Merger Synergies, Task Exposure, and Worker Careers*

Malin Gardberg
IFN

Fredrik Heyman
IFN

Joacim Tåg
IFN and Hanken

June 2026

ABSTRACT

We study the worker-level incidence of technology-related merger synergies by following incumbent target workers around acquisitions, regardless of whether they remain with the firm. Average acquisition wage effects are close to zero, but acquisitions redistribute gains and losses across task groups if the acquirer has a technological advantage. Workers in substitutable occupations experience a relative wage decline. Workers in complementary occupations gain. Acquired firms also receive more foreign managers, and IT spending rises. Thus, technology-related merger synergies can appear neutral for workers on average while changing the relative value of tasks within firms and having long-run spillovers on workers' careers.

Keywords: careers, human capital, mergers and acquisitions, technology-related synergies, wages.
JEL Codes: G34, J30, J31, O39.

*E-mails: malin.gardberg@ifn.se, fredrik.heyman@ifn.se, joacim.tag@ifn.se. We thank Ashwini Agrawal, Alex Xi He, Ernst Maug, Andreas Moxnes, Pehr-Johan Norbäck, Oskar Nordström Skans, Martin Olsson, Anna Salomons, Per Strömberg, Joonas Tuhkuri, Ramona Westermann and seminar and conference participants at the CSEF-RCFS Conference on Finance, Labor and Inequality, IFN, NICE and the Swedish conference in Economics at the Stockholm School of Economics for valuable comments and suggestions. We are grateful for financial support from the Jan Wallanders och Tom Hedelius stiftelse samt Tore Browaldhs stiftelse (W19-0030, P22-0094), Marianne and Marcus Wallenberg Foundation (P2020.0049; P2024.0019; P2025.0107) and Torsten Söderbergs Stiftelse (ET2/20).

1 Introduction

Mergers and acquisitions (M&As) reallocate control over productive assets, organizational practices, and technological capabilities (Shleifer and Vishny, 1988; Maksimovic, Phillips, and Prabhala, 2011). They are commonly justified by operational improvements and synergies: the acquiring firm can combine its assets, technologies, and organizational practices with those of the target (Devos, Kadapakkam, and Krishnamurthy, 2009). Yet synergies are implemented inside firms by changing tasks, work processes, and the relative value of human capital. This means that the labor-market incidence of a merger need not be captured by the average wage effect on target workers. A technology-related acquisition can be neutral for workers on average while making some incumbent workers' skills less valuable and others more valuable.

This paper studies the worker-level incidence of technology-related merger synergies. We ask whether acquisitions by technologically advanced acquirers reprice tasks inside target firms, and whether the resulting gains and losses persist in workers' careers over the long run. The question matters because M&A is a major channel through which control over technologies and organizational capabilities is reallocated (Lagaras, 2021; Ma, Ouimet, and Simintzi, 2025). If these capabilities substitute for some tasks and complement others, then the market for corporate control is also a channel through which technological change is transmitted to worker careers in the long run.

Specifically, the hypothesis we investigate is that technology-related acquisitions change the demand for tasks inside the target. When the acquirer is intensive in software and data systems, integration may automate, standardize, or monitor tasks that overlap with software technologies (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011; Kogan, Papanikolaou, Schmidt, and Seegmiller, 2023). Workers performing such tasks may experience relative wage losses. At the same time, integration may raise the value of workers who perform tasks needed to implement, adapt, manage, or complement the acquirer's technologies. These workers should gain. The effects should be strongest when the acquirer is more technologically advanced than the target. The key prediction is therefore not a positive or negative average wage effect, but a task-specific redistribution that is strongest when acquirer capabilities overlap with target-worker task exposure.

We test this prediction using Swedish matched employer–employee data on cross-border acquisitions from 1997 to 2015. We focus on cross-border acquisitions because they provide variation in acquirer technological capabilities. The data allow us to identify workers employed at target firms before acquisition and follow them across employers for up to eight years. This distinction is central. Establishment-level wage effects may reflect worker turnover or changes in job composition (Ma et al., 2025). By following the original target workers, we estimate the career incidence of acquisition exposure on the workers initially affected by the transaction.

Our empirical design combines matched worker panels with stacked difference-in-differences and triple-difference specifications (Olsson and Tåg, 2017; Baker, Larcker, and Wang, 2022; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Olden and Møen, 2022). We compare workers at acquired firms with matched workers at non-acquired firms before and after acquisition. We then ask whether post-acquisition wage changes differ between workers in occupations with high and low exposure to software technologies, and whether this difference is concentrated when the acquirer is more software-intensive than the corresponding Swedish target industry. The identifying assumption is that, absent acquisition, the relative wage gap between exposed and less exposed workers would have evolved similarly in acquired and matched non-acquired firms. The main threat is selection: software-intensive acquirers may buy targets whose software-exposed workers would have experienced relative wage declines even without the transaction (Harford, 2005; Hershbein and Kahn, 2018; Jaimovich and Siu, 2020; Tuzel and Zhang, 2021; Bena, Ortiz-Molina, and Simintzi, 2022). We address this concern using dynamic event studies, pre-trend tests, splits by acquirer software intensity, earlier matching, offshoring falsification tests, robot-exposure checks, and estimates for workers who remain employed at the target firm.

The first result is that average wage effects are close to zero. This null average conceals substantial redistribution across tasks. Workers in software-exposed occupations experience persistent relative wage declines after acquisition. The baseline triple-difference estimate implies a wage decline of about 3 percent. The decline is concentrated in acquisitions by software-intensive acquirers, where the estimate is about 4 percent, and is small and statistically insignificant for less software-intensive acquirers. The dynamic estimates show no systematic differential pre-trends and a widening relative wage gap after acquisition.

The second result is that technology-related acquisitions also create winners. Workers in occupations exposed to earlier AI-related technologies, which proxy for cognitive and non-routine tasks that may complement digital upgrading, gain about 3-4% in acquisitions by software-intensive acquirers. Similar patterns and magnitudes appear for managers and professionals. These gains help explain why average wage effects remain close to zero: technology-related merger synergies redistribute gains and losses across workers rather than shifting all workers in the same direction.

Firm-level evidence supports this interpretation. Acquired firms employ more foreign managers after ownership change, consistent with organizational integration. Acquired firms also increase spending on data and software services, with the increase most visible among software-intensive acquirers. Among workers who remain employed at the target, software-exposed occupations experience even larger relative wage declines when the acquirer is software intensive. These patterns are consistent with persistent changes in task assignment, monitoring, bargaining, or pay setting inside the acquired firm.

Our paper contributes to the corporate finance literature on mergers by estimating the worker-level incidence of technology-related merger synergies. Prior work studies effects on firm performance, productivity, employment, and average worker outcomes, with less attention to the distribution of long-run career effects across incumbent target workers.[1] Closely related studies estimate worker effects of acquisitions (Prager and Schmitt, 2021; Arnold, 2025; He and Le Maire, 2025;

[1]For post-M&A labor reallocation, see Gehrke, Maug, Obernberger, and Schneider (2025); for privatizations, Arnold (2022); Olsson and Tåg (2025); for foreign acquisitions, Heyman, Sjöholm, and Tingvall (2007, 2011); Setzler and Tintelnot (2021); and for private equity buyouts, Agrawal and Tambe (2016); Olsson and Tåg (2017, 2018); Antoni, Maug, and Obernberger (2019); Cohn, Nestoriak, and Wardlaw (2021); Garcia-Gomez, Maug, and Obernberger (2026); Fang, Goldman, and Roulet (2026). Although they study buyouts rather than M&As, Agrawal and Tambe (2016) and Olsson and Tåg (2017) link IT investment or technological modernization to worker careers; the evidence is weaker in Antoni et al. (2019) and Fang et al. (2026).

Lagaras, 2026; Gehrke et al., 2025; Bach, Baghai, Bos, and Silva, 2025). We differ by isolating a technology-synergy mechanism: software-intensive acquirers reprice the value of tasks within the target workforce, generating persistent wage losses for workers in substitutable occupations and gains for workers in complementary occupations. The results show that technology-related merger synergies can leave average worker outcomes unchanged while persistently redistributing gains and losses across incumbent workers with different task exposures.

Our mechanism builds on Lagaras (2021) and Ma et al. (2025), who show that M&As can transmit technological and organizational change to target firms. At the *firm level*, Lagaras (2021) documents post-merger shifts toward high-skilled, knowledge-intensive, innovation-intensive, and IT-related labor and away from routine tasks, with evidence consistent with knowledge transfer, management changes, scale economies, and stronger effects for technologically advanced or foreign acquirers. At the *establishment level*, Ma et al. (2025) provide evidence that horizontal M&As reduce routine employment, increase technology-oriented employment and IT investment, and raise within-establishment wage dispersion. We estimate the *worker-level* incidence of technology-related merger synergies, which is important because firm- and establishment-level outcomes cannot distinguish changes in the composition of jobs inside target firms from changes in the careers of the incumbent workers initially exposed to the acquisition. By following these workers whether or not they remain at the target, we show that technology-related synergies reprice incumbent workers' tasks and generate persistent gains and losses across task groups.

Our paper also contributes to the international economics literature on technology transfer through multinational activity, foreign ownership, and cross-border acquisitions (Branstetter, 2006; Keller, 2010; Guadalupe, Kuzmina, and Thomas, 2012; Bloom, Sadun, and Van Reenen, 2012; Baziki, Norbäck, Persson, and Tåg, 2017; Davidson, Heyman, Matusz, Sjöholm, and Zhu, 2026), and to the broader trade-and-innovation literature surveyed by Keller (2010) and Akcigit and Melitz (2022). Related work studies how FDI and foreign ownership affect wages and wage inequality, including heterogeneity by worker type, source country, management practices, and multinational production networks (Helpman, 2017; Hale and Xu, 2020; Heyman et al., 2011; Conyon, Girma, Thompson, and Wright, 2002; Griffith and Simpson, 2004; Heyman, Norbäck, and Hammarberg, 2019; Setzler and Tintelnot, 2021). We connect these literatures by showing that cross-border acquisitions transmit digital capabilities in ways that generate task-specific and persistent wage effects for incumbent domestic target workers, shifting attention from firm-level productivity, innovation, and average wage premia to the distribution of gains and losses among initially exposed domestic incumbent workers.

The remainder of the paper is organized as follows. Section 2 describes the data and measurement of acquisitions, acquirer technology intensity, and occupational exposure. Section 3 presents the empirical design. Section 4 reports the worker-level results. Section 5 examines firm-level changes and within-firm wage effects among stayers. Section 6 reports additional analyses and robustness checks. Section 7 concludes.

2 Data

2.1 Data sources and classifying mergers and acquisitions

We use Swedish administrative data from Statistics Sweden for 1996–2015. The firm registers cover Swedish firms and report value added, book capital, employment, ownership status, sales, and industry. We also use Regional Labor Market Statistics (RAMS) to measure plant-level education and demographics, which we aggregate to the firm level. We additionally use Statistics Sweden survey data on firm-level expenditures on data and software services for a smaller subsample of firms and years in the analyses of post-acquisition IT investment.

The worker data come from the Salary Structure Statistics (*Lönestrukturstatistiken*), an annual survey administered by Statistics Sweden. The survey covers roughly half of private-sector employment and provides individual-level information on full-time equivalent monthly wages, education, occupation, gender, and employer identifiers.[2] Occupations use the Swedish Standard Classification

[2] The survey is drawn from Statistics Sweden's firm database and is stratified by industry and firm size. It includes

of Occupations (SSYK96), which is based on ISCO-88. Firms are legally required to respond, so the data contain reliable occupation and wage information.

We identify acquisitions by matching the firm registers to ownership data from the Swedish Agency for Economic and Regional Growth. The ownership data report the nationality of foreign multinational enterprises operating in Sweden. A firm is foreign-owned if foreign owners hold more than 50% of equity, and the largest owner's country defines the firm's nationality. We define a *foreign acquisition* as a switch from domestic to foreign ownership between two years. Some firms are later reacquired by Swedish owners or acquired by another foreign owner; after a second nationality change, we drop the firm and its workers so that the estimates capture the first foreign acquisition.

Figure 1 describes the matched acquisition sample by year, industry, and acquirer country. The worker-level sample contains 149,894 treated workers and an equal number of matched controls, while the firm-level corroborating analyses use 617 treated firms and 617 matched controls. Panel A of Table 1 reports balance statistics for the matched worker sample used in the wage regressions, while Panel B reports balance statistics for the separate matched firm-level sample used in the corroborating firm-level analyses. Manufacturing accounts for the largest share of acquisitions and affected workers. Foreign acquisitions are procyclical, with a spike in 2001 after the dot-com bust (Lerner and Tåg, 2013).

2.2 Measuring technological intensity and occupational exposure

We measure acquirer technology intensity with EU KLEMS 2019. EU KLEMS reports annual industry-level capital and labor statistics for EU countries, the United States, the United Kingdom, and Japan. We construct the ratio of software and database capital to total capital at the country–industry–year level. An acquirer is defined as software-intensive if this ratio in the acquirer's country, industry, and year is at least as high as the corresponding ratio in the target industry in

a representative annual sample of approximately 8,000–11,000 firms with more than ten employees, while firms with at least 500 employees are sampled with probability one.

Sweden. We observe the acquirer's nationality but not its industry, so we assign the acquirer to the target firm's industry. This introduces measurement error because the acquirer's true industry may differ from the target's industry. Classical misclassification would attenuate differences between high- and low-intensity acquirers, although technologically advanced acquirers may still select into particular Swedish targets.

We classify workers by occupational exposure to software, robots, and early (pre-ChatGPT) artificial intelligence (AI) using Webb (2020). These measures combine O*NET data on occupations and tasks with patent descriptions from the Google Patents Public Data. They capture the overlap between occupational tasks and patent text in each technology class. An occupation's score is the average of its task scores. A higher software-exposure score means that an occupation's tasks overlap more with software patents, which we interpret as greater exposure to technologies that may substitute for these tasks. Webb's AI-exposure measure similarly captures the overlap between occupational tasks and AI-related patents. Because the data end in 2015, the AI measure captures exposure to earlier generations of AI-related technologies rather than generative AI.

The Webb (2020) occupations use the U.S. SOC 2010 classification. We map SOC occupations to ISCO08 and then to two-digit SSYK96 occupations in the Swedish data. We express exposure as occupation-level percentiles. We classify workers as highly exposed if their occupation lies in the top decile of the baseline exposure distribution. The comparison group is all other occupations. Appendix Table A1 provides detailed variable definitions.

3 Empirical design

We estimate acquisition effects with stacked difference-in-differences and triple-difference designs. The design compares workers in acquired firms to matched workers in non-acquired firms with similar observable characteristics and pre-acquisition wage dynamics (Olsson and Tåg, 2017, 2018; Baker et al., 2022). The stacked design also avoids some of the heterogeneous-treatment-effect problems that arise in standard staggered two-way fixed effects models (Sun and Abraham, 2021;

Callaway and Sant'Anna, 2021; Baker et al., 2022).

3.1 Creating the treated and control groups

We construct the treated group from workers aged 26–54 who are employed at the target firm in the year before acquisition. We require non-missing information on wage, occupation, education, firm affiliation, and industry, and exclude workers employed at firms with fewer than 10 employees. The control pool consists of workers satisfying the same sample restrictions and employed at firms that are never acquired during the sample period.

For each acquisition cohort, we use coarsened exact matching (CEM) to match treated workers to comparable workers in non-acquired firms. Matching strata are defined using occupation, employment location in a metropolitan area, Swedish multinational (MNE) status, and calendar year. Within matched strata, we implement one-to-one matching without replacement.

We implement the matching procedure separately for each acquisition cohort. Once matched, control workers are removed from the control pool and cannot be reused in later cohorts. We then construct cohort-specific panels for treated and matched control workers spanning from four years before to eight years after the cohort-specific event year. Worker characteristics used as controls are measured in the year prior to acquisition. Finally, we stack the cohort-specific panels into a single event-time dataset centered on the acquisition year. This structure allows us to estimate difference-in-differences and triple-difference models within a common event-time framework across acquisition cohorts.

Panel A in Table 1 compares treated and matched control observations before acquisition. We assess balance with normalized t-values because conventional t-statistics are mechanically sensitive to sample size.[3] Treated and control workers are similar on the reported characteristics, with normalized t-values well below 0.25. The matched worker sample contains 149,894 treated workers and 149,894 matched controls. The typical worker in the matched sample is male, employed in a

[3] The normalized t-value divides the mean difference by the square root of the sum of the two variances (Imbens and Wooldridge, 2009). Absolute values above 0.25 indicate substantial imbalance.

metropolitan area, and has several years of pre-acquisition tenure.

The firm-level analyses use a separate matched sample constructed at the firm level. For each acquisition cohort, we retain firms observed in the year before acquisition and exclude firms with fewer than 10 employees, firms acquired more than once, and acquisitions without EU KLEMS information on acquirer technology intensity. Potential control firms are subject to the same restrictions, cannot have been previously acquired, and cannot be reused across cohorts. We then apply one-to-one coarsened exact matching within baseline-year cells defined by target industry, Swedish MNE status, and metropolitan area. After matching, we construct firm-level panels spanning from four years before to eight years after the event year and measure baseline firm characteristics in the year prior to acquisition for the firm-level regressions.

Panel B of Table 1 shows that the firm-level match achieves good balance on the reported baseline characteristics. The matched sample contains 617 treated firm-by-acquisition-cohort observations and 617 matched controls. Although treated firms are somewhat smaller than controls on average, with 374 versus 436 employees, and have slightly lower value added per worker, 0.577 versus 0.629, the corresponding normalized t-values are close to zero (-0.027 and -0.028, respectively), indicating only minor imbalance relative to the dispersion of the variables. We therefore treat the firm-level match as providing comparable treated and control firms on observables. The firm-level regressions additionally control for baseline firm size, value added per worker, skill share, and sales.

3.2 The stacked difference-in-differences model

We first estimate a standard difference-in-differences model for log wages of worker i in event year k and calendar year t:

$$w_{ikt} = \alpha_0 + \beta\,(Post_k \times Acq_i) + \alpha_1 Acq_i + \alpha_2 Post_k + \omega_t + \delta_{j(i)} + \lambda_{m(i)} + \psi_{s(i)} + X_i'\Gamma + X_f'\Pi + \epsilon_{ikt}. \quad (1)$$

Here, $Post_k$ equals one in the acquisition year ($k = 0$) and later years. Acq_i equals one for workers employed at a firm acquired one year later and zero for matched workers in firms that are never acquired. The coefficient β on $Post_k \times Acq_i$ is the average intention-to-treat effect of acquisition on worker wages.

The model includes calendar-year fixed effects ω_t, baseline industry fixed effects $\delta_{j(i)}$, baseline municipality fixed effects $\lambda_{m(i)}$, baseline Swedish-MNE fixed effects $\psi_{s(i)}$, worker controls X_i, and firm controls X_f. Worker controls are age, gender, education, experience, experience squared, an indicator for unemployment in years $k-4$ to $k-2$, and an indicator for at least three years of tenure at the target firm. Firm controls are log firm size, value added per worker, the share of high-skilled workers, and log sales. We measure all controls in the year before acquisition.

We then allow acquisition effects to vary with occupational exposure and acquirer technology intensity. The triple-difference model is

$$\begin{aligned} w_{ikt} = \alpha_0 &+ \beta_1 \left(Post_k \times Acq_i\right) + \beta_2 \left(Post_k \times Acq_i \times High_i\right) \\ &+ \alpha_1 Acq_i + \alpha_2 Post_k + \mu_1 High_i + \mu_2 \left(High_i \times Acq_i\right) + \mu_3 \left(High_i \times Post_k\right) \\ &+ \omega_t + \delta_{j(i)} + \lambda_{m(i)} + \psi_{s(i)} + X_i'\Gamma + X_f'\Pi + \epsilon_{ikt}. \quad (2) \end{aligned}$$

The triple difference compares wage changes across three margins: before versus after acquisition, treated versus control workers, and high-exposure versus other occupations. $High_i$ equals one if worker i is employed in a high-exposure occupation. The coefficient of interest is β_2 on $Post_k \times Acq_i \times High_i$. It measures whether high-exposure workers in acquired firms experience a different post-acquisition wage change than other workers, relative to the same difference in matched non-acquired firms. To study source-country technology heterogeneity, we also estimate equation (2) separately for acquisitions by high- and low-technology-intensity acquirers.

The DDD coefficient captures the relative post-acquisition wage change of high-exposure workers in acquired firms, net of (i) the average wage change of acquired-firm workers, (ii) the wage change of high-exposure workers in matched non-acquired firms, and (iii) common calendar-year, industry, municipality, and baseline firm-type shocks. We interpret the coefficient as evidence consistent with a technology-related merger synergy channel when it aligns with acquirer technology intensity and survives the falsification and robustness checks described below.

For the event-study figures and pre-trend tests, we replace $Post_k$ in equations (1) and (2) with event-time indicators for $k = -4, \ldots, 8$, omitting $k = -1$ as the baseline year. In the difference-in-differences event studies, the plotted coefficients are the event-time analogues of β. In the triple-difference event studies, they are the event-time analogues of β_2. The pre-trend tests jointly test whether the pre-acquisition coefficients for $k = -4, -3, -2$ equal zero. We cluster standard errors at the firm-by-acquisition-cohort level in the individual-level analysis to allow for common shocks among workers within the same acquisition cohort.[4] Standard errors in the firm-level analysis are clustered at the industry by event-year level.

3.3 Internal validity

The triple-difference design requires parallel trends in the triple difference, not parallel trends in each of the two underlying difference-in-differences comparisons (Olden and Møen, 2022). The central concern is that high-exposure workers in firms acquired by software-intensive foreign owners were already on different wage trajectories prior to acquisition.

We address five potential threats. First, acquisition selection could target firms whose exposed workers were already on different wage trajectories; the dynamic DDD estimates and joint pre-trend tests directly assess this concern. Second, country-industry technology shocks unrelated to the acquisition could affect exposed occupations at the same time as ownership change; the stacked design, baseline industry and municipality fixed effects, and comparison with matched non-acquired firms

[4]Standard errors are clustered at the target-firm-by-acquisition-cohort level because the stacked matching design allows the same control firm to appear in multiple acquisition cohorts with different assigned event years. In the treated sample, each firm belongs to a single acquisition cohort defined by its own acquisition year. In the control sample, however, different workers from the same non-acquired firm may be matched to treated workers from different acquisition cohorts. Clustering at the firm-by-cohort level therefore aligns the dependence structure of treated and control observations within the stacked event-time design.

help account for common shocks, while the high- versus low-acquirer-intensity splits test whether the pattern is specific to the acquirer's relative technological advantage. Third, generic restructuring or offshoring could produce similar wage losses; the offshoring-exposure falsification does not replicate the software-exposure pattern. Fourth, anticipation could contaminate the baseline match at $k = -1$; matches constructed three years before acquisition produce similar point estimates and leave the main software-intensity results largely unchanged. Fifth, worker sorting could mechanically generate wage effects if exposed workers leave acquired firms; the stayer results indicate that the wage decline is not driven solely by selective worker exits from the target.

The design also requires limited spillovers between treated and control workers (SUTVA). We therefore draw controls from the full population of non-acquired workers rather than from the same narrow local labor market. This choice reduces the chance that an acquisition in one part of Sweden affects the matched control workers.

3.4 External validity

Sweden is a high-income economy with above-OECD-average GDP per capita, high public trust, low corruption, and strong labor market protections. These institutions shape the interpretation of the estimates, and several features make the setting informative beyond Sweden.

First, Swedish employment law does not impose special employment conditions after ownership changes. Rights and obligations toward workers transfer to the new owner, and existing contracts remain in place unless workers renegotiate. Ownership change alone is not grounds for dismissal unless substantial organizational or economic restructuring occurs. Severance agreements are not guaranteed and are usually negotiated individually.

Second, collective bargaining agreements cover almost 90% of the Swedish labor market (Saez, Schoefer, and Seim, 2019). These agreements are renegotiated roughly every three years, but most contracts leave firms room to adjust individual wages. Wage setting is therefore more flexible than Sweden's employment protections might suggest.

Third, evidence suggests that Swedish labor markets behave similarly to labor markets in other developed countries, including other Scandinavian countries, Belgium, France, Germany, Italy, the Netherlands, and the United States (Lazear and Shaw, 2009). The magnitudes may depend on institutions, but the task-based incidence mechanism we study is likely not specific to Sweden.

4 The worker-level incidence of merger synergies

4.1 Average wage effects

We begin by examining average wage effects before turning to heterogeneity across worker tasks and acquirer technology intensity. Figure 2, Panel A plots the dynamic difference-in-differences estimates from equation (1) relative to the year before acquisition. Treated and control workers exhibit similar wage paths before acquisition, and the gap remains small afterward, showing that average wage effects are close to zero. Column 1 of Table 2 reports an average DiD estimate of 0.002, which is economically small and statistically insignificant. The absence of an average wage effect is consistent with earlier Swedish evidence on foreign acquisitions (Heyman et al., 2007).

If acquisitions transfer new technologies and organizational capabilities, wages may change in opposite directions for different task groups. We therefore next estimate whether wage effects differ between workers whose tasks are more likely to be substituted by software and workers in AI-exposed occupations, whose tasks may complement digital upgrading.

4.2 Substitutable tasks

We begin with occupations whose tasks are most exposed to software substitution. A large literature on task-biased technological change shows that digital technologies can substitute for workers performing routine and codifiable tasks (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Aghion, Antonin, Bunel, and Jaravel, 2022).

In the M&A setting, Ma et al. (2025) provide closely related establishment-level evidence that

U.S. horizontal M&As catalyze technology adoption at targets: routine-employment shares and routine-employment percentiles decline, technology-oriented employment rises, IT investment increases in deals where acquirers are more IT intensive than targets, and within-establishment wage dispersion increases. Their wage evidence is necessarily establishment-level and partly compositional because they do not observe individual wage histories. Our employer–employee data allow us to estimate persistent effects on incumbent target workers' careers. Olsson and Tåg (2017) additionally find that private equity buyouts similarly displace workers in routine and offshorable occupations.

The key prediction we want to test is that wage effects depend on both worker tasks and acquirer technology. When a software-intensive firm acquires a less software-intensive target, workers in software-exposed occupations should lose relative to other workers. We expect smaller effects when the acquirer is not software intensive or when workers' tasks are less exposed to software substitution.

Columns 2 and 3 of Table 2 and Panels B and C in Figure 2 show that the average masks heterogeneity: workers in software-exposed occupations experience a 2.8% relative wage decline (-0.028, column 2), while other workers do not experience a significant wage change.

The triple-difference estimates show that this heterogeneity is tied to software exposure. Figure 2, Panel D plots dynamic DDD coefficients comparing treated versus control workers, before versus after acquisition, and high versus lower software exposure. Before acquisition, relative wages of high-exposure and other workers evolve similarly in treated and control firms. After acquisition, workers in high-software-exposure occupations experience a clear relative wage decline. The decline begins immediately and widens over time, stabilizing near 5% after three to four years. Column 4 of Table 2 reports a post-period DDD estimate of -0.031, significant at the 1% level. This implies a 3.1% relative wage decline for workers in the most software-exposed occupations.

The effect is concentrated among acquisitions by software-intensive acquirers. Figure 2, Panels E and F split the dynamic DDD estimates by acquirer software intensity. For software-intensive acquirers, workers in software-exposed occupations experience a pronounced relative wage decline. For less software-intensive acquirers, the relative wage gap remains close to zero. Table 2, column 5 reports a DDD estimate of -0.041 for software-intensive acquirers, significant at the 1% level. Relative to treated workers' pre-acquisition mean wage in Table 1, the estimate corresponds to roughly SEK 920 (about EUR 85 or USD 100) per month, or SEK 11,000 (about EUR 1,015 or USD 1,190) per year.[5] Column 6 reports a smaller and insignificant coefficient of -0.009 for less software-intensive acquirers. The alignment between occupational exposure and acquirer technology intensity is difficult to explain with generic restructuring alone and is consistent with technology-related merger synergies leading to repricing of tasks inside firms.

4.3 Complementary tasks

Technology-related acquisitions can also create winners. Digital technologies may raise demand for workers performing tasks that complement new software capabilities (Acemoglu and Autor, 2011). Bloom, Garicano, Sadun, and Van Reenen (2014) show that IT investment facilitates delegation and creates demand for specialized managerial and professional roles that complement new technologies. In the M&A setting, acquirer technologies may raise productivity and wages for workers whose tasks complement the incoming digital infrastructure. Unlike software exposure, the expected wage effects of exposure to early AI technologies are less clear ex ante because early AI-related tasks may either be displaced or complemented by digital upgrading. If digital upgrading complements the tasks performed in AI-exposed occupations, these workers should benefit most when the acquirer is software intensive.

Figure 3, Panel D plots dynamic DDD estimates comparing treated versus control workers, before versus after acquisition, and AI-exposed versus other occupations. Pre-acquisition trends are similar. After acquisition, relative wages of workers in AI-exposed occupations begin to rise,

[5] Currency conversions use European Central Bank reference rates from May 8, 2026: EUR 1 = SEK 10.8420 and EUR 1 = USD 1.1761. Conversions are rounded to the nearest five units for monthly amounts and the nearest five or ten units for annual amounts.

reaching about 5% above their pre-acquisition trend six to eight years later. Table 3, column 4 reports a full post-period DDD estimate of 0.015, which is positive but imprecisely estimated.

The gains become clearer when we condition on acquirer technology intensity. Figure 3, Panel E shows that workers in AI-exposed occupations gain in acquisitions by software-intensive acquirers. Panel F shows a small and insignificant average effect for less software-intensive acquirers, although later event-year estimates are imprecise. Table 3, column 5 reports a DDD estimate of 0.036 for software-intensive acquirers, significant at the 5% level. This implies a 3.6% relative wage gain, corresponding to roughly SEK 9,500 (about EUR 876 or USD 1,030) per year at the pre-acquisition wage mean. Column 6 reports an insignificant estimate of -0.010 for less software-intensive acquirers. Appendix Table A2 shows similar results when we instead focus on managers and professionals. As in the software-exposure results, the relevant heterogeneity comes from the interaction between worker tasks and acquirer technology intensity.

The positive effects for AI-exposed workers help offset the negative effects for software-exposed workers, helping explain why average wage effects remain close to zero. The main high-intensity software estimate implies an annual loss of about SEK 11,000 for exposed workers, while the AI-exposure estimate implies an annual gain of about SEK 9,500. Both effects are concentrated among acquisitions by technologically advanced firms. Together, these patterns suggest that cross-border acquisitions redistribute the labor-market incidence of technology-related merger synergies within the target workforce.

5 Firm-level changes after acquisition

The worker-level results show that wage effects depend on the alignment between occupational task exposure and acquirer technology intensity. We next examine firm-level patterns that are consistent with the technology-related merger synergies interpretation. These analyses ask whether acquired firms receive managerial inputs, increase IT-related spending, and exhibit within-firm wage changes among employees who remain at the target firm.

5.1 Managerial integration

Cross-border acquisitions may transfer organizational capabilities through managers and technical personnel. Bloom et al. (2012) document that US multinationals bring more structured people-management practices to acquired plants, helping those plants benefit from IT investment. Gehrke et al. (2025) show that post-merger restructuring often changes the management layer, with managerial turnover exceeding turnover among operational employees. New managers may implement acquirer IT systems, introduce management practices, retrain workers, or reorganize teams. If this channel matters, acquired targets should experience an increase in foreign managers.

Figure 4, Panel A plots dynamic DiD estimates for the share of foreign managers at the target firm. Treated and control firms have similar pre-acquisition trends. After acquisition, the foreign-manager share rises persistently at acquired firms. Table 4, column 1 reports an average increase of 0.023, significant at the 1% level, relative to a pre-acquisition treated-firm mean of 0.100 in Table 1. The increase appears in both subsamples: column 2 reports 0.023 for software-intensive acquirers, and column 3 reports 0.018 for less software-intensive acquirers. Both estimates are significant at the 10% level. The high-versus-low DDD estimate in column 4 is 0.003 and statistically insignificant. The inflow of foreign managers therefore appears to be a general feature of cross-border acquisitions. This broader increase likely reflects standard post-merger integration and organizational coordination between the acquirer and target and is consistent with cross-border acquisitions bringing organizational inputs into target firms (Bloom et al., 2012; Heyman et al., 2019).

5.2 IT Investment

A second corroborating pattern is IT investment. Acquirers may bring financial resources or strategic priorities that help targets modernize IT infrastructure and production processes. Ma et al. (2025) show that acquirers increase IT investment at acquired establishments, and similar patterns appear after private equity buyouts (Agrawal and Tambe, 2016; Olsson and Tåg, 2017). These investments could automate routine tasks and change labor demand, especially when they take place in connection with new managers entering the organization.

We analyze expenditures on data and software services in a smaller firm-level subsample for which IT expenditure data are available. These data cover fewer firms and a shorter time period than the main firm-level sample, so the event-study window extends only four years after acquisition. Because the IT-expenditure variables contain many zero observations, we estimate these specifications using Poisson QMLE rather than log-linear OLS models. Figure 4, Panel E shows dynamic DiD estimates for these expenditures. Pre-acquisition trends are similar in the full sample and the high-intensity subsample, although the low-intensity subsample shows differential pre-trends. After acquisition, expenditures increase. Table 4, column 5 reports a DiD coefficient of 0.863, significant at the 1% level.

The within-sample increase is larger for software-intensive acquirers. Figure 4, Panel F and Table 4, column 6 show a coefficient of 1.264 for high-intensity acquirers, significant at the 1% level. Panel G and column 7 show a smaller and insignificant estimate for less software-intensive acquirers (0.122). The direct high-versus-low DDD estimate in column 8 is 0.558 and statistically insignificant. Thus, the within-sample patterns point toward stronger IT-spending increases among technologically advanced acquirers, although the direct high-versus-low comparison is imprecisely estimated. We therefore interpret these results as corroborating indications consistent with post-acquisition digital upgrading.

5.3 Within-firm wage effects (stayers)

Taken together, the foreign-manager and IT-expenditure results are consistent with organizational integration and digital upgrading as a result of technology-specific merger synergies. We next examine internal changes in task and pay setting. In particular, the wage estimates in Section 4 combine wage changes within firms with worker moves across firms. If the main results reflect changes in task assignment or pay setting inside the target firm, this should also appear among workers who stay. If the main results instead reflect selective worker mobility, the estimated wage effects should weaken among firm stayers. We use the stayer sample to assess whether the wage decline is also visible among workers who remain at the target, although remaining at the target firm is itself an endogenous outcome of the acquisition.

We therefore estimate wage effects among the workers who stay within the acquired firm. Table A3 reports DiD and DDD estimates for workers who remain employed at the target firm after acquisition. Column 4 reports a coefficient of -0.027, significant at the 5% level, implying a 2.7% relative wage decline for software-exposed stayers. The effect is concentrated in acquisitions by software-intensive acquirers. Column 5 reports -0.049, significant at the 1% level, implying a 4.9% relative wage decline for the stayers in high software-exposed occupations at targets acquired by technologically advanced firms. Column 6 reports no corresponding negative effect for less software-intensive acquirers (0.014, insignificant). The corresponding event-study specification for less software-intensive acquirers, however, exhibits evidence of differential pre-trends.

The stayer estimate of 4.9% exceeds the full-sample estimate of 4.1% in Table 2, column 5. This pattern suggests that selective worker exits are unlikely to fully explain the wage decline. The estimates are consistent with persistent within-firm changes in task assignment or pay setting after acquisition.

6 Additional analyses

We report robustness checks for anticipation, functional form, and alternative explanations.

6.1 Anticipation effects

Anticipation could bias the baseline match at $k = -1$ if workers or firms adjust before ownership changes. We address this concern by matching three years before acquisition ($k = -3$). Table A4 reports the results. The DDD estimates remain close to the baseline. Column 5 reports -0.035 for software-intensive acquirers, significant at the 1% level, compared with -0.041 in the baseline specification. Column 6 reports no significant effect for less software-intensive acquirers (0.002). The earlier match produces somewhat smaller but still economically meaningful estimates. While the low-intensity specification exhibits evidence of differential pre-trends, the main high-intensity estimate remains stable, suggesting that anticipation does not drive the baseline results.

6.2 Poisson QMLE regressions

Our baseline analysis uses full-time equivalent monthly wages, which provide a clean measure of wage rates but abstract from changes in working hours and labor-force attachment. An important question is therefore whether the results also hold when considering workers' actual annual labor-market earnings. To examine this, we replace monthly wages with annual labor income, which captures both changes in wage rates and changes in labor supply, including transitions into part-time work or unemployment. Because annual labor income also includes some zero-valued observations, log-linear OLS specifications are less suitable. Following the recommendations in Chen and Roth (2024), we instead estimate Poisson QMLE models using labor income in levels.

Table A5 reports the results. Columns 1–3 show DiD estimates. The average acquisition effect is economically small (column 1: 0.002). Workers in software-exposed occupations experience a 1.8% decline in labor income (column 2: −0.018), while other workers do not (column 3: 0.005). Columns 4–6 report DDD estimates. The triple interaction is negative and significant in the full sample (column 4: −0.027). The estimate is also negative for software-intensive acquirers (column 5: −0.036), though imprecisely estimated, and remains small and insignificant for less software-intensive acquirers (column 6: −0.015).

Overall, the Poisson QMLE estimates closely mirror the baseline wage results. This suggests that the findings are not specific to full-time equivalent monthly wages but also extend to realized annual labor income, a broader measure that incorporates variation in working hours and part-time employment.

6.3 Alternative technology and exposure: Robotics

We next examine whether similar patterns appear for industrial robots as an additional test of whether the results are tied to technology-specific task exposure. If the main results primarily reflected generic restructuring, they would be less likely to depend on technology-specific alignment between acquirer robot intensity and worker robot exposure.

We measure acquirer robot intensity using the IFR Robot Database, which reports annual industry-level robot stocks by country and industry. We define an acquirer as robot intensive if robot stocks per worker in the acquirer country-industry-year are at least as high as in the corresponding Swedish target industry. Because IFR excludes industries with very limited robot use, the robot sample covers fewer industries than the software sample.

Table A6 shows that treated and control workers remain well balanced in this restricted sample. Table A7 reports the wage estimates. Workers in robot-exposed occupations experience a 2.1% relative wage decline after acquisition (column 2: -0.021), while other workers do not. The DDD estimate in column 4 is -0.021, significant at the 10% level, indicating that robot-exposed workers fare worse relative to less exposed workers after acquisition.

The negative estimates are larger for robot-intensive acquirers. Column 5 reports -0.037 for robot-intensive acquirers, whereas column 6 reports -0.018 for less robot-intensive acquirers. The stronger negative estimate for robot-intensive acquirers is consistent with the technology-related synergies interpretation. Although the difference between the two subsamples is imprecisely estimated, the robot results are directionally similar to the software results and provide an additional technology-specific check on the interpretation.

6.4 Alternative exposure: Offshoring

Offshoring is a natural alternative explanation. If technologically advanced acquirers relocate routine tasks to lower-cost locations, workers in offshorable occupations may experience wage losses for reasons unrelated to the technology-specific mechanisms emphasized in the paper.

We test this explanation by replacing software exposure with occupational offshoring exposure, which captures whether tasks can be performed remotely and relocated abroad. Table A8 reports the estimates. Columns 1–3 show no significant average or subsample differences between workers in high- and low-offshoring-exposure occupations, although the high-offshoring DiD column shows evidence of differential pre-trends. Columns 4–6 report DDD estimates that interact treatment with offshoring exposure and acquirer technology intensity. Unlike the baseline software-exposure results, the triple-difference estimates are economically small and statistically insignificant in all specifications. The offshoring specifications, however, exhibit evidence of differential pre-trends, so these results should be interpreted cautiously. The main wage effects, however, appear more closely tied to technology-specific task exposure than to a broader offshoring channel.

7 Conclusion

This paper studies the worker-level incidence of technology-related merger synergies. Using Swedish matched employer–employee data on cross-border acquisitions from 1997 to 2015, we follow incumbent target workers before and after acquisition, whether or not they remain with the acquired firm. Average wage effects are close to zero. Beneath this average, however, acquisitions reprice tasks. Workers in substitutable occupations experience persistent relative wage declines, especially when the acquirer is software intensive. Workers in complementary occupations gain in the same type of acquisition, suggesting that some exposed tasks complement rather than substitute for acquirer technological capabilities.

The evidence points to task exposure as a central margin through which merger synergies affect worker careers. The negative effects for software-exposed workers are larger among stayers, which suggests that the results are not driven only by selective exits or changing workforce composition. Acquired firms also receive more foreign managers, while software-intensive acquisitions are associated with higher spending on data and software services. These firm-level patterns are consistent with a mechanism in which acquirer capabilities change task assignment, monitoring, bargaining, or pay setting inside the target firm. Thus, technology-related merger synergies can appear neutral for workers on average while changing the relative value of tasks within firms and having persistent effects on workers' careers.

References

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2022, June). The effects of automation on labor demand: A survey of the recent literature. In L. Y. Ing and G. M. Grossman (Eds.), *Robots and AI: A New Economic Era* (1 ed.)., pp. 15–39. London: Routledge.

Agrawal, A. and P. Tambe (2016, September). Private equity and workers' career paths: The role of technological change. *The Review of Financial Studies 29*(9), 2455–2489.

Akcigit, U. and M. Melitz (2022). International trade and innovation. In *Handbook of International Economics*, Volume 5, pp. 377–404. Elsevier.

Antoni, M., E. Maug, and S. Obernberger (2019, September). Private equity and human capital risk. *Journal of Financial Economics 133*(3), 634–657.

Arnold, D. (2022, October). The impact of privatization of state-owned enterprises on workers. *American Economic Journal: Applied Economics 14*(4), 343–380.

Arnold, D. (2025). Mergers and acquisitions, local labor market concentration, and worker outcomes. *SSRN Electronic Journal*.

Autor, D. H. and D. Dorn (2013, August). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., F. Levy, and R. J. Murnane (2003, November). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Bach, L., R. P. Baghai, M. Bos, and R. C. Silva (2025, October). How Do Mergers Affect the Mental Health of Employees? *Management Science*, mnsc.2023.04277.

Baker, A. C., D. F. Larcker, and C. C. Wang (2022, May). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics 144*(2), 370–395.

Baziki, S. B., P.-J. Norbäck, L. Persson, and J. Tåg (2017, May). Cross-border acquisitions and restructuring: Multinational enterprises and private equity firms. *European Economic Review 94*, 166–184.

Bena, J., H. Ortiz-Molina, and E. Simintzi (2022). Shielding firm value: Employment protection and process innovation. *Journal of Financial Economics 146*(2), 637–664.

Bloom, N., L. Garicano, R. Sadun, and J. Van Reenen (2014, December). The distinct effects of information technology and communication technology on firm organization. *Management Science 60*(12), 2859–2885.

Bloom, N., R. Sadun, and J. Van Reenen (2012, February). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review 102*(1), 167–201.

Branstetter, L. (2006, March). Is foreign direct investment a channel of knowledge spillovers? evidence from Japan's FDI in the United States. *Journal of International Economics 68*(2), 325–344.

Callaway, B. and P. H. Sant'Anna (2021, December). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Chen, J. and J. Roth (2024, May). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

Cohn, J., N. Nestoriak, and M. Wardlaw (2021, October). Private equity buyouts and workplace safety. *The Review of Financial Studies 34*(10), 4832–4875.

Conyon, M. J., S. Girma, S. Thompson, and P. W. Wright (2002, March). The productivity and wage effects of foreign acquisition in the united kingdom. *The Journal of Industrial Economics 50*(1), 85–102.

Davidson, C., F. Heyman, S. Matusz, F. Sjöholm, and S. C. Zhu (2026). From local to global: How foreign acquisitions reshape job mobility. *Journal of International Economics 162*, 104278.

Devos, E., P.-R. Kadapakkam, and S. Krishnamurthy (2009). How do mergers create value? a comparison of taxes, market power, and efficiency improvements as explanations for synergies. *Review of Financial Studies 22*(3), 1179–1211.

Fang, L., J. Goldman, and A. Roulet (2026). Private equity and pay gaps inside the firm. *Forthcoming in Journal of Finance*.

Garcia-Gomez, P., E. G. Maug, and S. Obernberger (2026). Private equity buyouts and employee health. *Forthcoming in Management Science*.

Gehrke, B., E. Maug, S. Obernberger, and C. Schneider (2025). Post-merger restructuring of the labor force. ECGI Finance Working Paper 753/2021, European Corporate Governance Institute. First posted 2021; last revised August 27, 2025.

Griffith, R. and H. Simpson (2004). Characteristics of foreign-owned firms in british manufacturing. In D. E. Card, R. Blundell, and R. B. Freeman (Eds.), *Seeking a Premier Economy: The Economic Effects of British Economic Reforms, 1980-2000*, National Bureau of Economic Research comparative labor markets series. Chicago: University of Chicago Press.

Guadalupe, M., O. Kuzmina, and C. Thomas (2012). Innovation and foreign ownership. *The American Economic Review 102*(7), 3594–3627.

Hale, G. and M. Xu (2020, May). Fdi effects on the labor market of host countries. In *Encyclopedia of International Economics and Global Trade*, pp. 285–304. World Scientific.

Harford, J. (2005). What drives merger waves? *Journal of Financial Economics 77*(3), 529–560.

He, A. X. and D. Le Maire (2025, July). Managing Inequality: Manager-Specific Wage Premiums and Selection in the Managerial Labor Market. *Review of Economics and Statistics*, 1–46.

Helpman, E. (2017, July). Globalisation and wage inequality. *Journal of the British Academy 5*, 125–162.

Hershbein, B. and L. B. Kahn (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review 108*(7), 1737–1772.

Heyman, F., P.-J. Norbäck, and R. Hammarberg (2019, April). Foreign direct investment, source country heterogeneity and management practices. *Economica 86*(342), 362–395.

Heyman, F., F. Sjöholm, and P. G. Tingvall (2007, November). Is there really a foreign ownership wage premium? evidence from matched employer–employee data. *Journal of International Economics 73*(2), 355–376.

Heyman, F., F. Sjöholm, and P. G. Tingvall (2011, May). Multinationals, cross-border acquisitions and wage dispersion: Multinationals, cross-border acquisitions. *Canadian Journal of Economics 44*(2), 627–650.

Imbens, G. W. and J. M. Wooldridge (2009, March). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature 47*(1), 5–86.

Jaimovich, N. and H. E. Siu (2020). Job polarization and jobless recoveries. *Review of Economics and Statistics 102*(1), 129–147.

Keller, W. (2010). International trade, foreign direct investment, and technology spillovers. In *Handbook of the Economics of Innovation*, Volume 2, pp. 793–829. Elsevier.

Kogan, L., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, National Bureau of Economic Research. Revised November 2025.

Lagaras, S. (2021). Corporate takeovers and labor restructuring. *SSRN Electronic Journal*.

Lagaras, S. (2026). M&As, employee costs, and labor reallocation. *Journal of Finance*. Forthcoming.

Lazear, E. P. and K. L. Shaw (2009, January). Wage structure, raises and mobility: An introduction to international comparisons of the structure of wages within and across firms. In *The Structure of Wages: An International Comparison*, pp. 1–57. University of Chicago Press.

Lerner, J. and J. Tåg (2013, February). Institutions and venture capital. *Industrial and Corporate Change 22*(1), 153–182.

Ma, W., P. P. Ouimet, and E. Simintzi (2025). Mergers and acquisitions, technological change and inequality. *Journal of Financial Economics 172*, 104136.

Maksimovic, V., G. Phillips, and N. R. Prabhala (2011). Post-merger restructuring and the boundaries of the firm. *Journal of Financial Economics 102*(2), 317–343.

Olden, A. and J. Møen (2022, March). The triple difference estimator. *The Econometrics Journal 25*(3), 531–553.

Olsson, M. and J. Tåg (2017, July). Private equity, layoffs, and job polarization. *Journal of Labor Economics 35*(3), 697–754.

Olsson, M. and J. Tåg (2018, November). Are foreign private equity buyouts bad for workers? *Economics Letters 172*, 1–4.

Olsson, M. and J. Tåg (2025, August). What Is the Cost of Privatization for Workers? *The Journal of Finance 80*(4), 2107–2151.

Prager, E. and M. Schmitt (2021, February). Employer consolidation and wages: Evidence from hospitals. *American Economic Review 111*(2), 397–427.

Saez, E., B. Schoefer, and D. Seim (2019, May). Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers' tax cut in Sweden. *American Economic Review 109*(5), 1717–1763.

Setzler, B. and F. Tintelnot (2021, June). The effects of foreign multinationals on workers and firms in the united states. *The Quarterly Journal of Economics 136*(3), 1943–1991.

Shleifer, A. and R. W. Vishny (1988). Value maximization and the acquisition process. *Journal of Economic Perspectives 2*(1), 7–20.

Sun, L. and S. Abraham (2021, December). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Tuzel, S. and M. B. Zhang (2021). Economic stimulus at the expense of routine-task jobs. *Journal of Finance 76*(6), 3347–3399.

Webb, M. (2020, January). The impact of artificial intelligence on the labor market. Working paper. Posted November 15, 2019; last revised January 11, 2020.

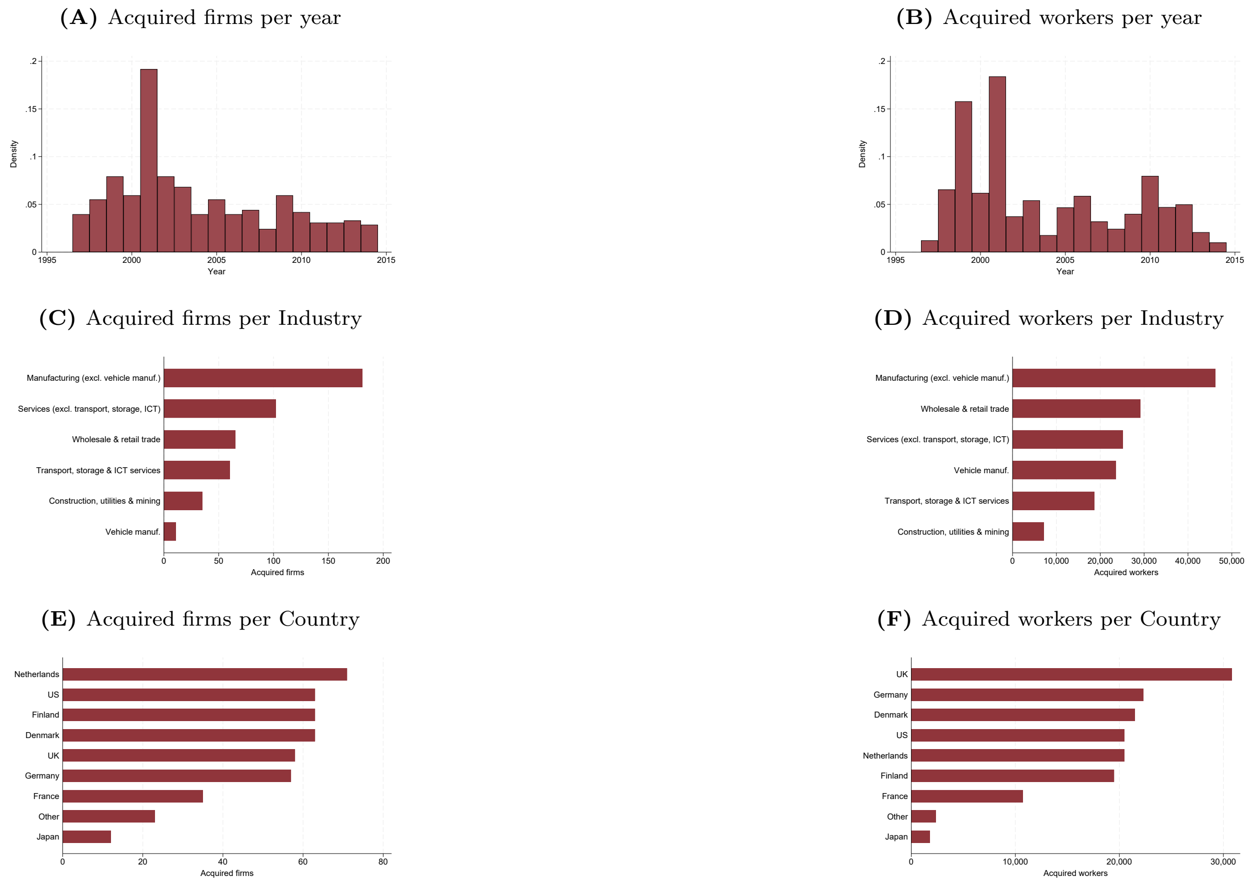

Figure 1: M&As in the sample

Panels A and B show the distribution of foreign acquisitions and treated workers across acquisition years. Panels C and D show the industry distribution at the firm and worker levels. Panels E and F show the most common acquirer countries. A firm is acquired when foreign ownership exceeds 50 percent and the largest owner's nationality changes from domestic to foreign. Worker counts refer to employees at the target firm in the year before acquisition.

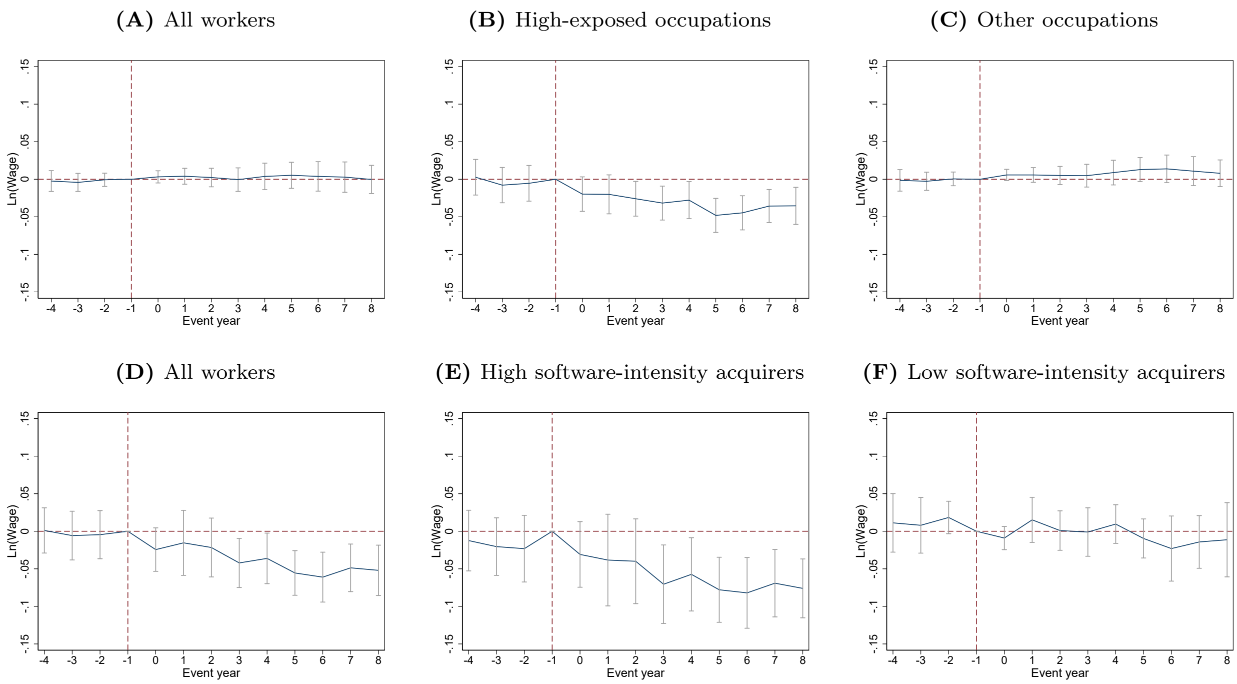

Figure 2: Wage effects by software-exposed occupations and acquirer technology intensity
The figure reports event-time coefficients from dynamic difference-in-differences (DiD) and triple-difference (DDD) regressions based on equations (1) and (2). The outcome is log wages. Event time is measured relative to the acquisition year ($k = 0$); the year before acquisition ($k = -1$) is omitted and marked by the vertical dashed line. The sample includes workers employed at target firms one year before acquisition and matched control workers in non-acquired firms. Panels A–C report DiD estimates for all workers, software-exposed occupations, and other occupations. Panels D–F report DDD estimates for all acquisitions and separately for high and low software-intensity acquirers. High-exposed occupations are occupations in the top decile of the software-exposure distribution. High software-intensity acquirers have country–industry software capital intensity at least as high as the target industry in Sweden. Points are coefficient estimates; vertical bars are 95% confidence intervals. Standard errors are clustered at the baseline target-firm-by-acquisition-cohort level.

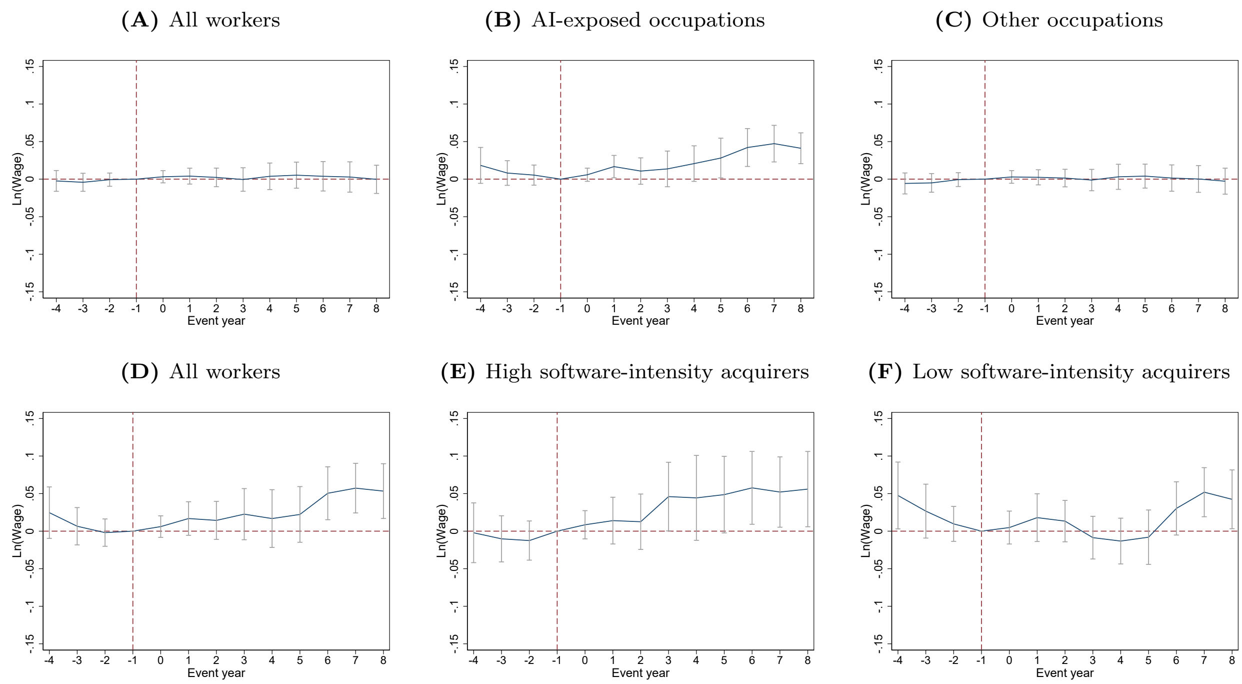

Figure 3: Wage effects by AI-exposed occupations and acquirer technology intensity
The figure reports event-time coefficients from dynamic difference-in-differences (DiD) and triple-difference (DDD) regressions based on equations (1) and (2). The outcome is log wages. Event time is measured relative to the acquisition year ($k = 0$); the year before acquisition ($k = -1$) is omitted and marked by the vertical dashed line. The sample includes workers employed at target firms one year before acquisition and matched control workers in non-acquired firms. Panels A–C report DiD estimates for all workers, AI-exposed occupations, and other occupations. Panels D–F report DDD estimates for all acquisitions and separately for high and low software-intensity acquirers. AI-exposed occupations are occupations in the top decile of the AI-exposure distribution. Positive wage effects for these workers in software-intensive acquisitions are interpreted as consistent with complementarity between AI-exposed tasks and acquirer digital capabilities. High software-intensity acquirers have country–industry software capital intensity at least as high as the target industry in Sweden. Points are coefficient estimates; vertical bars are 95% confidence intervals. Standard errors are clustered at the baseline target-firm-by-acquisition-cohort level.

(A) All acquisitions **(B)** High software-intensity acquirers **(C)** Low software-intensity acquirers **(D)** Technology intensity comparison

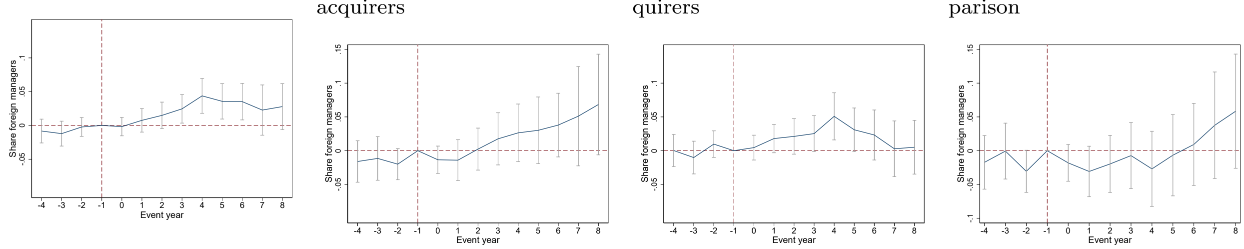

(E) All acquisitions **(F)** High software-intensity acquirers **(G)** Low software-intensity acquirers **(H)** Technology intensity comparison

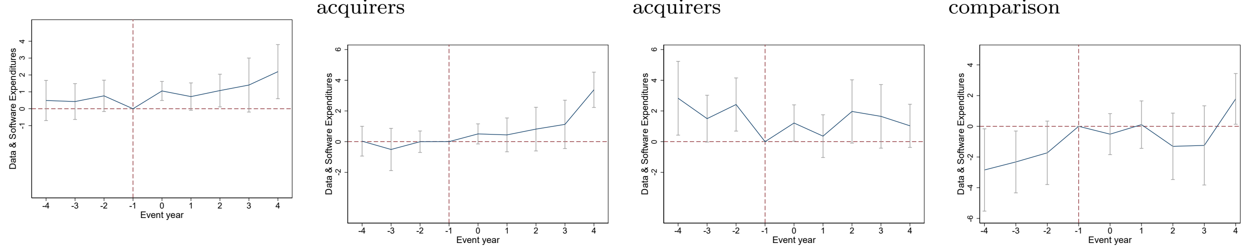

Figure 4: Corroborating firm-level patterns: foreign managers and IT investment

The figure reports event-time coefficients from dynamic difference-in-differences (DiD) and triple-difference (DDD) regressions corresponding to Table 4. The top row shows estimates for the share of foreign managers in the target firm. The bottom row shows Poisson QMLE estimates for firm expenditures on data and software services. Event time is measured relative to the acquisition year ($k = 0$); the year before acquisition ($k = -1$) is omitted and marked by the vertical dashed line. The sample is the matched firm-level sample. Each row reports estimates for all acquisitions, high software-intensity acquirers, low software-intensity acquirers, and the high-versus-low intensity comparison. Points are coefficient estimates; vertical bars are 95% confidence intervals. Standard errors are clustered at the baseline industry by event-year level.

Table 1: Pre-acquisition characteristics of treated and matched control samples

Panel A: Worker-level sample				
	Treated	Control	Difference	Norm. t-val
Log wage	9.996	9.990	0.005	0.012
Software exposure	0.539	0.539	0.000	0.000
AI exposure	0.544	0.544	0.000	0.000
Age	39.36	40.84	-1.484	-0.128
Female	0.348	0.341	0.007	0.010
Education (1–7)	3.739	3.689	0.050	0.026
Experience	20.60	22.15	-1.552	-0.124
Previous unemployment	0.294	0.274	0.020	0.031
Tenure $\geq$ 3 years	0.563	0.685	-0.122	-0.180
Major city	0.488	0.488	-0.000	0.000
Swedish MNE	0.519	0.519	-0.000	0.000
Observations	149,894	149,894	299,788	

Panel B: Firm-level sample				
	Treated	Control	Difference	Norm. t-val
Firm size (employees)	373.8	435.8	-62.1	-0.027
Value added per worker	0.577	0.629	-0.052	-0.028
Share high skilled	0.299	0.267	0.032	0.091
Software exposure	0.525	0.519	0.006	0.028
AI exposure	0.543	0.524	0.018	0.092
Swedish MNE	0.300	0.300	-0.000	0.000
Share foreign managers	0.100	0.079	0.020	0.094
Observations	617	617	1,234	

Notes: Panel A reports pre-acquisition means for the worker-level sample used in the wage regressions. Panel B reports pre-acquisition means for the separate firm-level sample used in the corroborating firm analyses. Treated observations in Panel A are workers employed at firms acquired in the following year. Treated observations in Panel B are firm-by-acquisition-cohort observations at $k = -1$, not unique acquisition events. Control observations are matched workers or firm-cohort observations not acquired. Column (3) reports treated minus control means. Column (4) reports the normalized t-statistic; absolute values above 0.25 indicate substantial imbalance (Imbens and Wooldridge, 2009). Software exposure and AI exposure are occupation-level exposure scores. The worker-level match uses occupation, employment location in a metropolitan area, Swedish multinational (MNE) status, and calendar year. The firm-level match uses target industry, Swedish MNE status, and metropolitan area.

Table 2: Wage effects by software-exposed occupations and acquirer technology intensity

	Difference-in-Differences			Triple Differences		
	All workers	High exposed	Other	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.002	-0.028***	0.005	0.005	0.002	0.009
	(0.007)	(0.008)	(0.007)	(0.007)	(0.011)	(0.007)
Post × Acq. × High Exp.				-0.031***	-0.041***	-0.009
				(0.011)	(0.014)	(0.012)
Pre-trend test (p-value)	0.869	0.687	0.903	0.925	0.732	0.322
R^2	0.498	0.508	0.494	0.505	0.513	0.502
Observations	2,204,859	193,182	2,011,676	2,204,859	1,069,998	1,134,861

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages. Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High Exp. denotes software-exposed occupations: occupations in the top decile of the software exposure distribution. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. Subsamples do not exactly sum to the full sample due to the included fixed effects. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 3: Wage effects by AI-exposed occupations and acquirer technology intensity

	Difference-in-Differences			Triple Differences		
	All workers	High exposed	Other	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.002	0.008	0.002	0.001	-0.007	0.009
	(0.007)	(0.010)	(0.006)	(0.006)	(0.010)	(0.007)
Post × Acq. × High Exp.				0.015	0.036**	-0.010
				(0.015)	(0.016)	(0.019)
Pre-trend test (p-value)	0.869	0.510	0.751	0.294	0.722	0.155
R^2	0.498	0.432	0.452	0.503	0.511	0.503
Observations	2,204,859	223,683	1,981,176	2,204,859	1,069,998	1,134,861

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages. Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High Exp. denotes AI-exposed occupations: occupations in the top decile of Webb's AI-exposure distribution, interpreted as a task-content proxy for cognitive, non-routine work that may complement digital upgrading. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 4: Corroborating firm-level patterns: foreign managers and IT investment

	Foreign Managers				IT Expenditures			
	DiD			DDD	DiD			DDD
	All	High tech	Low tech	All	All	High tech	Low tech	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × Acq.	0.023***	0.023*	0.018*	0.022**	0.863***	1.264***	0.122	0.537
	(0.008)	(0.013)	(0.010)	(0.010)	(0.312)	(0.473)	(0.468)	(0.432)
Post × Acq. × Int.				0.003				0.558
				(0.017)				(0.575)
Pre-trend test (p-value)	0.588	0.330	0.300	0.082	0.427	0.724	0.0098	0.074
R^2	0.156	0.245	0.180	0.156	0.858	0.924	0.871	0.866
Observations	10,368	3,862	6,505	10,368	516	243	266	516

Notes: The table reports firm-level DiD and DDD estimates after foreign acquisitions. Columns (1)–(4) report OLS estimates for the share of foreign managers in the target firm. Columns (5)–(8) report Poisson QMLE estimates for expenditures on data and software services from Statistics Sweden's IT usage survey; the outcome is available for base years 2009 onward. Int. denotes the high software-intensity acquirer indicator. High tech denotes acquirers whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline firm controls (log firm size, value added per worker, share high skilled, and log sales) and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the industry by event-year level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Appendix

Table A1: Variable definitions

Variable	Definition
Panel A: Worker-level variables	
Log wage	Full-time equivalent monthly real wage measured in November in the Salary Structure Statistics (Lönestrukturstatistiken).
Wage income in levels	Register-based wage-income variable used in the Poisson QMLE robustness table. The QMLE sample excludes observations with partial-year migration, students, pensioners, and reacquisition-related exclusions used in the wage analysis.
Software exposure	Occupation-level exposure to software technologies from Webb (2020). Exposure measures the overlap between occupational tasks and software-related patents. U.S. SOC occupations are mapped to ISCO08 and then to Swedish SSYK96 occupations.
High-exposed software occupations	Indicator equal to one if the worker's occupation is in the top decile of the baseline worker-level software exposure distribution, and zero otherwise.
AI-exposure	Occupation-level exposure to AI technologies from Webb (2020) used to proxy for cognitive and non-routine tasks that may complement digital upgrading. The score is based on Webb's AI-exposure measure.
High-exposed AI occupations	Indicator equal to one if the worker's occupation is in the top decile of the baseline worker-level AI exposure distribution, and zero otherwise.
Robot exposure	Occupation-level exposure to industrial robots from Webb (2020). Higher values indicate greater overlap between occupational tasks and robot-related patents.
Robot-exposed occupation	Indicator equal to one if the worker's occupation is in the top decile of the baseline worker-level robot exposure distribution, and zero otherwise.
Offshoring exposure	Occupational offshoring score assigned at the two-digit SSYK96 level in the pre-match data construction. High offshoring exposure denotes occupations in the top decile of the baseline worker-level offshoring score distribution.
Manager	Indicator equal to one for SSYK96 two-digit occupations 12 and 13.
Professional	Indicator equal to one for SSYK96 two-digit occupations 21, 22, and 24.
Age	Age from the Swedish population registry.
Female	Indicator equal to one for female workers and zero otherwise.
Education (1–7)	Highest completed education level from the Swedish Education Register, on a scale from 1 (primary education) to 7 (PhD).
Experience	Potential labor market experience constructed from age and education following standard labor economics conventions.
Previous unemployment	Indicator equal to one if the worker received unemployment benefits in years $t-4$ to $t-2$.
Tenure $\geq$ 3 years	Indicator equal to one if the worker has been employed at the firm for at least three years prior to acquisition.
Major city	Indicator for residence in one of Sweden's largest metropolitan areas, based on municipality classifications from Statistics Sweden.
Occupation	Worker occupation based on the two-digit SSYK96 classification.
Stayer	Indicator equal to one in worker-years in which the worker remains employed at the same firm as in the year before acquisition ($k=-1$), and zero otherwise.
Panel B: Firm-level variables	
Firm size	Number of employees in the firm in November of each year.
Log firm size	Natural logarithm of firm size.
Value added per worker	Firm value added divided by the number of employees.

Table A1 continued

Variable	Definition
Share high skilled	Share of workers in the firm with a university degree.
Swedish MNE	Indicator equal to one if the firm is a Swedish multinational enterprise.
Sales	Firm sales.
Log sales	Natural logarithm of firm sales.
Share foreign managers	Share of managers in the firm who are not Swedish nationals.
IT expenditures	Firm expenditures on data and software services from Statistics Sweden's IT usage survey, constructed as data services plus software services for base years 2009 onward.
Acquisition / Acq.	Indicator equal to one for treated firms or workers employed at treated firms, where a foreign acquisition occurs when foreign ownership exceeds 50 percent and the nationality of the largest owner changes from domestic to foreign.
High software-intensity acquirer / High tech	Indicator equal to one if the acquirer's country–industry software and database capital intensity is at least as high as the corresponding Swedish target-industry intensity in that year. Low tech is the complementary category.
High robot-intensity acquirer / High robot	Indicator equal to one if the acquirer's country–industry robot stock per worker is at least as high as the corresponding Swedish target-industry robot intensity in that year. Low robot is the complementary category.
Int.	High technology-intensity indicator used in firm-level triple-difference specifications. In the software tables, this is the high software-intensity acquirer indicator.
Post	Indicator equal to one in the acquisition year and all post-acquisition event years, and zero in pre-acquisition event years.
DiD / DID	Difference-in-differences interaction, $Post_k \times Acq_i$.
DDD / DIDID	Triple-difference interaction, $Post_k \times Acq_i \times High_i$, where $High_i$ denotes the relevant occupational exposure or acquirer-intensity indicator.

Table A2: Alternative definitions of complementary occupations

	Difference-in-Differences			Triple Differences		
	All workers	Managers	Others	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Managers						
Post × Acq.	0.002	0.008	0.001	0.001	-0.007	0.009
	(0.007)	(0.008)	(0.007)	(0.007)	(0.011)	(0.007)
Post × Acq. × Manager				0.011	0.029***	-0.008
				(0.009)	(0.010)	(0.013)
Pre-trend test (p-value)	0.869	0.103	0.907	0.513	0.237	0.040
R^2	0.498	0.491	0.519	0.564	0.579	0.557
Observations	2,204,859	165,513	2,039,346	2,204,859	1,069,998	1,134,861
Panel B: Professionals						
Post × Acq.	0.002	0.016*	0.000	-0.001	-0.009	0.007
	(0.007)	(0.008)	(0.006)	(0.006)	(0.010)	(0.007)
Post × Acq. × Professional				0.023*	0.037**	0.008
				(0.013)	(0.014)	(0.018)
Pre-trend test (p-value)	0.869	0.492	0.868	0.567	0.571	0.423
R^2	0.498	0.415	0.436	0.508	0.515	0.508
Observations	2,204,859	303,544	1,901,314	2,204,859	1,069,998	1,134,861

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages using alternative definitions of complementary occupations. Panel A defines complementary occupations as managers (SSYK96 two-digit occupations 12 and 13). Panel B defines complementary occupations as professionals (SSYK96 two-digit occupations 21, 22, and 24). Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A3: Within-firm wage effects for workers remaining after acquisition

	Difference-in-Differences			Triple Differences		
	All workers	High-exposed	Other	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	-0.006	-0.038***	-0.001	-0.002	0.004	-0.006
	(0.008)	(0.008)	(0.008)	(0.008)	(0.012)	(0.007)
Post × Acq. × High Exp.				-0.027**	-0.049***	0.014
				(0.012)	(0.013)	(0.014)
Pre-trend test (p-value)	0.983	0.777	0.998	0.962	0.756	0.009
R^2	0.501	0.499	0.498	0.507	0.515	0.503
Observations	1,825,827	150,246	1,675,577	1,825,827	925,912	899,915

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages for stayers. Stayers are worker-years in which the worker remains employed at the same firm as in the year before acquisition ($k = -1$). Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High Exp. denotes software-exposed occupations: occupations in the top decile of the software exposure distribution. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A4: Anticipation robustness by matching three years before acquisition

	Difference-in-Differences			Triple Differences		
	All workers	High-exposed	Other	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.003	-0.017**	0.007	0.007	0.009	0.005
	(0.009)	(0.009)	(0.009)	(0.009)	(0.014)	(0.009)
Post × Acq. × High Exp.				-0.025**	-0.035***	0.002
				(0.011)	(0.013)	(0.020)
Pre-trend test (p-value)	0.386	0.129	0.490	0.190	0.636	0.003
R^2	0.503	0.522	0.493	0.510	0.507	0.532
Observations	1,019,797	96,940	922,857	1,019,797	597,799	421,998

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages when treated and control workers are matched three years before acquisition rather than one year before. Treated workers are employed at firms acquired three years later and are matched to workers in non-acquired firms. High Exp. denotes software-exposed occupations: occupations in the top decile of the software exposure distribution. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A5: Robustness using Poisson QMLE estimates

	Difference-in-Differences			Triple Differences		
	All workers	High-exposed	Other	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.002	-0.018*	0.005	0.005	0.004	0.006
	(0.006)	(0.011)	(0.006)	(0.006)	(0.011)	(0.006)
Post × Acq. × High Exp.				-0.027*	-0.036	-0.015
				(0.014)	(0.023)	(0.011)
Pre-trend test (p-value)	0.566	0.546	0.576	0.569	0.174	0.995
Pseudo R^2	0.403	0.291	0.404	0.408	0.403	0.415
Observations	2,692,136	244,178	2,447,958	2,692,136	1,257,539	1,434,597

Notes: The table reports Poisson QMLE estimates of foreign-acquisition effects on register wage income measured in levels. The QMLE sample excludes observations with partial-year migration, students, pensioners, and reacquisition-related exclusions used in the wage analysis. Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High Exp. denotes software-exposed occupations: occupations in the top decile of the software exposure distribution. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A6: Comparison of treated and control workers in the IFR sample

	Treated	Control	Difference	Norm. t-val
Log wage	9.987	9.983	0.004	0.009
Robot exposure	0.524	0.524	0.000	0.000
Age	39.64	40.81	-1.174	-0.101
Female	0.316	0.309	0.007	0.011
Education (1–7)	3.688	3.619	0.069	0.036
Experience	20.95	22.22	-1.276	-0.102
Previous unemployment	0.281	0.264	0.016	0.026
Tenure $\geq$ 3 years	0.574	0.683	-0.109	-0.161
Major city	0.451	0.451	-0.000	0.000
Observations	111,921	111,921	223,842	

Notes: The table reports pre-acquisition means for treated and matched control workers in the IFR robot sample. Treated workers are employed at firms acquired in the following year. Column (3) reports treated minus control means. Column (4) reports the normalized t-statistic; absolute values above 0.25 indicate substantial imbalance (Imbens and Wooldridge, 2009).

Table A7: Wage effects by robot exposure and acquirer robot intensity

	Difference-in-Differences			Triple Differences		
	All workers	Robot-exposed	Other	All workers	High robot	Low robot
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.001	-0.021**	0.005	0.005	0.002	0.004
	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)
Post × Acq. × Robot-exposed				-0.021*	-0.037***	-0.018
				(0.012)	(0.012)	(0.014)
Pre-trend test (p-value)	0.756	0.129	0.488	0.210	0.466	0.667
R^2	0.488	0.505	0.478	0.498	0.514	0.495
Observations	1,684,374	210,311	1,474,033	1,684,344	451,063	1,233,281

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages using occupational exposure to industrial robots. Robot-exposed denotes occupations in the top decile of the robot exposure distribution from Webb (2020). High robot acquirers have country–industry robot intensity at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A8: Robustness using offshoring exposure

	Difference-in-Differences			Triple Differences		
	All workers	High offsh.	Low offsh.	All workers	High tech	Low tech
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Acq.	0.002	0.004	0.002	0.002	-0.006	0.009
	(0.007)	(0.008)	(0.007)	(0.007)	(0.012)	(0.007)
Post × Acq. × High Offsh.				0.000	0.009	-0.009
				(0.010)	(0.013)	(0.017)
Pre-trend test (p-value)	0.869	0.0125	0.378	0.000	0.005	0.004
R^2	0.498	0.471	0.503	0.504	0.516	0.501
Observations	2,204,859	257,319	1,939,611	2,196,932	1,069,000	1,127,932

Notes: The table reports DiD and DDD estimates of foreign-acquisition effects on log wages using occupational offshoring exposure as the heterogeneity dimension. Treated workers are employed at firms acquired in the following year and are matched to workers in non-acquired firms. High Offsh. denotes occupations in the top decile of the baseline worker-level offshoring exposure distribution. High tech denotes high software-intensity acquirers, whose country–industry software capital intensity is at least as high as the target industry in Sweden. All regressions include baseline worker controls (age, gender, education, experience, experience squared, unemployment history, and tenure), baseline firm controls (log firm size, value added per worker, share high skilled, and log sales), and calendar-year, baseline industry, baseline municipality, and baseline Swedish-MNE fixed effects. Standard errors, in parentheses, are clustered at the baseline target-firm-by-acquisition-cohort level. The pre-trend test reports the p-value from a joint F-test that all pre-acquisition event-time coefficients in the dynamic event-study specification equal zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.