# Prompting Change: Firm-Level AI Adoption and Worker-Level Outcomes in Linked Administrative Data

*Preliminary conference draft*

Simone Maria Bonin* Saman Darougheh* Andreas Kuchler*

This version: February 1, 2026

**Abstract**

We study the adoption of generative artificial intelligence (genAI) by Danish firms and its effects on workers. Combining a firm-level survey on AI adoption with administrative employer-employee data, we document which firms adopt genAI technologies and how adoption affects workers. Larger and more productive firms are substantially more likely to adopt AI, and these adopters also report higher profitability gains from AI use. Using novel measures of actual firm-level AI adoption, we find that occupation-level AI exposure roughly aligns with theoretical exposure measures but exhibits meaningful deviations. Our preliminary findings on worker outcomes indicate no detectable employment effects, but workers in AI-exposed occupations experience significant wage growth, even when their own employers have not adopted AI.

**JEL Classification:** J23, J24, J31, O33

**Keywords:** Generative AI, technology adoption, employment, wages, skills

*Danmarks Nationalbank. Emails: smbo@nationalbanken.dk (Bonin), sdma@nationalbanken.dk (Darougheh), aku@nationalbanken.dk (Kuchler). The views expressed in this paper are those of the authors and do not necessarily reflect the position of Danmarks Nationalbank.

# 1 Introduction

The rapid diffusion of generative artificial intelligence presents a potentially transformative shift in how firms produce goods and services. Unlike previous waves of automation that primarily affected routine manual tasks, genAI technologies—such as large language models and image generators—may reshape demand for cognitive and creative work. Understanding which firms adopt these technologies and how adoption affects workers is essential for anticipating macroeconomic adjustments and informing policy.

This paper provides an empirical analysis of genAI adoption and its labour market implications using comprehensive Danish data. We pursue two related objectives. First, we document the determinants of firm-level adoption, examining whether productivity differentials, labour shortages, or workforce characteristics predict early adoption. Second, we study how adoption affects workers by analysing employment dynamics, wage changes, and workforce composition at adopting firms relative to non-adopters.

Our empirical approach combines two data sources. We use Statistics Denmark's survey on IT usage in firms, which provides direct measures of AI tool adoption across a representative sample of Danish enterprises. We link these survey responses to the Danish matched employer-employee register (Beskæftigelse for Lønmodtagere, BFL), which allows us to characterise workforce composition at adopting and non-adopting firms and to track worker outcomes over time.

The Danish setting offers important advantages. Denmark is consistently identified as among the most AI-ready countries globally, ranking first in the EU on digital infrastructure and AI adoption by enterprises (McKinsey & Company, 2019; OECD, 2025). The 2024 EU Digital Decade report documents that 28% of Danish companies use AI—the highest share in Europe and double the EU average (European Commission, 2024b). Denmark's labour market is characterized by the "flexicurity" model, combining flexible hiring and firing rules with generous social protection and active labour market policies (Kreiner and Svarer, 2022), making it particularly adaptable to technological change. As we document

in our data, 41% of firms and 61% of workers in our sample are exposed to AI adoption. This combination—high AI adoption, adaptable labour markets, and comprehensive administrative data linking firm decisions to detailed worker characteristics—makes Denmark an ideal laboratory for studying the early labour market effects of generative AI.

We document three main findings. First, firm size and productivity are the strongest predictors of AI adoption, consistent with fixed costs of implementation and scaling potential. Notably, more productive firms that adopt AI also report significantly higher profitability gains from AI use, suggesting that adoption may amplify productivity differences and increase market concentration. Second, we compare actual AI usage patterns—measured directly from firm adoption surveys—with theoretical exposure measures from Eloundou et al. (2023). We find that occupation-level exposure to AI-adopting firms broadly aligns with these theoretical predictions, but with meaningful deviations that highlight the importance of using observed adoption data.[1] Third, examining worker outcomes, our preliminary findings indicate no detectable effects on employment but significant positive effects on wages. Workers in more AI-exposed occupations experience wage growth even when their own employer has not adopted AI, with particularly strong wage gains for exposed workers at AI-adopting firms. This pattern suggests either general equilibrium effects operating across firms or individual worker adoption of AI tools independent of firm-level policies. These findings are consistent with recent industry reports indicating that most companies remain in an experimentation phase rather than undertaking fundamental organizational restructuring (McKinsey & Company, 2024; Boston Consulting Group, 2024; Samans and Zahidi, 2025), which may explain why we observe wage effects without corresponding employment adjustments during our 2023–2025 ob-

[1]Our preliminary analysis suggests that theoretical exposure measures perform somewhat better at the two-digit occupation level ($R^2$ of 0.50) than at the more disaggregated three-digit level ($R^2$ of 0.30), possibly because AI substitutability is difficult to precisely define at finer occupation classifications. Researchers without access to firm-level microdata may prefer using two-digit occupation codes when constructing exposure measures.

servation window.[2]

Our analysis contributes to a rapidly growing literature on AI and the labour market. This literature spans several related questions: which workers and occupations are most exposed to AI, how AI adoption affects firm performance and worker productivity, and what the aggregate employment and wage impacts are likely to be.

*Measuring exposure to AI.* Recent work has developed task-based measures of occupational exposure to large language models. Eloundou et al. (2023) construct exposure measures based on the extent to which GPT-4 can perform tasks within each occupation, finding substantial exposure among high-skill cognitive occupations. Felten et al. (2023) develop alternative exposure measures that distinguish between different types of AI applications, showing heterogeneity in which occupations are exposed to different AI capabilities. Our paper contributes to this literature by comparing these theoretical exposure measures to actual firm-level AI adoption patterns, providing the first direct validation using observed adoption data linked to administrative employer-employee records.

*Firm-level AI adoption and productivity.* Several studies examine how AI adoption affects firm performance. Brynjolfsson et al. (2025), in a field experiment at a Fortune 500 customer support center, find that access to a generative AI tool increases worker productivity by 14% on average, with particularly large gains (34%) for novice and low-skilled workers. They show that AI disseminates best practices and helps newer workers move down the experience curve. Babina et al. (2024), using a novel measure of AI investments based on employee resumes, find that AI-investing firms experience higher sales growth, employment growth, and market valuations, with growth coming primarily through increased product innovation rather than labor displacement. However, they document that these benefits concentrate among larger firms, contributing to increased industry concentration. Our findings confirm this pattern of concentrated adoption and benefits:

[2]McKinsey reports that while 65% of organizations adopted AI in 2024, nearly two-thirds have not begun scaling AI across the enterprise, with most still experimenting or piloting without redesigning processes. BCG finds that 74% of companies struggle to unlock value from AI, and only about one in five companies has redesigned processes around AI (McKinsey & Company, 2024; Boston Consulting Group, 2024).

we similarly find that larger and more productive firms are substantially more likely to adopt AI, and that these same adopters report higher profitability gains from AI use. We complement this work by examining both self-reported effects of AI on firm profitability and linking adoption to detailed worker-level outcomes in administrative data.

*Adoption patterns.* Recent evidence documents rapid adoption of generative AI tools. Bick et al. (2024) conduct nationally representative surveys of US workers and find that by late 2024, nearly 40% of the US population aged 18-64 uses generative AI, with 23% of employed respondents using it for work. They document that work adoption of generative AI has been as fast as the personal computer following its introduction. Our paper provides complementary evidence on firm-level adoption patterns using representative firm surveys linked to administrative data, allowing us to identify which types of firms adopt AI and how workforce composition relates to adoption decisions.

*Labor market impacts and the automation-augmentation tradeoff.* A central question in this literature concerns whether AI will primarily automate workers out of jobs or augment their capabilities. Acemoglu (2024) develops a theoretical framework distinguishing automation (capital replacing labor in tasks) from the creation of new tasks in which labor has comparative advantage. He argues that much current AI development focuses on "so-so automation" that displaces workers without generating sufficient productivity gains, potentially reducing labor demand. Autor (2024) offers a more optimistic perspective, arguing that AI's unique capacity to weave information with acquired experience could enable mid-skill workers to perform higher-stakes decision-making currently restricted to elite experts, potentially rebuilding middle-class jobs. Acemoglu et al. (2022), using online vacancy postings, find that earlier AI technologies primarily substituted for workers in routine tasks. Our findings of wage gains without employment losses during 2023–2025 are consistent with Autor (2024)'s augmentation view, though our observation window captures a period where we believe most firms had not yet undertaken organizational restructuring.

Recent empirical work on generative AI's labor market impacts yields mixed findings. Noy and Zhang (2023) conduct a randomized experiment with professional writers and find that ChatGPT increases productivity by 40% for writing tasks. Several studies document positive wage effects in exposed occupations without widespread displacement. However, long-run impacts remain uncertain, and theoretical work suggests the direction of effects depends critically on whether AI development focuses on automation versus augmentation (Acemoglu, 2024).

We contribute to this literature in several ways. First, we combine direct measures of firm-level AI adoption from representative surveys with comprehensive administrative employer-employee data, allowing us to characterize both adoption patterns and worker outcomes in a setting with minimal measurement error. Second, we validate theoretical occupation-level exposure measures by comparing them to actual adoption patterns, documenting both broad alignment and important deviations. Third, we provide evidence on the early-stage impacts of AI adoption on employment and wages, documenting wage gains for workers in exposed occupations even at non-adopting firms—a pattern that may reflect either general equilibrium adjustments or individual worker adoption outside formal firm policies. Finally, our setting allows us to examine heterogeneity in outcomes by detailed occupation, firm characteristics, and worker demographics using population-level administrative data.

The remainder of the paper proceeds as follows. Section 2 describes our data sources and sample construction. Section 3 presents results on adoption determinants, exposure patterns, and worker outcomes, integrating empirical approaches with findings for each question. Section 4 concludes.

# 2 Data

Our analysis draws on two main data sources: a firm-level survey on IT usage and administrative employer-employee records. We begin by describing the Danish context that makes this setting particularly suitable for studying AI adoption and its labour market effects. We then describe each data source and our sample construction.

## 2.1 The Danish Context

Denmark's labour market is characterized by the "flexicurity" model, which combines three core elements: flexible employment protection legislation allowing firms to adjust their workforce, generous unemployment insurance providing income security for workers, and comprehensive active labour market policies with compulsory participation for benefit recipients (Kreiner and Svarer, 2022). Denmark has among the lowest employment protection legislation (EPL) levels in the OECD, with employers able to hire and fire with relatively few restrictions (Andersen and Svarer, 2005). This flexibility is balanced by a robust social safety net and high job-to-job mobility: approximately 25% of Danish private sector workers change jobs annually (Bredgaard et al., 2009). This institutional structure creates conditions where firms can readily adapt to technological change, while workers have both the security and incentives to transition across jobs and occupations.

Denmark's digital infrastructure and AI readiness provide further advantages for this study. Denmark consistently ranks first or second in the EU on the Digital Economy and Society Index (DESI), with particularly strong performance in connectivity (95% of households covered by very-high-capacity networks), digital public services (scoring 87.1 versus the EU average of 68.1), and digital skills (69.6% of the population possessing at least basic digital skills versus 55.6% EU average) (European Commission, 2024a). This digital foundation has translated into rapid AI adoption. The 2024 EU Digital Decade report documents that 28% of Danish companies use AI technologies—the highest share

in the EU and nearly double the European average of 13.5%—with Denmark ranking first among EU member states on AI adoption by both all enterprises and SMEs (European Commission, 2024b). Among large Danish firms (250+ employees), 63% report using AI. McKinsey characterizes Denmark as one of nine European "digital front-runners" with key strengths in AI readiness, including high technology adoption potential, solid digital and data infrastructure, and supportive policy environments (McKinsey & Company, 2019). Recent policy initiatives reinforce this trajectory, including a 2024–2027 national digitalization strategy committing 800 million DKK to AI, automation, and digital transformation, as well as partnerships to develop AI innovation centers with advanced computing infrastructure (OECD, 2025).

These institutional and technological features make Denmark an particularly informative setting for studying AI's labour market impacts. The combination of flexible labour markets, high digital readiness, rapid AI adoption, and comprehensive administrative data allows us to observe early-stage adjustments to AI in an environment conducive to both adoption and adaptation.

## 2.2 Administrative data on workers and firms

Our first data source are administrative records on workers and firms. For the firm-level analysis we primarily use data from the General Firm Statistics (FIRM) register maintained by Statistics Denmark. This register contains central economic variables such as investment, value added and employment. We link this data to data at the firm group level using a dataset of ownership structures based on an updated version of the data used by (Hviid et al., 2025). We primarily consider economic variables at the firm level, but for the classification into firm sizes we use total employment in the group that the firm belongs to.

For the analysis of workers, we use the Beskæftigelse for Lønmodtagere (BFL) register, which contains the universe of Danish wage payments since 2008 and allows us to

match workers to firms. BFL further contains the six-digit occupation code associated with each wage payment recorded in the DISCO-08 standard.[3] We further enrich this data with workers' demographic characteristics (age, gender, education) through links to population and education registers.

The employer-employee structure allows us to construct detailed measures of firm-level workforce composition. For each firm, we calculate the distribution of workers across occupations, education levels, age groups, and tenure categories. These measures serve both as potential determinants of AI adoption and as outcomes of interest when studying labour market adjustments. Importantly, BFL allows us to track individual workers over time, enabling analysis of worker-level outcomes such as job separations, wage growth, and occupational transitions.

## 2.3 IT Usage Survey

Our second data source is Statistics Denmark's survey on Information and Communication Technology (ICT) usage in enterprises. This survey, conducted annually and harmonised with Eurostat requirements, asks firms directly about their adoption of various digital technologies, including artificial intelligence tools.

The survey covers firms with 10 or more employees across most non-agricultural sectors. Relevant variables include indicators for AI adoption (distinguishing among different AI applications), the purposes of AI use, barriers to adoption, and the extent of implementation. The survey provides a representative picture of technology adoption across the Danish business sector.

[3]The Danish coding system, DISCO-08, is tightly linked to the to the 2008 version of the International Standard Classification of Occupations (ISCO).

## 2.4 Sample Construction

We construct our analysis sample by linking the IT usage survey to the administrative registers at the firm level or the worker level. For this, we use the 2024 wave of the IT usage survey.[4] The BFL is available with little delay, which allows us to study worker-level events up to 2025. Firm data (such as balance sheets) however are only available until 2023.

Table 1 presents summary statistics for key variables. The merged sample consists of 2,391 firms covering approximately 114,000 workers. Compared to the universe of firms with 10 or more employees, the merged sample is slightly larger on average and more likely to be exporters, reflecting the survey's sampling design. Panel A shows that the average firm in our sample employs 48 workers, with log value added per worker of 13.56 and an average age of 21 years. Panel B reports worker characteristics: the average worker is 40 years old with a log hourly wage of 5.53, and 82% work full-time.

Table 1: Summary statistics

| | Universe | Merged Sample |
|---|---|---|
| *Panel A: Firm Characteristics* | | |
| Number of employees | 49.93 | 47.74 |
| Log value added per worker | 13.55 | 13.56 |
| Firm age (years) | 20.12 | 20.65 |
| Export share | 0.12 | 0.21 |
| Exporter (indicator) | 0.35 | 0.54 |
| Number of firms | 45,445 | 2,391 |
| *Panel B: Worker Characteristics* | | |
| Age | 41.54 | 39.88 |
| Log hourly wage | 5.49 | 5.53 |
| Full-time (share) | 0.82 | 0.82 |
| Female (share) | 0.50 | 0.36 |
| Immigrant (share) | 0.13 | 0.17 |

*Notes:* Universe refers to firms with 10 or more employees.

[4] This survey has AI-related questions for earlier years. The potential of the technology however was much more limited in earlier years, which is why we focus on 2024.

Table 2 compares firm and workforce characteristics between AI adopters and non-adopters. AI adoption is substantial: 41% of firms in our sample report adopting AI tools, and these firms employ 61% of workers in the sample. Adopters are markedly larger (86 employees versus 33), more productive (log value added per worker of 13.74 versus 13.48), and more export-oriented (72% are exporters versus 47%). Their workforces are also somewhat younger (39 versus 41 years average age) and higher-paid (log hourly wage of 5.56 versus 5.49).

Table 2: Workforce Composition: Adopters vs. Non-Adopters

| | Adopters | Non-Adopters |
|---|---|---|
| *Panel A: AI Adoption Rates* | | |
| Share of firms | 0.41 | |
| Share of workers | 0.61 | |
| *Panel B: Firm Characteristics* | | |
| Number of employees | 86.08 | 32.82 |
| Log value added per worker | 13.74 | 13.48 |
| Firm age (years) | 21.03 | 20.50 |
| Export share | 0.29 | 0.18 |
| Exporter (indicator) | 0.72 | 0.47 |
| Number of firms | 970 | 1,421 |
| *Panel C: Workforce Composition* | | |
| Average worker age | 39.23 | 40.90 |
| Average log hourly wage | 5.56 | 5.49 |
| Full-time (share) | 0.83 | 0.81 |
| Share female | 0.37 | 0.33 |
| Share immigrant | 0.16 | 0.18 |

*Notes:* Comparison of firm and workforce characteristics between AI adopters and non-adopters in the merged sample. AI adopters are firms reporting AI tool usage in the 2024 IT usage survey. Firm characteristics from administrative data (2023 or latest available). Worker characteristics are employment-weighted averages from BFL register (2024).

# 3 Results

This section presents our main findings on AI adoption and its labour market implications. We proceed in four steps. First, in Section 3.1, we examine which firm characteristics predict AI adoption. Second, in Section 3.2, we analyse firms' self-reported assessments of AI's effects on profitability and other outcomes. Third, in Section 3.3, we measure worker exposure to AI-adopting firms and compare actual exposure patterns to theoretical predictions. Finally, in Sections 3.4 and 3.5, we investigate how AI exposure correlates with labour market outcomes, examining both occupation-level and worker-level evidence on employment and wages.

## 3.1 Determinants of AI Adoption

We first examine which firm characteristics predict genAI adoption. To do this, we estimate regressions of the form:

$$\text{AI}_f = \alpha + \mathbf{X}_f'\boldsymbol{\beta} + \mathbf{W}_f'\boldsymbol{\gamma} + \delta_s + \varepsilon_f, \tag{1}$$

where $\text{AI}_f$ is an indicator for whether firm $f$ reports adopting AI tools, $\mathbf{X}_f$ is a vector of firm characteristics (size, age, productivity, labour shortage indicators), $\mathbf{W}_f$ captures workforce composition (skill shares, average tenure, demographic composition), and $\delta_s$ denotes sector fixed effects.

Figure 1 reports estimates from this specification, showing coefficient estimates from multivariate regressions with industry fixed effects (not shown). We find that firm size and productivity are the strongest predictors of AI adoption: larger firms and more productive firms are substantially more likely to adopt AI. This pattern is consistent with fixed costs of implementing AI technologies combined with their scaling potential—larger and more efficient firms can more easily absorb upfront costs and spread benefits across more workers or outputs.

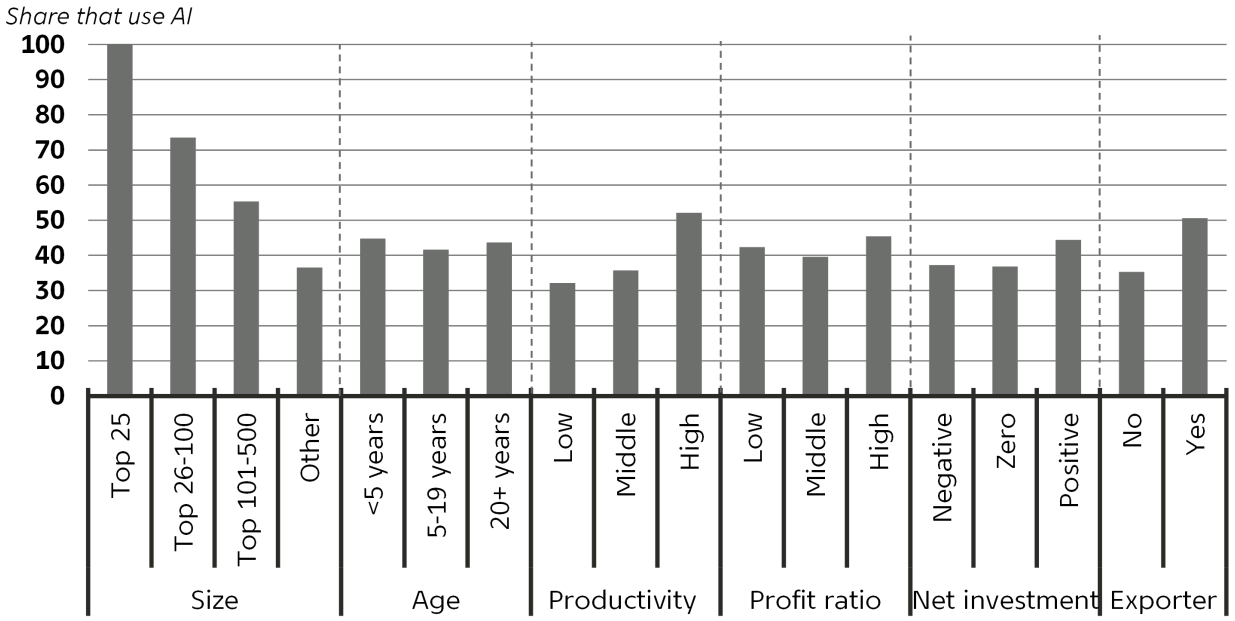



Figure 1: Determinants of AI adoption

*Notes:* Coefficient estimates from multivariate linear probability regressions of an AI adoption indicator on firm characteristics, with industry fixed effects (not shown). Sample includes firms from the 2024 IT usage survey merged with administrative data. AI adoption is an indicator for firms reporting use of AI tools in the survey.

The survey also asks firms about specific tasks for which they use AI. Figure 2 shows AI usage across different task categories, split by firm productivity. We find relatively high adoption rates for autonomous machines, text/code generation, text analysis, and data analysis. AI usage is notably lower for generating images, video, or sound, and for automating workflows. Importantly, while more productive firms adopt AI more intensively across all categories, there is no strong differential pattern in *which* types of AI these firms use—highly productive firms simply use more AI for all purposes. Appendix Figure A.1 shows substantial variation across industries, with Information and Communication being the largest adopter for most AI applications (especially data analysis), while Mining and Utilities lead in workflow automation.

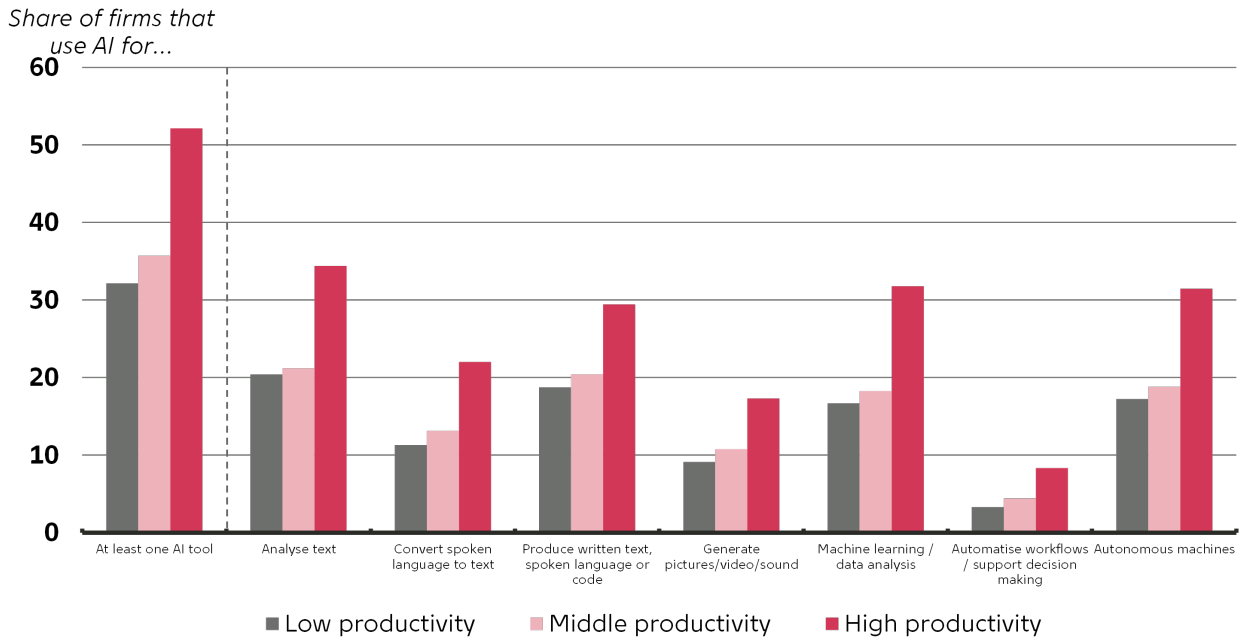



Figure 2: AI usage by task and firm productivity

*Notes:* Share of AI-adopting firms using AI for each task category, split by firm productivity terciles. Task categories based on survey questions about AI applications. Productivity measured as log value added per worker. Source: 2024 IT usage survey.

## 3.2 Self-Reported Effects of AI on Firm Outcomes

Having documented adoption patterns, we next examine firms' own assessments of AI's impact. Due to data limitations, we cannot yet observe objective firm outcomes post-adoption in the administrative records (recall that firm balance sheet data are only available through 2023, while the AI adoption survey is from 2024). The survey does, however, ask AI adopters about the extent to which AI has affected several firm-level outcomes.

Figure 3 reports these self-reported effects, split by firm productivity. These are subjective responses at the firm level and come with the usual caveats—not only must AI adoption genuinely affect outcomes, but firm management and survey respondents must be aware of these effects. With these limitations in mind, we find that on average 12% of AI adopters report that AI use has had a positive effect on their profits. Importantly, this is 12% of adopters, not 12% of all firms. This relatively modest figure seems reasonable given that adoption and implementation take time and involve significant upfront costs.

More strikingly, we find a clear gradient by firm productivity: more productive firms are substantially more likely to report positive profit effects from AI. This suggests not only that productive firms are more likely to adopt AI (as shown above), but also that they are better at implementing and extracting value from it. This pattern raises the possibility that AI adoption may amplify existing productivity differences and further increase market concentration. Appendix Figure A.2 shows large industry differences consistent with this interpretation: Information and Communication reports the highest share of adopters with positive profit effects (20%), while Construction reports the lowest (6%).

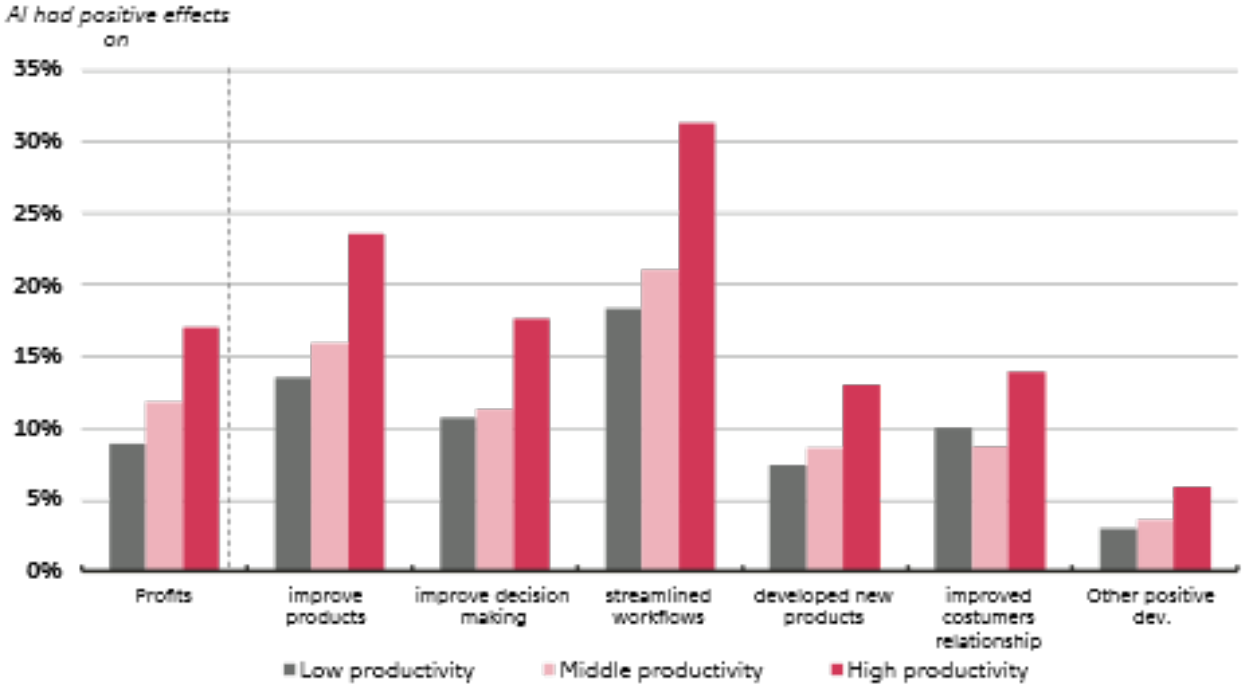



Figure 3: Self-reported firm-level effects of AI adoption

*Notes:* Share of AI-adopting firms reporting positive effects of AI on various firm outcomes, split by firm productivity terciles. Based on survey questions asking AI adopters whether AI has had positive, negative, or no effect on profits, employment, productivity, etc. Productivity measured as log value added per worker. Source: 2024 IT usage survey.

## 3.3 Measuring Worker Exposure to AI-Adopting Firms

We now turn to worker-level implications. A central question is which workers are most exposed to firms adopting AI technologies. We characterise worker exposure by measur-

ing the share of workers in each occupation employed at AI-adopting firms. For each occupation $o$, we compute:

$$\text{Exposure}_o = \frac{\sum_{f \in \mathcal{A}} L_{of}}{\sum_f L_{of}}, \tag{2}$$

where $L_{of}$ is employment in occupation $o$ at firm $f$, and $\mathcal{A}$ denotes the set of firms reporting AI adoption in the 2024 IT usage survey. This measure captures the extent to which workers in a given occupation are concentrated at firms that have adopted genAI technologies. It provides an empirical measure of actual exposure based on observed firm adoption decisions, in contrast to theoretical exposure measures based on task content analysis.

A key advantage of our empirical setting is that we can compare this actual exposure to theoretical exposure measures. Eloundou et al. (2023) develop occupation-level exposure measures based on the extent to which tasks in each occupation could be affected by large language models. These theoretical measures come in several variants: three different weighting schemes ($\alpha$, $\beta$, and $\gamma$) that vary in how they weight and aggregate task-level exposure to occupation-level measures,[5] and two different approaches to assessing task exposure—human expert ratings and AI-generated (deep vision) assessments. We merge these measures to our data using occupation codes and focus primarily on the human-rated $\alpha$ variant (which emphasizes direct task exposure), though we examine robustness to alternative specifications in the appendix.

Figure 4 shows the relationship between theoretical AI exposure (human-rated $\alpha$ measure from Eloundou et al.) and actual AI adoption among workers in each three-digit occupation. The figure reveals two important patterns. First, there is a clear positive relationship: occupations with higher theoretical exposure do indeed have higher actual exposure to AI-adopting firms. This validates the usefulness of task-based exposure measures for predicting which workers will be most affected by AI diffusion. Second, there is

[5]The $\alpha$ measure uses direct exposure weighting, $\beta$ incorporates complementarities between exposed and non-exposed tasks, and $\gamma$ accounts for potential task automation. See Eloundou et al. (2023) for details.

substantial dispersion around the trend line, with some occupations being considerably more or less exposed than the theoretical measure predicts.

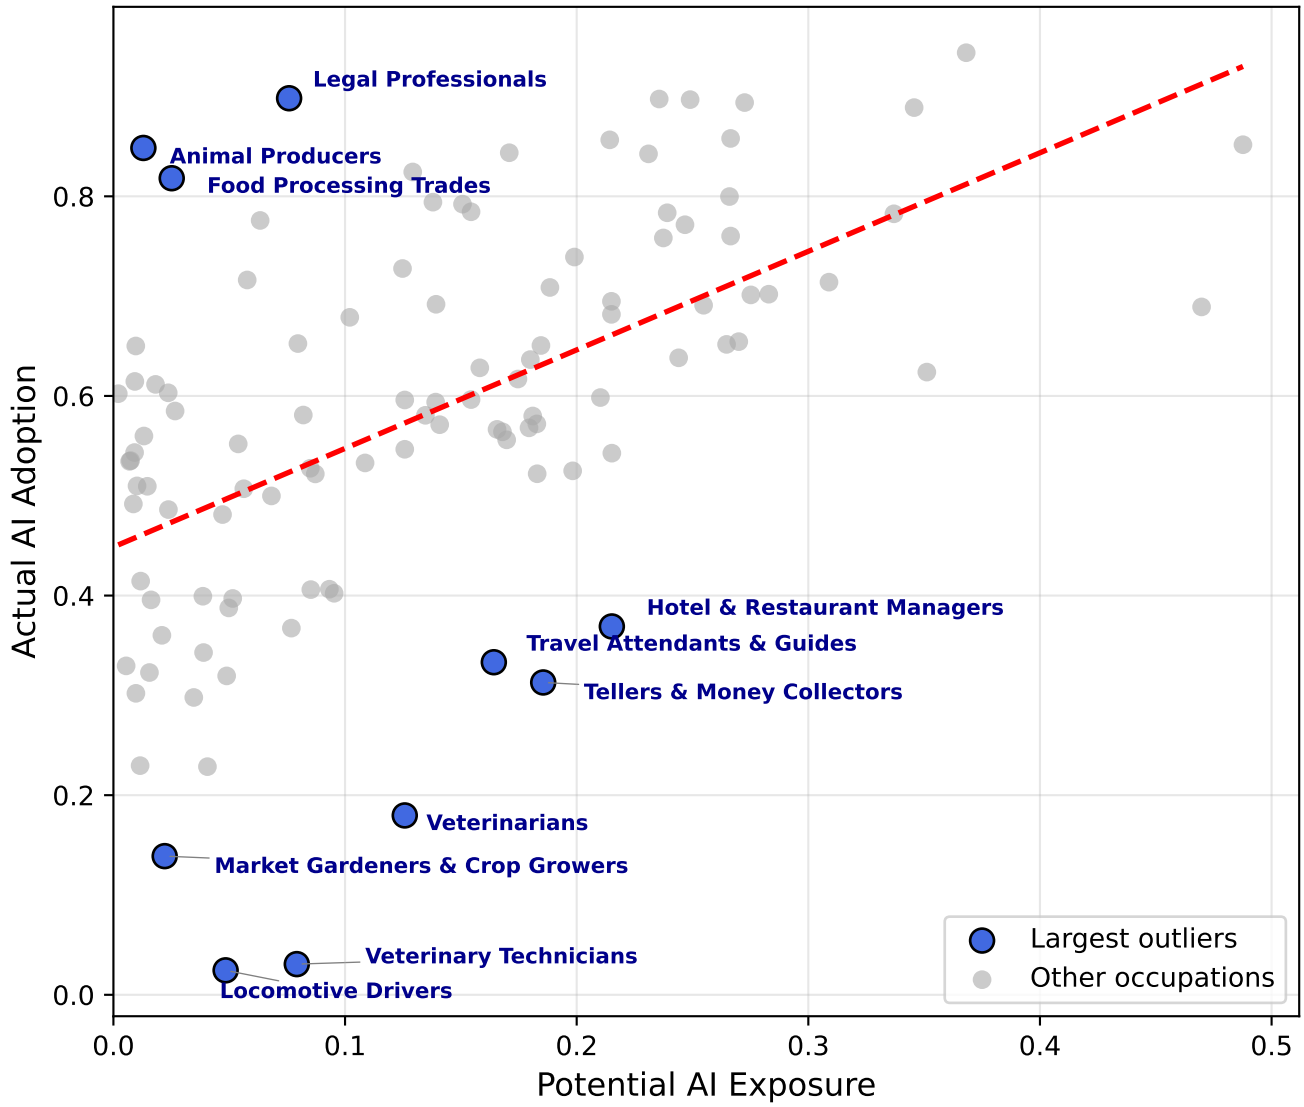
Several notable outliers merit discussion. Legal professionals appear substantially more exposed to AI-adopting firms than their theoretical exposure score would suggest. This may reflect early enthusiasm for AI tools in legal work—such as document review, contract analysis, and legal research—where large language models have demonstrated concrete applications even if theoretical task-based measures underestimate their potential impact. The legal profession's relatively high wages and document-intensive workflows may make AI adoption particularly attractive despite the specialized nature of legal reasoning. Conversely, some highly-exposed occupations according to theoretical measures show lower actual firm-level adoption, potentially reflecting implementation challenges, regulatory barriers, or lags between technological capability and organizational deployment.

This heterogeneity highlights the value of using observed adoption data: actual AI implementation may differ from theoretical predictions due to complementarities, firm-specific factors, or differences between technological potential and realized adoption.[6]

This result is robust to alternative exposure measures. Appendix Table A.1 reports regressions using all three theoretical exposure variants ($\alpha$, $\beta$, $\gamma$) with human ratings at the three-digit occupation level (N=109), showing that the positive relationship holds across all specifications with $R^2$ values ranging from 0.30 to 0.33. Interestingly, the theoretical exposure measure performs notably better when occupations are aggregated to the two-digit level. Appendix Table A.2 shows that at the two-digit level (N=38), the $R^2$ increases to 0.50 for the $\alpha$ variant—substantially higher than the 0.30 observed at the three-digit level. This pattern may reflect difficulties in precisely defining AI substitutability at very disaggregated occupation classifications. When tasks within broader occupation groups are heterogeneous, finer classifications may introduce noise that obscures the relationship

[6]The point estimates for mean AI usage come with a standard error of approximately 0.002 due to the large number of observations, which is why we do not display confidence bands in this figure.

Figure 4: Theoretical AI exposure and actual AI usage




*Notes:* Relationship between theoretical AI exposure and actual firm-level AI adoption by three-digit occupation (N=109). Theoretical exposure is the human-rated $\alpha$ measure from Eloundou et al. (2023), based on task-level exposure to GPT-4 capabilities aggregated to occupations. Actual AI usage is the share of workers in each occupation employed at firms reporting AI adoption in the 2024 IT usage survey. Occupations coded using DISCO-08 (Danish version of ISCO-08).

between theoretical and actual exposure. For researchers without access to firm-level microdata, this suggests that two-digit occupation codes may be preferable for constructing AI exposure measures. Appendix Table A.3 presents additional robustness checks using AI-generated (deep vision) exposure assessments instead of human ratings, confirming that the positive relationship between theoretical and actual exposure is consistent across different measurement approaches.

We also examine which specific types of AI usage best predict overall worker exposure. The average AI usage indicator naturally best approximates overall population exposure to AI. Appendix Table A.4 shows that overall AI usage and text generation usage are the best predictors of worker exposure—other more specific usage indicators (such as image generation or workflow automation) do less well at predicting overall population exposure.

### 3.4 Worker Outcomes: Employment and Wages

Having established which occupations are most exposed to AI adoption, we examine whether this exposure correlates with changes in labour market outcomes. We measure occupation-level changes in employment and wages between 2023 and 2025 and relate them to exposure measured in 2024. This analysis is descriptive: occupations differ along many dimensions beyond AI exposure, and we do not claim a causal interpretation. Rather, we document whether occupations with greater concentration of workers at AI-adopting firms experienced systematically different trajectories during this period.

Table 3 presents results from simple bivariate regressions of the form:

$$\Delta Y_o = \alpha + \beta \cdot \text{Exposure}_o + \varepsilon_o, \tag{3}$$

where $\Delta Y_o$ denotes the change in outcome $Y$ for occupation $o$ between 2023 and 2025, measured at the three-digit occupation level (109 occupations). The table reports results

for employment levels, log employment, and log hourly wages. We consider two exposure measures: actual firm-level AI usage and the theoretical exposure measure $\alpha$ from Eloundou et al. (2023).

The results show no systematic relationship between AI exposure and employment changes. Neither exposure measure predicts significant changes in employment levels or log employment, with small and statistically insignificant coefficients. However, we find a positive and marginally significant relationship between AI exposure and wage growth: the occupation-level AI usage measure predicts approximately 3.9 percentage points higher wage growth (column 7), significant at the 5% level. Given that the AI usage measure ranges from roughly 0.3 to 0.7 across occupations in our sample, this coefficient implies meaningful wage differences between high- and low-exposure occupations. The theoretical exposure measure $\alpha$ also predicts positive wage growth, though the coefficient is only marginally significant (column 8).

When both measures are included simultaneously (column 9), the actual AI usage measure remains statistically significant while the theoretical exposure measure loses significance. This pattern suggests that the deviations between theoretical and actual exposure documented in Figure 4 are meaningful for predicting labour market outcomes. While both measures capture relevant dimensions of AI exposure, our observed measure based on actual firm adoption may better reflect the aspects of AI implementation that translate into wage adjustments. This could occur if firms adopt AI differentially across occupations in ways that theoretical task-based measures do not fully capture—for instance, if complementarities between AI and certain skills, organizational factors, or implementation constraints lead actual adoption patterns to diverge from predictions based solely on task content. We interpret this cautiously: theoretical exposure measures remain valuable, particularly for researchers without access to firm-level adoption data, but our results suggest that observed adoption patterns contain additional information relevant for understanding early labour market impacts.

| | Employment | | | Log employment | | | Log hourly wages | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Occ: AI Usage | 673.4 | | 560.8 | 0.0501 | | 0.106 | 0.0394** | | 0.0348* |
| | (0.56) | | (0.39) | (0.61) | | (1.09) | (2.56) | | (1.88) |
| Exposure ($\alpha$) | | 923.3 | 369.6 | | -0.0789 | -0.184 | | 0.0497* | 0.0153 |
| | | (0.43) | (0.14) | | (-0.54) | (-1.05) | | (1.76) | (0.46) |
| R-squared | 0.00295 | 0.00171 | 0.00314 | 0.00351 | 0.00269 | 0.0137 | 0.0576 | 0.0282 | 0.0595 |
| N | 109 | 109 | 109 | 109 | 109 | 109 | 109 | 109 | 109 |

$t$ statistics in parentheses
$^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$

Table 3: AI exposure and occupation-level outcomes, 2023–2025

*Notes:* Bivariate regressions of occupation-level changes in outcomes (2023–2025) on AI exposure measures at three-digit occupation level (N=109). Occupation AI usage is the mean firm-level AI adoption rate within the occupation. Exposure ($\alpha$) is the theoretical exposure measure from Eloundou et al. (2023). Employment measured in levels and logs; log hourly wages from BFL register. Robust standard errors; $t$ statistics in parentheses.

These occupation-level patterns suggest that AI exposure may be associated with wage gains without detectable employment losses, at least in the short run covered by our data. This combination of findings—positive wage effects without employment losses—is consistent with broader evidence on the early stages of generative AI adoption. Recent surveys indicate that while AI adoption accelerated rapidly in 2024, with usage reaching 65% of organizations (McKinsey & Company, 2024), most companies remain in an experimentation phase rather than undertaking fundamental organizational restructuring. McKinsey reports that nearly two-thirds of organizations have not yet begun scaling AI across the enterprise, with the majority still experimenting or piloting AI tools without redesigning processes. Only about one in five companies has redesigned processes around AI (McKinsey & Company, 2024), and three-quarters of companies report not yet unlocking substantial value from AI (Boston Consulting Group, 2024). This "wait and see" approach, with companies testing AI capabilities before committing to workforce restructuring (Samans and Zahidi, 2025), may explain why we observe productivity-linked wage gains for exposed workers without corresponding employment adjustments. As AI adoption matures and firms move beyond experimentation to deeper organizational

integration, employment effects may emerge—but our 2023–2025 observation window captures a period when most firms had not yet undertaken such restructuring.

However, these are aggregate occupation-level correlations that cannot distinguish between different margins of adjustment or control for confounding occupation characteristics.

### 3.5 Worker-Level Wage Effects and Spillovers

To examine worker outcomes more granularly, we turn to individual-level regressions that exploit both occupation-level exposure and firm-level adoption. Table 4 presents estimates from worker-level regressions of hourly wage growth on occupation AI exposure, firm AI use, and their interaction. The specification is:

$$\Delta \log(wage)_{ifo} = \alpha + \beta_1 \cdot \text{Exposure}_o + \beta_2 \cdot \text{AI Use}_f + \beta_3 \cdot (\text{Exposure}_o \times \text{AI Use}_f) + \delta_s + \varepsilon_{ifo}, \tag{4}$$

where $i$ indexes workers, $f$ indexes firms, $o$ indexes occupations, and $\delta_s$ denotes industry fixed effects.

The results reveal several important patterns. First, workers in more AI-exposed occupations experience higher wage growth, even controlling for whether their specific employer adopted AI (Panel A, columns 1 and 3). The coefficient on occupation AI exposure is positive and highly significant across all specifications. Notably, this wage premium exists even for workers whose firms have not adopted AI, suggesting that workers in exposed occupations may benefit from AI adoption through channels beyond their own firm's decisions. One plausible explanation is that workers in these occupations may be using AI tools individually—perhaps tools they access independently or through informal channels—even when their employers have not formally adopted AI technologies. Alternatively, this pattern could reflect general equilibrium wage adjustments as AI adoption spreads across the economy.

Second, the direct effect of firm AI adoption is negative in these regressions—workers at AI-adopting firms have slightly lower wage growth on average, holding occupation exposure constant. However, this negative coefficient must be interpreted carefully in the presence of the interaction term.

Third, the interaction term between occupation exposure and firm AI use is large, positive, and highly significant. Workers in high-exposure occupations at AI-adopting firms experience substantially higher wage growth than would be predicted by either factor alone. Combined with the positive main effect of occupation exposure, this means that workers in AI-exposed occupations see wage gains whether or not their firm has adopted AI, but gains are particularly large when their firm has adopted.

These patterns are robust to controlling for industry fixed effects (columns 2 and 4) and to using alternative wage measures (narrow versus broad wage definitions). The effects are similar whether exposure is measured at the three-digit or four-digit occupation level (Panels A and B), though point estimates are somewhat smaller at the more disaggregated four-digit level.

Taken together, these worker-level results indicate that AI adoption is associated with wage gains for workers in exposed occupations, regardless of whether their own firm has adopted AI. The positive main effect of occupation exposure suggests either general equilibrium wage adjustments or individual-level AI usage even at non-adopting firms. The positive interaction term shows that wage gains are particularly strong when workers in exposed occupations are also employed at AI-adopting firms. These findings provide micro-level evidence consistent with the occupation-level patterns documented earlier, while revealing heterogeneity in how AI exposure affects different workers depending on their occupation and employer characteristics.

Table 4: AI Exposure and Hourly Wage Growth: Worker-Level Evidence

| | Broad Wage | | Narrow Wage | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| *Panel A: 3-Digit Occupation* | | | | |
| Occupation AI Exposure | $0.112^{***}$ | $0.031^{***}$ | $0.110^{***}$ | $0.029^{***}$ |
| | (0.004) | (0.004) | (0.004) | (0.004) |
| Firm AI Use | $-0.044^{***}$ | $-0.044^{***}$ | $-0.047^{***}$ | $-0.047^{***}$ |
| | (0.003) | (0.003) | (0.003) | (0.003) |
| Firm AI Use × Occ. AI Exposure | $0.097^{***}$ | $0.087^{***}$ | $0.101^{***}$ | $0.090^{***}$ |
| | (0.004) | (0.004) | (0.005) | (0.005) |
| Constant | $0.043^{***}$ | $0.071^{***}$ | $0.043^{***}$ | $0.070^{***}$ |
| | (0.002) | (0.003) | (0.002) | (0.003) |
| Industry FE | | ✓ | | ✓ |
| Observations | 499,739 | 499,739 | 499,741 | 499,741 |
| *Panel B: 4-Digit Occupation* | | | | |
| Occupation AI Exposure | $0.074^{***}$ | $0.008^{**}$ | $0.072^{***}$ | $0.007^{**}$ |
| | (0.003) | (0.003) | (0.003) | (0.003) |
| Firm AI Use | $-0.043^{***}$ | $-0.032^{***}$ | $-0.046^{***}$ | $-0.034^{***}$ |
| | (0.003) | (0.003) | (0.003) | (0.003) |
| Firm AI Use × Occ. AI Exposure | $0.096^{***}$ | $0.069^{***}$ | $0.099^{***}$ | $0.071^{***}$ |
| | (0.004) | (0.004) | (0.004) | (0.004) |
| Constant | $0.068^{***}$ | $0.085^{***}$ | $0.068^{***}$ | $0.085^{***}$ |
| | (0.002) | (0.002) | (0.002) | (0.002) |
| Industry FE | | ✓ | | ✓ |
| Observations | 499,739 | 499,739 | 499,741 | 499,741 |

*Notes:* Worker-level regressions of log hourly wage growth (2023–2025) on occupation AI exposure, firm AI use, and their interaction. Sample includes workers present in BFL register in both 2023 and 2025. Occupation AI Exposure is the mean firm-level AI adoption rate within the occupation (three-digit DISCO-08 in Panel A, four-digit in Panel B). Firm AI Use is an indicator for the worker's employer reporting AI adoption in the 2024 IT usage survey. Broad wage includes all wage payments; narrow wage excludes irregular payments. Industry fixed effects at two-digit level. Standard errors clustered at firm level; robust standard errors in parentheses. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

# 4 Conclusion

This paper documents patterns of generative AI adoption by Danish firms and correlates in worker outcomes using linked survey and administrative data. Three findings emerge. First, larger and more productive firms are substantially more likely to adopt AI, and these same firms report higher profitability gains from adoption. Second, theoretical task-based exposure measures broadly align with actual adoption patterns, though with meaningful deviations that underscore the value of observed adoption data. Third, workers in AI-exposed occupations experience differential wage growth regardless of whether their own employer has adopted AI, with particularly strong wage differences at adopting firms. We find no corresponding differences in employment.

These patterns carry several broader implications, though our descriptive approach warrants caution in interpretation.

*On inequality and market structure.* The concentration of both adoption and self-reported benefits among already-productive firms suggests AI may widen firm-level productivity dispersion. If these patterns persist, AI diffusion could accelerate trends toward market concentration documented in other contexts. At the worker level, differential wage growth favoring AI-exposed occupations—predominantly cognitive, high-skill work—may exacerbate occupational wage inequality, even as it raises living standards for affected workers.

*On the role of institutions.* Denmark's flexicurity model provides an unusual laboratory: flexible hiring and firing should facilitate reallocation, while generous social insurance reduces worker resistance to technological change. That we observe wage differentiation without employment disruption may partly reflect this institutional configuration. Whether similar patterns emerge in economies with stronger employment protection or weaker safety nets remains an open question. The Danish experience may represent an upper bound on how smoothly labour markets can absorb AI in the short run.

*On measurement.* Our comparison of theoretical and observed exposure measures

offers practical guidance. Theoretical measures perform better at two-digit than three-digit occupation levels, suggesting researchers without firm-level microdata should prefer broader occupation classifications when constructing exposure measures. The substantial residual variation between theoretical and actual exposure—even at $R^2$ of 0.50—highlights that task content alone incompletely predicts which workers encounter AI-adopting employers.

*On timing.* Our 2023–2025 window captures what appears to be an experimentation phase, with widespread adoption but limited organizational restructuring. The absence of employment effects may reflect this timing rather than a fundamental feature of the technology. As firms move from experimentation to process redesign, employment adjustments may materialize. Continued monitoring with administrative data will be essential for distinguishing transitional dynamics from longer-run equilibrium effects.

Several limitations bear emphasis. Our analysis documents correlations, not causal effects. Firms and workers selecting into AI adoption differ on observable and likely unobservable dimensions, and the occupations most exposed to AI were plausibly on different trajectories before these technologies arrived. The finding that wages rise in exposed occupations even at non-adopting firms is consistent with multiple interpretations—general equilibrium reallocation, individual worker adoption outside firm policies, or spurious correlation with pre-existing occupation-level trends—that we cannot distinguish. Future work exploiting plausibly exogenous variation in AI adoption would sharpen these conclusions.

Despite these limitations, the linked administrative data and direct adoption measures allow us to characterize early-stage patterns with unusual precision. The Danish experience suggests that the initial correlates of generative AI adoption include wage differentiation favoring exposed occupations and concentrated adoption among productive firms, without—so far—the employment displacement that some have anticipated.

# Acknowledgements

The authors have no financial interests or conflicts of interest to disclose. The views expressed in this paper are those of the authors and do not necessarily reflect the position of Danmarks Nationalbank.


# References


**Acemoglu, Daron**, "The Simple Macroeconomics of AI," *Working Paper*, April 2024.

_ , **David Autor, Jonathon Hazell, and Pascual Restrepo**, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

**Andersen, Torben M and Michael Svarer**, "The Danish Flexicurity Model," Technical Report, Employment Relations Research Centre, University of Copenhagen 2005.

**Autor, David H**, "Applying AI to Rebuild Middle Class Jobs," *NBER Working Paper*, February 2024, (32140).

**Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson**, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

**Bick, Alexander, Adam Blandin, and David J Deming**, "The Rapid Adoption of Generative AI," *NBER Working Paper*, September 2024, (32966).

**Boston Consulting Group**, "AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value," Technical Report, Boston Consulting Group October 2024.

**Bredgaard, Thomas, Flemming Larsen, and Per Kongshøj Madsen**, "Denmark: Flexibility, Security and Labour Market Success," Technical Report 53, International Labour Organization 2009.

**Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond**, "Generative AI at Work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–942. NBER Working Paper 31161, April 2023.

**Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock**, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," *arXiv preprint arXiv:2303.10130*, 2023.

**European Commission**, "Digital Economy and Society Index 2024: Denmark," Technical Report, European Commission 2024.

_ , "The State of the Digital Decade 2024: Denmark," Technical Report, European Commission 2024.

**Felten, Edward, Manav Raj, and Robert Seamans**, "Occupational Heterogeneity in Exposure to Generative AI," *AEA Papers and Proceedings*, 2023, *113*, 282–287.

**Hviid, Simon Juul, Morten Spange, Tobias Renkin, Mia Renee Herløv Jørgensen, and Rasmus Rold Sørensen**, "The Increasing Importance of the Largest Firms," March 2025. Analysis No. 8.

**Kreiner, Claus Thustrup and Michael Svarer**, "Danish Flexicurity: Rights and Duties," *Journal of Economic Perspectives*, 2022, *36* (4), 81–102.

**McKinsey & Company**, "An AI Nation: Harnessing the Opportunity of Artificial Intelligence in Denmark," Technical Report, McKinsey & Company October 2019.

_ , "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value," Technical Report, McKinsey & Company May 2024.

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

**OECD**, "Progress in Implementing the European Union Coordinated Plan on Artificial Intelligence: Denmark," Technical Report, OECD 2025.

**Samans, Richard and Saadia Zahidi**, "2025: The Year Companies Prepare to Disrupt How Work Gets Done," *World Economic Forum*, January 2025.

# A Additional Tables and Figures

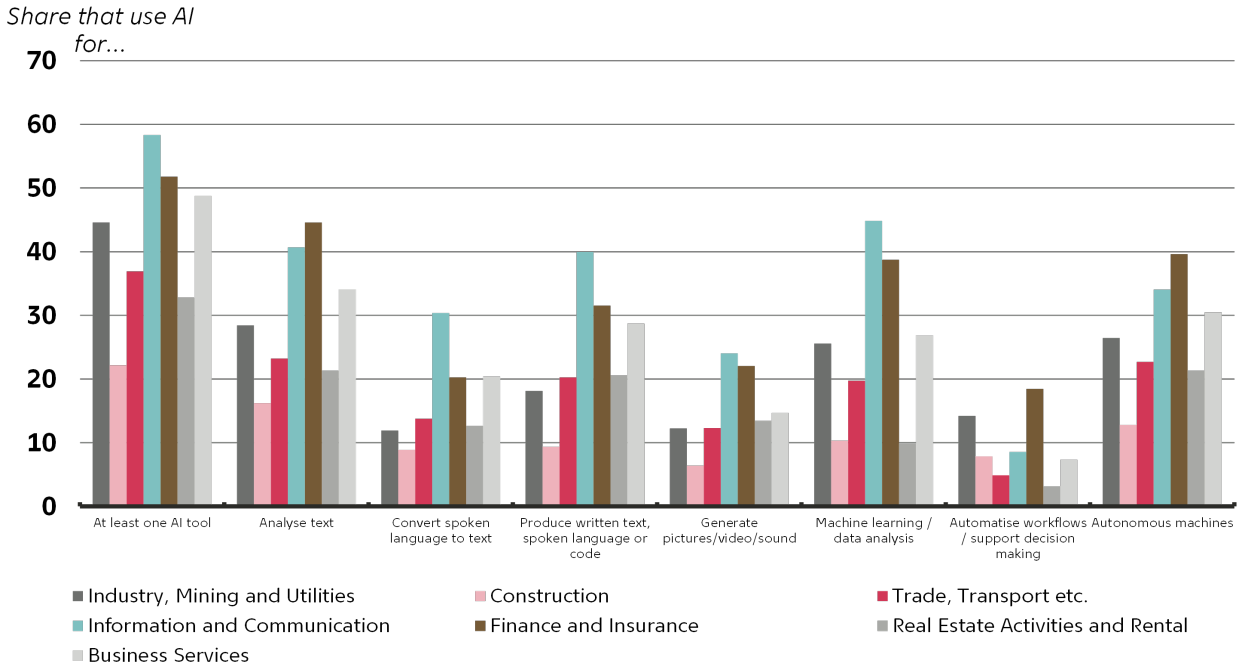



Figure A.1: AI usage by industry and task

*Notes:* Share of AI-adopting firms using AI for each task category, by industry. Industries classified using two-digit NACE codes. Task categories based on survey questions about AI applications. Source: 2024 IT usage survey.

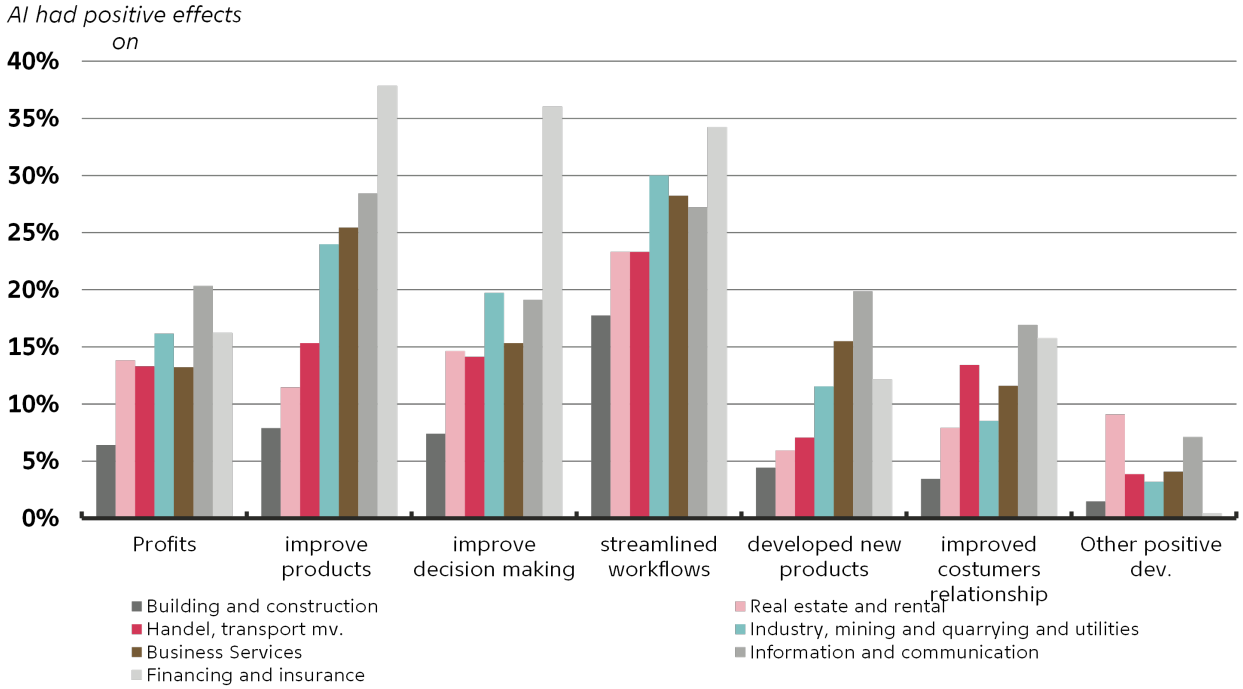



Figure A.2: Self-reported firm-level effects of AI adoption by industry

*Notes:* Share of AI-adopting firms reporting positive effects of AI on various firm outcomes, by industry. Industries classified using two-digit NACE codes. Based on survey questions asking AI adopters whether AI has had positive, negative, or no effect on profits, employment, productivity, etc. Source: 2024 IT usage survey.

| | (1) | (2) | (3) | (4) | (5) | |
|---|---|---|---|---|---|---|
| | mean_ai_use | mean_ai_use | mean_ai_use | mean_adj_ai_use | mean_adj_ai_use | m |
| exposure | 0.987*** | 0.642*** | 0.445*** | 0.264** | 0.152* | |
| | (6.78) | (7.33) | (7.22) | (2.64) | (2.46) | |
| _cons | 0.449*** | 0.397*** | 0.386*** | 0.566*** | 0.558*** | |
| | (17.72) | (13.35) | (12.26) | (32.58) | (26.59) | |
| AI Measure | $\alpha$ | $\beta$ | $\gamma$ | $\alpha$ | $\beta$ | |
| Industry FE | | | | ✓ | ✓ | |
| R-squared | 0.301 | 0.334 | 0.328 | 0.0613 | 0.0534 | |
| N | 109 | 109 | 109 | 109 | 109 | |

$t$ statistics in parentheses

$^{*}$ $p < 0.05$, $^{**}$ $p < 0.01$, $^{***}$ $p < 0.001$

Table A.1: Theoretical exposure and actual AI usage: human ratings, 3-digit occupations

*Notes:* Dependent variable is mean firm-level AI usage within each three-digit occupation (N=109). The three theoretical exposure measures ($\alpha$, $\beta$, $\gamma$) are from Eloundou et al. (2023) based on human expert ratings. Columns 1–3 report bivariate regressions; columns 4–6 include industry fixed effects.

| | (1) mean_ai_use | (2) mean_ai_use | (3) mean_ai_use | (4) mean_adj_ai_use | (5) mean_adj_ai_use | m |
|---|---|---|---|---|---|---|
| exposure | 1.205*** | 0.670*** | 0.450*** | 0.295* | 0.146 | |
| | (6.02) | (5.93) | (5.71) | (2.28) | (1.98) | |
| _cons | 0.428*** | 0.398*** | 0.393*** | 0.568*** | 0.566*** | |
| | (13.03) | (10.60) | (9.93) | (26.82) | (23.18) | |
| AI Measure | $\alpha$ | $\beta$ | $\gamma$ | $\alpha$ | $\beta$ | |
| Industry FE | | | | ✓ | ✓ | |
| R-squared | 0.502 | 0.494 | 0.476 | 0.127 | 0.0984 | |
| N | 38 | 38 | 38 | 38 | 38 | |

$t$ statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2: Theoretical exposure and actual AI usage: human ratings, 2-digit occupations

*Notes:* Dependent variable is mean firm-level AI usage within each two-digit occupation (N=38). The three theoretical exposure measures ($\alpha$, $\beta$, $\gamma$) are from Eloundou et al. (2023) based on human expert ratings. Columns 1–3 report bivariate regressions; columns 4–6 include industry fixed effects. The $R^2$ values at the two-digit level (0.50 for $\alpha$) are substantially higher than at the three-digit level (0.30), suggesting that theoretical exposure measures perform better at broader occupation aggregations.

| | (1) mean_adj_ai_use | (2) mean_adj_ai_use | (3) mean_adj_ai_use |
|---|---|---|---|
| exposure | 0.101 | 0.119* | 0.0802* |
| | (1.13) | (2.06) | (2.17) |
| _cons | 0.587*** | 0.562*** | 0.560*** |
| | (34.43) | (25.32) | (25.19) |
| AI Measure | $\alpha$ | $\beta$ | $\gamma$ |
| Industry FE | ✓ | ✓ | ✓ |
| R-squared | 0.0118 | 0.0383 | 0.0422 |
| N | 109 | 109 | 109 |

$t$ statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.3: Theoretical exposure and actual AI usage: AI-generated ratings, 3-digit occupations

*Notes:* Dependent variable is mean firm-level AI usage within each three-digit occupation (N=109). The three theoretical exposure measures ($\alpha$, $\beta$, $\gamma$) are from Eloundou et al. (2023) based on AI-generated (deep vision) assessments rather than human expert ratings. Results are similar to those using human ratings (Appendix Table A.1).

| | (1) All AI Usage | (2) Text | (3) Speech | (4) Language | (5) Pictures | (6) Drawing | (7) AIAUTR |
|---|---|---|---|---|---|---|---|
| human_rating_alpha | 0.987*** | 1.106*** | 0.772*** | 0.982*** | 0.576*** | 1.062*** | 0.336* |
| | (6.78) | (7.78) | (6.97) | (6.51) | (4.81) | (6.76) | (3.09) |
| Constant | 0.449*** | 0.227*** | 0.103*** | 0.248*** | 0.122*** | 0.242*** | 0.0866* |
| | (17.72) | (9.17) | (5.34) | (9.45) | (5.87) | (8.87) | (4.58) |
| R-squared | 0.301 | 0.361 | 0.312 | 0.284 | 0.178 | 0.299 | 0.0821 |
| N | 109 | 109 | 109 | 109 | 109 | 109 | 109 |

*t* statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.4: AI usage types and worker exposure

*Notes:* Regressions of mean occupation-level AI usage (various types) on theoretical exposure measure ($\alpha$ from Eloundou et al. 2023). Each column represents a different dependent variable: overall AI use, or specific AI application types (text generation, image generation, data analysis, etc.). Sample includes three-digit DISCO-08 occupations (N=109). Robust standard errors; *t* statistics in parentheses.

| | (1) mean_AIPOSINDT | (2) mean_AIPOSINDT | (3) mean_AIPOSINDT |
|---|---|---|---|
| mean_ai_use | 0.187 | | -0.0497 |
| | (1.32) | | (-0.25) |
| human_rating_alpha | | 0.506* | 0.567 |
| | | (2.15) | (1.65) |
| _cons | 0.331*** | 0.374*** | 0.395*** |
| | (3.79) | (9.78) | (4.22) |
| R-squared | 0.0473 | 0.116 | 0.118 |
| N | 37 | 37 | 37 |

*t* statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.5: Self-reported profit effects by occupation (2-digit)

*Notes:* Regressions of mean occupation-level self-reported profit effects on theoretical AI exposure. Dependent variable is the share of AI-adopting firms within each occupation reporting positive effects of AI on profits. Exposure measures ($\alpha$, $\beta$, $\gamma$) from Eloundou et al. (2023). Sample aggregated to two-digit DISCO-08 occupations (N=38). Robust standard errors; *t* statistics in parentheses.

| | (1) mean_AIPOSINDT | (2) mean_AIPOSINDT | (3) mean_AIPOSINDT |
|---|---|---|---|
| mean_ai_use | 0.366*** | | 0.292* |
| | (3.60) | | (2.34) |
| human_rating_alpha | | 0.471** | 0.203 |
| | | (2.86) | (1.02) |
| _cons | 0.222*** | 0.377*** | 0.238*** |
| | (3.47) | (12.96) | (3.62) |
| R-squared | 0.114 | 0.0748 | 0.123 |
| N | 103 | 103 | 103 |

*t* statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.6: Self-reported profit effects by occupation (3-digit)

*Notes:* Regressions of mean occupation-level self-reported profit effects on theoretical AI exposure. Dependent variable is the share of AI-adopting firms within each occupation reporting positive effects of AI on profits. Exposure measures ($\alpha$, $\beta$, $\gamma$) from Eloundou et al. (2023). Sample includes three-digit DISCO-08 occupations (N=109). Robust standard errors; *t* statistics in parentheses.