DANMARKS NATIONALBANK

WORKING PAPER | Labour market
16 September 2026 | No. 223

AI Adoption and Firm Size: Employment Effects in Monthly Administrative Data

How do firms adjust their employment after adopting artificial intelligence (AI)? Using a unique dataset of Danish firm-level AI adoption data matched to monthly employer–employee records, we find that firms that adopted AI in 2023 reduced employment growth by 11 per cent relative to non-adopters by late 2025. Effects are concentrated in smaller firms and reflect lower recruitment, particularly in AI-exposed occupations. Economy-wide employment is so far relatively unaffected.

Written by

Simone Maria Bonin
Senior Economist

smbo@nationalbanken.dk
+45 3363 6527

Saman Darougheh
Senior Research Economist

sdma@nationalbanken.dk
+45 3363 6126

Andreas Kuchler
Advisor

aku@nationalbanken.dk
+45 3363 6561

Keywords

Danish economy | Households and corporations | Labour market | Wages

AI Adoption and Firm Size: Employment Effects in Monthly Administrative Data*

Simone Maria Bonin† Saman Darougheh* Andreas Kuchler*

September 16, 2026

Abstract

How do firms change their employment after adopting AI? We answer this using two unique Danish data sources: the official firm-level IT usage survey, which includes AI-related questions since 2022, and a matched employer–employee register that records every wage payment at monthly frequency. Together, they identify the firms that first used AI in 2023—the earliest measurable cohort—and follow them month by month. Adopters are larger than non-adopters and employ more workers in AI-exposed occupations, so we compare each firm not against other firms but against the employment path it was already on. We find that adopters fall behind that path significantly more than non-adopters do: 6% two years on, 11% by late 2025. The gap comes entirely from adopters with fewer than 100 FTEs, and runs through reduced hiring rather than layoffs, concentrated in AI-exposed occupations—25% below trend, against 7% elsewhere. Wages do not respond: the entire adjustment runs through employment. These effects are so far too small to show up in economy-wide employment: shrinking AI adopters employ too few workers to move the aggregate, and while AI-exposed industries do shed employment, an exact decomposition attributes only about a quarter of that to adoption.

Keywords: artificial intelligence, technology adoption, firm employment, hiring, occupational exposure, firm size, matched employer–employee data.
JEL classification: D22, E24, J23, J24, O33.

---

*The views expressed are those of the authors and not necessarily those of Danmarks Nationalbank.

†Danmarks Nationalbank. Corresponding author: Saman Darougheh (sdma@nationalbanken.dk). The views expressed in this paper are those of the authors and do not necessarily reflect the position of Danmarks Nationalbank. We thank Tobias Renkin and other colleagues for valuable input.

1 Introduction

Artificial Intelligence (AI) is diffusing through the economy faster than previous automation waves: firms adopt AI because it raises worker productivity on cognitive and knowledge tasks (Brynjolfsson, Li, and Raymond, 2025; Noy and Zhang, 2023). How do adopting firms change their employment structure? The answer to this question is key to understanding how fast the transition is taking place, to what extent it affects workers, and which workers' job security is most threatened.

This requires knowing which firms started using AI, and knowing it early. Employment adjusts slowly, so only firms that adopted years ago have had the time to show a response; by 2025 AI use is so widespread that the firms still not using it are no longer a credible comparison group. Few data sources identify adopters that early, which is why the question has largely gone unanswered.

We can answer this question because Denmark provides us with a unique combination of two data sources. The first is Statistics Denmark's IT Usage survey, the official survey of firms' technology use, harmonized across EU countries. Eurostat dropped its AI module in 2022; Statistics Denmark kept asking the question of its own accord, and it is the continuity of the survey that makes first-time adoption in 2023 identifiable at all—elsewhere in Europe the missing year breaks the series.[1] The second is the matched employer–employee register, which records every wage payment in the Danish economy at monthly frequency. The combination—an official adoption survey read at monthly frequency—is what the adoption literature has so far lacked. It is not a technical detail: our results cannot be reproduced for the same period using balance-sheet data alone.[2]

This allows us to compare firms that first report using AI in 2023 with firms reporting no AI use, within detailed industry-by-month cells.[3] Because adopters and non-adopters differ from the outset, we do not compare their employment levels against each other: we first estimate the employment path each firm was already on before 2023, and then ask how far each firm departs from its own path. Our key finding is that the employment responses to adoption differ sharply by firm size. Adopters with fewer than 100 FTEs fall about 16 log points—roughly 15%—below their own pre-trend by late 2025, while large adopters stay statistically flat. On average, adopters fall about 6% below trend two years after adoption and about 11% by late 2025, relative to comparable non-adopters. The decline runs through reduced hiring rather than separations, and concentrates in the occupations most exposed to AI. We now discuss our findings in more detail.

In the first part of the paper, we describe AI adoption in Denmark and establish that our survey-based measure is meaningful. AI usage among surveyed firms rose steeply, from 15% of firms in 2023 to 43% in 2025, driven almost entirely by text and code generation—the hallmark of the 2023 chatbot wave. Adoption is far from random: larger firms, firms employing more educated

[1] A change in question framing limits comparability with the 2021 wave, so we date our adoption series from 2022.

[2] Appendix F repeats our exercise on the same firms' annual balance-sheet data, the data which most adoption studies use, and returns confidence bands spanning roughly −8% to +5%—wide enough to contain everything we report below.

[3] Common to the adoption literature, we do not claim to find causal effects of AI adoption on employment. Instead, we document employment differences between firms that choose to employ AI—where AI usage and employment are likely jointly determined by an overarching strategy—and similar firms that do not.

workers, and firms in industries considered more AI-exposed are more likely to adopt. Several patterns support the measure's credibility: firm-level adoption is highly persistent across survey waves, occupational AI exposure at adopting firms maps tightly onto task-based exposure indices such as Felten, Raj, and Seamans (2023), and about half of adopters report that AI has not yet been profitable, suggesting that responses are not merely enthusiastic self-promotion.

In the second part of the paper, we compare the post-adoption employment paths of 2023 adopters and firms reporting no AI use. We estimate a firm-specific pre-trend from each firm's 2021–2022 employment and study deviations from it in monthly matched employer–employee data. Adopters and non-adopters track their trends almost identically before 2023; from 2023 onward, a gap opens that widens roughly linearly, reaching about 6% below trend after two years and roughly 11% by late 2025. The decline operates through reduced hiring rather than separations: the stock of workers with under two years of tenure falls roughly 28 log points below trend, and the workforce measurably ages.

The decline, however, is not homogeneous across firms: it is driven entirely by adopters with fewer than 100 FTEs—about −16 log points by late 2025—while large adopters stay on their own pre-adoption trend.

Among these smaller firms, we further split the entire distribution of firm employment growth (relative to pre-trends) for both adopters and non-adopters. We show that the fall in employment is not homogeneous across the growth distribution, nor is it driven by the bottom tail alone: the top quantiles have similar growth among adopting and non-adopting firms; it is firms in the median and below that grow significantly less among adopting firms. The bottom quantiles also appear to grow less among adopting firms, but with confidence bands too wide to interpret confidently.

Turning from firm heterogeneity to worker heterogeneity, we show that the fall is sharper among young, university-educated workers. It is most concentrated in AI-exposed occupations, where employment at adopters falls roughly 25% below trend, compared with 7% elsewhere. Among small adopters, this is driven by a sharp decline in exposed-occupation employment. Large adopters also shift their employment from exposed to non-exposed occupations, while their overall headcount stays on trend.

Section 5 then asks what this firm-level decline implies beyond our estimation sample. Aggregated to the industry level, an exact decomposition attributes only about a quarter of the industry exposure gradient to the adoption channel and roughly three quarters to composition—AI-exposed industries simply contain more firms shrinking across the board, rather than adopters shrinking more than their industry peers.

Overall, our results point to significant restructuring of employment at adopting firms already in 2024 and 2025. Small firms use AI to shrink exposed occupations, whereas large firms keep their head count on trend but tilt employment towards less-exposed occupations. So far, these adoption effects are too small to have aggregate effects. Yet, there are reasons for concern: these effects come from early adopters, with most firms only shallowly integrating AI while future costs are still uncertain. Already finding employment effects under these circumstances lends weight to concerns that labor markets need to be more resilient for a potentially large restructuring ahead.

Literature. A large and still-growing literature has studied the effects of AI from a theoretical and empirical perspective, for a review of the theoretical papers see del Río-Chanona et al. (2025). The task-based approach (studied in the context of automation by Acemoglu and Restrepo (2019); see Restrepo (2024) for a recent review) lends itself to studying which workers are most affected by AI. Felten, Raj, and Seamans (2023), Eloundou, Manning, Mishkin, and Rock (2023), and Webb (2020) measure the AI exposure of each occupation by estimating how productive artificial intelligence is at performing its underlying tasks. They find that workers in white-collar occupations performing cognitive tasks are—in theory—most exposed to AI. It is ex-ante theoretically ambiguous whether that higher exposure implies higher productivity, higher employment and wages, or whether AI-exposed workers are more likely to be replaced by AI and thus face lower wages and lower employment. Handa et al. (2025) measure AI-exposure of occupations by linking actual AI usage data to estimated worker-specific occupations. They attempt to cut through the ambiguous effects of AI by classifying each occupation as AI-augmentable or AI-substitutable depending on whether the AI requires much user intervention or can mostly complete tasks independently.

Exposure-based studies ask whether workers in exposed labor markets (mostly occupations) experience declining employment rates or wages and find that employment has indeed begun weakening (Brynjolfsson, Chandar, and Chen, 2025; Lodefalk, Löthman, Koch, and Engberg, 2026; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025). Beyond the occupation-level employment studies cited above, the exposure strand's evidence spans job postings—junior software-developer postings fell by about 14% after ChatGPT (Westby, Sasser Modestino, and Cheng, 2025)—online freelancing platforms (Demirci, Hannane, and Zhu, 2025; Hui, Reshef, and Zhou, 2024), and regional exposure designs (Bonfiglioli, Crinò, Gancia, and Papadakis, 2025). Timing is a shared caveat: AI-exposed occupations began deteriorating months before ChatGPT (Frank et al., 2026), cautioning against attributing all post-2022 divergence to AI.

Turning to AI adoption, prior work has documented that AI use by firms is strongly size-dependent (Bonney et al., 2024; Calvino and Fontanelli, 2023). Among firm-level adoption designs, the central difficulty is measuring actual AI use. For instance, some studies infer adoption from firms' direct survey responses on AI use (McElheran et al., 2024), others from worker-reported establishment initiatives (Humlum and Vestergaard, 2025), and others from job-posting or vacancy text (Acemoglu, Autor, Hazell, and Restrepo, 2022; Babina, Fedyk, He, and Hodson, 2024; Alekseeva et al., 2021).

Firm-level adoption designs ask whether adopting firms change their employment structure, and typically find that they do not (e.g., Yotzov et al., 2026). The two features of our design set out above—the earliest measurable adoption cohort, and monthly rather than annual employment data—can explain why our findings differ from theirs. Generative AI has diffused faster than earlier technologies (Bick, Blandin, and Deming, 2025): later cohorts leave no credible comparison group, and time aggregation, together with the reporting lags of annual accounts, smears away a ramp-shaped decline of this size.

In our Danish context, Humlum and Vestergaard (2025) provide pioneering evidence on chatbots at work, linking a large worker survey to the administrative registers and focusing on the workplaces where chatbot use matters most—those employing eleven highly exposed occupations. They find that workplaces that encourage AI use do not change their employment differentially relative to non-encouraging workplaces. The two designs are complementary and answer different questions. Their treatment is a standing encouragement policy among already-exposed workplaces—an intensive margin. In contrast, we compare first-time adopters against firms reporting no AI use at all: the extensive margin, in a representative cross-section of Danish firms and within finely defined industries. Where the two papers measure the same objects, they agree.

Closest to our design, Aghion et al. (2025) link France's official adoption survey to administrative data and find—for the 2017–2020 pre-chatbot vintage of AI—that employment *increases* after adoption. Indeed, the effects of pre-chatbot AI appear quite different, since U.S. firms investing in AI through the 2010s tilted toward educated and STEM workers with *rising* junior shares (Babina, Fedyk, He, and Hodson, 2023).

Our findings refine the incidence picture of the exposure literature. The decline at adopting firms concentrates in theoretically exposed occupations—validating task-based measures such as Felten, Raj, and Seamans (2023) both as predictors of adoption and of employment losses.

Our aggregation results bridge the firm-level adoption evidence and the industry-level exposure findings. They are qualitatively consistent with Tucker (2026)—declines concentrated in the more AI-exposed industries and operating through reduced hiring rather than separations—though our composition decomposition cautions that these cross-sectional gradients do not purely pick up effects from AI usage.

2 Setting and Data

2.1 Institutional context

Denmark is a well-suited laboratory for studying the early labor-market effects of AI, for two reasons: its labor market is unusually fluid, and its firms are among the most AI-ready in Europe.

Denmark's labor market is characterized by the "flexicurity" model, combining flexible employment protection with generous unemployment insurance and active labor market policies (Kreiner and Svarer, 2022). Danish employment protection is among the lowest in the OECD—firms can adjust their workforce with few restrictions (Andersen and Svarer, 2005)—and worker mobility is correspondingly high: roughly a quarter of private-sector workers change jobs each year (Bredgaard, Larsen, and Madsen, 2009). This matters for interpreting our results. In a high-churn labor market, a firm that stops hiring need not shed incumbents, and workers not hired by one firm are readily absorbed elsewhere. The flexicurity setting is thus key to how our two central findings—an adjustment that runs through the hiring margin rather than separations, and a firm-level employment decline that does not (yet) show up at the aggregate—should be interpreted. It plausibly makes

Denmark an upper bound on how smoothly a labor market absorbs AI in the short run.

Denmark is also a frontrunner in AI adoption. It consistently ranks at or near the top of the EU on the Digital Economy and Society Index, with strong connectivity, digital public services, and digital skills (European Commission, 2024a; McKinsey & Company, 2019). This has translated into fast take-up of AI: the 2024 EU Digital Decade report finds that 28% of Danish enterprises use AI—the highest share in the EU and roughly double the European average of 13.5%—rising to 63% among large firms (European Commission, 2024b), a lead reinforced by recent national digitalization initiatives (OECD, 2025). These externally measured figures corroborate the high adoption rates we document in our own survey data in section 3 and suggest that Denmark offers an early view of adjustments that other economies may encounter later.

2.2 Data and sample

We combine three linked sources: Statistics Denmark's IT Usage Survey; a matched employer-employee register which contains the universe of wage payments (BFL); and firm balance-sheet and background data, available through 2024 in the FIRM and FIRE registers.

At the heart of the analysis is Statistics Denmark's IT survey. This EU-harmonized survey has been conducted annually since 1998 and samples enterprises with at least 10 full-time employees in private non-financial industries.

In this paper, we will distinguish between AI users—any firm that employs AI—and AI adopters—firms that have reported using AI in a given year, but not the previous year. Unless otherwise specified, "AI adopter" refers to 2023-adopters (who had not been using AI in 2022). Their comparison group, which we call "non-adopters," are the firms reporting no AI use in either 2022 or 2023; we reserve "never-adopters" for firms reporting no use in any survey wave, a distinction the industry-level decomposition in section 5.2 relies on. This is the earliest year in which we can measure adoption, allowing us to trace the employment effects of adopters for the longest possible horizon.

We link the 4,000 surveyed firms to the administrative registers. To compare the surveyed firms with our register data, we restrict the "universe" of firms to those with at least 10 employees—comparable to the sampling restriction of the survey. Table 1 compares these surveyed firms with the universe of Danish firms with at least 10 employees; this 10% sample comprises 30% of the Danish labor force. Surveyed firms are also more export-oriented than the average Danish firm. Both facts are consistent with the survey design's oversampling of larger firms. Relative to the baseline sample, the surveyed firms employ slightly younger workers who are more likely to be male.

Table 1: Matched sample descriptives

	Universe	Matched Sample
Panel A: Firm Characteristics		
Number of employees	84.02	58.30
Log labor productivity	13.47	13.49
Investment rate	4.88	4.70
Firm age (years)	23.03	22.46
Export share	0.16	0.18
Exporter (indicator)	0.43	0.54
Number of firms	25,608	4,139
Panel B: Worker Characteristics		
Age	41.43	39.74
Log hourly wage	5.49	5.53
Full-time (share)	0.82	0.82
Female (share)	0.50	0.36
Non-Danish origin (share)	0.13	0.17
Basic education (share)	0.15	0.20
Upper secondary education (share)	0.11	0.12
Short higher education (share)	0.26	0.17
Long higher education (share)	0.14	0.12
Number of workers	2,346,654	749,111

Notes: All statistics refer to 2024. Universe: firms with $\geq$10 employees and their workers, unweighted. Matched sample: survey firms linked to registers and their workers; firm-level statistics are weighted using survey sampling weights, while worker-level statistics are unweighted. Worker counts are average monthly headcounts over the year.

3 Adoption

One concern might be that a surveyed firm's response is subject to noise—perhaps because the person in charge of responding to the survey is not well informed about the firm's AI strategy, or because they have little incentive to answer it correctly (which might cost effort).

Besides describing AI usage in Denmark, this section aims to verify that the AI usage responses are indeed meaningful.

We proceed in two steps. Section 3.1 describes what AI is used for—its subtypes, how usage evolves across survey waves, and whether adopters think it has paid off. Section 3.2 then asks who adopts, and shows that firms are selected both on their own characteristics and on employing occupations that theory deems AI-exposed. Along the way, we document two patterns that suggest the credibility of the AI measures: (i) respondents are not overly positive with respect to the effects of AI on firm profits, and (ii) firm-level AI adoption correlates highly with the extent to which its employees are in theoretically AI-exposed occupations.

3.1 What AI is used for

AI adoption in Denmark is both high and rising: In our sample, 43% of firms report using AI in 2024—consistent with Denmark's top place in the European rankings documented in section 2.1. Figure 1a aggregates the seven possible non-exclusive survey responses to three major groups: generation of text (and code), working with data, and automating processes. In 2025, most firms

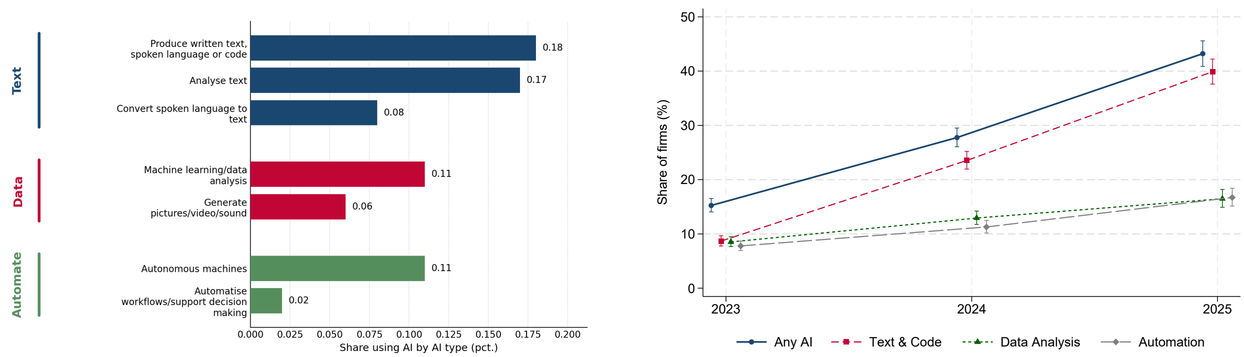

(a) By AI type

(b) Usage over time

Figure 1: Share of AI users over time

Notes: Share of surveyed firms using AI, by 2024 usage type (Text / Data / Automate) in panel (a), and over time in panel (b). The years refer to the reference period in the survey. Source: Statistics Denmark IT Usage Survey; firms with $\geq$10 employees; survey-weighted.

use AI for text generation and analysis, followed by data analysis and automation.

Figure 1b shows the evolution of AI usage over time: while it is at 43% in 2025, it was only at 15% in 2023, underlining the sharp increase throughout the past 3 years. The figure also breaks AI usage down into the three major subcategories. While usage increased across all categories, the overall rise in AI usage stems almost entirely from increases in Text & Code. This is consistent with the primary innovation in 2023 being chatbots, with a strong comparative advantage in generating text and code.

Figure 2a selects firms that responded in both 2023 and 2024 and plots AI usage in 2024 conditional on usage in 2023. If responses were not consistent with a true firm strategy and instead determined by noise, one would expect mean reversion: 2023 users would be less likely to be 2024 users. At the extreme, the share of 2024 users among 2023 users would be the same as the unconditional sample average. Instead, we observe that 2023 users are far more likely to be 2024 users than 2023 non-users: 83.2% versus 25.9%, a gap of nearly sixty percentage points around an unconditional share of 43.2%. When we turn to firms that were sampled in all three survey waves (not reported in the figure), we observe that 49% never use AI: if AI usage was an independent process purely determined by noise, we would expect that share to be much lower, at about 33%.

Usage composition also varies sharply across industries: information and communication is the most intensive user across every application, construction the least (Figure 2b).

Do firms think AI has paid off? Figure 3a reports users' own assessments, which we find are reassuring about the quality of the data for several reasons. First, only about a third of adopters report a positive effect on profits, suggesting candid truth-telling (rather than promotional responses that might be consistent with what firms put into their brochures). Second, the reported gains line up with economic priors across industries and usage types, in particular firms in Information & Communication being most likely to report positive profits. Finally, 3b reports profitability by usage type: firms using AI for Text & Code are approximately 15% less likely to report positive

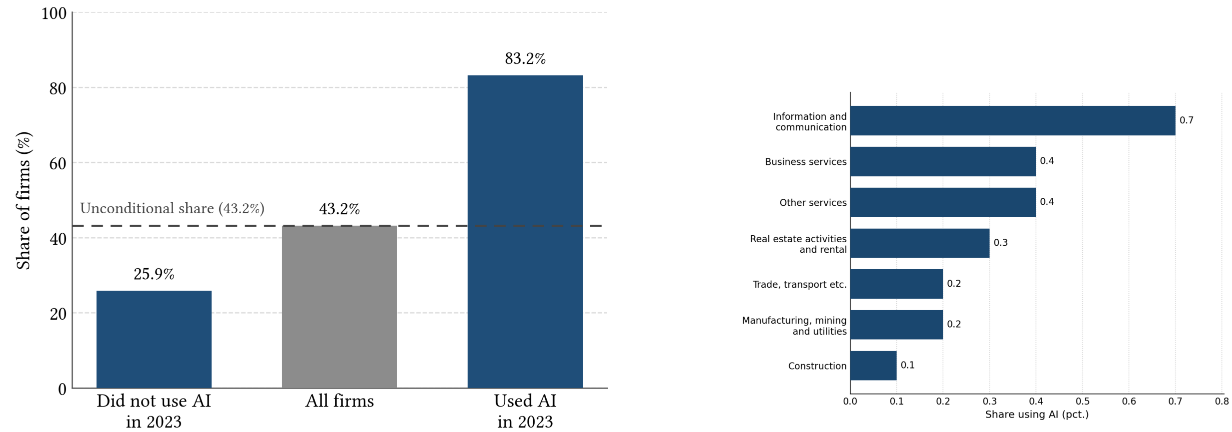

(a) Conditional usage in 2024

(b) By industry (2024)

Figure 2: What AI is used for, and where

Notes: Panel (a) reports the share of firms using AI in 2024, unconditionally and conditional on the firm's own answer in 2023. The two conditional shares are computed on the 3,117 firms responding in both the 2024 and 2025 survey waves (reference years 2023 and 2024), whereas the unconditional share is computed on the full 2024 cross-section of 4,046 firms; the dashed line marks that unconditional share, which is where all three bars would lie if the response were independent across waves. Panel (b) reports the share of surveyed firms using AI by industry in 2024. Source: Statistics Denmark IT Usage Survey; firms with $\geq$10 employees; survey-weighted.

profitability than for the other two usage types.[4]

[4] Appendix Figure A10 lists other self-reported outcomes beyond profitability.

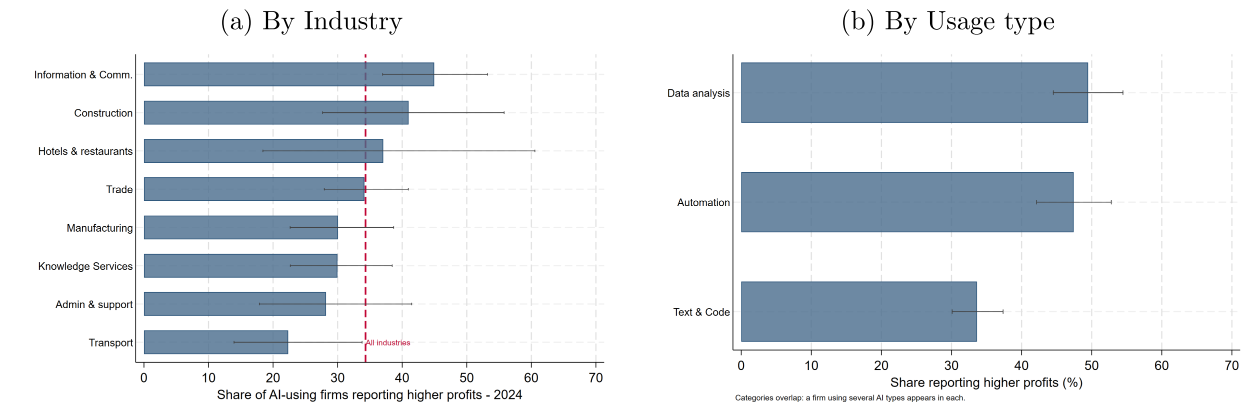

Figure 3: Profitability of AI Usage

Notes: Share of surveyed firms using AI that report it having a positive effect on profits, by industry in panel (a) and by usage type in panel (b), 2024. Source: Statistics Denmark IT Usage Survey; firms with $\geq$10 employees; survey-weighted.

3.2 Who adopts

Adoption is far from random, as Figure 4 shows. Panel (a) shows that within narrowly defined industries, firms that are larger, more productive, more skill-intensive, and more export-oriented are systematically more likely to adopt. It moreover highlights industry composition, which absorbs away a significant part of the relationship between adoption and both education and productivity.

Panel (b) cautions us that some of our predictors are highly correlated: in particular, the relationship between AI usage and value added per worker and firm age depends on what other covariates are included in the regression.

The firm's workforce education and size however remain strong predictors of AI usage, consistent with what the literature has found. The fact that large firms disproportionately adopt AI suggests fixed AI adaptation costs. This suggests that many firms are not merely giving workers access to publicly available chatbots, but are instead investing significantly in integrating artificial intelligence.

Firms differ not only in size, but also in what they produce and how they produce it. Artificial intelligence is more productive in some of these steps than in others. Most importantly, it—currently—can perform cognitive but not physical labor. However, even among cognitive tasks, it is more competent in some than others (for instance, stronger at drafting or generating code than at tasks demanding novel judgment): a good test of the validity of our adoption data would be to test whether firms that perform more AI-competent tasks are also more likely to adopt AI.

We implement this test with the task-based exposure measure of Felten, Raj, and Seamans (2023), which scores each occupation i by how much of its work AI can perform—a purely *theoretical* measure of exposure. Against it, we set *actual* exposure: the employment-weighted share of occupation i's workers employed at AI-using firms. Figure 5 plots the two.

As shown in the figure, theoretical AI exposure and actual AI adoption are highly correlated. The theoretical AI-exposure measure is standardized: we find that workers in occupations with one standard deviation higher theoretical exposure to AI are 11pp more exposed to actual AI usage.

That firms adopt precisely where the technology is theoretically most applicable is an economically sensible selection pattern and an external validation of our survey measure, benchmarked against indices built entirely outside the Danish data.

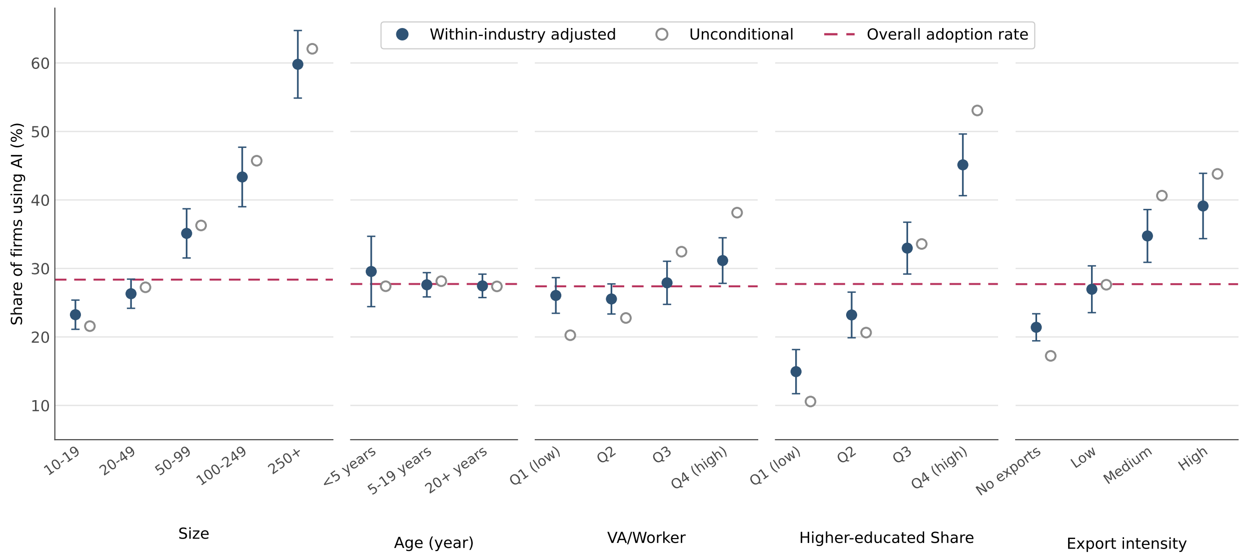

(a) Adoption rates by firm characteristics

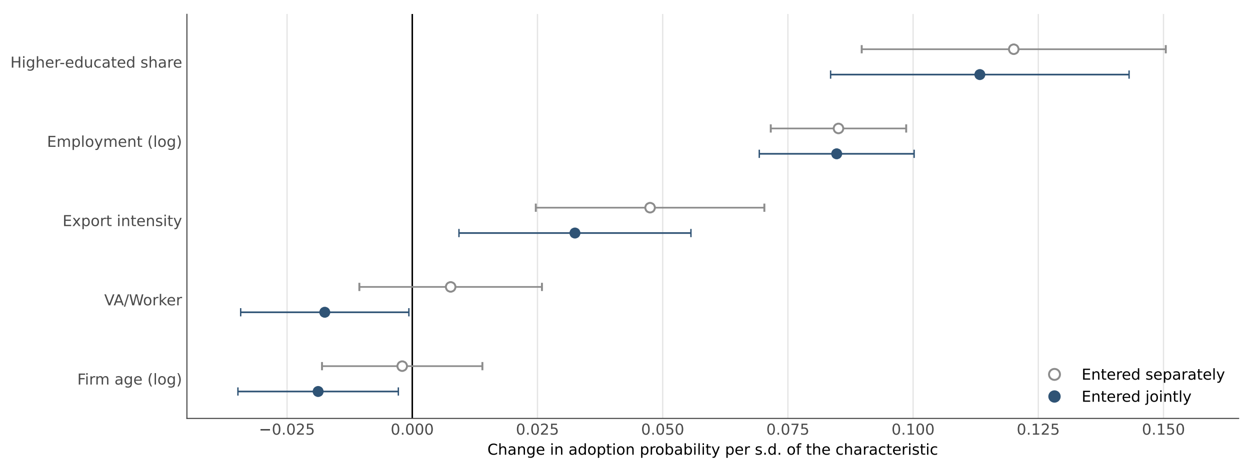

(b) Standardized regression coefficients, separate vs. joint

Figure 4: Within industries, larger and more skill-intensive firms adopt

Notes: AI adoption and firm characteristics, 2024 wave; firms with ≥10 employees; survey-weighted; three-digit-industry fixed effects; standard errors clustered by industry. *VA/Worker* is the logarithm of value added per FTE employee. Panel (a): within-industry adjusted adoption shares by bins of each characteristic (filled markers, whiskers are 95% CIs on the contrast against the weighted-average bin), against unconditional shares (open markers) and the overall adoption rate (dashed line). Panel (b): linear-probability-model coefficients on standardized characteristics, entered one at a time (open markers; N = 4,015–4,095) versus all five jointly (filled markers; N = 3,994); whiskers are 95% CIs.

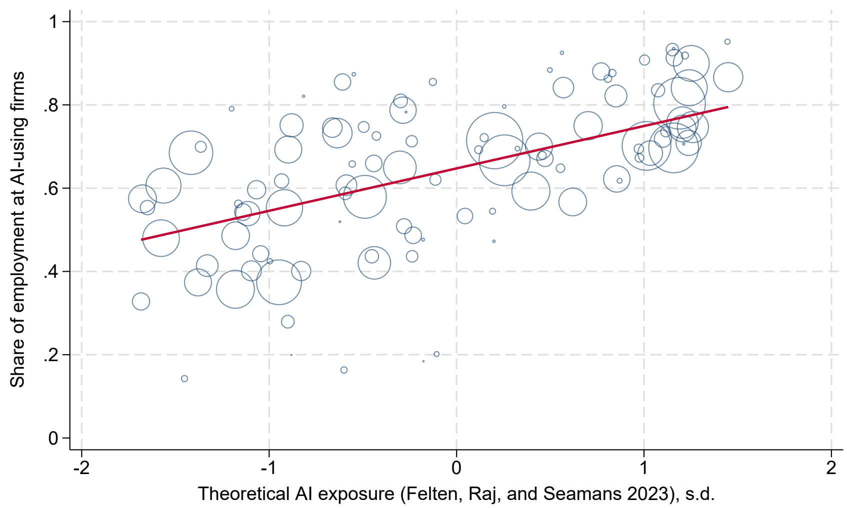

Figure 5: AI-using firms over-employ AI-exposed occupations

Notes: Each point is a three-digit occupation. Horizontal axis: standardized task-based AI-exposure score of Felten, Raj, and Seamans (2023); vertical axis: share of the occupation's employment in surveyed firms that is located at firms reporting AI use in the 2025 wave. Marker area is proportional to the occupation's employment; the red line is the employment-weighted fitted regression of the vertical on the horizontal axis.

4 Firm-level effects of AI adoption

Throughout the remainder of this section, we will now compare AI adopters—firms that report using AI in 2023 but not 2022—to non-adopters—firms that report not using AI in either year: do AI-adopting firms significantly change their employment strategy relative to firms that do not adopt AI?

Figure 6 plots firm-level total (log) employment relative to January 2023 by adoption status. Before 2023, the two groups move broadly together—both fall between 2018 and 2021 and rise thereafter—but not identically: their paths already diverge in places, notably over 2020–2022. As discussed above, adopters and non-adopters differ systematically, particularly in size, which places them on different employment trajectories. This makes the raw comparison hard to interpret: adopting firms might have grown less than non-adopters after 2023 regardless of AI, simply because their larger size put them on a slower-growing path.

The rest of the section proceeds as follows. Section 4.1 lays out the identification strategy; section 4.2 presents the average decline and shows that it operates through the hiring margin rather than separations; section 4.3 shows how the decline is distributed across adopters and how it varies by firm size; and section 4.4 examines which workers it falls on. Throughout this section we will study exclusively changes in employment: similar to Brynjolfsson, Chandar, and Chen (2025), we find that the entirety of the adjustment is coming from employment and not from wages.[5]

[5] Both adopting firm's average hourly wages and those of new hires are flat. For more information see Appendix E.

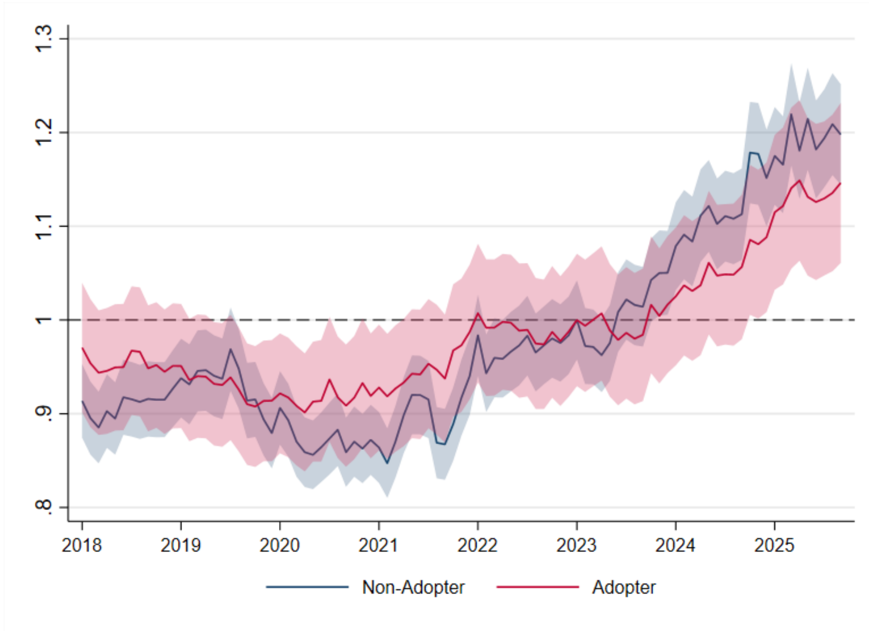

Figure 6: Raw employment growth, adopters vs. non-adopters

Notes: Unconditional (log) employment of 2023 first-time adopters versus non-adopters, without controls or detrending, relative to January 2023.

4.1 Identification

The central identification problem is that AI rolls out as an aggregate, time-varying shock while adopters are selected, so neither raw time series nor naive cross-sections isolate the effect of adoption—we therefore read our estimates as descriptive rather than necessarily causal.[6]

Our solution is to assume (and verify) that each firm's employment, on average, follows a firm-specific trend that does not change in the short-to-medium term. These trends could be defined by the firms' production functions and current state, defined for instance by industry, size, or capital intensity. While a firm might deviate from its trend in the short run (for instance, when a worker quits), it will attempt to stay on its trend in the medium term.

To operationalize this approach, we estimate each firm's trend using the years 2021 and 2022. We chose this window to exclude larger employment fluctuations around the COVID lockdowns. We then residualize each firm's monthly employment against that trend, in a regression where we further include firm-by-calendar-month fixed effects to purge away firm-specific seasonal employment effects, see equation (1).[7]

[6] We understand AI adoption as the profit maximizing decision jointly applied to adoption and employment strategy. If one were to randomly "assign" AI adoption to a different firm, the optimal employment strategy of that firm need not have been the same.

[7] Estimating the trend on the pre-adoption window, rather than including a firm-specific linear trend directly in the event study, keeps the post-2023 employment path from contaminating the trend estimate: a full-sample trend control would be partly run on post-adoption data and would absorb part of the treatment response, attenuating the estimated effect. The attenuation is severe in our setting because the estimated effect itself grows approximately linearly after adoption, so a full-sample firm trend is nearly collinear with the treatment path—the classic critique of unit-specific trends under dynamic treatment effects (Wolfers, 2006). Our two-step procedure—fit the firm-specific

$$\ln E_{f,t} = b_f\, t + \delta_{f,m(t)} + \epsilon_{f,t}, \tag{1}$$

where $b_f\, t$ a firm-specific linear time trend, and $\delta_{f,m(t)}$ a firm-by-calendar-month fixed effect capturing firm-specific seasonality. The coefficients are estimated on the pre-adoption years 2021 and 2022. One convention follows from this and applies throughout the paper: every employment effect we report is a deviation from the firm's own trend, not a change in its level. A coefficient of zero means a firm stayed on the path it was already on—typically still growing and still hiring—not that its headcount stood still.

We then estimate the standard event study on the resulting trend-residualized log employment $\tilde{E}_{f,t}$ as indicated by equation (2), where AI_f is the first-time-2023 adopter indicator and January 2023 is the reference period. The specification absorbs three-digit-industry-by-month fixed effects $\delta_{\mathrm{ind}(f)\times t}$ alongside size fixed effects δ_{size}, so that each adopter is compared only with non-adopters in the same three-digit industry—one of 196—in the same month. This differences out any shock hitting an industry in a given month—for instance, the post-2023 contraction of IT-intensive sectors—so that identification runs off variation in adoption within an industry-month. Observations are population-survey-weighted, and standard errors are clustered at the firm.

$$\tilde{E}_{f,t} = \delta_{\mathrm{ind}(f)\times t} + \delta_{\mathrm{size}(f)} + \sum_{\tau \neq 2023m1} \beta_\tau\, \mathrm{AI}_f \cdot \mathbf{1}[t = \tau] + \varepsilon_{f,t} \tag{2}$$

4.2 The average decline and its margin

Figure 7 plots the event-study coefficients associated with regression (2). Adopters and non-adopters had somewhat different employment experiences during the turbulent COVID years 2020 and 2021. However, the pre-2023 coefficients are broadly centered around zero, indicating parallel trends and providing key support for the empirical design.

Turning to the post-2023 evolution, we find that adopters' employment falls roughly 6% below pre-trend over the first two years post-adoption; by late 2025, the end of our sample, the gap reaches roughly 11%. As highlighted in Figure 6, this does not mean that absolute employment is falling at adopting firms, but it is increasing less than the pre-2023 employment trends would have implied.[8] This decline is a feature of high-frequency register data: re-estimated on annual balance-sheet employment—the measure most firm-level studies rely on—it is hard to find, and the other balance-sheet outcomes (revenue, profits, productivity, and investment) likewise show no significant effects (Appendix F).

trend and seasonal effects on untreated (pre-adoption) observations only, impute the counterfactual, and average the deviations—is the imputation estimator of Borusyak, Jaravel, and Spiess (2024), which is robust to this contamination and efficient. Note also that the residualization and the event-study fixed effects absorb distinct objects: the first removes within-firm pre-period structure, while the industry-by-month effects compare firms within an industry-month.

[8] For a numerical illustration, consider a firm at the sample average of about 50 employees, for whom its pre-2023 trend implied growth to roughly 60 employees by late 2025. This firm instead grew to about 53—some 11% below the trend-predicted level, yet still above its January-2023 headcount.

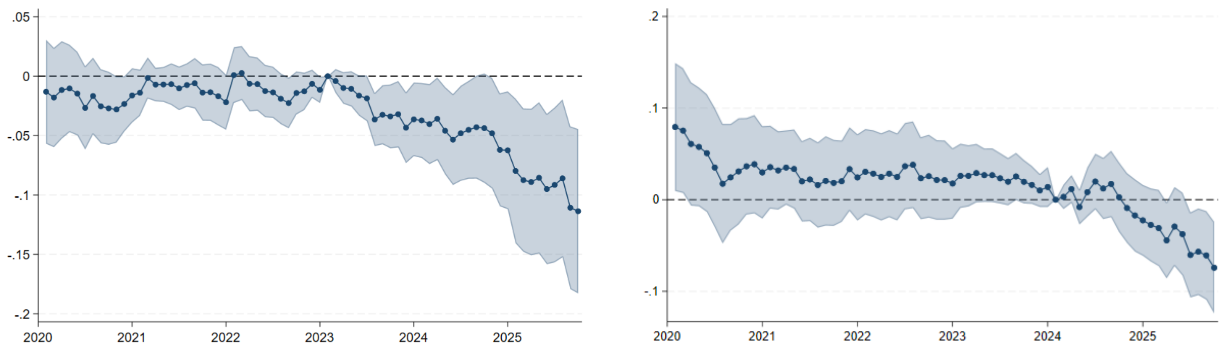

(a) 2023 adopters

(b) 2024 adopters

Figure 7: Employment around AI adoption (event study)

Notes: Event-study coefficients β_τ from equation (2); outcome is log FTE employment, reference period 2023m1. Panel (a): first-time 2023 adopters (baseline treatment). Panel (b): the same specification with first-time 2024 adopters as the treatment group, with reference period 2024m1, and firms with no reported use through 2024 as control. Three-digit-industry-by-month and size fixed effects plus a firm-specific linear trend fit on 2021–2022; survey-weighted; standard errors clustered by firm; shaded bands are 95% confidence intervals. Firms are weighted using survey weights, the results are similar to an entropy-balanced reweighting of non-adopters toward adopters' pre-2023 characteristics (Hainmueller, 2012) (see Appendix Figure A2).

Panel (b) repeats the exercise for the following cohort, first-time 2024 adopters: their coefficients stay near zero through 2023, drift downward only after 2024, and reach roughly 10% below trend by late 2025—the trajectory of the 2023 cohort at the same event time, shifted one year later. That the decline reproduces in a second adoption cohort, timed to that cohort's own adoption, is difficult to attribute to any common shock.

These estimates are likely a lower bound on the employment differences between adopters and non-adopters, because control firms increasingly adopt over the post-2023 window: the share of never-adopters and 2023 non-adopters that adopt in 2024 is 24%.

These main results were for the aggregated AI usage treatment. Appendix Figure A11 splits the result into the three main categories of AI usage: text and code generation, data, and automation, and finds that data and automation have very wide confidence bands (text generation makes up 92% of AI-using firms). Indeed, text generation has coefficients identical to our "AI usage" category, suggesting that most of our effect stems from chatbot adoption.

Two natural concerns arise that we can address. First, mean reversion: AI-exposed occupations such as programmers may have been over-hired during the COVID boom, so that their later decline reflects a correction rather than AI. This concern is most acute for aggregate occupation-exposure designs; in our cross-sectional comparison, it would bite only if over-hiring was systematically stronger at the firms that later adopted AI—within narrowly defined industries. In Appendix Figure A5 we test this by estimating pre-trends on an earlier sample, which allows us to test mean reversion already in 2022, which we do not find. [9] Another concern is industry-level shocks, which,

[9]In Appendix Figure A5 we perform an out of sample test where we refit the trends on 2020–2021 which turns 2022 into a held-out year. Indeed, the 2022 coefficients stay flat while the post-2023 gap re-emerges at −9.2 log points. Mean reversion from overfitted trends would instead show up as drift in the held-out year.

since adopters are not uniformly spread across industries, could contribute to the estimated gap. But because we difference within three-digit industry-by-month cells, such concerns are mostly mitigated. Sub-industry shocks correlated with adoption within three-digit industries cannot be ruled out, but should have a limited quantitative effect.

On average, adopting firms shrink. To what extent is this shrinkage stemming from fewer hires, and to what extent from additional layoffs? As before, we have to render our results comparable across adopters and non-adopters since they have different employment patterns. Previously, we detrended firm-level employment rates; this is not possible with hiring and separation rates which are rare events at the monthly and even the annual frequency.

We therefore reframe the question from hiring and separation rates to low- and high-tenured workers: are the changes in employment driven by changes to the stock of low-tenured workers?

Figure 8 splits firm employment into three tenure stocks—workers with under one year, under two years, and two years or more of tenure—each detrended exactly as the baseline employment outcome. The two shorter-tenure stocks are flat through 2020–2022 and then fall steadily after January 2023: the under-two-year stock reaches about -28 log points by late 2025, and the under-one-year stock falls further, to about -48. This confirms that AI-adopting firms are reducing their hiring rates significantly relative to non-adopting firms, and that the shortfall is sharpest among the most recent hires. This is consistent with economic intuition: since AI adoption is gradual and a choice, forward-looking firms (that intend to do only minor workforce adjustments) can respond to it via reduced hiring, and need not turn to layoffs.

Finally, we find that adopting firm's employment of high-tenure workers increases relative to its pre-trend. This is not mechanically due to an aging workforce but implies that AI-adopting firms retain tenured workers for longer. It is consistent with AI-adoption making high-tenure workers more productive, so adopters retain more of them than non-adopters. It also implies that the overall fall in employment understates the risk for exposed junior workers, since it is partly attenuated by additional retention of senior workers.

Can the hiring margin explain the entire fall in adoption? The stock of workers with under two years of tenure is close to a 24-month moving sum of hires, so its trend deviation approximates the average hiring shortfall over the preceding two years: the -28 log point endpoint says adopters hired roughly a quarter less than their own pre-trend implies over 2024–2025. Because the coefficients are still falling at the end of the sample—close to 0.8 log points per month throughout—the shortfall in the current flow runs ahead of the shortfall in the stock by about half a moving window, or some 10 log points, putting hiring at the end of 2025 nearer 35–40 log points below trend; this attributes the whole movement to hiring, so it is an upper bound to the extent that separations of low-tenure workers also fell.

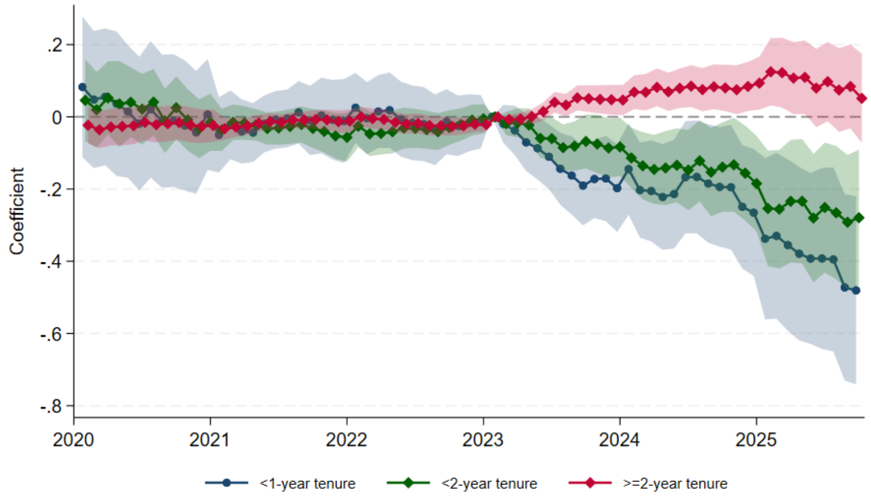

Figure 8: Employment by worker tenure of 2023 adoption (event study)

Notes: Firm employment in each month is split into three FTE stocks by worker tenure at the firm: under 12 months, under 24 months, and 24 months or more; the first two are nested cumulative stocks, and spells left-censored at the start of the worker panel are assigned to the highest-tenure stock. Each stock is taken in logs and residualized exactly as the baseline employment outcome—firm-by-calendar-month means plus a firm-specific linear trend estimated on 2021–2022—and then used as the outcome in equation (2), absorbing industry-by-month and size-class fixed effects, weighted by survey weights, with standard errors clustered by firm, reference month January 2023 and 95% confidence bands. The three lines—under-one-year, under-two-year, and two-year-or-more tenure—are the three tenure stocks. Specification, weighting, and inference otherwise as in Figure 7.

4.3 The distribution of the decline

We now show that the decline in employment is not homogeneous across all adopting firms. To this end, we split the baseline event study at 100 FTEs, measured at the pre-period base. For ease of exposition, we will hereinafter refer to firms with less (more) than 100 FTEs as "small" ("large").[10] Panel (a) in Figure 9 repeats our baseline analysis separately for the two size classes. As before, employment of small and large firms is insignificant from their pre-trends prior to 2023, after which they split: employment of small firms falls steadily below trend after 2023, whereas that of large firms continues along its trend.[11]

Through 2024, the confidence bands in Figure 9 exclude declines larger than about 2 log points at large adopters, and about 7 log points by the end of the sample: a fraction of the small adopters' 16-log-point decline by late 2025. Panel (b) plots the triple-difference event studies and finds a statistically significant average difference of 20% by late 2025.[12].

The average decline could describe two different economies: one in which every adopter runs

[10] In the 2024 survey wave, 2,295 responding firms have fewer than 100 employees and 1,555 have 100 or more.

[11] Under covariate-balancing weights, the small-firm decline is qualitatively the same but attenuates to about −9 log points by the end of the sample (Appendix Figure A3).

[12] A formal triple-difference test agrees (Appendix table 5): the small-adopter post-2023 coefficient is −0.051 (s.e. 0.024), and the large-firm interaction offsets it almost exactly—+0.044 (s.e. 0.035) in the detrended baseline and +0.064 (s.e. 0.035, significant at the ten-percent level) with industry-by-month-by-size effects—leaving large adopters' post-adoption path indistinguishable from zero

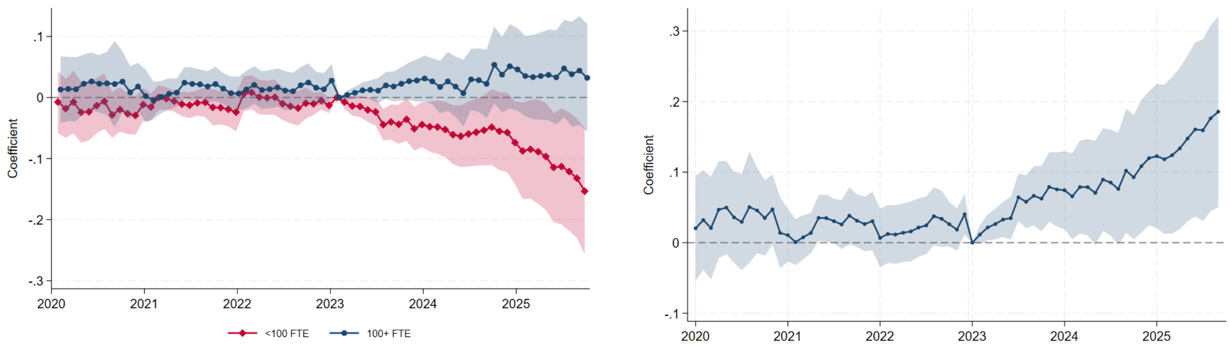

(a) By size class

(b) Large-minus-small difference

Figure 9: Employment around AI adoption, by firm size (event study)

Notes: Panel (a): event-study coefficients from equation (2) for the baseline log-FTE-employment outcome, estimated separately for 2023 adopters with fewer than 100 FTEs and with 100 or more FTEs at the pre-period base. Panel (b): the month-by-month difference between the two—the triple-interaction coefficient of adopter status, event time, and the large-firm indicator—with 95% confidence bands; it corresponds to the average post-period contrast in Appendix table 5. Survey-weighted; specification and inference as in Figure 7.

modestly below trend, and one in which a few adopters collapse while the rest are unaffected. We now test for this heterogeneity.

We first compute for each firm its average post-2023 change in its residualized employment, Δ_f:

$$\Delta_f \;=\; \frac{1}{|\mathcal{T}_{\text{post}}|}\sum_{t\in\mathcal{T}_{\text{post}}}\tilde{E}_{f,t} \;-\; \frac{1}{|\mathcal{T}_{\text{pre}}|}\sum_{t\in\mathcal{T}_{\text{pre}}}\tilde{E}_{f,t}, \tag{3}$$

where $\tilde{E}_{f,t}$ is the trend-residualized log employment of equation (1). We will loosely refer to Δ_f as the firm's employment growth (relative to pre-trends).

We then ask: how does the distribution of growth rates Δ_f differ across adopting and non-adopting firms? Since adopting and non-adopting firms are not distributed equally across industries (and differ in size), we first net out industry-by-month and size-fixed effects. We then plot the quantiles of the residualized average growth rates of adopting firms against non-adopting firms in Figure 10.

The figure shows a broad shift: employment growth is lower among most quantiles of adopters. Table 2 further documents that the divergence is not driven by the extreme tails.[13] Fast-growing adopters grow statistically indistinguishably from fast-growing non-adopters. The gap is significant through the median and closes only in the top quartile, implying that the mean gap of -0.072 is generated overwhelmingly below the 75th percentile. The same leftward shift is visible in the histograms of Δ_f: adopters lie below non-adopters through the lower and middle of the distribution, with the top quartile pinned near zero (Appendix Figure A7).[14]

[13] Note also that trimming the bottom fifth of both distributions preserves roughly three-quarters of the gap (-0.055 of -0.072 in the late window).

[14] The shift is concentrated among adopters with a high base-period young-worker share: split into young-share quartiles, the top quartile shows a 25th-percentile adopter-minus-non-adopter gap of about -0.25 against roughly

Table 2: Distribution of the firm-level decline: quantiles of Δ_f and gaps

	Non-adopters	Adopters	Gap
Panel A. Quantiles of Δ_f			
p10	-0.565^{***}	-0.661^{***}	-0.096
	(0.036)	(0.074)	(0.082)
p25	-0.166^{***}	-0.267^{***}	-0.101^{***}
	(0.016)	(0.020)	(0.025)
p50	0.063^{***}	-0.008	-0.072^{***}
	(0.010)	(0.018)	(0.021)
p75	0.272^{***}	0.253^{***}	-0.019
	(0.013)	(0.033)	(0.035)
p90	0.502^{***}	0.505^{***}	0.003
	(0.019)	(0.035)	(0.040)
Panel B. Summary			
Gap in means (headline)			-0.072^{**}
			(0.031)
Fan test: gap(p10)−gap(p90)			-0.099
			(0.087)
Firms			2968

Notes: Δ_f is a firm's average post-window deviation of detrended log FTE employment from its own 2021–2022 pre-trend, residualized on industry-by-month and size before collapsing, exactly as in the event study of Figure 7. Panel A reports quantiles of Δ_f for non-adopters and adopters and their difference (*Gap*); Panel B reports the headline *Gap in means* and the *Fan test* (the p10-minus-p90 gap contrast). The late post window is 2024m10–2025m9. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends, which are estimated once outside the bootstrap. Full-window, unresidualized, by-size, and placebo variants are in Appendix B (tables 6, 7, 8, and 9).

As with our baseline results, one could be concerned about other factors, such as mean reversion from the estimated pre-trends, being behind the gap. To address this, we conduct a placebo exercise in which we estimate pre-trends in 2020 and then compute Δ_f over 2021–20222. The resulting quantile gaps are indistinguishable from zero and reported in Appendix table 8.

What explains the distributional decline? First, we document that it is not stemming from size: Figure 11 repeats the quantile exercise separately for small and large firms plots the resulting quantiles of adopters and non-adopters against each other.[15] We see that the entirety of the effect stems from smaller firms, consistent with our earlier size splits.

A further investigation into what other characteristics beyond size could explain the divergent gap among adopters with lower growth rates is limited by the small number of adopters in each size quantile and could not be answered satisfactorily.[16]

−0.05 for the bottom quartile, while ninetieth-percentile gaps are near zero throughout (Appendix Table 10); because the cells are thin we read this as suggestive.

[15]As before, we control for industry composition by first residualizing Δ_f against an industry-by-month fixed effect (but no longer a size fixed effect).

[16]In our investigation, we tested the predictive power of worker turnover, two-sided worker throughput, the age of the work force, tertiary-educated shares, firm size, and firm age. A joint test of whether these characteristics predict

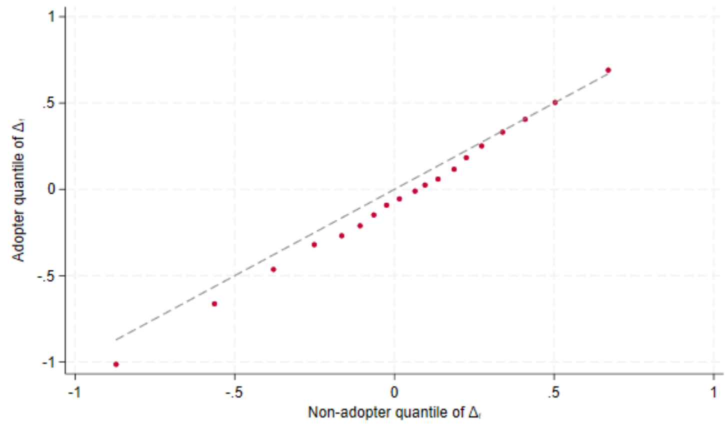

Figure 10: Quantile–quantile plot of the firm-level decline

Notes: Adopter quantiles of Δ_f plotted against non-adopter quantiles, with the 45-degree line. Points fall below the line through the lower and middle range and rejoin it in the top quartile, so the gap is concentrated below the 75th percentile while the upper tail is unmoved at the point estimates. Late post window (2024m10–2025m9), residualized on industry-by-month and size; the full-window version is Figure A8 in the appendix.

Our takeaway from this is twofold. First, a (relative-to-trend) reduction in the firm size of the average adopting firm need not imply any aggregate reductions in employment. Indeed, because the average employment decline of adopting firms is carried by small adopters, the overall employment-weighted effect of AI adoption is much smaller: weighting each firm by its employment rather than by its survey weight, the relative decline is only about 5% over two and a half years, with confidence bands that span zero throughout the post-2023 window

Second, small and large firms adopt AI in different ways, which will be relevant for the aggregate employment effects of AI adoption.

left-tail membership differently for adopters than for non-adopters cannot reject the null ($p = 0.45$), and the gap profiles across quartiles of each characteristic are non-monotone.

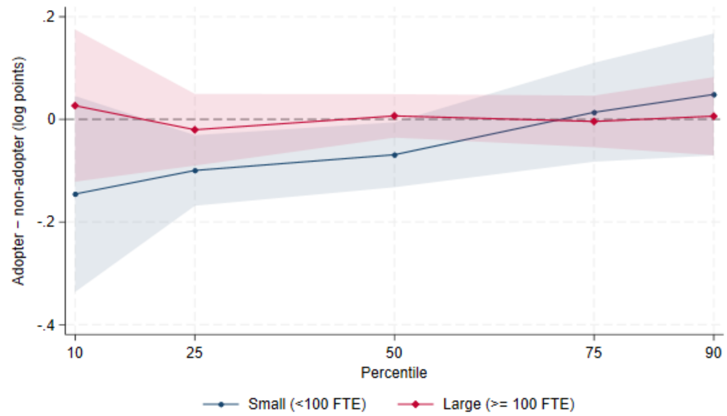

Figure 11: The distribution of the decline by size class

Notes: Adopter-minus-non-adopter gap in Δ_f at each percentile, one line per size class (small <100 FTE, large ≥100 FTE), with 95% bootstrap bands, over the late post window (2024m10–2025m9); residualized within class as class×industry×month. The small-class profile fans below zero through the lower and middle percentiles and rejoins zero at the top—p25 gap −0.100 (se 0.035), median −0.069 (se 0.032)—while the large-class profile is flat at zero throughout (p25 −0.020, se 0.036; median +0.007, se 0.022; mean +0.016, se 0.033). By-size quantile gaps are tabulated in Appendix table 9.

4.4 Where the decline falls

We now turn from firms to workers: which worker types are most affected by the employment changes? Consistent with theory and prior literature, we study age, education, and AI substitutability as relevant dimensions. We implement this by repeating the event study in equation (2) on subsets of the data separately.

Figure 12 displays the results, separately by age (Figure 12a), education (Figure 12b), and occupational AI exposure (Figures 12c and 12d). We find that most employment reductions fall on young workers (aged below 30), with a 20% decrease relative to trend by 2025. For older workers, the same decrease is only about 5%.

Turning to education, we cannot reject the hypothesis that workers with less than a tertiary education are not affected at all: the fall is concentrated on university graduates, with a relative employment fall of 10% by 2025.

Finally, we turn to occupational AI exposure. As shown in Figure 5, firms with a more AI-exposed workforce were more likely to adopt AI; we now test whether the employment decline also concentrates on those workers. We classify an occupation as AI-exposed if it falls in the top quartile of a given exposure index, and re-estimate the event study separately for exposed and non-exposed workers. Figures 12c and 12d use the indices of Felten, Raj, and Seamans (2023) and Eloundou, Manning, Mishkin, and Rock (2023): in both, the relative employment decline concentrates among AI-exposed workers—the exposed-worker path reaches roughly 30% below trend by late 2025 for the Felten measure, against about 6% for non-exposed workers. [17]

[17] We also performed the split using the use, augmentation, and automation shares of the Anthropic Economic Index (Appendix Figure A9). Against our expectations, the fall in employment is not concentrated in automation-exposed

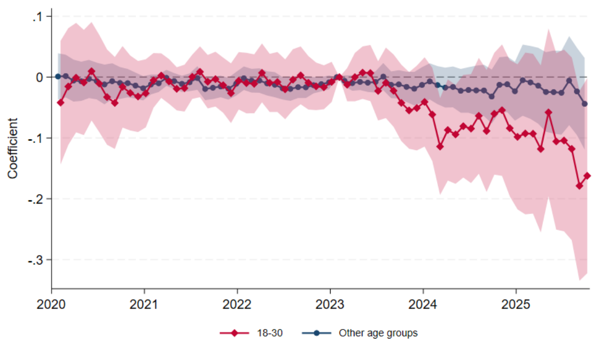

(a) By age

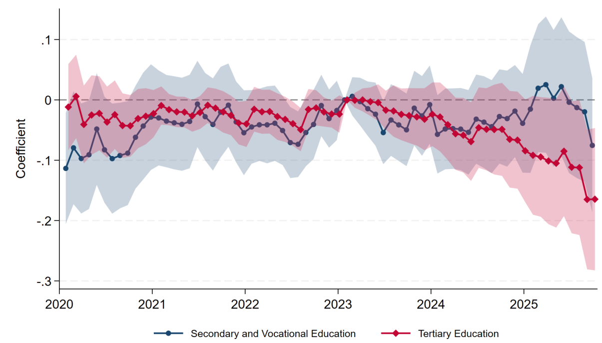

(b) By education

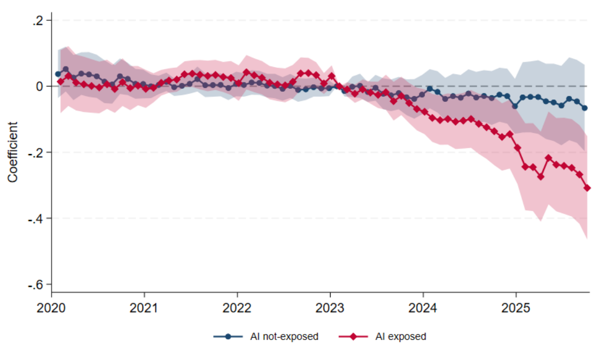

(c) By AI exposure (Felten)

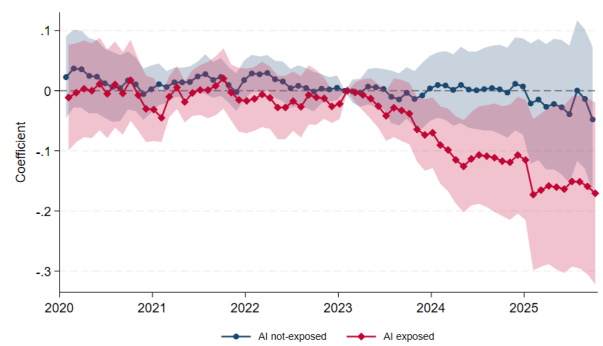

(d) By AI exposure (Eloundou)

Figure 12: Heterogeneity in the employment decline

Notes: Event-study coefficients from equation (2) estimated separately by worker age (a), education (b), and occupational AI exposure (c, d). In panels (c) and (d) an occupation is "AI-exposed" if it lies in the top quartile of the exposure index of Felten, Raj, and Seamans (2023) and Eloundou, Manning, Mishkin, and Rock (2023), respectively. Specification, weighting, and inference as in Figure 7.

The decline thus shows up along all three dimensions—age, education, and occupational exposure. The three margins overlap heavily, however: young workers are more often university educated, and university graduates are far more often in exposed occupations. Is one of them the driver? In Appendix C we race the three margins against each other, collapsing each firm's employment into the eight age-by-education-by-exposure cells. Occupational exposure emerges as the strongest predictor of where the decline falls: by late 2025, exposed employment at adopters is roughly 25% below trend against 7% for non-exposed employment, and exposure is the only margin whose contribution is significant at the one-percent level, while the effects of age and education are slim, albeit with large confidence bands.[18]

Since the employment decline is quite heterogeneous by firm size, we find it informative to split the occupation-exposure result of Figure 12 also by firm size. Figure 13 shows that small

occupations.

[18]Does occupational mobility within adopting firms also mirror the occupational hiring patterns? A within-firm decomposition (Appendix Table 11) finds only small and statistically insignificant recomposition inside firms—consistent with the decline operating through the firm-level hiring margin rather than through active reshuffling of the existing workforce, though the decomposition is underpowered to settle this.

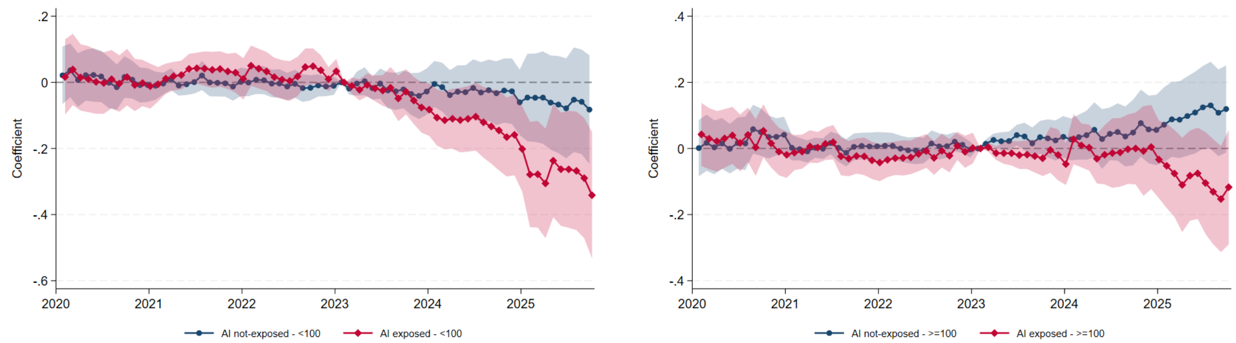

(a) Small adopters (<100 FTE) (b) Large adopters ($\geq$100 FTE)

Figure 13: The occupation-exposure decline by firm size (event study)

Notes: Event-study coefficients from equation (2) estimated separately for AI-exposed occupations (red) and non-exposed occupations (blue), shown for adopters with fewer than 100 FTEs (a) and 100 or more FTEs (b) at the pre-period base. An occupation is "AI-exposed" if it lies in the top quartile of the Felten, Raj, and Seamans (2023) exposure index. Survey-weighted; specification and inference as in Figure 7.

and large firms have very different adoption behavior. As everywhere in the paper, the changes below are relative to each firm's own trend. Small firms significantly shrink their employment in AI-exposed occupations, and their non-exposed occupation much less so (and the change is statistically insignificant from zero). Large firms shrink employment in exposed occupations much less than small firms and increase employment in non-exposed occupations. Statistically, both evolutions are insignificant from zero; their increase in non-exposed occupations is, however, significantly larger than that in exposed occupations.

Together, these results suggest that all AI-adopting firms change their employment mix away from exposed occupations: small adopters shrink their employment (predominantly in exposed occupations), whereas large firms keep their head count on trend when adopting—slowly reshuffling employment from AI-exposed to non-exposed occupations.[19]

5 Aggregate implications

Having established a firm-level employment decline, we now ask what it implies beyond our estimation sample.

5.1 The industry gradient

Our firm-level estimates yielded clean estimates by studying the subset of firms that adopted in 2023 and comparing them with non-adopters. This estimate should diverge from the overall employment effects of AI usage for two reasons. First, firms already using AI in 2022 are excluded by design; these are likely the firms with the largest gains from AI, and perhaps the largest employment

[19]We see some evidence that adopters already move already-employed workers across occupations, rather than only adjusting through hiring. However, these adjustments occur already in 2022, which is consistent with both reorganization ahead of adoption or selection. For further information, see Appendix D.

response. Second, we exclude any firm that was not surveyed; in particular tiny firms with fewer than ten FTE that are below the sampling threshold. This is worrisome because the entirety of the employment shortfalls was carried by smaller firms.

We now turn to an industry-level approach that allows us to incorporate all firms. Specifically, we study whether AI adoption has differential effects also across industries: do AI-adopting industries (i.e., those with more adopters) shrink more than non-exposed industries? These industry-level employment measures will incorporate all employment, including AI incumbents and small firms, and thus correct the shortfall of the clean difference-in-difference design from section 4.

For each of 196 three-digit industries, we measure two objects. First, s_i, its 2023 employment share in AI-adopting firms. Second, its pre-trend residualized employment. This residualization can be done either at the firm- or at the industry-level. We develop both in detail in Appendix G and plot the resulting event studies in Figure 14.

Panel (a) is the firm-level exercise of section 4, aggregated. We find that industries with more AI adoption show no employment changes (relative to pre-trends) before 2023, and a significant fall in employment post 2023, which qualitatively aligns with the firm-level results. The coefficient declines smoothly from January 2023 onward, with confidence bands excluding zero from mid-2024, and reaches about -0.16 by late 2025—an employment-growth gap of about 1.6 percentage points per 10 points of adopter exposure.

Panel (b) is the regression a researcher without firm-level data must run: raw full-register industry employment, detrended only at the industry level. Since no firm-level pre-trends are estimated, panel (b)—unlike panel (a)—retains firms with no January-2023 employment: panel (b) thus also includes potential differential effects from firm entry.

Because industry-level detrending cannot absorb firm-level seasonality, panel (b) is visibly noisier—a seasonal sawtooth in the point estimates and wider bands—but it tells the same story, declining to about -0.17 by late 2025.

The two routes agree: whether the firm-level residualization is carried through the aggregation (panel (a)) or only an industry trend is taken out (panel (b)), AI-exposed industries lose employment relative to trend after 2023, the event-study path reaching about -0.16 (panel (a)) to -0.17 (panel (b)) per unit of adopter share by September 2025, the end of our sample.

The employment weight of small 2023-AI-adopting firms is only 1.2% of total employment (1.8% extrapolating adoption to survey non-respondents): a 16% shortfall concentrated there is 0.19–0.29% of aggregate employment over more than two years—several times below the industry design's standard errors.[20] Relative to an annual hire rate of about 45% of employment (the

[20] The lower bound stems from assuming that there are no non-respondents that adopt AI, whereas the upper bound assumes that AI adoption among respondents is representative for the entire population. We estimate these bounds as follows. For the lower bound: Firms under 100 FTEs hold 17.7% of identified adopters' employment—slightly more than the 14% of employment they hold in the matched sample as a whole. Identified adopters—sampled respondents with the adoption pattern—hold 6.5% of full-register employment in the January-2023 base month (125,632 of 1.92 million FTEs), so small-adopter employment is about 1.15% of the total, or 22,229 FTEs; a 16% shortfall on that base is about 3,600 FTEs, or 0.19% of aggregate employment. For the upper bound, we assume that survey response is random within the frame, so that the survey weights extrapolate adoption to non-respondents: adopters then hold about 10.4% of register employment (roughly 200,000 FTEs), small adopters 1.8% (about 35,400 FTEs), and a 16%

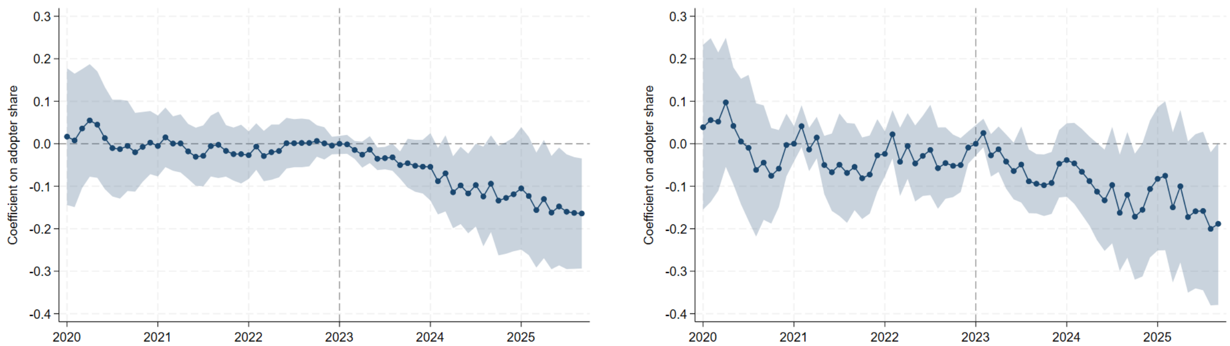

(a) Firm-level trend correction, aggregated

(b) Industry-level detrending

Figure 14: Industry employment growth by AI exposure, full register

Notes: Coefficients from an industry-month event-study regression of detrended full-register employment growth on adopter exposure s_i interacted with month indicators, normalized to January 2023 (equation (5)). Panel (a) residualizes each firm's employment exactly as the firm-level design—a firm-specific linear trend and firm-by-calendar-month effects estimated on 2021–2022, equation (1)—before aggregating to three-digit industries; entrants, which have no pre-period, drop out of this panel. Panel (b) applies no firm-level correction and instead removes an industry-level log-employment trend fitted over January 2021–January 2023 (allocated to incumbent groups by January-2023 employment shares for the group outcomes); it retains firms with no January-2023 employment, so entry enters industry growth. Both panels are weighted by January-2023 industry employment, absorb industry and month fixed effects, and cluster standard errors by industry; shaded bands are 95% confidence intervals. The baseline exposure measure is constructed with survey weights; constructing the shares with employment weights instead yields a very similar path (Appendix Figure A1).

2023–2024 average), hires over the two-year window amount to roughly 90% of employment, so the cumulative footprint is about 0.21–0.33% of the hires made over the same period. This already indicates that the industry-level exercise captures components beyond 2023 adoption alone. The same point shows up in the full population: young-worker employment rates and wages show no divergence after 2023 (Appendix Figure A12), as an effect of this magnitude would be well inside the noise of an aggregate series.

5.2 Adoption versus composition

So, how do the industry-level coefficients compare to the previously estimated (employment-weighted) firm-level coefficients? The two are in the same units: if adoption lowers adopters' employment by τ relative to non-adopters within an industry, and industries differ only in their adopter share, industry growth inherits $s_i \cdot \tau$ and the gradient estimates $\gamma \approx \tau$. At the January-2025 cross-section, the gradients are -0.105 (panel a) and -0.083 (panel b), while the employment-weighted firm coefficient stands at about -1.5 log points (with confidence bands spanning zero). In other words, we estimate industry-level effects that are much stronger than their firm-level counterparts, albeit not statistically significant.

To sharpen this comparison, we ask what exposure gradient the firm-level adoption effects alone would generate: holding each industry's baseline size-by-adoption composition fixed and projecting

shortfall is about 5,700 FTEs—0.29% of aggregate employment.

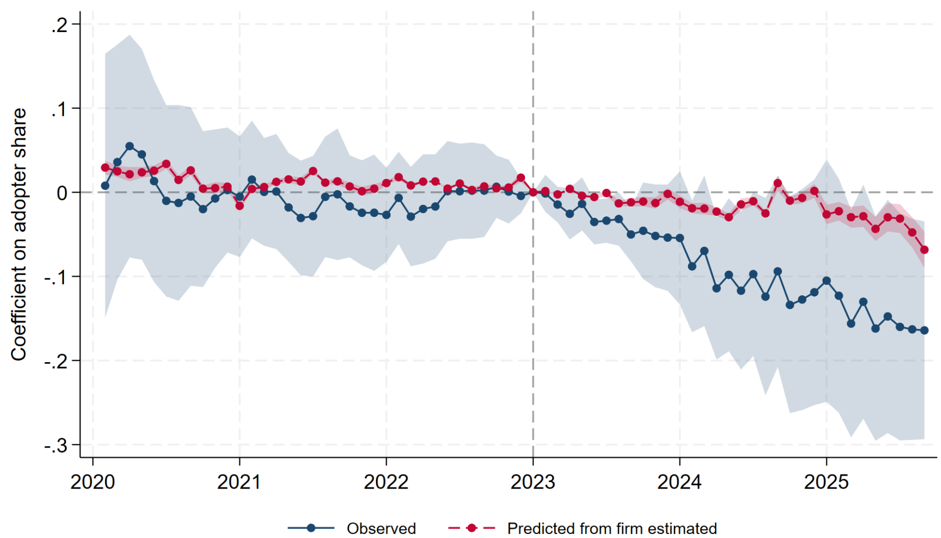

Figure 15: Micro-to-macro reconciliation: observed versus predicted exposure gradient

Notes: The observed series is the industry event-study gradient on adopter exposure s_i, estimated on the full-register, firm-level-detrended industry series of panel (a) of Figure 14 (reference January 2023). The predicted series regresses the constructed component C_{it}^{AI} of equation (6)—size-specific adopter employment shares times the size-specific within-industry adopter contrasts $\widehat{\tau}_{Ak,t}$ from the firm event studies—on s_i in the same specification. The predicted path is zero before 2023 by construction. Shaded areas are 95% confidence bands; industry and month fixed effects; weighted by January-2023 industry employment; standard errors clustered by industry.

it with the within-industry adopter contrasts yields a predicted component C^{AI} (constructed in Appendix G.2, equations (6)–(7)).

Indeed, we find that the estimated coefficient γ_τ is about four times larger than the adoption-consistent coefficient (Figure 15 and table 3). At the January-2025 horizon, the size-specific contrasts are $\widehat{\tau}_{AS} = -0.090$ for small and $\widehat{\tau}_{AL} = +0.016$ for large adopters, and the projection of the predicted component on s_i weights them 0.40 and 0.60.[21] The predicted gradient is -0.026, against an observed gradient of -0.105—the January-2025 cross-section of the same full-register series featured in panel (a) of Figure 14: the adoption effects, correctly aggregated, account for about a quarter of the exposure gradient. The overlay makes the point visually: the predicted path is flat at zero before 2023 by construction and reaches only -0.07 by the end of the sample, against -0.16 observed, because the group that dominates exposure—large adopters—has no employment response to predict.

A quarter of the industry-level gradient is implied by 2023 adoption. If the wider net cast by the industry exercise is working as intended, the remaining three quarters should be the broader AI us-

[21] The projection weights are the slopes of the size-specific adopter shares on total exposure, which sum to one; small adopters account for 40% of the cross-industry *variation* in exposure despite holding only about 17.7% of adopter employment. We do not pursue a by-size version of the industry exercise. Splitting the exposure share s_i by adopter size is uninformative: small- and large-adopter exposure are nearly unrelated across industries, the only non-mechanical term—the adopter differential—is approximately zero everywhere, and unclassified firms, which are 46.2% of base employment and mostly fall below the survey size frame, load heavily on the small-adopter exposure measure, so it proxies industry small-firm intensity rather than small-adopter exposure.

Table 3: Micro-to-macro reconciliation at January 2025

Panel A: Firm-level inputs		
	Firm contrast $\hat{\tau}_{Ak}$	Projection weight π_k
Small ($<$ 100 FTE)	-0.090	0.397
Large ($\geq$ 100 FTE)	0.016	0.603
Panel B: Reconciliation		
Predicted gradient $\pi_S\hat{\tau}_{AS} + \pi_L\hat{\tau}_{AL}$	-0.026 (0.021)	
Observed gradient	-0.105	
Share explained	0.250	

Notes: The predicted gradient is $\pi_S\hat{\tau}_{AS} + \pi_L\hat{\tau}_{AL}$, where the projection weights π_k are the slopes of the size-k adopter employment share on total exposure s_i across industries (they sum to one) and $\hat{\tau}_{Ak}$ are the January-2025 within-industry adopter contrasts from the firm event studies, estimated separately for firms below and above 100 FTEs. The observed gradient is the January-2025 cross-section of the full-register, firm-level-detrended industry series of panel (a) of Figure 14—the same object featured there; at a single cross-section the with- and without-industry-fixed-effects variants coincide. Weighted by January-2023 industry employment; standard errors clustered by industry.

age the experiment excludes—incumbent 2022 users and the unsurveyed firms. The decomposition lets us check each account against its own cell.

An exact accounting We split industry employment growth into the additive contributions of five adoption-status groups—never-adopters, 2023 adopters, 2022 users, firms the survey does not classify (below the size frame or non-responding), and net entry—which sum exactly to the gradient. We again focus on the January-2025 cross-section, so every coefficient below is directly comparable to the -0.083 gradient quoted above, and ask how the five adoption-status groups contribute to the overall coefficient.

With this machinery in place, we can now decompose the industry-level regression results for January 2025. Table 4 lists the decomposition for the five groups; its first row reports the total contributions γ_g^c (denoted β_g in the table). It assigns -0.075 to 2023 adopters—about 91% of the net cross-sectional coefficient—with unclassified firms contributing -0.036, 2022 users -0.010, and net entry -0.004, against an offsetting $+0.042$ from never-adopters. The two shares are consistent: the 91% is the adopters' *total* contribution, which includes the mechanical fact that exposed industries simply contain more adopter employment (composition), whereas the one-quarter of Figure 15 is the part that *adoption effects*—the within-industry adopter-versus-non-adopter contrasts—can generate.

Do 2023-adopters contribute to the gradient because there are more of them in AI-exposed industries (by construction), or because adopters in AI-exposed industries shrink more than adopters in non-exposed industries? The remaining rows of table 4 answer this by splitting each group's total contribution into three parts: a *differential* term—the group grows differently in more-exposed industries—a *composition* term—more-exposed industries contain more of the group—and a *re-*

mainder, the interaction of the two.[22]

The lower rows show that the adopters' differential term is −0.006, essentially zero: adopters do not shrink more in exposed industries than elsewhere—their contribution is composition, more of them sitting in exposed industries.

We now turn to firms that, by design, are excluded from the firm-level exercise. First, firms that were already using AI in 2022. These firms make up a significant share (27% of employment), and we do find that they have differential employment falls in AI-exposed industries—consistent with the overall finding—a differential term of −0.043. Their net addition to the aggregate coefficient is nonetheless only −0.010: incumbent users are underweight in exposed industries, so an offsetting composition term of +0.059 absorbs most of their differential.

Second, firms that we could not classify since they did not respond to the survey (or were not surveyed to begin with). These make up 46% of overall employment and lower the aggregate gradient by 0.0355. Can these effects be driven by AI? The decomposition shows that their differential is in fact positive (0.03): unclassified firms grow, on average, more in AI-exposed industries than in non-exposed industries. These firms contribute negatively to the overall gradient primarily because of their composition: AI-exposed industries host more unclassified firms, which on average have a more negative growth rate (-0.2 relative to pre-trends) than the average firm.

The measured slope for unclassified firms (0.03) is a net slope; the group might be a mixture of small firms that truly adopt AI (consistent with our size-based findings, where the employment response concentrates at small adopters) and firms whose decline has nothing to do with AI.[23] Because only the net is observable, the cell cannot separate the two. The industry variation thus neither supports counting the unsurveyed block as AI nor rules out AI effects within it.

What about absorption from never-adopters? Their total contribution is positive (+0.042), but the absorption-relevant cell is their differential growth in more-exposed industries, which only amounts to +0.013. Given the large uncertainty around the total contribution, we conclude that there is no evidence that non-adopters grow more in the industries where adopters cut more.[24]

Overall, the industry-level design measures a four times larger gradient than the firm-level exercise. One-eighth of this is already-AI-using firms; half of this is unclassified firms, which may or may not be due to AI usage. Of equal magnitude is an offsetting positive contribution

[22]The technical details behind the split are in Appendix G.3.

[23]This is not a remote possibility: if survey response is random, the survey-weighted adoption rates imply that unidentified adopters account for just under a tenth of the unclassified block's employment. Such a mixture, however, would leave a testable trace. If hidden adopters among the unsurveyed were located the way measured small-firm adoption is located—and measured small-adopter exposure alone accounts for 40% of the cross-industry variation in s_i—their decline would pull the group's growth slope on s_i negative. The observed slope is positive—so any such negative response is either small or netted out by the rest of the group.

[24]The decomposition can only detect absorption that happens inside the same industry—a never-adopter in an exposed industry growing faster. However, absorption need not happen there. The margin of adjustment is foregone hiring rather than layoffs (section 4.2), so there is no pool of displaced workers that enters any particular industry; the workers otherwise hired at adopters would have taken jobs elsewhere, spread across the whole economy. Any such economy-wide absorption raises employment a little everywhere, and the monthly fixed effects absorb exactly that kind of common movement. So even if every foregone hire were absorbed one-for-one elsewhere, the exercise would still register nothing: the design is informative about within-industry reallocation only.

Table 4: Contributions to industry employment growth, Jan. 2023–Jan. 2025

	Never Adopter	2023 Adopter	2022 User	Unclassified	Net Entry
Total contribution β_g	0.0419	-0.0749	-0.0102	-0.0355	-0.0039
	(0.0175)	(0.0152)	(0.0594)	(0.0560)	(0.0272)
Differential	0.0132	-0.0062	-0.0430	0.0307	.
Composition	0.0233	-0.0590	0.0588	-0.0399	.
Remainder	0.0055	-0.0098	-0.0261	-0.0262	.
Memo: base weight	0.1640	0.0981	0.2760	0.4619	.
Memo: mean growth	-0.1109	-0.1063	-0.1094	-0.2068	.
Memo: weight slope	-0.2099	0.5545	-0.5377	0.1931	.
Aggregate coefficient			-0.0826		
			(0.0857)		

Notes: Each column reports how one group of firms contributes to the employment growth of the industries it operates in, as a function of those industries' AI-adopter employment share s_i. Industry employment growth is the sum of five group contributions, $g_{i,t} = \sum_g c_{ig,t}$ with $c_{ig,t} = (E_{ig,t} - E_{ig,b})/E_{i,b}$ and b the base month January 2023 (equation (8)). The groups are firms never reporting AI use, firms first adopting in 2023, firms already using AI in 2022, firms in the register that the survey does not classify, and firms with no employment in the base month (net entry). *Total contribution* is the coefficient β_g from regressing that group's contribution on s_i; because the contributions add up and all five regressions use the same sample and weights, the β_g sum to the aggregate coefficient in the final row, which is the coefficient from regressing total industry growth on s_i. The next three rows split each β_g as in equation (10). *Differential* is the group's mean base employment weight times the slope of its own growth rate on s_i: the part arising because the group grows differently in more exposed industries than in less exposed ones. *Composition* is the group's mean growth rate times the slope of its base employment weight on s_i: the part arising because more exposed industries simply contain more of that group. *Remainder* is the interaction of the two deviations. Net entry has no base-month employment and hence no weight to split. The memo rows report the three ingredients of that split—the mean base weight $\bar{\omega}_g$, the mean growth rate $\Delta\bar{e}_g$, and the slope of the weight on s_i. The mean base weight is the employment-weighted mean across industries of the group's within-industry base-employment share, not the group's economy-wide share of base FTE (a ratio of sums): the two differ—9.8% versus 6.5% for 2023 adopters—because the industry weights used in the average are not proportional to the group's own base employment. All regressions are weighted by base (January-2023) industry employment; robust standard errors in parentheses. Industry trends are fitted in logs over January 2021–January 2023 and allocated to incumbent groups by January-2023 base employment shares.

from never-adopters. We conclude that industry-level designs—Tucker (2026), for instance, who finds that regression-adjusted early-career employment in the most AI-exposed quintile of U.S. industry-state cells fell 12% over the ten quarters following ChatGPT's release—are a reasonable approximation to the overall usage of AI: they find larger effects than firm-level estimates because they include some AI usage that clean experiments may not; at the same time, some of their effects may be driven by pure composition across industries.

6 Conclusion

We find that 2023 and 2024 AI adopters shrink relative to non-adopters. These results are descriptive comparisons between self-selected groups of firms. Adoption is a choice, so our results have to be understood through a framework in which firms jointly decide to adopt AI and change their employment structure.

Our findings are consistent with occupation-exposure designs, which find declines consistent with ours in sign and incidence: Brynjolfsson, Chandar, and Chen (2025) find about a 19% relative decline for the youngest workers in the most exposed U.S. occupations, against our roughly 20% for young workers at adopters—with the estimand caveat that ours is an adopter-versus-non-adopter partial-equilibrium contrast.

Our industry-level exercise shows that the measured adoption effects are small relative to the aggregate labor market churn, mostly because the effects are largely driven by small firms. This size heterogeneity contrasts the adoption decision: larger firms are more likely to adopt, but smaller adopters are more likely to shrink (relative to pre-trends).

Finally, Denmark provides an upper bound on how quietly these firm-level differences may pass through a labor market. Its flexicurity system combines low firing costs with high churn, an environment in which adjustment through reduced hiring can be absorbed particularly smoothly. Less fluid labor markets may meet the same firm-level differences with louder aggregate consequences.

However, even in Denmark, widespread adoption could strain its existing absorption capacity. These effects stem from early adopters, in times when future cost and productivity of AI are still uncertain. It remains important to monitor the employment effects of AI going forward and, if necessary, strengthen labor markets with active interventions such as occupational retraining.

References

Acemoglu, Daron and Pascual Restrepo (2019). “Automation and New Tasks: How Technology Displaces and Reinstates Labor”. In: *Journal of Economic Perspectives* 33.2, pp. 3–30.

Acemoglu, Daron et al. (2022). “Artificial Intelligence and Jobs: Evidence from Online Vacancies”. In: *Journal of Labor Economics* 40.S1, S293–S340.

Aghion, Philippe et al. (2025). “How Different Uses of AI Shape Labor Demand: Evidence from France”. In: *AEA Papers and Proceedings* 115, pp. 62–67.

Alekseeva, Liudmila et al. (2021). “The Demand for AI Skills in the Labor Market”. In: *Labour Economics* 71, p. 102002.

Andersen, Torben M and Michael Svarer (2005). *The Danish Flexicurity Model*. Tech. rep. Employment Relations Research Centre, University of Copenhagen.

Babina, Tania et al. (2023). *Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition*. Working Paper 31325. National Bureau of Economic Research.

Babina, Tania et al. (2024). “Artificial Intelligence, Firm Growth, and Product Innovation”. In: *Journal of Financial Economics* 151, p. 103745.

Bick, Alexander, Adam Blandin, and David J. Deming (2025). “The Rapid Adoption of Generative AI”. In: *Management Science*. Forthcoming.

Bonfiglioli, Alessandra et al. (2025). “Artificial Intelligence and Jobs: Evidence from US Commuting Zones”. In: *Economic Policy* 40.121, pp. 145–194.

Bonney, Kathryn et al. (Apr. 2024). *Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey*. Working Paper 32319. National Bureau of Economic Research.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024). “Revisiting Event-Study Designs: Robust and Efficient Estimation”. In: *Review of Economic Studies* 91.6, pp. 3253–3285.

Bredgaard, Thomas, Flemming Larsen, and Per Kongshøj Madsen (2009). *Denmark: Flexibility, Security and Labour Market Success*. Tech. rep. 53. International Labour Organization.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025). “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence”. In: *Stanford Digital Economy Lab Working Paper*.

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond (2025). “Generative AI at Work”. In: *Quarterly Journal of Economics* 140.2. NBER Working Paper 31161, April 2023, pp. 889–942.

Calvino, Flavio and Luca Fontanelli (2023). *A Portrait of AI Adopters across Countries: Firm Characteristics, Assets and Productivity*. OECD Science, Technology and Industry Working Papers 2023/02. OECD.

del Río-Chanona, R. Maria et al. (2025). *AI and Jobs: A Review of Theory, Estimates, and Evidence*. Working Paper arXiv:2509.15265. arXiv.

Demirci, Ozge, Jonas Hannane, and Xinrong Zhu (2025). “Who Is AI Replacing? The Impact of Generative AI on Online Freelancing Platforms”. In: *Management Science* 71.10, pp. 8097–8108.

Eloundou, Tyna et al. (2023). “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models”. In: *arXiv preprint arXiv:2303.10130*.

European Commission (2024a). *Digital Economy and Society Index 2024: Denmark*. Tech. rep. European Commission.

— (2024b). *The State of the Digital Decade 2024: Denmark*. Tech. rep. European Commission.

Felten, Edward, Manav Raj, and Robert Seamans (2023). “Occupational Heterogeneity in Exposure to Generative AI”. In: *AEA Papers and Proceedings* 113, pp. 282–287.

Frank, Morgan R. et al. (2026). *AI-exposed Jobs Deteriorated before ChatGPT*. Working Paper arXiv:2601.02554. arXiv.

Hainmueller, Jens (2012). “Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies”. In: *Political Analysis* 20.1, pp. 25–46.

Hampole, Menaka et al. (Feb. 2025). *Artificial Intelligence and the Labor Market*. NBER Working Paper 33509. Revised September 2025. National Bureau of Economic Research.

Handa, Kunal et al. (2025). *Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations*. Anthropic Economic Index Working Paper arXiv:2503.04761. Anthropic.

Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024). “The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market”. In: *Organization Science* 35.6, pp. 1977–1989.

Humlum, Anders and Emilie Vestergaard (May 2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. Working Paper 33777. National Bureau of Economic Research.

Kreiner, Claus Thustrup and Michael Svarer (2022). "Danish Flexicurity: Rights and Duties". In: *Journal of Economic Perspectives* 36.4, pp. 81–102.

Lodefalk, Magnus et al. (Mar. 2026). *Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring*. Ratio Working Papers 388. Ratio Institute.

McElheran, Kristina et al. (2024). "AI Adoption in America: Who, What, and Where". In: *Journal of Economics & Management Strategy* 33.2, pp. 375–415.

McKinsey & Company (Oct. 2019). *An AI Nation: Harnessing the Opportunity of Artificial Intelligence in Denmark*. Tech. rep. McKinsey & Company.

Noy, Shakked and Whitney Zhang (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence". In: *Science* 381.6654, pp. 187–192.

OECD (2025). *Progress in Implementing the European Union Coordinated Plan on Artificial Intelligence: Denmark*. Tech. rep. OECD.

Restrepo, Pascual (2024). "Automation: Theory, Evidence, and Outlook". In: *Annual Review of Economics* 16, pp. 1–25.

Tucker, Lee C. (Apr. 2026). *You're (not) Hired: Artificial Intelligence and Early Career Hiring in the Quarterly Workforce Indicators*. Working Paper CES-WP-26-27. U.S. Census Bureau, Center for Economic Studies.

Webb, Michael (2020). "The Impact of Artificial Intelligence on the Labor Market". In: *Working Paper*.

Westby, Samuel, Alicia Sasser Modestino, and Peiran Cheng (2025). *Generative AI and the Redefinition of Entry-Level Software Work*. IZA Discussion Paper 18723. IZA – Institute of Labor Economics.

Wolfers, Justin (2006). "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results". In: *American Economic Review* 96.5, pp. 1802–1820.

Yotzov, Ivan et al. (Feb. 2026). *Firm Data on AI*. NBER Working Paper 34836. National Bureau of Economic Research.

A Additional figures

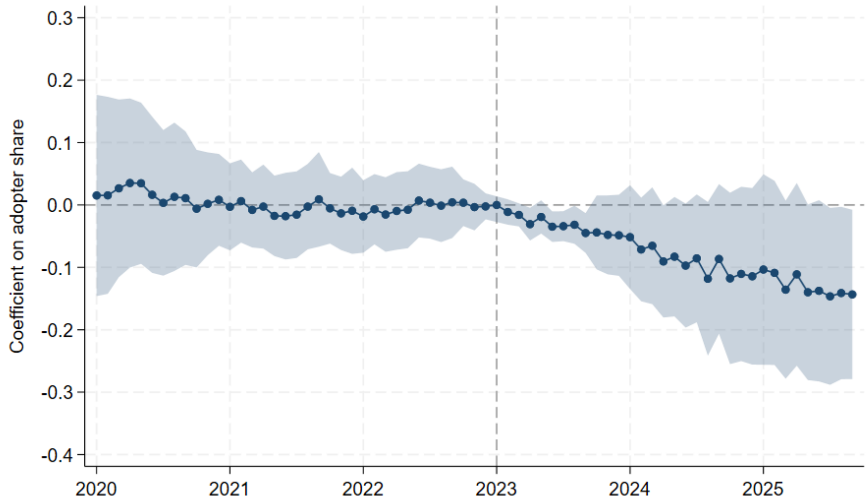

Appendix Figure A1: Industry event study with employment-weighted exposure shares

Notes: The industry-level event study of Figure 14(a) with the adopter-exposure shares s_i constructed from employment weights instead of the baseline survey weights. Coefficients on s_i interacted with month indicators, reference January 2023; industry and month fixed effects; weighted by January-2023 industry employment; standard errors clustered by industry; shaded bands are 95% confidence intervals. The path is very similar to the baseline: flat before 2023 and declining to about -0.14 by late 2025.

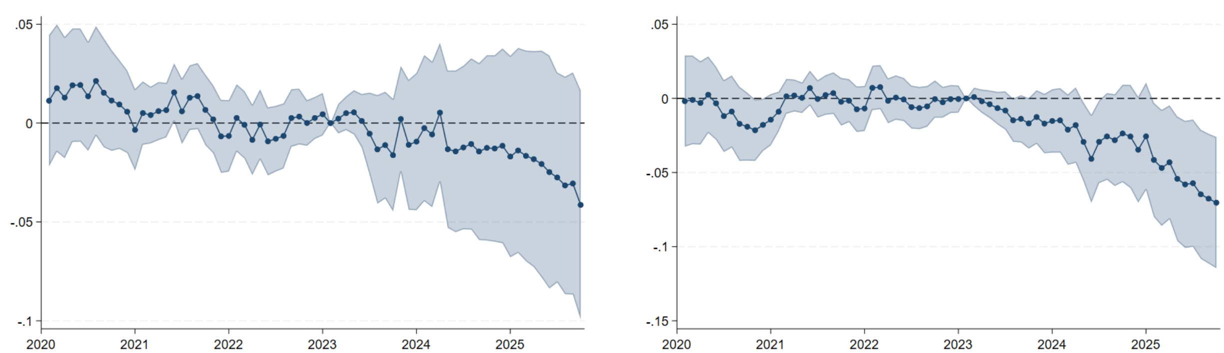

(a) Employment weights

(b) Entropy-balanced weights

Appendix Figure A2: Employment decline, alternative weightings

Notes: Employment event study from equation (2) under employment weights (a) and synthetic entropy-balanced weights (b; Hainmueller 2012). Specification and inference as in Figure 7.

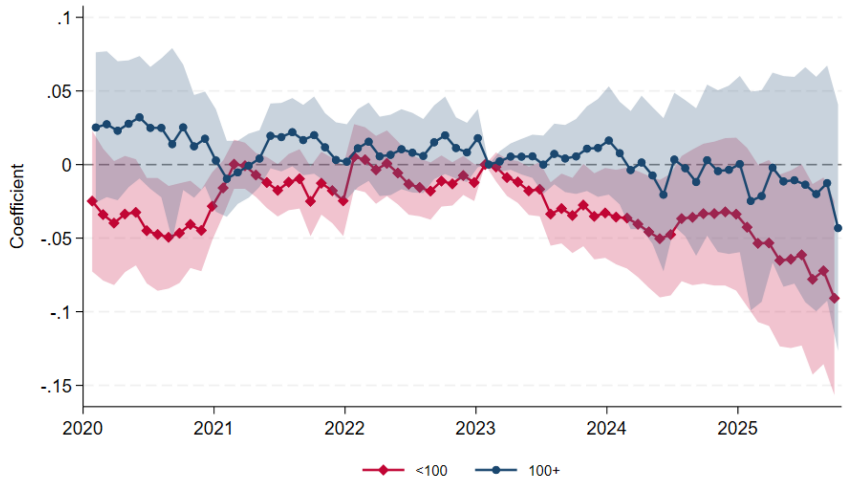

Appendix Figure A3: Employment around AI adoption, by firm size, entropy-balanced weights

Notes: Event-study coefficients from equation (2) for the baseline log-FTE-employment outcome, estimated separately for 2023 adopters with fewer than 100 and 100 or more FTEs at the pre-period base, under entropy-balanced weights (Hainmueller 2012). Specification and inference as in Figure 7.

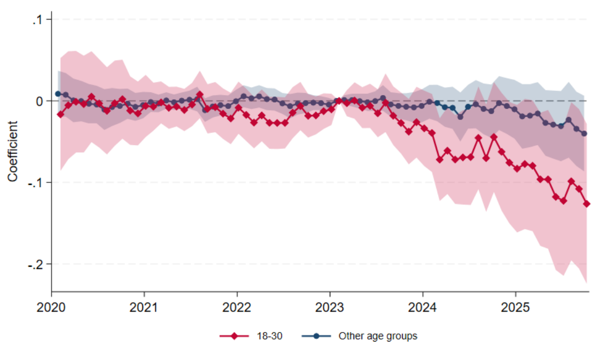

(a) By age

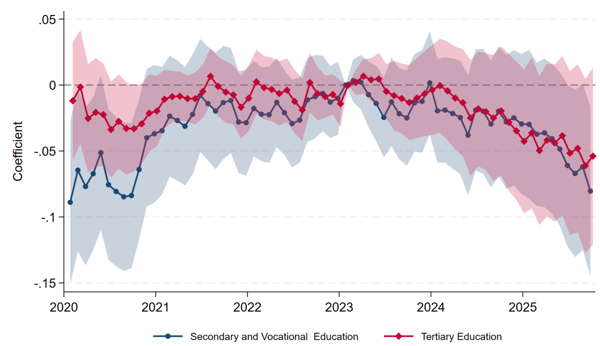

(b) By education

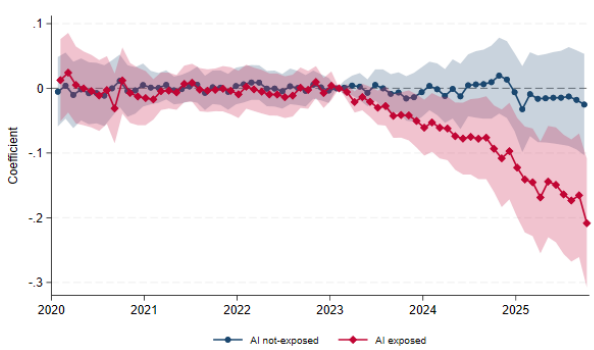

(c) By AI exposure (Felten)

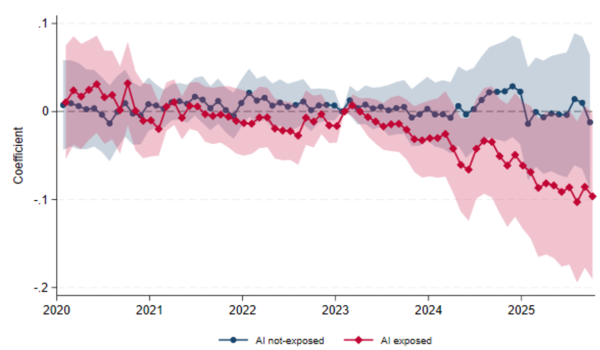

(d) By AI exposure (Eloundou)

Appendix Figure A4: Heterogeneity, entropy-balanced weights

Notes: Heterogeneity event study from equation (2) under entropy-balanced weights, mirroring the panels of Figure 12; employment-weighted estimates are analogous. Specification and inference as in Figure 7. The age, education, and occupational-exposure heterogeneity is robust to the weighting scheme (entropy-balanced shown; employment-weighted analogous).

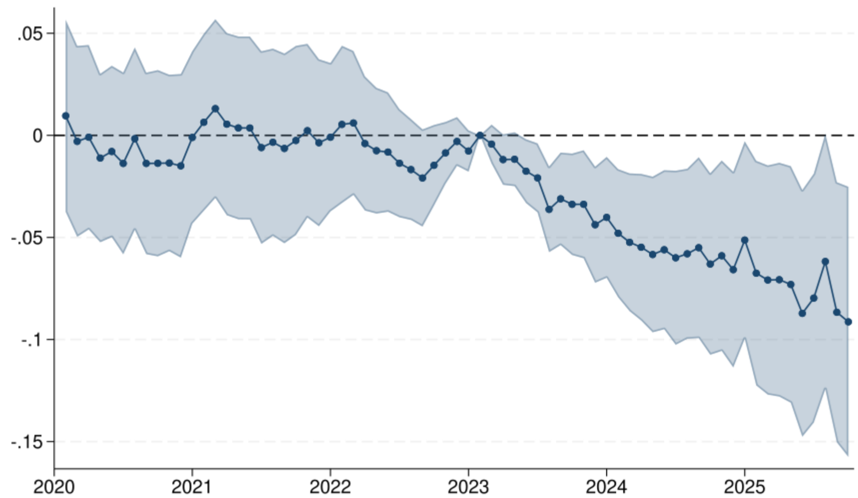

Appendix Figure A5: Baseline employment event study with firm trends fit on 2020–2021

Notes: Event-study coefficients from equation (2); outcome is log FTE employment residualized as in equation (1), but with the firm-specific linear trend and calendar-month effects estimated on 2020–2021 instead of 2021–2022, so that calendar 2022 is out of sample. Reference period 2023m1; three-digit-industry-by-month and size fixed effects; survey-weighted; standard errors clustered by firm; 95% confidence bands. Specification otherwise as in Figure 7. If adopters were merely selected on unusually steep growth (e.g. overhiring during covid), the extrapolated trends would miss in the held-out year and open a spurious gap there. Instead the 2022 coefficients remain flat within about 1.5 log points of zero, and the post-2023 gap re-emerges at -9.2 log points by late 2025, against -11.3 in the baseline.

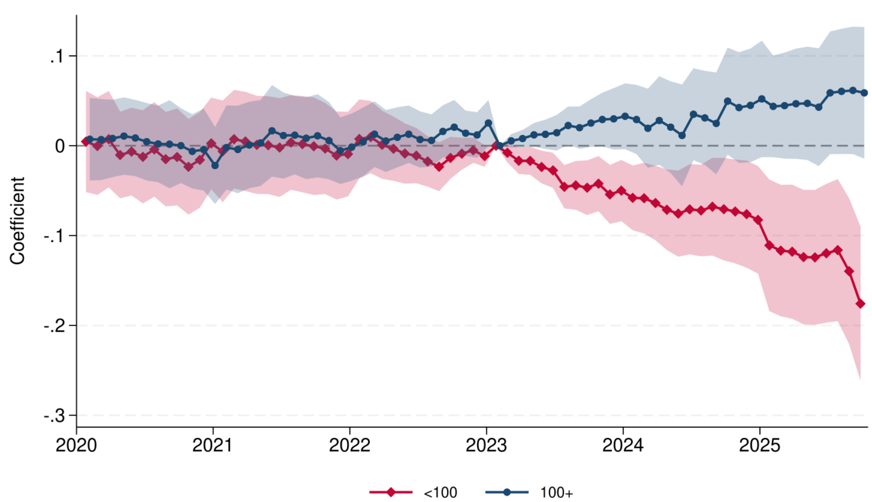

Appendix Figure A6: Employment by firm size with firm trends fit on 2020–2021

Notes: Event-study coefficients from equation (2) for the baseline log-FTE-employment outcome, estimated separately for 2023 adopters with fewer than 100 and with 100 or more FTEs at the pre-period base, with the firm-specific linear trend and calendar-month effects estimated on 2020–2021 instead of 2021–2022, so that calendar 2022 is out of sample. Reference period 2023m1; three-digit-industry-by-month and size fixed effects; survey-weighted; standard errors clustered by firm; 95% confidence bands. Specification otherwise as in Figure 9. Refitting the trends on 2020–2021, the 2022 coefficients of both size classes stay flat, and after 2023 the small-adopter decline re-emerges—reaching roughly -18 log points by late 2025—while large adopters remain flat to slightly positive.

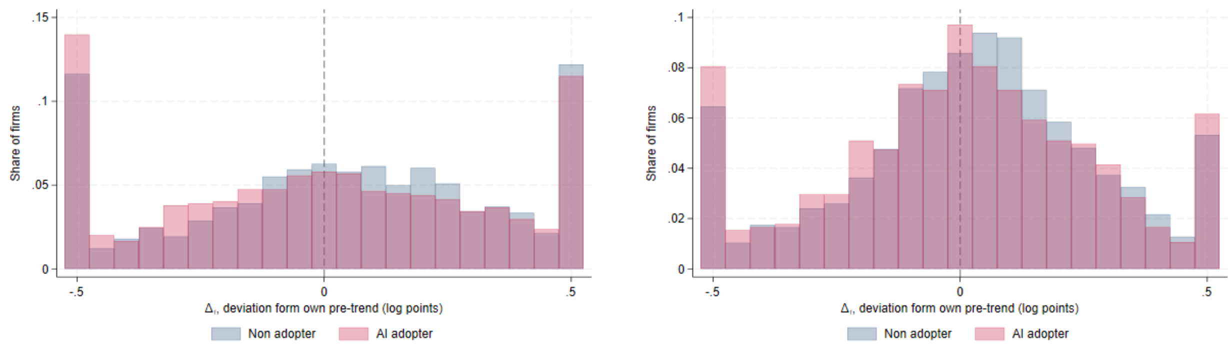

(a) Late window

(b) Full window

Appendix Figure A7: Distribution of the firm-level decline, adopters versus non-adopters

Notes: Histograms of Δ_f for adopters and non-adopters, residualized on industry-by-month and size, over the late (a) and full (b) post windows. The distributions are top-coded at ± 0.5, and between 8 and 14% of firms fall in the two end bins, so the extreme tails are not readable here; the quantile–quantile plots of Figures 10 and A8 are the cleaner view of the tails.

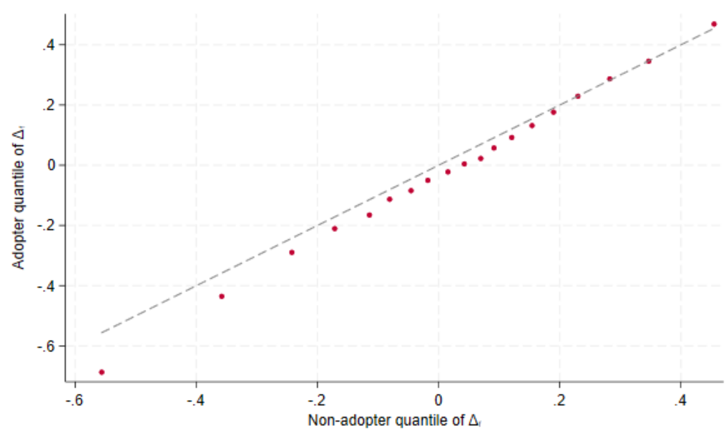

Appendix Figure A8: Quantile–quantile plot of the firm-level decline, full post window

Notes: Full-window (2023m2–2025m9) counterpart of Figure 10: adopter quantiles of Δ_f against non-adopter quantiles, with the 45-degree line; residualized on industry-by-month and size.

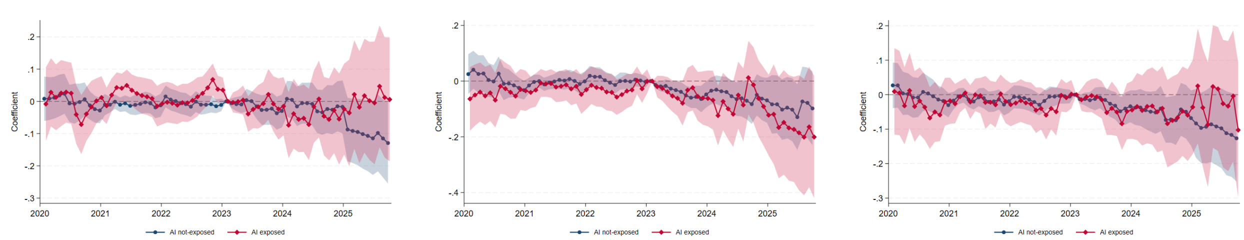

(a) Use

(b) Augmentation

(c) Automation

Appendix Figure A9: Employment decline by occupational AI exposure, Anthropic Economic Index

Notes: The task-based indices in section 4.4 rate occupations by how much of their work AI can *in principle* do. We contrast them with the usage-based measures of the Anthropic Economic Index (Handa et al., 2025), which instead rate occupations by how their tasks are *actually* handled in Claude conversations—by overall use, and by whether that use augments or automates the task. Event-study coefficients from equation (2), splitting occupations at the top quartile of the use, augmentation, and automation shares of the Anthropic Economic Index (Handa et al., 2025) (AI-exposed vs. not exposed). Specification, weighting, and inference as in Figure 7.

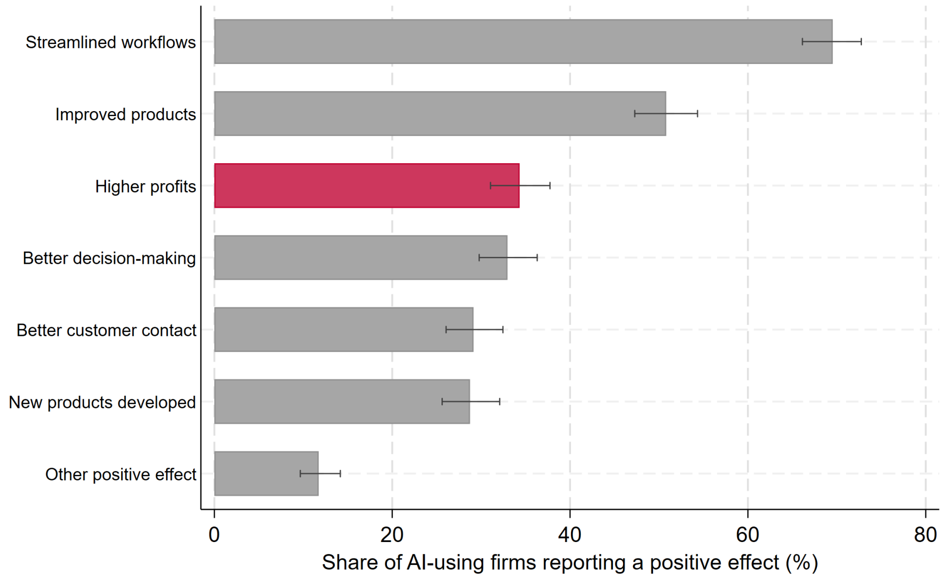

Appendix Figure A10: Self-reported benefits of AI usage beyond profitability

Notes: Shares of AI-using firms reporting each benefit from the survey's benefit battery, complementing the profitability assessments in Figure 3a.

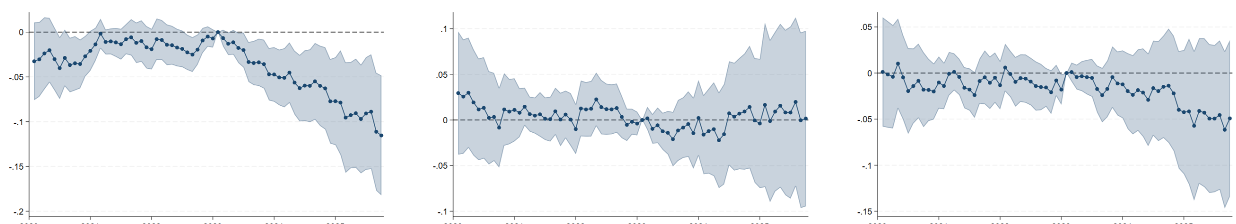

(a) Text and code (b) Image, video, sound (c) Automation

Appendix Figure A11: Employment around 2023 adoption, by type of AI use

Notes: Event-study coefficients from equation (2), re-estimated with first-time adoption of a specific AI-use type as the treatment: text and code generation (a), image, video, and sound generation (b), and automation (c). The three types are not mutually exclusive—a firm may adopt more than one—so these are separate contrasts, not a decomposition. Survey-weighted; specification and inference as in Figure 7.

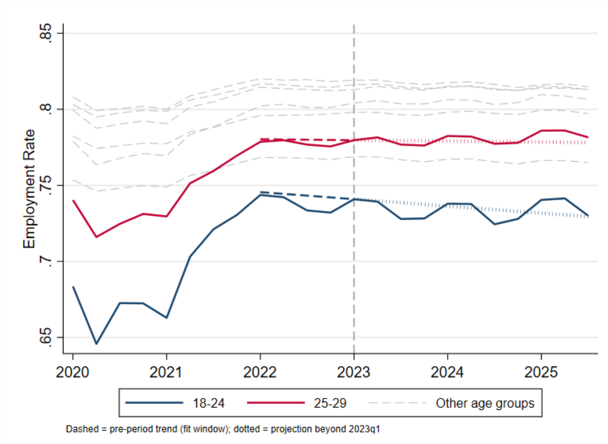

(a) Employment rate, by age (2023Q1= 100)

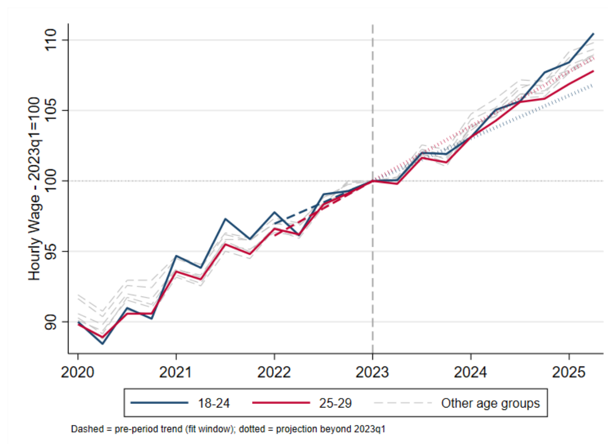

(b) Hourly wage, by age (2023m1= 100)

Appendix Figure A12: Aggregate young-worker outcomes

Notes: Full-population young-worker employment rate (a) and hourly wage (b) from the BFL register, by age group, indexed to 2023 with pre-period trend projections. There is no clear divergence in young-worker outcomes after 2023. This null is not in tension with the firm-level decline and is not evidence that it was offset: an effect of the size we estimate—under a percent of employment, concentrated in firms holding a tenth of it—is well inside the noise of an aggregate series.

B Additional tables

Table 5: Employment by firm size: triple-difference estimates

	(1) ln FTE employment (detrended)	(2) ln FTE employment (detrended)
Adopter	0.0371**	0.0364**
	(0.0179)	(0.0185)
Adopter × Post	-0.0510**	-0.0520**
	(0.0238)	(0.0246)
Adopter × Large	-0.0082	0.0017
	(0.0296)	(0.0325)
Adopter × Post × Large	0.0438	0.0643*
	(0.0350)	(0.0350)
Observations	131668	129751
Weights	Survey	Survey
Post window	2023+	2023+
Specification	Detrended	Ind×m×size FE

Notes: Triple-difference estimates on detrended log FTE employment. *Adopter* is the first-time-2023-adopter indicator; *Post* indicates months from January 2023; *Large* indicates 100 or more FTEs at the pre-period base. Column (1) uses the baseline detrended specification; column (2) absorbs industry-by-month-by-size fixed effects. The Adopter×Post coefficient is the average post-adoption effect at small adopters; the Adopter×Post×Large coefficient is the small-versus-large difference. Survey-weighted; standard errors clustered by firm in parentheses.

Table 6: Distribution of the firm-level decline: quantile gaps, adopters minus non-adopters (late window)

	(1)
Non-adopters, p10	-0.565***
	(0.036)
p25	-0.166***
	(0.016)
p50	0.063***
	(0.010)
p75	0.272***
	(0.013)
p90	0.502***
	(0.019)
Adopters, p10	-0.661***
	(0.074)
p25	-0.267***
	(0.020)
p50	-0.008
	(0.018)
p75	0.253***
	(0.033)
p90	0.505***
	(0.035)
Gap, p10	-0.096
	(0.082)
p25	-0.101***
	(0.025)
p50	-0.072***
	(0.021)
p75	-0.019
	(0.035)
p90	0.003
	(0.040)
Gap in means (the headline)	-0.072**
	(0.031)
Fan test: gap(p10) − gap(p90)	-0.099
	(0.087)
Firms	2968

Notes: Full version of the compact main-text table 2. Δ_f is a firm's average post-window deviation of detrended log FTE employment from its own 2021–2022 pre-trend, residualized on industry-by-month and size before collapsing, exactly as in the event study of Figure 7. The table reports quantiles of Δ_f for non-adopters and adopters and their difference (*Gap*); *Gap in means* is the headline average effect and *Fan test* is the p10-minus-p90 gap contrast. The late post window is 2024m10–2025m9. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends, which are estimated once outside the bootstrap. Full-window and unresidualized variants are reported in table 7.

Table 7: Distribution of the firm-level decline: quantile gaps, full post window

	(1)
Non-adopters, p10	-0.358***
	(0.022)
p25	-0.115***
	(0.010)
p50	0.042***
	(0.009)
p75	0.189***
	(0.009)
p90	0.346***
	(0.011)
Adopters, p10	-0.434***
	(0.043)
p25	-0.164***
	(0.020)
p50	0.006
	(0.012)
p75	0.177***
	(0.020)
p90	0.347***
	(0.021)
Gap, p10	-0.076
	(0.048)
p25	-0.049**
	(0.022)
p50	-0.036**
	(0.015)
p75	-0.012
	(0.022)
p90	0.001
	(0.024)
Gap in means (the headline)	-0.044**
	(0.021)
Fan test: gap(p10) − gap(p90)	-0.076
	(0.051)
Firms	2964

Notes: As in table 6 but over the full post window 2023m2–2025m9; residualized on industry-by-month and size. Unresidualized (raw) gaps are uniformly larger—roughly one-third of the raw gap is composition—and are not tabulated here. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends.

Table 8: Placebo: the firm-level decline over a held-out pre-period

	(1)
Non-adopters, p10	-0.323***
	(0.016)
p25	-0.148***
	(0.009)
p50	-0.010
	(0.006)
p75	0.137***
	(0.011)
p90	0.349***
	(0.017)
Adopters, p10	-0.384***
	(0.037)
p25	-0.147***
	(0.018)
p50	0.005
	(0.011)
p75	0.150***
	(0.016)
p90	0.320***
	(0.032)
Gap, p10	-0.061
	(0.041)
p25	0.000
	(0.020)
p50	0.015
	(0.013)
p75	0.013
	(0.019)
p90	-0.029
	(0.036)
Gap in means (the headline)	-0.004
	(0.019)
Fan test: gap(p10) − gap(p90)	-0.032
	(0.051)
Firms	3009

Notes: Placebo for the slope-noise artifact of table 6: the firm-specific pre-trends are re-estimated on an early window (calendar 2020), and Δ_f is then formed over a held-out pre-period (2021m7–2022m12)—before any adoption—using the identical residualization on industry-by-month and size, the same quantiles, and the same bootstrap. Because nothing happens in this window, every adopter-minus-non-adopter gap should be zero under the null of no artifact from estimation error in the firm slopes. It is: the 25th-percentile and median gaps are 0.000 (se 0.020) and +0.015 (se 0.013), the headline gap in means is −0.004 (se 0.019), and the fan test is −0.032 (se 0.051)—all null, in contrast to the significant negative gaps of table 6. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends.

Table 9: The firm-level decline by firm-size class (late window)

	(1) Small	(2) Large
Gap, p10	-0.145	0.027
	(0.097)	(0.076)
p25	-0.100***	-0.020
	(0.035)	(0.036)
p50	-0.069**	0.007
	(0.032)	(0.022)
p75	0.014	-0.004
	(0.049)	(0.026)
p90	0.049	0.006
	(0.061)	(0.039)
Gap in means	-0.060	0.016
	(0.040)	(0.033)
Fan test	-0.194*	0.021
	(0.108)	(0.081)
Mean base FTE	27.0	247.6
Adopters	353	309
Firms	1353	865

Notes: Quantile gaps of Δ_f (adopters minus non-adopters) estimated separately within the small (<100 FTE) and large ($\geq$100 FTE) size classes, split at the pre-period base. Within each class the residualization is class-specific—industry-by-month effects are absorbed as class$\times$industry$\times$month—so size is partialled out within class by construction; late post window (2024m10–2025m9). The small class carries the entire result—a fanning left shift with a 25th-percentile gap of -0.100 (se 0.035) and a median of -0.069 (se 0.032)—while the large class is a precise zero at every quantile (p25 -0.020, se 0.036; p50 $+0.007$, se 0.022; mean $+0.016$, se 0.033), matching the large-adopter contrast $\widehat{\tau}_{AL} = +0.016$ of section 5.2. The bottom panel reports mean base-period FTEs and adopter counts by class. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends.

Table 10: The firm-level decline by base-period young-worker share (late window)

	(1) Young q1	(2) q2	(3) q3	(4) q4
Gap, p10	0.062	0.003	-0.092	-0.230
	(0.122)	(0.221)	(0.118)	(0.221)
p25	-0.051	-0.076	-0.113	-0.248**
	(0.040)	(0.072)	(0.081)	(0.109)
p50	-0.077**	-0.058	-0.104**	-0.090
	(0.038)	(0.044)	(0.042)	(0.063)
p75	-0.001	0.011	-0.124**	0.009
	(0.059)	(0.064)	(0.049)	(0.060)
p90	-0.027	0.005	-0.082	0.016
	(0.065)	(0.144)	(0.068)	(0.092)
Gap in means	-0.061	-0.021	-0.099**	-0.123*
	(0.067)	(0.060)	(0.049)	(0.068)
Fan test	0.089	-0.002	-0.011	-0.246
	(0.133)	(0.254)	(0.130)	(0.230)
Difference in spread	-0.089	0.002	0.011	0.246
	(0.133)	(0.254)	(0.130)	(0.230)
Gap in means, bottom 20% trimmed	-0.056	-0.008	-0.106**	-0.094*
	(0.046)	(0.049)	(0.045)	(0.057)
Mean young share (pp)	3.8	12.9	24.0	51.5
Firms	740	739	739	739

Notes: Adopters split into quartiles of their base-period young-worker employment share; each column re-estimates the quantile gaps of table 6 against the common non-adopter distribution over the late post window (2024m10–2025m9), residualized on industry-by-month and size. Cells hold roughly 740 firms each and are correspondingly thin, so the gradient is suggestive. Bootstrap standard errors in parentheses, clustered at the firm and conditional on the first-stage firm pre-trends.

C Which worker margin drives the decline: a joint heterogeneity exercise

The main text shows that AI-adopting firms reduce employment (relative to pre-trends) especially among workers that are educated, young, and employed in AI-exposed occupations, albeit showing one margin at a time. We now decompose this exposure further: since the margins are correlated, we ask which (if any) of the three factors is the primary determinant of the relative employment fall between adopters and non-adopters. To answer this question, we reframe our exercise in a manner that allows us to race the three margins in a regression format.

To this end, we construct eight firm-level cells by collapsing worker-months at surveyed firms across the combinations of young, tertiary educated, and working in an AI-exposed occupation.[25] In line with the main exercise, we focus on log employment counts in each cell, deseasonalized and detrended within firm-cell on the years 2021 and 2022.

We then study the September 2025 employment of each firm cell relative to its pre-trend: each coefficient is a log deviation from the cell's own extrapolated pre-trend for adopters relative to non-adopters. The specification absorbs firm-by-cell, cell-by-industry-by-month, and size-by-month fixed effects, is survey weighted, clusters standard errors by firm, and uses 381,020 firm-cell-month observations.

Panel A of table 12 reports the results. For comparison, the first column reports the results without any further splits: the estimated pooled adopter effect at September 2025 is -0.09, consistent with the main difference-in-differences results in Figure 7.

Columns (2)–(4) then test whether the three different margins are significant in isolation, and finds AI exposure as the only statistically significant and meaningful determinant at -0.14. In that specification, the coefficient on adoption falls to -0.03: two thirds of the employment effect can be explained by occupational AI exposure. The joint specification (column (5)) confirms AI-exposure as the only statistically significant margin.

Panel B of table 12 extends the joint specification with further interactions. We find that in particular the interaction of AI-exposed and tertiary education is significant—in this specification, the effects from AI-exposure themselves are largely absorbed into the doubly-interacted term (columns 3 and 4). In the fully saturated specification (column 5), only the triple interaction is significant.

Why are age and education insignificant by themselves while they are significant in the main figure (Figure 12)? This is because we now split the data into 8 cells, and estimate pre-trends for each firm-cell combination. Because of this, and the wide confidence bands in the coefficients on young and education, we surmise that we cannot reject that these margins are still significant.

This issue is particularly salient since the finer granularity might lead to many zeros which our log-based outcome does not deal well with. We therefore now turn to an alternative specification whose outcome is each cell's employment as a share of firm base employment—defined at zero for

[25] We define each worker's characteristics according to their base month, January 2023, and hold them fixed thereafter.

cells that empty out, and additive across the eight cells so that they sum to the firm total.

Table 14 decomposes the firm-level decline this way: the exposed-cell contribution is the only component significant at the one-percent level, and the decline concentrates there rather than in the age or education splits. Because the three margins are correlated and none of the pairwise contrasts is sharp, we read this as locating where the decline falls across firms, not as isolating a single mechanism. The implied exposed decline is larger here than the log specification's marginal effect, which is itself the point: cells that empty out entirely are invisible to $\ln E$, and their weight grows with the horizon.

How much of this concentration is recomposition within firms, rather than exposed-heavy firms shrinking more as a whole? Table 11 absorbs firm-by-month fixed effects, leaving only within-firm composition tilts. The within-firm exposed tilts are not distinguishable from zero, and are small against the exposed contribution that the share specification measures without that absorption. At this sample size the decomposition therefore cannot say how the exposure concentration divides between recomposition inside firms and differential shrinkage across them.

The one within-firm margin that approaches significance is generational: the young share of the workforce tilts down, at the ten-percent level. This is the composition signature of the hiring freeze of Section 4.2—adopters stop hiring, and the workers who would have arrived are young—and it is where the within-firm evidence is strongest.

A direct test of the across-firm channel is likewise null: interacting the firm-level adopter effect with the firm's base-period exposed employment share leaves it flat, and the firm's young-worker share dominates exposure when both are entered.

The age margin is not rejected so much as unresolved: the young marginal effect is imprecise, with an interval that contains zero but also effects as large as the age gradient of Section 4.4. A single post month and cells covering a fifth of base employment cannot adjudicate it.

Table 13 reports the cell-implied effects from the saturated model, which locate the decline in the tertiary exposed cells and deepest among young tertiary exposed workers. At this horizon the saturated interactions are marginally significant—weak evidence that the concentration is a genuine interaction rather than the sum of additive margins—but the additive joint test does not reject, so we read the pattern as suggestive rather than settled.

Aggregating the saturated cell effects by adopter base composition recovers a decline in line with the pooled estimate, but it is not comparable to the firm-level employment fall: cells that empty out drop from the log outcome, so the extensive margin is invisible. The share specification of Table 14, which retains those cells, remains our magnitude benchmark.

Table 11: Within-firm versus between-firm components of the exposure and age margins

	AI-exposed (pp of workforce)	Young (pp)
Between-firm run, level	-8.040	-5.990
Between-firm run, centred	-1.780	0.270
Within-firm run	-1.327	-1.590
	(0.862)	(0.905)
Exposed − non-exposed	-2.653	
	(1.724)	

Notes: Units are percentage points of the firm's base workforce (share coefficients ×100). *Between-firm run, level* is the group contribution from the share specification of table 14 at September 2025; *centred* subtracts the firm's average cell change, the part a within-firm design could in principle recover. *Within-firm run* re-estimates the additive specification absorbing firm-by-month fixed effects, so only within-firm composition tilts remain; the final row is the within-firm exposed-minus-non-exposed contrast. Survey (VGT) weights—the only weighting for which the within–between comparison is valid; standard errors clustered by firm in parentheses. Do-file `stata_codes/4l_jointregression_share_additive.do`. Reading the two runs against each other requires care. If every adopter shrank at a uniform rate across all eight cells, each cell's contribution in the between-firm run would be exactly its base employment share times the firm total, and the within-firm run would return zero by construction. The exposed cell's contribution exceeds that uniform benchmark, so something beyond uniform shrinkage is present, and only two things can produce it: exposed cells falling faster than the firm average, which the within-firm run measures, or firm-level shrinkage covarying with exposure composition across adopters, which it does not. At these standard errors the table attributes the excess to neither, and it should not be read as evidence for either. Two further cautions. The eight cells are not independent—exposed cells are disproportionately young and tertiary educated—so the exposed-cell contribution absorbs whatever the age and education margins are doing and is not by itself evidence of an exposure mechanism. And exposure enters here as cell membership in the top Felten quartile with occupation frozen at the start of the worker's firm spell, which is a different object from a firm's average exposure measured relative to its industry; the two need not rank adopters the same way.

Table 12: Racing the three margins: one at a time, jointly, and interacted

Panel A: one split at a time and jointly

	(1) Pooled	(2) Age	(3) Education	(4) AI Exposure	(5) Joint
Adopter x Post	-0.0938**	-0.1000**	-0.0935*	-0.0380	-0.0391
	(0.0422)	(0.0482)	(0.0561)	(0.0533)	(0.0677)
Adopter x Post x Young		0.0199			-0.0095
		(0.0853)			(0.0840)
Adopter x Post x Tertiary			-0.0007		0.0115
			(0.0752)		(0.0748)
Adopter x Post x AI-exposed				-0.1413*	-0.1441**
				(0.0732)	(0.0717)
p-value: all three splits = 0					0.256
Post month	2025m9	2025m9	2025m9	2025m9	2025m9
Observations	381020	381020	381020	381020	381020
Firm x cell FE	Yes	Yes	Yes	Yes	Yes
Cell x Industry x month FE	Yes	Yes	Yes	Yes	Yes
Size x month FE	Yes	Yes	Yes	Yes	Yes

Panel B: pairwise and saturated interactions

	(1) Age x Edu	(2) Age x Exp	(3) Edu x Exp	(4) All pairs	(5) Saturated
Adopter x Post	-0.0170	-0.0569	-0.0993	-0.0966	-0.0599
	(0.0698)	(0.0683)	(0.0747)	(0.0767)	(0.0782)
Adopter x Post x Young	-0.0707	0.0305	0.0079	0.0016	-0.0853
	(0.1131)	(0.1007)	(0.0847)	(0.1228)	(0.1358)
Adopter x Post x Tertiary	-0.0308	0.0173	0.1414	0.1061	0.0237
	(0.0819)	(0.0749)	(0.0936)	(0.0969)	(0.1002)
Adopter x Post x AI-exposed	-0.1501**	-0.1110	0.0093	0.0276	-0.0545
	(0.0720)	(0.0753)	(0.0979)	(0.0989)	(0.1077)
Adopter x Post x Young x Tertiary	0.1381			0.1035	0.3266
	(0.1704)			(0.1685)	(0.2162)
Adopter x Post x Young x AI-exposed		-0.1304		-0.1334	0.2553
		(0.1714)		(0.1719)	(0.2425)
Adopter x Post x Tertiary x AI-exposed			-0.3224**	-0.2991**	-0.1203
			(0.1425)	(0.1405)	(0.1654)
Adopter x Post x Young x Tertiary x AI-exposed					-0.7281**
					(0.3454)
p-value: interactions = 0				0.126	0.045
Post month	2025m9	2025m9	2025m9	2025m9	2025m9
Observations	381020	381020	381020	381020	381020
Firm x cell FE	Yes	Yes	Yes	Yes	Yes
Cell x Industry x month FE	Yes	Yes	Yes	Yes	Yes
Size x month FE	Yes	Yes	Yes	Yes	Yes

Notes: Worker-months at surveyed firms are collapsed into eight firm-level cells, the combinations of young (aged under 30 at the base month, held fixed thereafter), tertiary educated, and employed in an AI-exposed occupation (top quartile of the Felten, Raj, and Seamans (2023) index, occupation frozen at the start of the worker's firm spell). The outcome is log cell employment, deseasonalized and detrended within firm-cell on 2021m1–2022m12; the estimation sample is that pre-window plus September 2025 alone, with 2023m2–2025m8 dropped, so each coefficient is a log deviation from the cell's own extrapolated pre-trend for adopters relative to non-adopters. The specification absorbs firm-by-cell, cell-by-industry-by-month, and size-by-month fixed effects, is survey weighted, clusters standard errors by firm, and uses 381,020 firm-cell-month observations. *Panel A* enters the three splits one at a time in columns (1)–(4) and all three additively in column (5), so each triple-interaction coefficient is that margin's contribution holding the other two fixed; its reported p-value is for the joint test that all three split coefficients are zero. *Panel B* extends

Table 13: Cell-implied effects and contributions to the adopter aggregate

	Implied Effect $\tilde{\beta}_c$	Base emp. share ω_c	Contribution $\omega_c\beta_c$	Share of Total (%)
Old, non tertiary, not exposed	-0.0599	0.407	-0.0244	24.7
	(0.0782)		(0.0318)	
Old, non tertiary, exposed	-0.1144	0.108	-0.0123	12.5
	(0.0860)		(0.0093)	
Old, tertiary, not exposed	-0.0362	0.124	-0.0045	4.6
	(0.0748)		(0.0093)	
Old, tertiary, exposed	-0.2110**	0.153	-0.0323	32.7
	(0.0986)		(0.0151)	
Young, non tertiary, not exposed	-0.1452	0.121	-0.0176	17.8
	(0.1180)		(0.0143)	
Young, non tertiary, exposed	0.0556	0.018	0.0010	-1.0
	(0.1806)		(0.0032)	
Young, tertiary, not exposed	0.2051	0.034	0.0069	-7.0
	(0.1451)		(0.0049)	
Young, tertiary, exposed	-0.4425***	0.035	-0.0156	15.8
	(0.1604)		(0.0056)	
Adopter aggregated	-0.0988**	1.000	-0.0988	100.0
	(0.0469)		(0.0469)	

Notes: Effects implied for each of the eight firm-level cells by the saturated specification in column (5) of Panel B of table 12, whose interactions are jointly insignificant. ω_c is the cell's share of adopter base employment and $\omega_c\beta_c$ its contribution to the composition-weighted adopter aggregate in the final row. Cells, outcome, sample, fixed effects, weighting, and inference as in table 12. Delta-method standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Group contributions to the firm-level decline, share specification

	Contribution	Std. err.	p-value
Firm Total	-0.1252	(0.0589)	0.034
By age:			
Young (age < 30)	-0.0599	(0.0295)	0.043
Older	-0.0653	(0.0438)	0.136
By education:			
Tertiary educated	-0.0480	(0.0294)	0.103
Non-tertiary educated	-0.0773	(0.0464)	0.096
By exposure:			
AI-exposed occupation	-0.0804	(0.0259)	0.002
Non AI-exposed occupation	-0.0448	(0.0487)	0.357
Memo: base share young:		0.217	
Memo: base share tertiary:		0.351	
Memo: base share AI-exposed:		0.319	

Notes: Share-based analogue of the horse-race exercise. The outcome is each cell's FTE employment as a share of firm base employment, deseasonalized and detrended within firm-cell on the pre-window as in table 12; unlike the log outcome, the share is defined (at zero) for cell-months that empty out, so such cells remain in the sample, and contributions are additive—the two groups within each split sum to the firm total in the first row. Survey weighted; winsorized; evaluated at September 2025; coefficients are in raw share units (multiply by 100 for percentage points of base employment); standard errors clustered by firm. The memo rows give adopter base employment shares. Do-file `stata_codes/4l_jointregression_share_noFMfe.do`.

D Worker reallocation: occupational mobility

We have previously shown that adopting firms change their employment composition towards less AI-exposed occupations. Is this purely coming from differences in hiring, or also from within-firm reallocation of workers across tasks?

To answer this question, we turn to the occupation codes attached to each salary payment. These codes are maintained by the firms' internal staff and are updated at infrequent intervals. As a consequence, within-firm occupational mobility is a rare event subject to bunching, and we measure it by estimate for each firm the share of high-tenured workers that has a changed occupation code relative to an anchor date. We aggregate occupations to four digits, further aggregating 6 codes that significantly differ between the two occupational coding regimes ISCO08 and ISCO88.[26]

We then estimate the difference in occupational mobility between adopting and non-adopting firms in our standard regression framework (equation (2)). Figure D1, panel (a), plots the relative occupational mobility of workers in adopting firms relative to 2022. We chose this reference point since we noticed that occupational mobility already increases ahead of 2023.

An increase in occupational mobility ahead of AI adoption has two primary explanations. For one, AI-adopting firms could be forward looking and shift workers across tasks ahead of the actual reduced hiring. Alternatively, this could be driven by inherent differences between adopting and non-adopting firms.

Panel (b) repeats the same alternative for 2021; we do not find the same differential drift in occupational mobility for that placebo date. Panels (c) and (d) decompose occupational mobility into that towards more and less AI-exposed occupations (according to Felten, Raj, and Seamans, 2023). We find increased occupational mobility towards less AI-exposed occupations among AI-adopting firms, and no increased occupational mobility towards more AI-exposed occupations.

We read these patterns—the null placebo, and the drift towards less-exposed occupations—as suggesting that the increased occupational mobility is not incidental, even though it already occurs in 2022: firms reorganize their workforce ahead of deploying the technology.[27]

Finally, we find no significant differences in excess occupational mobility by size, as documented in Online Appendix table 15. This suggests that all adopting firms restructure employment across occupations in anticipation of adoption, regardless of size.

[26]Formally, ISCO08 was introduced in 2010. We notice that a significant share of workers change occupations in January 2025 across codes that differ within the ISCO08 regime, but have identical labels across ISCO08 and ISCO88, the previous coding regime. Our hypothesis is that some firms had not updated their occupational coding to the new regime; we aggregate these codes to prevent spurious occupational mobility reflecting such coding updates.

[27]The design cannot separate anticipation from selection. Firms already reallocating occupations during 2022 are exactly the firms that report adopting in 2023, and the event study cannot tell an anticipatory reorganization apart from a pre-existing difference between eventual adopters and everyone else. We therefore read the run-up as one more descriptive fact about who adopts, consistent with the paper's descriptive, non-causal framing, and not as evidence that adoption caused the reallocation.

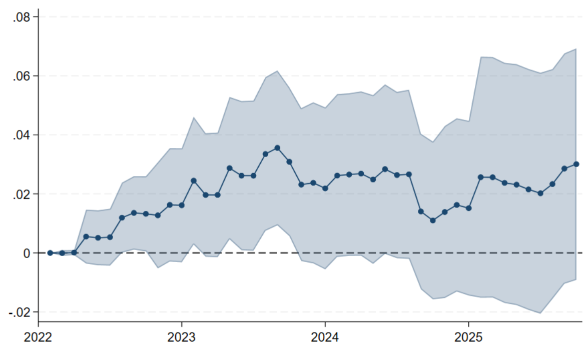

(a) Total switched share

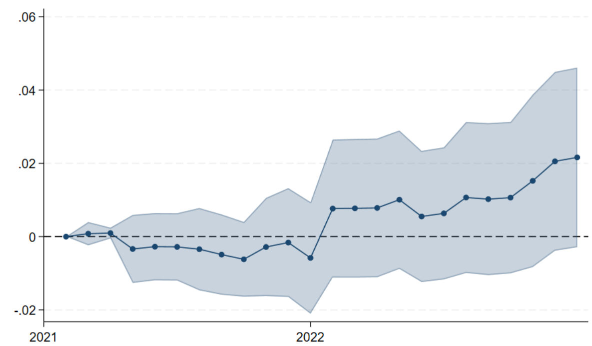

(b) Placebo: January-2021 anchor

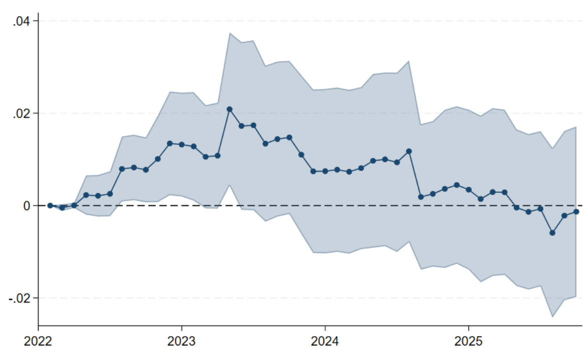

(c) Toward less-exposed occupations

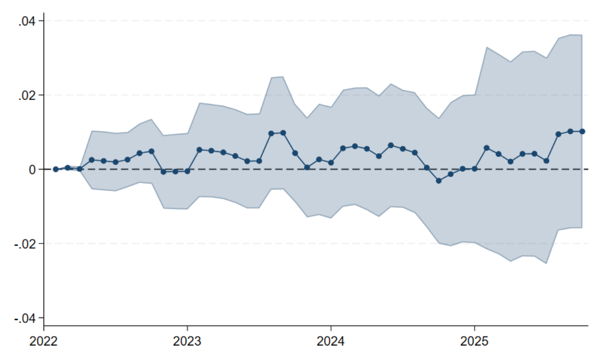

(d) Toward more-exposed occupations

Appendix Figure D1: Occupational mobility of incumbents around AI adoption (event studies)

Notes: The switched share is the share of surviving roster workers—incumbents with 24 months or more of tenure at the anchor date—recorded under a three-digit DISCO occupation code different from their frozen anchor-date code, among those still at the firm. Panels (a), (c), and (d) anchor the roster in January 2022 and are estimated over 2022–2025: panel (a) is the overall switched share, and panels (c) and (d) split post-adoption moves by direction, toward three-digit occupations with lower (panel (c)) or higher (panel (d)) average AI-exposure than the origin, using Felten, Raj, and Seamans (2023). Panel (b) is a placebo that anchors the roster in January 2021 and is estimated over 2021–2022. The specification, weighting (survey weights), and firm-clustered inference are otherwise as in Figure 7, with 95% confidence bands.

Table 15: Occupational switching by firm size: triple-interaction estimates

	(1) shBD_sw	(2) shBD_sw	(3) shBD_sw	(4) shBD_swlo	(5) shBD_swlo	(6) shBD_swlo	(7) shBD_swhi	(8) shBD_swhi	(9) shBD_swhi
Adopter	0.0085*	0.0072*	0.0058	0.0059**	0.0038**	0.0039*	0.0007	0.0010	0.0001
	(0.0046)	(0.0038)	(0.0038)	(0.0027)	(0.0019)	(0.0021)	(0.0028)	(0.0022)	(0.0022)
Adopter × Post	0.0145	-0.0046	0.0045	0.0017	-0.0064	-0.0038	0.0014	-0.0000	0.0030
	(0.0121)	(0.0112)	(0.0097)	(0.0069)	(0.0048)	(0.0049)	(0.0070)	(0.0069)	(0.0059)
Adopter × Large	-0.0014	0.0024	0.0019	-0.0046	-0.0017	-0.0023	0.0040	0.0052	0.0051
	(0.0066)	(0.0051)	(0.0055)	(0.0035)	(0.0024)	(0.0030)	(0.0041)	(0.0034)	(0.0033)
Adopter × Post × Large	-0.0029	-0.0096	0.0017	0.0030	0.0032	0.0054	0.0042	-0.0044	0.0031
	(0.0168)	(0.0169)	(0.0156)	(0.0080)	(0.0069)	(0.0073)	(0.0092)	(0.0104)	(0.0083)
Observations	118159	123602	123602	118159	123602	123602	118159	123602	123602
Weights	Survey	FTE	Balanced	Survey	FTE	Balanced	Survey	FTE	Balanced
Post window	2023+	2023+	2023+	2023+	2023+	2023+	2023+	2023+	2023+

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Triple-difference estimates on the January-2022-anchored switched share. Columns (1)–(3) use the overall switched share, columns (4)–(6) the share switching toward less-exposed three-digit occupations, and columns (7)–(9) the share switching toward more-exposed occupations, each estimated under survey, FTE, and covariate-balancing weights. *Adopter* is the level gap over the January-2022 reference window; *Post* indicates months from January 2023; *Large* indicates 100 or more FTEs at the pre-period base. The Adopter×Post×Large coefficient tests whether any post-adoption switching excess differs by size class: it is insignificant in all nine columns and flips sign across weighting schemes, and Adopter×Post is itself null. The only stable coefficient is the adopter level gap over the January-2022 reference window—the anticipation run-up—which is positive and marginally significant for overall and less-exposed switching but does not grow after adoption. Exposure direction uses Felten, Raj, and Seamans (2023). Standard errors in parentheses; specification otherwise as in Figure D1.

E Wages

We ask whether adoption changed the price of labor at adopting firms, and study two outcomes. The first is the firm's average hourly wage—its total wage bill divided by total hours across all workers employed that month—deseasonalized and detrended with the firm-specific pre-period trend, exactly as for employment (panel (a)). The second is the entry wage of newly hired workers: the log hourly wage in the first month of each non-left-censored spell, with age-bin and education controls and three-digit occupation fixed effects (panel (b)). Both are estimated via (2) with base month 2023m1, absorbing industry-by-month and size-class fixed effects, weighted by survey weights and clustered by firm, as in Figure 7.

Both are a clean null. The firm average is flat around zero (in kroner per hour; the outcome in panel (a) is a level, not a log) across the whole 2020–2025 window, with no pre-trend and no break at January 2023 and wide bands throughout (panel (a)); entry wages are likewise flat, with a noisy, insignificant uptick in 2025 but nothing systematic before or after adoption (panel (b)). Neither null depends on weighting: both hold under employment and covariate-balancing weights. Adoption does not change the price of labor at adopting firms, neither for the existing workforce on average nor for new hires. In the appendix, we also show the wages of stayers only: adoption lowers the wage of stayers, which is driven by compositional changes. Once we control for those, the wage of stayers no longer differentially drifts at adopting firms than at non-adopting firms.[28]

Three explanations are consistent with the absence of a wage response at adopting firms. First, we are comparing firms within the same labor market—adopters and non-adopters in the same industry, with very similar composition of employees. In a neoclassical framework, differentially productive firms would still pay the same price for the same labor (and the extent to which the labor differs is minimal). Second, although the Danish labor market is flexible on the quantity margin, wages are less responsive to firm-level productivity: they are set through collective bargaining, with industry- and firm-level unions establishing wage floors and only a limited role for individual bargaining. Third, the wage effect of AI is theoretically ambiguous—exposed workers may earn more as AI raises their productivity, or less as greater substitutability weakens their bargaining position. Therefore, the absence of a wage response at adopters is not evidence of AI having no wage effect economy-wide.

[28] Figure E1 in the appendix decomposes it into stayer margins. A chained index of stayers' within-worker wage growth does drift below trend after 2023, but a version that freezes the contributing cohort at January 2023 is flat and null (panels (c) and (d) of Figure E1, respectively). The two differ only in whether the stayer pool is allowed to change, so the drift is purely compositional: it reflects the stayer pool tilting toward longer-tenured workers with flatter wage profiles—the retention pattern of section 4.2 and Figure 8—not a repricing of any given worker.

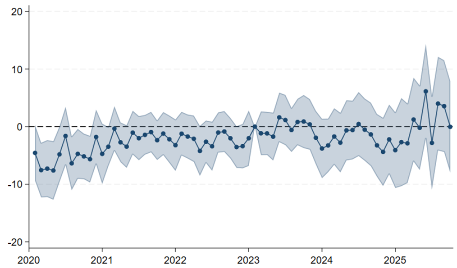

(a) Average hourly wage (all workers)

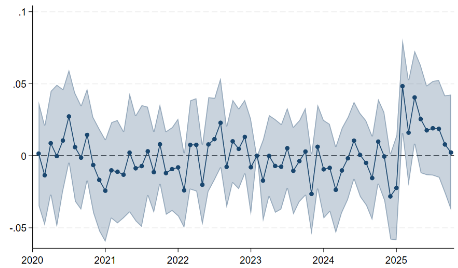

(b) New hires (entry wage)

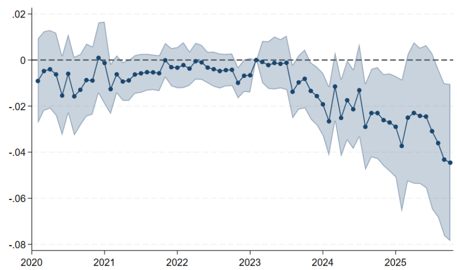

(c) Chained stayer index

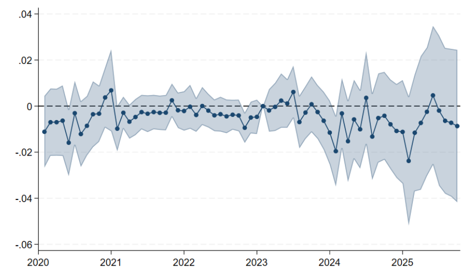

(d) Stayer cohort fixed at January 2023

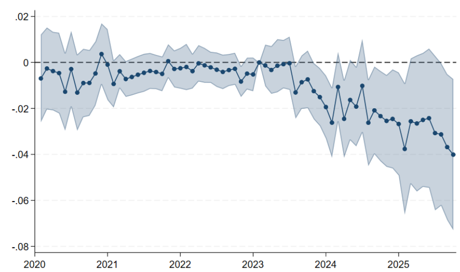

(e) Stayer wage index, firms with more than five stayers

Appendix Figure E1: Wages around 2023 adoption (event studies)

Notes: Event-study coefficients from equation (2) for wage margins. Panel (a): the firm's total wage bill divided by total hours—the average hourly wage across all workers employed at the firm—deseasonalized against firm-by-calendar-month means and residualized against a firm-specific linear trend fit on the pre-period, then run through the same event study as the other outcomes. Its vertical axis is in kroner per hour (a level), not log points; firm-months in the bottom percentile of the detrended hourly wage variable are trimmed. Panel (b): the log hourly wage of newly hired workers in the first month of a non-left-censored spell, additionally controlling for age bin and education group and absorbing occupation (3-digit) fixed effects. Panels (c)–(e) show the stayer wage margins underlying the flat firm average in panels (a) and (b). Panel (c): the chained stayer index, defined as the average one-month within-worker log hourly wage growth of workers employed at the same firm in consecutive months, averaged across stayers within each firm-month, cumulated over time into a firm-level index, rebased to zero at January 2023, and residualized against firm-specific calendar-month means and a firm-specific linear trend estimated on 2021–2022, exactly as the baseline employment outcome; the pre-period coefficients are therefore a standard pre-trend test, which passes. Panel (d) differs only in freezing the set of contributing workers at the January 2023 cohort rather than chaining it month to month, flat throughout. The chained decline is therefore compositional—the stayer pool tilting toward longer-tenured workers with flatter wage profiles—and once composition is held fixed there is nothing. Panel (e) restricts the sample to firms with more than five stayers and is otherwise identical to the chained stayer index in panel (c); the chained path is essentially unchanged. All panels use base month 2023m1, absorb industry-by-month and size-class fixed effects, are weighted by survey weights, and cluster standard errors by firm.

F Balance-sheet outcomes

We complement the monthly register outcomes with annual firm balance-sheet data, which add revenue, profits, value added, and investment. These accounts are reported on a fiscal-year basis and are available only annually and only through 2024, making them considerably coarser and noisier than our monthly register; because fiscal years do not align across firms, the annual figures are also not exactly comparable across firms.

F.1 The annual employment analogue

To gauge whether these coarser data can detect the effect the monthly baseline identifies, Figure F1 repeats the employment exercise on the balance-sheet employment measure. Panel (a) shows the aggregate; panel (b) the trend-residualized difference-in-differences with industry-by-year and size controls, the annual analogue of equation (2). For the trend-residualization, we now have to estimate the linear trend using only two data points for the years 2021 and 2022.

In the aggregate (panel a), adopters appear to have shrunk in employment relative to their trend—consistent with the monthly baseline (Figure 7). In the difference-in-differences specification (panel b), however, the confidence bands are quite wide—spanning a 6% decline to a 5% increase—and contain both zero and the two-year decline that the monthly baseline identifies (section 4.2); at this effect size, the annual data are simply uninformative.

We conclude that the annual balance-sheet data are too sparse to detect employment effects before 2024. Appendix Figure F2 repeats the exercise for the other balance-sheet outcomes—log revenue, log profit, value added per worker, and investment—and likewise finds no significant effects.

(a) Aggregate employment (indexed)

(b) Adopters vs. non-adopters

Appendix Figure F1: Employment around 2023 adoption, annual balance-sheet data

Notes: Balance-sheet (annual accounts) employment. Panel (a): aggregate employment of adopters and non-adopters, indexed to 2023. Panel (b): trend-residualized event study with industry-by-year and size fixed effects and a pre-trend fit on 2021–2022; the annual analogue of equation (2).

F.2 Other balance-sheet outcomes

Beyond employment, the remaining balance-sheet outcomes tell the same story. Appendix Figure F2 repeats the exercise for log revenue, log profit, value added per worker, and investment, and likewise finds no significant effects.

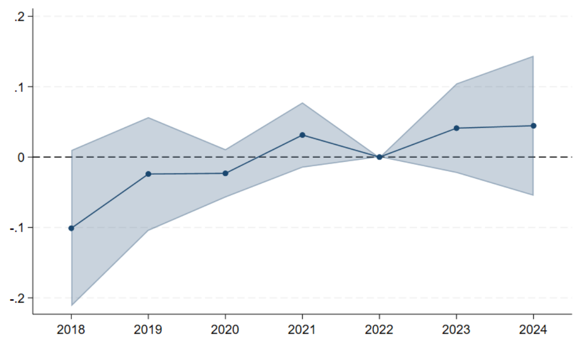

(a) Log revenue

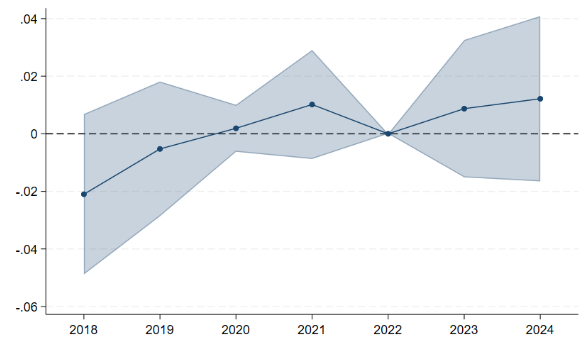

(b) Log profit

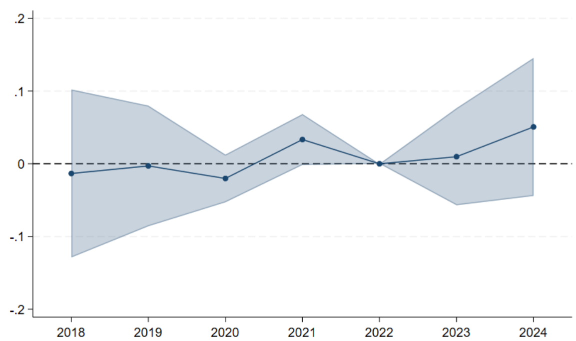

(c) Value added per worker

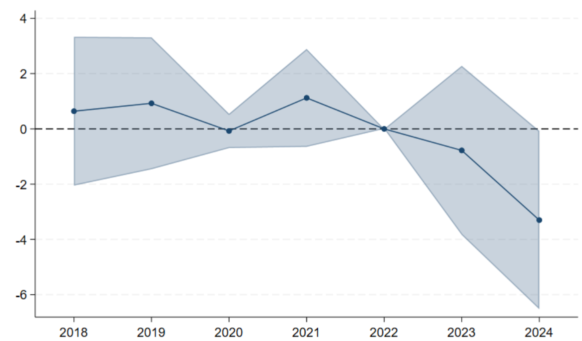

(d) Investment rate

Appendix Figure F2: Balance-sheet outcomes of 2023 adoption (event studies)

Notes: Trend-residualized annual event studies for balance-sheet outcomes, adopters vs. non-adopters, with industry-by-year and size fixed effects and a pre-trend fit on 2021–2022. Value added per worker is the firm's labor-productivity series; the fourth panel is the investment rate.

G Industry-level design: construction and decomposition

This appendix develops the construction behind the industry-level results of Section 5: the exposure measure s_i, the two industry event-study routes, the adoption-consistent component C^{AI}, and the exact contribution decomposition, together with its robustness variants.

G.1 The industry gradient

In order to regress industry-level employment changes on industry-level AI adoption, we need two ingredients. First, industry-level AI adoption. For each of our 196 three-digit industries, we compute its AI-usage intensity as the share of its 2023-employment that is among 2023-AI-adopting firms:

$$s_i = \left. \frac{\text{employment at 2023 adopters in industry } i}{\text{total industry } i \text{ employment}} \right|_{2023m1}. \tag{4}$$

Weighted by industry employment, s_i averages 0.19 across the 196 three-digit industries, with a median of 0.15 and a range from 0.04 at the 10th percentile to 0.33 at the 90th.

For industry-level employment, we will go two routes. In the first approach, we will compute the exact industry-level analogue of (2) by first residualizing each firm's employment against the pre-2023 trend and a firm-by-calendar-month fixed effect. We then aggregate these residuals up to each industry. In the second approach, we refrain from using firm-level information: we aggregate FTE employment to the industry level and then residualize industry-level pre-trends and industry-by-calendar-month fixed effects. This yields a detrended employment variable that a researcher without access to firm-level identifiers can measure.

We then estimate the industry-level analogue of our firm-level regression (equation (5)). Each industry is weighted by its 2023-employment, and standard errors are clustered by industry. Because the unit of observation is the industry-month, industry and month fixed effects enter separately, unlike the industry-by-month absorption of the firm design. We run this regression on two constructions of $g_{i,t}$ that differ in how much of the firm-level structure is removed before aggregation, shown as panels (a) and (b) of Figure 14.

$$g_{i,t} = \alpha_i + \mu_t + \sum_{\tau \neq 2023m1} \gamma_\tau \, s_i \cdot \mathbf{1}[t = \tau] + \varepsilon_{i,t}. \tag{5}$$

G.2 Decomposing the gradient

This wedge might stem from an overly crude back-of-the envelope calculation: after all, we highlighted that essentially all of the employment response is driven by small and medium-sized firms. This matters because industries differ starkly in their share of small firms: even though small-adopter shares are small on average, they account for 40% of the cross-industry variation. Employment-weighted across the 196 industries, the small-adopter share $s_{i,\text{small}}$ averages 0.09 (sd 0.08), ranging from zero at the 10th percentile through a median of 0.07 to 0.19 at the 90th. To improve the back-of-the-envelope calculation we repeat the industry-level exercise on an employment series that is built on the firm-level estimates from section 4.3. We ask: if the adoption effects estimated at the firm level were the only source of cross-industry differences in detrended employment growth—adopters deviating by $\hat{\tau}_k$, every other firm staying on trend—what exposure gradient would we measure?

To construct our firm-level consistent employment growth, we hold the base-month employment share of adopters in each size class k in their industry fixed and project it forward using the within-industry size-k adopter contrast $\widehat{\tau}_{Ak,t}$. For each industry, we then obtain the industry-level employment series by aggregating across size classes:

$$C_{it}^{AI} \;=\; \sum_{k\in\{\text{small},\,\text{large}\}} \frac{E_{Aki,0}}{E_{i,0}}\,\widehat{\tau}_{Ak,t}. \tag{6}$$

To be clear, C_{it}^{AI} is the deviation of detrended growth in industry i stemming purely from its baseline composition of firm sizes and adoption status, together with the trend-differential employment growth we estimated for the adoption-size group in the baseline exercise (equation (2), estimated separately by size class). We then regress this predicted component on the industry-level adopter share s_i:

$$C_{it}^{AI} \;=\; \alpha_i + \mu_t + \sum_{\tau\neq 2023m1} \gamma_\tau^{\text{pred}}\, s_i \cdot \mathbf{1}[t=\tau] + \varepsilon_{i,t} \tag{7}$$

and compare the predicted gradient $\gamma_\tau^{\text{pred}}$ with the observed gradient γ_τ estimated on actual industry employment. This comparison is not mechanical; nothing forces the two to agree—the only construction is that the predicted path is zero before 2023. There is no restriction between the magnitudes of γ_τ and $\gamma_\tau^{\text{pred}}$: factors unrelated to adoption can either raise or lower γ_τ relative to what adoption itself would imply.

We turn to growth rates $g_{i,t}$ in order to allow entry and exit and denote by $c_{ig,t}$ the contribution of group g to the employment growth of industry i:

$$c_{ig,t} = \frac{E_{ig,t} - E_{ig,\text{base}}}{E_{i,\text{base}}}. \tag{8}$$

By construction, industry-level employment growth is consistent with group-level contributions: $g_{i,t} \;=\; \sum_g c_{ig,t}$. We then estimate the total contribution of group g, γ_g^c, as the coefficient from regressing group g's January-2025 contribution c_{ig} (equation (8)) on s_i:

$$c_{ig} = \alpha_g + \gamma_g^c\, s_i + \epsilon_{ig}. \tag{9}$$

Because the contributions sum to industry growth and all five regressions share the same sample and weights, the γ_g^c sum exactly to γ, the aggregate cross-sectional gradient of -0.083. This gradient is estimated on the same industry-detrended series as panel (b) of Figure 14, but as a single January-2025 cross-section; it therefore differs slightly from the event-study coefficient at that date (-0.102), since the cross-section holds one observation per industry and cannot absorb

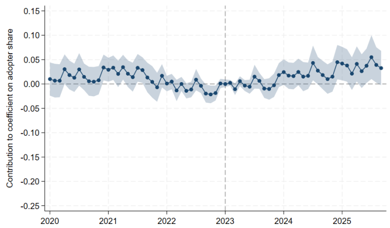

(a) Never-adopters

(b) 2023 adopters

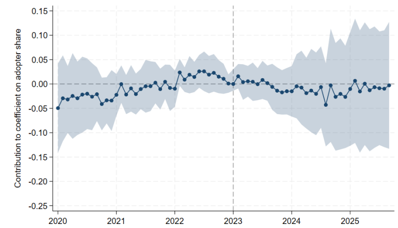

(c) 2022 users

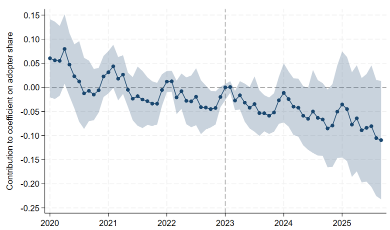

(d) Unclassified firms

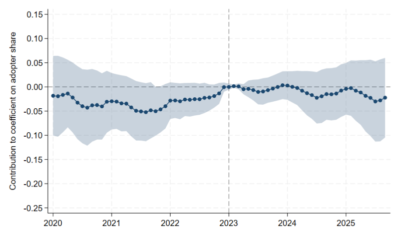

(e) Net entry

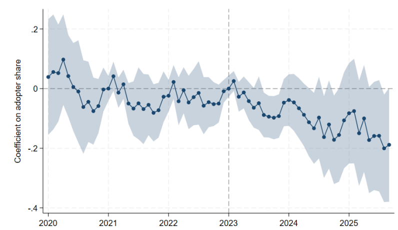

(f) Industry total

Appendix Figure G1: Contributions to industry employment growth by adoption status

Notes: Event-study coefficients on the industry adopter-exposure share s_i, estimated separately on each group's contribution $c_{ig,t}$ from equation (8). Panels (a)–(e) sum to panel (f) at every event time by linearity of least squares. Reference January 2023; industry and month fixed effects; weighted by base industry employment; standard errors clustered by industry; note the differing vertical scales. Shaded bands are 95% confidence intervals. The pre-2023 path in panel (e) reflects firms that exited before the base month rather than entrants.

industry fixed effects.

Per-group event studies and alternative specifications behind the decomposition in table 4: the month-by-month contributions (Figure G1), the post-2023 monthly panel (table 16), and the survey-respondent estimate of the differential/composition split (table 17).

Table 16: Contributions to industry employment growth (monthly panel, post-2023)

	(1) Industry total	(2) Never Adopter	(3) 2023 Adopter	(4) 2022 User	(5) Unclassified	(6) Net Entry
Adopter employment share X Post	-0.0775	0.0100	-0.0546***	-0.0019	-0.0519	0.0209
	(0.0731)	(0.0124)	(0.0138)	(0.0460)	(0.0408)	(0.0174)
Constant	0.0095	-0.0046***	0.0021*	-0.0035	-0.0205***	0.0360***
	(0.0066)	(0.0011)	(0.0012)	(0.0041)	(0.0037)	(0.0016)
Observations	13328	13328	13328	13328	13328	13328
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

Outcomes in col. 1 is detrendend industry employment growth from the full register; ...

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Adopter exposure share interacted with a post-January-2023 indicator. Outcomes are detrended full-register employment-growth contributions; column 1 is the industry total and columns 2–6 are the group contributions, which sum to column 1 by linearity of least squares. Industry trends are estimated in logs over January 2021–January 2023 and allocated using January-2023 base employment shares. Regressions use January-2023 industry-employment weights, industry and month fixed effects, and industry-clustered standard errors. Baseline adopter shares use survey weights; employment-weighted shares are a robustness check.

Table 17: Contributions to industry employment, survey-respondent panel

	Never Adopter	2023 Adopter	2022 User
Total contribution β_g	0.1846	-0.2183	-0.1079
	(0.0681)	(0.0514)	(0.0832)
Differential $\bar{\omega}_g \gamma_g$	0.0328	-0.0416	-0.0806
Composition $\Delta\bar{e}_g\, b_g^{\omega}$	0.0999	-0.1750	0.0655
Remainder	0.0519	-0.0017	-0.0929
Memo: base weight $\bar{\omega}_g$	0.3411	0.2004	0.4585
Memo: mean change $\Delta\bar{e}_g$	-0.1581	-0.1733	-0.1734
Memo: weight slope b_g^{ω}	-0.6318	1.0098	-0.3780
Aggregate coefficient		-0.1415	
		(0.1341)	

Notes: The decomposition of table 4 estimated on survey respondents only, where the adopter employment weight coincides with the exposure measure s_i and the composition term therefore has a known benchmark. Group totals sum to the aggregate coefficient in the final row. Weighted by base industry employment; robust standard errors in parentheses.

G.3 Decomposition details

This appendix gives the formal split behind the differential, composition, and remainder rows of table 4, and spells out the arithmetic by which the composition of firm groups across industries, combined with a below-trend shortfall shared by all groups, produces a negative exposure gradient even though no group grows differently in exposed than in unexposed industries. Each group's total contribution γ_g^c splits into three parts:

$$\gamma_g^c \;=\; \underbrace{\bar{\omega}_g\,\gamma_g^{\text{diff}}}_{\text{differential}} \;+\; \underbrace{\Delta\bar{e}_g\,\frac{\partial\omega_{ig}}{\partial s_i}}_{\text{composition}} + \text{cov}_g. \tag{10}$$

The *differential* term $\bar{\omega}_g\gamma_g^{\text{diff}}$ captures the part of the contribution that arises because the group grows differently in more-exposed industries; the *composition* term $\Delta\bar{e}_g\,\partial\omega_{ig}/\partial s_i$ the part that arises because more-exposed industries simply contain more of the group; and cov_g is the remainder, the interaction of the two deviations.[29] In terms of our question: the differential term is the "adopters in exposed industries shrink more" channel, and the composition term is the "there are more of them" channel.

To spell out what that unexplained component is: composition here means the composition of *firm groups across industries*. Industries with high s_i mechanically contain more adopter employment and less of the other groups—the memo rows of table 4 show that per unit of s_i, the adopter base-employment weight rises by 0.55 while the never-adopter and 2022-user weights fall by 0.21 and 0.54. The common shortfall is the fact that *every* group's average detrended growth is negative—between −0.11 and −0.21 across the four incumbent groups—because the whole economy decelerated relative to trends extrapolated from the 2021–2022 boom; this below-trend level is shared, not adoption-specific. Multiplying each group's mix shift by its below-trend average growth and summing produces a negative gradient on s_i even though no group grows differently in exposed than in unexposed industries (the differentials are near zero). The predicted component C^{AI} excludes this arithmetic by design: the firm contrasts $\widehat{\tau}_{Ak,t}$ are within-industry comparisons of adopters to non-adopters, which difference out the shared below-trend level, so C^{AI} carries only the part of the gradient that adoption effects can generate—and that part is about a quarter of what the exposure regression measures.

[29] ω_{ig} is group g's share of industry i's January-2023 base employment, and $\bar{\omega}_g$ is its employment-weighted mean across industries; γ_g^{diff} is the slope of group g's *own* growth rate—its employment change per unit of its own base employment—on s_i; $\Delta\bar{e}_g$ is group g's average detrended growth rate across industries; and $\partial\omega_{ig}/\partial s_i$ is the slope of the group's base weight on exposure.

Publications

ANALYSIS

Analyses focus on current issues of particular relevance to Danmarks Nationalbank's objectives. The analyses may also contain Danmarks Nationalbank's recommendations. They include our projections for the Danish economy and our assessment of financial stability. Analyses are targeted at people with a broad interest in economic and financial matters.

STAFF PAPER

Staff Papers provide insights into the analytical work being done by employees of Danmarks Nationalbank. Staff Papers include, for example, background analyses and method descriptions. They are primarily targeted at people who already have knowledge of economic and financial analyses. The conclusions expressed in Staff Papers are the authors' own and should not be interpreted as views held by Danmarks Nationalbank.

WORKING PAPER

Working Papers present research work by both Danmarks Nationalbank's employees and our partners. Working Papers are primarily targeted at professionals and people with an interest in central banking research as well as economics and finance in a broader sense. The conclusions expressed in Working Papers are the authors' own and should not be interpreted as views held by Danmarks Nationalbank.

NEWS

News may be a news article or an appetiser offering quick insight into one of Danmarks Nationalbank's more extensive publications. The series is targeted at people who need an easy overview and like a clear angle.

STATISTICAL NEWS

Statistical news focuses on the latest figures and trends in Danmarks Nationalbank's statistics. Statistical news is targeted at people who want quick insight into current financial data.

REPORT

Reports consist of recurring reports on Danmarks Nationalbank's areas of work and activities. Here you will find Danmarks Nationalbank's annual report, among other documents. Reports are targeted at people who need a status and update on the past period.

Danmarks Nationalbank
Langelinie Allé 47
2100 Copenhagen Ø
+45 3363 6363

DANMARKS NATIONALBANK