the average reported benefits among users provide a lower bound on the causal effect of employer initiatives on *individual* benefits, due to the presence of a negative *selection* effect.

In this light, the fact that Figure 5 and Table 2 show higher reported benefits among users in encouraged settings is striking. It suggests two things: (i) employer encouragement primarily raises perceived benefits (α_b), as reductions in adoption costs (α_c) would only generate the negative selection effect; and (ii) the employer-induced increases in benefits outweigh the adverse selection effect—implying that employer encouragement is powerful relative to degree of individual heterogeneity ($\gamma < 1$).

By contrast, the finding that enterprise chatbots and training—when implemented in isolation—are associated with lower reported benefits may reflect either that these initiatives reduce benefits for all users ($\alpha_{b1} < 0$), or that negative selection effects outweigh any positive individual-level gains.

A.1.3 Discussion

The preceding section examines the implications of workers self-selecting into chatbot adoption based on their individual benefits. Beyond this margin of selection, at least two additional forms of selection may be at play.

First, firms that choose to encourage chatbots may be the ones with the greatest benefits from the tools—even absent any encouragement initiatives—introducing potential reverse causality between employer encouragement and reported benefits. In the framework above, this corresponds to E being positively related to α_{0b}. Second, employers may target their encouragement toward employees who are expected to benefit the most individually. This would resemble a "rotation" of the cost and benefit schedules in Figure A.1, rather than simple vertical shifts.

Section F addresses both of these additional selection margins empirically. First, Section F.1 examines whether employers select into encouragement based on their underlying benefits from chatbot use. We show that the effects of encouragement remain robust when controlling for a range of firm- and worker-level characteristics, including firm age, size, and productivity, as well as workers' detailed task mixes and other attributes. These robustness checks show that the association between employer encouragement and reported benefits is not driven by observable confounders.

Second, Section F.2 implements a "coworker encouragement" design, using the average encouragement reported by a worker's colleagues to measure workplace-wide encouragement rates. Because coworker encouragement rates capture only *workplace*-level variation, this approach helps address potential bias arising from employers targeting encouragement toward *individual* employees with the highest idiosyncratic benefits. As shown in Table F.3 and Figure F.2, the findings remain largely robust when using coworker encouragement as the source of variation.

A.2 Labor Market Outcomes

This section examines how AI chatbots affect labor market outcomes and how these effects can be identified empirically. To this end, we embed the Roy adoption model from Section A.1 into a simple equilibrium framework in which heterogeneous firms adopt chatbot initiatives and hire workers to maximize profits, while workers sort across firms and adopt chatbots to maximize utility. For clarity of exposition, we focus on encouraged use as the employer initiative, reflecting its central role in the empirical analyses in Section 4.

A.2.1 Setup

Productivity and earnings. Worker i's earnings from working at firm j is the product of a firm-specific skill price W_j and the worker's efficiency units of human capital H_{ij}:

$$Y_{ij}(A_{ij}) = W_j H_{ij}(A_{ij}). \tag{10}$$

Workers are endowed with a unit of baseline human capital, which may be augmented through chatbot use, $A_{ij} \in \{0, 1\}$:

$$H_{ij}(A_{ij}) = \exp(B_{ij} A_{ij}), \tag{11}$$

where B_{ij} captures the productivity benefit of AI chatbots.

As in Section A.1, the productivity benefit and adoption cost depends on worker i's type and firm j's chatbot initiatives:

$$B_{ij} = \alpha_{b0} + \alpha_{b1} E_j + \beta_b \times i \tag{12}$$

$$C_{ij} = \alpha_{c0} + \alpha_{c1} E_j + \beta_c \times i. \tag{13}$$

Firm production and demand. Firms produce using efficiency units of labor:

$$Y_j = H_j = \int_{i \in \mathcal{I}(j)} H_{ij} di, \tag{14}$$

where $\mathcal{I}(j)$ denotes the set of workers employed at firm j.

The production function in Equation (14) models AI chatbots as tools that affect output only when actively deployed by workers. While we view this as a realistic representation of current chatbot technology—which depends on direct human prompting—Section A.2.4 examines the implications of a more general technology where chatbots may affect production beyond individual use.

Product markets are characterized by monopolistic competition, where atomistic and horizontally differentiated firms hire labor and produce output to maximize profits. Firms face the product demand curves

$$Y_j = P_j^{-\epsilon}, \tag{15}$$

where ϵ is the elasticity of demand across firms within the sector.

Labor demand. Labor markets are perfectly competitive, so workers are paid according to the marginal revenue product of labor:

$$W_j = \frac{\partial P_j Y_j}{\partial H_j} = (1 - 1/\epsilon) Y_j^{-\frac{1}{\epsilon}}. \tag{16}$$

Since firms' production and demand curves are symmetric, and workers derive the same net benefit $\hat{\alpha}_1$ from employer encouragement across employers, two skill prices W_j exist in equilibrium: one at encouraged employers W_1 and one at non-encouraged employers W_0. Let $j \in \{0, 1\}$ indicate whether a firm encourages chatbot use.

Equilibrium skill prices adjust to equate the supply and demand for human capital across encouraged and non-encouraged employers. These supply and demand forces reflect workers' sorting and adoption decisions, as well as firms' encouragement and hiring policies. We characterize equilibrium skill prices in Section A.2.2.

Employer encouragement. Employers derive log utility from flow profits and encourage AI chatbots if the utility from the resulting profit gain exceeds their encouragement cost κ_j:

$$E_j = \mathbf{1}\left[\kappa_j \leq \log \Pi_j(1) - \log \Pi_j(0)\right], \tag{17}$$

where $\Pi_j(E)$ denotes firm j's flow profits under encouragement policy $E \in \{0, 1\}$. Let F_κ denote the CDF of encouragement costs.

Constant returns to scale (Equation (14)) and CES demand (Equation (15)) imply that a firm's optimal flow profit is a constant fraction $\frac{1}{\epsilon}$ of revenues:

$$\Pi_j = \frac{1}{\epsilon} P_j Y_j = \frac{1}{\epsilon} Y_j^{1-1/\epsilon}. \tag{18}$$

Combining Equations (16) and (18), the profit gain from encouraging chatbots equals:

$$\log \Pi_j(1) - \log \Pi_j(0) = (1-\epsilon)(\log W_1 - \log W_0) \tag{19}$$

Since the profit gain is common across employers, the encouragement policy follows a cutoff rule based on encouragement costs:

$$E_j = \mathbf{1}\left[\kappa_j \leq \kappa\right], \quad \text{with} \quad \kappa = (1-\epsilon)(\log W_1 - \log W_0). \tag{20}$$

Adoption. Workers derive log-utility from labor earnings and adopt AI chatbots if the monetary benefit exceeds the adoption cost C_{ij}. Fixing the worker's employer, the adoption rule is:

$$A_{ij} = \mathbf{1}\left[\log Y_{ij}(1) - \log Y_{ij}(0) \geq C_{ij}\right] \tag{21}$$

$$= \mathbf{1}\left[B_{ij} \geq C_{ij}\right] \tag{22}$$

$$= \mathbf{1}\left[\hat{\alpha}_0 + \hat{\alpha}_1 E_j + \hat{\beta} \times i \geq 0\right], \tag{23}$$

where $\hat{x} = x_b - x_c$ denotes net benefits, resulting in the adoption rule from Section A.1.2. This implies a strict cutoff rule for adoption:

$$A_{ij} = \mathbf{1}\left[i \leq i^*(E_j)\right], \quad \text{with} \quad i^*(E) = -\frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 E). \tag{24}$$

Sorting. Workers not only decide whether to adopt chatbots, but also which firms to sort into. Since all adopters receive the same net benefit from encouragement, $\hat{\alpha}_1$, workers effectively face a choice between sorting into encouraged firms and adopting AI chatbots, or sorting into non-encouraged firms and not adopting. The condition for sorting into encouraged firms is given by:

$$E_{j(i)} = \mathbf{1}\left[\log Y_{i1}(1) - \log Y_{i0}(0) \geq C_{i1}\right] \tag{25}$$

$$= \mathbf{1}\left[\hat{\alpha}_0 + \hat{\alpha}_1 + \hat{\beta} \times i + \log W_1 - \log W_0 \geq 0\right]. \tag{26}$$

This implies a strict cutoff rule for worker sorting:

$$E_{j(i)} = \mathbf{1}\left[i \leq i^*\right], \quad \text{with} \quad i^* = -\frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_1 - \log W_0). \tag{27}$$

A.2.2 Equilibrium

An equilibrium is a set of encouragement policies $\{E_j\}_{j\in\mathcal{J}}$, skill prices $\{W_j\}_{j\in\mathcal{J}}$, allocation of workers to firms $\{\mathcal{I}_j\}_{j\in\mathcal{J}}$, and adoption decisions $\{\{A_{ij}\}_{i\in\mathcal{I}_j}\}_{j\in\mathcal{J}}$, such that firms encourage AI chatbots and hire workers to maximize profits (Equations (16) and (17)), workers sort into firms and adopt AI chatbots to maximize utility (Equations (24) and (27)), and labor and product markets clear (Equations (14) and (15)).

The initial equilibrium without AI chatbots is simple: because all firms are symmetric, they each employ an equal share of the baseline human capital stock. The following paragraphs characterize the equilibrium when AI chatbots are available.

Sorting. Worker sorting across firms when AI chatbots are available is straightforward: Assuming adopters receive higher benefits from encouraged firms ($\hat{\alpha}_1 \geq 0$), the equilibrium features full segregation, with all adopters in encouraged firms, all non-adopters in non-encouraged firms, or both. We focus on the empirically relevant case in which some, but not all, firms encourage AI chatbot use, and some, but not all, workers adopt.

Skill prices. Equilibrium skill prices adjust to equalize the supply and demand of human capital across encouraged and non-encouraged employers, which in turn depend on workers' sorting and adoption decision rules as well as firms' encouragement and hiring policies.

Let κ denote the encouragement cost of the marginal adopting employer. From Equation (20), equilibrium skill price differentials satisfy:

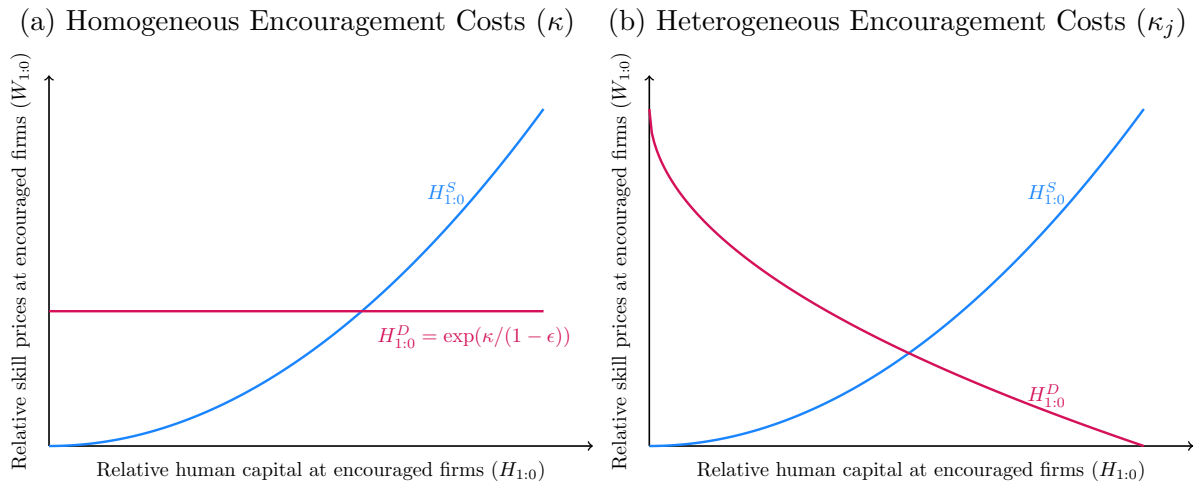
$$\log W_1 - \log W_0 = \exp\left(\kappa/(1-\epsilon)\right). \tag{28}$$

Thus, the cost of employer encouragement is priced into skill differentials: workers are not rewarded for these costly investments.

A simple case arises when all employers face the same encouragement cost, $\kappa_j = \kappa$ for all j. In this case, all firms are marginal to encouraging use, and the share of encouraged employers adjusts to match the relative supply of human capital by adopting workers. For analytical simplicity, the model predictions in Section A.2.3 focus first on the case with homogeneous encouragement costs. Section A.2.7 extends the analysis to the general case with heterogeneous costs.

Figure A.2 illustrates the inverse aggregate supply and demand curves for relative human capital at encouraged firms. Panel (a) depicts the case where firms face the same encouragement cost κ, while Panel (b) illustrates the more general case with heterogeneous encouragement costs.

Figure A.2: Equilibrium Skill Price Differentials



Notes: This figure illustrates illustrates the labor market equilibrium across encouraged and non-encouraged firms. Equilibrium skill price differentials, $W_{1:0} = W_1/W_0$, is determined by the aggregate supply $H^S_{1:0} = H^S_1/H^S_0$ and aggregate demand $H^D_{1:0} = H^D_1/H^D_0$ for relative human capital at encouraged vs. non-encouraged firms.

A.2.3 Model Predictions

In this section, we use the model to derive predictions for how the arrival of AI chatbots affects labor market outcomes, including worker sorting across firms and earnings.

Worker earnings. Combining Equations (10)–(11) yields the following expression for the log change in earnings for worker i at firm j:

$$d\log Y_{ij} = \underbrace{d\log W_j}_{\text{skill–price effect}} + \underbrace{A_i B_{ij}}_{\text{productivity effect}} \tag{29}$$

$$= d\log W_0 + E_j \frac{\kappa}{1-\epsilon} + A_i(\alpha_{b0} + \beta_b \times i) + E_j A_i \alpha_{b1}, \tag{30}$$

Equation (30) yields three key predictions about how AI chatbots affect worker earnings.

1. Workers who do not use chatbots experience an earnings loss at encouraged firms:

$$\log Y_{i1}(A_i = 0) - \log Y_{i0}(A_i = 0) = -\frac{\kappa}{\epsilon - 1} \leq 0 \tag{31}$$

2. Holding encouragement status fixed, adoption increases worker earnings by the

productivity effect:

$$\log Y_{iE}(A_i = 1) - \log Y_{iE}(A_i = 0) = \alpha_{b0} + \alpha_{b1} \times E + \beta_b \times i \tag{32}$$

3. The impact of employer encouragement on adopter earnings is theoretically ambiguous, as it depends on the relative magnitudes of the productivity and skill price effects:

$$\log Y_{i1}(A_i = 1) - \log Y_{i0}(A_i = 1) = \alpha_{b1} - \frac{\kappa}{\epsilon - 1} \gtrless 0. \tag{33}$$

Even if employer encouragement enhances the productivity benefits of chatbots ($\alpha_{b1} > 0$, as documented in Section 3), these gains need not translate into higher earnings for adopters. When encouragement is costly (high κ), product demand is less elastic (low ϵ), and productivity gains are modest (low α_{b1}), employer encouragement may reduce adopter earnings. Adopters may still prefer sorting into encouraged firms if encouragement sufficiently reduces adoption costs (i.e., if $\alpha_{c1} < \alpha_{b1} - \frac{\kappa}{\epsilon-1}$). However, these benefits will not be reflected in their earnings.

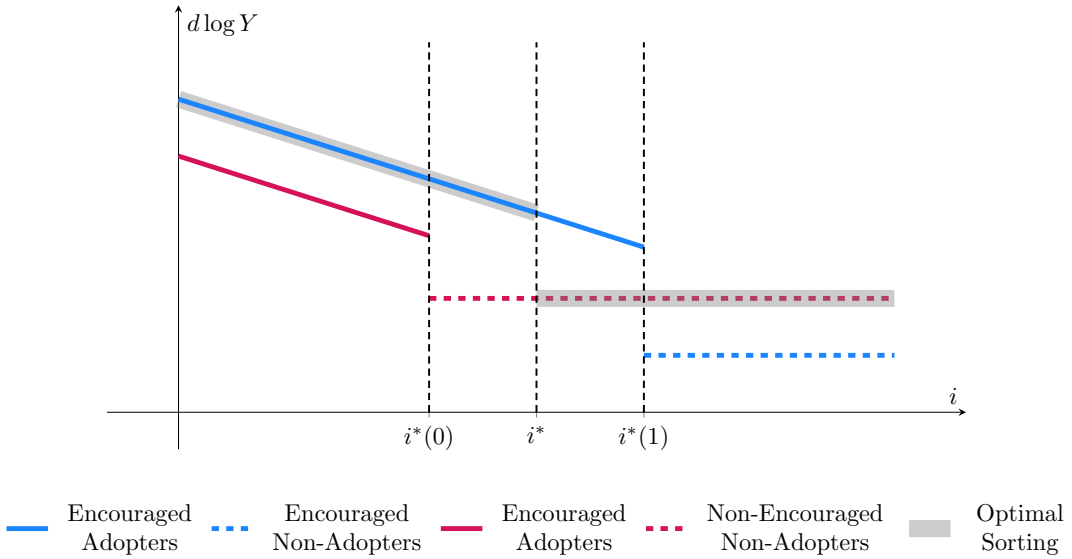
Worker-firm sorting. Following Equation (27), AI chatbots should trigger a reallocation of workers who would like to adopt chatbots (low-i types) into firms that encourage use ($E = 1$).

Worker outcomes. Figure A.3 summarizes how AI chatbots affect workers' sorting, adoption, and earnings. The grey line highlights the equilibrium sorting and adoption patterns: workers with $i < i^*$ sort into encouraged firms and adopt chatbots, while those with $i > i^*$ sort into non-encouraged firms and do not adopt.

Figure A.3 illustrates the case where employer encouragement increases adopter earnings. However, as shown above, encouragement can also reduce adopter earnings—reversing

the solid lines—depending on model parameters. In fact, if employer-encouraged chatbot use provides enough amenity value (i.e., if α_{c1} is low), some encouraged adopters may even experience earning losses relative to non-encouraged non-adopters, with the solid blue line crossing the dashed red one.

Figure A.3: The Impact of AI Chatbots on Worker Sorting, Adoption, and Earnings



Notes: This figure illustrates how AI chatbots affect workers' sorting, adoption, and earnings across the distribution of AI advantage. The x-axis ranks workers by their chatbot advantage i; the y-axis shows the resulting impact on earnings. The blue and red lines represent earnings effects for each combination of worker and firm types. The gray line highlights the optimal sorting and adoption patterns.

Workplace employment. The effect of employer encouragement on workplace employment (measured in headcounts, N) depends on whether the increased demand for human capital is offset by the productivity gains from adoption.

The increase in human capital demand is given by:

$$\log H_1 - \log H_0 = \frac{\epsilon}{\epsilon - 1}\kappa. \tag{34}$$

The average productivity gain among encouraged adopters is given by:

$$\mathbb{E}[B_{ij}(\alpha) \mid i \leq i^*, E_j = 1] = (1-\gamma)(\alpha_{b0} + \alpha_{b1}) + \gamma(\alpha_{c0} + \alpha_{c1} - (\log W_1 - \log W_0)) \quad (35)$$

$$= (1-\gamma)(\alpha_{b0} + \alpha_{b1}) + \gamma(\alpha_{c0} + \alpha_{c1} + \kappa/(\epsilon - 1)) \quad (36)$$

where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$. Equation (35) generalizes Equation (9) to account for worker sorting across firms.

Combining Equations (34) and (36), the effect of employer encouragement on workplace employment is:

$$\log N_1 - \log N_0 = \frac{\epsilon}{\epsilon - 1}\kappa - (1-\gamma)(\alpha_{b0} + \alpha_{b1}) - \gamma(\alpha_{c0} + \alpha_{c1} + \kappa/(\epsilon - 1)) \quad (37)$$

$$= \frac{\epsilon - \gamma}{\epsilon - 1}\kappa - (1-\gamma)(\alpha_{b0} + \alpha_{b1}) - \gamma(\alpha_{c0} + \alpha_{c1}) \quad (38)$$

In particular, if encouragement is relatively costless (κ is low), product demand is elastic (ϵ is high), workers are relatively homogeneous in their individual chatbot productivity (γ is low), and chatbots offer large productivity gains (α_b is high) but are costly to adopt (α_c is high), then employer encouragement will tend to reduce workplace employment.

Workplace wage bills. Encouraged firms increase their wage bills by:

$$\log(W_1 H_1) - \log(W_0 H_0) = (1-\epsilon)(\log W_1 - \log W_0) = \kappa. \quad (39)$$

The prediction that employer encouragement unambiguously raises wage bills is a consequence of modeling AI chatbots solely as tools that augment worker productivity (see Equation (14)). Section A.2.4 relaxes this assumption by considering more general production technologies, thereby breaking this unambiguous result.

A.2.4 Discussion

Our model relies on a series of simplifying assumptions to gain analytical tractability. In this section, we discuss how plausible deviations from these assumptions affect our predictions.

Labor market competition. The model assumes perfectly competitive labor markets, where workers are paid their individual marginal product of labor (Equation (16)). A commonly considered alternative is the monopsony model of Card et al. (2018), in which firms exert monopsony power due to workers' heterogeneous preferences over employers but cannot price discriminate across workers. In this setting, adopters do not experience within-firm earnings differentials. Instead, any earnings effects from AI chatbots arise at the firm level as they adjust labor demand. Thus, while we should not expect within-firm earnings gaps between adopters and non-adopters (as analyzed in Figure 9) in this setup, the effects of AI chatbots should still appear in firm-level outcomes such as total wage bills, employment, and average wages (as studied in Figure 13). In this framework, the extent to which firm-level employment effects are reflected in wages depends on how heterogeneous workers' preferences for firms are—that is, on the elasticity of their labor supply across employers. Finally, even in the monopsony setup, adopters should still reallocate toward encouraged firms, as examined in Figure 13.

Chatbot technology. Equations (11) and (14) assume that AI chatbots contribute to output only when deployed by a worker. We view this as a realistic representation of current chatbot technology, which relies on direct human prompting. The assumption also aligns with our survey design, which measures deployment through individual use.

Nonetheless, this modeling choice for how chatbots enter production omits other potential channels through which technologies can affect labor demand. On one hand, some autonomous technologies—such as industrial robots or future agentic AI systems—may

operate without human involvement. In the terminology of Acemoglu and Restrepo (2018), this corresponds to a *task displacement effect*, which reduces labor demand without directly benefiting any individual worker. On the other hand, chatbots may generate new tasks for humans, such as oversight or implementation work, even among non-users. This constitutes a *task reinstatement effect*, increasing labor demand. As Acemoglu and Restrepo (2018) show, the adoption of new technologies (e.g., encouraging AI chatbot use) will only raise the wage bill if the combined scale and reinstatement effects outweigh the negative displacement effect. Our model predictions in Section A.2.3 only capture the positive scale effect.

Mobility. The sorting patterns above assume workers move frictionlessly across firms. In the short run, however, some adopters may remain in non-encouraged firms, and some non-adopters in encouraged ones. Even so, the other lines in Figure A.3 still capture the adoption and earnings effects across worker and firm types.

A.2.5 Empirical Identification

This section demonstrates that the causal effects of AI chatbots on worker and workplace outcomes can be identified using a series of difference-in-differences estimators. We further use the model to structurally interpret these estimates.

The analytical derivations above, for simplicity, assumed that differences across workers and firms arise solely from chatbot use and employer encouragement. A more realistic setting allows for baseline heterogeneity in worker productivity ($H_{ij}(A_{ij}) = H_{i0}\exp(B_{ij}A_{ij})$) and firm-level product demand ($Y_j = Y_{j0}P_j^{-\epsilon}$). For example, Section 2 shows that workers who adopted AI chatbots were already earning more before chatbots became available, and that employers who adopted enterprise solutions were already larger prior to adoption.

Even with such heterogeneity, the causal effects of chatbot adoption can still be identified using difference-in-differences. The key assumption is that baseline heterogeneity

is time-invariant and enters the outcome equations in a log-additive way, such that it cancels out when examining log changes over time. For example:

$$\Delta H_{ij} = \log(H_{i0}\exp(B_{ij}A_{ij}^{post})) - \log(H_{i0}\exp(B_{ij}A_{ij}^{pre})) = B_{ij}(A_{ij}^{post} - A_{ij}^{pre}) \quad (40)$$

$$\Delta Y_j = -\epsilon(\log(Y_{j0}P_j^{post}) - \log(Y_{j0}P_j^{pre})) = -\epsilon(\log P_j^{post} - \log P_j^{pre}), \quad (41)$$

where $\Delta X = \log X^{post} - \log X^{pre}$ denotes log changes over time (after vs. before the arrival of AI chatbots).

More generally, if baseline heterogeneity varies over time, the identifying assumption is that workers (firms) do not adopt (encourage) chatbots in response to anticipated shocks to their baseline human capital H_{i0} (or product demand Y_{j0}). This assumption is consistent with the policy rules for adoption and encouragement (Equations (17) and (23)), which do not depend on baseline human capital or product demand.

Worker-level adoption–encouragement difference-in-differences. The earnings equation in (29) motivates the three difference-in-differences estimators used in Section 4.2.1.

1. **Productivity effect among non-encouraged adopters** (*Non-encouraged adopters vs. non-encouraged non-adopters, before vs. after AI*)

$$\begin{aligned}\Delta^{Y}_{\text{NE,A}-\text{NE,NA}} &= \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=1,\ E_j=0\Big] - \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=0,\ E_j=0\Big] \\ &= \gamma\alpha_{c0} + (1-\gamma)\alpha_{b0},\end{aligned} \quad (42)$$

   where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$ (see Equation (9)).[41]

2. **Differential skill–price effect in encouraged firms** (*Encouraged non-adopters vs.*

[41] The frictionless sorting model (Equation (27)) predicts that all adopters sort into encouraged firms, while all non-adopters sort into non-encouraged ones, leaving no workers left to identify Equation (42). In practice, however, Section 4.2 finds little evidence of worker resorting in response to chatbot adoption. For this reason, Equation (42) derives the structural interpretation without worker mobility.

non-encouraged non-adopters, before vs. after AI)

$$\begin{aligned}\Delta^{Y}_{\text{E,NA}-\text{NE,NA}} &= \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=0,\ E_j=1\Big] - \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=0,\ E_j=0\Big] \\ &= \mathbb{E}\Big[d\log W_j \;\Big|\; E_j=1\Big] - \mathbb{E}\Big[d\log W_j \;\Big|\; E_j=0\Big] \\ &= -\kappa/(\epsilon-1).\end{aligned} \tag{43}$$

3. **Combined effects for encouraged adopters** (*Encouraged adopters vs. non-encouraged non-adopters, before vs. after AI*)

$$\begin{aligned}\Delta^{Y}_{\text{E,A}-\text{NE,NA}} &= \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=1,\ E_j=1\Big] - \mathbb{E}\Big[\Delta\log Y_{ij} \;\Big|\; A_{ij}=0,\ E_j=0\Big] \\ &= -\kappa/(\epsilon-1) + \gamma(\alpha_{c0}+\alpha_{c1}) + (1-\gamma)(\alpha_{b0}+\alpha_{b1}).\end{aligned} \tag{44}$$

By contrast, when workers are perfectly mobile, Equation (36) implies that:

$$\Delta^{Y}_{\text{E,A}-\text{NE,NA}} = -\kappa/(\epsilon-1) + \gamma(\alpha_{c0}+\alpha_{c1}+\kappa/(\epsilon-1)) + (1-\gamma)(\alpha_{b0}+\alpha_{b1}). \tag{45}$$

Workplace-level difference-in-differences. Section 4.2.3 analyzes the difference-in-differences of employer encouragement on workplace wage bills and employment. Following Section A.2.3, these estimates identify:

1. **Workplace employment effect of employer encouragement.** Following Equation (38), the difference-in-differences estimate of employer encouragement on workplace employment identifies:

$$\begin{aligned}\Delta^{N}_{\text{E}-\text{NE}} &= \mathbb{E}\Big[\Delta\log N_j \;\Big|\; E_j=1\Big] - \mathbb{E}\Big[\Delta\log N_j \;\Big|\; E_j=0\Big] \\ &= \frac{\epsilon-\gamma}{\epsilon-1}\kappa - (1-\gamma)(\alpha_{b0}+\alpha_{b1}) - \gamma(\alpha_{c0}+\alpha_{c1}),\end{aligned} \tag{46}$$

2. **Workplace wage bill effect of employer encouragement.** Following Equation (39), the difference-in-differences estimate of employer encouragement on the workplace

wage bill identifies:

$$\begin{aligned}\Delta^{WH}_{\text{E}-\text{NE}} &= \mathbb{E}\Big[\Delta \log W_j H_j \;\Big|\; E_j = 1\Big] - \mathbb{E}\Big[\Delta \log W_j H_j \;\Big|\; E_j = 0\Big] \\ &= \kappa\end{aligned} \tag{47}$$

A.2.6 Missing Intercept

The difference-in-differences estimators in Equations (42)–(44) identify earnings effects up to the *skill price effect* in non-encouraged firms, $d\log W_0$. This is the familiar *missing intercept* problem: within-sector difference-in-differences do not capture general equilibrium (sector-wide) effects.

In this section, we our model's structure to characterize this residual term. To do so, we need to take a stand on how firms compete in product markets. We assume a nested CES demand function with elasticity η across and elasticity ϵ within. The product demand curve (Equation (15)) then becomes:

$$Y_j = P_j^{-\epsilon} P^{\epsilon-\eta}, \tag{48}$$

where $P = (\int P_j^{1-\epsilon} dj)^{\frac{1}{1-\epsilon}}$ is the sectoral price index.

In particular, combining this with Equations (11), (14), (16), and (48) implies that:

$$d\log W_0 = \underbrace{-\frac{1}{\epsilon} d\log Y_0}_{\text{Residual demand effect}} + \underbrace{\frac{\eta-\epsilon}{\eta} d\log Y}_{\text{Sector competition effect}}, \tag{49}$$

where the *residual demand effect* reflects that expanding firms move down their residual demand curves, and the *sector competition effect* reflects that competing firms expand by stealing sectoral demand.

The change in output of non-encouraged firms, $d\log Y_0$, reflects both productivity and

employment responses:

$$d\log Y_0 = \underbrace{\mathbb{E}[A_{i0}B_{i0}]}_{\text{Adopter productivity}} + \underbrace{\mathbb{E}\left[d\log N_0\right]}_{\text{Employment effect}}, \tag{50}$$

where the *adopter productivity* and *productivity* effects are identified by:

$$\mathbb{E}[A_{i0}B_{i0}] = A_0 \times \Delta^{Y}_{\text{NE,A}-\text{NE,NA}} \tag{51}$$

$$\mathbb{E}\left[d\log N_0\right] = -E \times \Delta^{N}_{\text{E}-\text{NE}}, \tag{52}$$

where A_0 is the adoption rate in non-encouraged firms and E is the share of firms that encourage AI chatbots. Equation (52) utilizes the labor market clearing condition $E \times d\log N_1 + (1-E)d\log N_0 = 0$.

Similarly, total output growth, $d\log Y$, reflects aggregate adopter productivity:

$$d\log Y = (1-E) \times A_0 \times \Delta^{Y}_{\text{NE,A}-\text{NE,NA}} + E \times A_1 \times \Delta^{Y}_{\text{E,A}-\text{E,NA}}, \tag{53}$$

noting that total employment effects cancel in aggregate (due to labor market clearing).

Putting these together, the skill-price effect in non-encouraged firms is:

$$\mathbb{E}[d\log W_0] = -\frac{1}{\eta}\left[A_0 \cdot \Delta^{Y}_{\text{NE,A}-\text{NE,NA}} - E \times \Delta^{N}_{\text{E}-\text{NE}}\right] \tag{54}$$

$$+ \frac{\eta-\epsilon}{\eta}\left[(1-E)A_0 \cdot \Delta^{Y}_{\text{NE,A}-\text{NE,NA}} + EA_1 \cdot \Delta^{Y}_{\text{E,A}-\text{E,NA}}\right], \tag{55}$$

where all terms are identified by our existing difference-in-differences estimates, up to calibrated values of the demand elasticities η and ϵ.

Importantly, when both the firm-level employment difference-in-differences Δ^N and the worker-level earnings difference-in-differences Δ^Y are approximately zero (as in Figures 9 and 13), this implies that the missing intercept, $\mathbb{E}[d\log W_0]$, is also close to zero.

A.2.7 Heterogeneous Encouragement Costs

For analytical simplicity, our main analysis focused on the case where firms face the same encouragement cost, $\kappa_j = \kappa$ for all j. In this section, we consider the general case with heterogeneous encouragement costs, letting F_κ denote the CDF of employer encouragement costs.

When encouragement costs are heterogeneous, equilibrium skill price differentials reflect both the supply of and demand for human capital across encouraged and non-encouraged employers. These, in turn, depend on workers' sorting and adoption decisions, as well as firms' encouragement and hiring policies. We begin by characterizing these decision rules in terms of the relative skill price, $W_{1:0} = W_1/W_0$, and then combine supply and demand to solve for the resulting equilibrium skill differentials.

Employer encouragement. Following the cutoff rule in Equation (20), the share of employers encouraging AI chatbots satisfies:

$$\mathbb{P}[E_j = 1] = F_\kappa \left[(1-\epsilon) \log W_{1:0}\right], \tag{56}$$

where F_κ denotes the CDF of encouragement costs.

Workplace labor demand. Encouraged firms gain a marginal productivity advantage through lower skill prices, increasing their demand for human capital by:

$$\log H_{1:0} = -\epsilon \log W_{1:0}. \tag{57}$$

Demand for human capital. Combining Equations (56)-(57), the aggregate demand for human capital by encouraged firms relative to non-encouraged firms satisfies

$$H^D_{1:0}(\underbrace{W_{1:0}}_{-} \mid \underbrace{F_\kappa}_{+}, \underbrace{\epsilon}_{+}) = \frac{\mathbb{P}[E_j = 1] H_{1:0}}{1 - \mathbb{P}[E_j = 1]} \tag{58}$$

$$= \frac{F_\kappa \left[(1-\epsilon) \log W_{1:0}\right] W_{1:0}^{-\epsilon}}{1 - F_\kappa \left[(1-\epsilon) \log W_{1:0}\right]}, \tag{59}$$

which is strictly decreasing in the relative skill price $W_{1:0}$, strictly increasing in the elasticity of product demand ϵ (since $W_{1:0} < 1$ in equilibrium), and increasing in F_κ. Intuitively, the relative demand for human capital at encouraged firms is higher when encouragement costs are low (i.e., $F_\kappa(x)$ is high for any x), the relative skill price $W_{1:0}$ is low (enhancing the productivity advantage of encouraged firms), and product demand is highly elastic (allowing firms to scale more easily, recoup encouragement costs, and hire more workers).

Adoption. The share of workers who sort into encouraged firms and adopt chatbots satisfies (following Equation (23)):

$$\mathbb{P}[A_i = 1] = i^* = -\frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_{1:0}). \tag{60}$$

Supply of human capital. The aggregate supply of human capital to encouraged firms satisfies:

$$H^S_1(\underbrace{W_{1:0}}_{+} \mid \underbrace{\alpha_b}_{+}, \underbrace{\alpha_c}_{-}, \underbrace{\beta_b}_{+}, \underbrace{\beta_c}_{-}) = \int_0^{i^*} \exp(\alpha_{b0} + \alpha_{b1} + \beta_b i) di \tag{61}$$

$$= \frac{\exp(\alpha_{b0} + \alpha_{b1})}{\beta_b} \left[\exp(\beta_b i)\right]_0^{i^*} \tag{62}$$

$$= \frac{\exp(\alpha_{b0} + \alpha_{b1})}{\beta_b} \left[\exp(-\frac{\beta_b}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_{1:0})) - 1\right], \tag{63}$$

which is strictly increasing in the relative skill price ($W_{1:0}$), strictly increasing in both the common and idiosyncratic productivity benefits of AI chatbots (α_b and β_b), and strictly decreasing in both the common and idiosyncratic adoption costs (α_c and β_c)—under the

assumption that chatbots are never counterproductive when encouraged (i.e., $B_{ij}(i = 1, E_j = 1) = \alpha_{b0} + \alpha_{b1} + \beta_b > 0$).

The aggregate supply of human capital to non-encouraged firms satisfies:

$$H_0^S(\underbrace{W_{1:0}}_{-} \,|\, \underbrace{\alpha_b}_{-}, \underbrace{\beta_b}_{-}, \underbrace{\alpha_c}_{+}, \underbrace{\beta_c}_{+}) = \int_{i^*}^{1} 1di = 1 + \frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_{1:0}), \tag{64}$$

which is strictly decreasing in the relative skill price $W_{1:0}$, strictly increasing in both the common and idiosyncratic productivity benefits of AI chatbots (α_b and β_b), and strictly decreasing in both the common and idiosyncratic adoption costs (α_c and β_c).

Combining Equations (60)-(63), the aggregate supply of human capital to encouraged firms relative to non-encouraged firms satisfies

$$H_{1:0}^S(\underbrace{W_{1:0}}_{+} \,|\, \underbrace{\alpha_b}_{+}, \underbrace{\beta_b}_{+}, \underbrace{\alpha_c}_{-}, \underbrace{\beta_c}_{-}) = \frac{\frac{\exp(\alpha_{b0}+\alpha_{b1})}{\beta_b}\left[\exp(-\frac{\beta_b}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_{1:0})) - 1\right]}{1 + \frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_{1:0})}, \tag{65}$$

Intuitively, the relative supply of human capital to encouraged firms is higher when the relative skill price is high (inducing more workers to sort into these firms), the productivity benefits of chatbot adoption (α_b and β_b) are large, and adoption costs (α_c and β_c) are low—provided that chatbots are not counterproductive for encouraged adopters.

Skill price differentials. Equilibrium skill price differentials equate the relative supply and demand for human capital at encouraged firms:

$$H_{1:0}^S(W_{1:0}) = H_{1:0}^D(W_{1:0}) \tag{66}$$

Since the supply curve $H_{1:0}^S$ is strictly increasing in $W_{1:0}$ and the demand curve $H_{1:0}^D$ is strictly decreasing, Equation (66) pins down a unique equilibrium skill price differential.

Panel (b) of Figure A.2 illustrates the inverse aggregate supply and demand curves for relative human capital at encouraged firms. The equilibrium skill price differential corresponds to their intersection.

Equilibrium skill price differentials are high when relative demand for human capital, $H^D_{1:0}$, exceeds relative supply, $H^S_{1:0}$. From Equations (59) and (65), this occurs when encouragement costs are low (F_κ is high), product demand is elastic (ϵ is high), the productivity benefits of chatbot adoption (α_b, β_b) are small, and adoption costs (α_c, β_c) are high:

$$W_{1:0}(\underbrace{F_\kappa}_{+}, \underbrace{\epsilon}_{+}, \underbrace{\alpha_b}_{-}, \underbrace{\beta_b}_{-}, \underbrace{\alpha_c}_{+}, \underbrace{\beta_c}_{+}). \tag{67}$$

We may use the policy rules for encouragement, adoption, and sorting to further bound the equilibrium skill price differentials.

First, in any equilibrium where some but not all firms encourage chatbot use, Equation (20) implies that the relative skill price must lie within:

$$\log W_{1:0} \in [\kappa_{\max}/(1-\epsilon), \kappa_{\min}/(1-\epsilon)]\,. \tag{68}$$

Similarly, in any equilibrium where only some workers sort into encouraged firms, Equation (23) implies that:

$$\log W_{1:0} \in [-\hat{\alpha}_1, 0]\,. \tag{69}$$

That is, when $\log W_{1:0} > 0$, even non-adopters prefer encouraged firms; when $\log W_{1:0} < -\hat{\alpha}_1$, even adopters prefer non-encouraged firms.

B Data and Institutional Setting

B.1 Wage Setting Systems in Denmark

The Danish labor market is characterized by both high flexibility and strong collective institutions. Union density is relatively high at 65%, and the coverage rate of collective bargaining agreements (CBAs) is even higher at 81% (Jäger, Naidu and Schoefer, 2024; Munch and Olney, 2024). While unions negotiate CBAs, coverage is not dependent on union membership. A worker may be a union member without being covered by a CBA, and conversely, a non-union member may still be covered. Whether a worker is covered depends on whether their employer is a member of an employer organization.

Beginning in the 1980s and continuing through the 1990s, the Danish wage-setting system underwent significant decentralization (Dahl, Le Maire and Munch, 2013). Today, many workers negotiate at least part of their wages individually, even when they are formally covered by a CBA. When a CBA is negotiated, the resulting wage-setting structure typically falls into one of the following three categories:

- **Sector-Level Bargaining (Centralized):** Wages are negotiated centrally, with wages largely determined by the combination of industry and occupation. These agreements may include provisions for returns to tenure or education.
- **Two-Tiered Bargaining:** The union and employer organization negotiate either a minimum hourly wage or a minimum income level for each industry-occupation category. Actual wages are then negotiated at the firm level, potentially with union support.
- **Firm-Level Bargaining (Decentralized):** No specific wages are set at the collective level. Instead, all wage negotiations occur directly between the firm and its employees. The precise form of firm-level bargaining can vary.

Table B.1 provides an overview of the wage-setting systems in our 11 study occupations. Decentralized (two-tiered or firm-level) bargaining systems are the most prevalent in our sample of 11 occupations. The main exception is teachers, who primarily work in the public sector and are typically subject to centralized wage-setting agreements.

However, not all workers in our sample are necessarily covered by CBAs. According to Danish union representatives, CBA coverage is especially limited in small private firms.[42] To provide an indication of expected coverage, Columns (1)-(3) of Table B.1 show the share of our sample employed in either large private firms (defined as firms with at least 25 employees) or in the public sector, as proxies for likely CBA coverage. Columns (4)-(6) report the type of wage-setting system in place among workers who are covered by a CBA, utilizing the data collection from Dahl, Le Maire and Munch (2013).

Table B.1: Wage Setting Systems in Denmark

	Proxies for CBA coverage			Wage setting systems under CBAs		
	Public sector	Large private firms	Total (1+2)	Sector-Level	Two-Tiered	Firm-Level
	(1)	(2)	(3)	(4)	(5)	(6)
Acccountants	.055	.729	.784	.115	.035	.839
Customer Service	.090	.826	.916	.228	.447	.304
Financial Advisors	.019	.942	.961	.906	.011	.080
HR Prof.	.395	.548	.943	.173	.039	.781
IT Support	.211	.681	.892	.170	.125	.697
Journalists	.416	.544	.960	0.000	.540	.457
Legal Prof.	.561	.379	.940	.043	.002	.954
Marketing	.023	.851	.874	.156	.123	.711
Office Clerks	.493	.385	.878	.355	.045	.583
Software Dev.	.005	.766	.771	0.000	0.000	.999
Teacher	.787	.148	.935	.956	0.000	.043
All	.278	.618	.896	.282	.124	.586

Notes: This table summarizes the prevalence of wage-setting systems across our study occupations. Columns (1)–(2) report proxies for collective bargaining agreement (CBA) coverage: the share of workers in the public sector and in large (25+ employee) private firms. Column (3) sums these as a proxy for overall CBA coverage. Columns (4)–(6) show the share of CBA-covered workers under centralized, two-tiered, and decentralized bargaining systems, based on industry-by-occupation classifications from Dahl, Le Maire and Munch (2013).

[42]We thank David Rosenqvist from Dansk Metal, the largest union for workers in the manufacturing sector, for sharing his insights with us.

B.2 Sampling Protocol

Our 2024 survey invited 115,000 workers across 11 occupations. Ideally, we would sample an equal number from each—i.e., 10,450 journalists, 10,450 software developers, etc. However, some occupations in Denmark employ fewer than 10,450 workers. To address this, we follow these steps:

1. If an occupation has fewer than 10,450 workers, we sample all available workers.
2. The remaining invitations are redistributed equally among the other occupations.
3. Workplaces are randomly selected for sampling, and if chosen, all relevant workers (i.e., those in the target occupation) within the workplace are included.
4. Large workplaces can distort the sample balance. To mitigate this, we apply individual-level sampling to the top 2.5% of workplaces (ranked by the number of employees in the relevant occupation), randomly selecting employees using the same sampling probability as in Step 3.
5. To precisely reach our target of 115,000 workers, we make final adjustments by randomly including or excluding workers, independent of their workplace.

The 2023 survey followed a similar protocol, with minor modifications; see Humlum and Vestergaard (2025) for details.

B.3 Survey Sample

Table B.2 outlines how successive sample restrictions define our analysis sample. In total, we obtained about 25,000 complete and valid responses per survey round that can be linked to registry data. The attrition and response rates in our survey are comparable to those obtained in previous Danish surveys (Hvidberg, Kreiner and Stantcheva, 2023). While our main analysis focuses on responses from the 2024 round, we use the 2023 round to examine the dynamics of our estimated effects.

Table B.2: Sample Construction

	2024 Survey		*2023 Survey*			
			Main Survey		*Follow-Up Only*	
	Individuals	Percent of invitees	Individuals	Percent of invitees	Individuals	Percent of invitees
1. Invitees	115,000	100.0	100,000	100.0	15,000	100.0
2. Respondents	30,411	26.4	29,067	29.1	4,094	27.3
3. In target occupation(s)	26,925	23.4	25,121	25.1	3,504	23.4
4. Complete responses	25,241	21.9	18,109	18.1	2,561	17.1
5. Linked to registers	24,796	21.6	17,907	17.9	2,559	17.1

Notes: This table outlines how successive sample restrictions define our analysis sample. We conducted two survey rounds in November 2023 and 2024, each inviting 115,000 workers to participate. The 2023 survey included both a main survey and a two-week follow-up, with 15,000 workers invited only to the follow-up. Row 2 reports the number of individuals who responded to the survey. Row 3 shows the respondents who were still employed in one of our 11 target occupations at the time of the surveys. Row 4 presents the respondents who fully completed the survey questionnaire. Row 5 indicates the complete responses that could be linked to registry data.

B.3.1 Representativeness and Response Quality

In this section, we extend the checks of representativeness and response quality provided in Humlum and Vestergaard (2025) to include the 2024 survey round.

Table B.3 show that our survey respondents resemble our survey populations on observable characteristics.

Table B.3: Balance Table for Survey Respondents

	2024 Survey			*2023 Main Survey*		
	Population	Sampled	Responded	Population	Sampled	Responded
	(1)	(2)	(3)	(1)	(2)	(3)
Age	42.93	42.94	46.11	42.41	42.40	45.38
	(11.54)	(11.52)	(11.50)	(11.57)	(11.57)	(11.51)
Female	0.56	0.56	0.56	0.52	0.52	0.49
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
log(Earnings)	12.98	12.98	13.01	13.07	13.07	13.11
	(0.70)	(0.70)	(0.64)	(0.58)	(0.59)	(0.53)
Experience	6.11	6.11	7.24	6.05	6.05	7.12
	(4.80)	(4.80)	(4.92)	(4.58)	(4.57)	(4.67)
Wealth / Earnings	10.92	6.50	6.74	4.09	4.87	4.10
	(2,148.09)	(286.36)	(204.16)	(157.40)	(262.31)	(39.57)
Observations	284,439	115,000	25,241	283,806	100,000	18,109

Notes: This table compares the mean characteristics of workers in our population (Column 1), our sampled survey invitees (Column 2), and survey respondents with complete responses (Column 3) for each survey round. The *Sampled* columns correspond to line 1 of Table B.2. The *Responded* columns correspond to line 4 of Table B.2. *Population* columns (1) show a difference in the female share between the 2023 and 2024 survey rounds that warrants explanation. This difference arises from a slight modification to the sampling protocol in 2023, in which some sampled workplaces had only (a random) 50% of their relevant workers invited to the survey. This altered the weight each of our 11 occupations received in the invite population, leading to the shifts in gender share observed in columns (1). Importantly, and as expected, the gender composition of the survey population *within* each of our 11 occupations remains virtually unchanged between survey rounds, as does the total unweighted worker population (i.e., without reweighting occupations to reflect our sampling protocol). Since all analyses include occupation fixed effects, this change in occupational composition across survey rounds does not affect our results. Moreover, nearly all analyses in this paper rely on the 2024 survey round, which did not involve the sampling protocol modification. *Sample*: The table includes all individuals in our survey population.

Table B.3 shows that complete respondents (who form the basis of our main analysis sample) and partial respondents have similar characteristics and give similar responses to the survey (before partial respondents drop out).

Table B.4: Balance Table for Complete vs. Partial Responses

	2024 Survey		*2023 Main Survey*	
	Completed	Drop Out	Completed	Drop Out
	(1)	(2)	(1)	(2)
Panel A: Characteristics				
Age	46.11	44.46	45.38	45.00
	(11.50)	(11.98)	(11.51)	(11.53)
log(Earnings)	13.01	13.00	13.11	13.10
	(0.64)	(0.71)	(0.53)	(0.53)
Experience	7.24	6.52	7.12	6.88
	(4.92)	(4.82)	(4.67)	(4.63)
Net Wealth/Earnings	6.74	3.77	4.10	3.75
	(204.16)	(10.08)	(39.57)	(16.43)
Female	0.56	0.57	0.49	0.60
	(0.50)	(0.49)	(0.50)	(0.49)
Panel B: Adoption				
Used	0.69	0.75	0.55	0.51
	(0.46)	(0.43)	(0.50)	(0.50)
Used for Work	0.49	0.58	0.40	0.38
	(0.50)	(0.49)	(0.49)	(0.48)
Used for Core Task	0.31	0.15	0.21	0.17
	(0.46)	(0.35)	(0.41)	(0.38)
Observations	25,241	1,773	18,109	7,012

Notes: This table compares the mean characteristics and adoption behaviors of workers who fully completed (Column 1) and partially completed (Column 2) our surveys. The *Completed* columns correspond to line 4 of Table B.2. Standard deviations are shown in parentheses. See the note of Table B.3 for an explanation of the difference in the female population shares between the 2023 and 2024 survey rounds. *Sample*: All individuals with partial survey responses.

Table B.5 shows that workers who are randomly offered a higher participation prize are more likely to take part in our surveys but do not systematically differ in their responses. Dutz et al. (2025) develop an econometric framework that uses this variation to reweight the sample based on workers' latent willingness to participate; see Humlum and Vestergaard (2025) for its application to our survey.

Table B.5: Balance Table for Participation Prize Categories

	2024 Survey					*2023 Main Survey*				
	Levels	Differences to 1000 DKK			p-value	Levels	Differences to 1000 DKK			p-value
	1000 DKK	2500 DKK	5000 DKK	10000 DKK		1000 DKK	2500 DKK	5000 DKK	10000 DKK	
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Panel A: Characteristics										
Age	46.11	-0.25	-0.18	-0.44	0.18	45.38	-0.46	-0.42	-0.49	0.15
		(0.20)	(0.90)	(0.91)			(0.24)	(0.96)	(0.97)	
log(Earnings)	13.01	-0.00	-0.01	-0.01	0.48	13.11	-0.03	-0.00	-0.01	0.04
		(0.01)	(0.11)	(0.11)			(0.01)	(0.06)	(0.06)	
Experience	7.24	-0.09	-0.11	-0.06	0.55	7.12	-0.01	-0.01	-0.05	0.95
		(0.08)	(0.40)	(0.40)			(0.09)	(0.35)	(0.35)	
Net Wealth/Earnings	6.74	-1.96	0.62	0.94	0.50	4.10	-0.05	0.87	0.38	0.55
		(1.71)	(36.33)	(35.44)			(0.27)	(0.42)	(0.42)	
Female	0.56	-0.02	-0.00	-0.02	0.04	0.49	0.00	-0.01	-0.00	0.41
		(0.01)	(0.03)	(0.04)			(0.01)	(0.04)	(0.04)	
Panel B: Adoption										
Used	0.69	0.00	-0.00	-0.01	0.69	0.55	-0.02	-0.01	-0.01	0.40
		(0.01)	(0.03)	(0.03)			(0.01)	(0.03)	(0.03)	
Used for Work	0.49	-0.00	-0.02	-0.02	0.02	0.40	-0.01	-0.00	-0.00	0.61
		(0.01)	(0.03)	(0.03)			(0.01)	(0.04)	(0.04)	
Used for Core Task	0.31	-0.00	-0.01	-0.00	0.75	0.21	-0.01	0.00	-0.00	0.59
		(0.01)	(0.03)	(0.03)			(0.01)	(0.04)	(0.04)	
Response Rate	0.20	0.02	0.02	0.03	0.00	0.16	0.02	0.02	0.04	0.00
		(0.00)	(0.00)	(0.00)			(0.00)	(0.00)	(0.00)	
Observations	5,787	6,351	6,432	6,671		4,026	4,525	4,549	5,009	

Notes: This table shows that individuals assigned to different participation prize categories (1,000 DKK, 2,500 DKK, 5,000 DKK, and 10,000 DKK) have similar characteristics (Panel A) and adoption behaviors (Panel B) but differ in their rates of completed responses (last row). Column (5) reports p-values from a joint test of whether mean outcomes are equal across the four prize categories. The total number of observations corresponds to line 4 of Table B.2. See the note of Table B.3 for an explanation of the difference in the female population shares between the 2023 and 2024 survey rounds. *Sample*: All complete survey responses.

As an external validation of our survey responses, we cross-check variables workers' reported occupations with those recorded in the administrative registers. Table B.6 shows that the survey and registers agree on the occupation of 87% of our respondents.

Table B.6: Correlation Between Occupation in Survey vs. Register, P(Survey|Register)

	In Survey											
	Journalists	Software Developers	Paralegals	Accountants and Auditors	Customer Service Rep.	Marketing Professionals	Financial Advisors	HR Professionals	Office Clerks	Teachers	IT Support	Observations
Panel A: 2024 Survey												
Journalists	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	325.00
Software Developers	0.00	0.86	0.00	0.00	0.01	0.01	0.00	0.00	0.01	0.00	0.08	2,799.00
Paralegals	0.01	0.03	0.81	0.02	0.00	0.00	0.02	0.02	0.07	0.01	0.01	2,106.00
Accountants and Auditors	0.00	0.02	0.01	0.86	0.01	0.01	0.02	0.01	0.05	0.00	0.01	2,793.00
Customer Service Rep.	0.00	0.02	0.01	0.01	0.79	0.03	0.00	0.01	0.09	0.01	0.01	631.00
Marketing Professionals	0.00	0.07	0.01	0.01	0.09	0.69	0.01	0.01	0.07	0.01	0.03	1,781.00
Financial Advisors	0.00	0.00	0.00	0.00	0.01	0.00	0.96	0.00	0.01	0.00	0.00	1,243.00
HR Professionals	0.01	0.03	0.03	0.01	0.00	0.02	0.02	0.73	0.12	0.01	0.01	849.00
Office Clerks	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.97	0.00	0.01	6,488.00
Teachers	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	6,440.00
IT Support	0.00	0.13	0.00	0.00	0.02	0.02	0.00	0.00	0.03	0.00	0.79	1,470.00
Panel B: 2023 Survey												
Journalists	0.97	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.00	555.00
Software Developers	0.00	0.87	0.00	0.00	0.01	0.02	0.00	0.00	0.01	0.00	0.08	3,185.00
Paralegals	0.01	0.03	0.79	0.02	0.01	0.00	0.01	0.02	0.08	0.01	0.01	2,518.00
Accountants and Auditors	0.00	0.02	0.01	0.85	0.01	0.01	0.02	0.02	0.05	0.00	0.01	2,710.00
Customer Service Rep.	0.01	0.03	0.01	0.01	0.79	0.04	0.01	0.01	0.07	0.01	0.01	869.00
Marketing Professionals	0.00	0.05	0.00	0.00	0.09	0.74	0.01	0.01	0.06	0.00	0.03	2,125.00
Financial Advisors	0.00	0.00	0.00	0.00	0.01	0.00	0.95	0.00	0.02	0.00	0.00	1,918.00
HR Professionals	0.01	0.03	0.06	0.01	0.00	0.01	0.02	0.68	0.14	0.01	0.02	1,434.00
Office Clerks	0.00	0.01	0.00	0.00	0.00	0.01	0.00	0.00	0.96	0.00	0.01	3,395.00
Teachers	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	4,135.00
IT Support	0.00	0.15	0.00	0.00	0.02	0.02	0.00	0.01	0.03	0.00	0.76	2,277.00

Notes: This table presents the correlation between occupational codes reported in the survey and those recorded in the administrative data of Statistics Denmark. Each cell represents the probability of reporting the column occupation in the survey, conditional on having the row occupation registered with Statistics Denmark. The average agreement rate (diagonal elements) is 87%. *Sample*: All completed survey responses.

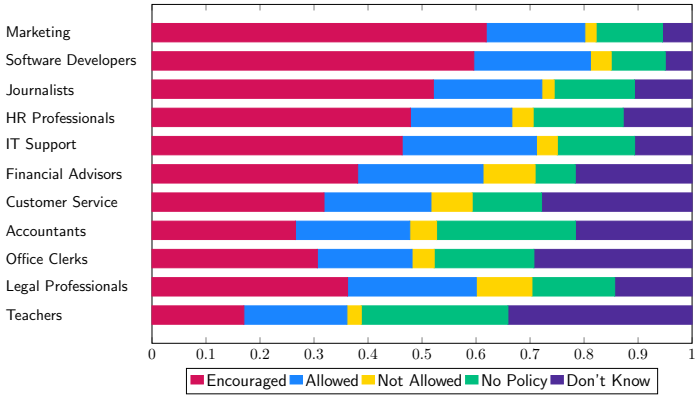
The disagreements in Table B.6 likely reflect measurement error in the registers because firms generally do not update occupational switches of existing employees (Groes, Kircher and Manovskii, 2015). Furthermore, some workers may have switched jobs between June 2024 (our latest month of register data) and November 2024 (the launch of our survey). Table B.6 shows the disagreements occur in cells that reflect likely switches, such as (IT Support, Software Developer). By contrast, the survey and register data agree on the occupation of 100% of our school teachers.

C Adoption

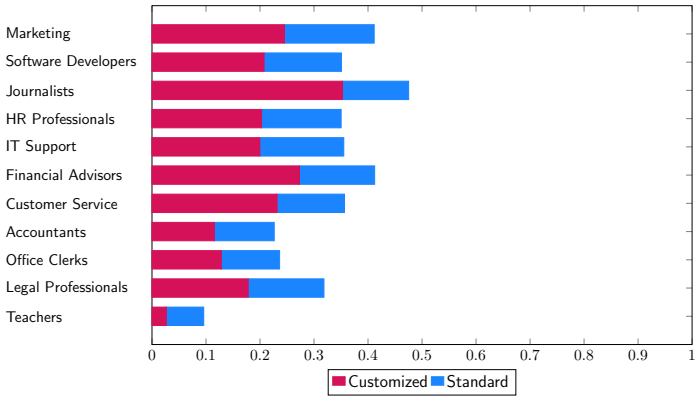
C.1 Employer Initiatives

Figure C.1: The Prevalence of Employer Initiatives for AI Chatbot Adoption (Workplaces)

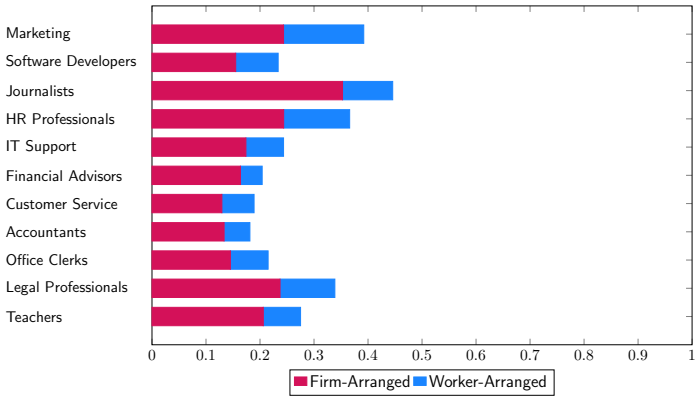
(a) Usage Policies



(b) Have Enterprise AI Chatbot



(c) Have Received Training in AI Chatbots



Notes: This figure shows the share of workplaces affected by various employer initiatives related to AI chatbot adoption. Workers in our sample have been reweighted so all workplaces have the same weight. Panel (a) shows employers' policies on AI chatbot usage for work. Panel (b) indicates whether the employer has its own AI chatbot. Panel (c) reports whether workers have participated in AI chatbot training courses. *Sample*: All completed responses from the 2024 survey.

Table C.1: Which Workplaces Have Adopted AI Chatbot Initiatives?

	Encouraged (1)	Allowed (2)	Not Allowed (3)	No Policy (4)	Don't Know (5)	Firm-Arranged Training (6)	Enterprise Chatbot (7)
Panel A: Within Occupations (Table 1)							
Firm Age (10 Years)	-0.0140**	0.0030	0.0028	0.0038*	0.0044	-0.0122**	-0.0091*
	(0.0044)	(0.0021)	(0.0026)	(0.0016)	(0.0025)	(0.0039)	(0.0043)
log(Firm Employment)	0.0036	-0.0082***	0.0036	-0.0195***	0.0204***	-0.0048	0.0362***
	(0.0048)	(0.0023)	(0.0023)	(0.0022)	(0.0029)	(0.0043)	(0.0056)
log(Firm Labor Productivity)	0.0468**	0.0275***	0.0051	-0.0374***	-0.0420***	0.0765***	0.0697***
	(0.0153)	(0.0077)	(0.0051)	(0.0077)	(0.0084)	(0.0135)	(0.0176)
Private Firm	0.0414*	-0.0076	-0.0141*	-0.0104	-0.0092	0.0257	0.0554**
	(0.0168)	(0.0094)	(0.0072)	(0.0066)	(0.0096)	(0.0144)	(0.0180)
Occupation FE's	✓	✓	✓	✓	✓	✓	✓
Mean of Outcome	0.419	0.209	0.061	0.132	0.178	0.231	0.375
Within ^{2}	0.013	0.003	0.009	0.016	0.016	0.010	0.042
^{2}	0.106	0.005	0.032	0.046	0.071	0.023	0.123
Observations	24184	24184	24184	24184	24184	24184	24184
Panel B: Within Occupations, Worker Controls							
Firm Age (10 Years)	-0.0135**	0.0026	0.0029	0.0039*	0.0041	-0.0119**	-0.0088*
	(0.0045)	(0.0021)	(0.0026)	(0.0016)	(0.0025)	(0.0039)	(0.0043)
log(Firm Employment)	0.0035	-0.0081***	0.0037	-0.0197***	0.0206***	-0.0050	0.0362***
	(0.0047)	(0.0023)	(0.0023)	(0.0022)	(0.0028)	(0.0043)	(0.0055)
log(Firm Labor Productivity)	0.0415**	0.0251**	0.0051	-0.0363***	-0.0355***	0.0716***	0.0640***
	(0.0145)	(0.0077)	(0.0050)	(0.0077)	(0.0077)	(0.0131)	(0.0170)
Private Firm	0.0396*	-0.0094	-0.0140*	-0.0098	-0.0064	0.0238	0.0533**
	(0.0167)	(0.0092)	(0.0071)	(0.0067)	(0.0094)	(0.0144)	(0.0178)
Worker Controls	✓	✓	✓	✓	✓	✓	✓
Occupation FE's	✓	✓	✓	✓	✓	✓	✓
Mean of Outcome	0.419	0.209	0.061	0.132	0.178	0.231	0.375
Within R^2	0.024	0.006	0.010	0.018	0.039	0.017	0.053
R^2	0.116	0.008	0.033	0.047	0.093	0.029	0.133
Observations	24184	24184	24184	24184	24184	24184	24184

Notes: This table examines which firm characteristics predict the adoption of various employer initiatives to promote AI chatbot use. Columns (1)–(5) report results for employer usage policies, while Columns (6)–(7) report results for enterprise chatbots and firm-provided training. Firm characteristics are measured in 2021, our latest data year available. Labor productivity is measured as value added per full-time equivalent worker. The regressions control for whether the firm reports value added. Panel A controls for occupation fixed effects (corresponding to Table 1), whereas Panel B additionally controls for worker characteristics. Standard errors, reported in parentheses, are clustered at the firm level. *Sample*: The table is based on all completed responses from the 2024 survey that can be linked to the registry data.

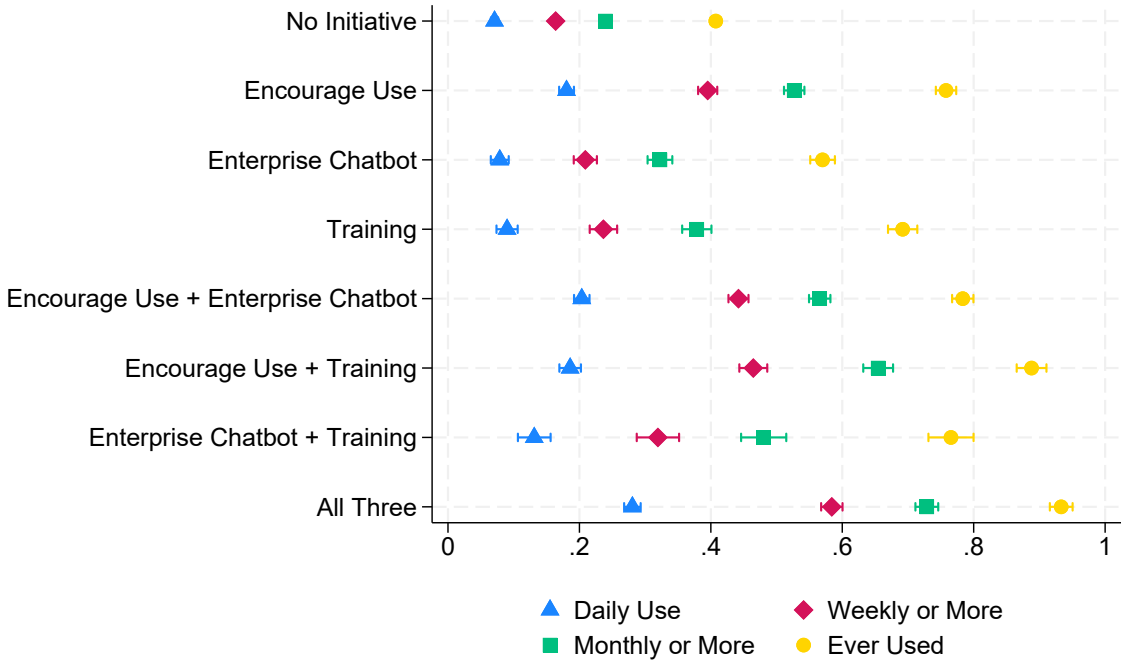
C.2 Worker Adoption

Table C.2: Employer Initiatives and Chatbot Adoption

	(1) Daily Use	(2) Weekly or More	(3) Monthly or More	(4) Ever Used
Encouraged	0.110***	0.231***	0.287***	0.350***
	(0.006)	(0.008)	(0.009)	(0.009)
Enterprise Chatbot	0.008	0.045***	0.083***	0.163***
	(0.008)	(0.010)	(0.010)	(0.010)
Employer Training	0.019*	0.073***	0.139***	0.284***
	(0.009)	(0.011)	(0.012)	(0.012)
Encouraged × Enterprise Chatbot	0.015	0.002	-0.044**	-0.137***
	(0.011)	(0.014)	(0.015)	(0.015)
Encouraged × Employer Training	-0.014	-0.003	-0.011	-0.154***
	(0.013)	(0.017)	(0.018)	(0.018)
Enterprise Chatbot × Employer Training	0.033*	0.038	0.019	-0.089***
	(0.017)	(0.022)	(0.023)	(0.023)
Encouraged × Enterprise Chatbot × Employer Training	0.038	0.035	0.017	0.109***
	(0.021)	(0.028)	(0.030)	(0.029)
No Initiative, Level	0.061	0.155	0.235	0.399
Worker Controls	✓	✓	✓	✓
Occupation FEs	✓	✓	✓	✓

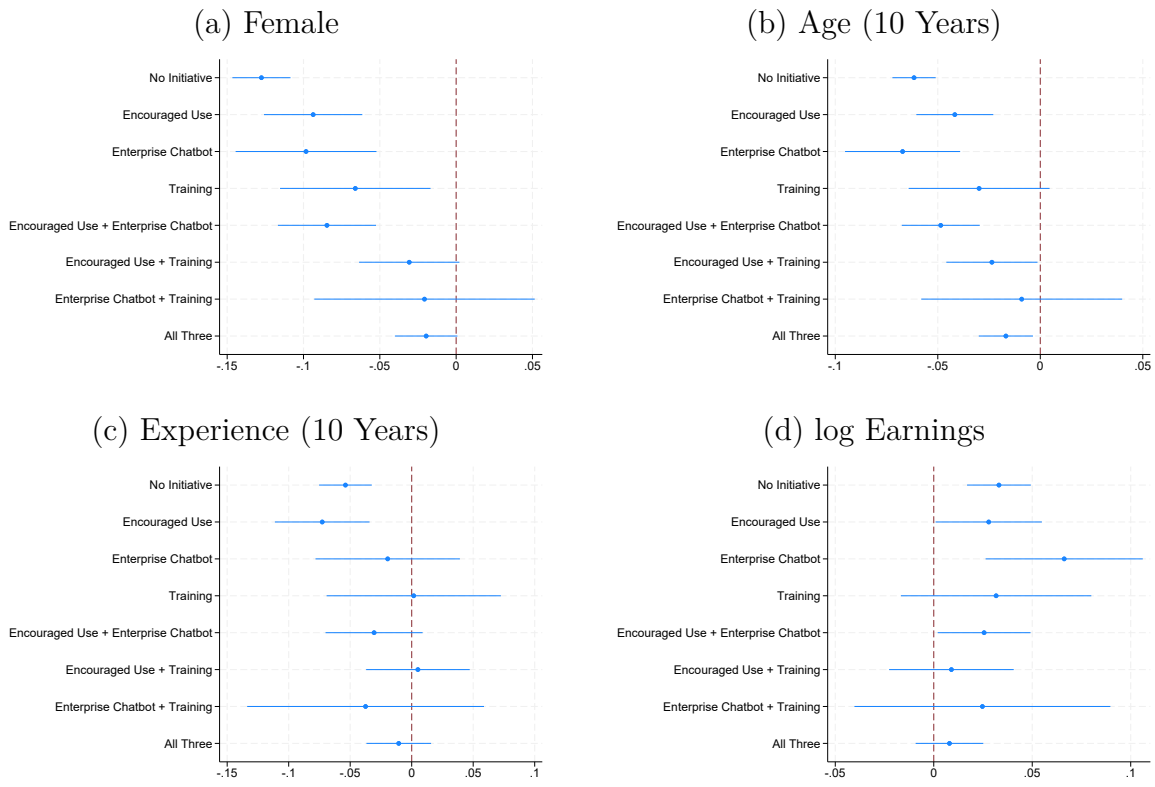
Notes: This table shows how workers' adoption of AI chatbots varies with their employer initiatives. Estimates are based on Equation (1), with standard errors in parentheses. *Sample*: All completed 2024 survey responses linked to registry data.

Figure C.2: Employer Initiatives and Chatbot Adoption



Notes: This figure illustrates how workers' use of AI chatbots vary with their employer initiatives. The estimates are based on predicted values from Equation (1), varying EmployerInitiatives while holding workers' characteristics X at their mean values; see Table C.2 for the full set of regression estimates. Whiskers represent 95% confidence bands of the predicted values. Some confidence intervals are too narrow to be distinguished from the point estimates at this scale. *Sample*: All completed responses from the 2024 survey linked to registry data.

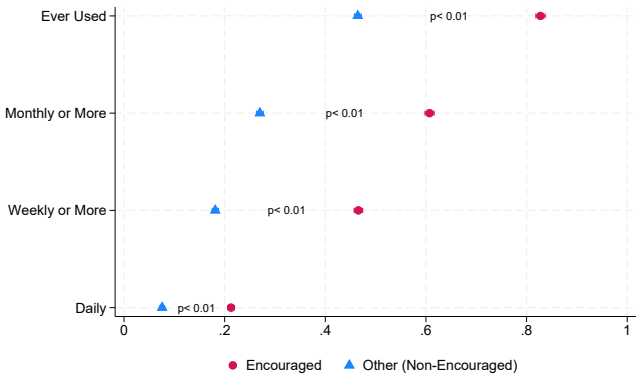
Figure C.3: Employer Initiatives and Worker Gaps in Adoption



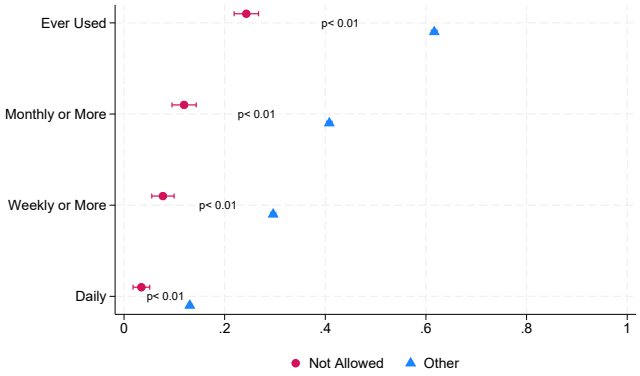
Notes: This figure illustrates the impact of employer usage policies on worker disparities in AI chatbot adoption. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives (Encouraged = 1 vs. Encouraged = 0). Whiskers represent 95% confidence intervals. The reported p-values test whether the coefficients differ between the two groups. *Sample*: All completed responses from the 2024 survey linked to registry data.

Figure C.4: Importance of Employer Policies in AI Chatbot Adoption

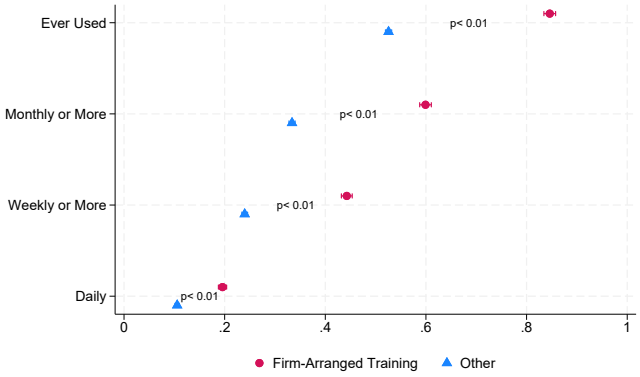
(i) Usage Policy: Encouraged



(ii) Usage Policy: Not Allowed



(iii) Firm-Arranged Training

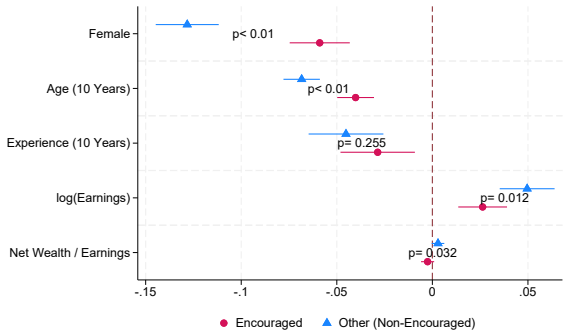


Notes: This figure illustrates the impact of employer policies on workers' use of AI chatbots. The estimates are based on predicted values from Equation (1), varying employer initiatives (EmployerInitiative = 1 vs. EmployerInitiative = 0) while holding workers' characteristics X at their mean values. Panel (i) splits workers based on whether their employer encourages AI chatbot use, Panel (ii) splits workers based on whether their employer allows AI chatbot use, and Panel (iii) splits workers based on whether they have participated in firm-arranged AI chatbot training. Panel (i) is identical to Figure 1. Whiskers represent 95% confidence intervals. The reported p-values test whether the coefficients differ between the two groups. *Sample*: All completed responses from the 2024 survey linked to registry data.

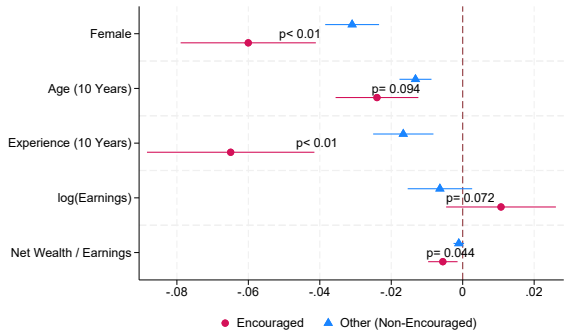
Figure C.5: Influence of Employer Initiatives on Worker Gaps in AI Chatbot Adoption

(i) Usage Policy: Encouraged

(a) Ever Used

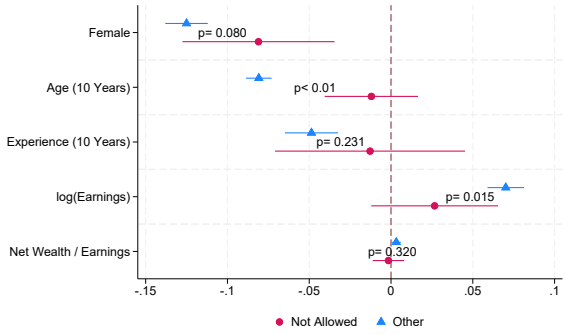


(b) Daily Use

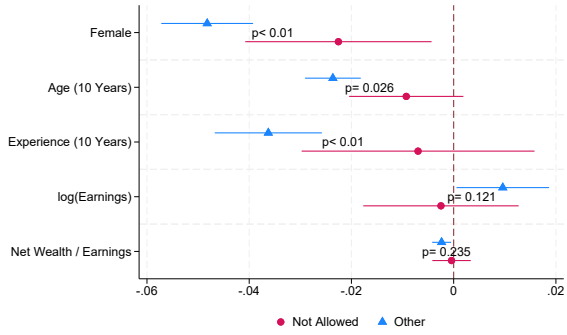


(ii) Usage Policy: Not Allowed

(a) Ever Used

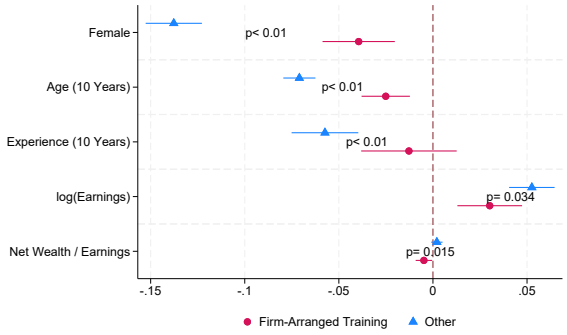


(b) Daily Use

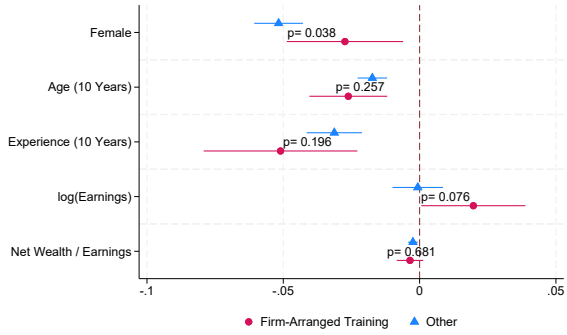


(iii) Firm-Arranged Training

(a) Ever Used



(b) Daily Use



Notes: This figure illustrates the impact of employer initiatives on worker disparities in AI chatbot adoption. The estimates are obtained from regressions of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives. Panel (i) splits workers based on whether their employer encourages AI chatbot use, Panel (ii) splits workers based on whether their employer allows AI chatbot use, and Panel (iii) splits workers based on whether they have participated in firm-arranged AI chatbot training. For each of these, subpanels (a) predicts whether workers have ever used AI chatbots for work, while subpanels (b) predicts whether workers use AI chatbots daily. Panel (i) is identical to Figure C.3. Whiskers represent 95% confidence intervals. The reported p-values test whether the coefficients differ between the two groups. *Sample*: All completed responses from the 2024 survey linked to registry data.

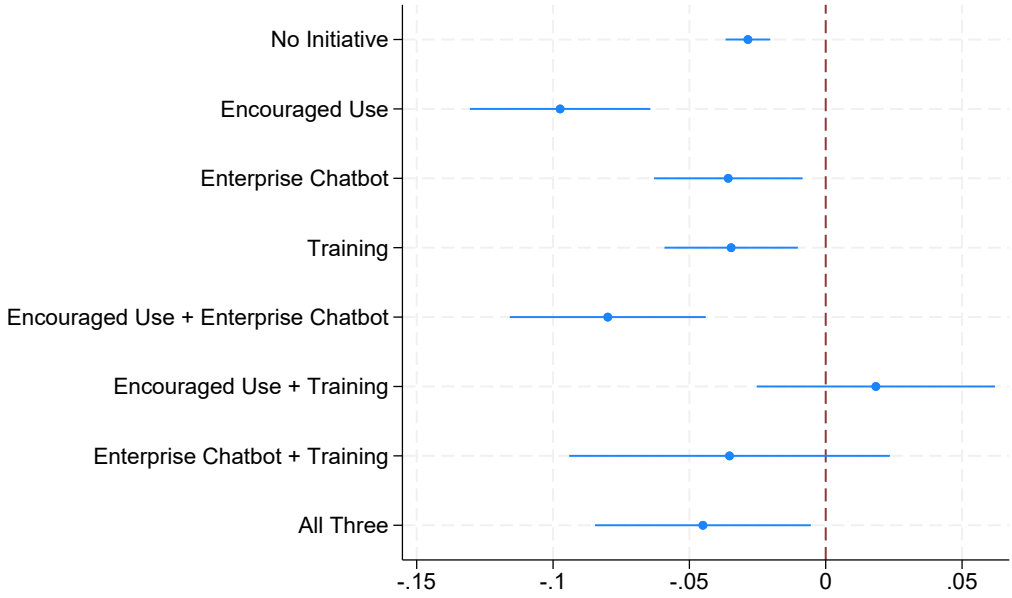
Table C.3: Employer Initiatives and Worker Gaps in Adoption

Panel A. Employer Initiatives in Isolation								
	No Initiative		Encouraged		Enterprise Chatbot		Training	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.128***	-0.123***	-0.094***	-0.069**	-0.098***	-0.092**	-0.066**	-0.040
	(0.009)	(0.012)	(0.016)	(0.023)	(0.024)	(0.034)	(0.025)	(0.035)
Age (10 Years)	-0.062***	-0.053***	-0.042***	-0.034**	-0.067***	-0.084***	-0.030	-0.037
	(0.005)	(0.007)	(0.010)	(0.013)	(0.014)	(0.021)	(0.018)	(0.025)
Experience (10 Years)	-0.054***	-0.035*	-0.073***	-0.070*	-0.020	0.043	0.002	0.074
	(0.011)	(0.014)	(0.020)	(0.028)	(0.030)	(0.044)	(0.036)	(0.053)
log(Earnings)	0.033***	0.020*	0.028*	0.022	0.066**	0.081**	0.032	0.033
	(0.008)	(0.010)	(0.014)	(0.019)	(0.020)	(0.029)	(0.025)	(0.039)
Net Wealth / Earnings	0.002	0.002	-0.009**	-0.006	-0.000	-0.001	-0.006	-0.004
	(0.002)	(0.002)	(0.003)	(0.005)	(0.005)	(0.007)	(0.006)	(0.008)
Occupation FE's	✓	✓	✓	✓	✓	✓	✓	✓
Firm Characteristics Control		✓		✓		✓		✓
Worker Task Mix FEs		✓		✓		✓		✓
Panel B. Employer Initiatives in Combination								
	Encouraged + Enterprise Chatbot		Encouraged + Training		Enterprise Chatbot + Training		All Three	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.085***	-0.085***	-0.031	-0.033	-0.021	0.061	-0.020	0.002
	(0.016)	(0.023)	(0.017)	(0.024)	(0.037)	(0.076)	(0.010)	(0.014)
Age (10 Years)	-0.049***	-0.048***	-0.024*	-0.019	-0.009	-0.009	-0.017*	-0.014
	(0.010)	(0.014)	(0.011)	(0.017)	(0.025)	(0.053)	(0.007)	(0.009)
Experience (10 Years)	-0.031	-0.009	0.005	0.012	-0.037	0.001	-0.011	-0.024
	(0.020)	(0.029)	(0.021)	(0.032)	(0.049)	(0.107)	(0.013)	(0.017)
log(Earnings)	0.026*	0.011	0.009	0.033	0.025	0.038	0.008	0.007
	(0.012)	(0.016)	(0.016)	(0.029)	(0.033)	(0.086)	(0.009)	(0.011)
Net Wealth / Earnings	0.001	-0.002	-0.004	-0.005	-0.016	-0.027	-0.002	0.000
	(0.003)	(0.005)	(0.004)	(0.005)	(0.008)	(0.017)	(0.002)	(0.003)
Occupation FE's	✓	✓	✓	✓	✓	✓	✓	✓
Firm Characteristics Control		✓		✓		✓		✓
Worker Task Mix FEs		✓		✓		✓		✓

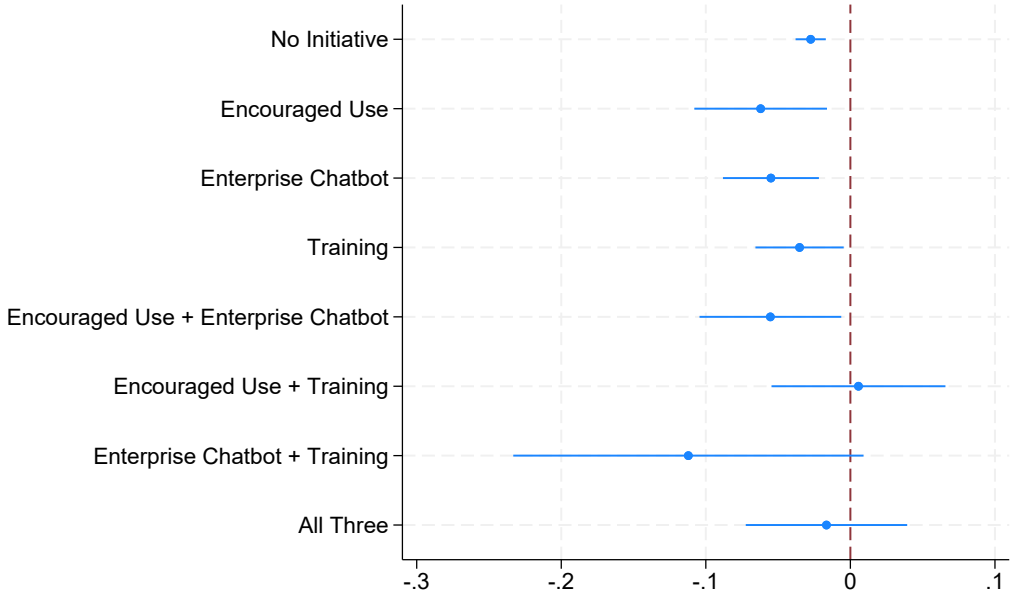
Notes: This table estimates how demographics gaps in AI chatbot adoption varies with employer initiatives. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives. All other characteristics are based on register variables from 2022. *Experience* is the years of employment in the relevant occupation. *Earnings* are total labor income. *Net Wealth* is the sum of real assets, financial assets, and pension savings minus the sum of priority debt, other private debt, and public debt, winsorized at the 5th and 95th percentiles. Standard errors in parentheses. *Sample*: The table is based on all complete responses from the 2024 survey that can be linked to the registry data.

Figure C.6: Employer Initiatives and Gender Gaps in Adoption: Daily Use

(a) Main Specification



(a) Additional Controls



Notes: This figure illustrates how gender gaps in AI chatbot adoption varies with employer initiatives. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives. Worker characteristics include gender, age, experience, earnings, and net wealth; see Table C.3 for the full set of regression estimates. Whiskers represent 95% confidence intervals. *Sample*: All completed responses from the 2024 survey linked to registry data.

C.2.1 AI Chatbot Products

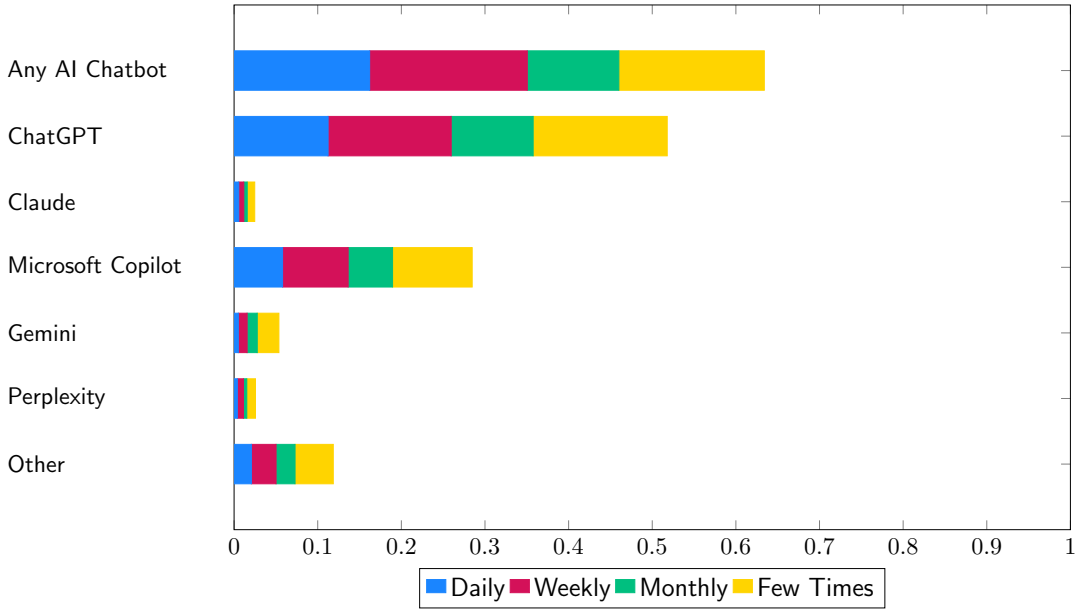
While our main analysis focuses on the adoption of any AI chatbot, our survey also measures usage of specific products. This section provides details on product-level adoption.

The main takeaway is that ChatGPT remains the dominant tool. Figures C.7–C.9 show that approximately 80% of all adopters use ChatGPT, and its dominance holds across all occupations.

Table C.4 further reveals that users of other chatbots often multihome—that is, they also use ChatGPT. For example, 84% of daily Gemini users also use ChatGPT daily. An exception is Microsoft Copilot: only 39% of its daily users also use ChatGPT. Still, this figure is substantially higher than the overall share of daily ChatGPT users in the population (11%), suggesting that workers are generally not exclusive to a single chatbot—those who use one are more likely to use others as well.

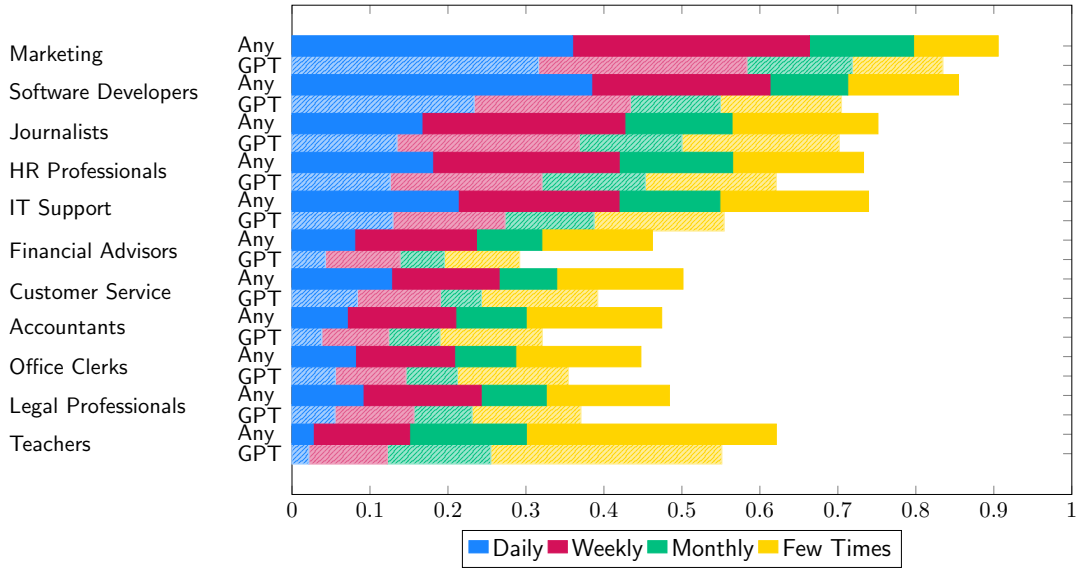
Table C.5 characterizes workers' use of AI chatbot products based on whether their employers provide an enterprise chatbot. The patterns of daily and weekly usage in Panels (a) and (b) are especially informative about ongoing, active use. The table shows that even customized enterprise chatbots are most often versions of ChatGPT. These are typically thin wrappers around ChatGPT Enterprise or the OpenAI API, adapted for workplace use through features such as data security compliance, training on internal data, and custom prompt creation.

Figure C.7: The Prevalence of AI Chatbot Products



Notes: This figure displays the share of workers who have used various AI chatbots for work, categorized by frequency of usage. *Sample:* All completed responses from our 2024 survey round.

Figure C.8: The Dominance of ChatGPT



Notes: This figure shows the share of workers in our study occupations who have used AI chatbots for work, distinguishing between those who have used any AI chatbot and those who have specifically used ChatGPT. *Sample:* All completed responses from our 2024 survey round.

Table C.4: Correlation Between AI Chatbot Product Usage, P(Column|Row)

(a) Ever Used for Work

	Any	ChatGPT	Claude	Copilot	Gemini	Perplexity	Other
Any	1.00	0.82	0.04	0.45	0.09	0.04	0.19
ChatGPT	1.00	1.00	0.05	0.39	0.10	0.05	0.17
Claude	1.00	0.94	1.00	0.66	0.39	0.23	0.39
Copilot	1.00	0.72	0.06	1.00	0.11	0.05	0.16
Gemini	1.00	0.91	0.18	0.56	1.00	0.15	0.32
Perplexity	1.00	0.90	0.22	0.56	0.32	1.00	0.36
Other	1.00	0.73	0.08	0.39	0.15	0.08	1.00
All Workers	0.64	0.53	0.03	0.29	0.06	0.03	0.12

(b) Daily Use for Work

	Any	ChatGPT	Claude	Copilot	Gemini	Perplexity	Other
Any	1.00	0.69	0.03	0.36	0.03	0.02	0.13
ChatGPT	1.00	1.00	0.03	0.20	0.04	0.02	0.07
Claude	1.00	0.63	1.00	0.38	0.09	0.10	0.20
Copilot	1.00	0.39	0.04	1.00	0.04	0.01	0.08
Gemini	1.00	0.84	0.10	0.43	1.00	0.08	0.26
Perplexity	1.00	0.66	0.14	0.22	0.10	1.00	0.20
Other	1.00	0.38	0.05	0.21	0.06	0.04	1.00
All Workers	0.16	0.11	0.01	0.06	0.01	0.00	0.02

Notes: This table presents the correlation between the use of different AI chatbot products. Each cell represents the probability of using the column product, conditional on using the row product. Panel (a) reports probabilities for any work-related use, while Panel (b) focuses on daily work-related use. *All Workers* reports the unconditional usage rate of the chatbot in the given column, weighting all occupations equally. *Sample:* All completed survey responses from our 2024 survey round.

Table C.5: AI Chatbot Product Usage by Availability of Enterprise Chatbot

(a) Daily Use for Work

	Any	ChatGPT	Claude	Copilot	Gemini	Perplexity	Other
Customized Enterprise Chatbot	0.25	0.16	0.01	0.08	0.01	0.01	0.06
Standard Enterprise Chatbot	0.22	0.13	0.01	0.12	0.01	0.00	0.02
No Enterprise Chatbot	0.13	0.10	0.00	0.04	0.00	0.00	0.01
All Workers	0.16	0.11	0.01	0.06	0.01	0.00	0.02

(b) Weekly Use for Work

	Any	ChatGPT	Claude	Copilot	Gemini	Perplexity	Other
Customized Enterprise Chatbot	0.52	0.36	0.02	0.18	0.02	0.02	0.14
Standard Enterprise Chatbot	0.47	0.30	0.01	0.28	0.02	0.01	0.04
No Enterprise Chatbot	0.29	0.24	0.01	0.10	0.01	0.01	0.02
All Workers	0.35	0.26	0.01	0.14	0.02	0.01	0.05

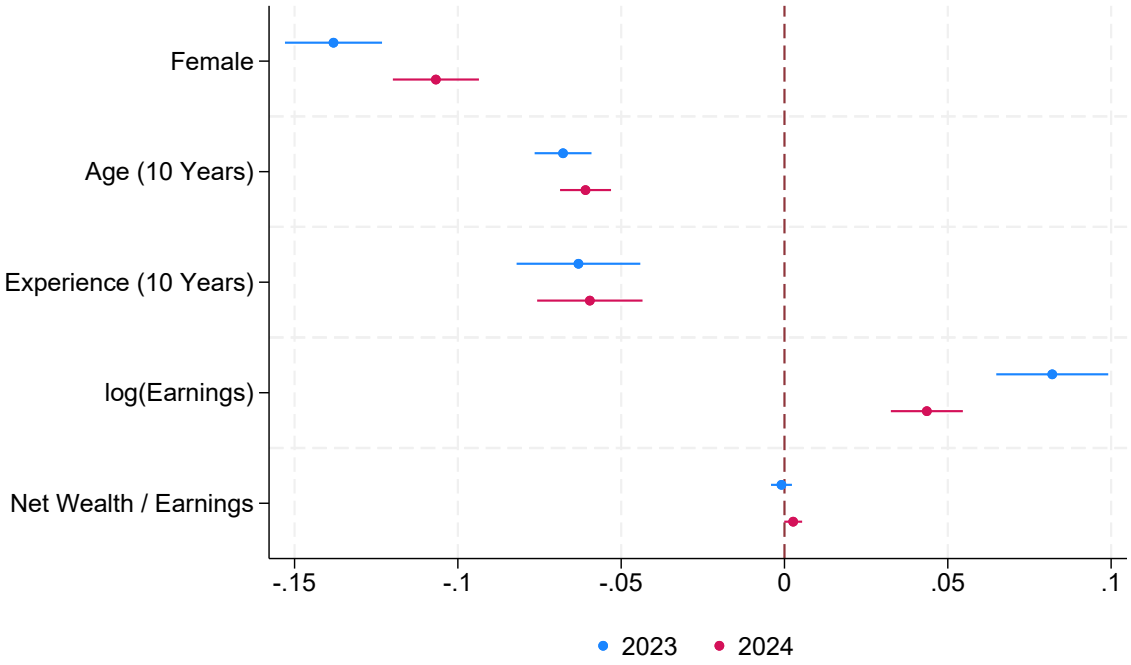
(c) Ever Used for Work

	Any	ChatGPT	Claude	Copilot	Gemini	Perplexity	Other
Customized Enterprise Chatbot	0.80	0.63	0.04	0.34	0.07	0.03	0.29
Standard Enterprise Chatbot	0.80	0.59	0.03	0.54	0.06	0.03	0.11
No Enterprise Chatbot	0.59	0.52	0.02	0.24	0.06	0.03	0.07
All Workers	0.64	0.53	0.03	0.29	0.06	0.03	0.12

Notes: This table reports usage rates of different AI chatbot products among adopters, split by whether their employers offer their own enterprise chatbot. Panel (a) presents daily usage rates, Panel (b) focuses on weekly usage rates, while Panel (c) shows rate of any usage for work. *Sample:* All completed survey responses from our 2024 survey round.

C.2.2 Comparison to 2023 Survey

Figure C.9: Use of ChatGPT for Work (Nov '23, '24)



Notes: This figure presents estimates from a multivariate regression of ChatGPT usage for work on worker characteristics, controlling for occupation fixed effects. The regressions are estimated separately for the 2023 and 2024 survey rounds. *Sample*: All completed responses from the 2023 and 2024 survey rounds.

D Work

D.1 Benefits for Users

Table D.1: Perceived Benefits of AI Chatbots by Occupation

Occupation	Time Savings (1)	Quality (2)	Creativity (3)	Job Satisfaction (4)	No Benefits (5)	Don't Know (6)
Accountants	.709	.524	.418	.156	.08	.059
Customer Service	.723	.572	.413	.142	.083	.041
Financial Adv.	.769	.55	.488	.173	.082	.021
HR Prof.	.844	.638	.623	.253	.028	.03
IT Support	.76	.526	.423	.239	.097	.034
Journalists	.67	.394	.467	.126	.137	.042
Legal Prof.	.783	.544	.45	.204	.083	.027
Marketing	.898	.698	.625	.289	.023	.014
Office Clerks	.689	.563	.489	.185	.073	.067
Software Dev.	.838	.528	.446	.305	.076	.022
Teachers	.637	.321	.458	.132	.15	.071

Notes: This table presents the share of adopters who report various benefits from using AI chatbots for work, broken down by occupation and by whether their employer encourages AI chatbot use. *Sample*: All completed responses from the 2024 survey round.

Table D.2: Employer Initiatives and Reported Time Savings

	(1) Any Time Saving	(2) 15+ min/day	(3) 60+ min/day
Encouraged	0.087***	0.107***	0.034***
	(0.011)	(0.013)	(0.009)
Enterprise Chatbot	-0.000	-0.039*	-0.024
	(0.015)	(0.017)	(0.012)
Employer Training	-0.065***	-0.042*	-0.006
	(0.015)	(0.017)	(0.012)
Encouraged × Enterprise Chatbot	0.009	0.031	0.014
	(0.019)	(0.022)	(0.016)
Encouraged × Employer Training	0.028	-0.014	-0.013
	(0.021)	(0.024)	(0.018)
Enterprise Chatbot × Employer Training	0.040	0.023	-0.011
	(0.028)	(0.032)	(0.024)
Encouraged × Enterprise Chatbot × Employer Training	0.023	0.068	0.043
	(0.035)	(0.039)	(0.029)
No Initiative, Level	0.702	0.478	0.148
Worker Controls	✓	✓	✓
Occupation FEs	✓	✓	✓

Notes: This table shows how workers' time savings from using AI chatbots vary with their employer initiatives. Estimates are based on Equation (1), with standard errors in parentheses. *Sample*: All completed 2024 survey responses linked to registry data.

Table D.3: Employer Initiatives and Allocation of Time Savings

	(1)	(2)	(3)	(4)
	More Leisure	More Breaks	More of Diff. Tasks	More of Same Tasks
Encouraged	-0.014*	-0.005	0.005	0.048***
	(0.007)	(0.006)	(0.010)	(0.012)
Enterprise Chatbot, Standard	-0.009	-0.001	0.021	-0.003
	(0.010)	(0.008)	(0.013)	(0.016)
Enterprise Chatbot, Custom	-0.003	0.003	0.015	0.021
	(0.010)	(0.009)	(0.013)	(0.016)
Employer Training	-0.015	-0.008	-0.004	0.009
	(0.009)	(0.008)	(0.013)	(0.015)
Worker Training	0.004	0.007	0.030**	0.024
	(0.008)	(0.007)	(0.011)	(0.013)
Encouraged × Enterprise Chatbot	0.007	0.002	0.003	-0.011
	(0.012)	(0.010)	(0.016)	(0.020)
Encouraged × Employer Training	0.010	0.001	0.032	-0.032
	(0.013)	(0.011)	(0.018)	(0.022)
Enterprise Chatbot × Employer Training	0.014	-0.001	-0.000	-0.037
	(0.017)	(0.015)	(0.024)	(0.029)
Encouraged × Enterprise Chatbot × Employer Training	-0.013	0.005	-0.023	0.052
	(0.021)	(0.018)	(0.029)	(0.036)
No Initiative, Level	0.091	0.073	0.839	0.249
Worker Controls	✓	✓	✓	✓
Occupation FEs	✓	✓	✓	✓

Notes: This table shows how workers time allocations in response to time savings from AI chatbots vary by their employer chatbot initiatives. Estimates are based on Equation (1), with standard errors in parentheses. *Sample*: All completed 2024 survey responses linked to registry data.

D.2 Workloads and Task Creation

Table D.4: Employer Initiatives and New Workloads From AI Chatbots

Panel A: Adopters			
	(1)	(2)	(3)
	Same Tasks	New Tasks	Same and New Tasks
Encouraged	0.016**	0.036***	0.003
	(0.005)	(0.008)	(0.003)
Enterprise Chatbot	0.002	0.024*	0.003
	(0.007)	(0.011)	(0.005)
Employer Training	0.013	0.075***	0.002
	(0.007)	(0.011)	(0.005)
Encouraged × Enterprise Chatbot	0.002	-0.006	0.007
	(0.009)	(0.014)	(0.006)
Encouraged × Employer Training	-0.021*	-0.034*	-0.001
	(0.010)	(0.015)	(0.007)
Enterprise Chatbot × Employer Training	0.001	-0.051*	-0.001
	(0.014)	(0.021)	(0.009)
Encouraged × Enterprise Chatbot × Employer Training	0.012	0.042	-0.000
	(0.017)	(0.025)	(0.011)
No Initiative, Level	0.039	0.069	0.012
Worker Controls	✓	✓	✓
Occupation FEs	✓	✓	✓
Panel B: Non-Adopters			
	(1)	(2)	(3)
	Same Tasks	New Tasks	Same and New Tasks
Encouraged	-0.002	0.019*	-0.001
	(0.004)	(0.008)	(0.002)
Enterprise Chatbot	0.020***	0.024***	-0.001
	(0.003)	(0.007)	(0.001)
Employer Training	-0.003	0.048***	0.001
	(0.005)	(0.010)	(0.002)
Encouraged × Enterprise Chatbot	-0.016*	-0.007	0.009**
	(0.007)	(0.014)	(0.003)
Encouraged × Employer Training	0.010	-0.026	0.008
	(0.011)	(0.022)	(0.005)
Enterprise Chatbot × Employer Training	0.041***	-0.047*	0.015***
	(0.010)	(0.021)	(0.004)
Encouraged × Enterprise Chatbot × Employer Training	-0.024	-0.018	-0.032***
	(0.017)	(0.035)	(0.007)
No Initiative, Level	0.006	0.027	0.001
Worker Controls	✓	✓	✓
Occupation FEs	✓	✓	✓

Notes: This table shows how new workloads from AI chatbots vary by their employer chatbot initiatives. Estimates are based on Equation (1), with standard errors in parentheses. *Sample*: All completed 2024 survey responses linked to registry data.

Our survey includes free-text responses about the new tasks workers have recived due to AI chatbots. We categorize these responses into six broad AI-related categories, listed in Table D.5, as well as into more granular, occupation-specific subtasks.

Table D.5: Categories of New Tasks from AI Chatbots

Task	**Description**
AI Ideation	Leveraging AI to spark or expand creative ideas—such as concepts, strategies, or solutions. The human selects and builds on the most promising suggestions.
AI Content Drafting	Using AI tools to generate initial drafts of text or media (e.g., documents, emails, code). The human professional prompts the AI, then edits and refines the output for accuracy and tone.
AI Quality Review	Reviewing AI-generated content for accuracy, clarity, and relevance. The human fact-checks, corrects errors, and ensures the output meets required standards.
AI Data Insights	Using AI to analyze data or documents and surface patterns, summaries, or key insights. The human then interprets and applies these findings to decisions.
AI Integration	Embedding AI into workflows to automate or enhance tasks. Professionals design prompts, refine workflows, correct outputs, and fine-tune systems based on feedback.
AI Ethics & Compliance	Ensuring AI use follows ethical, legal, and institutional standards. This includes setting guidelines, monitoring for bias or misuse, and reviewing outputs for compliance.

Notes: This table describes our six broad categories of AI-related tasks.

Examples of occupation-specific tasks, along with their corresponding general task categories, include:

1. **Accountants:** Brainstorming budget plans or tax strategies with AI suggestions (*AI Ideation*), Drafting financial statements and reports using AI for initial content (*AI Content Drafting*), Reviewing AI-generated financial outputs for accuracy and

completeness (*AI Quality Review*), Analyzing financial data with AI tools to identify trends or anomalies (*AI Data Insights*), Ensuring AI-driven accounting processes comply with financial regulations and standards (*AI Ethics & Compliance*)

2. **Customer Support:** Using AI to draft responses to common customer queries or emails (*AI Content Drafting*), Reviewing AI-suggested responses to ensure accuracy and proper tone (*AI Quality Review*), Analyzing customer interactions with AI to identify common pain points and FAQs (*AI Data Insights*), Creating and refining prompts for AI chatbots to handle customer questions (*AI Integration*), Training the AI customer service chatbot by feeding it new Q&As from resolved issues (*AI Integration*), Ensuring the AI chatbot adheres to customer privacy and service guidelines (*AI Ethics & Compliance*)
3. **Financial Advisors:** Brainstorming investment strategies or portfolio ideas using AI insights (*AI Ideation*), Generating draft financial plans and investment recommendations with AI assistance (*AI Content Drafting*), Reviewing AI-suggested investment recommendations for accuracy and client suitability (*AI Quality Review*), Analyzing market trends and client data with AI to inform advice (*AI Data Insights*), Ensuring AI-driven financial advice complies with regulations and ethical standards (*AI Ethics & Compliance*)
4. **HR Professionals:** Brainstorming employee training, development, or wellness program ideas using AI (*AI Ideation*), Drafting job postings, policy documents, or employee communications using AI (*AI Content Drafting*), Reviewing AI-generated candidate evaluations or HR reports for accuracy and bias (*AI Quality Review*), Analyzing employee survey results or HR data with AI to gain insights (*AI Data Insights*), Integrating AI tools for resume screening, interview scheduling, and answering candidate inquiries in recruitment (*AI Integration*), Ensuring AI recruitment and evaluation tools are fair, unbiased, and legally compliant (*AI Ethics & Compliance*)

5. **IT Support Specialists:** Generating technical troubleshooting guides and FAQs using AI (*AI Content Drafting*), Validating AI-proposed solutions to ensure they resolve issues without risk (*AI Quality Review*), Crafting effective queries/prompts for AI tools to diagnose IT issues (*AI Integration*), Integrating AI assistants into support systems to automate routine help requests (*AI Integration*)
6. **Journalists:** Brainstorming story ideas, angles, or interview questions with AI (*AI Ideation*), Using AI to draft article outlines, summaries, or initial news reports (*AI Content Drafting*), Fact-checking and editing AI-generated content to ensure accuracy and clarity (*AI Quality Review*), Summarizing research materials or interview transcripts for quick insight with AI (*AI Data Insights*), Ensuring AI-generated content abides by journalistic ethics and standards (*AI Ethics & Compliance*)
7. **Legal Professionals:** Brainstorming legal arguments, interpretations, or negotiation strategies with AI (*AI Ideation*), Drafting contracts, briefs, or other legal documents with AI providing initial content (*AI Content Drafting*), Reviewing AI-generated legal documents or analyses for accuracy and compliance (*AI Quality Review*), Using AI to research and summarize case law, statutes, or legal documents (*AI Data Insights*), Developing organizational AI usage policies and guidelines (*AI Integration, AI Ethics & Compliance*), Ensuring AI tools and outputs uphold legal ethics and confidentiality (*AI Ethics & Compliance*)
8. **Marketing Professionals:** Brainstorming campaign themes, slogans, or creative concepts with AI (*AI Ideation*), Generating marketing copy, social media posts, or product descriptions with AI (*AI Content Drafting*), Reviewing AI-created marketing content for quality and brand consistency (*AI Quality Review*), Analyzing consumer data and campaign results with AI to derive marketing insights (*AI Data Insights*)
9. **Office Clerks:** Drafting routine emails, letters, or documents using AI assistance (*AI Content Drafting*), Performing quality control on AI-generated text and doc-

uments (*AI Quality Review*), Using AI to extract information from documents or summarize data for reports (*AI Data Insights*), Ensuring no confidential information is inappropriately shared with AI tools (*AI Ethics & Compliance*)

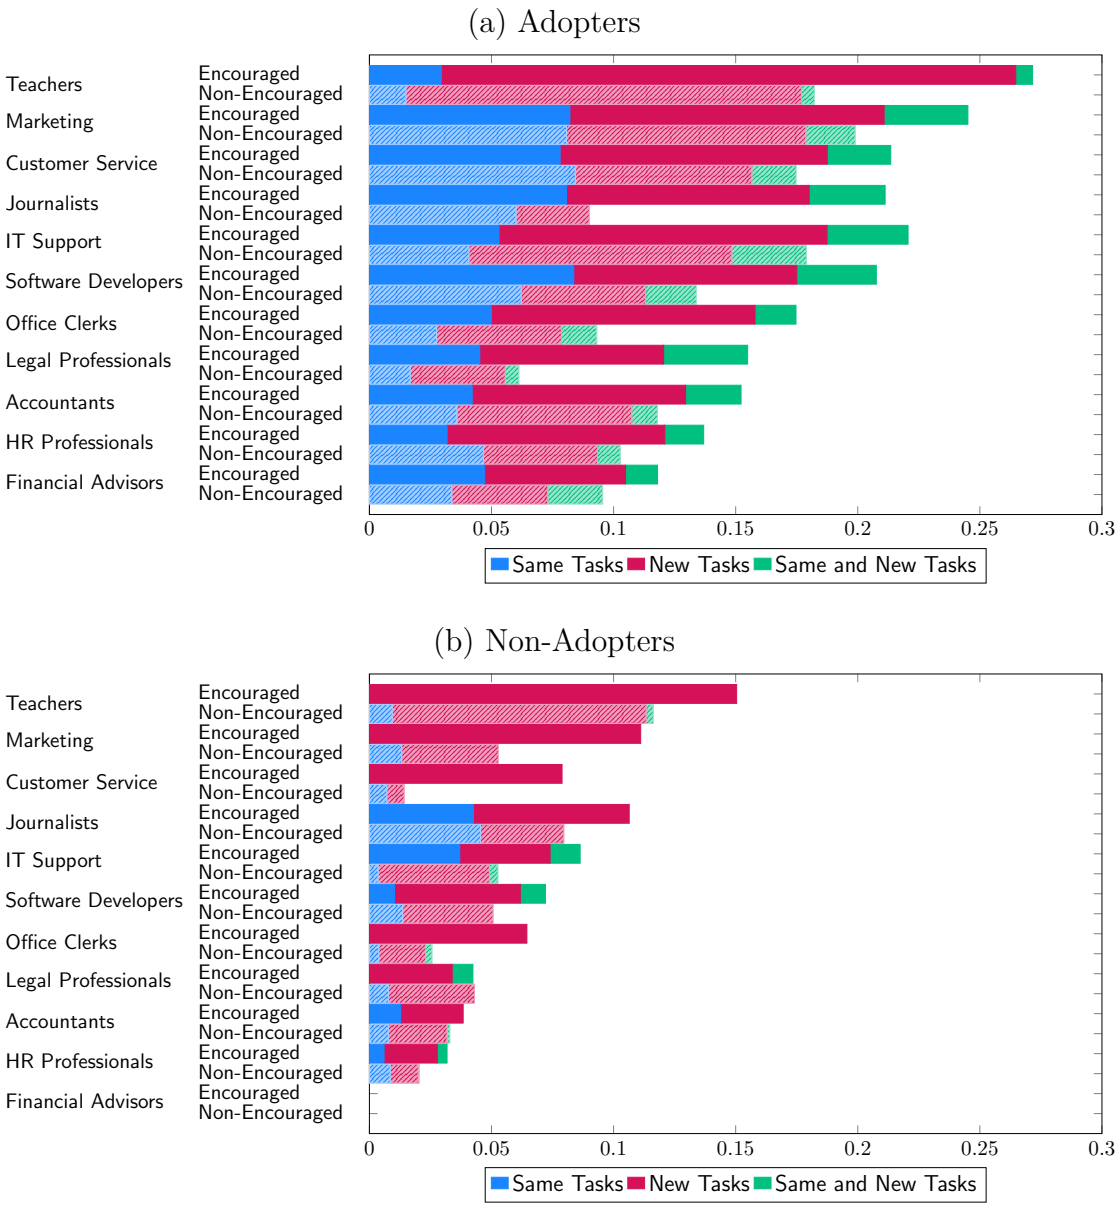
10. **Software Developers:** Using AI to generate code snippets, boilerplate code, or documentation (*AI Content Drafting*), Reviewing and testing AI-generated code to ensure correctness and security (*AI Quality Review*), Formulating specific prompts to guide AI in debugging or coding tasks (*AI Integration*), Fine-tuning the AI coding assistant by providing feedback and project-specific examples (*AI Integration*), Writing prompts for code generation (*AI Integration*)
11. **Teachers:** AI-assisted development of new course material and lesson plans (*AI Ideation*), Personalizing learning materials or feedback using AI insights from student performance (*AI Data Insights*), Adapting exams and assignments to account for AI tool usage (*AI Integration*), Integrating chatbots into lessons (*AI Integration*), Detecting AI-generated homework submissions (*AI Ethics & Compliance*)

For all broad categories, we include an "Other" subtask for each occupation (e.g., *AI Integration, Other*) to ensure that every broad category is represented across all occupations. In addition, we include a "non-AI" task category, which typically captures new assignments for the worker that are not novel within the broader workplace or profession. Examples include "meeting with customers," "taking over tasks due to freed-up time," or the more ambiguous "handling more complex tasks GenAI cannot solve."

To categorize the free-text responses, we divided them between two independent coders. Each coder then cross-checked a random sample of the other's work, with near-complete agreement.

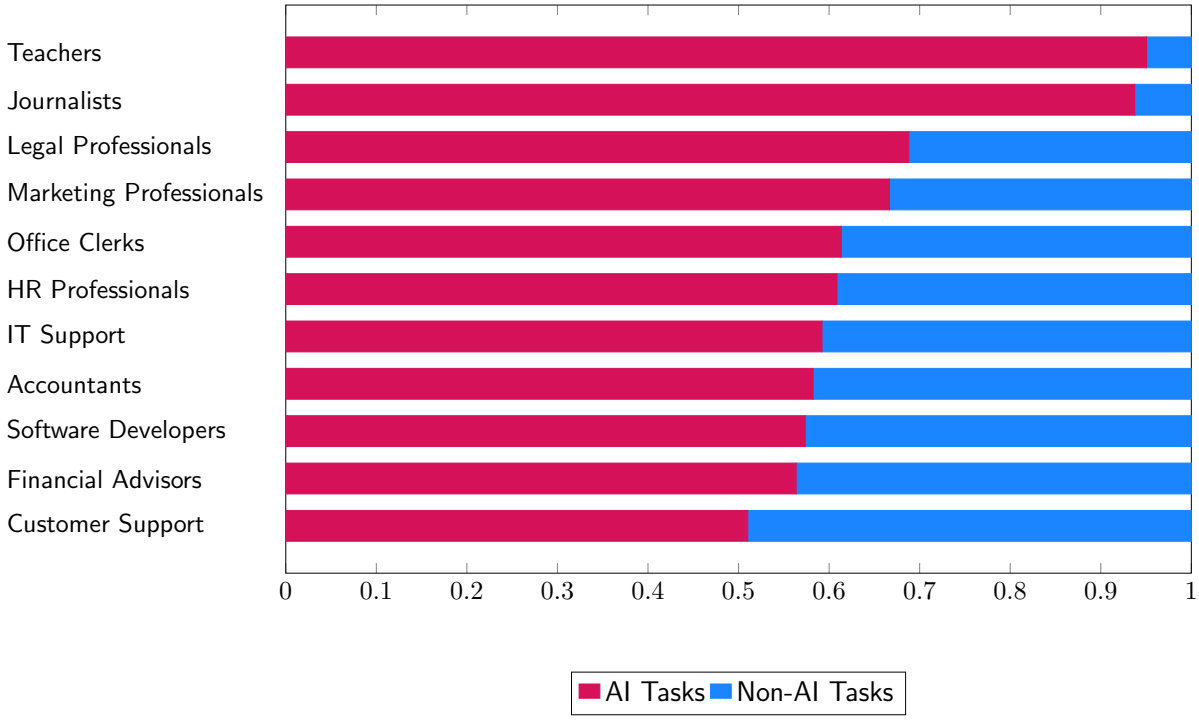
Figures D.1 illustrate that AI chatbots have led to the creation of new tasks across all 11 occupations in our study. Figure D.2 shows that 50% to 95% of these tasks are directly linked to AI use.

Figure D.1: Workloads from AI Chatbots



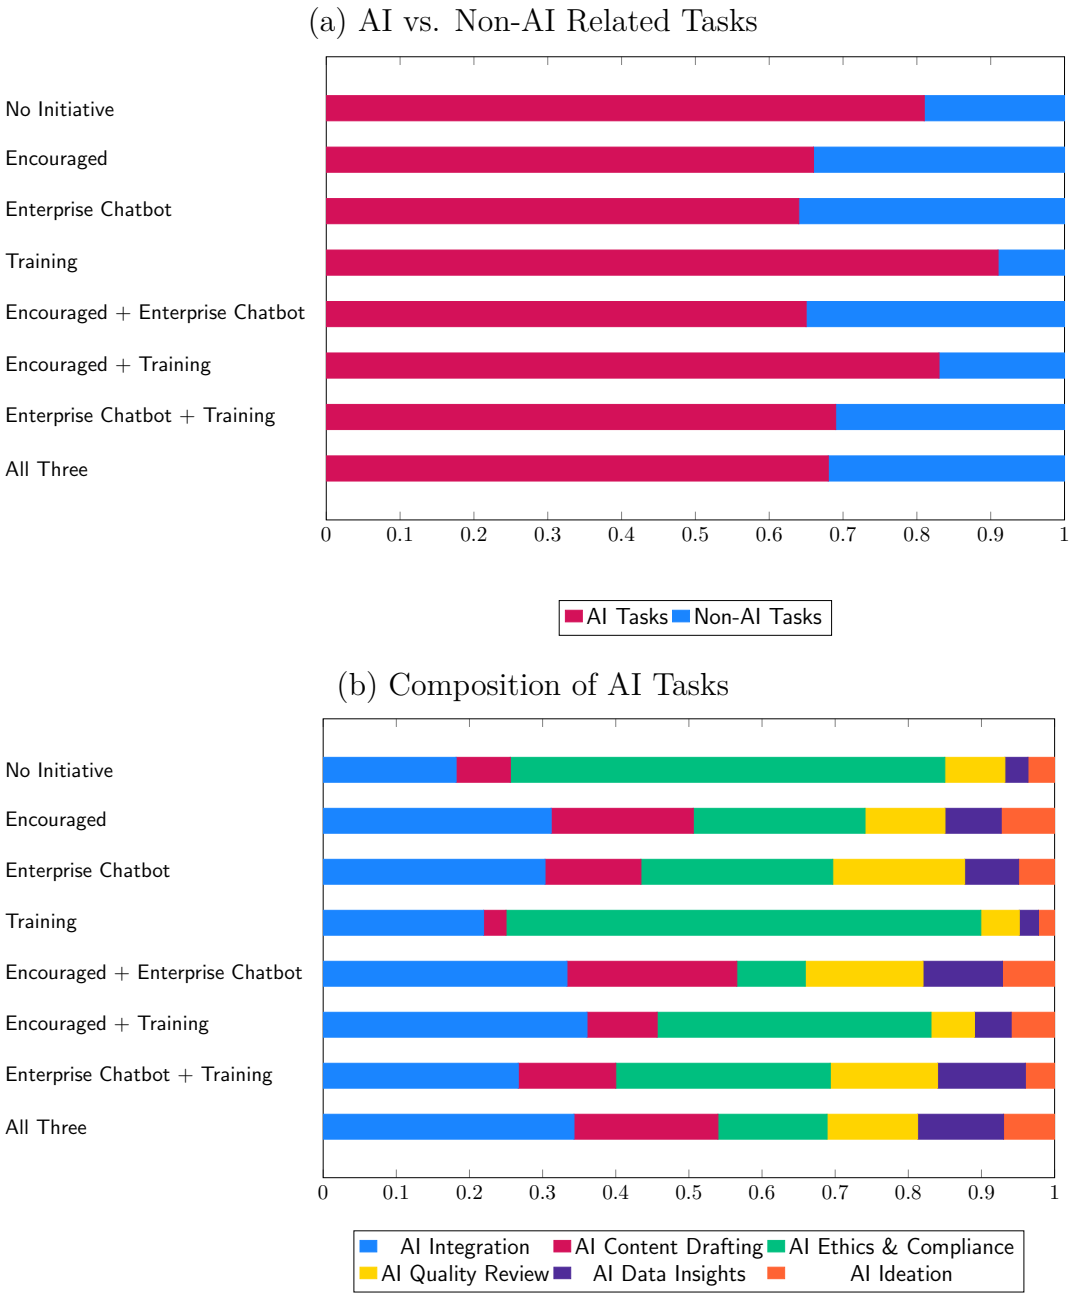
Notes: This figure presents the share of workers who report increased workloads due to AI chatbots, distinguishing between additional tasks of the same type, new job tasks, or both. The responses are broken down by occupation and by whether employers encourage AI chatbot use. Panel (a) focuses on adopters (workers who have ever used AI chatbots for work), while Panel (b) examines non-adopters. *Sample*: All completed responses from the 2024 survey round linked to registry data.

Figure D.2: Composition of New Job Tasks



Notes: This figure shows the share of new job tasks that are directly linked to AI chatbot use. Occupations are ordered according to their shares of AI tasks *Sample*: All completed responses from the 2024 survey who reported new job tasks due to AI chatbots.

Figure D.3: Composition of New Job Tasks by Employer Initiatives



Notes: Panel (a) shows the share of new job tasks that are directly linked to AI chatbot use. Panel (b) shows the distribution of reported new job tasks across major task categories for each occupation. *AI Ideation* refers to "using AI to generate or expand creative ideas, such as concepts, strategies, or solutions," *AI Content Drafting* refers to "using AI tools to produce initial drafts of text or media," *AI Quality Review* refers to "reviewing and correcting AI-generated content for accuracy, clarity, and relevance," *AI Data Insights* refers to "using AI to analyze data or documents and extract key patterns or insights," *AI Integration* refers to "embedding AI into workflows to automate or enhance tasks," and *AI Ethics & Compliance* refers to "ensuring AI use follows ethical, legal, and organizational standards." See Appendix D.2 for details and occupation-specific examples. Tasks are ordered according to their average shares among the eleven occupations. Occupations are ordered according to their shares on *AI Integration*, the most frequent tasks across the occupations. *Sample*: All completed responses from the 2024 survey who reported new job tasks due to AI chatbots.

D.3 Total Time Savings

Table D.6: What Drives Adopters' Time Savings from AI Chatbots?

Occupation		Time Savings / Work Hours (1)	Time Saved Per Day Used / Daily Work Hours (2)	Days Used / Work Days (3)	Covariance (4)
Marketing Professionals	Encouraged	.068	.109	.515	.012
Marketing Professionals	Non-Encouraged	.046	.087	.388	.012
Software Developers	Encouraged	.065	.093	.567	.012
Software Developers	Non-Encouraged	.039	.074	.378	.011
HR Professionals	Encouraged	.042	.081	.405	.009
HR Professionals	Non-Encouraged	.026	.074	.237	.009
IT Support	Encouraged	.041	.075	.407	.011
IT Support	Non-Encouraged	.028	.06	.283	.011
Customer Service Rep.	Encouraged	.039	.066	.439	.010
Customer Service Rep.	Non-Encouraged	.020	.050	.240	.008
Legal Professionals	Encouraged	.033	.074	.336	.008
Legal Professionals	Non-Encouraged	.018	.065	.166	.007
Office Clerks	Encouraged	.029	.065	.309	.009
Office Clerks	Non-Encouraged	.015	.047	.177	.007
Journalists	Encouraged	.028	.050	.374	.010
Journalists	Non-Encouraged	.021	.048	.221	.011
Financial Advisors	Encouraged	.022	.053	.288	.007
Financial Advisors	Non-Encouraged	.015	.040	.179	.008
Accountants and Auditors	Encouraged	.022	.055	.282	.006
Accountants and Auditors	Non-Encouraged	.009	.041	.143	.003
Teachers	Encouraged	.010	.049	.140	.003
Teachers	Non-Encouraged	.006	.044	.088	.002
All	Encouraged	.036	.070	.369	.010
All	Non-Encouraged	.022	.057	.227	.009

Notes: This table presents workers' time savings from AI chatbots, categorized by workers' occupations and employer policies. The table focuses on workers who have ever used AI chatbots for work. Column (1) reports the average time savings as a percentage of total work hours. The remaining columns decompose these time savings into three components: time savings per day of use (Column (2)), the share of workdays with AI chatbot usage (Column (3)), and the covariance between Columns (2) and (3) (Column (4)). These components satisfy the relationship: Column (1) = Column (2) $\times$ Column (3) + Column (4), which follows from the identity: $E[XY] = E[X]E[Y] + \text{Cov}(X, Y)$. Occupations are sorted by their average time savings (Column (1)) among encouraged users. We code daily time savings (Column (2)) as follows: 0–15 minutes/day as 7.5 minutes, 15–60 minutes as 37.5 minutes, and 60+ minutes as 90 minutes. We code frequency of use (Column (3)) as follows: daily use is divided by 1; weekly use by 5 (corresponding to 5 workdays per week); monthly use by $21\frac{2}{3}$ (corresponding to $4\frac{1}{3}$ weeks per month); and use a few times by 65 (corresponding to 4 uses over 12 months). We set daily work hours to 8, as nearly all sampled workers are full-time employees. *Sample*: The table is based on all completed responses from the 2024 survey that can be linked to the registry data.

Table D.7: Decomposition of Adopters' Time Use with AI Chatbots

Occupation		Time Use / Work Hours (1)	Time Used Per Day Used / Daily Work Hours (2)	Days Used / Work Days (3)	Covariance (4)
Software Developers	Encouraged	.053	.074	.567	.011
Software Developers	Non-Encouraged	.033	.06	.378	.01
Marketing Professionals	Encouraged	.049	.076	.515	.01
Marketing Professionals	Non-Encouraged	.036	.062	.388	.011
Customer Service Rep.	Encouraged	.035	.058	.439	.01
Customer Service Rep.	Non-Encouraged	.014	.04	.24	.005
IT Support	Encouraged	.034	.06	.407	.009
IT Support	Non-Encouraged	.023	.051	.283	.008
HR Professionals	Encouraged	.031	.056	.405	.009
HR Professionals	Non-Encouraged	.016	.049	.237	.004
Legal Professionals	Encouraged	.027	.057	.336	.007
Legal Professionals	Non-Encouraged	.01	.044	.166	.003
Journalists	Encouraged	.025	.048	.374	.007
Journalists	Non-Encouraged	.013	.038	.221	.005
Office Clerks	Encouraged	.023	.054	.309	.006
Office Clerks	Non-Encouraged	.012	.044	.177	.004
Accountants and Auditors	Encouraged	.018	.044	.282	.006
Accountants and Auditors	Non-Encouraged	.007	.038	.143	.002
Financial Advisors	Encouraged	.016	.038	.288	.005
Financial Advisors	Non-Encouraged	.01	.031	.179	.005
Teachers	Encouraged	.009	.052	.14	.002
Teachers	Non-Encouraged	.005	.043	.088	.001
All	Encouraged	.029	.056	.369	.008
All	Non-Encouraged	.016	.046	.227	.006

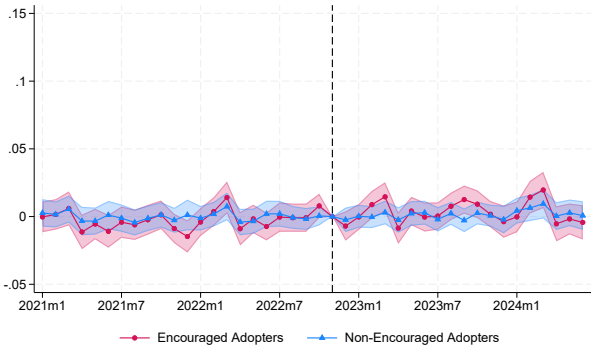
Notes: This table presents workers' time use with AI chatbots, categorized by workers' occupations and employer policies. The table focuses on workers who have ever used AI chatbots for work. Column (1) reports the average time used as a percentage of total work hours. The remaining columns decompose this time use into three components: time use per day of use (Column (2)), the share of workdays with AI chatbot usage (Column (3)), and the covariance between Columns (2) and (3) (Column (4)). These components satisfy the relationship: Column (1) = Column (2) × Column (3) + Column (4), which follows from the identity: $E[XY] = E[X]E[Y] + \mathrm{Cov}(X, Y)$. We code daily time savings (Column (2)) as follows: 0–15 minutes/day as 7.5 minutes, 15–60 minutes as 37.5 minutes, and 60+ minutes as 90 minutes. We code frequency of use (Column (3)) as follows: daily use is divided by 1; weekly use by 5 (corresponding to 5 workdays per week); monthly use by $21\frac{2}{3}$ (corresponding to $4\frac{1}{3}$ weeks per month); and use a few times by 65 (corresponding to 4 uses over 12 months). We set daily work hours to 8, as nearly all sampled workers are full-time employees. *Sample*: The table is based on all completed responses from the 2024 survey that can be linked to the registry data.

E Labor Market Outcomes

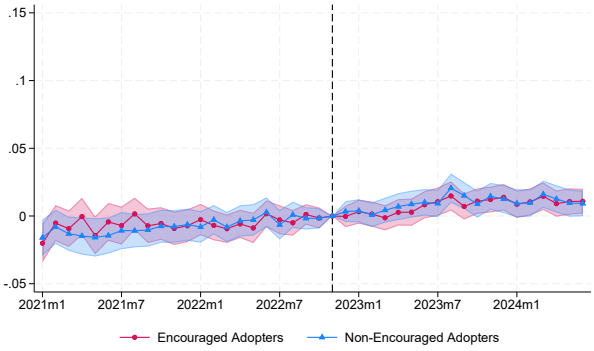
E.1 Results

Figure E.1: Have Adopting Workers Experienced Earnings Gains?
(Log Earnings Relative to Non-Encouraged Non-Adopters)

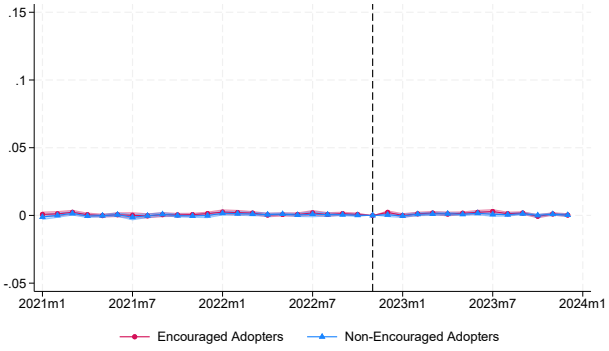
(a) log Hourly Wage



(b) log Hours



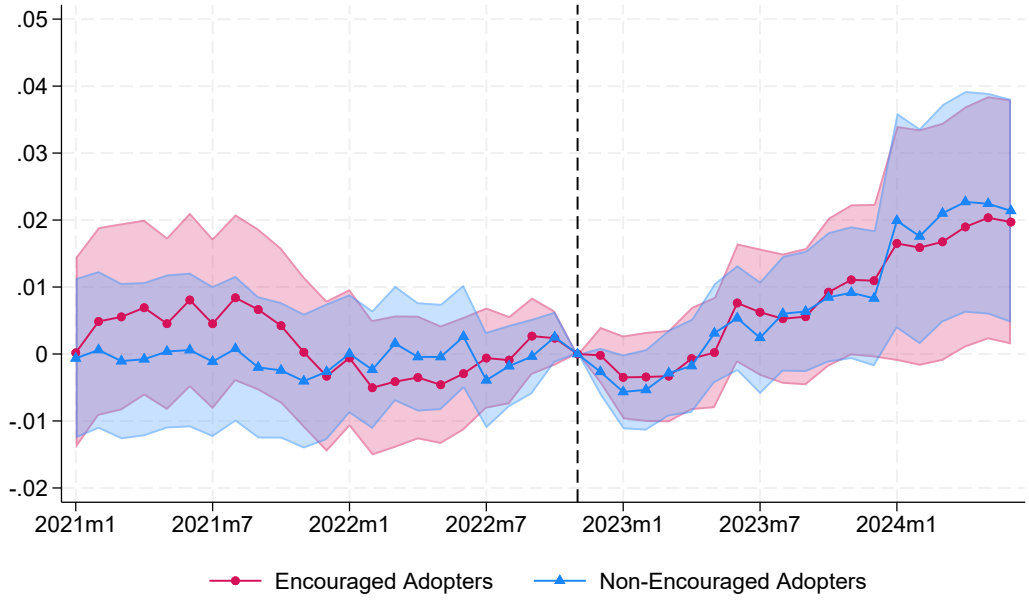
(c) Extensive-Margin Employment



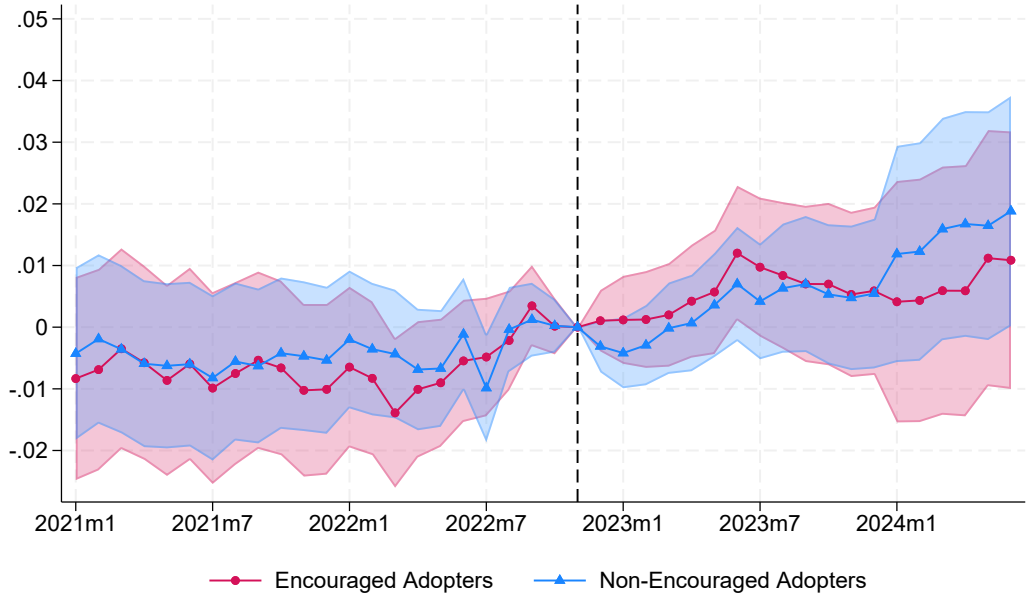
Notes: This figure presents the differential labor market outcomes of AI chatbot adopters relative to non-adopters, indexed to the launch of ChatGPT in November 2022. Effects are estimated separately for adopters whose employers encourage AI chatbot use ("Encouraged") and those without encouragement ("Non-Encouraged"). Each group is compared to non-encouraged non-adopters. Estimates are based on the dynamic difference-in-differences specification in Equation (2). Shaded areas represent 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

Figure E.2: Have Adopters Experienced Greater Job Mobility?
(DiD to Non-Encouraged Non-Adopters)

(a) FTE Employment in Latest Occupation

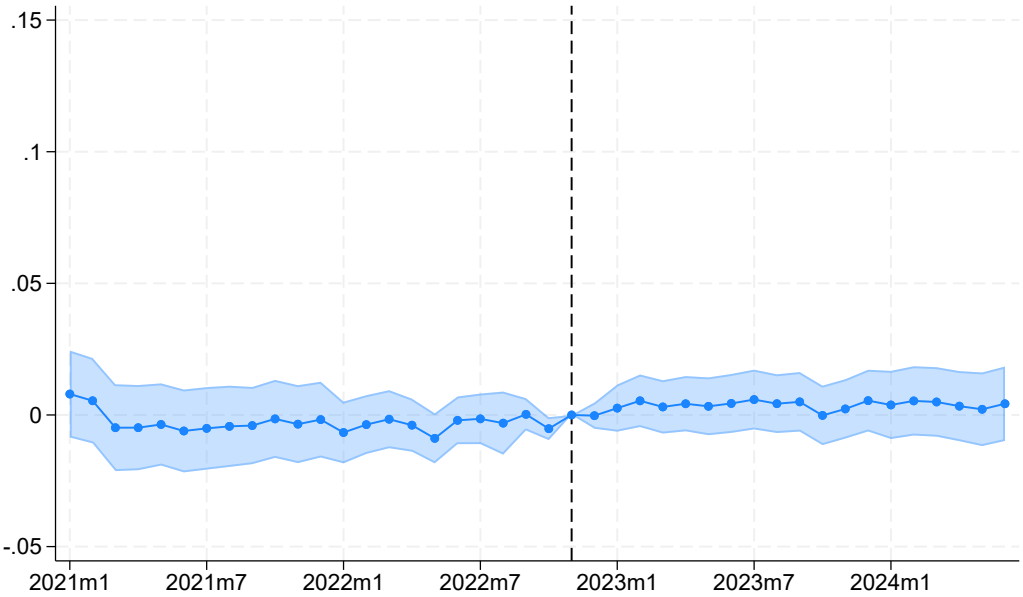


(a) FTE Employment in Latest Workplace



Notes: This figure presents difference-in-differences estimates of work hours in workers' latest jobs. Panel (a) focuses on occupations, while Panel (b) focuses on workplaces. The estimates compare changes for adopters—both encouraged and non-encouraged—relative to non-encouraged non-adopters, indexed to the launch of ChatGPT in November 2022. Estimates are based on the dynamic difference-in-differences specification in Equation (2). Shaded areas represent 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

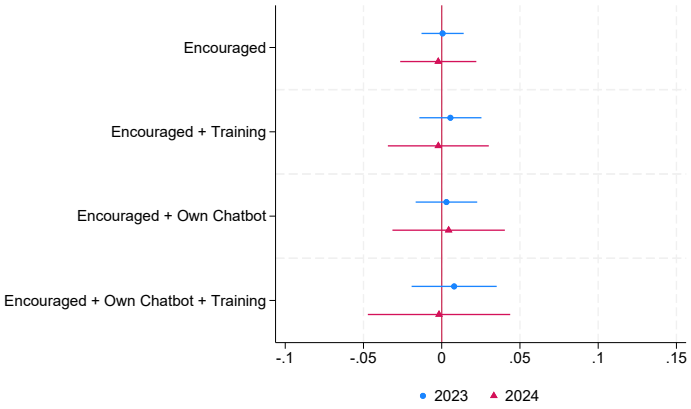
Figure E.3: Log Total Workplace Employment (Encouraged DiD)



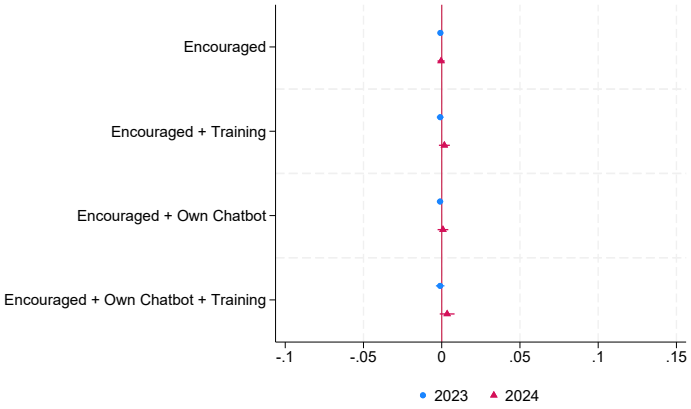
Notes: This figure shows total workplace employment for workers in encouraged-use workplaces relative to those without encouragement, indexed to November 2022. Estimates are based on the dynamic difference-in-differences specification in Equation (2), with shaded areas representing 95% confidence intervals. *Sample:* All completed 2024 survey responses linked to registry data.

Figure E.4: Workplace Effects of Chatbot Initiatives:
Heterogeneity by Strength of Treatment (DiD to Non-Encouraged)

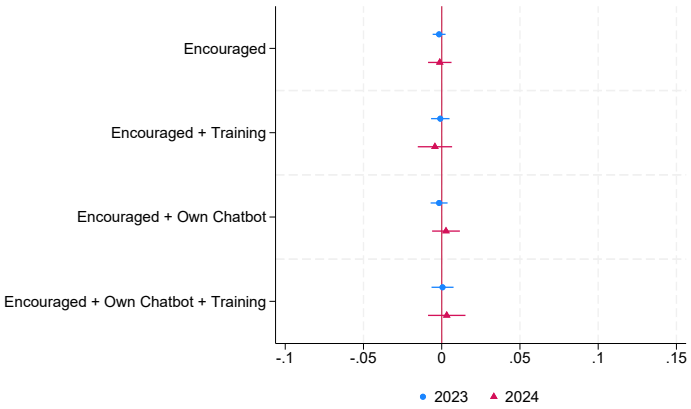
(a) Log Wage Bill



(b) Job Churn

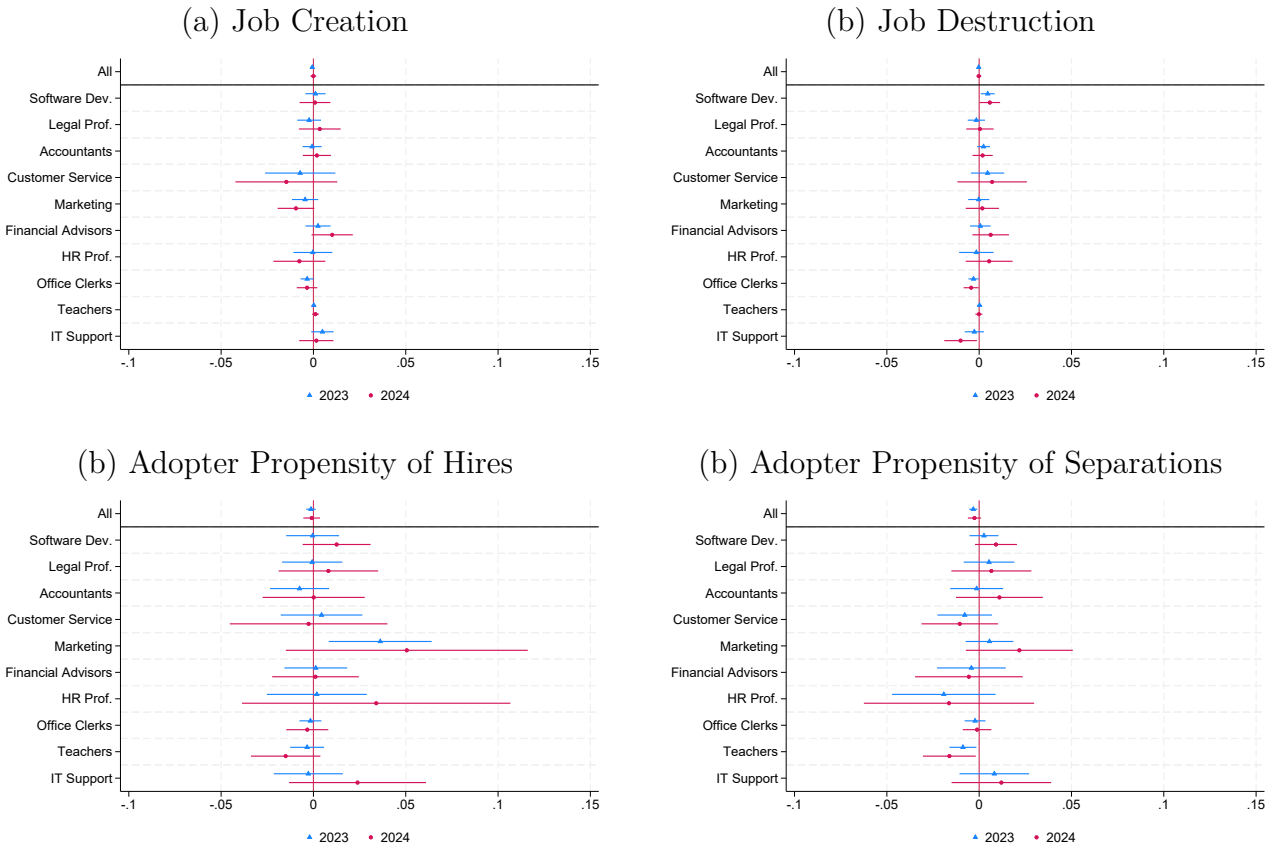


(c) Adopter Propensity of Hires/Separations



Notes: This figure shows workplace outcomes (job churn in Panel (a) and job sorting in Panel (b)) for workers who report encouraged usage policies relative to those without encouragement, indexed to November 2022. Panel (a) shows the log wage bill. Panel (b) shows job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment). Panel (c) shows the adopter propensity of hires and separations (see Appendix G.1 for details). Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Encouraged + Training* restricts the treatment group to encouraged workers who received employer-provided training; *Encouraged + Enterprise Chatbot*, to those whose employer deployed an enterprise chatbot; and *Encouraged + Enterprise Chatbot + Training*, to those with both. *Sample:* All completed 2024 survey responses linked to registry data.

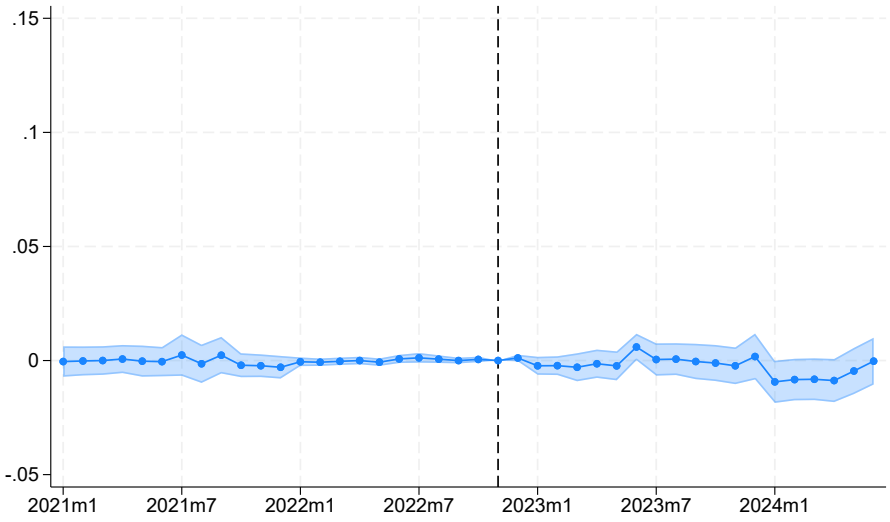
Figure E.5: Workplace Effects of AI Chatbots (Encouraged DiD): Heterogeneity by Occupation



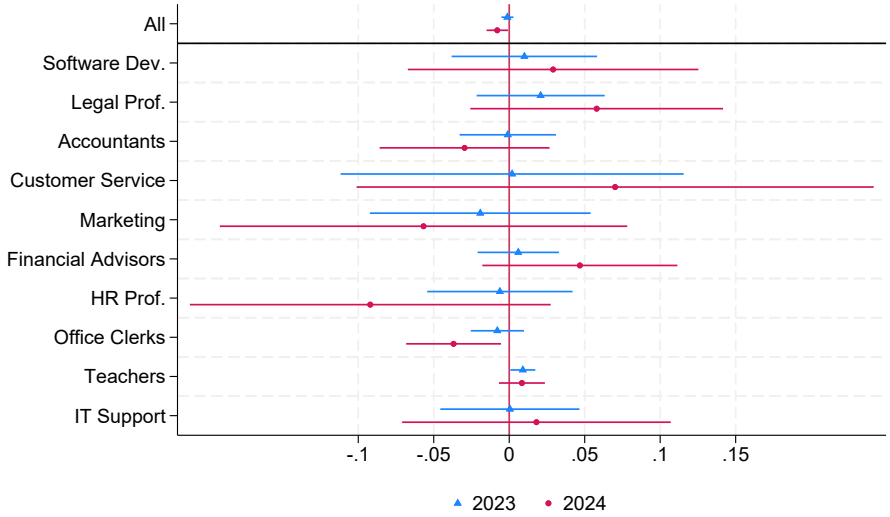
Notes: This figure shows occupation-specific estimates for the workplace effects of encouraged-use policies. Panels (a) and (b) show job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment). Panels (c) and (d) show the adopter propensity of hires and separations (see Appendix G.1 for details). Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Sample:* All completed 2024 survey responses linked to registry data.

Figure E.6: Incumbent Worker Outcomes by Adoption Rates (Dynamic Diff-in-Diff)

(a) Average Impacts (Dynamic Diff-in-Diff)

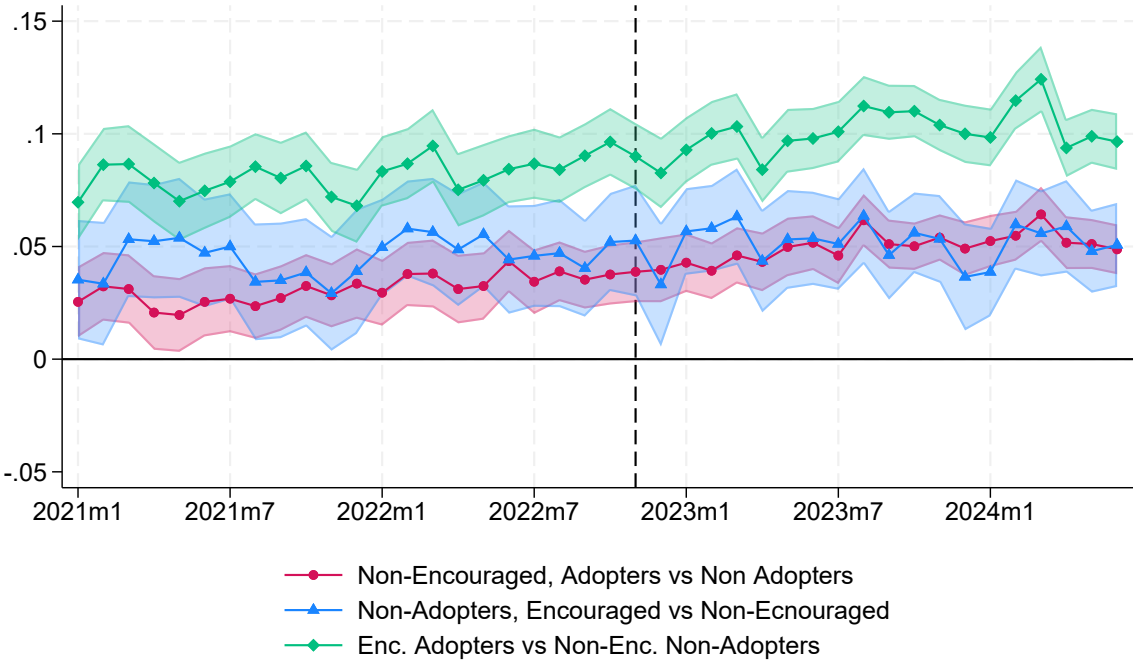


(b) Occupation-Level Impacts (Pooled Diff-in-Diff)



Notes: This figure shows changes in incumbent workers' outcomes by adoption rates, indexed to the launch of ChatGPT in November 2022, at the workplace-occupational level. Incumbent workers are defined as those employed continuously from January to November 2022, leading up to the arrival of ChatGPT. Adoption rates are measured as the share of employees in a worker's initial workplace who have used AI chatbots for work, adjusted for measurement error using the Empirical Bayes shrinkage procedure described in Appendix G.2. The standard deviation of workplace adoption rates within occupations, after shrinkage, is 20 percentage points. Panel (a) is based on the dynamic difference-in-differences specification in Equation (2), weighted by the number of respondents at each workplace; shaded areas indicate 95% confidence intervals. Panel (b) shows the effects separately by occupation, based on the pooled difference-in-differences specification in Equation (3), also weighted by the number of respondents at each workplace; whiskers indicate 95% confidence intervals. *Sample*: All completed responses from workplaces with at least two respondents in the 2024 survey linked to registry data.

Figure E.7: Have Adopting Workers Experienced Earnings Gains?
(Log Earnings Relative to Non-Encouraged Non-Adopters)



Notes: This figure shows the earnings gap of workers according to their encouragement and adoption status. Each group is compared to non-encouraged non-adopters. The figure reports the difference in log earnings, controlling for occupation fixed effects, pre-determined worker characteristics, and seasonality. Shaded areas represent 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

E.2 Supplementary Evidence

E.2.1 Perceived Impacts

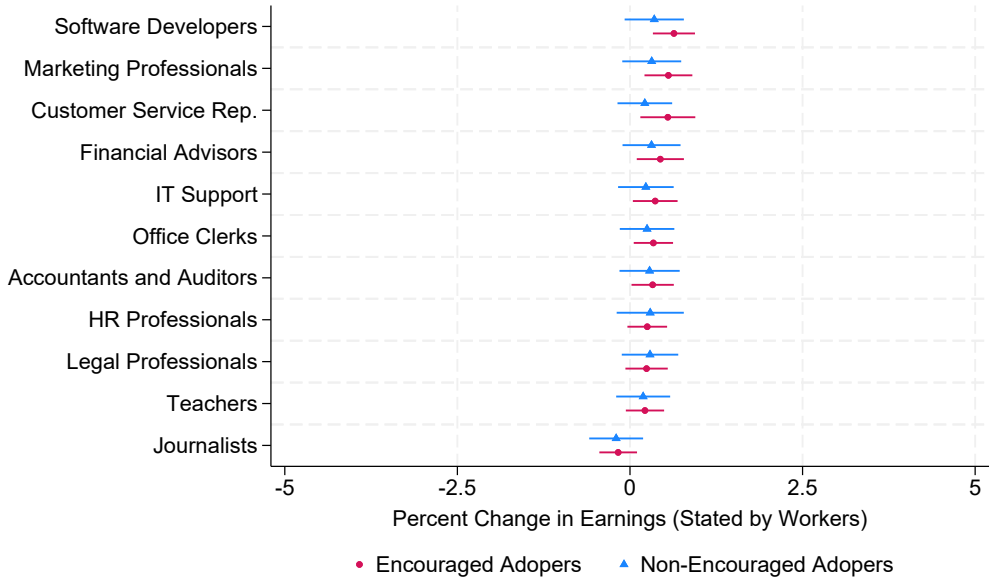
Table E.1: Decomposition of Perceived Earnings Effects of AI Chatbots Among Adopters

			Share of Impacts			Conditional Average Impacts (in %)		
Occupation	Policy	Average Impact (in %) (1)	Decreased Earnings (2)	Unchanged Earnings (3)	Increased Earnings (4)	Decreased Earnings (5)	Unchanged Earnings (6)	Increased Earnings (7)
Journalists	Non-Encouraged	-.118		.994			0	
Journalists	Encouraged	-.159		.983			0	
Software Developers	Non-Encouraged	.092	.008	.974	.015	-12.187	0	12.5
Software Developers	Encouraged	.428		.96	.037		0	11.944
Legal Professionals	Non-Encouraged	.031		.978			0	
Legal Professionals	Encouraged	.057		.987	.009		0	10.5
Accountants and Auditors	Non-Encouraged	0	.004	.915	.004	-13.75	0	13.75
Accountants and Auditors	Encouraged	.126		.973	.015		0	10
Customer Service Rep.	Non-Encouraged	-.11	.013	.913		-9.167	0	
Customer Service Rep.	Encouraged	.303		.942	.041		0	7.25
Marketing Professionals	Non-Encouraged	.112		.956	.023		0	7.25
Marketing Professionals	Encouraged	.37	.006	.955	.038	-16.5	0	12.422
Financial Advisors	Non-Encouraged	.022		.964			0	
Financial Advisors	Encouraged	.214		.979	.019		0	11.667
HR Professionals	Non-Encouraged	0		.964			0	
HR Professionals	Encouraged	.071		.986			0	
Office Clerks	Non-Encouraged	.015	.002	.901	.003	-7.143	0	8.462
Office Clerks	Encouraged	.113		.978	.014		0	10.403
Teachers	Non-Encouraged	-.007	.001	.975	.001	-9.167	0	4
Teachers	Encouraged	.041		.992	.007		0	7.188
IT Support	Non-Encouraged	.039	.006	.981		-4	0	
IT Support	Encouraged	.172		.974	.02		0	11.25

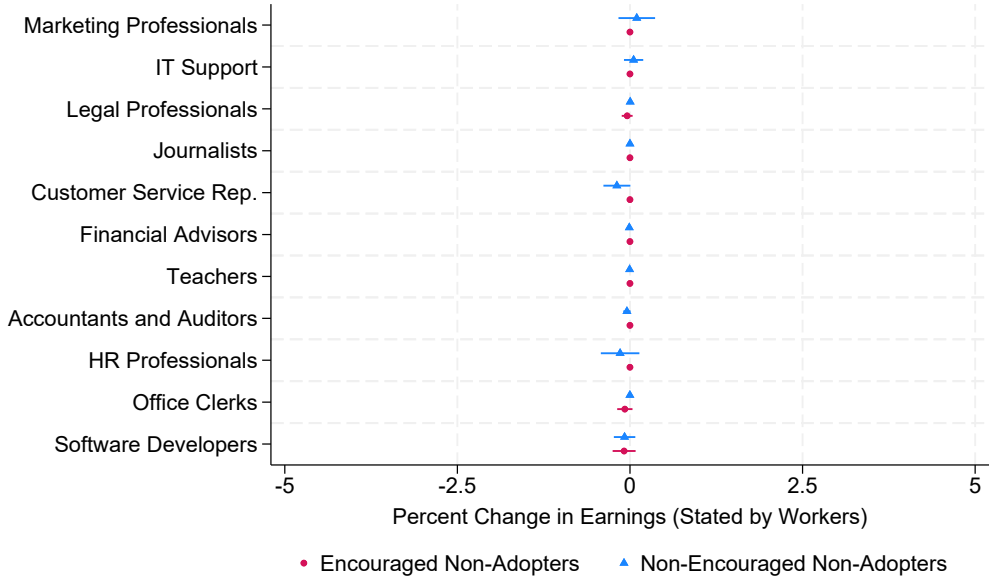
Notes: This table decomposes the average perceived earnings impact (in percent) of AI chatbots among adopters (Column 1) into contributions from workers reporting decreased earnings (Columns 2 and 5), unchanged earnings (Columns 3 and 6), and increased earnings (Columns 4 and 7). Responses are categorized by workers' occupations and whether their employers encourage AI chatbot use. Empty cells represent fewer than five individuals and are blanked out in accordance with confidentiality rules. We code reported earnings impacts as follows: "No change" as 0%, "0–5% change" as 2.5%, "5–15% change" as 10%, and "Over 15% change" as 20%. Our conclusions are robust to alternative codings. *Sample*: All completed responses from the 2024 survey round linked to registry data.

Figure E.8: Perceived Labor Market Impacts among Workers

(a) Adopters

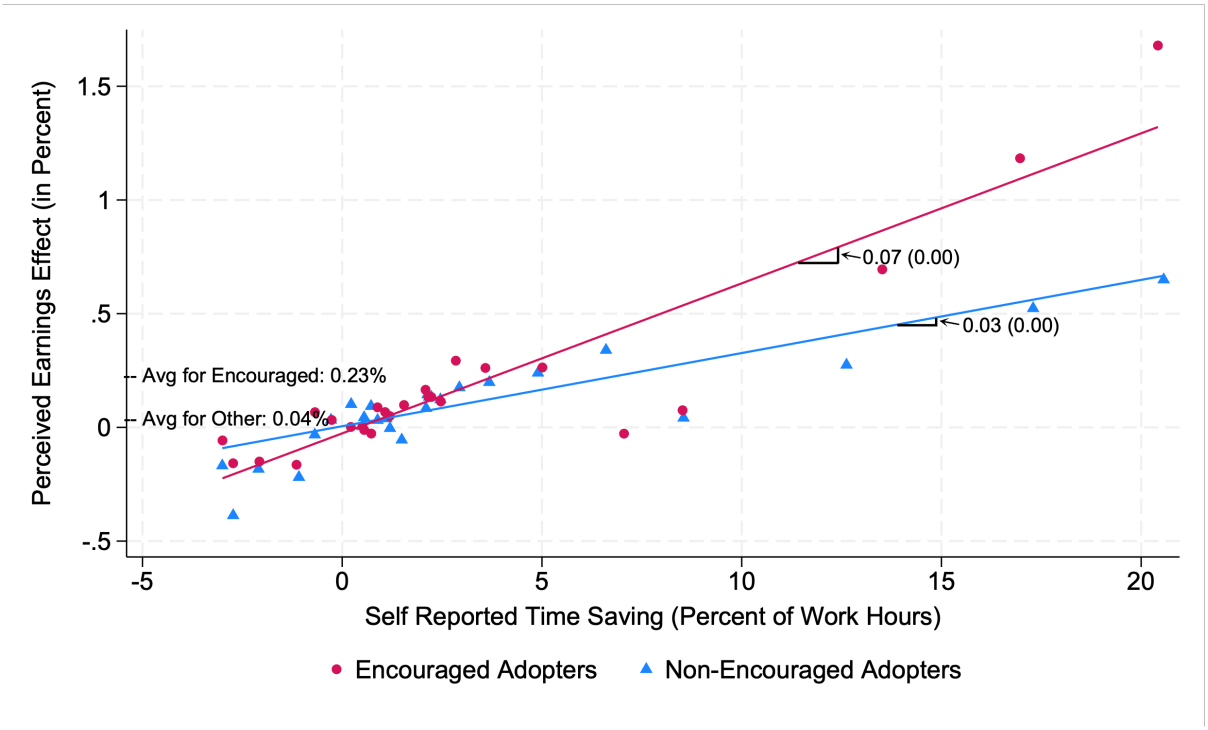


(b) Non-Adopters



Notes:This figure shows the average perceived impact of AI chatbots on earnings, broken down by workers' occupations and whether their employers encourage chatbot use. Panel (a) focuses on adopters—workers who have used AI chatbots for work. Panel (b) focuses on non-adopters—those who have not. See the notes of Table E.1 for how we code perceived earnings impacts. Whiskers represent 95% confidence intervals. *Sample*: All completed responses from the 2024 survey round linked to registry data.

Figure E.9: How Do Workers' Earnings Impacts Relate to Their Time Savings?

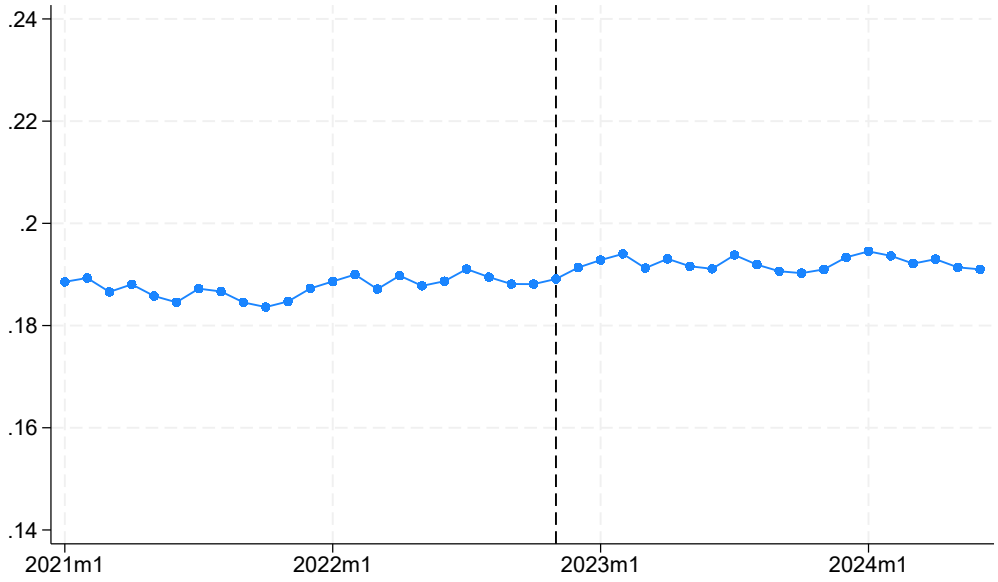


Notes: This figure presents a binned scatterplot of workers' perceived earnings impacts from AI chatbots against their estimated time savings from these tools, absorbing occupation fixed effects. The sample is divided based on whether employers encourage AI chatbot use. The regression line represents the line of best fit for each group, controlling for occupation fixed effects. *Sample*: All completed responses from the 2024 survey.

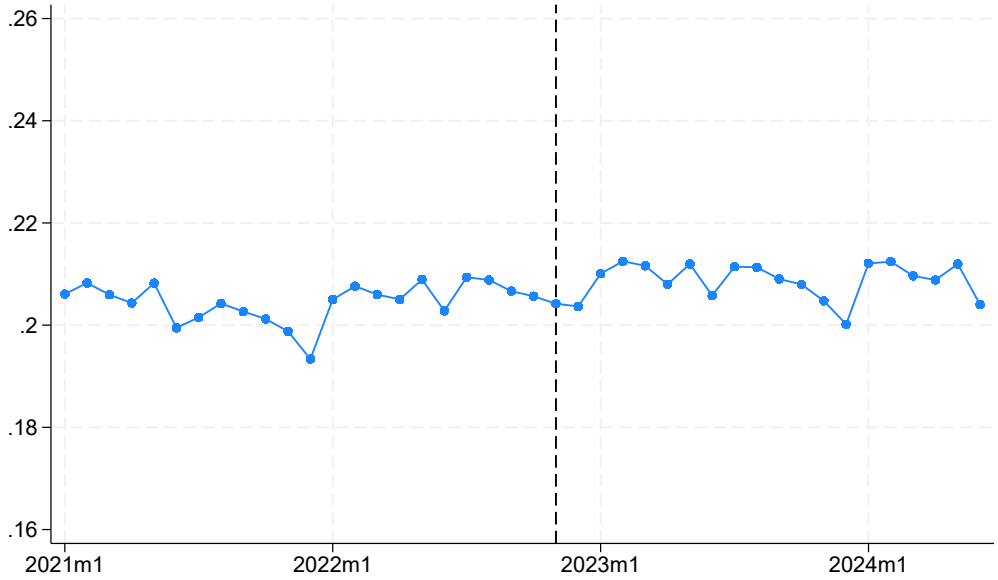
E.2.2 Spillovers to Non-Adopters

Figure E.10: Share of Study Occupations in Aggregate Danish Employment

(a) Hours



(b) Earnings



Notes: This figure shows the share of our 11 survey occupations in aggregate employment in Denmark. Panel (a) focuses on employment in hours, while Panel (b) focuses on total wage bill in the occupations. *Sample*: All registry data.

F Robustness Analysis

Sections 2 and 3 examine how employer encouragement affects workers' adoption and benefits of AI chatbots. This section assesses the robustness of our results to potential confounding factors. First, in Section F.1, we show that our findings are robust to controlling for a wide range of observable firm and worker characteristics, mitigating concerns that omitted variables drive the results. Second, in Section F.2, we compare individual-level and coworker-wide measures of encouragement. We find strong alignment between the two, supporting our interpretation of employer encouragement as a workplace-level treatment.

F.1 Additional Controls

This section evaluates the robustness of the relationships between employer encouragement of AI chatbots and workers' adoption and reported benefits. We leverage the richness of our data to examine potential confounders on both the firm and worker sides.

On the firm side, we show that all results remain robust when controlling for firm characteristics (Table 1), ensuring that observed differences across workplaces are not driven by variation in firm age, size, or productivity. On the worker side, we show that the results are similarly robust to controlling for workers' detailed task mixes within occupations, ensuring that the effects of employer encouragement are not merely driven by differences in task types more amenable to AI chatbot use.[43]

Table F.1 and Figure F.1 summarize these robustness checks, showing that our estimates of the impact of employer encouragement on chatbot adoption and perceived benefits remain virtually unchanged after accounting for firm characteristics and worker task mixes.

[43]Task importances are derived from our survey, which asked workers to rate the importance of six representative O*NET job tasks in their occupations; see Appendix I and Humlum and Vestergaard (2025, SI6) for details.

The robustness of our estimates to the inclusion of this rich set of controls provides reassurance that observable confounders are not driving the results. Moreover, Oster (2019) and Altonji, Elder and Taber (2005) provide conditions under which the stability of coefficients to the inclusion of observable controls can also help rule out selection on *un*observables.

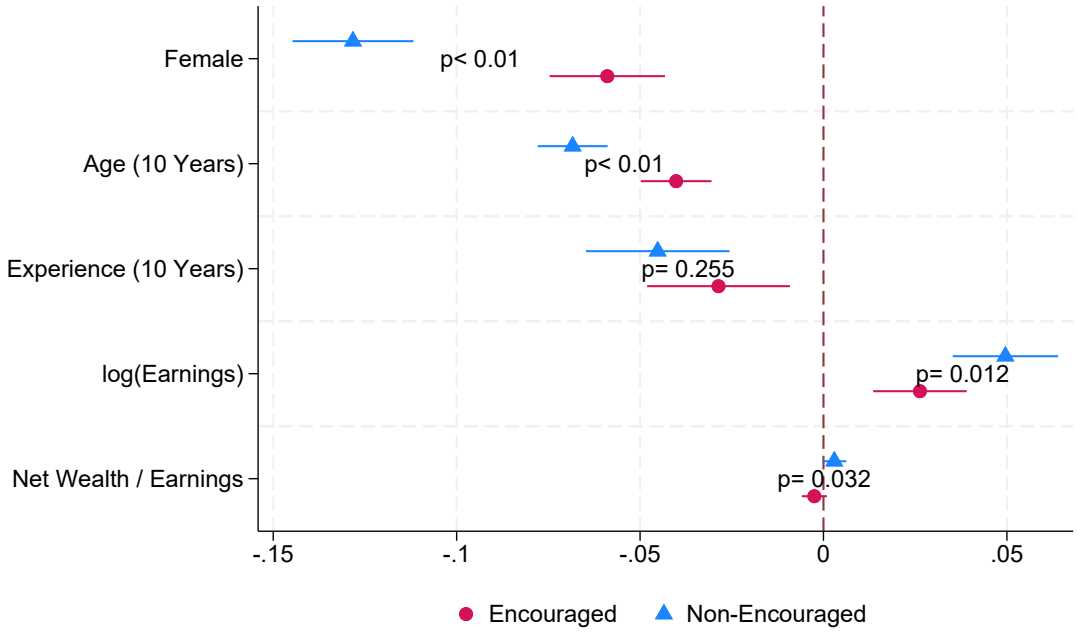
Table F.1: Encouragement, Adoption, and Work (Additional Controls)

	(a) Adoption (Figure 3)							
	Ever Used		Monthly or More		Weekly or More		Daily	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Encouraged	0.365	0.347	0.337	0.321	0.283	0.268	0.136	0.129
	(0.006)	(0.006)	(0.007)	(0.007)	(0.006)	(0.006)	(0.005)	(0.005)
Worker characteristics controls	✓	✓	✓	✓	✓	✓	✓	✓
Firm characteristics controls		✓		✓		✓		✓
Worker task mix controls		✓		✓		✓		✓
	(b) Reported Benefits (Table 2)							
	Time Savings		Quality		Creativity		Job Satisfaction	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Encouraged	0.098	0.099	0.115	0.112	0.094	0.094	0.071	0.069
	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.007)	(0.007)
Worker characteristics controls	✓	✓	✓	✓	✓	✓	✓	✓
Firm characteristics controls		✓		✓		✓		✓
Worker task mix controls		✓		✓		✓		✓
	(c) New Workloads from AI Chatbots for Adopters (Figure 6.(a))							
	Same Tasks		New Tasks		Same and New Tasks			
	(1)	(2)	(3)	(4)	(5)	(6)		
Encouraged	0.016	0.016	0.046	0.040	0.008	0.008		
	(0.004)	(0.004)	(0.005)	(0.005)	(0.002)	(0.002)		
Worker characteristics controls	✓	✓	✓	✓	✓	✓		
Firm characteristics controls		✓		✓		✓		
Worker task mix controls		✓		✓		✓		

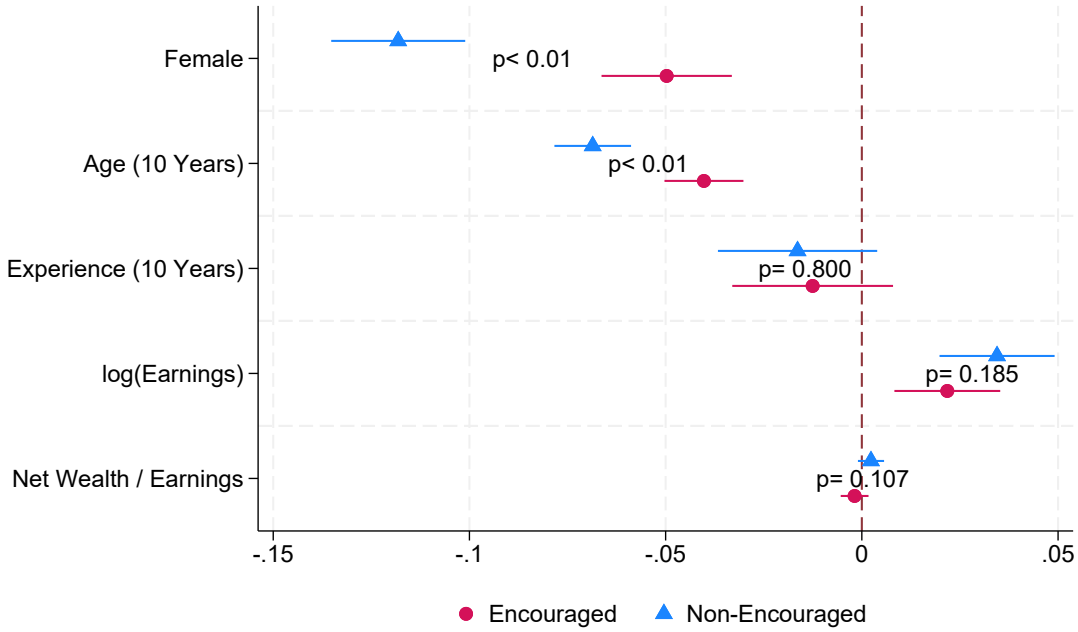
Notes: This table presents estimates of the impact of employer encouragement on AI chatbot adoption (Panel a), the reported benefits of adoption (Panel b), the allocation of time savings (Panel c), and new workloads resulting from chatbot use among adopters (Panel d). Odd-numbered columns report our main estimates from Equation (1). Even-numbered columns report specifications where we augment controls for firms' characteristics and workers' detailed task mixes. Robust standard errors are shown in parentheses. *Sample*: All complete responses from the 2024 survey that can be linked to the registry data.

Figure F.1: Employer Encouragement and Worker Gaps in Adoption (Figure C.3)

(a) Main Specification



(b) Additional Controls



Notes: This figure illustrates the impact of employer usage policies on worker disparities in AI chatbot adoption. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation fixed effects, and are estimated separately based on employers' AI chatbot initiatives (Encouraged = 1 vs. Encouraged = 0). Panel (a) presents our main specification from Figure C.3, while Panel (b) adds controls for firm characteristics (Table 1) and workers' detailed task mixes. Whiskers represent 95% confidence intervals. The reported p-values test whether the coefficients differ between the two groups. *Sample*: All completed responses from the 2024 survey linked to registry data.

F.2 Coworker Encouragement

A recurring theme in our analysis is the relationship between employer encouragement and workers' use of AI chatbots.

While our main analysis relies on workers' self-reports of employer encouragement, this section examines whether the results change when using coworker reports instead. This helps clarify whether the estimated effects reflect workplace-level treatments or individual-targeted interventions.

F.2.1 Empirical Strategy

To operationalize these ideas, we construct leave-one-out averages of coworkers' reported employer encouragement for each worker i:

$$\text{Encouraged}_{j,-i} = \frac{1}{N_{j(i)} - 1} \sum_{k \in j(i) \setminus i} \text{Encouraged}_k, \tag{70}$$

where $N_{j(i)}$ denotes the number of workers in the same workplace-occupation cell $j(i)$ as worker i.

The coworker measurements are made possible by our workplace-based sampling design (described in Section 1.5), which ensures that most respondents have coworkers who also participated in the survey. To mitigate measurement error due to incomplete sample coverage, we apply an empirical shrinkage procedure to the leave-out encouragement rates, $\text{Encouraged}_{j,-i}$ (see Appendix G.2 for details on this method). Importantly, our results in Section F.2.2 are robust to using the raw leave-out means instead.

Table F.2 shows that coworker encouragement strongly predicts an individual worker's own reported encouragement: the first-stage coefficient is approximately 1, with an F-statistic of 3,645. The strong correlation in reported employer encouragement at the workplace level supports our interpretation of these initiatives as centralized policies.

Table F.2: Relationship Between Self- and Coworker-Reported Encouragement

	Encouraged
	(1)
Coworker Encouragement (EBS)	1.181***
	(0.020)
Coworker Female Share	-0.016
	(0.011)
Coworker Age	0.001
	(0.001)
Coworker Potential Experience	-0.001
	(0.001)
Occupation FEs	✓
F-Stat (Partial)	3644.56
Observations	16974

Notes: This table reports the relationship between workers' self-reported employer encouragement and those reported by their coworkers, estimated using the specification in Equation (71). Standard errors (in parentheses) are clustered at the workplace level. *F-Stat (Partial)* reports the F-statistic on the partial effect of Coworker Encouragement (EBS). *Sample*: The table is based on all completed responses from workplaces with at least two respondents from the 2024 survey linked to registry data.

In Section F.2.2, we estimate how employer encouragement affects the impact of AI chatbots, contrasting the effects of workers' self-reported encouragement with those of their coworkers'. To ensure the effects are measured on comparable scales, and not attenuated by measurement error, we present 2SLS estimates for the impact of coworker-reported employer encouragement:

$$\text{Encouraged}_i = \pi' X_{j,-i} + \alpha \times \text{Encouraged}_{j,-i} + \varepsilon_{1i} \tag{71}$$

$$Y_i = \gamma' X_{j,-i} + \beta \times \widehat{\text{Encouraged}}_i + \varepsilon_{2i}, \tag{72}$$

where $X_{j,-i}$ denotes the leave-one-out mean of the characteristics X of worker i's coworkers. Table F.2 reports estimates of Equation (71).

F.2.2 Results

Adoption and Work. Table F.3 and Figure F.2 compare the effects of self- and coworker-reported employer encouragement on AI chatbot adoption and associated benefits. The results broadly align, but with interesting differences.

The impact of coworker-reported encouragement on adoption (Panel (a)) is 1.5–2 times larger than the effect of self-reported encouragement. Furthermore, Figure F.2.(b) shows that coworker encouragement can entirely eliminate the gender gap in AI chatbot adoption. By contrast, coworker encouragement has a more muted effect on reported benefits among adopters (Panel (b)). For example, self-reported encouragement increases the share of workers reporting time savings by 10 percentage points, whereas the effect of coworker encouragement is only 5.3 percentage points. Finally, Panel (d) shows that coworker encouragement has roughly twice the effect of self-reported encouragement on the creation of new tasks among adopters.

Taken together, these results suggest that workplace-wide encouragement (captured by coworker reports) significantly boosts chatbot adoption and leads to broader organizational changes that induce task creation. However, individual-level benefits—such as perceived time savings—appear more sensitive to whether workers personally feel encouraged, as measured by self-reports.

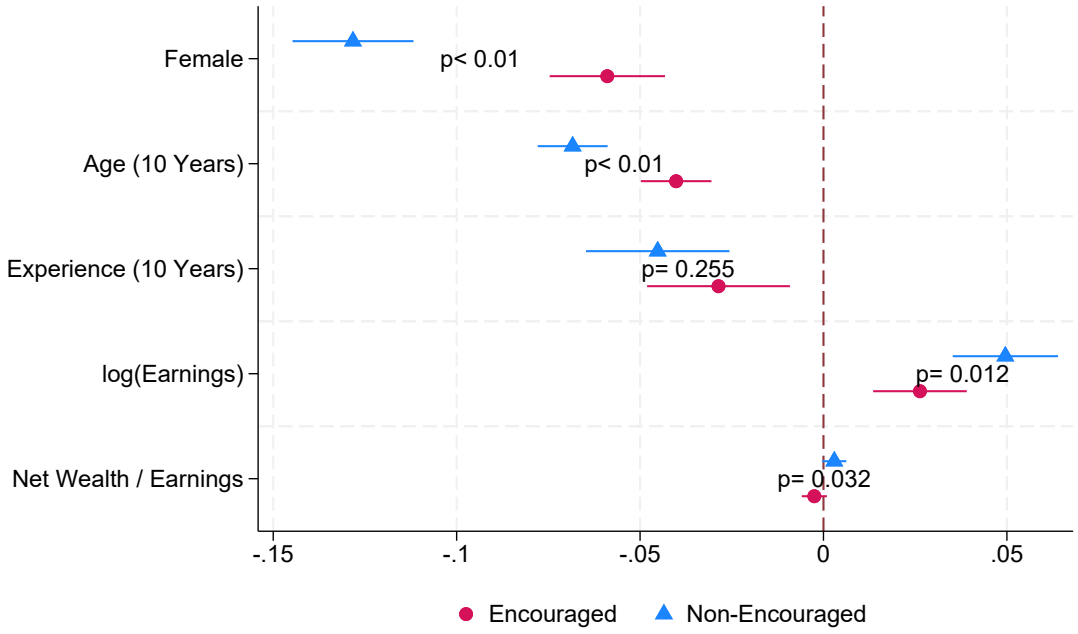
Table F.3: Encouragement, Adoption, and Work (Own vs. Coworker Reports)

	(a) Adoption							
	Ever Used		Monthly or More		Weekly or More		Daily	
	Own (1)	CWs (2)	Own (3)	CWs (4)	Own (5)	CWs (6)	Own (7)	CWs (8)
Encouraged	0.363	0.725	0.337	0.640	0.285	0.506	0.137	0.228
	(0.007)	(0.022)	(0.007)	(0.022)	(0.007)	(0.021)	(0.005)	(0.017)
	(b) Reported Benefits							
	Time Savings		Quality		Creativity		Job Satisfaction	
	Own (1)	CWs (2)	Own (3)	CWs (4)	Own (5)	CWs (6)	Own (7)	CWs (8)
Encouraged	0.100	0.053	0.116	0.090	0.091	0.088	0.071	0.026
	(0.008)	(0.031)	(0.009)	(0.033)	(0.009)	(0.031)	(0.007)	(0.025)
	(c) Allocation of Time Savings							
	More of Same Tasks		More of Diff. Tasks		More Breaks		More Leisure	
	Own (1)	CWs (2)	Own (3)	CWs (4)	Own (5)	CWs (6)	Own (7)	CWs (8)
Encouraged	0.031	0.035	0.016	0.037	-0.007	-0.016	-0.011	-0.010
	(0.009)	(0.035)	(0.007)	(0.024)	(0.005)	(0.016)	(0.006)	(0.019)
	(d) New Workloads from AI Chatbots (Adopters)							
	Same Tasks		New Tasks		Same and New Tasks			
	Own (1)	CWs (2)	Own (3)	CWs (4)	Own (5)	CWs (6)		
Encouraged	0.016	0.018	0.045	0.106	0.008	0.014		
	(0.004)	(0.012)	(0.006)	(0.021)	(0.002)	(0.008)		

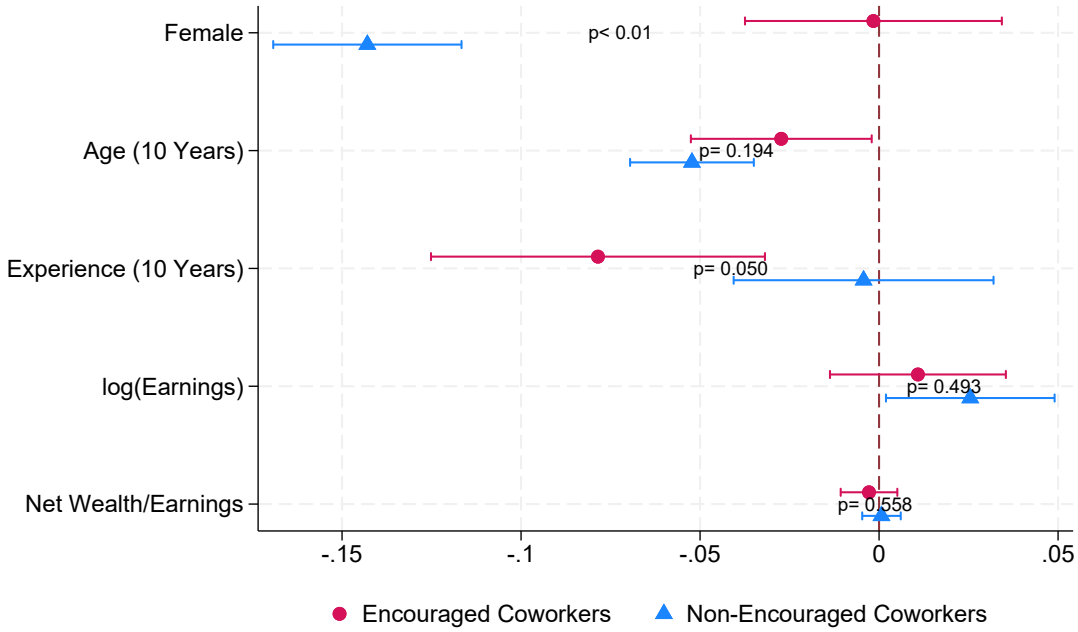
Notes: This table presents estimates of the impact of employer encouragement on: AI chatbot adoption (Panel a), reported benefits of adoption (Panel b), the allocation of time savings (Panel c), and new workloads resulting from chatbot use among adopters (Panel d). Odd-numbered columns ("Own") show the impact of self-reported encouragement, estimated using the specification in Equation (1). Even-numbered columns ("CWs") show the impact of coworker-reported encouragement, estimated using the specifications in Equations (71)–(72). Standard errors (in parentheses) are clustered at the workplace level. *Sample:* "Own" columns use all complete responses from the 2024 survey that can be linked to registry data. "CWs" columns include complete responses from workplaces with at least two linked respondents in the 2024 survey.

Figure F.2: Encouragement and Worker Gaps in Adoption (Own vs. Coworker Reports)

(a) Main Specification



(b) Coworker Encouragement

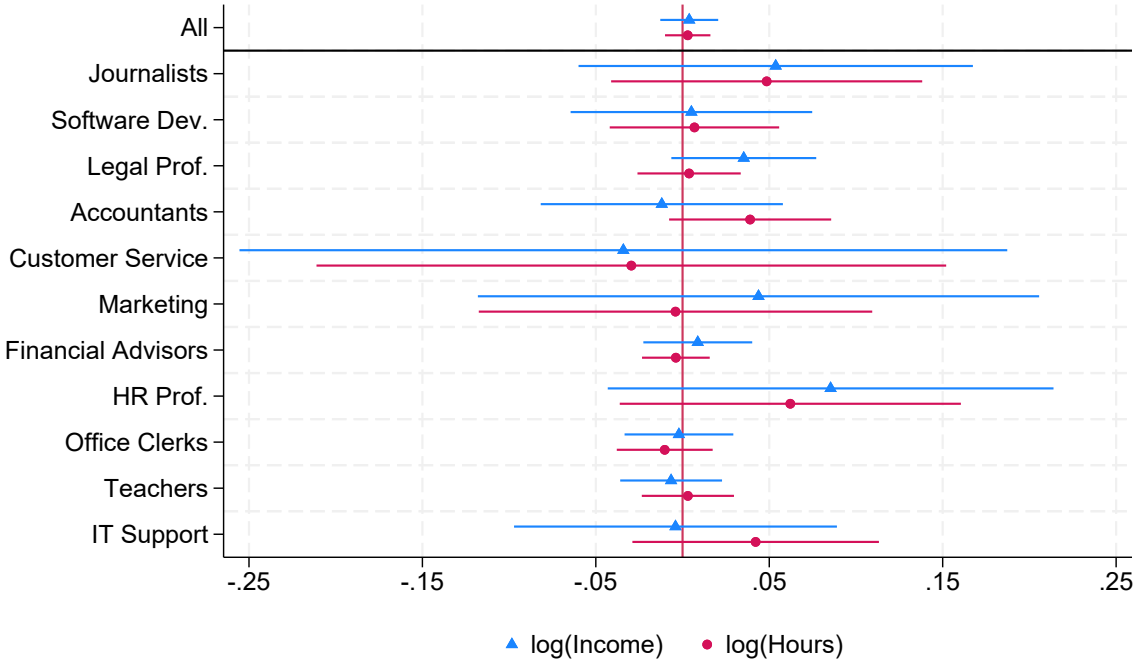


Notes: This figure presents the impact of coworker encouragement on worker disparities in AI chatbot adoption. The estimates are obtained from a multivariate regression of AI chatbot adoption on worker characteristics X, controlling for occupation-fixed effects, and all interacted with an indicator for employer encouragement. In Panel (b), we instrument these interaction effects using the corresponding interactions with our coworker encouragement variable. The point estimates show predicted adoption rates for an average worker, $\bar{X}$, when Encouraged = 0 or 1, respectively. Whiskers denote 95% confidence intervals. Reported p-values test whether the coefficients differ between the two groups. *Sample*: All completed responses from the 2024 survey linked to registry data.

Labor Market Outcomes. Figure F.3 revisits our analysis of how employer encouragement affects individual labor market outcomes, this time using coworker reports to measure encouragement. The results confirm that employer encouragement policies have had no measurable impact on workers' labor market outcomes.

A key takeaway from our earlier analysis is that the modest labor market effects are partly explained by weak pass-through from time savings to labor earnings—though pass-through is somewhat stronger when employers actively encourage chatbot use. Table F.4 replicates this result using coworker reports to measure encouragement. Pass-through remains modest (ranging from 2% to 6%) but is approximately three times higher when employers are reported to encourage chatbot use.

Figure F.3: Have Encouragements Affected Workers' Outcomes? (Coworker Reports)



Notes: This figure presents the differential labor market outcomes of workers who are encouraged to use AI chatbots by their employers, compared to all other workers, indexed to the launch of ChatGPT in November 2022. The figure is based on the pooled difference-in-differences specification in Equation (3), with whiskers representing 95% confidence intervals. The estimates come from a 2SLS specification where we instrument all effects of Encouraged using the corresponding coworker encouragement rates described in Section F.2.1. *Sample:* All completed responses from the 2024 survey linked to registry data.

Table F.4: Pass-Through of Time Savings to Earnings (Own vs. Coworker Encouragement)

	Baseline		Encouraged	
	Own	CWs	Own	CWs
	(1)	(2)	(3)	(4)
Pass-through rate	0.031	0.018	0.066	0.059
	(0.005)	(0.032)	(0.003)	(0.016)

Notes: This table reports estimates of the pass-through rate from workers' perceived time savings due to AI chatbots to their perceived earnings impacts. Estimates are reported separately based on whether employers encourage chatbot use (Encouraged = 0 or 1). The "Own" coefficients correspond to the slopes of the best-fit lines in Figure E.9. The "CWs" estimates instrument Encouraged using the coworker encouragement rates described in Section F.2.1. *Sample*: The table is based on all complete responses from the 2024 survey that can be linked to the registry data.

G Estimation Procedures

G.1 Adopter Propensity

We measure adopter propensity of job creation/destruction as the predicted adoption rate among hires/separations at workplace j in month t based solely on workers' pre-determined individual characteristics:

1. Estimate adoption propensity:

$$\text{Adopt}_i = \gamma' X_i + \beta' \text{EmployerInitiatives}_i + \varepsilon_i, \tag{73}$$

   where X_i includes predetermined worker characteristics (occupation, gender, age, experience, and earnings)

2. Among workplaces j with positive hires/separations in month t, measure the adopter propensity of their job creation/destruction as:

$$\text{JobCreationAdopterPropensity}_{jt} = \hat{\gamma}' \mathbb{E}[X_i \mid j_t(i) = j, j_{t-1}(i) \neq j] \tag{74}$$

$$\text{JobDestructionAdopterPropensity}_{jt} = \hat{\gamma}' \mathbb{E}[X_i \mid j_t(i) \neq j, j_{t-1}(i) = j], \tag{75}$$

where $j_t(i)$ is the workplace of worker i in month t.

Workers' adopter propensities $\hat{\gamma}'X_i$ correspond to their willingness to adopt AI chatbots i in our theoretical framework (Section A).

G.2 Empirical Bayes Shrinkage

As a robustness check, we estimate workplace rates of encouragement using Empirical Bayes shrinkage with a Beta-Binomial model; see Walters (2024) for a detailed introduction to Empirical Bayes methods. The shrinkage is performed separately for each occupation, allowing underlying encouragement rates to vary systematically across occupations.

We assume the adoption rate at each workplace, p_i, follows a Beta prior:

$$x_i \mid p_i \sim \text{Binomial}(n_i, p_i), \quad p_i \sim \text{Beta}(\alpha_0, \beta_0). \tag{76}$$

The Beta prior captures workplace-level variation. We estimate α_0, β_0 via Method of Moments, matching the Beta distribution's first two moments to observed data:

$$\bar{p} = \frac{1}{m}\sum_{i=1}^{m}\frac{x_i}{n_i}, \quad s^2 = \frac{1}{m}\sum_{i=1}^{m}\left(\frac{x_i}{n_i} - \bar{p}\right)^2. \tag{77}$$

From the Beta mean and variance formulas:

$$\alpha_0 = \frac{\bar{p}(1-\bar{p})}{s^2} - 1, \quad \beta_0 = \alpha_0\frac{1-\bar{p}}{\bar{p}}.$$

With these, we compute the posterior mean:

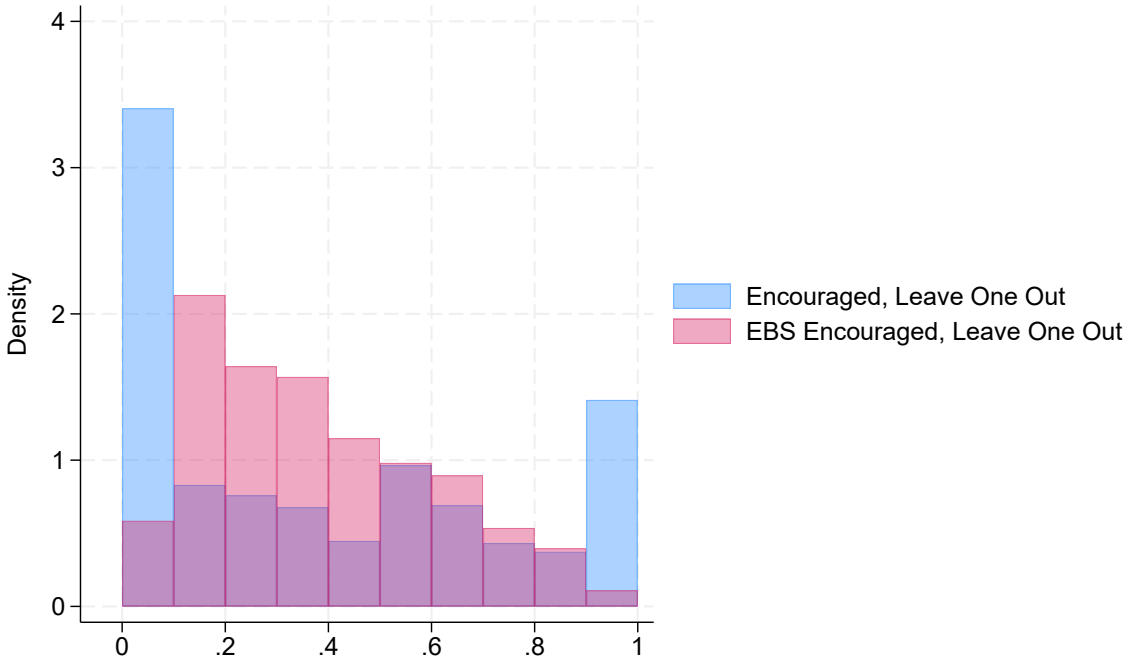
$$\mathbb{E}[p_i \mid x_i] = \frac{\alpha_0 + x_i}{\alpha_0 + \beta_0 + n_i}.$$

This shrinks estimates toward the overall mean, especially for small n_i.

G.2.1 Coworker Encouragement Rates

Figure G.1 compares the raw and adjusted distributions of coworker encouragement rates, while Table G.1 presents summary statistics for the adjusted rates of workplaces. The typical standard deviation within occupations is 17 percentage points. Importantly, our results in Section F.2.2 remain robust when using the raw coworker encouragement rates instead.

Figure G.1: Coworker Encouragement Rates (Raw vs. Shrinkage)



Notes: This figure compares the raw and adjusted distributions of coworker encouragement rates. The adjusted estimates are derived using an Empirical Bayes shrinkage procedure, as described in Section G.2. *Sample:* All completed responses from our 2024 survey round linked to registry data.

Table G.1: Coworker Encouragement Rates (Empirical Bayes Shrinkage)

	p25	p50	p75	sd
Journalists	.414	.457	.58	.138
Software Developers	.53	.68	.768	.16
Legal Professionals	.092	.178	.465	.23
Accountants	.289	.409	.555	.181
Customer Service Rep.	.291	.358	.466	.116
Marketing Professionals	.607	.823	.869	.18
Financial Advisors	.196	.35	.605	.237
HR Professionals	.289	.509	.737	.244
Office Clerks	.232	.31	.448	.149
Teachers	.116	.157	.243	.11
IT Support	.359	.48	.593	.162
All	.310	.428	.575	.173

Notes: This table presents summary statistics for the adjusted distributions of coworker encouragement rates, categorized by occupation. The adjusted estimates are derived using an Empirical Bayes shrinkage procedure, as described in Section G.2. *Sample:* All completed responses from our 2024 survey round linked to registry data.

H Invitation Letter

This section includes the invitation letter for our survey. We sent three reminders: two via email (Digital Post) and one via text message (SMS).

The section below focuses on our 2024 survey round. The invitation letter for the 2023 round follows the same format and is documented in Humlum and Vestergaard (2025).

The English translation begins on page 87, followed by the original Danish version on page 89.

November 2024

Artificial intelligence and your job tasks

Dear [name]

Statistics Denmark is inviting you to participate in a research project about AI chatbots and your job tasks. You can participate by clicking the link below and completing the questionnaire.

AI chatbots use artificial intelligence to read and write text. You have been selected because you work in an occupation where AI chatbots may be relevant.

Your responses are important for research on new technology in the labor market. Everyone who completes the questionnaire will automatically participate in a lottery with **a prize of [X,XXX] DKK tax-free**.

Statistics Denmark is conducting the survey on behalf of researchers at the University of Copenhagen and the University of Chicago. The questionnaire takes **about 10 minutes** to complete.

Start the survey [url]

Or access www.dst.dk/ditsvar and enter your response code **[code]**.

Statistics Denmark handles your data confidentially. Results are presented in a way that prevents individual answers from being identified, and the data is used solely for statistical and scientific purposes.

Participation is voluntary. If you do not wish to participate, you can indicate this here: [refusal_link]

If you have any questions, you can e-mail info@dstsurvey.dk or call on 7777 7708 (every day between 9am and 4pm). Please provide your response code when contacting us.

Best regards,

Marie Fuglsang
Head of Division, DST Survey

Anders Humlum
Assistant Professor, University of Chicago

Information about Statistics Denmark's surveys and your rights

Who is invited to Statistics Denmark's surveys?

Anyone residing in Denmark may be invited to participate in one of Statistics Denmark's surveys. Participants are randomly selected. Our surveys aim to reflect the opinions and attitudes of the entire population, across gender, age, education, and place of residence.

Why is Statistics Denmark allowed to contact you?

Statistics Denmark can use its statistical production and related activities to carry out tasks under the rules for revenue-financed activities. This is stipulated in §1, section 3, no. 5, of the Act on Statistics Denmark.

How do we process your information?

The responses you provide in the survey are handled in accordance with the European General Data Protection Regulation (GDPR) and the Danish Data Protection Act.

The University of Copenhagen is the data controller for this survey. You can read more about the data controller and find contact information here: https://informationssikkerhed.ku.dk/persondatabeskyttelse/publikation-af-videnskab/

Statistics Denmark is the data processor and is responsible for data collection on behalf of the data controller.

Your responses will only be used for statistical and scientific purposes in this survey. Your answers will be deleted or archived in accordance with applicable laws when they are no longer needed for the study.

You can read more about how we process your data at: https://www.dst.dk/privatlivspolitik-i-en-frivillig-undersoegelse.

If you have any other questions regarding the processing of your personal data, you are welcome to contact Statistics Denmark's Data Protection Officer at databeskyttelse@dst.dk.

November 2024

Kunstig intelligens og dine arbejdsopgaver

Kære [navn]

Danmarks Statistik inviterer dig til at deltage i et forskningsprojekt om AI chatbots og dine arbejdsopgaver. Du deltager ved at klikke på nedenstående link og svare på spørgeskemaet.

AI chatbots bruger kunstig intelligens til at læse og skrive tekst. Du er blevet udvalgt til at deltage i denne undersøgelse, fordi du arbejder i et erhverv, hvor det kan være relevant at bruge AI chatbots.
Dine svar er vigtige for forskning i ny teknologi på arbejdsmarkedet. Alle der gennemfører spørgeskemaet, deltager automatisk i lodtrækningen om **en præmie på [X.XXX] kr. skattefrit**.

Danmarks Statistik gennemfører spørgeskemaet for forskere på Københavns Universitet og University of Chicago. Det tager **ca. 10 minutter** at besvare spørgeskemaet.

Start undersøgelsen [url]

Eller gå ind på www.dst.dk/ditsvar og tast svarkoden **[kode]**

Danmarks Statistik behandler dine svar fortroligt. Vi formidler resultaterne på en måde, så ingen kan se, hvad den enkelte har svaret og data anvendes alene til statistiske og videnskabelige formål.

Det er frivilligt at deltage. Ønsker du ikke at deltage, kan du tilkendegive det: [refusal_link]

Har du spørgsmål, kan du skrive til info@dstsurvey.dk eller ringe på tlf. 7777 7708 (alle dage ml. kl. 9-16). Oplys venligst din svarkode ved henvendelse.

Med venlig hilsen

Marie Fuglsang
Kontorchef, DST Survey

Anders Humlum
Adjunkt, University of Chicago

Information om Danmarks Statistiks undersøgelser og dine rettigheder

Hvem bliver inviteret til Danmarks Statistiks undersøgelser?

Alle, der har bopæl i Danmark, har mulighed for at blive inviteret til at deltage i en af Danmarks Statistiks undersøgelser. Udvælgelse af personer til undersøgelsen sker tilfældigt. I vores undersøgelser er det vigtigt at kende meninger og holdninger fra hele befolkningen på tværs af køn, alder, uddannelse og bopæl.

Hvorfor må Danmarks Statistik kontakte dig?

Danmarks Statistiks kan bruge den statistiske produktion og afledte aktiviteter til at udføre opgaver efter reglerne for indtægtsdækket virksomhed. Det følger af § 1, stk. 3, nr. 5, i lov om Danmarks Statistik.

Hvordan behandler vi oplysninger om dig?

De svar, du afgiver ved deltagelse i spørgeskemaundersøgelsen, bliver behandlet i overensstemmelse med reglerne i den europæiske databeskyttelsesforordning (GDPR) og den danske databeskyttelseslov.

Københavns universitet er dataansvarlig for undersøgelsen. Du kan læse mere om den dataansvarlige og finde kontaktoplysninger her: https://informationssikkerhed.ku.dk/persondatabeskyttelse/publikation-af-videnskab/

Danmarks Statistik er databehandler og står for dataindsamlingen på vegne af den dataansvarlige.

Dine svar bruges udelukkende til statistiske og videnskabelige formål i denne undersøgelse. Dine svar slettes eller arkiveres efter gældende lovgivning, når oplysningerne ikke længere har et formål i undersøgelsen.

På linket https://www.dst.dk/privatlivspolitik-i-en-frivillig-undersoegelse kan du læse mere om, hvordan vi behandler oplysninger om dig.

Har du andre spørgsmål til behandling af dine personoplysninger, er du velkommen til at kontakte Danmarks Statistiks databeskyttelsesrådgiver på databeskyttelse@dst.dk

I Survey Questionnaire

This section contains our survey questionnaire. The questionnaire follows a common structure for the different occupations but with job tasks and titles tailored to each specific occupation.

For the sake of brevity, the questionnaire below focuses on one occupation (journalism), listing one of their six job tasks (write commentaries, columns, or scripts).

The section below focuses on our 2024 survey round. The survey questionnaire for the 2023 round follows a similar format and is documented in Humlum and Vestergaard (2025).

The English translation starts on page 92, with the original Danish version on page 97.

1. Introduction

AI chatbots use artificial intelligence to read and write text. You have been selected to participate in this survey because you work in a profession where AI chatbots may be relevant.

Your participation is important regardless of your knowledge of chatbots or artificial intelligence.

Block 1: Occupation and tasks

2.a Occupation

Are you employed in [journalism]?

- Yes
- No

2.b Occupation [if 2.a='No']

Are you employed in one of the following areas?

If you are employed in multiple areas, please select your primary work area.

- HR work
- IT support
- Office and secretarial work
- Customer support
- Legal work
- Marketing
- Auditing and accounting work
- Software development
- Teaching
- Financial consulting
- I am not employed in any of the above work areas

3. Task Importance [if 2.b!= 'I am not employed in any of the above work areas'; all tasks]

We will first ask about some typical work tasks among [journalists].

For each task, please assess how **important the task is for your work**.

Extremely important means that the task is critical for performing your current job.

[Write commentaries, columns, or scripts]

- Not important
- Slightly important
- Important
- Very important
- Extremely important

Block 2: Adoption

4. Awareness of AI chatbots

AI chatbots use artificial intelligence to read and write text. We will now ask about your experiences with AI chatbots.

Had you heard of the following chatbots before this survey?

Mark all tools you had heard of before this survey.

- ChatGPT (developed by OpenAI)
- Claude (developed by Anthropic)
- Copilot (developed by Microsoft)
- Gemini (developed by Google)
- Perplexity (developed by Perplexity AI)
- Other AI chatbots
- Had not heard of AI chatbots before this survey

5. Prior Use of AI Chatbots [if 4 = 'Yes']

Have you used the following AI chatbots?

[ChatGPT / Claude / Copilot / Gemini / Perplexity / Other AI chatbots]

- Yes, only for work
- Yes, only for leisure
- Yes, for work and leisure

6.a Purposes of Prior Use [if 5='Yes, only for leisure' or 'Yes, for work and leisure']

How often have you used the following AI chatbots **for leisure**?

[ChatGPT / Claude / Copilot / Gemini / Perplexity / Other AI chatbots]

- Never
- A few times
- Monthly
- Weekly
- Daily

6.b Purposes of Prior Use [if 5='Yes, only for work or 'Yes, for work and leisure']

How often have you used the following AI chatbots **for work**?

[ChatGPT / Claude / Copilot / Gemini / Perplexity / Other AI chatbots]

- Never
- A few times
- Monthly
- Weekly
- Daily

6.c Purposes of Prior Use [if 5='Yes, only for work or 'Yes, for work and leisure' for any option]
Have you used an AI chatbot to perform the following work tasks?

- [Job task 1-6]
- None of the above

7. Time use on AI chatbots [if 5='Yes, only for work or 'Yes, for work and leisure' for any option]
Think back to the days when you used AI chatbots for your work. How much time did you spend using AI chatbots on average?

- Less than 15 minutes per day
- Between 15 minutes and an hour per day
- More than an hour per day

8. Paid subscription [if 5='Yes, only for work or 'Yes, for work and leisure' for 'ChatGPT']
Do you have an active Plus subscription for ChatGPT?

- Yes
- No

Block 3: Employer Initiatives

9. Employer policies
What is your employer's policy regarding the use of AI chatbots?

- Allowed and encouraged
- Allowed but not encouraged
- Not allowed
- No policy
- Don't know

10. Enterprise AI chatbot
Does your workplace have its own AI chatbot?

- Yes, a custom-designed product
- Yes, a standard product
- No
- Don't know

11.a Training courses
Have you participated in courses on using AI chatbots?

- Yes
- No

11.b Training courses [if 11.a = 'Yes']
Was your AI chatbot course organized by your employer?

- Yes
- No

Block 4: Effects of AI chatbots

12. Benefits from AI chatbots

Have you experienced any of these benefits from using AI chatbots in your work?
Please select all that apply.

- Saved time at work
- Improved work quality
- Increased creativity
- Higher job satisfaction
- Have not experienced benefits
- Don't know

13. Time savings from AI chatbots [if 12='Saved time at work']

Think back to the days when you used AI chatbots for your work. How much time did you save using AI chatbots on average?

- Less than 15 minutes per day
- Between 15 minutes and an hour per day
- More than an hour per day

14.a Earnings impact of AI chatbots

Have AI chatbots affected how much you earn today?

- I earn more today as a result of AI chatbots
- AI chatbots have not affected my income
- I earn less today as a result of AI chatbots

14.b Earnings impact of AI chatbots [if 14.a=' I earn more today as a result of AI chatbots']

How much have AI chatbots increased your earnings?

- Under 5 percent
- Between 5 and 15 percent
- Over 15 percent

14.c Earnings impact of AI chatbots [if 14.a=' I earn less today as a result of AI chatbots']

How much have AI chatbots reduced your earnings?

- Under 5 percent
- Between 5 and 15 percent
- Over 15 percent

15. Allocation of time savings from AI chatbots

If AI chatbots save time on a task, do you expect to:

- Complete more of the same tasks
- Spend more time on other tasks
- Take more breaks
- Take more leisure time

16. Workloads from AI chatbots

Do you find that AI chatbots have increased your workload?

- Yes, more of the same job tasks
- Yes, new types of job tasks
- No

17. New job tasks [if 16='Yes, new types of job tasks']

What types of new tasks have you experienced after using AI chatbots?

- [open text field]

18. End of survey

Thank you for participating in the survey.

If you win one of the prizes, you will be notified directly in your e-Boks.

1. Introduction

AI chatbots bruger kunstig intelligens til at læse og skrive tekst. Du er blevet udvalgt til at deltage i denne undersøgelse, fordi du arbejder i et erhverv, hvor det kan være relevant at bruge AI chatbots. Din deltagelse er vigtig uanset dit kendskab til chatbots eller kunstig intelligens.

Block 1: Occupation and tasks

2.a Occupation

Er du beskæftiget med [journalistik]?

- Ja
- Nej

2.b Occupation [if 2.a='Nej']

Er du beskæftiget inden for et af følgende områder?

Hvis du er beskæftiget indenfor flere områder, vælg da dit primære arbejdsområde.

- HR-arbejde
- IT-support
- Kontor- og sekretærarbejde
- Kundesupport
- Juridisk arbejde
- Marketing
- Revisions- og regnskabsarbejde
- Softwareudvikling
- Undervisning
- Økonomisk rådgivning
- Jeg er ikke beskæftiget inden for ovenstående arbejdsområder

3. Task Importance [if 2.b!= 'Jeg er ikke beskæftiget inden for ovenstående arbejdsområder'; all tasks]

Vi vil først spørge ind til nogle typiske arbejdsopgaver blandt [journalister].

Til hver opgave bedes du vurdere, hvor **vigtig opgaven er for dit arbejde**.

Ekstremt vigtig betyder, at opgaven er kritisk for varetagelsen af dit nuværende job.

[Skrive kommentarer, klummer eller artikler]

- Ikke vigtig
- Lidt vigtig
- Vigtig
- Meget vigtig
- Ekstremt vigtig

Block 2: Adoption

4. Awareness of AI chatbots

AI chatbots bruger kunstig intelligens til at læse og skrive tekst. Vi vil nu spørge ind til dine erfaringer med AI chatbots.

Havde du hørt om følgende chatbots før denne undersøgelse?

Markér alle værktøjer, du havde hørt om før denne undersøgelse.

- ChatGPT (udviklet af OpenAI)
- Claude (udviklet af Anthropic)
- Copilot (udviklet af Microsoft)
- Gemini (udviklet af Google)
- Perplexity (udviklet af Perplexity AI)"
- Andre AI chatbots
- Havde ikke hørt om AI chatbots før denne undersøgelse

5. Prior Use of AI Chatbots [if 4 = 'Ja']

Har du benyttet følgende AI chatbots?

[ChatGPT / Claude / Copilot / Gemini /Perplexity / Andre AI chatbots]

- Ja, kun til arbejde
- Ja, kun til fritid
- Ja, til arbejde og fritid

6.a Purposes of Prior Use [if 5='Ja, kun til fritid' or 'Ja, til arbejde og fritid']

Hvor ofte har du benyttet følgende AI chatbots **til fritid**?

[ChatGPT / Claude / Copilot / Gemini /Perplexity / Andre AI chatbots]

- Aldrig
- Et par gange
- Månedligt
- Ugentligt
- Dagligt

6.b Purposes of Prior Use [if 5='Ja, kun til arbejde or 'Ja, til arbejde og fritid']

Hvor ofte har du benyttet følgende AI chatbots **til arbejde**?

[ChatGPT / Claude / Copilot / Gemini /Perplexity / Andre AI chatbots]

- Aldrig
- Et par gange
- Månedligt
- Ugentligt
- Dagligt

6.c Purposes of Prior Use [if 5='Ja, kun til arbejde or 'Ja, til arbejde og fritid' for any option]
Har du benyttet en AI chatbot til at udføre følgende arbejdsopgaver?

- [Arbejdsopgave 1-6]
- Ingen af ovennævnte

7. Time use on AI chatbots [if 5='Ja, kun til arbejde or 'Ja, til arbejde og fritid' for any option]
Tænk tilbage på de dage, hvor du har brugt AI chatbots til dit arbejde. Hvor meget tid brugte du med AI chatbots i gennemsnit?

- Mindre end 15 minutter per dag
- Mellem 15 minutter og en time per dag
- Mere end en time per dag

8. Paid subscription [if 5='Ja, kun til arbejde or 'Ja, til arbejde og fritid' for 'ChatGPT']
Har du et aktivt Plus-abonnement på ChatGPT?

- Ja
- Nej

Block 3: Employer Initiatives

9. Employer policies
Hvad er din arbejdsgivers politik ift. brugen af AI chatbots?

- Tilladt og tilskyndet
- Tilladt men ikke tilskyndet
- Ikke tilladt
- Ingen politik
- Ved ikke

10. Enterprise AI chatbot
Har din arbejdsplads sin egen AI chatbot?

- Ja, et specialdesignet produkt
- Ja, et standardprodukt
- Nej
- Ved ikke

11.a Training courses
Har du deltaget i kurser om brugen af AI chatbots?

- Ja
- Nej

11.b Training courses [if 11.a = 'Ja']
Var dit kursus i AI chatbots arrangeret af din arbejdsgiver?

- Ja
- Nej

Block 4: Effects of AI chatbots

12. Benefits from AI chatbots

Har du oplevet nogle af disse fordele ved brugen af AI chatbots i dit arbejde?
Markér gerne flere

- Sparet tid i arbejdet
- Forbedret kvalitet af arbejdet
- Øget kreativitet
- Højere arbejdsglæde
- Har ikke oplevet fordele
- Ved ikke

13. Time savings from AI chatbots [if 12='Sparet tid i arbejdet']

Tænk tilbage på de dage, hvor du har brugt AI chatbots til dit arbejde. Hvor meget tid sparede AI chatbots dig i gennemsnit?

- Mindre end 15 minutter per dag
- Mellem 15 minutter og en time per dag
- Mere end en time per dag

14.a Earnings impact of AI chatbots

Har AI chatbots påvirket hvor meget du tjener i dag?

- Jeg tjener mere i dag som følge af AI chatbots
- AI chatbots har ikke påvirket min indtjening
- Jeg tjener mindre i dag som følge af AI chatbots

14.b Earnings impact of AI chatbots [if 14.a=' Jeg tjener mere i dag som følge af AI chatbots']

Hvor meget har AI chatbots øget din indtjening?

- Under 5 procent
- Mellem 5 og 15 procent
- Over 15 procent

14.c Earnings impact of AI chatbots [if 14.a=' Jeg tjener mindre i dag som følge af AI chatbots']

Hvor meget har AI chatbots reduceret din indtjening?

- Under 5 procent
- Mellem 5 og 15 procent
- Over 15 procent

15. Allocation of time savings from AI chatbots

Hvis AI chatbots sparer tid i løsningen af en opgave, forventer du så, at

- Løse flere af samme opgaver
- Bruge mere tid på andre opgaver
- Tage flere pauser
- Tage mere fritid

16. Workloads from AI chatbots

Oplever du, at AI chatbots har øget din arbejdsmængde?

- Ja, flere af de samme arbejdsopgaver
- Ja, nye slags arbejdsopgaver
- Nej

17. New job tasks [if 16='Ja, nye slags arbejdsopgaver']

Hvilke slags nye arbejdsopgaver oplever du at have fået efter brugen af AI chatbots?

- [Fritekstfelt]

17. End of survey

Mange tak for at deltage i undersøgelsen.

Hvis du vinder en af præmierne, vil du få direkte besked i din e-Boks.