Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*

Anders Humlum† Emilie Vestergaard‡

March 13, 2026

Abstract

We study the early labor market impacts of AI chatbots by linking large-scale adoption surveys to administrative labor market records in Denmark. We document rapid currents: most employers in exposed occupations have adopted chatbot initiatives, workers report productivity benefits, and new AI-related tasks are widespread. Yet these currents have not broken the surface: using difference-in-differences, we estimate precise null effects on earnings and recorded hours at both the worker and workplace levels, ruling out effects larger than 2% two years after the launch of ChatGPT. What moves is the structure of work: employers absorb AI through task reorganization—including new tasks in content generation, AI oversight, and AI integration—and adopters transition into higher-paying occupations where AI chatbots are more relevant, though still too few to move average earnings. Technological change reshapes work well before it surfaces in earnings or hours.

*This paper was previously circulated under the title "Large Language Models, Small Labor Market Effects." We thank David Autor, Marianne Bertrand, Jacob Conway, Matthew Notowidigdo, Lindsey Raymond (discussant), Chad Syverson, and seminar participants at the 2025 NBER SI: Digital Economics and Artificial Intelligence, the 2025 SED Meeting: Automation, the Becker Friedman Institute, UC Berkeley, Boston College, Chicago Booth, the Chicago Fed, ETUI, Northwestern Kellogg, OpenAI Economic Research Conference, INSEAD, Microsoft Research, and MIT Sloan FutureTech 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 and Claude Code (Opus 4.6) 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

AI chatbots mark the rise of generative artificial intelligence (AI): these tools have seen the fastest worker take-up of any new technology (Bick, Blandin and Deming, 2025), and controlled experiments demonstrate large productivity gains for users (Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023). Yet even as generative AI tools have become widespread, substantial disagreement remains about their labor market effects, with forecasts ranging from dramatic displacement within a few years (Amodei, 2025) to modest impacts over a decade (Acemoglu, 2025).[1] This disagreement partly reflects how little evidence exists on how labor markets actually absorb new technologies in their early stages. How does work reorganize at the onset of major technological change?

This paper argues that early labor market transformation under generative AI is best understood as *rapid currents under still waters*: substantial reorganization of work occurs as workplaces integrate AI chatbots and workers take on new AI-related tasks, absorbing change before it surfaces in earnings or recorded hours.

To study this process, we develop a unique data infrastructure linking workplace adoption and work reorganization to administrative labor market outcomes. In collaboration with Statistics Denmark, we conducted a series of large-scale, representative surveys of AI chatbot adoption in 11 highly exposed occupations—where labor market effects are likely to emerge earliest and strongest—linking responses to administrative labor market records.[2] Our latest survey round, conducted in late 2024, includes responses from 25,000 workers across 7,000 workplaces. These linked survey–register data allow us to examine both recorded labor market outcomes—such as earnings and hours—and adjustments that typically go unrecorded, including changes in job tasks and work organization. Crucially, even within these highly exposed occupations, many otherwise similar workplaces and workers have taken markedly different adoption paths, allowing us to examine how chatbot

[1] Anthropic CEO Dario Amodei predicts that AI could eliminate half of all entry-level white-collar jobs, driving unemployment to 10–20% within the next one to five years (Amodei, 2025). In contrast, Acemoglu (2025) predicts that AI will increase annual TFP growth by less than 0.07 percentage points over the next ten years.

[2] 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.

adoption relates to work reorganization and labor market outcomes.

We start beneath the surface by examining how workers' use of, and reported benefits from, AI chatbots vary with their employer initiatives. Even in workplaces that neither encourage use nor provide enterprise tools or training, about 40% of workers have used AI chatbots at work, and roughly 7% use them daily. Despite this substantial worker-driven baseline, take-up rates almost double in workplaces with active employer initiatives. Encouraged-use policies, in particular, are associated with greater regular use and higher reported benefits among users, including time savings, quality, creativity, and job satisfaction. Overall, adoption and reported benefits peak in workplaces that combine encouraged use with enterprise chatbots and training: 93% of workers in such settings report having used AI chatbots at work, 28% use them daily, and 19% report saving more than one hour per day—patterns consistent with complementarities between new tools and organizational investments (Bloom, Sadun and Reenen, 2012; Brynjolfsson and Hitt, 2000).

AI chatbots not only affect existing work—they also create new tasks. Even in the absence of employer initiatives, about 8% of chatbot users report taking on entirely new tasks arising from chatbot use, rising to roughly 17% in workplaces with active employer initiatives. These tasks span content *generation* (e.g., drafting and ideation), *oversight* of AI outputs (e.g., quality review and compliance), and *integration* of AI tools into workflows (e.g., developing usage policies). Consistent with this reorganization of work, most chatbot users (85%) report reallocating time savings from AI chatbots to other job tasks. New tasks extend even to workers who have never used AI chatbots, signaling broader workplace transformation. These patterns support task-based theories of how new technologies reinstate labor demand (Acemoglu and Restrepo, 2018*a*; Autor et al., 2024).

We then rise to the surface, asking whether and how these rapid currents of chatbot adoption and work reorganization register in earnings and recorded hours. Our analysis boils down to a simple question: have chatbot adopters fared differently from comparable non-adopters? To address this question, we link our surveys to administrative records

on monthly earnings, hours worked, and occupations through December 2024—two 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, estimating impacts at both the worker and workplace levels.

We first analyze workers who use AI chatbots. If chatbot use enhances individual productivity, a natural question is whether and how such gains affect their earnings (Krueger, 1993). The answer is theoretically ambiguous, depending on whether adoption is worker- or firm-driven (Becker, 1964), whether it displaces users from higher-value tasks (Autor and Thompson, 2025), whether chatbot use is valued as a nonpecuniary amenity (Rosen, 1986), and on the structure of wage setting (Card et al., 2018).

We find precisely null effects of AI chatbot use on earnings and hours. Difference-in-differences estimates for earnings, recorded hours, and wages all center on zero, with confidence intervals ruling out average effects larger than 2%. These null results hold even among workers who use chatbots daily, report substantial productivity gains, perform new AI-related tasks, or work in firms that encourage and invest in chatbots. Moreover, we find no significant effects in any of our 11 occupations, including those with flexible, decentralized wage-setting.

While chatbots have not led to earnings gains for users on average, we find a clear link to occupational mobility: since the arrival of AI chatbots, adopters are significantly more likely to have switched occupations, working about 4% of a full-time equivalent more in their most recent occupation by December 2024. These transitions are most pronounced into IT support specialist and clerical (office support) roles—occupations in which workers have greater discretion over their toolsets and that do not require formal educational credentials. These patterns align with experimental evidence that chatbots are especially beneficial for workers with less prior expertise (Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023) and support the view that generative AI helps workers move into new occupations by granting access to otherwise scarce expertise (Autor, 2024). In particular, adopters who switch occupations tend to move into roles with higher wage premia and

greater AI chatbot relevance. On average, these workers see earnings grow 12 percentage points faster than other employed workers. While these switchers are still too few to drive average adopter earnings, this pattern suggests that occupational mobility may be a channel through which AI chatbots deliver broader earnings effects over time (Kleven et al., 2025).

Next, we zoom out to the workplace level and examine whether employers that encourage and invest in AI chatbots have experienced differential changes in labor demand. If chatbots enhance labor productivity, firms may expand hiring; if they substitute for labor, they may slow hiring or trigger layoffs (Acemoglu and Restrepo, 2019). Recent declines in early-career employment in AI-exposed occupations have raised concerns that generative AI is already displacing entry-level jobs (Brynjolfsson, Chandar and Chen, 2025). Crucially, our data allow us to split these employment trends by whether firms have actually adopted generative AI. While we replicate the pattern of declining early-career employment in exposed occupations in Denmark, our difference-in-differences analysis reveals that the declines are not driven by firms adopting AI chatbots. Our estimates are sufficiently precise to rule out even modest changes. Workplaces that encourage chatbot use exhibit no differential changes in employment or wage bills, job creation or destruction, or the composition of hires or separations, including among early-career workers. These null results hold across occupations and persist even when encouragement is paired with enterprise chatbot solutions and training.

Taken together, our findings show that early labor market transformation under generative AI is occurring beneath the surface. Even as wages and recorded hours remain stable, workplaces adopt chatbot initiatives, workers take on new AI-related tasks, and adopters transition into new occupations.

Contributions to the Literature. This paper makes two linked contributions. The first is conceptual: we show how labor markets adjust at the onset of major technological change, arguing that substantial reorganization of work absorbs change before it surfaces in economic statistics. The second is measurement: we develop a unique data infrastructure

linking large-scale, representative data on AI chatbot adoption and work reorganization to administrative labor market records.

A contemporaneous set of studies uses exposure measures—technical assessments of which job tasks are feasible to assist with AI chatbots—to examine whether exposed occupations have fared differently after the release of ChatGPT. Brynjolfsson, Chandar and Chen (2025) document a sharp decline in employment of early-career workers in AI-exposed occupations in the United States since 2022, suggesting that generative AI may already be displacing entry-level jobs. Our key contribution is to split these trends by whether firms have actually adopted generative AI, revealing that the declines are not driven by firms adopting AI chatbots.

A small but growing number of studies field worker-level surveys on the adoption of generative AI (Bick, Blandin and Deming, 2025; Hartley et al., 2025; Humlum and Vestergaard, 2025). While our adoption patterns and self-reported time savings align with these studies, we link these survey measures to administrative labor market outcomes, including earnings and hours at the worker and workplace levels, allowing us to move the analysis beyond self-reported effects.[3]

While several experiments document sizable productivity gains from chatbot use on specific tasks (Brynjolfsson, Li and Raymond, 2025; Dillon et al., 2025; Noy and Zhang, 2023), it remains unclear how these gains translate into workers' labor market outcomes, as high-quality microdata are rarely available. While AI chatbot use has not measurably affected workers' earnings and hours, we find a clear association between adoption and occupational mobility. This aligns with experimental evidence that chatbots provide greater benefits to workers with less prior expertise. Moreover, adopters who switch move into occupations with higher wage premia and where using AI chatbots is more relevant, suggesting that chatbots help workers access otherwise scarce expertise (Autor, 2024; Autor and Thompson, 2025).

[3]AI labs have recently released descriptive statistics from their internal usage data (Chatterji et al., 2025; Handa et al., 2025; Tomlinson et al., 2025). Our survey data complement these by capturing the full range of chatbot products, measuring complementary workplace investments, linking to administrative records, and sampling both users and non-users from a well-defined population frame.

We provide the first representative account of employer adoption of AI chatbots. Whereas surveys on traditional "pre-generative" AI show relatively low uptake (Acemoglu et al., 2022; Bonney et al., 2024), we find that most employers in exposed occupations have now embraced AI chatbots. Our findings align with Yotzov et al. (2026), who document widespread firm-level adoption of generative AI but limited perceived impacts on employment or revenue productivity. Alternative firm-level data proxy adoption through generative AI skills in job postings (Lichtinger and Hosseini Maasoum, 2025; Schubert, 2025), but it is unclear how well these capture actual tool use, especially for non-specialized tools like chatbots. We address this gap by surveying workers on their chatbot use and employer initiatives, linking responses to administrative firm data. Our survey also allows us to measure complementary workplace investments that otherwise would go unrecorded (Brynjolfsson, Rock and Syverson, 2021).

While employer chatbot initiatives have not led to net changes in earnings or hours, they strongly predict the emergence of new AI-related job tasks (Acemoglu and Restrepo, 2018*a*; Autor et al., 2024). This points to the mechanism behind the still-waters pattern: employers are absorbing generative AI through task reorganization rather than adjusting on the margin of total labor input. This is consistent with recent work on "job transformation," which emphasizes that new technologies alter the task content of jobs well before they displace them (Autor and Thompson, 2025; Freund and Mann, 2025).

Finally, our account of "still waters over rapid currents" echoes Robert Solow's famous observation about the IT revolution: "You can see the computer age everywhere but in the productivity statistics" (Solow, 1987)—and is consistent with the "productivity J-curves" observed at the onset of general-purpose technologies (Brynjolfsson, Rock and Syverson, 2021; McElheran et al., 2025). We document the hallmarks of such a trough: widespread adoption accompanied by substantial employer investments in training, integration, and oversight that have yet to translate into hard economic outcomes. Our unique data infrastructure reveals what lies beneath the still waters: the reorganization of work that absorbs technological change before it surfaces in hours or earnings.

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—similar to those of the U.S.—which allow firms and workers to adjust employment in response to technological change (Botero et al., 2004; Dahl, Le Maire and Munch, 2013). Appendix C.1 details the wage-setting systems in our study occupations, showing that most workers engage in annual negotiations with their employers, providing regular opportunities to adjust earnings and hours in response to AI chatbot adoption.

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 and work reorganization.

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 U.S.—and an unparalleled data infrastructure linking workplace adjustments to administrative labor market outcomes.

1.1 Survey Data

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, hours worked, 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 employer-led adoption initiatives, and (*iii*) documents changes in work organization alongside workers' actual usage and perceived benefits of these tools. Appendix I outlines the 2024 survey, which is the primary focus of this paper.

1.1.1 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 comprise the segments of the labor market where the impacts of AI chatbots are likely to emerge first and most strongly. Our analysis focuses on this set.

1.2 Register Data

We use several administrative registers at Statistics Denmark. Our matched employer-employee data come from the *Employment Statistics of Employees* (BFL), 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. We complement this with demographics data on individuals from the *Population Register* (BEF), and wealth information from 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.

1.3 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 each occupation and then sampling all relevant workers at these workplaces. This sampling procedure maximizes the statistical power of the workplace-level analyses while keeping our sample representative. Appendix B.1 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 survey. Appendix B.2 details our survey response rates. While our main analysis focuses on responses from the 2024 round, we use the 2023 round to examine early adopters.

1.3.1 Representativeness and Response Quality

Appendix B.2.1 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 The Currents: Adoption and Work Reorganization

In this section, we study how AI chatbots are integrated into the workplace. Section 2.1 describes employers' chatbot initiatives, showing that even as chatbots have become widespread, otherwise similar firms often adopt markedly different approaches. In Section 2.2, we examine how these strategies reflect in workers' chatbot use and its reported benefits. Finally, in Section 2.3, we study how AI chatbots reshape the organization of work, including the creation of new job tasks. Together, these results document the rapid currents of adoption and work reorganization set in motion by the arrival of AI chatbots.

2.1 Employer Initiatives

Figure 1, Panel (a), shows the prevalence of employer usage policies for AI chatbots across our 11 occupations. Firms now broadly embrace AI chatbots: about 43% explicitly encourage their use, another 21% allow it, while only about 6% explicitly prohibit it.[4] This widespread adoption of AI chatbots in the workplace is consistent with their general-purpose features (Eloundou et al., 2024).

Employers not only set usage policies but also invest directly in chatbot adoption. Panel (b) shows the chatbot packages adopted by different employers. Among firms that encourage chatbot use, 61% have an enterprise chatbot solution, 39% provide training, and 29% offer both. Notably, enterprise chatbots and training are present even in some workplaces that do not encourage use.

Which firms have adopted different AI chatbot initiatives? Table E.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). Yotzov et al.

[4]The widespread embrace of AI chatbots marks a shift from early responses to ChatGPT, when many employers restricted its use due to concerns about data confidentiality and output accuracy (Humlum and Vestergaard, 2025). Although outright bans are now rare, they persist in some occupations involving sensitive data (e.g., financial advisors) or requiring high factual accuracy (e.g., legal professionals).

(2026) find similar adoption patterns for generative AI across the US, UK, Germany, and Australia.[5] While firm size is unrelated to encouraged-use policies—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.

2.1.1 Employer Initiatives and Worker Adoption

We next examine how workers' adoption of and benefits from AI chatbots vary with the employer-led chatbot initiatives. 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 of or reported benefits from AI chatbots; X_i includes survey occupation fixed effects and worker pre-determined characteristics, including age, gender, and potential experience;[6] and $\text{EmployerInitiatives}_i$ 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 at the workplace-occupation level.

The specification in Equation (1) is motivated by a large literature on complementarities between new technologies and organizational investments (Bloom, Sadun and Reenen, 2012;

[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.

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

Brynjolfsson and Hitt, 2000), which emphasizes that the returns to technology adoption may depend critically on firms' broader organizational investments and management practices. Moreover, because early adoption of AI chatbots was largely worker-driven, this analysis sheds light on how adoption and work-related benefits differ in workplaces where employers also invest in the technology.

Appendix A.1 presents a Roy-style model of chatbot adoption to help interpret the estimates in Equation (1). The model shows that β for user-reported benefits combines an *individual*-level effect of the initiative on each user's gains with a *selection* effect arising from shifts in which workers opt into use. When workers adopt based on their individual gains, selection will tend to be negative: initiatives that expand adoption draw in workers who benefit less from the technology. We do not attempt to separate these two effects, as we view the bundled impact of employer initiatives as informative in itself.

While we interpret Equation (1) as descriptive, our rich firm- and worker-level data allow us to assess the robustness of these patterns to a range of potential confounders. Appendix F.5 presents these robustness checks. First, we show that all results hold when controlling for firm characteristics, 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] In addition, the panel analyses in Section 3 confirm that adopters of different 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.

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.

[7]We do not include task controls in our main specification, as task mix may itself be endogenous to chatbot adoption. Indeed, Section 2.3 shows that AI chatbots have created new job tasks for many workers.

Nonetheless, coworker reports are tightly correlated within workplaces, and Appendix F.6 shows that our findings are robust to using coworker reports to measure employer initiatives, supporting their interpretation as workplace-wide policies.

2.2 Worker Adoption and Reported Benefits

Figure 2 shows how AI chatbot adoption varies with employer initiatives for a typical worker. The top estimate (*No Initiative*) highlights the bottom-up nature of chatbot diffusion: 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 least once, and 7% use them daily.

Despite this high baseline, employer initiatives are associated with substantially higher adoption. Implemented individually, *Encouraged use*, *Enterprise chatbots*, and *Training* each raise extensive-margin take-up rates to 76%, 57%, and 69%, respectively. Encouraged-use policies are associated with especially high intensive usage, with daily use reaching 18%.

Table 1 shows that encouraged-use policies are not only associated with higher chatbot usage (Column 1-2), they are also consistently associated with users reporting greater time savings, improved quality, and enhanced creativity (Columns 3-6).[8]

Finally, as the bottom row of Figure 2 shows, 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.

Discussion. The finding that adoption and reported benefits are greatest when AI chatbots are combined with employer initiatives aligns with a large literature on complementarities between new technologies and organizational investments, including people management practices (Bloom, Sadun and Reenen, 2012; Brynjolfsson and Hitt, 2000).

[8]By contrast, training and enterprise tools—when implemented without encouragement—are associated with smaller reported gains. This pattern is consistent with these initiatives being aimed at mitigating misuse rather than enhancing productivity. Supporting this interpretation, Section 2.3 shows that *AI Ethics & Compliance* tasks are especially prevalent among workers who have received employer-provided training.

Still, the relative importance of encouraged-use policies may seem surprising, as such policies might appear inexpensive in monetary terms. But the low capital costs of chatbot solutions (Appendix E.2 details adoption by chatbot products) heighten the importance of complementary organizational practices, such as explicitly encouraging use. For example, Shopify has publicly reframed generative AI from something employees may use into a baseline expectation; Chicago Public Schools explicitly encourages faculty to engage with GenAI tools; and Denmark's Agency for Digital Government similarly emphasizes that adoption requires clear usage guidelines and supportive organizational frameworks even when publicly available tools are used (Chicago Public Schools, 2024; Digitaliseringsstyrelsen, 2024; Lütke, 2025). Consistent with this interpretation, Dillon et al. (2025) show that firm effects are a key determinant of workers' uptake of a randomized offer of Microsoft Copilot.

2.3 New Workloads and Task Creation

AI chatbots not only affect existing tasks—they transform the nature of work itself. Figure 3 shows how the emergence of new workloads from AI chatbots varies across employer initiatives. These new workloads may stem from individual workers taking on more work, from shifts in task allocation within teams, or from new demands generated by the tools themselves, such as oversight and integration.

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, about 5 percentage points report doing more of the same tasks, and 8 percentage points have taken on entirely new tasks.[9]

Notably, the share of users performing new tasks in Figure 3 increases to roughly 18% 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

[9]Consistent with this reorganization of work, most chatbot users (85%) report reallocating time savings from AI chatbots to other job tasks (Table E.3). By contrast, fewer than 10% report taking additional breaks or leisure, and about 30% spend more time on the same tasks they initially saved time on.

workers handle new responsibilities caused by AI chatbots.

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 B.3, we 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 53% to 91% of new tasks directly linked to AI use.

Figure 4 breaks down the composition of AI-related job tasks by occupation and task category. Appendix B.3 provides example tasks for each pair. About 42% of new AI-related tasks relate to generating new content, including *AI Ideation*, *AI Content Drafting*, and *AI Data Insights*. For example, marketing professionals report new tasks related to "prompting and iterating with AI to produce marketing copy, social media posts, and product descriptions." Notably, 35% of new AI-related work is dedicated to *AI Quality Review* and *AI Ethics & Compliance*, suggesting that AI adoption shifts workers not only toward content generation but also toward supervisory roles overseeing AI outputs.[10] For example, a large share of teachers report new work related to "detecting AI-generated homework." Finally, the single most common new AI task category is *AI Integration*—that is, the integration of AI chatbots into the workplace, which accounts for about 26% of all new tasks, with the highest share observed in IT support and software development. For example, software developers report "fine-tuning AI coding assistants by providing feedback and project-specific examples," and legal professionals describe "developing organizational AI usage policies and guidelines."

AI chatbot workloads extend beyond direct users: Figure 3, Panel (b) shows that 4% of non-users report new workloads resulting from AI chatbots, with effects most pronounced in workplaces that have adopted chatbot initiatives. In particular, Figure B.1, Panel (b) shows that 10% of teachers who have not used AI chatbots report new AI-related

[10]Appendix Figure E.3 shows that *AI Ethics & Compliance* tasks are especially common among workers who have received employer-provided AI chatbot training, helping to contextualize the finding in Section 2.2 that training—when implemented in isolation—is associated with lower productivity gains from AI chatbots. In contrast, workplaces that only encourage chatbot use exhibit a higher share of "generative tasks" focused on *AI Ideation* and *AI Content Drafting*.

workloads, largely consisting of monitoring and responding to students' AI use.

Discussion. The widespread creation of new tasks and spillovers to non-users highlight the broader workplace transformations caused by AI chatbots. These patterns support task-based frameworks of technological change, in which automation may reinstate labor demand (Acemoglu and Restrepo, 2018*a*; Autor et al., 2024), including models of "job transformation" where technologies reshape the content of jobs well before they displace them (Autor and Thompson, 2025; Freund and Mann, 2025). Workers shift not only toward content generation but also toward supervisory roles overseeing AI outputs—a pattern consistent with AI chatbots altering the hierarchical structure of production, with humans increasingly serving as monitors and validators of AI-generated work (Ide and Talamas, 2025). Finally, the high prevalence of integration tasks likely reflects an early phase of major technological change, with workers and firms still in the process of embedding AI into their workflows. These patterns could suggest that some organizations remain in the trough of a productivity J-curve (Brynjolfsson, Rock and Syverson, 2021; McElheran et al., 2025). On the other hand, as shown in Section 2.2, a substantial share of workers already report large time savings and quality improvements from using AI chatbots—especially when employers encourage and invest in the tools. In the next section, we assess whether and how these workplace transformations surface in administrative labor market outcomes.

3 The Surface: Earnings and Employment

Our preceding analysis shows that AI chatbots have generated strong currents within workplaces: employers are adopting chatbot initiatives, workers report substantial work-related benefits, and new AI-related tasks are emerging broadly. Do these currents surface in "hard" labor market outcomes such as earnings or hours? To answer this question, we link our survey responses to administrative labor market records.

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

Motivated by the dual nature of adoption—with variation across workplaces depending on their chatbot initiatives and within workplaces depending on individual use—we analyze impacts at both the worker and workplace levels.

Appendix A.2 presents a theoretical framework for how adoption translates into labor market outcomes at the worker and workplace levels, guiding the following analyses:

1. *Worker earnings:* If chatbots raise individual productivity, a natural question is whether and how these improvements translate into earnings for users relative to comparable non-users (Krueger, 1993). Wage effects depend on whether adoption is worker- or firm-driven, whether chatbot use displaces users from higher-value tasks or is valued as a nonpecuniary amenity, and the structure of wage setting (Autor and Thompson, 2025; Becker, 1964; Card et al., 2018; Rosen, 1986). Section 3.2 investigates these questions by estimating the impacts of chatbot use on workers' earnings by their intensity of use, employer initiatives, reported productivity gains, new AI-related workloads, and wage-setting systems.

2. *Workplace employment:* If chatbots reduce the need for labor, adopting workplaces might slow hiring or lay off workers. Conversely, if the tools boost productivity and spur demand, adopting workplaces could expand employment. In particular, concerns have been raised that generative AI could displace entry-level positions (Amodei, 2025; Brynjolfsson, Chandar and Chen, 2025). Section 3.4 analyzes these possibilities by estimating the effects of employer chatbot initiatives on workplace employment, early-career jobs, job churn, and incumbent worker outcomes.

3. *Job mobility:* If chatbots are especially useful for newcomers to an occupation,

adopters may be more likely to switch occupations (Autor, 2024). Section 3.3 tests this prediction, exploring heterogeneity across occupations and firm- and worker-level drivers of adoption.

3.1.1 Identification Strategy and Challenges

We use a series of difference-in-differences analyses to identify the relevant effects, comparing adopters and non-adopters before and after the introduction of AI chatbots and examining outcomes at both the worker and workplace levels. At its core, our analysis asks a simple question: have adopters fared differently from non-adopters since the arrival of chatbots?

The key identifying assumption is that, absent AI chatbots, adopters and non-adopters would have experienced similar changes in labor market outcomes. We apply this assumption at both the worker level (e.g., earnings of users versus non-users) and the workplace level (e.g., labor demand at encouraged versus non-encouraged workplaces). Appendix A.2 draws on our theoretical framework to clarify these identifying assumptions and to provide a structural interpretation of the difference-in-differences estimands.

A first concern in our analysis is that adoption may remain a relatively weak treatment: workers may struggle to use the tools effectively, and firms may provide limited guidance or infrastructure to support their implementation. Section 2.2 showed that employer encouragement is critical for unlocking both higher use and greater reported benefits. We use these "AI front-runner" workplaces—those with encouraged use—as our primary workplace treatment, since 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. At the individual level, we also analyze heterogeneity by intensity of use and reported gains.

A second concern is that adopters may have fared differently in the labor market even without AI chatbots. We take several 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 the launch of ChatGPT. Crucially, the event-study design allows us to assess whether adopters were on distinct labor market trends before AI chatbots became available. Finally, we verify that our conclusions about the effects of worker-level use are robust when we rely only on the variation in chatbot use that stems from employer policies. Employer initiatives may be more plausibly unrelated to individual worker trajectories, as they reflect firm-level choices rather than worker self-selection into chatbot use.

A third challenge 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.[11] We also separately examine effects for early adopters, identified as workers already using chatbots in our 2023 survey round, to allow more time for effects to materialize.[12] Even with rapid adoption, however, 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 vary in the flexibility of labor rules.

A final concern is that difference-in-differences identifies only *differential* effects for

[11]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 E.2 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., Brynjolfsson, Chandar and Chen (2025); Eisfeldt et al. (2024); Lichtinger and Hosseini Maasoum (2025); Schubert (2025); Teutloff et al. (2025)), facilitating comparison across studies.

[12]Our survey does not record the exact timing of adoption, precluding event studies indexed to workers' month of first use. However, even if such information were available, adoption is unlikely to represent a sharp before-and-after event, as workers typically ramp up usage gradually. This is another reason why we expect difference-in-differences effects to emerge over time rather than immediately.

adopters, leaving any level effect on non-adopters unidentified—the "missing intercept." For example, when we compare earnings relative to non-encouraged non-adopters in Section 3.2, any general equilibrium effects that also affect these workers are differenced out. Indeed, Section 2.3 showed that AI chatbots have created new job tasks even for non-adopters. The key question for our analysis below is whether such spillovers also surface in earnings or recorded hours. For instance, adopting firms might expand by poaching workers from non-adopting competitors, biasing our workplace employment estimates upward. Alternatively, if firms adjust labor demand in anticipation of AI chatbots—rather than based on their actual adoption—the estimates would be biased toward zero. We assess the potential magnitude of these spillover effects from several complementary angles.

First, we assess potential spillovers at different levels of aggregation by estimating difference-in-differences models at the worker, workplace, and occupation levels. For instance, the fact that encouraged workplaces do not experience differential changes in labor demand in Section 3.4 suggests that spillovers in employment outcomes within workplaces are limited. Consistent with this interpretation, Figure F.3 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 D.1 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 before and after the arrival of AI chatbots.

Second, Appendix A.2.6 uses a market equilibrium model to show that our difference-in-differences estimates are informative about the magnitude of the "missing intercept." In particular, under reasonable labor demand elasticities, null effects at both the workplace and worker levels imply that the occupation-wide equilibrium effect is also negligible. The key intuition is that adopters generate spillovers to non-adopters only when their own market outcomes change. Thus, the absence of adoption effects at both levels implies that spillovers to non-adopters are likewise limited.

Finally, as discussed in Section 3.5, 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 spillover impacts of AI chatbots remain minimal.

Taken together, these different analyses—difference-in-differences at varying levels of aggregation, equilibrium calibration of micro-level estimates, and perceived impacts among workers—all point to the conclusion that AI chatbots have not had any meaningful impact on non-adopters' earnings or employment thus far.

3.1.2 Regression Specifications

Sampling Design and Data. Our data and sampling are designed to answer the questions described above. By linking survey responses from workers in the most highly exposed occupations to comprehensive matched employer–employee panel data covering the entire Danish economy, we can study labor market outcomes before and after the arrival of AI chatbots. Crucially, the matched data allow us to track not only the surveyed workers themselves but also other workers at the same workplaces, including those who have since left.

Our survey sampling design also imposes restrictions that we explicitly account for in the analysis. In particular, because we survey workers in exposed occupations, all respondents are, by construction, employed in those occupations at the time of the survey sampling. Accordingly, all specifications include survey-occupation $\times$ time fixed effects, ensuring that all comparisons are made within survey strata and time period, between workers (or workplaces) that differ in their adoption of AI chatbots. We also study heterogeneity in effects by survey occupation.

Let Y_{it} denote a labor market outcome (e.g., earnings) for worker i in month-year t, let X_i denote survey-occupation fixed effects and worker pre-determined characteristics—age, gender, and potential experience in 2022—and let A_i indicate whether the worker has adopted AI chatbots. We examine both worker-level effects (e.g., worker earnings by

individual usage) and workplace-level effects (e.g., workplace employment by employer initiatives).

Dynamic Difference-in-Differences. To assess how impacts evolve over time, we employ a dynamic difference-in-differences specification that traces the monthly evolution of outcomes for adopters relative to non-adopters around the introduction of AI chatbots:

$$Y_{it} = \underbrace{\sum_{\tau} \lambda_{1\tau} X_i \mathbf{1}_{\{t=\tau\}}}_{\text{Controls}} + \underbrace{\sum_{\tau \neq 2022M11} \beta_\tau A_i \mathbf{1}_{\{t=\tau\}}}_{\text{Dynamic Diff-in-Diffs}} + \alpha_i + \varepsilon_{it}, \tag{2}$$

where α_i are individual (workplace) fixed effects. The parameters of interest, β_τ, capture the differential changes in labor market outcomes for adopters, indexed to November 2022, the release date of ChatGPT. We estimate Equation (2) by OLS and cluster the standard errors at the survey workplace-occupation level.

Pooled Difference-in-Differences. To examine heterogeneous impacts across multiple dimensions, we also estimate a pooled specification comparing average outcomes of adopters and non-adopters before and after the introduction of AI chatbots. To capture dynamics, we estimate separate effects for each year following the launch of ChatGPT:[13]

$$Y_{it} = \underbrace{\sum_{\tau} \lambda_{1\tau} X_i \mathbf{1}_{\{t=\tau\}}}_{\text{Controls}} + \underbrace{\lambda_2 A_i t}_{\text{Time Trends}} + \underbrace{\sum_{y=2023}^{2024} \beta_y A_i \mathbf{1}_{\{y(t)=y\}}}_{\text{Pooled Diff-in-Diffs}} + \alpha_i + \varepsilon_{it}, \tag{3}$$

where adopter-specific time trends guard against spurious pre-trends driving the estimates (which the dynamic specification instead reveals nonparametrically). The parameters of interest, β_y, measure how labor market outcomes for adopters changed in year y after the introduction of AI chatbots. We estimate Equation (3) by OLS and cluster the standard errors at the survey workplace-occupation level.

[13] 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.

3.2 Worker Earnings

Figure 5, Panels (a) and (b) examine earnings of chatbot adopters, distinguishing between those with and without employer encouragement (i.e., encouraged-use policies). Each group is compared to non-encouraged non-adopters—workers insulated from AI chatbots through both personal non-use and the absence of employer encouragement.

Panel (a) begins by presenting differences in means, controlling for workers' survey occupation and pre-determined characteristics. 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. Indexing to November 2022, the difference-in-differences estimates in Panel (b) show no differential changes for either group of adopters following the introduction of AI chatbots. Appendix F.2 shows that these null results hold separately for workers' hourly wages, intensive-margin hours, and extensive-margin employment. Moreover, Figure F.4 shows that the null results on worker earnings hold if we rely only on the variation in chatbot use that stems from employer policies—variation that is arguably more plausibly unrelated to individuals' earnings trajectories.

Figure 5, Panel (b) provides three key insights. First, two years after their launch, AI chatbots have had minimal earnings impact on workers who use them. The confidence intervals of our dynamic difference-in-differences estimates rule out changes larger than 3%, and our pooled difference-in-differences estimates rule out effects larger than 2%.[14] These effects are also small relative to the substantial variation in earnings changes observed for similar workers over time: the standard deviation of residual earnings changes is about 30 percentage points in December 2024. Second, these effects show no differential trends after the introduction of AI chatbots, indicating that the minimal impacts are not

[14]The dynamic difference-in-differences in Figure 5, Panel (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 7, Panel (a) (which control for pre-trends; see Equation (3)) are precise zeros, with confidence bands ruling out effects larger than 2%.

merely a short-term phenomenon. Third, despite evidence in Section 2.2 that employer encouragement boosts work-related benefits of AI chatbots, we find no differential changes in earnings for workers whose employers encourage their chatbot use.

Heterogeneity by Intensity of Adoption. Figure 6, Panel (a) examines the earnings impacts of chatbot use by the intensity of workers' adoption.

We first consider the effects of additional investments in employer-provided chatbots and training, offering insights into the impacts of AI chatbot use when firms pursue a more comprehensive strategy. Yet even when the "full package" of encouraged use, enterprise chatbots, and training substantially increases reported benefits (Section 2.2), we find no detectable impact on adopter earnings.

Next, we consider whether the null effects reflect the limited intensity of individual use. In particular, even if most adopters do not use the tools daily, a meaningful fraction of workers report more substantial benefits—such as daily time savings exceeding one hour. Turning to workers who use AI chatbots daily or who adopted within the first year after ChatGPT's launch, we find no differential earnings changes since the arrival of AI chatbots, despite their substantial engagement with the tools. Similarly, among workers who report greater benefits from chatbot use—time savings exceeding one hour per day, higher-quality output, and enhanced creativity—we find no differential earnings changes in the administrative data. Taken together, although "super users" (those with the most intensive adoption or highest reported benefits) certainly appear in our data, even their more intensive engagement with AI chatbots has not translated into detectable earnings gains.

Finally, we consider workers who report new workloads from AI chatbots—and new AI-related tasks in particular. Even if chatbot use does not directly affect adopter earnings, it could influence them indirectly through the creation of new work. Yet neither workers with new workloads nor those with new AI-related tasks have experienced differential earnings changes since the arrival of AI chatbots.

Occupational Heterogeneity. The overall zero effects could mask important heterogeneity across occupations. For instance, since teachers are covered by collective bargaining with centralized wage setting, it may be less surprising that individual productivity gains do not translate into higher earnings. By contrast, occupations such as marketing and software development feature decentralized wage setting, allowing more flexibility to adjust pay based on individual productivity.[15] Appendix C.1 provides an overview of the wage-setting systems in Denmark, showing rich variation across our study occupations. Occupational heterogeneity may also arise from whether chatbots are used primarily for core versus peripheral tasks within an occupation (Autor and Thompson, 2025).

Figure 7, Panel (a) shows the estimated effects of AI chatbots on worker earnings across the 11 occupations in our sample. Confidence bands are wider for these occupation-specific estimates and should therefore be interpreted with caution. Nonetheless, we find no statistically significant effects in any occupation, and point estimates are generally near zero. Even for our most contrasting groups—software developers and marketers, with more flexible wages and larger reported time savings, versus teachers, with centralized wage-setting and smaller reported savings—we find no systematic differences in earnings impacts.

Discussion. Our administrative data show that users of AI chatbots have not experienced differential changes in earnings or hours following the arrival of these tools. Even among workers who use chatbots daily and report substantial benefits, we find no net changes—not even a break in trends. These patterns echo those documented for computer use in the 1980s by Krueger (1993) and critiqued by DiNardo and Pischke (1997): although chatbot use is correlated with sizable cross-sectional earnings premia, these differentials vanish once accounting for positive selection.

Are the null wage effects of chatbot use theoretically surprising? Competitive labor markets (e.g., our framework in Appendix A.2) imply that workers should be compensated

[15] Furthermore, marketing professionals and software developers are exactly the occupations with the highest reported use and productivity gains from the tools (Appendix E.3).

for productivity gains that are not driven by employer initiatives (Acemoglu and Pischke, 1998; Becker, 1964). From this perspective, it is surprising that individual-specific benefits from AI chatbots—such as those among non-encouraged or high-use, high-benefit adopters—do not translate into higher earnings. On the other hand, if employers do not endorse or invest in the tools, they may be reluctant to reward their use. From this perspective, it is surprising that adopters supported by employer initiatives have also not experienced earnings gains. One possibility is that workers value chatbot access as a nonpecuniary amenity, accepting unchanged wages in exchange for easier or more engaging work (Mas, 2025). The null effects also persist across occupations with very different wage-setting institutions—from centralized collective bargaining for teachers to decentralized individual negotiation for software developers and marketers—suggesting that wage rigidity alone cannot explain the pattern. Indeed, the standard deviation of residual earnings changes in Equation (2) is about 30 percentage points in December 2024.

Another possibility is that workers overestimate the benefits they derive from these tools (Becker et al., 2025; Edelman, Ngwe and Peng, 2023).[16] A key contribution of our study is to link these self-reports to third-party administrative data, allowing us to move beyond purely self-reported effects. That said, inflated self-reports are unlikely to be the full explanation: Section 3.5 shows that even workers themselves recognize that AI chatbots have not materially affected their earnings.

3.3 Occupational Mobility

Figure 5, Panel (c) shows that adopters have experienced greater occupational mobility since the arrival of AI chatbots. We measure occupational mobility by tracking hours worked in each worker's most recent occupation as of December 2024. If adopters are more likely to switch into new occupations, their hours in that new occupation will rise relative to non-adopters, who are more likely to remain in their original occupation. Because the figure reports difference-in-differences estimates indexed to November 2022, the estimates

[16]Consistent with experimental evidence from Becker et al. (2025), Appendix E.4 shows that workers who report greater time savings from chatbots are also more likely to report taking on new tasks.

capture these changes relative to the launch of ChatGPT. By December 2024, adopters have increased their hours in their most recent occupation by about 4% of a full-time equivalent relative to non-adopters. The figure provides four insights.

First, the figure suggests that chatbots appeal to workers who are new to their occupations. This relationship may reflect that chatbots facilitate occupational transitions—for instance, by granting access to otherwise scarce expertise—or that workers who have recently switched find chatbots especially valuable as they adapt to new roles.[17]

Second, the effects on occupational mobility grow steadily, likely reflecting both the deepening of adoption over time and the emergence of new labor market opportunities. As we show below, these transitions tend to be into occupations with higher wage premia and greater AI chatbot relevance, suggesting that this mobility channel may gain importance over time.

Third, the flat pre-trend prior to the arrival of chatbots suggests that this association 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 Panel (a), which were entirely attributable to pre-existing differences between adopters and non-adopters.

Finally, the significant estimates provide reassurance that our empirical specification has the power to detect effects when they exist—a useful contrast to the null results for earnings and employment.

Heterogeneity by Intensity of Adoption. Figure 6, Panel (b) examines how the relationship between chatbot adoption and occupational mobility varies with the intensity of adoption. The figure shows that the association between chatbot adoption and occupational switching is driven primarily by individual-level factors rather than firm-level initiatives. For example, our main estimate on occupational switching rises from 1.3%

[17]In contrast, Appendix F.3 shows that the association between chatbot adoption and workplace mobility is weaker, and occupational mobility is no stronger at workplaces that encourage chatbot use. These patterns align with Figure 6, Panel (b), which indicates that the relationship between chatbot adoption and occupational switching is driven primarily by individual rather than firm-level factors.

to 4.1% for individuals who use the tools daily, and to 4.0% for those who report large time savings.[18] By contrast, the relationship between adoption and occupation switching is fairly similar across employer chatbot initiatives, ranging in 2024 from about 1.1% to 1.9% across initiatives.

Occupational Heterogeneity. Figure 7, Panel (b), splits the effects of chatbot adoption on mobility by survey occupation.[19] Positive effects appear in multiple occupations, but are strongest and most statistically significant among IT supporters, and to some extent, office clerks. These are occupations where individual workers likely have more flexibility in choosing their toolset, allowing newcomers to take advantage of AI chatbots. 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 on an individual basis. 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, the absence of switching effects in licensed occupations—such as teachers and accountants, which require several years of prior education—strengthens the interpretation that the mobility effects observed elsewhere reflect a genuine link between chatbot use and newcomers acquiring occupation-relevant skills.

Job Characteristics for Adopters Who Switch Occupations. Adopters who switch occupations are particularly informative for two reasons. First, occupational switching is the only administrative outcome that moves with workers' chatbot adoption. Second, since job switching is a core source of labor market progression, the characteristics of these

[18]The effect magnitudes in Figure 6, Panel (b) are generally smaller than the endline estimates in Figure 5, Panel (c). This reflects that the pooled specifications estimate trends over the full sample period, absorbing part of the post-period treatment variation (see Equation (3)).

[19]Because survey respondents, by construction, are employed in the survey occupations at the time of sampling (June 2024), the worker-level adoption analysis is primarily suited for studying occupational switching *into* these occupations. In Section 3.4, we examine potential displacement effects *out of* the survey occupations by studying incumbent workers—those employed in the survey occupations prior to the arrival of AI chatbots—based on their employers' subsequent adoption of the tools.

switches may serve as early indicators of the longer-run impacts of AI chatbots (Kleven et al., 2025). To understand what occupational mobility means for adopters, Table 2 compares the characteristics of origin and destination jobs for these switchers, measured in November 2022 and December 2024, respectively.

Adopters who switch occupations see 12 percentage points faster earnings growth between November 2022 and December 2024 relative to other Danish workers employed in both months. These switchers move into occupations with higher wage premia—rising from 0.46 to 0.53 standard deviations above the economy average, contributing a 3 percentage point boost to earnings growth. Even though these workers already come from occupations with above-average exposure to AI chatbots (0.92 standard deviations above the economy average, as measured by Eloundou et al. (2024)), they switch into occupations with 0.16 standard deviations higher exposure. In other words, adopters who switch occupations tend to move into ones where using AI chatbots is more relevant and where wages are higher—consistent with workers leveraging AI chatbots to access better-paying roles.

Discussion. 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., 2024; Noy and Zhang, 2023). This section presents evidence supporting a key implication of those studies: chatbot adoption is associated with occupational switching, and adopters who switch move into occupations with higher wage premia and where using AI chatbots is more relevant. This is consistent with the thesis of Autor (2024) that generative AI helps workers expand into new work by granting access to otherwise scarce expertise. Indeed, actual usage data from chatbot developers show that *information seeking* and *practical guidance* are among the most common tasks workers use chatbots for (Chatterji et al., 2025)—precisely the kinds of tasks where lack of prior expertise could otherwise be a barrier.

Occupational mobility is thus the one margin where the rapid currents of AI adoption break through to administrative labor market outcomes. While these switchers remain too few to move average adopter earnings, the substantial gains for those who do switch

suggest a channel through which AI chatbots may deliver broader earnings effects as adoption deepens (Kleven et al., 2025).

3.4 Workplace Employment

Concerns have been raised that generative AI is already displacing workers, particularly those early in their careers. Indeed, Appendix D.2 documents aggregate declines in early-career employment in several of our exposed occupations in Denmark, mirroring patterns documented in the United States (Brynjolfsson, Chandar and Chen, 2025). A key advantage of our data, however, is that we can split these employment trends by whether firms have actually adopted generative AI.

Panel (a) shows that workplaces encouraging chatbot use have seen no differential change in total hours after the arrival of AI chatbots.[20] Our estimates are precise, ruling out effects larger than 2%.

Crucially, Panel (b) shows that the decline in early-career jobs is not driven by firms adopting AI chatbots. Difference-in-differences estimates show no differential changes in the employment share of early-career workers at adopting workplaces, with estimates precise enough to rule out effects greater than a third of a percentage point. These bounds imply that AI adoption can account for at most a small fraction of the aggregate declines in early-career employment documented in Appendix D.2—declines whose connection to generative AI is itself debated (Frank et al., 2026; Iscenko and Curto Millet, 2026).[21]

Beyond levels of employment, AI chatbots may alter the composition of labor demand. Even with constant employment levels, adopting and operating new tools may require workers with different skills, leading firms to hire new workers at the expense of incumbents (Bessen et al., 2023). At the same time, assistance from AI chatbots could instead help retain workers (Brynjolfsson, Li and Raymond, 2025).

[20] Appendix F.4 shows that wage bills are similarly flat.

[21] About 43% of workplaces encourage chatbot use (Figure 1). Even if all encouraged workplaces experienced the largest decline consistent with our confidence intervals, chatbots would account for at most a $0.43 \times 0.33 \approx 0.14$ percentage point decline in the aggregate early-career share—small relative to the baseline share of about 8% (Table D.1) and the aggregate decline of about 0.6 percentage points between October 2022 and December 2024 implied by Figure D.2.

Panel (c) considers incumbent workers—those employed at the workplace before the arrival of AI chatbots—whose labor market outcomes we can track regardless of whether they remain with or separate from their original firm.[22] These workers have not experienced differential changes in hours. Figure F.9 shows that they have also not experienced differential occupational mobility. Hence, while Section 3.3 shows that workers' use of AI chatbots has helped them move *into* new occupations, employers' adoption of AI chatbots has not displaced workers *out of* their original occupations.

Panel (d) examines overall rates of job creation and destruction, showing that job churn at workplaces that encourage AI chatbots has remained stable, with confidence bands ruling out effects on monthly job churn rates greater than 0.5% of employment.

Heterogeneity by Intensity of Adoption. Figure 6, Panels (c) and (d) examine the effects of additional initiatives in enterprise chatbots and training. Although Section 2.2 shows that uptake and reported benefits are highest when the full suite of initiatives is implemented, we find no detectable impacts on overall employment or early-career jobs at these workplaces. Similarly, Appendix F.4 shows that incumbent workers and job churn have not changed differentially at workplaces with these investments.

Occupational Heterogeneity. Figure 7, Panels (c) and (d) report the estimated effects of AI chatbots on workplace employment and early-career jobs across the 11 occupations in our sample. Confidence bands are again wide for these occupation-specific estimates and should therefore be interpreted with caution. Nonetheless, we find no statistically significant effects in any occupation, and point estimates are generally near zero.[23] Even between our most contrasting groups—software developers and marketers (with flexible labor rules) versus teachers (with collective bargaining)—we observe no systematic differences in employment impacts. Moreover, while Appendix D.2 documents aggregate declines in early-career jobs in several of our exposed occupations, including

[22]This analysis of incumbent workers resembles the research designs of Autor et al. (2014); Walker (2013).

[23]Appendix F.4 similarly shows null results for job churn and worker sorting.

software development, customer support, legal professions, and marketing, these declines are not driven by firms adopting generative AI chatbots. This suggests that other forces, such as changes in the macroeconomic environment or anticipatory hiring freezes unrelated to actual adoption, may be driving the aggregate pattern.

Discussion. Workplaces that have adopted AI chatbots have not fared differently in terms of employment, early-career jobs, job churn, or incumbent worker outcomes. How do we reconcile this "still surface" with the "rapid currents" documented in Section 2? The task creation evidence in Section 2.3 provides a key part of the answer: employers are absorbing generative AI through task reorganization—creating new roles, reassigning responsibilities, investing in training and oversight—rather than adjusting on the margin of total labor input. Consistent with this, workers who report larger time savings are also more likely to have taken on new tasks (Appendix Figure E.2), suggesting that individual productivity gains are absorbed into task reorganization. In this sense, the null employment results are not evidence of inaction but of a specific mode of adjustment: one that operates through the internal reorganization of work rather than through the external labor market. This pattern is consistent with the productivity J-curves observed at the onset of prior general-purpose technologies, where substantial organizational investments precede measurable changes in economic outcomes (Brynjolfsson, Rock and Syverson, 2021; David, 1990; Solow, 1987).

3.5 Perceived Impacts

To complement the difference-in-differences analysis, we asked respondents directly: *"Have AI chatbots affected your labor earnings?"* If so, we followed up with, *"By how much?"* If AI chatbots were meaningfully affecting labor market conditions, workers in exposed occupations would likely notice.

Appendix E.5 reports workers' perceived earnings impacts from AI chatbots. 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. While individuals may struggle to assess

such counterfactuals precisely, these responses are informative in three ways.

First, the fact that virtually all non-adopters report no earnings impact suggests that occupation-wide spillovers are minimal. Despite broad adoption and changes in work content, AI chatbots do not appear to have shifted equilibrium wages or hours—reinforcing the conclusion from our identification analysis in Section 3.1 that the "missing intercept" remains small.

Second, even adopters report small average earnings impacts, ranging from 0% to 0.5% depending on occupation and employer chatbot initiatives. These estimates are an order of magnitude smaller than the time savings reported in Section 2.2 and fall within the confidence intervals of the difference-in-differences estimates for earnings in Section 3.2.

Third, the limited average earnings impacts primarily reflect the overwhelming share of workers reporting no change—rather than large positive and negative effects offsetting each other. This supports the conclusion that AI chatbots have not had significant heterogeneous impacts on workers' earnings.

Taken together, these self-reported perceptions from workers in the most exposed occupations reinforce our finding that AI chatbots have so far had limited impacts on earnings in the labor market.

4 Conclusion

This paper studies how labor markets adjust at the onset of major technological change, providing large-scale evidence on the early impacts of AI chatbots—the most widely adopted generative AI tool to date. By linking extensive surveys on AI chatbot adoption and work organization in 11 exposed occupations to administrative labor market records in Denmark, we document both what moves and what does not.

What moves is the structure of work. Most employers in exposed occupations have implemented chatbot initiatives, workers report substantial productivity benefits, and new AI-related tasks have emerged broadly—especially in workplaces that encourage and invest in the tools. These patterns are consistent with task-based theories of technology adoption

(Acemoglu and Restrepo, 2018*a*) and with evidence that complementary organizational investments are critical to realizing the productivity gains from new technologies (Bloom, Sadun and Reenen, 2012; Brynjolfsson and Hitt, 2000). Occupational mobility is the one margin where these rapid currents break through to administrative labor market outcomes: workers who use AI chatbots are more likely to have transitioned into new occupations, particularly into roles where workers have greater discretion over their toolsets. These workers move into better-paid occupations where using AI chatbots is more relevant. This supports the view that generative AI may help workers access otherwise scarce expertise, with broader implications for inequality in the labor market (Autor, 2024; Autor and Thompson, 2025).

What does not move is the surface—recorded hours and earnings. Using difference-in-differences, we estimate precise null effects at both the worker and workplace levels, ruling out effects larger than 2% two years after ChatGPT's launch. These null results hold across occupations, across intensities of adoption, and even among workers who report substantial daily time savings. While aggregate early-career employment has declined in several exposed occupations, firms' adoption of AI chatbots does not drive these declines.

We conjecture that these two patterns may be mutually reinforcing: when earnings and recorded hours remain stable, workers and workplaces absorb technological change through task reallocation and other unrecorded adjustments, reducing pressure on pay or hours. If so, the absence of measurable labor market effects is not evidence that nothing is happening—but rather that transformation is occurring on margins that conventional economic statistics do not capture. This pattern echoes the productivity J-curves observed at the onset of prior general-purpose technologies (Brynjolfsson, Rock and Syverson, 2021; David, 1990; Solow, 1987), and our data reveal what lies within the trough: the reorganization of work that absorbs technological change before it surfaces in hours or earnings. Understanding this process requires the kind of linked data infrastructure we develop here, capturing both the reorganization of work within firms and its translation into administrative labor market outcomes.

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*a*. "Artificial intelligence, automation, and work." In *The economics of artificial intelligence: An agenda*. 197–236. University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2018*b*. "The race between man and machine: Implications of technology for growth, factor shares, and employment." *American economic review*, 108(6): 1488–1542.

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.

Amodei, Dario. 2025. "Behind the Curtain: A White-Collar Bloodbath." *Interview reported by Axios*, Reported in Axios; accessed 2025-05-28.

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

Autor, David, and Neil Thompson. 2025. "Expertise." *Journal of the European Economic Association*, jvaf023.

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.

Autor, David H, David Dorn, Gordon H Hanson, and Jae Song. 2014. "Trade Adjustment: Worker-Level Evidence." *The Quarterly Journal of Economics*, 129(4): 1799–1860.

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

Becker, Joel, Nate Rush, Elizabeth Barnes, and David Rein. 2025. "Measuring the impact of early-2025 AI on experienced open-source developer productivity." *arXiv preprint arXiv:2507.09089*.

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.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans do IT better: US multinationals and the productivity miracle." *American Economic Review*, 102(1): 167–201.

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.

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, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford Digital Economy Lab.

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.

Chandar, Bharat. 2025. "Tracking employment changes in AI-exposed jobs." *Available at SSRN 5384519*.

Chatterji, Aaron, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. "How People Use ChatGPT." National Bureau of Economic Research Working Paper 34255.

Chicago Public Schools. 2024. "AI Guidebook." *PDF*, Accessed 2026-02-15.

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. 2024. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *Forthcoming, Organization Science*.

Digitaliseringsstyrelsen. 2024. "Guide til offentlige myndigheder om ansvarlig anvendelse af generativ kunstig intelligens." *PDF*, Version 1 (Jan. 2024); note dated 11 Mar. 2024. Accessed 2026-02-15.

Dillon, Eleanor W, Sonia Jaffe, Nicole Immorlica, and Christopher T Stanton. 2025. "Shifting Work Patterns with Generative AI." National Bureau of Economic Research. Conditionally accepted, American Economic Review: Insights.

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.

Frank, Morgan R., Alireza Javadian Sabet, Lisa Simon, Sarah H. Bana, and Renzhe Yu. 2026. "AI-Exposed Jobs Deteriorated before ChatGPT." *arXiv preprint arXiv:2601.02554*.

Freund, Lukas B, and Lukas F Mann. 2025. "Job Transformation, Specialization, and the Labor Market Effects of AI."

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

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, et al. 2025. "Which economic tasks are performed with ai? evidence from millions of claude conversations." *arXiv preprint arXiv:2503.04761*.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore. 2025. "The labor market effects of generative artificial intelligence." *Available at SSRN*.

Humlum, Anders. 2021. "Robot Adoption and Labor Market Dynamics."

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.

Ide, Enrique, and Eduard Talamas. 2025. "Artificial Intelligence in the Knowledge Economy." *Journal of Political Economy*, 133(12).

Iscenko, Zanna, and Fabien Curto Millet. 2026. "Looking for the Ladder: Is AI Impacting Entry-Level Jobs?" Economic Innovation Group. The American Worker Project.

Jäger, Simon, Suresh Naidu, and Benjamin Schoefer. 2025. "Collective Bargaining, Unions, and the Wage Structure: An International Perspective." In *Handbook of Labor Economics*. Vol. 6. Elsevier.

Kauhanen, Antti, and Petri Rouvinen. 2025. "Assessing early labour market effects of generative AI: evidence from population data." *Applied Economics Letters*, 1–4.

Kleven, Henrik, Claus Kreiner, Kristian Larsen, and Jakob Søgaard. 2025. "Micro versus macro labor supply elasticities: The role of dynamic returns to effort." *American Economic Review*, 115(9): 2849–2890.

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.

Lichtinger, Guy, and Seyed Mahdi Hosseini Maasoum. 2025. "Generative AI as Seniority-Biased Technological Change: Evidence from US Résumé and Job Posting Data." *Available at SSRN*.

Lütke, Tobi. 2025. "Reflexive AI usage is now a baseline expectation at Shopify." *X post (sharing internal memo)*, Accessed 2026-02-15.

Mas, Alexandre. 2025. "Non-Wage Amenities." National Bureau of Economic Research 33643.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. "The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve(s)." Center for Economic Studies, U.S. Census Bureau Working Paper CES-25-27. Also available via SSRN: SSRN 5036270.

Munch, Jakob R, and William W Olney. 2024. "Offshoring and the Decline of Unions." IZA Discussion Paper 17116.

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.

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.

Robinson, Joan. 1969. *The economics of imperfect competition.* Springer.

Rosen, Sherwin. 1986. "The theory of equalizing differences." *Handbook of labor economics*, 1: 641–692.

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.

The Budget Lab at Yale. 2026. "Evaluating the Impact of AI on the Labor Market: November/December CPS Update." *https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-novemberdecember-cps-update*, Accessed 2026-02-15.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. "Working with AI: Measuring the Occupational Implications of Generative AI." *arXiv preprint arXiv:2507.07935.*

Walker, W Reed. 2013. "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce." *The Quarterly journal of economics*, 128(4): 1787–1835.

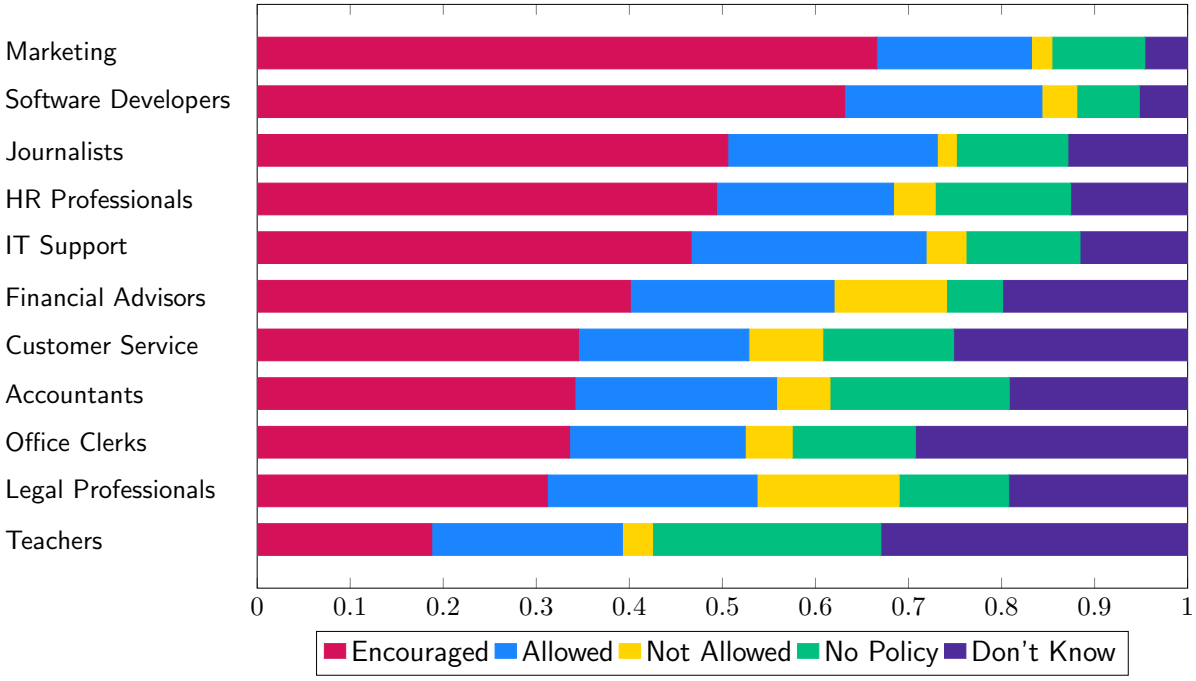
Walters, Christopher. 2024. "Empirical Bayes Methods in Labor Economics." In *Handbook of Labor Economics.* Vol. 5, 183–260. Elsevier.

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J Davis, Kevin M Foster, Aaron Jalca, Brent H Meyer, Paul Mizen, Michael A Navarrete, et al. 2026. "Firm Data on AI." National Bureau of Economic Research.

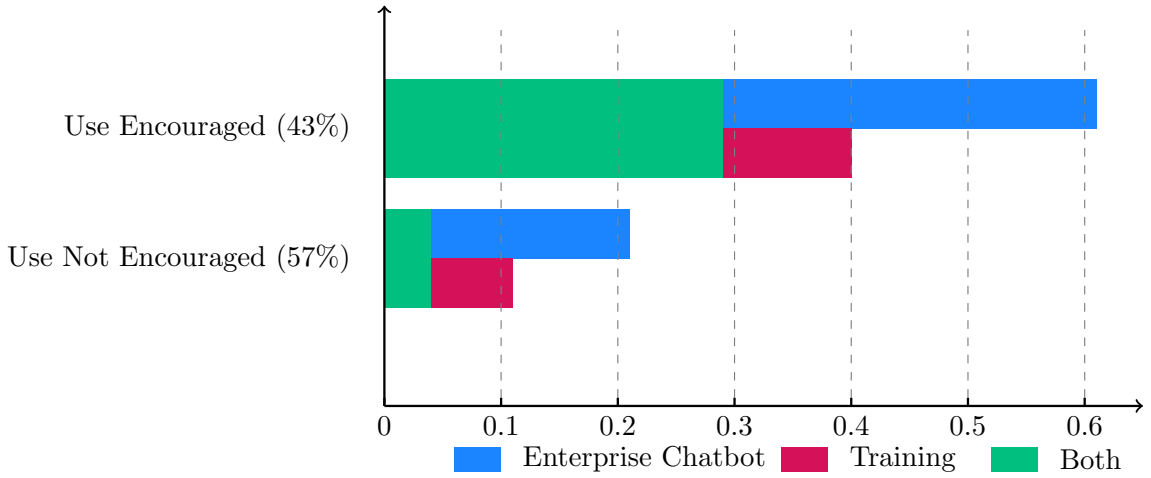
Figures and Tables

Figure 1: Employer Initiatives for AI Chatbot Use

(a) Employer Usage Policies

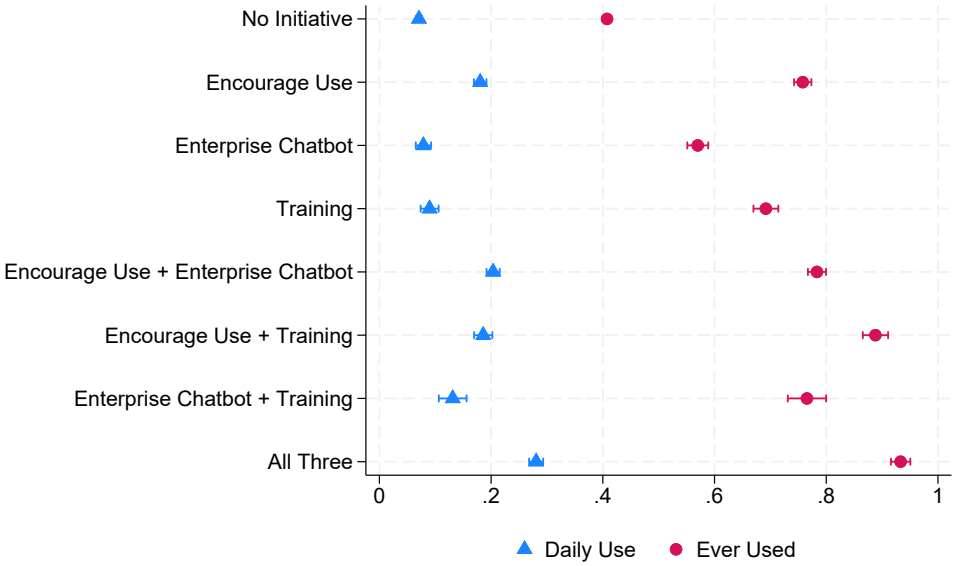


(b) Bundles of Employer Initiatives



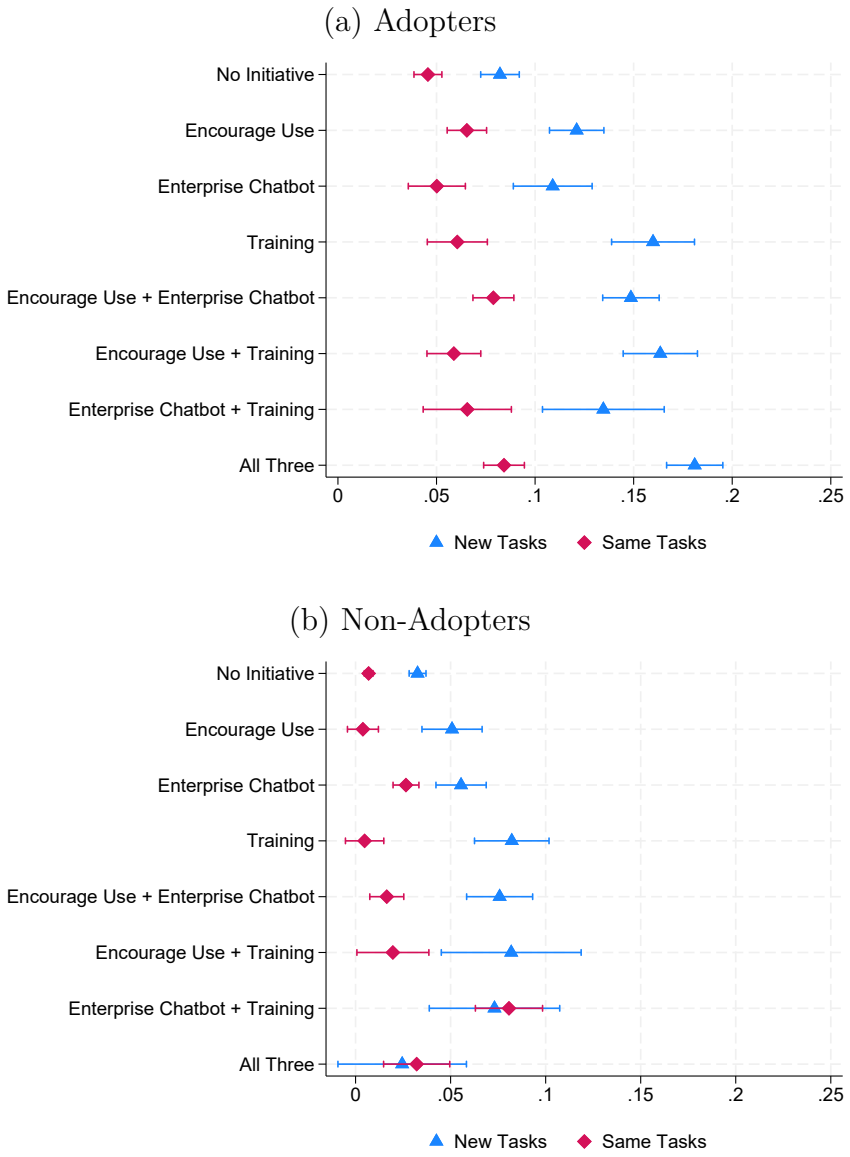
Notes: Panel (a) shows, by occupation, the share of workers subject to different employer policies on the use of AI chatbots for work. Figure F.1 provides a workplace-level version of this graph, yielding similar results. Panel (b) shows the share of workers affected by different combinations of employer-led 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 2: Chatbot Adoption by Employer Initiatives



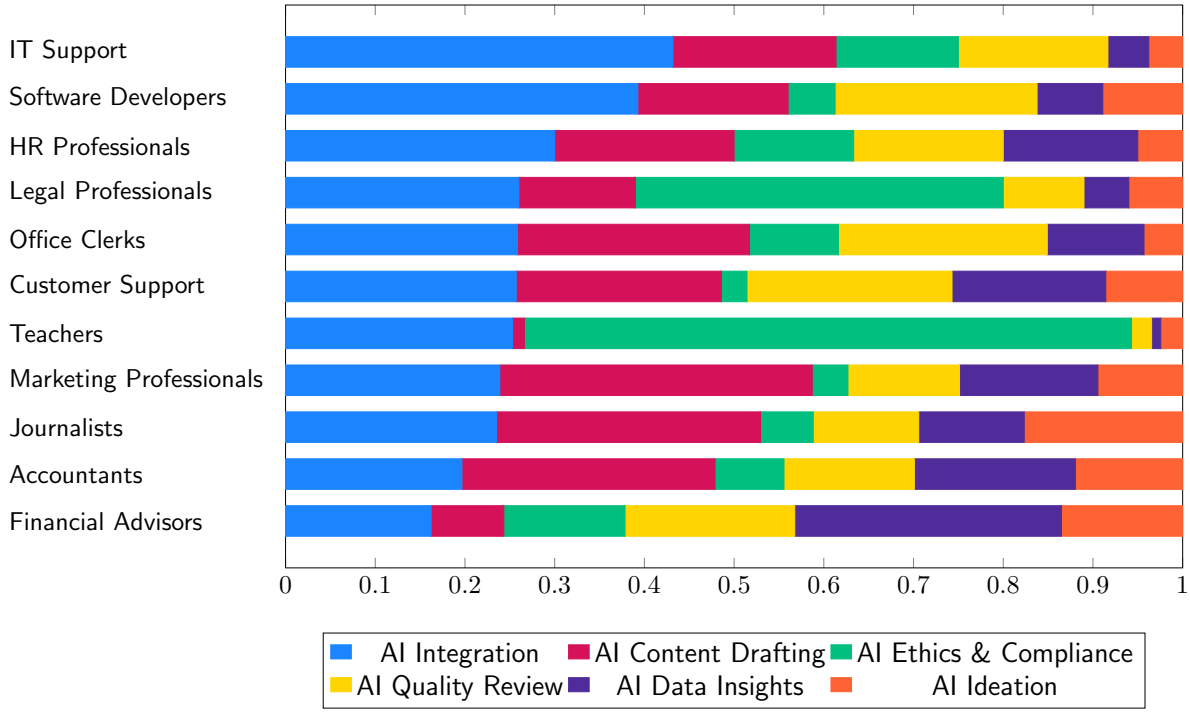
Notes: This figure illustrates how workers' use of AI chatbots vary with their employer initiatives. The estimates are based on predicted values from Equation (1), varying EmployerInitiatives while holding workers' characteristics X at their mean values. Whiskers represent 95% confidence bands of the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

Figure 3: New Workloads From AI Chatbots by Employer Initiatives



Notes: This figure illustrates how new workloads created by AI chatbots vary with employer initiatives. Panel (a) presents results for adopters (i.e., workers who have ever used AI chatbots for work), while Panel (b) focuses on non-adopters. Estimates are 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 1, Columns (7)-(10) report the full set of regression estimates. *Sample:* All completed responses from the 2024 survey linked to registry data.

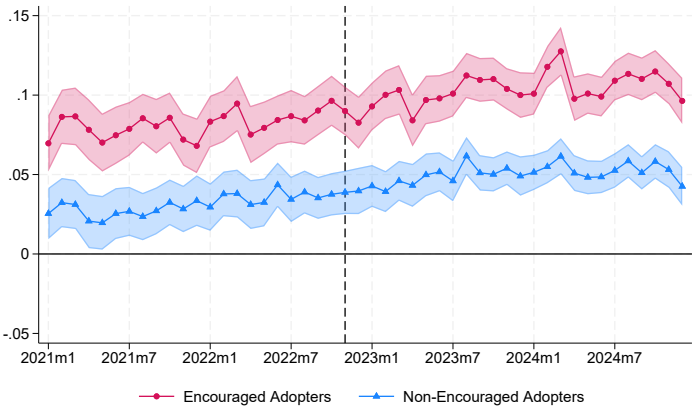
Figure 4: Composition of AI Tasks



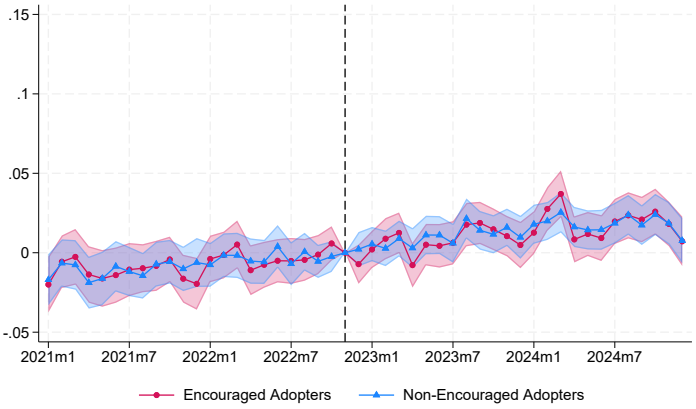
Notes: This figure shows the distribution of reported new job tasks across major task categories for each occupation. *AI Ideation* refers to "using AI to generate or expand creative ideas, such as concepts, strategies, or solutions," *AI Content Drafting* refers to "using AI tools to produce initial 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 B.3 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 of *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 5: Have Adopting Workers Fared Differently? (Adopter DiD)

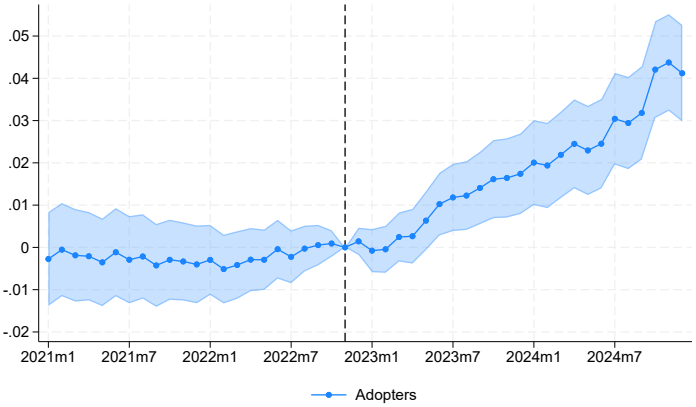
(a) Log Earnings: Differences in Means



(b) Log Earnings: Difference-in-Differences (DiD)



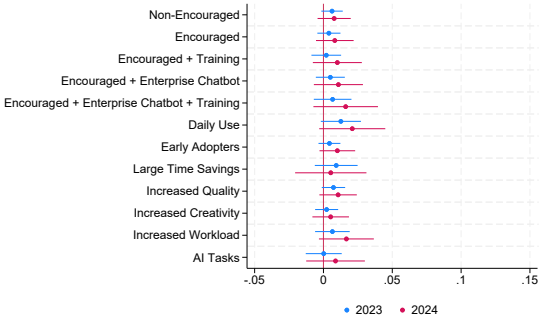
(c) Occupational Mobility (DiD)



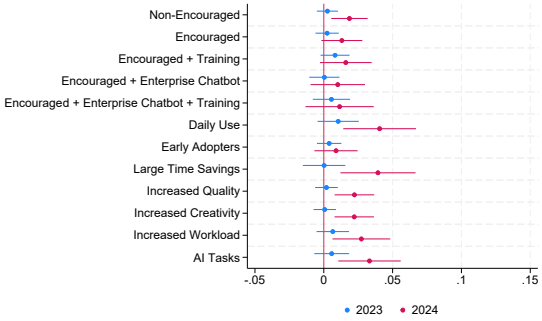
Notes: This figure shows labor market outcomes for chatbot adopters. Panels (a) and (b) separate adopters with and without employer encouragement, while Panel (c) pools the two (Figure F.5, Panel (b) shows the groups fare similarly). In all panels, adopters are compared to non-encouraged non-adopters. Panel (a) reports the difference in log earnings between AI chatbot adopters and non-adopters, controlling for survey occupation fixed effects and predetermined worker characteristics. Panel (b) presents the difference-in-differences corresponding to Panel (a), indexed to November 2022, the launch of ChatGPT. Panel (c) shows occupational mobility measured as full-time-equivalent (FTE) employment in workers' latest (December 2024) occupations. The difference-in-differences estimates are based on the specification in Equation (2). Shaded areas represent 95% confidence intervals. In December 2024, the standard deviation of residual earnings changes (Panel b) is 0.285 for encouraged adopters and 0.287 for non-encouraged adopters; the standard deviation of residual occupational mobility changes (Panel c) is 0.217. *Sample:* All completed responses from the 2024 survey linked to registry data.

Figure 6: Difference-in-Differences Estimates: Heterogeneity by Adoption Intensity

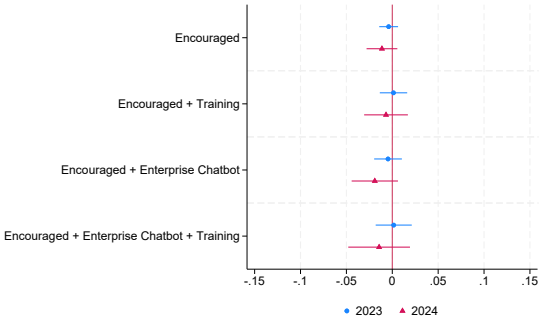
(a) Worker Earnings (Adopter DiD)



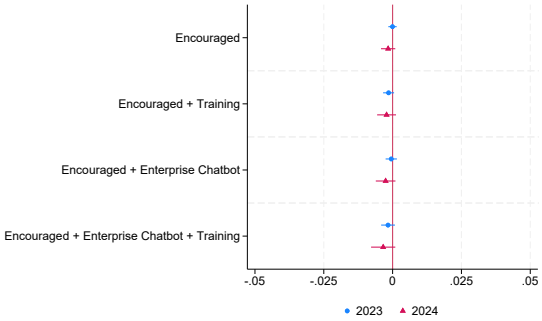
(b) Occupational Mobility (Adopter DiD)



(c) Workplace Employment (Encouraged DiD)



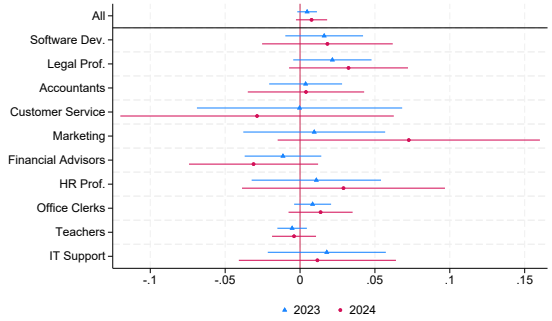
(d) Early-Career Share (Encouraged DiD)



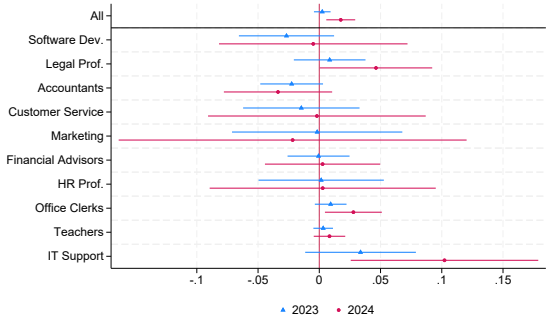
Notes: This figure shows how our difference-in-differences estimates vary with the intensity of adoption. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. Panels (a) and (b) focus on log earnings and occupational mobility for different adopters relative to non-encouraged non-adopters (Figure 5). Occupational mobility is defined as full-time-equivalent (FTE) employment in workers' latest (December 2024) occupations. *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 the time of our 2023 survey. *Large Time Savings* indicates time savings of more than 60 minutes per day of usage. *Increased Quality / Creativity* refers to users reporting improved output quality or creativity. *Increased Workload* refers to users reporting more workloads due to AI chatbots. *AI Task* refers to workers who report new AI-related job tasks. Panels (c) and (d) show workplace outcomes (log employment and early-career share) for workers who report encouraged usage policies relative to those without encouragement (Figure 8). Early-career workers are defined as workers aged 22-25; see Appendix D.2 for details. *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 7: Difference-in-Differences Estimates: Heterogeneity by Occupation

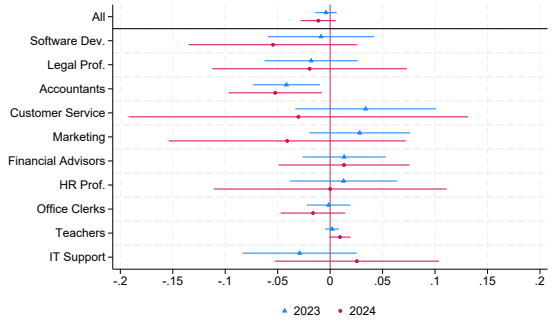
(a) Worker Earnings (Adopter DiD)



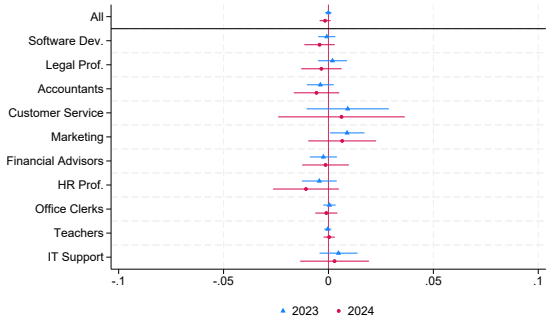
(b) Occupational Mobility (Adopter DiD)



(c) Workplace Employment (Encouraged DiD)



(d) Early-Career Share (Encouraged DiD)



Notes: This figure shows occupation-specific effects of AI chatbot adoption. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. Panels (a) and (b) focus on log earnings and occupational mobility for adopters relative to non-encouraged non-adopters (Figure 5). Occupational mobility is defined as full-time-equivalent (FTE) employment in workers' latest (December 2024) occupations. Panels (c) and (d) show workplace outcomes (log employment and early-career share) for workers who report encouraged usage policies relative to those without encouragement (Figure 8). Early-career workers are defined as workers aged 22-25; see Appendix D.2 for details. *Sample:* All completed 2024 survey responses linked to registry data.

Figure 8: Have Adopting 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 employment; Panel (b), the employment share of early-career workers (aged 22-25); Panel (c), log hours of incumbent workers (those employed at the workplace throughout the pre-period, 2021M1–2022M11); and Panel (d), job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment). Estimates are based on the dynamic difference-in-differences specification in Equation (2), with shaded areas representing 95% confidence intervals. In December 2024, the standard deviations of the residuals are 0.56 for employment changes (Panel a), 0.04 for Early-Career workers (Panel b), 0.08 for incumbent worker hours (Panel c), and 0.026 for job creation and 0.01 for job destruction (Panel d). *Sample:* All completed 2024 survey responses linked to registry data.

Table 1: Worker Adoption, Reported Benefits, and New Workloads by Employer Initiatives

	Adoption		Reported Benefits (Among Ever Used)				New Workloads (Among Ever Used)		New Workloads (Among Never Used)	
	Ever Used (1)	Daily Use (2)	Any Time Saving (3)	60+ min/day (4)	Quality (5)	Creativity (6)	Same Tasks (7)	New Tasks (8)	Same Tasks (9)	New Tasks (10)
Encouraged	0.350***	0.110***	0.087***	0.034***	0.076***	0.065***	0.020***	0.039***	-0.003	0.018**
	(0.009)	(0.008)	(0.011)	(0.010)	(0.014)	(0.013)	(0.006)	(0.008)	(0.003)	(0.009)
Enterprise Chatbot	0.163***	0.008	-0.000	-0.024**	-0.033*	-0.016	0.005	0.027***	0.020***	0.023***
	(0.012)	(0.007)	(0.015)	(0.011)	(0.017)	(0.016)	(0.008)	(0.010)	(0.006)	(0.007)
Training	0.284***	0.019***	-0.065***	-0.006	-0.041**	-0.031*	0.015**	0.078***	-0.002	0.050***
	(0.014)	(0.007)	(0.018)	(0.011)	(0.017)	(0.017)	(0.007)	(0.016)	(0.004)	(0.016)
Encouraged × Enterprise Chatbot	-0.137***	0.015	0.009	0.014	0.045**	0.009	0.009	0.001	-0.007	0.002
	(0.016)	(0.013)	(0.019)	(0.016)	(0.023)	(0.021)	(0.011)	(0.014)	(0.009)	(0.015)
Encouraged × Training	-0.154***	-0.014	0.028	-0.013	0.066***	0.079***	-0.022**	-0.035*	0.018	-0.018
	(0.017)	(0.014)	(0.023)	(0.016)	(0.024)	(0.024)	(0.011)	(0.021)	(0.014)	(0.032)
Enterprise Chatbot × Training	-0.089***	0.033**	0.040	-0.011	0.095***	0.022	0.001	-0.052**	0.056**	-0.032
	(0.023)	(0.016)	(0.032)	(0.020)	(0.033)	(0.031)	(0.015)	(0.025)	(0.026)	(0.029)
Encouraged × Enterprise Chatbot × Training	0.109***	0.038	0.023	0.043*	-0.047	-0.020	0.011	0.042	-0.056*	-0.050
	(0.028)	(0.025)	(0.037)	(0.026)	(0.039)	(0.039)	(0.020)	(0.030)	(0.033)	(0.043)
Worker Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
No Initiative, Level	0.399	0.061	0.702	0.148	0.460	0.448	0.051	0.080	0.007	0.027
Within R^2	0.217	0.068	0.041	0.016	0.021	0.019	0.012	0.014	0.011	0.009
R^2	0.292	0.183	0.077	0.063	0.073	0.032	0.027	0.029	0.015	0.043
Observations	24796	24796	14790	14790	14790	14790	14790	14790	10006	10006

Notes: This table reports how workers' chatbot adoption (Columns 1–2), user-reported benefits (Columns 3–6), and new workloads from AI chatbots (Columns 7–10) vary with employer chatbot initiatives. *60+ min/day* indicates time savings of more than one hour per day of use. The regressions control for survey occupation fixed effects and workers' predetermined characteristics (age, gender, experience); see Equation (1). *Sample*: All completed 2024 survey responses linked to registry data.

Table 2: Job Characteristics for Adopters Who Switch Occupations

	Standard deviation from economy average		Change in %
	November 2022	December 2024	Nov '22 to Dec '24
	(1)	(2)	(3)
Earnings	0.27	0.29	0.12
	(0.02)	(0.02)	(0.01)
Occupational Wage Premium	0.46	0.53	0.03
	(0.02)	(0.02)	(0.00)
LLM Exposure	0.92	1.08	
	(0.02)	(0.02)	

Notes: This table reports characteristics of origin and destination jobs for adopters who switch occupations between November 2022 and December 2024, measured relative to all Danish workers employed in both months. Adopters who switch occupations (`AdoptSwitch`) are defined as workers who (i) have used AI chatbots for work, (ii) since November 2022, have increased their hours worked in their December 2024 occupations, and (iii) were employed in both months. Of workers satisfying (i) and (ii), 99% also satisfy (iii). We compare these to all Danish workers who were employed in both November 2022 and December 2024 (i.e., satisfied requirement (iii)). The estimates are OLS estimates of β_1 from the regressions $Y_i = \beta_0 + \beta_1 \texttt{AdoptSwitch}_i$, with standard errors in parentheses. Outcomes in Columns (1) and (2) are measured in economy-wide standard deviations; Column (3) reports within-worker changes in percentage points. *Occupational Wage Premium* is the average hourly wage in the occupation. *LLM Exposure* is the "Direct Exposure (E1)" measure from Eloundou et al. (2024). *Sample*: All completed 2024 survey responses linked to registry data.

Online Appendix

Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI

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 . . . 7

B Data Construction **20**
B.1 Sampling Protocol . . . 20
B.2 Survey Sample . . . 21
B.3 Classification of New Job Tasks . . . 26

C Institutional Setting **31**
C.1 Wage Setting Systems in Denmark . . . 31

D Labor Market Trends **33**
D.1 Aggregate Employment . . . 33
D.2 By Age Groups (Brynjolfsson, Chandar and Chen, 2025) . . . 34

E Additional Results **38**
E.1 Employer Initiatives . . . 38
E.2 AI Chatbot Products . . . 38
E.3 Reported Benefits by Occupation . . . 40
E.4 Workloads and Task Creation . . . 42
E.5 Perceived Impacts . . . 44

F Robustness Analysis **45**
F.1 Employer Initiatives . . . 45
F.2 Worker Earnings . . . 46
F.3 Job Mobility . . . 50
F.4 Workplace Employment . . . 52
F.5 Additional Controls . . . 56
F.6 Coworker Encouragement . . . 58

G Estimation Procedures **64**
G.1 Empirical Bayes Shrinkage 64

H Invitation Letter **66**
H.1 English Translation 67
H.2 Original Danish Version 69

I Survey Questionnaire **71**
I.1 English Translation 72
I.2 Original Danish Version 77

A Theoretical Framework

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

In Section A.1, we present a simple partial-equilibrium model of chatbot adoption to interpret the descriptive patterns in Section 2.2. We posit a Roy (1951)-style selection model in which employer initiatives shift both the costs and benefits of adoption, and analyze how these shifts affect workers' adoption decisions and the work-related benefits they report.

In Section A.2, we embed this adoption model within a labor market equilibrium framework 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 3.

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 ordered by their benefits, such that $\beta_b \leq 0$. While we assume linear functional forms for analytical simplicity, our predictions 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.[24]

We conceptualize employer initiatives $E \in \{0,1\}$ as shifting workers' costs and benefits

[24] A common finding from RCTs is that AI chatbots yield greater benefits for less experienced (and thus typically younger) workers (Brynjolfsson, Li and Raymond, 2025; Dell'Acqua et al., 2024; Noy and Zhang, 2023). If these workers also find it easier to adopt new tools, then $\beta_c \geq 0$.

through the parameters α_c and α_b:

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

While we focus on a binary initiative for analytical simplicity, our propositions extend naturally to combinations of initiatives—such as use policies, enterprise chatbots, and training.

Workers adopt AI chatbots if their perceived benefits exceed the associated 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 average reported benefits; formal derivations follow below.

Adoption rates. Table 1, Columns 1–2 show how chatbot adoption rates vary with employer initiatives. The optimal adoption rate is:

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

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 initiatives raising net benefits ($\hat{\alpha}$) increase adoption, with stronger effects when $\hat{\beta}$ is numerically small—i.e., when workers are relatively homogeneous in their net benefits from AI chatbots.

In this light, the fact that employer-encouraged use in Table 1 is associated with a *doubling* of adoption rates suggests that encouragement is powerful relative to the degree of individual heterogeneity in net benefits.

Benefits for users. Table 1, Columns 3–6 show how reported benefits among chatbot users vary with employer initiatives. The average benefit reported by users is:

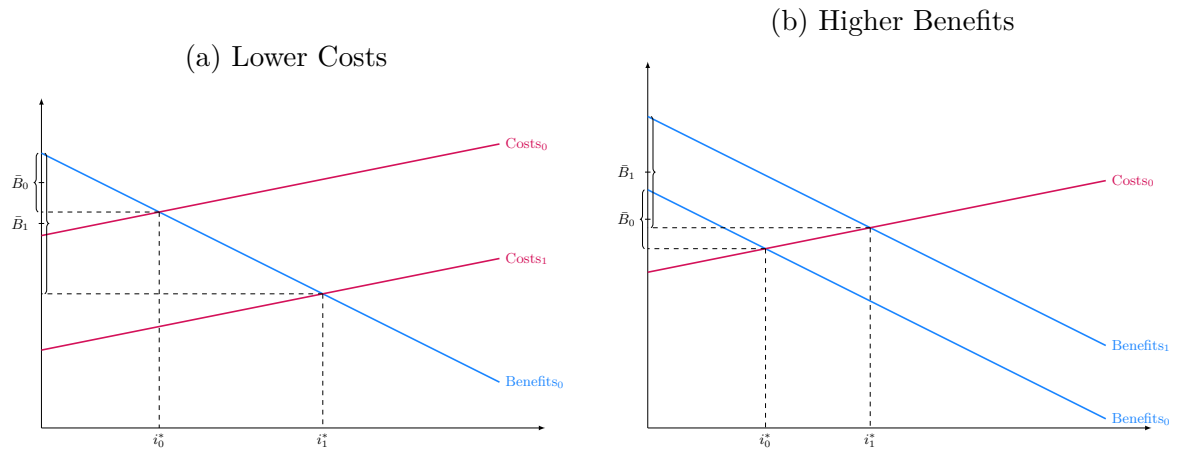
$$\bar{B}(E) = \mathbb{E}[B_i(\alpha) \mid i \leq i^*(\alpha)] = \gamma\alpha_c + (1-\gamma)\alpha_b, \tag{9}$$

where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$. Equation (9) yields two key insights.

First, initiatives that reduce adoption costs (α_c) *lower* average reported benefits through a *selection effect*: lower costs induce adoption by workers with lower idiosyncratic benefits (higher i). This negative selection effect is stronger when workers are more homogeneous in net benefits ($\hat{\beta}$ numerically small, so initiatives generate larger adoption increases; cf. Equation (8)) but more heterogeneous in gross benefits (β_b large in absolute value, so marginal adopters have substantially lower benefits).

Second, initiatives that increase benefits (α_b) have two offsetting effects. A higher α_b directly raises each user's benefit B_i, but also expands adoption to workers with lower idiosyncratic benefits, triggering the same negative selection effect. The net effect on $\bar{B}$ is positive if and only if $\gamma < 1$—a condition more likely when workers are similar in gross benefits (β_b small) but differ in adoption costs (β_c 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 initiatives that reduce adoption costs (α_c); Panel (b) depicts initiatives that enhance benefits (α_b). Lower costs unambiguously increase adoption and reduce average reported benefits via negative selection. Initiatives that raise benefits also increase adoption, but their impact on average reported benefits is theoretically ambiguous and, at most, equal to their individual-level effect.

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

The fact that Table 1 shows higher reported benefits in encouraged settings is therefore striking. It implies two things: (i) employer encouragement primarily raises perceived benefits (α_b), since reductions in adoption costs alone would generate only the negative selection effect; and (ii) these benefit increases outweigh the adverse selection—implying $\gamma < 1$, so that encouragement is powerful relative to individual heterogeneity.

By contrast, enterprise chatbots and training—when implemented in isolation—are associated with lower reported benefits, potentially reflecting either that these initiatives reduce individual benefits ($\alpha_{b1} < 0$) or that negative selection effects dominate any positive gains.

A.1.3 Discussion

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

First, firms that encourage chatbots may be those with the greatest underlying benefits from the tools—even absent any initiatives—introducing potential reverse causality. In our framework, this corresponds to E being positively related to α_{b0}. Second, employers may target encouragement toward employees expected to benefit most individually, resembling a rotation of the cost and benefit schedules in Figure A.1 rather than simple vertical shifts.

Section F addresses both margins empirically. First, Section F.5 examines whether employers select into encouragement based on underlying benefits from chatbot use. We show that the effects of encouragement remain robust when controlling for firm- and worker-level characteristics, including firm age, size, and productivity, as well as workers' detailed task mixes. Second, Section F.6 implements a "coworker encouragement" design, using the average encouragement reported by a worker's colleagues to capture *workplace-*

level variation. Because this approach nets out *individual*-level targeting, it addresses potential bias from employers directing encouragement toward workers with the highest idiosyncratic benefits. As shown in Table F.4, results remain largely robust under this design.

A.2 Labor Market Outcomes

This section provides a theoretical analysis of how AI chatbots affect labor markets and discusses how these effects can be identified empirically.

Modeling the labor market impacts of AI chatbots raises several key questions: How do chatbots enter the production process? What explains why some firms and workers adopt them while others do not? Are labor markets competitive, or do firms exercise monopsony power? Before committing to a formal framework, Section A.2.1 considers the theoretical implications of each issue.

Section A.2.2 then introduces a model of chatbot adoption and labor market outcomes, embedding the Roy adoption model from Section A.1 into an 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. Section A.2.4 derives predictions for how AI chatbots affect labor market outcomes, and Section A.2.5 applies the model to interpret our difference-in-differences estimates from Section 3.

For clarity of exposition, we focus on encouraged use as the employer initiative, reflecting its central role in the empirical analysis.

A.2.1 Modeling Choices

Chatbot technology. A first key question is how chatbots enter the production process: Are they "co-pilots," contributing to output only when prompted by humans, or closer to "co-workers," operating autonomously (Ide and Talamas, 2025)? More broadly, do AI chatbots complement or substitute for labor?

Acemoglu and Restrepo (2018*b*, 2019) provide a framework for how new technologies affect labor productivity through three channels: a negative *task displacement effect*, where

chatbots automate tasks previously performed by humans (e.g., automated copyediting); a positive *task reinstatement effect*, where new tasks for humans emerge (e.g., oversight or integration of AI outputs); and a positive *productivity effect*, where efficiency gains enable firms to scale up production, including labor demand—depending on the elasticity of product demand. Adoption of new technologies increases labor demand only when the combined productivity and reinstatement effects outweigh the displacement effect. A key consideration is whether chatbot use augments or displaces workers' expertise (Autor and Thompson, 2025); several experiments show that chatbots benefit less experienced workers, helping to close skill gaps within professions (Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023).

Modeling chatbots is further complicated by the fact that adoption is not solely a firm-level decision; it also depends on workers' choices to use the tools. Following Section A.1, workers use chatbots if the productivity benefits exceed the adoption costs. In particular, if the amenity value of chatbot use is sufficiently high, workers may adopt even when doing so reduces their productivity.

In sum, the implications of chatbot adoption for labor demand and productivity are theoretically ambiguous at both the firm and worker levels. At the firm level, the outcome depends on whether the productivity and reinstatement effects together outweigh displacement. At the worker level, it depends on the amenity value of chatbot use.

Firm and worker heterogeneity. Another key question is why some firms and workers adopt chatbots while others do not. Following Mas (2025); Rosen (1986), chatbots can be viewed as a "productive amenity": firms encourage chatbot use if the profit gains exceed costs, while workers adopt if the productivity benefits outweigh adoption costs (as in Section A.1). In competitive labor markets, Rosen (1986) shows that the wage effects of an amenity depend on the tradeoffs at the margin—the costs and benefits faced by firms and workers just indifferent to providing or accepting it. This implies that the wage effects of chatbot encouragement hinge on the marginal firm and worker.

Labor market competition. Who ultimately benefits from the productivity gains of AI chatbots depends crucially on labor market competition. In competitive markets, workers are compensated only for productivity gains not driven by firm-level investments (Acemoglu and Pischke, 1998; Becker, 1964). Worker-level adoption should then affect wages (Figure 5), whereas firm-level adoption should not. Instead, firm-level adoption should manifest in employment (Figure 8), as firms move down their product demand curves.

An important alternative is the non-discriminatory monopsony model of Card et al. (2018); Robinson (1969), in which firms exert monopsony power because workers have heterogeneous preferences over employers but firms cannot price discriminate across workers. In this framework, any earnings effects from chatbots arise at the firm level through labor demand adjustments, not through within-firm wage differentials between adopters and non-adopters. Thus, within-firm earnings gaps (Figure 5) should not emerge, but firm-level outcomes—total wage bills, employment, and average wages—should be affected (Figure 8). Whether employment effects translate into wages depends on the elasticity of workers' labor supply across employers—that is, how heterogeneous their firm preferences are.

A.2.2 Formal Setup

This section introduces a formal model of chatbot adoption and labor market outcomes. We make several simplifying assumptions. First, AI chatbots affect output only when actively deployed by workers, reflecting current technology that depends on direct human prompting. Second, workers' adoption decisions follow the Roy framework of Section A.1, with heterogeneity in both productivity benefits and adoption costs. We apply a parallel Roy model to firms' encouragement decisions. Finally, product markets are characterized by monopolistic competition and labor markets are perfectly competitive. While the individual components draw on standard frameworks, no existing model combines worker-level adoption decisions with firm-level encouragement policies and equilibrium sorting

across firms. Formalizing these interactions jointly allows us to derive predictions for equilibrium outcomes in Section A.2.4 and to provide structural interpretations of the difference-in-differences estimands in Section A.2.5.

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

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

where human capital may be augmented through chatbot use, $A_{ij} \in \{0, 1\}$:

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

and B_{ij} captures the productivity benefit of AI chatbots. As in Section A.1, productivity benefits and adoption costs depend on worker type and firm encouragement:

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

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

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 firm j's workforce. This formulation models chatbots as tools that augment worker output; Section A.2.1 discusses implications of more general technologies.

Product markets feature monopolistic competition among atomistic, horizontally differentiated firms facing demand curves:

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

where ϵ is the within-sector elasticity of demand.

Labor demand. With perfectly competitive labor markets, workers are paid their marginal revenue product:

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

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

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

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

Constant returns to scale and CES demand imply that optimal flow profit is $\Pi_j = \frac{1}{\epsilon} Y_j^{1-1/\epsilon}$, so the profit gain from encouragement equals:

$$\log \Pi_j(1) - \log \Pi_j(0) = (1 - \epsilon)(\log W_1 - \log W_0). \tag{18}$$

Since this gain is common across employers, encouragement follows a cutoff rule:

$$E_j = \mathbf{1}\left[\kappa_j \leq \kappa\right], \quad \text{with} \quad \kappa = (1 - \epsilon)(\log W_1 - \log W_0). \tag{19}$$

Adoption. Workers derive log utility from earnings and adopt chatbots if the monetary benefit exceeds the adoption cost. Fixing the employer, the adoption rule reduces to:

$$A_{ij} = \mathbf{1}\left[B_{ij} \geq C_{ij}\right] = \mathbf{1}\left[\hat{\alpha}_0 + \hat{\alpha}_1 E_j + \hat{\beta} \times i \geq 0\right], \tag{20}$$

where $\hat{x} = x_b - x_c$ denotes net benefits, yielding the cutoff:

$$A_{ij} = \mathbf{1}\left[i \leq i^*(E_j)\right], \quad \text{with} \quad i^*(E) = -\frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 E). \tag{21}$$

Sorting. Workers also choose which firms to sort into. Since all adopters receive the same net benefit $\hat{\alpha}_1$ from encouragement, workers effectively choose between sorting into encouraged firms and adopting, or sorting into non-encouraged firms and not adopting:

$$E_{j(i)} = \mathbf{1}\left[\hat{\alpha}_0 + \hat{\alpha}_1 + \hat{\beta} \times i + \log W_1 - \log W_0 \geq 0\right], \tag{22}$$

implying the cutoff:

$$E_{j(i)} = \mathbf{1}\left[i \leq i^*\right], \quad \text{with} \quad i^* = -\frac{1}{\hat{\beta}}(\hat{\alpha}_0 + \hat{\alpha}_1 + \log W_1 - \log W_0). \tag{23}$$

A.2.3 Equilibrium

An equilibrium consists of encouragement policies $\{E_j\}$, skill prices $\{W_j\}$, worker-firm allocations $\{\mathcal{I}_j\}$, and adoption decisions $\{A_{ij}\}$, such that firms and workers optimize (Equations (16), (17), (21), and (23)) and markets clear (Equations (14) and (15)). Without chatbots, all firms are symmetric and each employs an equal share of baseline human capital.

Sorting. When chatbots are available, the equilibrium features full segregation: assuming $\hat{\alpha}_1 \geq 0$, all adopters sort into encouraged firms and all non-adopters into non-encouraged firms, or both. We focus on the empirically relevant case where some but not all firms encourage and some but not all workers adopt.

Skill prices. From Equation (19), equilibrium skill price differentials satisfy:

$$\log W_1 - \log W_0 = \frac{\kappa}{1-\epsilon}, \tag{24}$$

where κ is the encouragement cost of the marginal firm. The cost of encouragement is thus priced into skill differentials: workers are not rewarded for these firm-level investments.

For analytical simplicity, we focus on the case where all employers face the same encouragement cost, $\kappa_j = \kappa$ for all j. All firms are then marginal, and the share of encouraged employers adjusts to match the relative supply of adopting workers' human

capital.[25]

A.2.4 Model Predictions

We now derive predictions for how chatbot arrival affects labor market outcomes.

Worker earnings. Combining Equations (10)–(11), the log change in earnings of worker i at firm j is:

$$d\log Y_{ij} = \underbrace{d\log W_j}_{\text{skill–price effect}} + \underbrace{A_i B_{ij}}_{\text{productivity effect}} \tag{25}$$

$$= d\log W_0 + E_j \frac{\kappa}{1-\epsilon} + A_i(\alpha_{b0} + \beta_b \times i) + E_j A_i \alpha_{b1}. \tag{26}$$

This yields three key predictions:

1. Non-adopters experience an earnings loss at encouraged firms:

$$\log Y_{i1}(A_i = 0) - \log Y_{i0}(A_i = 0) = -\frac{\kappa}{\epsilon - 1} \leq 0 \tag{27}$$

2. Holding encouragement fixed, adoption increases earnings by the productivity effect:

$$\log Y_{iE}(A_i = 1) - \log Y_{iE}(A_i = 0) = \alpha_{b0} + \alpha_{b1} \times E + \beta_b \times i \tag{28}$$

3. The effect of encouragement on adopter earnings is ambiguous, depending on the relative magnitudes of the productivity and skill price effects:

$$\log Y_{i1}(A_i = 1) - \log Y_{i0}(A_i = 1) = \alpha_{b1} - \frac{\kappa}{\epsilon - 1} \gtrless 0. \tag{29}$$

   Even if encouragement enhances productivity ($\alpha_{b1} > 0$, as documented in Section 2.2), these gains need not translate into higher earnings when encouragement is costly (high κ), demand is inelastic (low ϵ), or productivity gains are modest (low α_{b1}).

[25] With heterogeneous encouragement costs, skill price differentials also depend on the distribution of κ_j through the supply and demand for human capital across encouraged and non-encouraged firms. The key qualitative predictions are unchanged: encouragement is priced into skill differentials, and the sign of the earnings effect for encouraged adopters remains ambiguous.

Adopters may still prefer encouraged firms if encouragement sufficiently reduces adoption costs (i.e., $\alpha_{c1} < \alpha_{b1} - \frac{\kappa}{\epsilon-1}$), but these benefits will not appear in earnings.

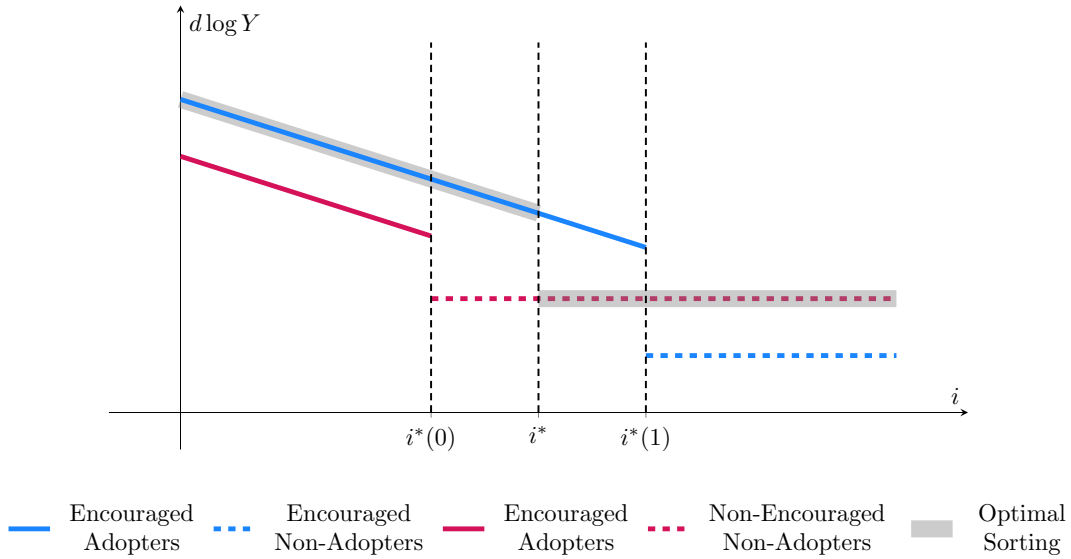
Occupational mobility. If chatbots create greater benefits B_{ij} for newcomers to an occupation, adoption will correlate with occupational switching (as documented in Section 3.3).[26]

Worker-firm sorting. Following Equation (23), chatbots trigger reallocation of workers who benefit from adoption (low-i types) into encouraged firms ($E = 1$).

Worker outcomes. Figure A.2 summarizes how chatbots affect sorting, adoption, and earnings. The gray line traces equilibrium patterns: workers with $i < i^*$ sort into encouraged firms and adopt, while those with $i > i^*$ sort into non-encouraged firms and do not adopt. The figure illustrates the case where encouragement increases adopter earnings; as shown above, encouragement can also reduce them, and if the amenity value of encouraged use is large enough (α_{c1} sufficiently low), some encouraged adopters may even earn less than non-encouraged non-adopters.

[26]This prediction reflects an effect running from occupational choice to adoption. Causality may also run the other way, with chatbots prompting moves into new occupations; our framework abstracts from this margin.

Figure A.2: 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 chatbot advantage i; the y-axis shows earnings impacts. Blue and red lines represent earnings effects for each worker-firm type combination. The gray line highlights equilibrium sorting and adoption patterns.

Workplace employment. The effect of encouragement on headcount employment (N) depends on whether increased human capital demand is offset by adoption-driven productivity gains. The increase in human capital demand is:

$$\log H_1 - \log H_0 = \frac{\epsilon}{\epsilon - 1}\kappa, \tag{30}$$

and the average productivity gain among encouraged adopters is:

$$\mathbb{E}[B_{ij} \mid i \leq i^*, E_j = 1] = (1-\gamma)(\alpha_{b0} + \alpha_{b1}) + \gamma(\alpha_{c0} + \alpha_{c1} + \kappa/(\epsilon - 1)), \tag{31}$$

where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$, generalizing Equation (9) to account for worker sorting. Combining these, the employment effect is:

$$\log N_1 - \log N_0 = \frac{\epsilon - \gamma}{\epsilon - 1}\kappa - (1-\gamma)(\alpha_{b0} + \alpha_{b1}) - \gamma(\alpha_{c0} + \alpha_{c1}). \tag{32}$$

Encouragement tends to *reduce* employment when encouragement is cheap (κ low), demand is elastic (ϵ high), workers are homogeneous in chatbot productivity (γ low), productivity

gains are large (α_b high), and adoption costs are high (α_c high).

Workplace wage bills. Encouraged firms increase wage bills by:

$$\log(W_1 H_1) - \log(W_0 H_0) = \kappa. \tag{33}$$

That encouragement unambiguously raises wage bills is a consequence of modeling chatbots solely as tools augmenting worker productivity (Equation (14)); as discussed in Section A.2.1, this need not hold under more general production technologies.

A.2.5 Empirical Identification

This section shows that causal effects of chatbots on worker and workplace outcomes can be identified using difference-in-differences, and provides structural interpretations of the estimands.

The derivations above assumed that differences across workers and firms arise solely from chatbot use and encouragement. More realistically, workers differ in baseline productivity ($H_{ij}(A_{ij}) = H_{i0}\exp(B_{ij}A_{ij})$) and firms in baseline demand ($Y_j = Y_{j0}P_j^{-\epsilon}$). For example, Figure 5(a) shows that adopters already earned more before chatbots became available.

Causal effects are still identified if baseline heterogeneity is time-invariant and log-additive, so that it differences out:

$$\Delta H_{ij} = B_{ij}(A_{ij}^{post} - A_{ij}^{pre}), \qquad \Delta Y_j = -\epsilon(\log P_j^{post} - \log P_j^{pre}), \tag{34}$$

where $\Delta X = \log X^{post} - \log X^{pre}$. If baseline heterogeneity varies over time, identification requires that adoption and encouragement decisions are uncorrelated with anticipated shocks to H_{i0} and Y_{j0}—consistent with the policy rules in Equations (17) and (20).

Worker-level difference-in-differences. Equation (25) motivates three estimators used in Section F.2:

1. **Productivity effect among non-encouraged adopters** (*Non-encouraged adopters vs. non-encouraged non-adopters*)

$$\Delta^{Y}_{\text{NE,A} - \text{NE,NA}} = \gamma\alpha_{c0} + (1-\gamma)\alpha_{b0}, \tag{35}$$

where $\gamma = \frac{\beta_b}{2\hat{\beta}} \geq 0$.[27]

2. **Skill-price effect in encouraged firms** (*Encouraged non-adopters vs. non-encouraged non-adopters*)

$$\Delta^{Y}_{\text{E,NA} - \text{NE,NA}} = -\kappa/(\epsilon - 1). \tag{36}$$

3. **Combined effects for encouraged adopters** (*Encouraged adopters vs. non-encouraged non-adopters*)

$$\Delta^{Y}_{\text{E,A} - \text{NE,NA}} = -\kappa/(\epsilon - 1) + \gamma(\alpha_{c0}+\alpha_{c1}) + (1-\gamma)(\alpha_{b0}+\alpha_{b1}). \tag{37}$$

With perfect worker mobility, Equation (31) implies an additional $\gamma\kappa/(\epsilon - 1)$ term.

Workplace-level difference-in-differences. Section 3.4 estimates the effect of encouragement on workplace outcomes:

1. **Employment effect:**

$$\Delta^{N}_{\text{E} - \text{NE}} = \frac{\epsilon - \gamma}{\epsilon - 1}\kappa - (1-\gamma)(\alpha_{b0} + \alpha_{b1}) - \gamma(\alpha_{c0} + \alpha_{c1}). \tag{38}$$

2. **Wage bill effect:**

$$\Delta^{WH}_{\text{E} - \text{NE}} = \kappa. \tag{39}$$

A.2.6 Missing Intercept

The estimators in Equations (35)–(37) identify earnings effects up to the skill price effect in non-encouraged firms, $d\log W_0$—the familiar missing intercept problem, since within-sector

[27] The frictionless sorting model predicts all adopters sort into encouraged firms, leaving no workers to identify Equation (35). In practice, Section 3 finds little worker resorting, so Equation (35) derives the structural interpretation without mobility.

difference-in-differences cannot capture sector-wide general equilibrium effects.

We use the model to characterize this residual. The key insight is that adopters generate spillovers to non-adopters only when their own market outcomes change; null effects on adopters' market outcomes thus imply limited spillovers.

To formalize, assume nested CES demand with elasticity η across and ϵ within sectors:

$$Y_j = P_j^{-\epsilon} P^{\epsilon-\eta}, \tag{40}$$

where $P = (\int P_j^{1-\epsilon} dj)^{\frac{1}{1-\epsilon}}$ is the sectoral price index. Combining with Equations (11), (14), (16), and (40):

$$d\log W_0 = \underbrace{-\frac{1}{\epsilon} d\log Y_0}_{\text{Residual demand}} + \underbrace{\frac{\eta-\epsilon}{\epsilon\eta} d\log Y}_{\text{Sector competition}}, \tag{41}$$

where the first term reflects firms moving down residual demand curves and the second reflects competing firms stealing sectoral demand.

The change in non-encouraged output reflects both productivity and employment:

$$d\log Y_0 = \underbrace{\mathbb{E}[A_{i0}B_{i0}]}_{\text{Adopter productivity}} + \underbrace{\mathbb{E}\left[d\log N_0\right]}_{\text{Employment}}, \tag{42}$$

where these terms are identified by:

$$\mathbb{E}[A_{i0}B_{i0}] = A_0 \times \Delta^Y_{\text{NE,A}-\text{NE,NA}}, \tag{43}$$

$$\mathbb{E}\left[d\log N_0\right] = -E \times \Delta^N_{\text{E}-\text{NE}}, \tag{44}$$

with A_0 the adoption rate in non-encouraged firms, E the share of encouraging firms, and Equation (44) using the labor market clearing condition. Similarly, total output growth reflects aggregate adopter productivity:

$$d\log Y = (1-E)A_0 \cdot \Delta^Y_{\text{NE,A}-\text{NE,NA}} + EA_1 \cdot \Delta^Y_{\text{E,A}-\text{E,NA}}, \tag{45}$$

since employment effects cancel in aggregate. Combining:

$$\mathbb{E}[d\log W_0] = -\frac{1}{\epsilon}\left[A_0 \cdot \Delta^Y_{\text{NE,A}-\text{NE,NA}} - E \cdot \Delta^N_{\text{E}-\text{NE}}\right] \tag{46}$$

$$+ \frac{\eta-\epsilon}{\epsilon\eta}\left[(1-E)A_0 \cdot \Delta^Y_{\text{NE,A}-\text{NE,NA}} + EA_1 \cdot \Delta^Y_{\text{E,A}-\text{E,NA}}\right], \tag{47}$$

where all terms are identified by existing difference-in-differences estimates, up to calibrated demand elasticities η and ϵ. Crucially, when both Δ^N and Δ^Y are approximately zero (as in Figures 5 and 8), this implies that the missing intercept $\mathbb{E}[d\log W_0]$ is also close to zero.

B Data Construction

B.1 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.2 Survey Sample

Table B.1 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.1: 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.2.1 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.2 shows that our survey respondents resemble our survey populations on observable characteristics.

Table B.2: 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.1. The *Responded* columns correspond to line 4 of Table B.1. *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.3: 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.1. Standard deviations are shown in parentheses. See the note of Table B.2 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.4 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.4: 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.1. See the note of Table B.2 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 workers' reported occupations against those recorded in the administrative registers. Table B.5 shows that the survey and registers agree on the occupation of 87% of our respondents.

Table B.5: 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.5 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.5 shows that 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.

B.3 Classification of New Job Tasks

Our survey includes free-text responses about the new tasks workers have received due to AI chatbots. We categorize these responses into six broad AI-related categories, listed in Table B.6, as well as into more granular, occupation-specific subtasks. The task categories fall into three groups: *Generation* of new content, *Oversight* of AI outputs, and *Integration* of AI tools.

Table B.6: Categories of New Tasks from AI Chatbots

Group	**Task**	**Description**
Generation	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 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.
Oversight	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 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.
Integration	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.

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 documents (*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*)

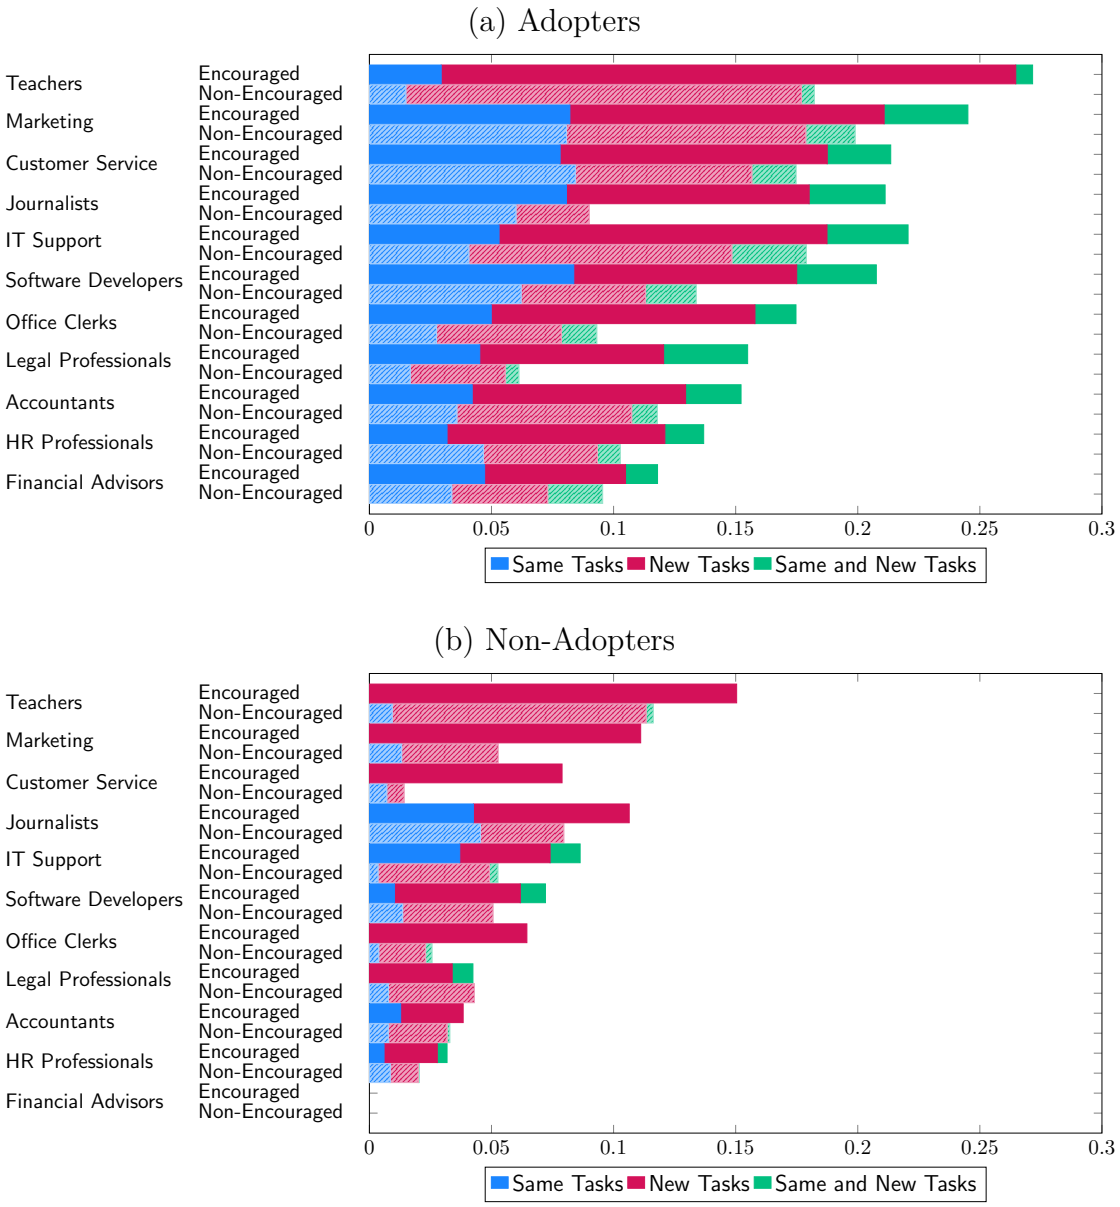
10. **Software Developers:** Using AI to generate code snippets, boilerplate code, or documentation (*AI Content Drafting*), Reviewing and testing AI-generated code to ensure correctness and security (*AI Quality Review*), Formulating specific prompts to guide AI in debugging or coding tasks (*AI Integration*), 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.

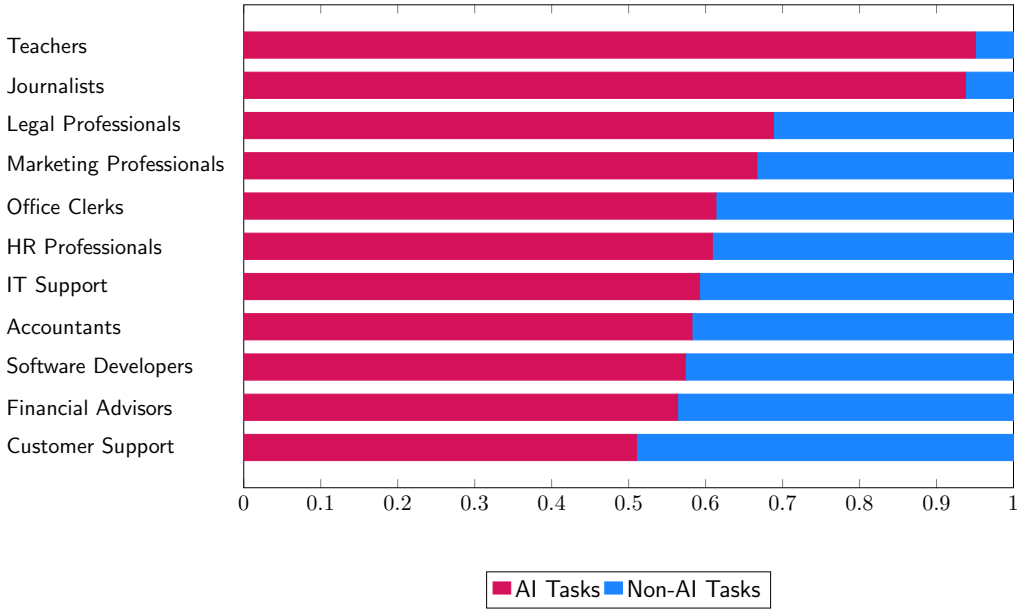
Figures B.1 illustrate that AI chatbots have led to the creation of new tasks across all 11 occupations in our study. Figure B.2 shows that 53% to 91% of these tasks are directly linked to AI use. Figure 4 in the main text breaks down the composition of AI-related job tasks by occupation and task category.

Figure B.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 B.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.

C Institutional Setting

C.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, 2025; 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 C.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.[28] To provide an indication of expected coverage, Columns (1)-(3) of Table C.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).

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

Table C.1: Wage Setting Systems in Denmark

	Proxies for CBA coverage			Wage setting systems under CBAs		
	Public sector	Large private firms	Total (1+2)	Central	TwoTier	Decentral
Accountants	.055	.729	.784	.092	.059	.839
Customer Service	.09	.826	.916	.228	.447	.304
Financial Advisors	.019	.942	.961	.008	.909	.08
HR Prof.	.395	.548	.943	.156	.056	.781
IT Support	.211	.681	.892	.124	.171	.697
Journalists	.416	.544	.96	0	.54	.457
Legal Prof.	.561	.379	.94	.043	.002	.954
Marketing	.023	.851	.874	.127	.152	.711
Office Clerks	.493	.385	.878	.351	.048	.583
Software Dev.	.005	.766	.771	0	0	.999
Teacher	.787	.148	.935	.956	0	.043
All	.278	.618	.896	.19	.217	.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).

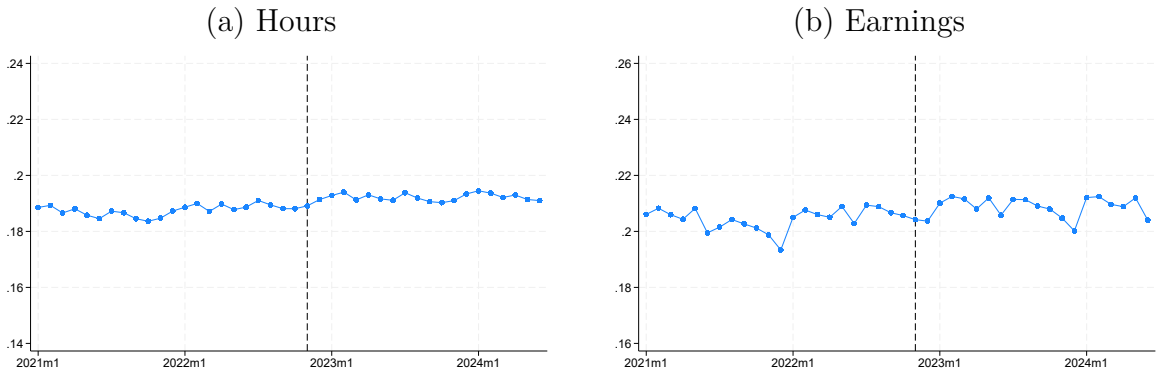
D Labor Market Trends

In this appendix, we present aggregate employment trends for our 11 highly-exposed study occupations.

D.1 Aggregate Employment

Figure D.1 plots the share of our 11 survey occupations in aggregate Danish employment. These occupations have tracked the rest of the economy, maintaining a stable 20% share of hours and earnings before and after the arrival of AI chatbots. This pattern aligns with Kauhanen and Rouvinen (2025), who use Finnish population data to show that occupations exposed to Generative AI have not experienced differential changes in hours or earnings. Similarly, occupations exposed to generative AI have not experienced differential changes in aggregate employment in the United States Brynjolfsson, Chandar and Chen (2025); Chandar (2025); The Budget Lab at Yale (2026).

Figure D.1: Share of Study Occupations in Aggregate Danish Employment



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.

D.2 By Age Groups (Brynjolfsson, Chandar and Chen, 2025)

In this section, we break down labor market trends by occupation and age group. Brynjolfsson, Chandar and Chen (2025) document that the employment of early-career workers (age 22-25) in AI-exposed occupations has fallen sharply in the United States since mid 2022. We use our population-wide administrative data to replicate these facts in the Danish market.

Table D.1 shows the employment headcount in each occupation and age group in October 2022. We use the same age groups as Brynjolfsson, Chandar and Chen (2025): early career 1 (age 22-25), early career 2 (age 26-30), developing (31-34), mid career 1 (35-40), mid career 2 (41-49), and senior (50+). Most occupational employment is concentrated in older categories, with the exception of customer support, which employs many early-career workers.

Figure D.2 shows the employment trends for each of the occupations and age groups, indexed to October 2022 as in Brynjolfsson, Chandar and Chen (2025). Panels (a) and (b) show aggregate trends inside and outside our 11 study occupations, with the remaining panels showing trends for each study occupation separately. The plots show that we replicate the declines in early-career jobs in many exposed occupations, including software developers, legal professionals, marketing professionals, IT support specialists,

journalists, HR professionals, and financial advisors. Importantly, however, our difference-in-differences analysis in Section 3.4 shows that the declines are not driven by firms adopting Generative AI.

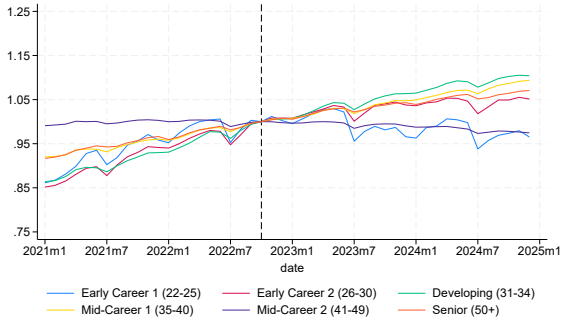
Table D.1: Employment by Occupation and Age Groups (October 2022)

Occupation	Early Career 1 (22-25)	Early Career 2 (26-30)	Developing (31-34)	Mid-Career 1 (35-40)	Mid-Career 2 (41-49)	Senior (50+)
IT Support	1157	2358	1800	2140	2804	3911
Teachers	4825	9253	8349	11925	23283	26562
Office Clerks	16954	17508	11632	16572	29232	54354
HR Prof.	399	970	919	1582	2897	3076
Financial Adv.	839	2361	1824	2845	3782	6602
Marketing	1470	4088	3201	3588	4767	4929
Customer Service	1625	1404	783	779	933	1148
Accountants	2227	4627	3318	4276	6566	12270
Legal Prof.	1479	5362	4080	4313	4735	4269
Software Dev.	2840	9748	7901	9104	12441	14862
Journalists	532	1616	1319	1623	2165	2365

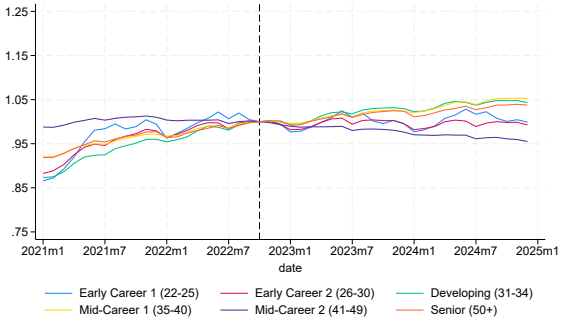
Notes: Table D.1 reports employment headcounts by occupation and age group in October 2022. We follow Brynjolfsson, Chandar and Chen (2025) in defining the age groups: early career 1 (ages 22–25), early career 2 (26–30), developing (31–34), mid-career 1 (35–40), mid-career 2 (41–49), and senior (50+). *Sample*: All registry data.

Figure D.2: Employment Trends by Occupation and Age Groups

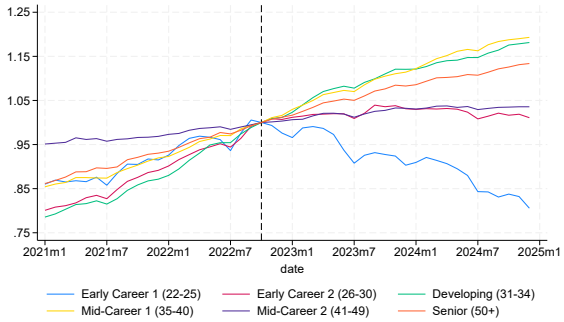
(a) Study Occupations (Exposed)



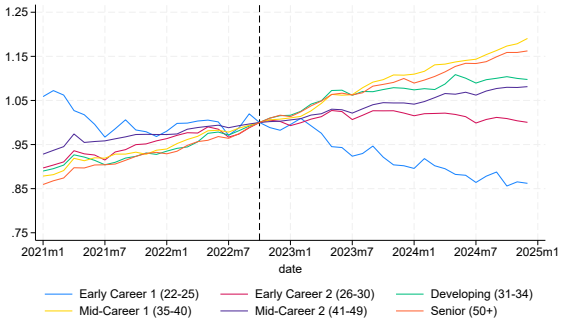
(b) Other Occupations (Non-Exposed)



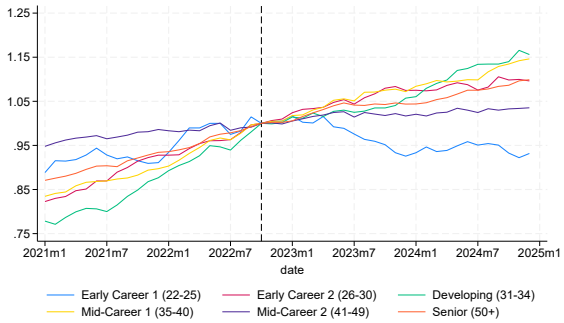
(c) Software Developers



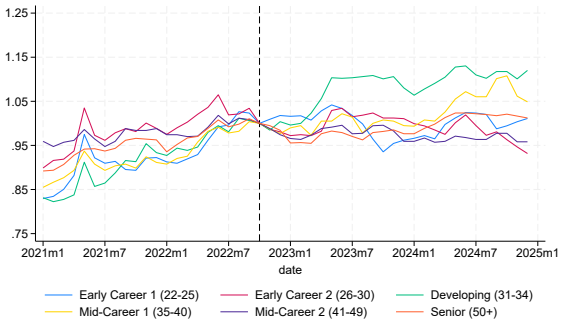
(d) Legal Professionals



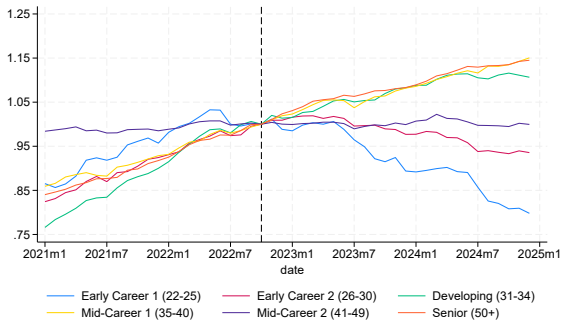
(e) IT Support



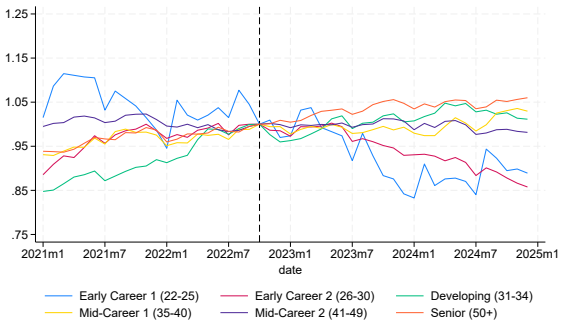
(f) Customer Service



(g) Marketing Professionals



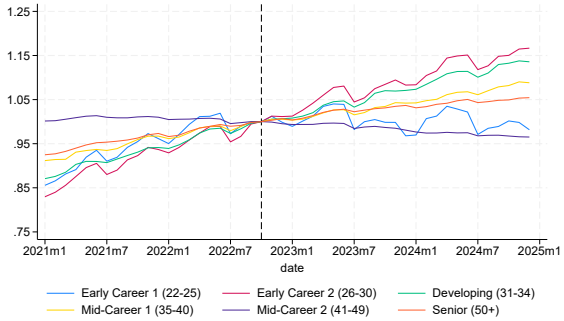
(h) Journalist



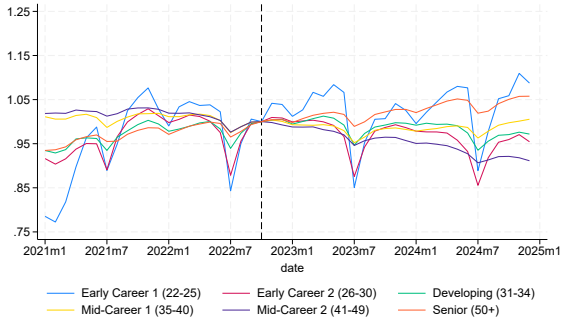
Notes: This figure is continued on the next page.

Figure D.2: Employment Trends by Occupation and Age Groups (Continued)

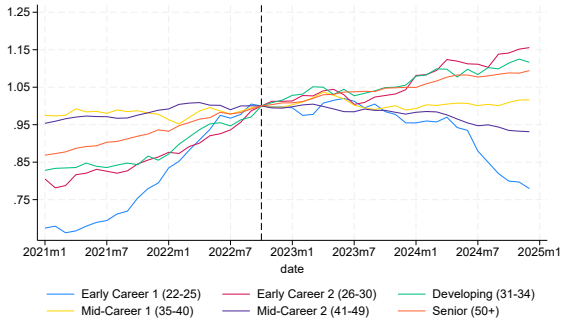
(i) Office Clerks



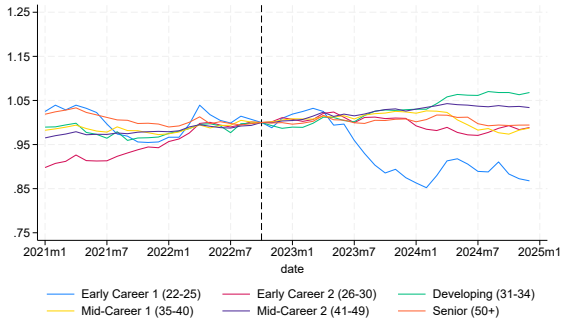
(j) Teachers



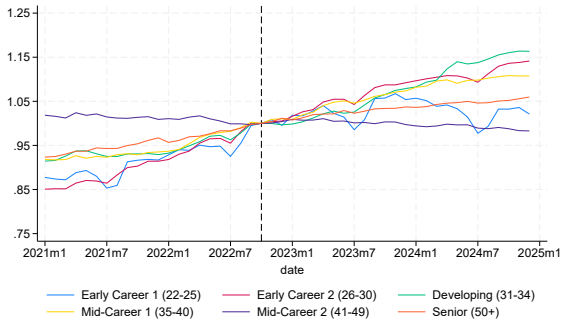
(k) HR Professionals



(l) Financial Advisors



(m) Accountants



Notes: Figure D.2 shows the employment trends for each of the occupations and age groups, indexing these to October 2022. Panels (a) and (b) show aggregate trends inside and outside our 11 study occupations, with the remaining panels showing trends for each study occupation separately. We follow Brynjolfsson, Chandar and Chen (2025) in defining the age groups: early career 1 (ages 22–25), early career 2 (26–30), developing (31–34), mid-career 1 (35–40), mid-career 2 (41–49), and senior (50+). We interpolate occupational codes within job-spells (worker-workplace pairs) to minimize measurement error resulting from the reclassification of occupational codes (Groes, Kircher and Manovskii, 2015; Humlum, 2021). *Sample*: All registry data.

E Additional Results

E.1 Employer Initiatives

Table E.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 F.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.

E.2 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.

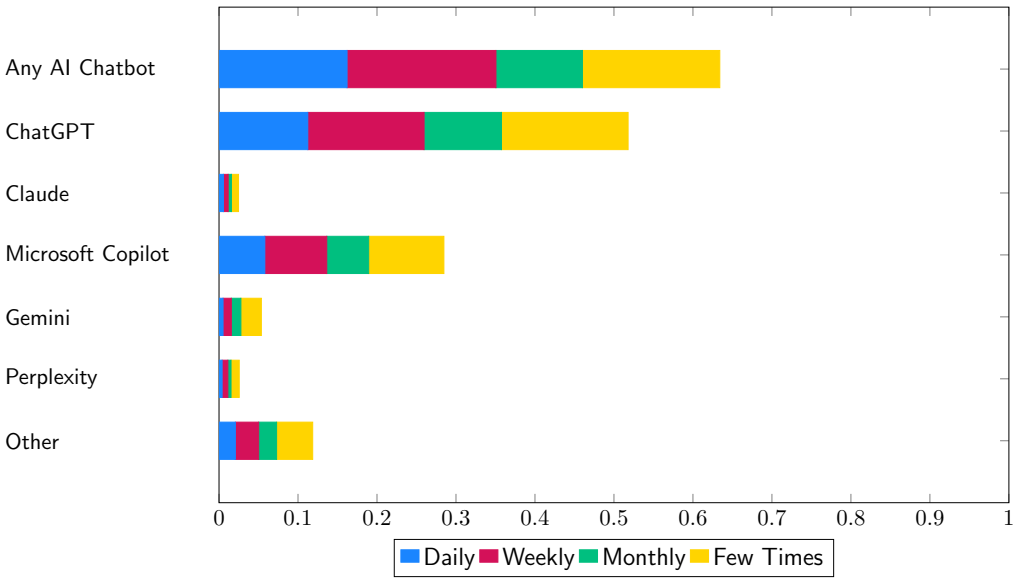
Figure E.1 shows that ChatGPT remains the dominant tool: approximately 80% of all chatbot adopters use ChatGPT, and its dominance holds across all occupations.

Table E.2 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

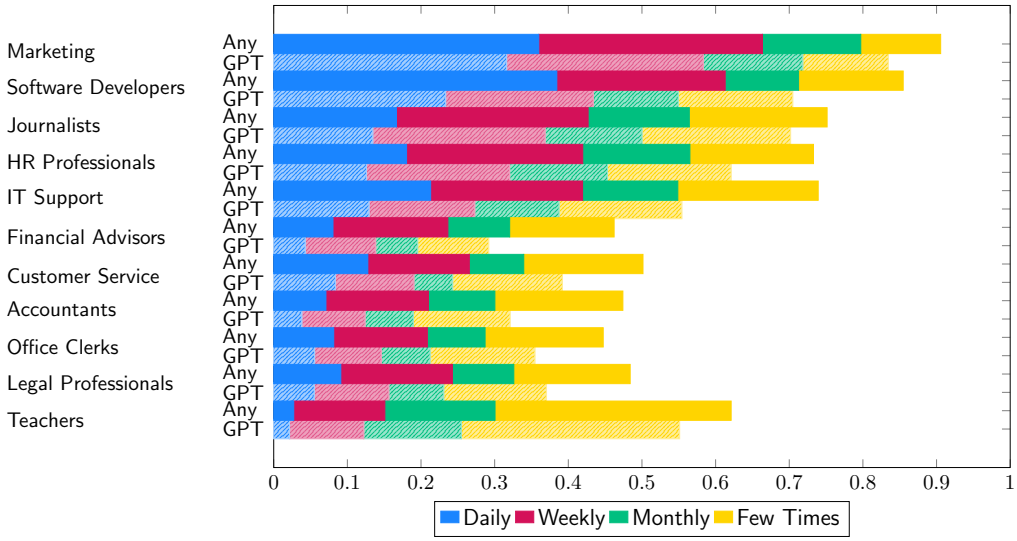
use through features such as data security compliance, training on internal data, and custom prompt creation.

Figure E.1: The Adoption of AI Chatbot Products

(a) Prevalence of AI Chatbot Products



(b) The Dominance of ChatGPT



Notes: Panel (a) displays the share of workers who have used various AI chatbots for work, categorized by frequency of usage. Panel (b) 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 E.2: 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.

E.3 Reported Benefits by Occupation

Table E.3 shows how reported benefits from using AI chatbots vary across occupations. Marketing professionals are more than twice as likely as teachers to report improved work quality (69.8% vs. 32.1%), and software developers are more than twice as likely as journalists to report higher job satisfaction (30.5% vs. 12.6%). Across all exposed occupations, however, 64–90% of users report time savings. This occupational heterogeneity supports the notion of a "jagged frontier" of AI chatbot costs and benefits (Dell'Acqua et al., 2024).

What is the economic significance of these reported benefits? In Column (6), we combine workers' frequency of use with their reported time savings per day of usage to estimate time savings as a percent of total work hours.[29] On average, adopters in our sample report savings of about 3% of their work hours. These estimates align with the nationally representative U.S. survey evidence in Bick, Blandin and Deming (2025).

What do workers do with these savings? Columns (7)–(10) show that the large majority (85%) of chatbot users reallocate saved time to other job tasks. By contrast, fewer than 10% report taking additional breaks or leisure, while about 30% spend more time on the same tasks from which they initially saved time. This pattern underscores the importance of task-based reorganization for understanding the effects of AI chatbots on work (Autor and Thompson, 2025; Freund and Mann, 2025).

Table E.3: Reported Benefits from AI Chatbots (Ever Used), by Occupation

	Reported Benefits						Allocation of Time Savings			
Occupation	Any Time Savings (1)	60+ min/day (2)	Quality (3)	Creativity (4)	Job Satisfaction (5)	Total Time Savings (%) (6)	Same Tasks (7)	Diff. Tasks (8)	Breaks (9)	Leisure (10)
Journalists	.67	.107	.394	.467	.126	2.5	.291	.846	.049	.072
Software Dev.	.838	.266	.528	.446	.305	5.6	.368	.837	.079	.082
Legal Prof.	.783	.178	.544	.45	.204	2.6	.306	.846	.082	.056
Accountants	.709	.086	.524	.418	.156	1.6	.289	.821	.067	.029
Customer Service	.723	.117	.572	.413	.142	3.0	.371	.776	.092	.078
Marketing	.898	.334	.698	.625	.289	6.1	.337	.858	.087	.079
Financial Adv.	.769	.084	.55	.488	.173	1.9	.304	.869	.033	.019
HR Prof.	.844	.185	.638	.623	.253	3.5	.219	.899	.094	.068
Office Clerks	.689	.134	.563	.489	.185	2.2	.288	.861	.05	.046
Teachers	.637	.081	.321	.458	.132	0.7	.152	.845	.095	.043
IT Support	.76	.182	.526	.423	.239	3.5	.313	.862	.063	.069
All	.756	.159	.532	.482	.2	3.0	.294	.847	.072	.058

Notes: This table presents the share of adopters who report various benefits from using AI chatbots for work, broken down by occupation. Columns (1)–(6) focus on workers who have ever used AI chatbots for work. Columns (7)–(10) further restrict to adopters who reported any time savings—that is, those coded as 1 in Column (1). *Sample*: All completed responses from the 2024 survey round.

[29]We code daily time savings 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 as follows: daily use is divided by 1; weekly use by 5 (corresponding to five 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 four uses over 12 months). We set daily work hours to 8, as nearly all sampled workers are full-time employees.

E.4 Workloads and Task Creation

Figure E.2 shows that workers who report greater time savings are also more likely to report new job tasks from chatbots.

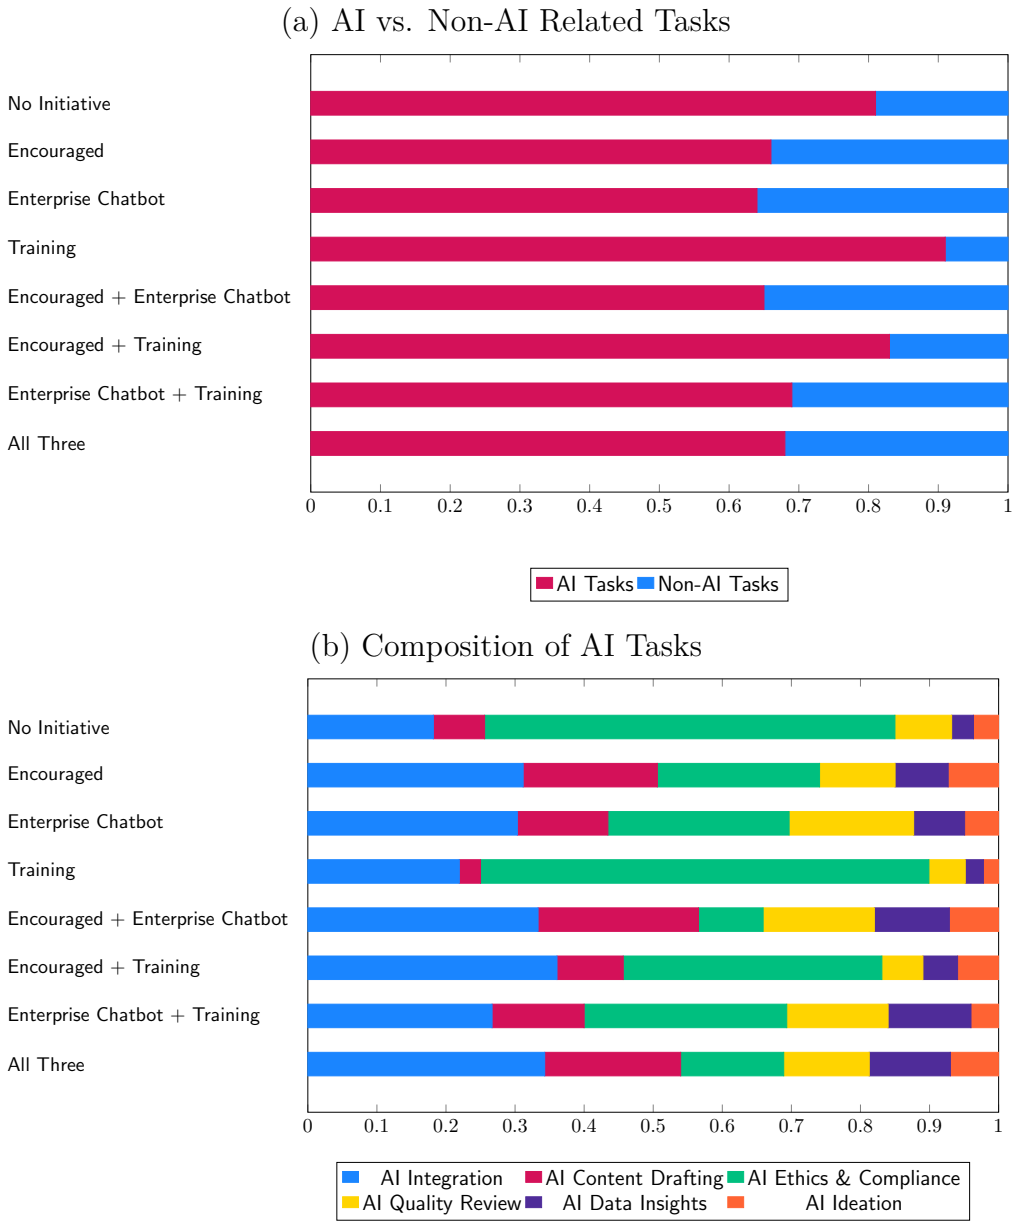
Figure E.3 shows the composition of new job tasks from AI chatbots by employer initiatives. Ethics and compliance tasks are especially common among workers who have received employer-provided training in AI chatbots. In contrast, workplaces that merely encourage chatbot use have a higher share of "productive AI tasks" focused on content generation and ideation.

Figure E.2: New Work and Reported Time Savings



Notes: This figure presents a binned scatterplot of workers' reported new tasks 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.

Figure E.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 by employer initiative. *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 B.3 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.

E.5 Perceived Impacts

Table E.4 reports workers' average perceived earnings impacts from AI chatbots, broken down by occupation, employer encouragement, and adoption status. Columns (1) and (3) report average impacts, while Columns (2) and (4) report the share of workers who perceive no change in earnings from AI chatbots.

Table E.4: Perceived Earnings Impacts of AI Chatbots

		Adopters		Non-Adopters	
Occupation	Usage Policy	Average Impact (%) (1)	No Impact (%) (2)	Average Impact (%) (3)	No Impact (%) (4)
Journalists	Non-Encouraged	-.2	99	0	100
Journalists	Encouraged	-.171	98.1	0	100
Software Developers	Non-Encouraged	.152	97.2	-.078	98.8
Software Developers	Encouraged	.466	95.8	-.085	99.2
Legal Professionals	Non-Encouraged	.09	99.2	.003	99.9
Legal Professionals	Encouraged	.07	98.7	-.04	98.4
Accountants and Auditors	Non-Encouraged	.084	98.1	-.044	99.7
Accountants and Auditors	Encouraged	.158	97.6	0	100
Customer Service Rep.	Non-Encouraged	.015	97	-.189	98.5
Customer Service Rep.	Encouraged	.378	94.8	0	100
Marketing Professionals	Non-Encouraged	.115	96.5	.099	97.4
Marketing Professionals	Encouraged	.386	95.5	0	100
Financial Advisors	Non-Encouraged	.112	98.9	-.01	99.6
Financial Advisors	Encouraged	.269	97.4	0	100
HR Professionals	Non-Encouraged	.093	98.1	-.143	99.3
HR Professionals	Encouraged	.08	98.7	0	100
Office Clerks	Non-Encouraged	.048	98.9	-.001	99.8
Office Clerks	Encouraged	.168	98.1	-.074	99.4
Teachers	Non-Encouraged	-.009	99.8	-.005	99.8
Teachers	Encouraged	.047	99.1	0	100
IT Support	Non-Encouraged	.03	99	.053	98.9
IT Support	Encouraged	.195	97.2	0	100
All	Non-Encouraged	.048	98.3	-.029	99.2
All	Encouraged	.186	97.4	-.018	99.7

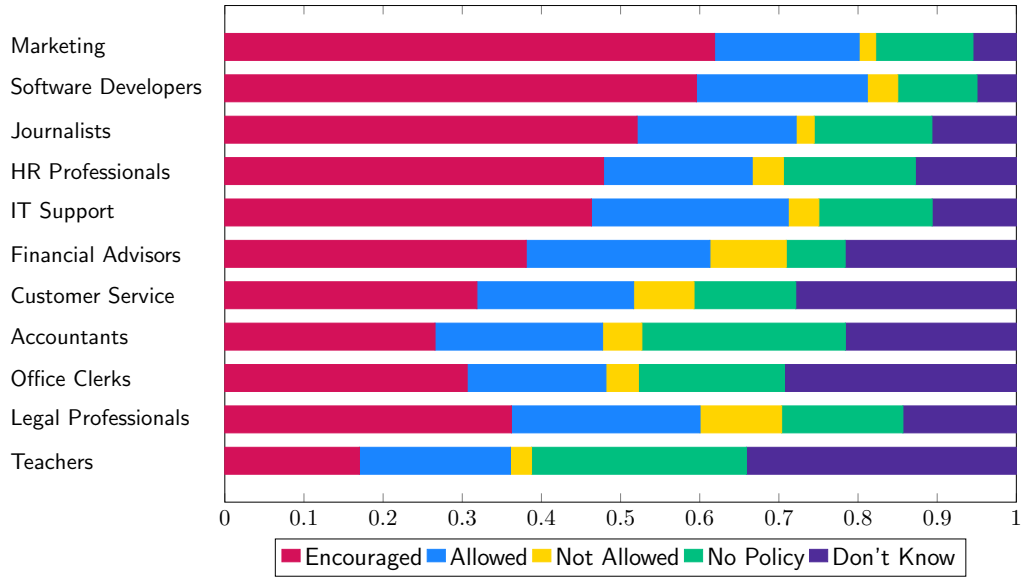
Notes: Columns (1) and (3) report the average perceived impact of AI chatbots on earnings, broken down by workers' occupation, employer encouragement, and adoption status (ever versus never used AI chatbots). Columns (2) and (4) show the share of workers who report no impact on earnings from AI chatbots. 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.

F Robustness Analysis

F.1 Employer Initiatives

Figure F.1 shows that the prevalence of employer initiatives (Figure 1, Panel (a)) is similar at the workplace level. Table F.1 shows that the correlations between employer initiatives and firm characteristics (Table E.1) are robust to controlling for worker characteristics.

Figure F.1: Employer Usage Policies (Workplaces)



Notes: This figure shows the share of workplaces affected by various employer usage initiatives for AI chatbots. Workers in our sample have been reweighted so all workplaces have the same weight. *Sample*: All completed responses from the 2024 survey.

Table F.1: Which Workplaces Have Adopted AI Chatbot Initiatives? (Worker Controls)

	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)	-0.0135**	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, occupation fixed effects, and worker pre-determined 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.

F.2 Worker Earnings

Figure F.2 shows that the null results for adopter earnings hold separately for workers' hourly wages, intensive-margin hours, and extensive-margin employment.

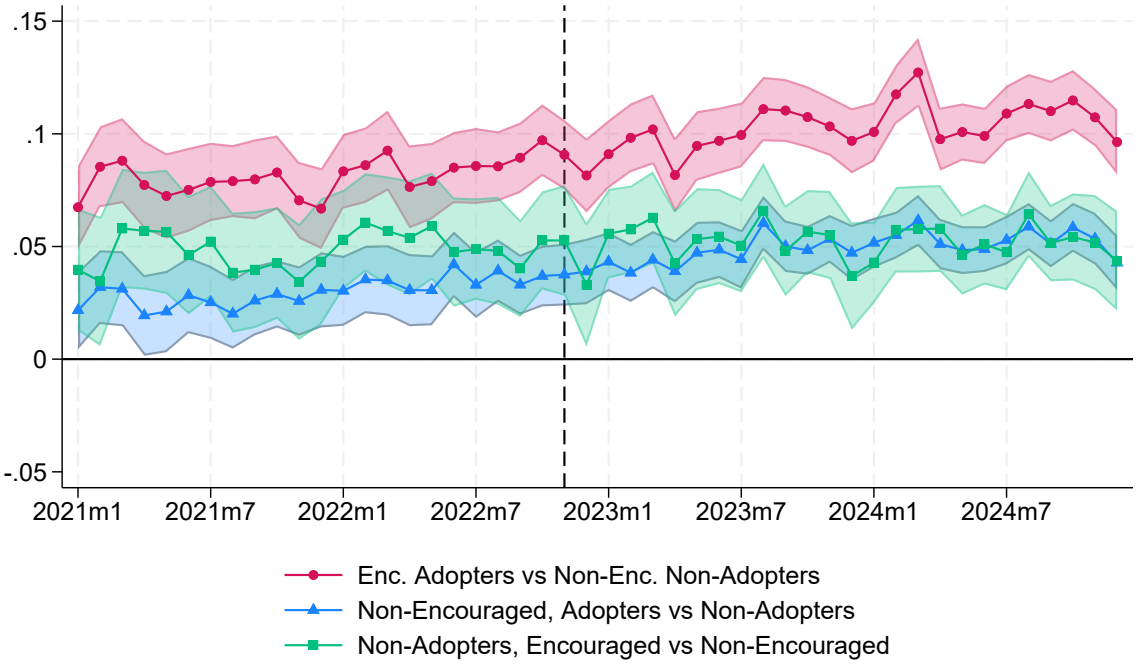
Figure F.3 adds non-adopters in encouraged workplaces to the main earnings plot (Figure 5, Panel (a)), showing they have also fared similarly to their counterparts in non-encouraged workplaces following the introduction of AI chatbots.

Figure F.2: Have Adopting Workers Experienced Earnings Gains?
(Log Earnings Relative to Non-Encouraged Non-Adopters)



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.

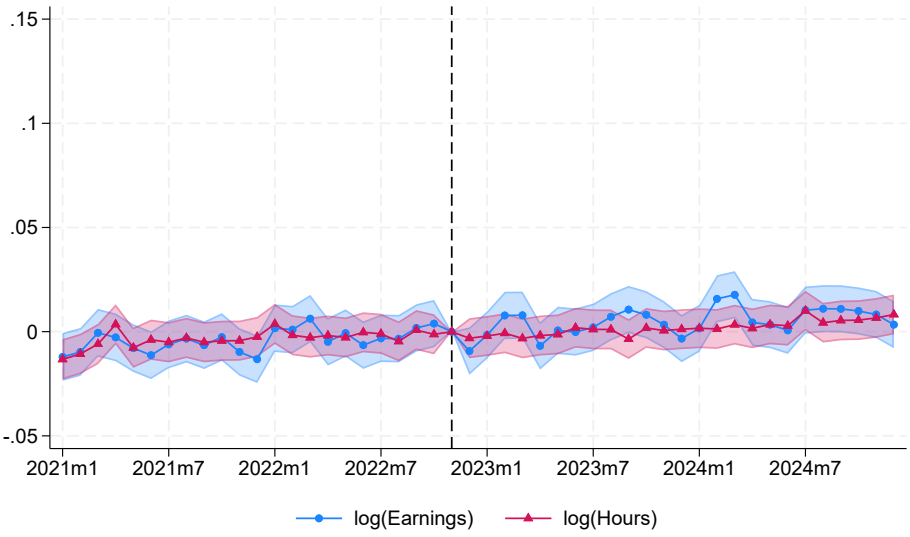
Figure F.3: Have Adopting Workers Experienced Earnings Gains?
(Log Earnings Relative to Non-Encouraged Non-Adopters)



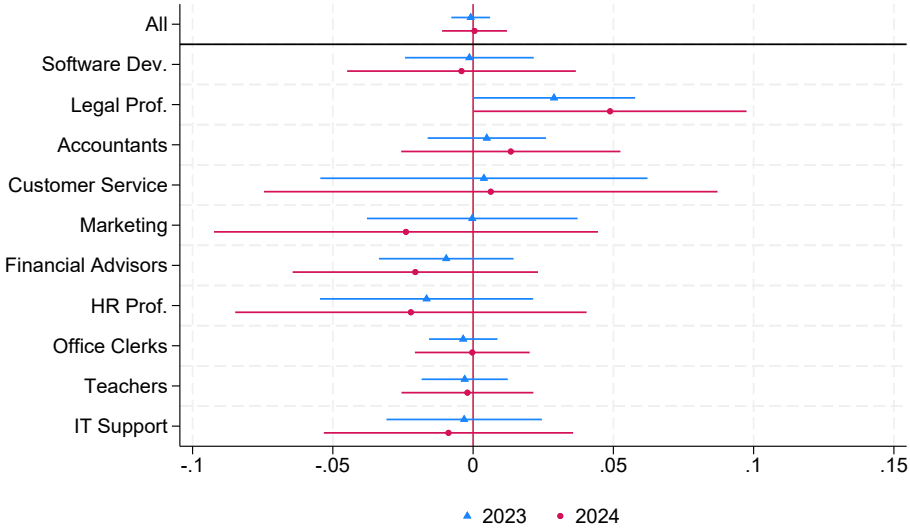
Notes: This figure shows the earnings gap of workers according to their encouragement and adoption status. Each group is compared to non-encouraged non-adopters. The figure reports the difference in log earnings, controlling for occupation fixed effects and pre-determined worker characteristics. Shaded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

Figure F.4: Have Encouraged Workers Fared Differently?

(a) Log Earnings and Log Hours



(b) Occupational Heterogeneity: Log Earnings



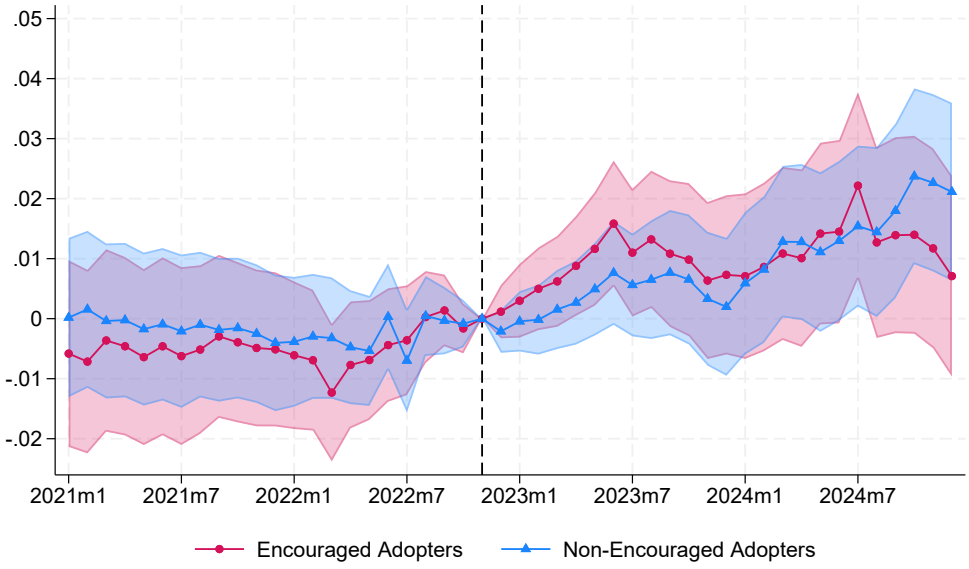
Notes: This figure presents the differential labor market outcomes of workers who are encouraged to use AI chatbots by their employers, compared to all other workers, indexed to the launch of ChatGPT in November 2022. Panel (a) is based on the dynamic difference-in-differences specification in Equation (2), with shaded areas representing 95% confidence intervals. The corresponding "first-stage" effect on adoption is 0.363 (0.006). Panel (b) first shows a pooled estimate (*All*) and then reports impacts separately for each of our 11 study occupations. These effects are based on the pooled difference-in-differences specification in Equation (3), with whiskers representing 95% confidence intervals. *Sample:* All completed responses from the 2024 survey linked to registry data.

F.3 Job Mobility

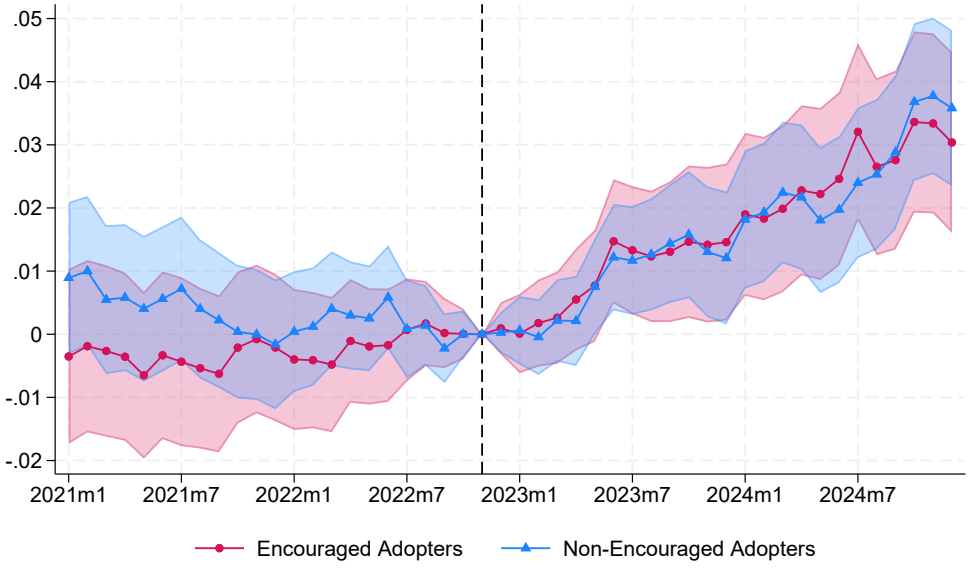
Our main analysis in Section 3.3 focuses on occupational mobility. In addition, AI chatbots may also induce mobility across workplaces. For instance, if chatbots help workers move up firm-specific learning curves, adopters may be more likely to switch employers. Alternatively, if encouraged-use policies enhance the utility of chatbots (as shown in Section 2.2), adopters may seek to sort into such workplaces. This prediction is formalized in our theoretical framework in Appendix A.2. Despite these possibilities, Figure F.5 shows only a weak link between chatbot adoption and workplace mobility (Panel a). Moreover, job mobility effects—across either workplaces or occupations—are not stronger among adopters whose use is encouraged by their endline employers (Panel b).

Figure F.5: Have Adopters Experienced Greater Job Mobility?
(DiD to Non-Encouraged Non-Adopters)

(a) FTE Employment in Latest Workplace



(b) FTE Employment in Latest Occupation



Notes: This figure presents difference-in-differences estimates of work hours in workers' latest jobs. Panel (a) focuses on workplaces, while Panel (b) focuses on occupations. 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.

F.4 Workplace Employment

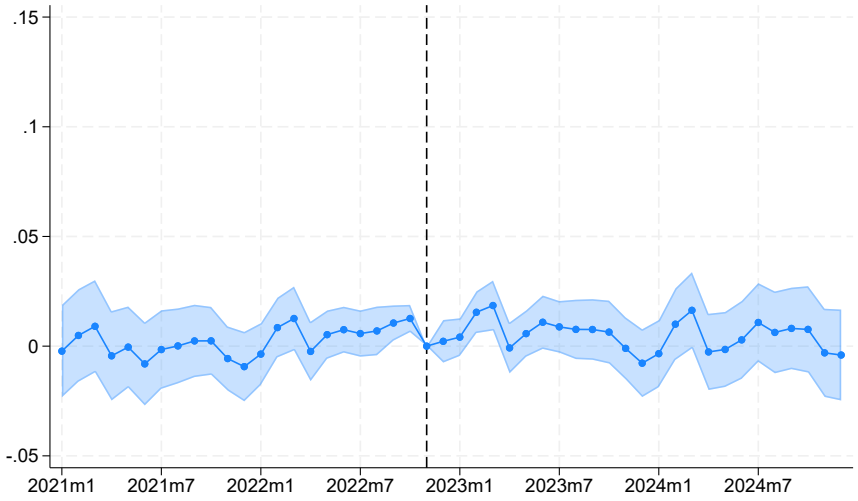
Figure F.6, Panel (a) shows that workplaces that encourage chatbot usage have seen no differential change in wage bill after the arrival of AI chatbots. While our main analysis focuses on employment within workers' respective occupations, Panel (b) shows that overall workplace employment remains flat as well.

Figure F.7 shows that incumbent worker hours and job churn have not changed differentially at workplaces with additional investments in enterprise chatbots and training. Figure F.8 shows that the null results for incumbent workers and job churn hold in each of our 11 occupations.

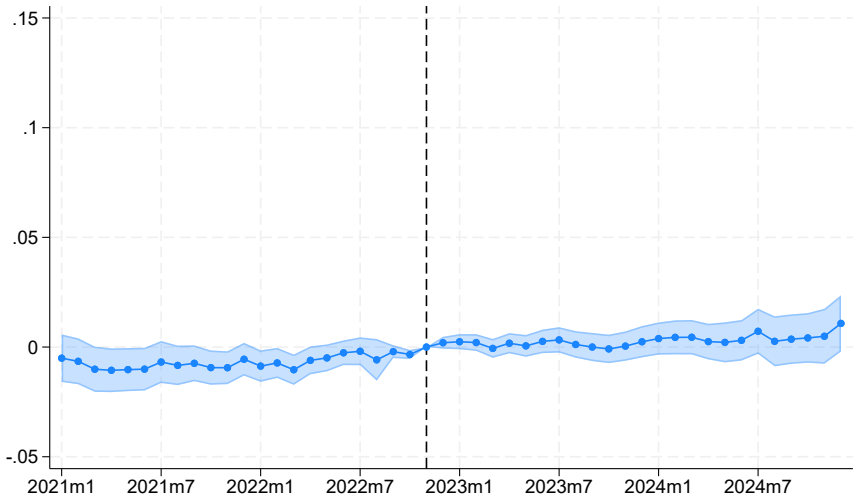
Figure F.9 shows that incumbent workers have also not experienced differential occupational mobility.

Figure F.6: Have Adopting Workplaces Fared Differently? (Encouraged DiD)

(a) Log Occupational Wage Bill



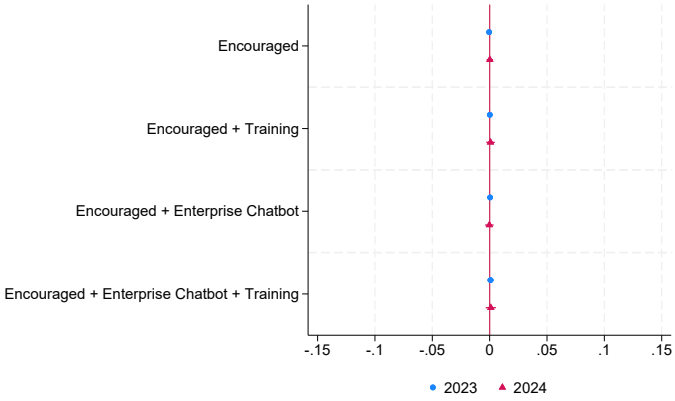
(b) Log Total Employment



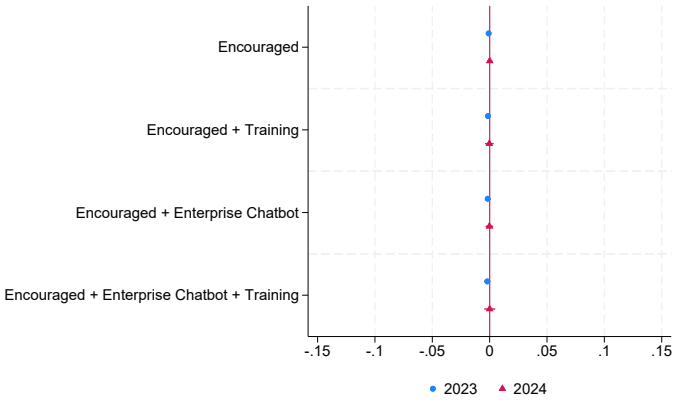
Notes: This figure shows the occupational wage bill (Panel a) and total workplace employment (Panel b) in encouraged-use workplaces relative to those without encouragement, indexed to November 2022. Estimates are based on the dynamic difference-in-differences specification in Equation (2), with shaded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

Figure F.7: Have Adopting Workplaces Fared Differently?
Heterogeneity by Adoption Intensity (DiD to Non-Encouraged)

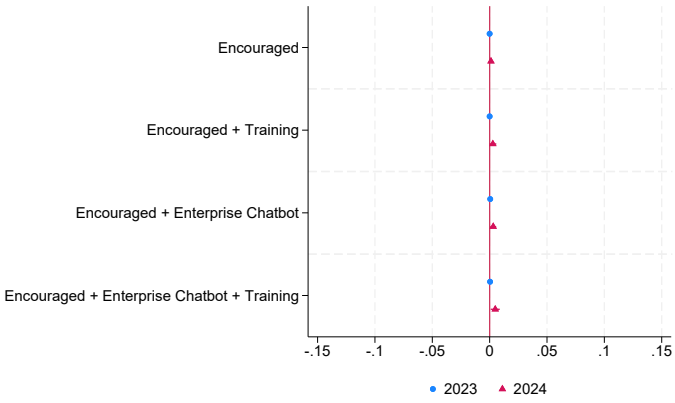
(a) Log Hours of Incumbent Workers



(b) Job Creation



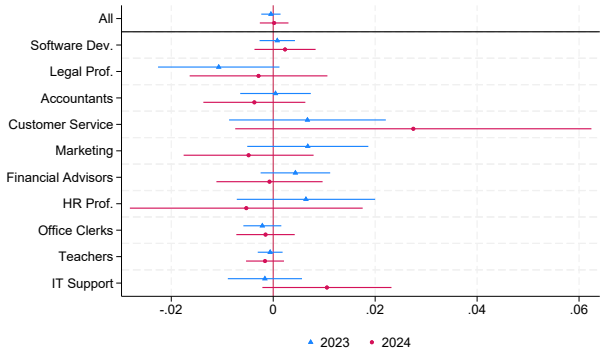
(c) Job Destruction



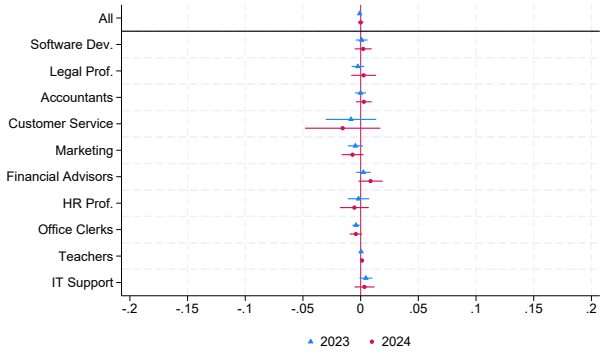
Notes: This figure shows how workplace outcomes vary with the intensity of adoption, for workers who report encouraged usage policies relative to those without encouragement (Figure 8). Panel (a) shows log hours of incumbent workers (those employed at the workplace throughout the pre-period, 2021M1–2022M11); Panels (b) and (c) show job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment). 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 F.8: Have Adopting Workplaces Fared Differently?
Heterogeneity by Occupation (DiD to Non-Encouraged)

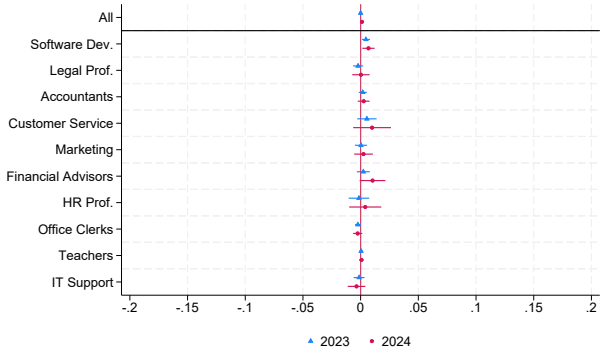
(a) Log Hours of Incumbent Workers



(b) Job Creation



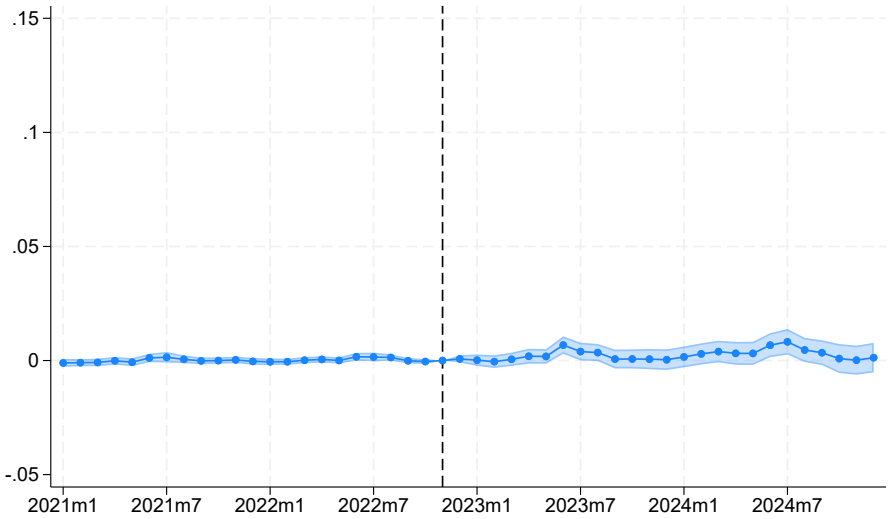
(c) Job Destruction



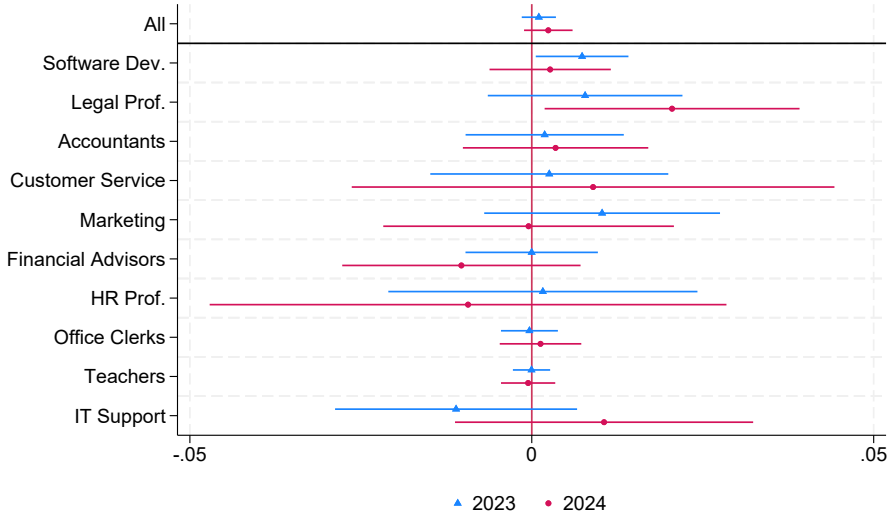
Notes: This figure shows occupation-specific estimates for the remaining workplace outcomes from Figure 8. Panel (a) shows log hours of incumbent workers (those employed at the workplace throughout the pre-period, 2021M1-2022M11); Panels (b) and (c) show job creation and destruction rates (measured as the share of hires and separations, respectively, in total employment). 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 F.9: Occupational Mobility of Incumbent Workers

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



(b) Occupational Heterogeneity (Pooled Diff-in-Diff)



Notes: This figure shows full-time equivalent employment of incumbent workers in their original occupations, in encouraged-use workplaces relative to those without encouragement, indexed to November 2022. Incumbent workers are those employed at the workplace throughout the pre-period (2021M1-2022M11). Panel (a) is based on the dynamic difference-in-differences specification in Equation (2), with shaded areas representing 95% confidence intervals. Panel (b) shows the effects separately by survey occupation, 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.

F.5 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 E.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.[30]

Table F.2 summarizes 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.

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.

[30]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.

Table F.2: Encouragement, Adoption, and Work (Additional Controls)

	(a) Adoption							
	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 (Among Ever Used)							
	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 (Among Ever Used)							
	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.

F.6 Coworker Encouragement

A key focus of our analysis is the link between employer chatbot initiatives and labor market outcomes. Section 2.2 shows how workers' adoption and reported work effects vary with these initiatives, and Section 3 uses employer-encouraged use as the main workplace treatment.

While our main analysis relies on self-reported employer encouragement, this appendix demonstrates robustness to using coworker reports. At the individual level, this helps

assess whether effects reflect workplace-wide initiatives or individual interventions; at the workplace level, it addresses concerns about attenuation bias from survey measurement error.

We focus the robustness analysis on encouraged use, our primary workplace treatment. Section F.6.1 describes how we construct coworker measures, and Section F.6.2 shows that our findings hold when using coworker encouragement, replicating first the adoption and self-reported work effects from Section 2.2, then the labor market results from Section 3.

F.6.1 Empirical Strategy

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

$$\text{Encouraged}_{j,-i} = \frac{1}{N_{j(i)} - 1} \sum_{k \in j(i) \setminus i} \text{Encouraged}_k, \tag{48}$$

where $N_{j(i)}$ denotes the number of respondents 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 1.3), 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.1 for details on this method). Importantly, our results in Section F.6.2 are robust to using the raw leave-out means instead.

Table F.3 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 3,116. The strong correlation in reported employer encouragement at the workplace level supports our interpretation of these initiatives as centralized policies.

Table F.3: Relationship Between Self- and Coworker-Reported Encouragement

	(1) Encouraged
Coworker IV (Leave Out, EBS)	0.877***
	(0.016)
Coworkers Share Female	-0.016
	(0.011)
Coworkers Age	0.001
	(0.001)
Coworkers Potential Experience	-0.001
	(0.001)
Occupation FEs	✓
F-Stat (Partial)	3116.352
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 (49). 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.6.2, 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:

$$\text{Encouraged}_i = \pi' X_{j,-i} + \alpha \times \text{Encouraged}_{j,-i} + \varepsilon_{1i} \tag{49}$$

$$Y_i = \gamma' X_{j,-i} + \beta \times \widehat{\text{Encouraged}}_i + \varepsilon_{2i}, \tag{50}$$

where $X_{j,-i}$ denotes the leave-one-out mean of the characteristics X of worker i's coworkers. Table F.3 reports estimates of Equation (49).

F.6.2 Results

Adoption and Work. Table F.4 compares the effects of self- and coworker-reported employer encouragement on AI chatbot adoption and associated benefits. The results broadly align, but with some 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. 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.

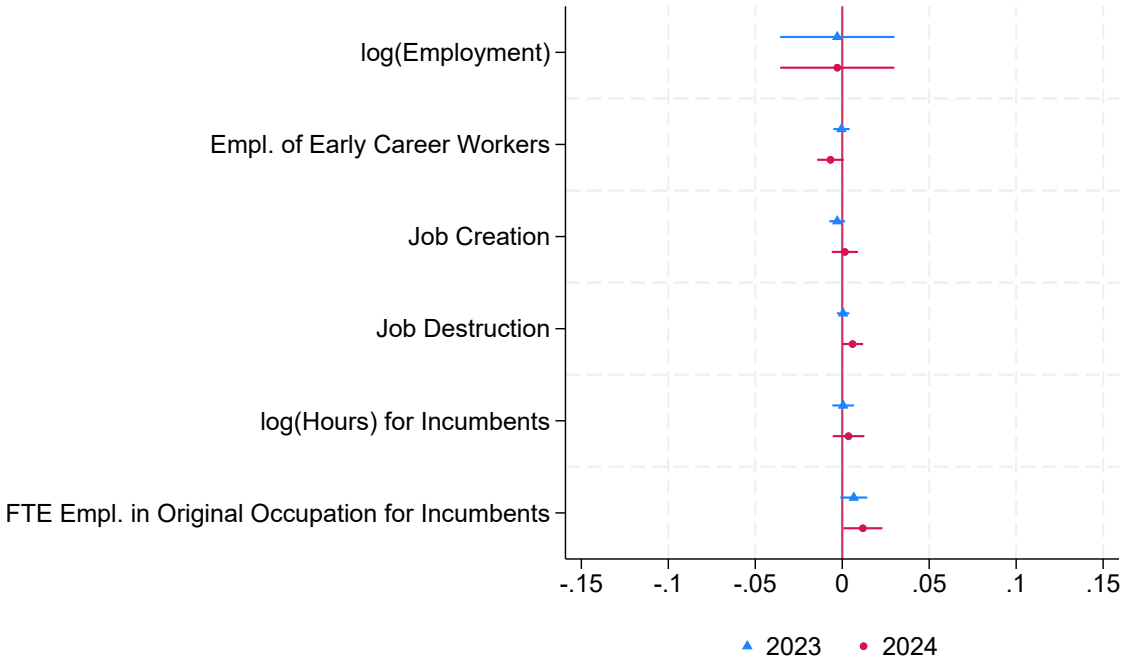
Table F.4: Encouragement, Adoption, and Work (Own vs. Coworker Reports)

	(a) Adoption							
	Ever Used		Monthly or More		Weekly or More		Daily	
	Own	CWs	Own	CWs	Own	CWs	Own	CWs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Encouraged	0.363	0.727	0.337	0.638	0.285	0.505	0.137	0.220
	(0.007)	(0.021)	(0.007)	(0.020)	(0.007)	(0.020)	(0.005)	(0.015)
	(b) Reported Benefits							
	Time Savings		Quality		Creativity		Job Satisfaction	
	Own	CWs	Own	CWs	Own	CWs	Own	CWs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Encouraged	0.100	0.049	0.116	0.079	0.091	0.091	0.071	0.035
	(0.008)	(0.030)	(0.009)	(0.031)	(0.009)	(0.030)	(0.007)	(0.024)
	(c) Allocation of Time Savings							
	More of Same Tasks		More of Diff. Tasks		More Breaks		More Leisure	
	Own	CWs	Own	CWs	Own	CWs	Own	CWs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Encouraged	0.031	0.040	0.016	0.037	-0.007	-0.020	-0.011	-0.017
	(0.009)	(0.034)	(0.007)	(0.024)	(0.005)	(0.015)	(0.006)	(0.018)
	(d) New Workloads from AI Chatbots (Adopters)							
	Same Tasks		New Tasks		Same and New Tasks			
	Own	CWs	Own	CWs	Own	CWs		
	(1)	(2)	(3)	(4)	(5)	(6)		
Encouraged	0.016	0.015	0.045	0.107	0.008	0.014		
	(0.004)	(0.011)	(0.006)	(0.020)	(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 (49)–(50). 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.

Workplace Employment. Figure F.10 revisits the impact of employer encouragement on workplace employment outcomes (Figure 8), now using our leave-out Bayes shrinkage measure for employer encouragement. The figure shows that all our difference-in-differences estimates are robust to using coworker reports to measure employer encouragement.

Figure F.10: Have Adopting Workplaces Fared Differently? (Coworker Encouraged DiD)



Notes: This figure shows workplace employment outcomes for encouraged-use workplaces relative to those without encouragement, indexed to November 2022. It replicates our "Encouraged" estimates from Figure 8, now using coworker reports to measure employer encouragement (see Section F.6.1). The estimates correspond to the following outcome variables (from top to bottom): log employment, employment share of early-career workers (aged 22–25), job creation and destruction rates (hires and separations as shares of total employment), log total hours of incumbent workers (those employed at the workplace throughout the pre-period, 2021M1–2022M11), and FTE hours in incumbent workers' original occupation. Estimates are based on the pooled specification in Equation (3), with 95% confidence intervals shown as whiskers. *Sample:* All completed responses from the 2024 survey linked to registry data.

G Estimation Procedures

G.1 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:

$$x_i \mid p_i \sim \text{Binomial}(n_i, p_i), \quad p_i \sim \text{Beta}(\alpha_0, \beta_0). \tag{51}$$

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:

$$\bar{p} = \frac{1}{m}\sum_{i=1}^{m} \frac{x_i}{n_i}, \quad s^2 = \frac{1}{m}\sum_{i=1}^{m}\left(\frac{x_i}{n_i} - \bar{p}\right)^2. \tag{52}$$

From the Beta mean and variance formulas:

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

With these, we compute the posterior mean:

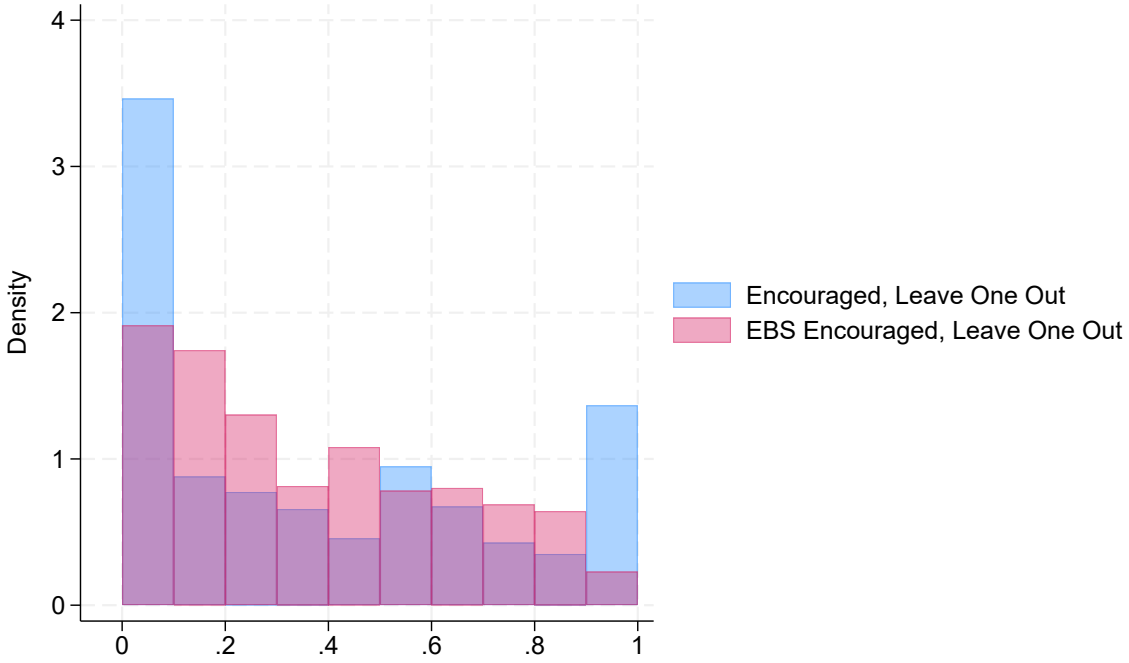
$$\mathbb{E}[p_i \mid x_i] = \frac{\alpha_0 + x_i}{\alpha_0 + \beta_0 + n_i}.$$

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

G.1.1 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.6.2 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.1. *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	.395	.45	.636	.186
Software Developers	.492	.686	.78	.181
Legal Professionals	.039	.127	.544	.306
Accountants & Auditors	.247	.41	.661	.247
Customer Service Rep.	.218	.354	.557	.179
Marketing Professionals	.589	.844	.891	.205
Financial Advisors	.111	.347	.74	.303
HR Professionals	.206	.506	.813	.299
Office Clerks	.177	.284	.54	.213
Teachers	.044	.121	.282	.178
IT Support	.284	.486	.678	.222
All	.127	.308	.587	.268

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.1. *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 67, followed by the original Danish version on page 69.

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

Our 2024 survey is organized into the four blocks summarized below. The 2023 round followed a similar structure.

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.

The section below 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 English translation starts on page 72, with the original Danish version on page 77.

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.

Survey Questionnaire – Danish Version

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.