

Organized Labor Versus Robots?

Evidence from Micro Data*

Sebastian Findeisen[†]  Wolfgang Dauth[‡]  Oliver Schlenker[§]

January 5, 2025.

Abstract

New technologies drive productivity growth, yet the distribution of gains may be unequal. We study how labor market institutions – specifically shop-floor worker representation – mediate the impact of automation. Combining German individual-level administrative records with plant-level data on industrial robot adoption, we find that works councils reduce the separation risk for incumbent workers during automation events. When labor markets are tight and replacement costs are high, incumbent workers become more valuable from the firm’s perspective. Consequently, we document that the moderating effects of works councils diminish. Older workers, who face greater challenges reallocating to new employers, benefit the most from organized labor in terms of wages and employment. Finally, we observe that works councils do not hinder robot adoption; rather, they spur the use of higher-quality robots, encourage more worker training during robot adoption, and foster higher productivity growth thereafter.

*We thank Fabrizio Colella, Thomas Cornelissen, Christian Dustmann, Bernd Fitzenberger, Guido Friebel, Georg Grätz, Simon Jäger, Philipp Kircher, Steffen Müller, Peter Nilsson, Torsten Persson, Pascual Restrepo, David Seim, Jens Stegmaier, Simon Trenkle, Joonas Tuhkuri, and numerous seminar participants for very helpful comments. We especially thank Mario Bossler and Martin Popp for sharing their measure of labor market tightness. Further, we acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG – German Research Foundation) under Germany’s Excellence Strategy – EXC-2035/1 – 390681379.

[†]University of Konstanz - Department of Economics & Cluster of Excellence "The Politics of Inequality" and CEPR

[‡]Institute for Employment Research (IAB) and University of Bamberg

[§]University of Konstanz - Department of Economics & Cluster of Excellence "The Politics of Inequality" and ZEW Mannheim

1 Introduction

Economists have long acknowledged that technological advances do not necessarily guarantee widely shared gains from productivity growth, especially in the short-run (Keynes, 1930; Gordon, 2016). History offers numerous examples of conflicts between workers and capital owners over how the benefits and costs of new technologies ought to be distributed (Acemoglu and Johnson, 2023).¹ Recently, the rise of automation through robotics and artificial intelligence (AI) has sparked renewed debates about how workers, employers, and governments can navigate future labor market disruptions (Furman and Seamans, 2019; Autor, 2024).

A recent literature studies the effects of industrial robots and automation technologies on employment and wages, uncovering strong heterogeneity across skill groups, occupations, industries, and firms (e.g. Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Bonfiglioli et al., 2022). Acemoglu and Restrepo (2022) find that rapid automation in the US can account for the largest share of wage declines among workers specialized in routine tasks. Dauth et al. (2021) detect that negative employment effects of robotization are concentrated among regions with low levels of unionization in Germany, providing a hint for the importance of labor market institutions. However, so far, the literature has paid very limited attention to the roles that labor relations and the relative bargaining power of workers and firms (Stansbury and Summers, 2020) play as mediators of technological change.

In this paper, we shed new light on the interaction between labor market institutions and automation with the goal of advancing the debate on whether and how policy responses could be deployed in light of ongoing technological disruption. We focus on co-determination, in the form of work councils, which grant co-decision-making rights to organized labor at the establishment level (Addison, 2009; Jäger et al., 2022). In Germany, works councils represent approximately 40% of the workforce and have potent means to protect workers' employment and working conditions. Their power ranges from veto rights against dismissals (that can only be overruled by labor courts) to co-determination in matters such as working hours and workplace monitoring.

It is typically assumed that firms automate as long as it is profitable, not internalizing

¹For instance, in Britain in the 18th century, the power loom increased productivity and profits for machine owners but massively replaced skilled weavers and let workers' wages deteriorate (Acemoglu and Johnson, 2024).

the consequences for displaced workers.² This can lead to only marginally profitable ('so-so') automation where the productivity gains from automation are small relative to the employment and earnings losses for some workers (Acemoglu and Restrepo, 2018). As works councils protect the interests of the incumbent workforce, their presence and rights to be involved in procedural changes associated with technology adoption may alter both the process and consequences of automation. This might lead to different wage and employment outcomes for workers during automation events in establishments with and without works councils.

In our analysis, we utilize detailed linked employer-employee administrative data combined with establishment surveys. The data and the institutional context in Germany allow us to leverage variation in robot usage between and within plants over time (Plümpe and Stegmaier, 2023), taking into account the presence of works councils (across plants).³ To this end, we match incumbent workers based on their own and their establishments' characteristics prior to robot adoption and additionally restricting matches to be within the same industry, occupation, and works council status. Then we compare the evolution of matched incumbent workers' employment outcomes in the wake of robot adoption using event studies, separately for workers in establishments with and without works council. By computing the difference between these two distinct 'difference-in-differences' estimators, we are able to examine the role of works councils during times of robot adoption. Our rich firm-level data allows us to perform extensive robustness checks to rule out that results on the effect of robot adoption are confounded by other events happening at the time of adoption, such as the accompanying adoption of information and communication technologies (ICT) and firm expansion. Furthermore, our fine-grained worker-level administrative records enable us to not only account for workers' careers before adoption but also examine which type of workers are most strongly affected by robotization and protected by works councils. The survey data allows us to complement the main analysis with an examination of the mechanisms at the plant level, including training for workers, productivity changes, as well as the direction and intensity of technology.

The first finding of the paper is that automation events increase retention of incumbent blue-collar workers by decreasing separation probabilities – but only in plants with active

²This can make government intervention welfare improving. Beraja and Zorzi (2024) show that under frictional reallocation with unemployment spells and the presence of borrowing constraints, the optimal policy for the government is to slow down automation.

³The German economy has one of the highest robot densities in the world, providing rich variation in adoption events across plants.

works councils. Consistent with the idea that works councils, as a form of shop-floor representation, act in the interest of incumbent workers, we estimate a positive retention effect of around 4 percentage points. The observed increase in employment stability during automation events is consistent with rent-sharing manifesting through enhanced job security rather than wage growth. This pattern aligns with the institutional mandate of works councils, which grants them direct co-determination rights over worker dismissals but offers only limited influence over wage increases.⁴

Older workers are the main beneficiaries of these policies as their employment rates increase by 2 percentage points in subsequent years. Following automation events, younger displaced workers adjust successfully and transition smoothly into new employment. In contrast, older workers are more likely to remain unemployed after an automation-induced separation, consistent with rising adjustment costs over the life cycle.⁵

Next, we investigate firms' incentives to shield incumbent workers from layoffs. In frictional labor markets, the value of retaining incumbent workers in automating plants increases in labor market tightness, since replacement and recruitment costs rise as labor becomes harder to find. Hiring difficulties, therefore, align the incentives of management and incumbent workers – represented by works councils. Indeed, sample splits by firm-specific labor market tightness reveal that works council representation only leads to higher retention when firm-specific replacement costs are low. In contrast, when replacement costs are high, automation leads to similar retention rates in plants with and without works councils.

Works councils may also engage in bargaining strategies that limit wage cuts for vulnerable workers, particularly those whose task profile is prone to automation or who face poor outside options.⁶ To test this, we examine wages for two groups: (1) older employees, who face documented difficulties in adjusting and thus have poorer outside options, and (2) workers performing routine-manual tasks, thus are confronted with high automation risk. We find that works councils have a sizeable positive wage effect of around 3.5% for these vulnerable workers, driven the prevention of wage cuts.

⁴Worker representation might affect wage decreases, by reducing or preventing them, a point we elaborate on below.

⁵Older workers may have also acquired more task, industry, or firm-specific human capital, which makes transitions across these categories more challenging.

⁶While works councils do not have an official mandate to negotiate wages, they can influence the pay groups within a collective agreement that individual workers are classified into. In addition, their powers in other fields are strong enough to provide incentives for employers to cooperate also in fields that are not covered by their statutory powers.

In the final part of the paper, we investigate by which means works councils mitigate the negative effects of automation for incumbent workers. When automation-related investment decisions internalize (part of) the displacement costs, they effectively raise the profitability threshold for marginal investments. All else being equal, this diminishes the incentives to implement 'so-so technologies,' i.e., automation investments that strongly displace workers while yielding only modest productivity gains (Acemoglu and Restrepo, 2018). Consequently, automation events in firms with worker representation should go along with larger productivity gains relative to displacement, as those must counterbalance the internalized displacement costs. In line with this argumentation, we find that robot-adopting plants with works councils experience greater post-adoption productivity growth than comparable plants without entrenched worker representation.

Using proxies for robot quality, we find that this productivity growth is primarily driven by higher investment in the quality, but not the quantity, of robots acquired. Furthermore, we find that adopters with works councils provide more training for their workers during robot adoption events. These investments into the human capital of incumbent workers likely contribute directly to increased retention. Exploiting the panel dimension of our data, we can detect 'training spikes' around the time of adoption, solidifying this interpretation. Before automation events, firms with works councils already exhibit a higher baseline share of workers who participate in training compared to their counterparts. During automation events, this gap widens from around 5 percentage points to about 15 percentage points, before ultimately reverting to around 5 percentage points in subsequent years. This could be due to co-determination rights enhancing job security in the face of automation, thereby making workers more willing to invest in firm-specific skills and engage in training, as argued by Freeman and Lazear (1995).

Our paper contributes to the literature linking firm- and worker-level data to understand how new automation technologies affect labor market outcomes. Robot-adopting plants typically expand employment (Koch et al., 2021; Hirvonen et al., 2023), often at the expense of competing firms. On average, workers performing replaceable tasks incur losses (Bessen et al., 2023), while other workers may experience gains (Acemoglu et al., 2023), with effects varying by age group (Deng et al., 2023). Acemoglu et al. (2023) also highlight the importance of controlling for worker sorting through worker fixed effects, which we incorporate into our event-study analysis. To the best of our knowledge, there are no previous studies that have explored the role of labor market institutions or

leveraged variations in the relative decision-making power between workers and firms.

Similarly, equilibrium studies at the local labor market level, such as those by Acemoglu and Restrepo (2020) or Adachi et al. (2024), focus on variation across sectors or skill levels. While Acemoglu and Restrepo (2020) document negative employment effects in the U.S., Adachi et al. (2024) find an expansion of employment in Japan. Dauth et al. (2021) identify an interaction between the displacement effect at the local labor market level and local unionization rates. In contrast, in this paper we use detailed micro data on firm-level adoption and individual worker trajectories, employing event studies to demonstrate the dynamic interaction between institutions and automation.

Jäger et al. (2021) show that board-level participation of workers can increase capital investment rather than decrease it. Consistent with this, our findings indicate that automation events are associated with greater productivity growth in establishments that have shop-floor worker representation. Relatedly, Addison et al. (2001), Jäger et al. (2022), and others have highlighted a positive correlation between productivity and the presence of works councils. The presence of works councils has been associated with reduced separations of (blue collar) workers in work by Hirsch et al. (2010) and a shift of bargaining power towards employees, as found in Dobbelaere et al. (2024). Budde et al. (2024) find that elected blue-collar workers as works council representatives place a strong weight on retention of incumbent workers. Dustmann et al. (2014) have argued that work councils have facilitated wage decentralization in the German labor market crisis in the mid 2000's, elevating employment in Germany. As noted by Jäger et al. (2022) in their survey on the effects of co-determination "due to a lack of sharp and exogenous variation, the effects of works councils on worker and firm outcomes remain an open research question." We contribute to this literature on shared governance by being able to identify an interaction effect between works councils and technology adoption.

In addition, our empirical strategy – employing event studies with matched worker-level control groups – is related to papers in the literature studying the cost of job loss (Bertheau et al. (2023) and Illing et al. (2024)).

In the next section, we describe the data and briefly discuss the institutional background. In Section 3, we present the empirical models and strategy. Sections 4 and 5 contain the main results regarding employment and wages, respectively. In Section 6, we present underlying mechanisms with a focus on technology adoption and worker training. Section 7 concludes.

2 Data

For our main analysis in Sections 4 and 5, we use a plant-level survey containing information on the presence of works councils and robot usage and link it with the employment biographies of the plants' entire workforce. Our plant data stem from the Establishment Panel (Bellmann et al., 2021) of the Institute for Employment Research (IAB), a representative annual survey of around 15,000 establishments in Germany. The survey comprises data on, among others, general information on the plant, workforce structure and trends, labor relations and co-determination, as well as information on the plant's technical endowment.

Automation and First-Time Robot Adoption Events. The Establishment Panel is augmented by questions on current topics on a yearly basis. Notably, the wave 2019 contains information on robot usage between 2014 and 2018 which we use to construct robot adoption events. We first distinguish between robot users and plants that have never used robots up until 2019. Among the set of robot users, we again distinguish between plants that newly adopted robots between 2015 and 2018 and incumbent users, i.e. plants that already used robots in 2014. We follow the seminal studies of Graetz and Michaels (2018) and Acemoglu and Restrepo (2020) and interpret the event of installing robots for the first time as an event where firms automate routine-manual tasks.

The recent literature on the plant-level effects of robot adoption emphasizes that robot adopters are inherently different from non-adopters. As in Koch et al. (2021) for Spanish manufacturing firms in the period of 1990 to 2016, Deng et al. (2024) show that newly robot-adopting German plants are positively selected regarding size and productivity. However, Appendix Table A.1 highlights that plants already using robots in 2014 and those that adopted them between 2015 and 2018 (i.e., our sample of firms) share common characteristics. Specifically, both groups are positively selected in terms of employment size, workforce composition, and productivity, and they face hiring constraints to a similar extent. Moreover, robot users are predominantly present in the manufacturing sector, though we also observe increasing robot adoption in non-manufacturing establishments (27% for new adopters compared to only 17% for early adopters). In contrast, establishments in the manufacturing sector that had not adopted robots by 2018 are, on average, smaller, employ fewer skilled workers, and are founded more recently. Thus, although we focus on more recent adopters, the same key factors remain strongly associated with robot

adoption as in prior periods and studies. Additionally, Plümpe and Stegmaier (2023) show that the robot density obtained from the survey correlates strongly with commonly used industry-level data from the International Federation of Robotics (IFR).

Works Councils. Another advantage of the IAB Establishment Panel is that it provides information on the presence of a works council. The German Works Constitution Act (BetrVG) stipulates that workers in plants with at least five permanent employees have the right to elect a works council (Jäger et al., 2022). Since the establishment of a works council requires an initiative by the employees, by far not all plants have one. 41 percent of all German workers in 2015 were employed in a plant that had a works council, but this share varies from nine percent in plants with between five and 50 employees and 89 percent in plants with more than 500 employees (Ellguth and Kohaut, 2015). Addison (2009) and Mohrenweiser (2022) provide extensive overviews of the powers of works councils in Germany and their economic implications. Those powers range from consultation in events of technology adoption over veto rights in cases of hirings, dismissals, and internal transfers (that can only be overruled by labor courts) to co-decision rights in matters that concern, e.g., working hours, workplace monitoring, or performance pay. While they have no mandate to bargain over wages directly, they can negotiate in which pay group an individual worker is classified within a firm’s collective agreement. This might be particularly important if employers plan to downgrade production workers in routine-manual occupations who directly compete with robots. Since they can stall or even prevent dismissals, they can also incentivize employers to pay efficiency wages. However, it is important to notice that works councils are usually interested in the success of their firms and may in fact be beneficial since they may raise worker satisfaction and awareness about the economic state of the firm, raise efficiency by improving the communication on work processes (Freeman and Lazear, 1995), and identify specific training needs.

Administrative Worker Data. Since the Establishment Panel is based on a random sample of plants with at least one employee subject to social security contributions, it can be matched with administrative data from the IAB.⁷ Specifically, we use the Integrated Employment Biographies (IEB v16.01.00) prepared by the IAB, which comprises the full universe of all individuals who have held a job subject to social security contributions,

⁷This link is only available for plants who have either given explicit consent in the 2020 wave of the survey or have dropped out of the survey earlier.

marginal employment, or received unemployment benefits.⁸ The IEB contains spell-level information on each individual's jobs, including precise start- and end-dates, occupation, region, industry, age, schooling. Gross wages, measured in Euro per day, are top-coded at the contribution ceiling for pension insurance. We employ the procedure introduced by Card et al. (2013) to impute those censored wages. From this data, we identify all individuals who were ever employed in one of the plants surveyed in the 2019 wave of the Establishment Panel. The resulting dataset allows us to examine the complete employment biographies (employment, wages, occupation, region, industry) and background characteristics (age, schooling, gender) of all workers who have been exposed to robot adoption even if they separate from the plant in subsequent years.

We use the occupation code of the current job to classify occupations according to the popular classification by Blossfeld (1987), which permits us to distinguish blue-collar production jobs from others. Within these blue-collar production jobs (Blossfeld occupation 2-5), we are further able to identify workers with a high share of routine-manual tasks according to the classification from Spitz-Oener (2006) based on the 1991 BIBB/IAB Employment Survey. As argued in Acemoglu and Restrepo (2018) and Acemoglu et al. (2023), these workers, who are classified as being in "simple manual blue-collar occupations" (Blossfeld occupation 2) are most prone to being directly affected by automation, as their job contains a high share of potentially replaceable routine tasks.⁹ We also use the occupation codes to quantify each establishment's occupational employment structure in order to merge a novel measure of plant-specific labor market tightness. This measure was provided by Bossler and Popp (2024) and is calculated as the ratio of the number of job seekers to the number of vacancies, both taken from official statistics. Since not all vacancies are registered with the employment agency, Bossler and Popp (2024) use plant-level survey data to correct for varying penetration rates by skill levels. The result is a measure of labor market tightness that varies by both the local labor market and detailed 5-digit occupation.

⁸Access to this data is regulated by Section 75 of the German Social Code (Book X).

⁹Compared to the overall mean share of around 51% routine tasks across all blue-collar production workers in our main sample (Blossfeld occupation 2-5), routine-manual production workers (Blossfeld occupation 2) have a 14 percentage points higher share of such tasks (see Table 2 in Section 3 for details).

3 Event-Study Models with Double and Triple Differences

Our empirical strategy examines the impact of firms' automation events on directly affected incumbent workers and investigates whether works councils can moderate these consequences. To this end, we borrow from the current literature on worker-level effects of job displacement due to mass layoffs (e.g. Lachowska et al., 2020; Bertheau et al., 2023). Our approach is motivated by recent papers like Schmieder et al. (2023) and Illing et al. (2024), which use propensity score matching to identify a control group of comparable never-treated workers in comparable plants prior to running an event-study analysis. This has the advantage that we obtain pairs consisting of a worker in a robot-adopting plant and a matched control worker in a never-adopting plant, who are both assigned a common event date. Schmieder et al. (2023) point out that this avoids the problems of two-way fixed effects models when treatment timing varies (as expounded by Goodman-Bacon, 2021).

To ensure comparability between workers in robot-adopting and non-adopting plants, we perform a 1-nearest neighbor propensity score matching with a caliper (Stuart and Rubin, 2008), based both on worker and plant characteristics prior to adoption. These characteristics are (log) daily wage, job experience, plant size, and pre-estimated AKM plant fixed effects (Abowd et al., 1999; Bellmann et al., 2020). Additionally, we force workers from matching pairs to have the same gender, nationality, contract status (full-time, part-time), and are in the same occupation of in-total 4 groups, according to Blossfeld (1987). Moreover, matched workers are from establishments in the same aggregate industry and with the same work council status. In this way, we can account for the sorting of workers into firms regarding firm-level institutions, such as co-determination. Table 1 contains an overview and additional details for all variables, separately by type of matching for which they are used (propensity score and exact matching).

To rule out that our results are influenced by workers who have been hired endogenously in the course of implementing the new robot technology (e.g. experts on robot use or maintenance), we restrict our sample to all incumbent workers employed at the same plant at least two years prior to robot adoption. Additionally, we restrict our sample to blue-collar production workers (as defined in Blossfeld (1987)), as those are most likely directly affected by robotization, and who are aged between 25 and 60 in the year of robot adoption. This leaves us with yearly observations of 17,721 individuals in 718 plants.

Table 1: Matching variables

Matching type	Variable	Additional description
Propensity score matching	(Log) Daily wage	Gross wages; Censored top-coded wages above the contribution ceiling for the pension insurance are imputed following Card et al. (2013).
	Job experience	In years
	Plant size	Number of regular employees
	AKM plant fixed effects	Pre-estimated in period 2003-2010 by Bellmann et al. (2020)
Exact matching	Sex	Male, female
	Nationality	German, non-German
	Contract type	Full-time, part-time
	Blossfeld occupation	simple manual (2), qualified manual (3), technicians (4), or engineers (5) according to Blossfeld (1987)
	Aggregate Industry	43 distinct industries (13 manufacturing industries)
	Missing AKM plant fixed effect	Missing, non-missing
	Works council status	Works council, no works council

Notes: This table contains variables used for the matching approach described in Section 3. The matching type refers to whether the variable is used to calculate the propensity score for the 1-nearest neighbor matching or the subsequent restriction of matches having identical characteristics, for instance, having the same gender. If not stated otherwise, variables are measured in the year prior to the event of robot adoption.

Table 2 shows descriptive statistics among matched workers – for all production workers, as well as separately for the subset of older and routine-manual workers. Reassuringly, the table shows that across groups, workers in the robot-adopting and non-adopting plants have comparable demographics, education profiles, and perform similar tasks. This leads us to conclude that the matching procedure accounts for a wide range of observable worker-level characteristics that could potentially bias the results. As discussed in Section 2 and in line with previous findings, adopting firms tend to be larger, which we account for in our estimation strategy described below.

In the second step, we quantify the effect of an automation event on incumbent workers. As a starting point, consider a difference-in-differences (DiD) design of the following form:

$$Y_{it}^g = \alpha^g + \sum_{\tau=-4;\tau \neq -1}^3 \beta_{\tau}^g \times I_{\tau} \times R_{j(i)} + X'_{jt} \phi^g + \eta_{\tau}^g + \eta_t^g + \eta_i^g + u_{it}^g \quad (1)$$

Table 2: Summary statistics of matched workers

	Production workers		Older workers		Routine-manual workers	
	Adopters	Non-Adopters	Adopters	Non-Adopters	Adopters	Non-Adopters
Age	43.79 (9.87)	45.02 (10.00)	57.29 (1.67)	57.40 (1.71)	44.20 (9.85)	45.83 (9.66)
No degree	0.05 (0.22)	0.06 (0.24)	0.06 (0.24)	0.08 (0.27)	0.09 (0.28)	0.11 (0.31)
Vocational degree	0.86 (0.35)	0.85 (0.36)	0.86 (0.35)	0.83 (0.37)	0.89 (0.31)	0.87 (0.34)
College degree	0.09 (0.28)	0.08 (0.27)	0.08 (0.27)	0.08 (0.28)	0.02 (0.14)	0.02 (0.15)
Simple manual	0.46 (0.50)	0.46 (0.50)	0.49 (0.50)	0.50 (0.50)	1.00 (0.00)	1.00 (0.00)
Qualified manual	0.35 (0.48)	0.36 (0.48)	0.33 (0.47)	0.33 (0.47)	0.00 (0.00)	0.00 (0.00)
High-skilled manual	0.19 (0.39)	0.18 (0.38)	0.17 (0.38)	0.17 (0.37)	0.00 (0.00)	0.00 (0.00)
Share routine tasks	0.53 (0.22)	0.50 (0.19)	0.56 (0.24)	0.51 (0.19)	0.67 (0.16)	0.62 (0.12)
Share routine manual tasks	0.49 (0.26)	0.45 (0.23)	0.52 (0.27)	0.46 (0.23)	0.64 (0.17)	0.59 (0.14)
Tenure	13.40 (8.34)	15.07 (8.65)	17.87 (9.26)	19.29 (8.93)	12.87 (8.03)	15.17 (8.47)
Log daily wage	4.68 (0.43)	4.70 (0.45)	4.64 (0.47)	4.69 (0.48)	4.55 (0.35)	4.58 (0.38)
Number of employees	451.69 (367.42)	361.48 (401.43)	424.42 (354.72)	329.41 (353.53)	519.21 (405.87)	409.84 (444.44)
Firm-level share of production workers	0.76 (0.11)	0.75 (0.14)	0.77 (0.11)	0.76 (0.14)	0.77 (0.09)	0.76 (0.12)
Observations	8823	8898	1409	1788	4076	4124

Notes: This table displays means and standard deviations for matched workers in robot adopting and non-robot adopting plants in the year prior to adoption. The first two columns refer to the entire sample of production workers aged 25 to 60 in the year of adoption and being employed at least two years prior to robot adoption in the plant. The middle (last) two columns refer to the sub-sample of workers aged 55 and older (routine-manual workers). The number of observations differs slightly between the treatment and control group due to the post-matching restrictions on firm attachment. Production workers are defined as being in Blossfeld occupation 2-5. Routine-manual workers are defined as being in Blossfeld occupation 2. Simple, qualified, and high-skilled manual refer to Blossfeld occupations 2, 3, and 4+5 respectively. The shares of routine and routine-manual tasks are calculated following Spitz-Oener (2006) based on the 1991 BIBB/IAB Employment Survey.

for individual i in calendar year t , period τ , and plant j . As outcome variables, we use an indicator variable that equals 1 if a worker is employed at least one day per calendar year (either at the initial plant or anywhere) and mean log daily wage.

To assess the differential consequences of the automation event for workers in plants

with ($g = WC$) and without works councils ($g = NWC$), we run separate regressions for both groups g . Although this already allows us to visually compare the effect of automation across both groups, the visual representation would suffer when conducting further sub-group analyses, such as distinguishing between age groups.

Thus, to directly show and quantify the differences between plants with and without works councils, we employ a triple difference (DiDiD) design (as discussed by Olden and Møen, 2022). This has the advantage that we obtain point estimates and confidence intervals for the differential effects of the automation event at each point in time before and after the event.

This estimation equation takes the form:

$$Y_{it} = \alpha + \sum_{\tau=-4;\tau \neq -1}^3 \delta_{\tau} \times I_{\tau} \times R_{j(i)} \times G_{j(i)} + X'_{jgt} \xi + \eta_{\tau} + \eta_t + \eta_i + \epsilon_{it} \quad (2)$$

where $G_{j(i)}$ indicates whether a plant has a works council. The observation period spans from four years before to three years after robot adoption, i.e., $\tau \in \{-4, 3\}$. R_j indicates robot adoption at $\tau = 0$, and I_{τ} denotes the relative time to automation. To account for differences between workers and plants, we include period η_{τ} , calendar year η_t , and individual fixed effects η_i , plus (age-45) squared as controls (X_{jt}). In the DiDiD design, X_{jgt} also includes all lower-order interaction terms between I_{τ} , $R_{j(i)}$, and $G_{j(i)}$. δ_{τ} shows the causal effect of automation in works council versus non-works council plants under the assumptions of (i) parallel trends and (ii) no anticipation. (i) requires the trend between plants with and without works councils to evolve similarly with and without robot adoption. Although not directly testable, we assess this by visually inspecting if our pre-trend coefficients are different from zero. By restricting our sample to incumbent workers with at least two years of tenure, we ensure our estimates do not suffer from bias due to anticipation and selection into treatment. Further, by restricting matches within works council groups, we account for the selection into both robot adoption and works council plants. We follow Abadie and Spiess (2022) and cluster standard errors at the level of matched worker pairs.

Remaining Threats to Identification. Our matched event study design ensures that we compare treated individuals with very similar never-treated individuals in terms of individual and plant characteristics. Inspecting pre-trends reveals that their careers evolved in a similar way before the event. However, this does not necessarily rule out that robot

adoption might be accompanied by other events that affect potential worker outcomes like investments, changes in management strategies, or demand shocks. While this may complicate the causal interpretation of robot adoption, our main interest lies in the moderating role of worker representation. This requires the somewhat weaker assumption that the accompanying events of robot adoption are similar between plants with or without a works council. Another concern is that firms with works councils differ in many ways from firms without works councils and that those differences also moderate the effects of automation. On average, plants with works councils are larger, more productive, and pay higher salaries than plants without (Mohrenweiser, 2022). The most important determinant of having a works council is size. Since many large plants had adopted robots already before 2015 and since we match workers based on the size of their plant, our sample consists of many medium-sized plants. Additionally, our matching approach ensures that we are not systematically comparing small and large firms. Still in a later section, we conduct a robustness check in which we account for a possibly confounding influence of firms’ size over time – which yields results that are very similar to the main results. In a similar vein, we control for contemporaneous investments in ICT and real estate, where the latter serves as a proxy for a general production expansion.

4 Retention and Employment Effects

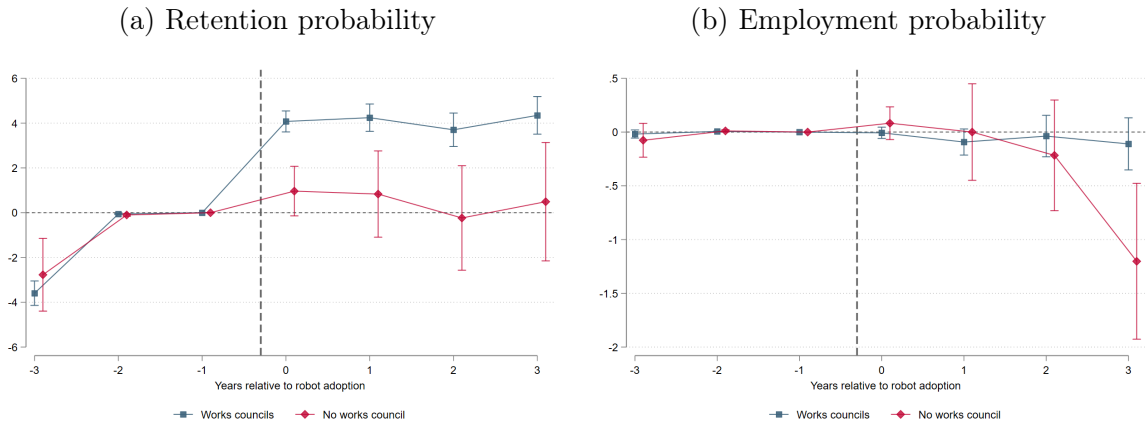
4.1 Main Effects

We begin by showing how automation – induced by the adoption of robots at the plant level – affects incumbent workers’ turnover and employment. Figure 1 displays the effect of robot adoption separately estimated for matched workers in plants with (blue) and without (red) works councils using the difference-in-differences specification in Equation 1.

Panel (a) shows that robot adoption has an, on average, positive effect on retention – but only in plants with a works council. For all types of plants, separation rates increase over time: after 3 years, on average 12% of the blue-collar workforce have left their initial plant. However, workers in automating plants with a works council are around 4 percentage points more likely to remain in their initial establishment relative to their matched counterparts in non-automating plants with a works council.

Previous studies, such as Acemoglu et al. (2023) and Dixon et al. (2021), find that robot-adopting firms expand employment. Our results complement these findings by

Figure 1: Employment effect of robotization by works council status



Notes: This figure shows the effect of robot adoption in plants on workers’ employment, either at the initial plant (in Panel (a)) or anywhere (in Panel (b)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences (DiD) estimates obtained from Equation 1, separately for workers in plants with and without works council. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

showing that, in Germany, incumbent workers benefit from this increase in employment prospects – but only when worker representation is present. These results are also consistent with Dauth et al. (2021), who find that industry-level robot exposure increases worker retention on average and more strongly in more unionized labor markets. Our plant-level results indicate that also works councils facilitate sharing the gains from automation by enhancing job security for incumbent workers. In the task-based automation framework developed by Acemoglu and Restrepo (2018), changes in labor demand can be decomposed into a negative displacement effect and a positive productivity effect stemming from cost reduction generated by automation. Worker representation appears to amplify the productivity effect, which offsets displacement and reduces net separations among incumbent workers. In line with Budde et al. (2024), we find that shop-floor worker representation enhances job security rather than wages. This reflects institutional features that grant works councils significant authority over restructuring and layoffs but limited influence on wage-setting.

In Panel (b), we investigate the consequences for total employment across all firms for workers exposed to automation. Starting with establishments without worker representation, it is visible that the employment prospects of workers exposed to firm-level robotization gradually decrease up to around 1 percentage point after 3 years, suggest-

ing that those workers separating from automating establishments are displaced by the robotization events. In contrast, with worker representation, exposure to firm-level automation is estimated to have no total employment effect for incumbent workers. For them, increased job stability cancels out separation effects from displacement.

4.2 Heterogeneity by Worker Characteristics

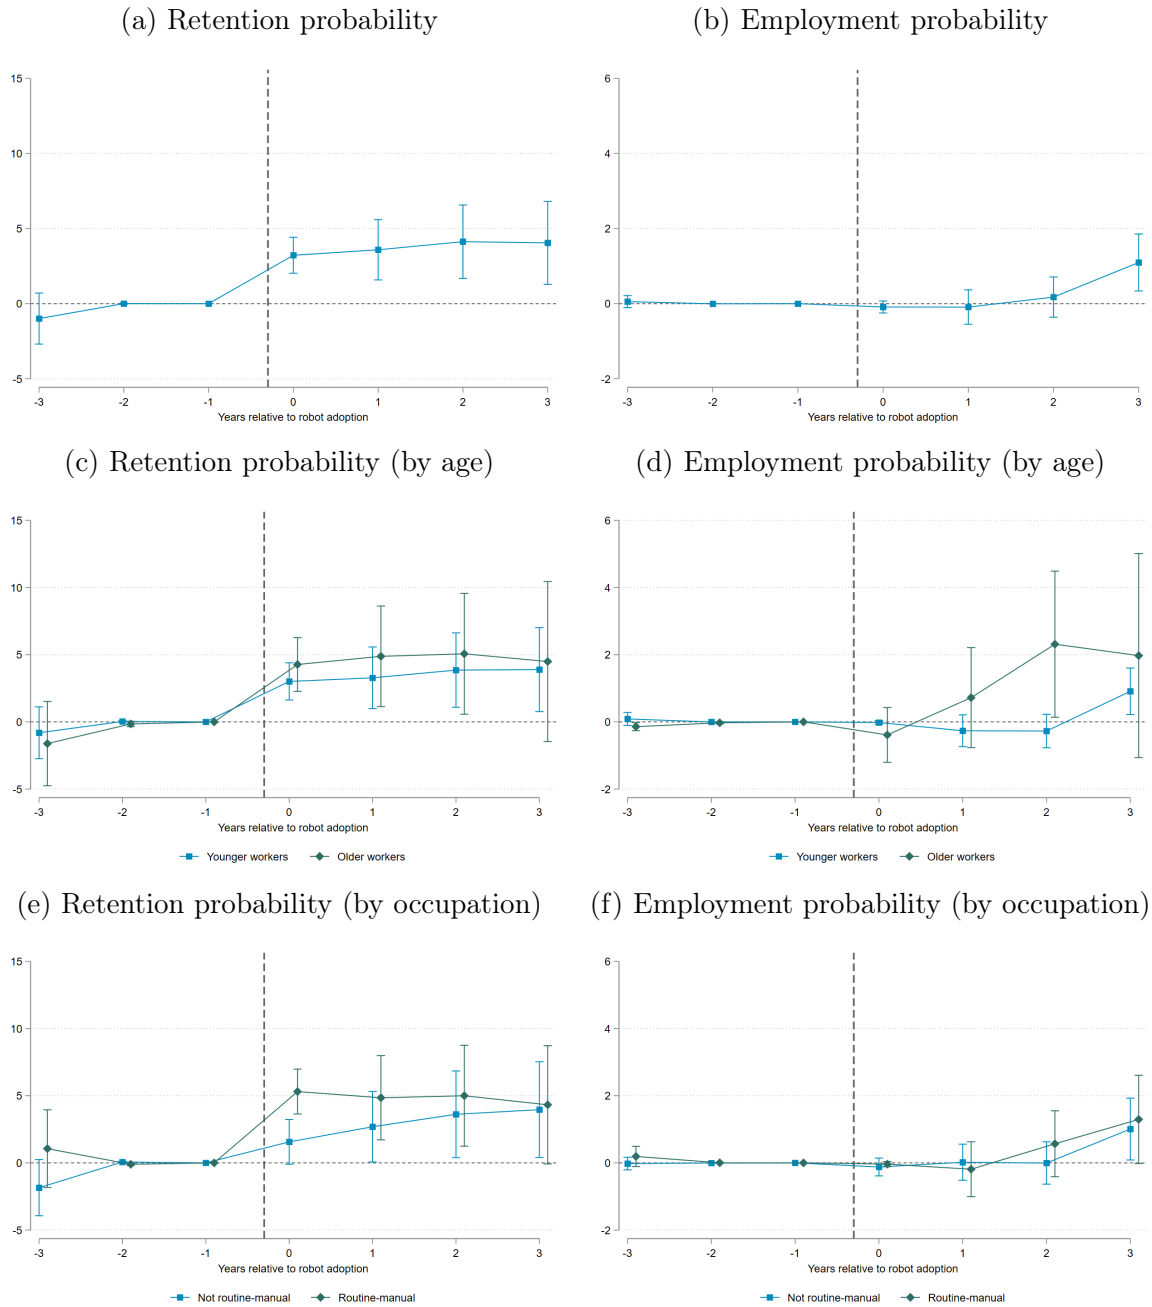
We continue by examining heterogeneity with respect to worker type. We begin by focusing on older workers, who typically adapt less easily to new technologies (Behaghel et al., 2014). Employers may prefer to substitute these workers with younger ones if new tasks favor younger workers (Battisti et al., 2023). Consequently, works councils may prioritize efforts to support older workers, who also face greater barriers to finding new employment (Aubert et al., 2006).

To study heterogeneity and present the results in a comprehensible way, we make use of the triple-difference (DiDiD) design specified in Equation (2). To benchmark the results and aid interpretation, Panels (a) and (b) in Figure 2 replicate the findings from Figure 1, showing the differing impacts of automation in firms with and without worker representation. Each line in Panels (c) to (f) reflects the results of a separately regressed DiDiD specification from Equation (2).

To see whether older employees benefit more from works councils, we divide the sample of production workers into employees under 55 years and those 55 and older. Panel (c) in Figure 2 shows that both age groups are similarly likely to remain in their initial plant, indicating that the efforts of works councils to reduce separations have no direct age bias. However, unlike young workers, who have low adjustment costs once displaced, works councils increase the probability of being employed for older workers by 2 percentage points in the long run. These effects are economically significant given that older workers are, on average, non-employed with a probability of around 2.5% three years after the adoption of robots.¹⁰

¹⁰These effects are not driven by early retirement as we only consider spells of individuals subject to social security contributions.

Figure 2: Employment effect of works council by workers' characteristics



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the probability of being employed at least one day per calendar year, either at the initial plant (in Panel (a), (c), and (e)) or anywhere (in Panel (b), (d), and (f)). All sub-figures display the triple differences (DiDiD) estimates obtained from Equation 2, either for the whole sample or separately estimated across worker groups. In Panel (c) and (d), workers are divided into groups based on whether they are below 55 or between 55 and 60 in the year of robot adoption. In Panel (e) and (f), the division is based on workers' occupation in the year prior to adoption. Routine manual workers (RMW) are defined as being in Blossfeld occupation 2. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3. Results from a difference-in-differences (DiD) estimation showing the effect of robot adoption separately by plant with and without works council can be found in Figures 1, as well as C.2, and C.3 in the Appendix.

In the Appendix in Figure C.2, we show the corresponding DiD results of this analysis. This decomposition reveals that the positive retention effect for older workers is partly driven by increased separations in automating establishments without worker representation, implying direct protection from layoffs by works councils for this group (Panel (b)). In addition, Panel (d) in Figure C.2 highlights the vulnerability of older workers to automation events. In the absence of works councils, their probability of being employed reduces by around 3 percentage points 3 years after robot adoption. In other words, older workers are the primary beneficiaries of works councils’ efforts to reduce separations.

Another group of workers who face greater challenges recovering from involuntary job loss are those in routine-intensive occupations (Blien et al., 2021; Cortes, 2016). Panels (e) and (f) in Figure 2 divide the sample of production workers into routine-intensive and non-routine occupations. Consistent with our previous results regarding age and the findings of Budde et al. (2024), who show that worker representatives benefit workers of all types¹¹, we find no inherent bias in the effect of works councils on job protection. Instead, we observe an equal increase in retention for both routine-intensive and non-routine workers. The corresponding DiD results in Figure C.3 further support this conclusion.

4.3 Labor Market Tightness


Next, we examine how labor market tightness from the employers’ perspective influences separations and the role of worker power. In models incorporating labor market tightness, the value of retaining workers to the firm increases as replacement and recruitment costs rise (Kline et al., 2019; Jäger and Heining, 2022). Consequently, labor market tightness is expected to enhance the alignment of goals between management and worker representation. Building on this reasoning, we test the hypothesis that works councils have a reduced retention effect in tight labor markets where labor is scarce. Put differently, the intuition is that works councils play a smaller role for incumbent workers when firms already have strong incentives to minimize separations.

To measure labor market tightness, we use a fine-grained, plant-specific measure developed by Bossler and Popp (2024). This measure is calculated as the ratio of vacancies to job seekers for each 5-digit occupation and local labor market, weighted by the plant’s employment structure. Using this measure, we categorize establishments based on their local labor market tightness – a proxy for replacement and hiring costs – in the year prior


¹¹Budde et al. (2024) estimate the effects of blue-collar worker representatives and find homogeneous retention effects across groups.

Figure 3: Employment effect of works council by labor market tightness

(a) Retention probability (25th percentile) (b) Employment probability (25th percentile)



(c) Retention probability (50th percentile) (d) Employment probability (50th percentile)



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the probability of being employed at least one day per calendar year, either at the initial plant (in Panel (a) and (c)) or anywhere (in Panel (b) and (d)). All sub-figures display the triple differences (DiDiD) estimates obtained from Equation 2, either for the whole sample or separately estimated across worker groups. Workers are divided into groups based on plants' local labor market tightness (LMT) in the year prior to adoption (cutoff in the 25th percentile in Panel (a) and (b), 50th percentile in Panel (c) and (d)). This measure is obtained from Bossler and Popp (2024) and defined as the ratio of the number of vacancies to job seekers at the occupation-region level, weighted by plants' employment shares. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3. Results from a difference-in-differences (DiD) estimation showing the effect of robot adoption separately by plant with and without works council can be found in Figures C.4 and C.5 in the Appendix.

to robot adoption. Plants are classified as operating in a slack labor market (labeled 'low LMT' in the upcoming figures) if their plant-specific labor market tightness is above the 75th percentile when ranked by how slack market conditions are. For robustness, we also perform the classification using the median as a cutoff.

Panel (a) in Figure 3 reveals that works councils increase worker retention during automation events only when labor markets are slack, that is when replacement costs for the plant are low. Thus, intuitively, when labor is abundant and firms face stronger

incentives to replace workers with robots, worker representation has the largest impact on job security. As argued in the previous section, the positive retention coefficients in slack labor markets suggest that worker representation enables productivity gains to dominate displacement effects, manifesting as enhanced job security for incumbent workers.¹²

Next, we investigate monotonicity using the median as a cutoff instead of the 75th percentile of plants’ labor market tightness. Panel (c) in Figure 3 and Figure C.5 in the Appendix demonstrate that while the moderating effect of works councils remains visible, it is less pronounced. This finding suggests a monotonic relationship, where the moderating effect of works councils is strongest in the slackest labor markets – consistent with the argumentation that replacement costs increase firms’ efforts to reduce turnover and thus align the incentives of employers and their workforce.

4.4 Sensitivity and Robustness

To ensure that our results are not driven by the choice of the outcome variable or confounding effects of robot adoption or the presence of works councils, we conduct several robustness checks which are presented in the Appendix and briefly discussed below. Our main dependent variable so far was the employment status of a person, defined as being employed for at least one day per calendar year, either at the initial plant or anywhere. We complement this by also taking into account the intensive margin and using the number of days employed in any year in Appendix Figure D.6 and Figure D.7.¹³ The baseline positive effect on retention probability of 4 percentage points corresponds to around 15 working days. As before, older workers seem to be the main beneficiaries of works councils during automation events. While the retention effect is again only slightly higher for older production workers (Panel (c)), they have stronger employment effects compared to their younger counterparts when using days in employment as an outcome variable (Panel (d)).

As mentioned before, the probability of plants having a works council and adopting robots strongly increases with size (Mohrenweiser, 2022; Deng et al., 2024). To rule out the possibility that this influences our results, we account for differential trends by establishment size. More specifically, we address this by including a dummy variable for above- versus below-median-sized plants, fully interacted with event time indicators,

¹²The decomposition in Appendix Figure C.4 in Panel (d) shows that the total employment effects of automation are similar across establishments with and without works councils in tight labor markets, consistent with workers having better outside options in tight markets regardless of works council protection.

¹³Non-employment is defined as zero days employed.

in our main specification. This takes into account the potentially confounding effect of being employed in a large firm over the course of robot adoption. Panel (b) in Appendix Figures D.8 and D.9 demonstrates that the results remain consistent with the baseline.

Furthermore, we leverage additional information from the establishment survey to account for events that are likely to accompany robot adoption events and thus potentially confound our results by influencing worker turnover. First, we test robustness by including a dummy variable for accompanying investments in real estate (Panel (c)) – a proxy for firm expansions – and investments in information and communication technologies (ICT) (Panel (d)). Again, both dummy variables are fully interacted with event time indicators. Panels (c) and (d) in Figure D.8 show that these adjustments lead the positive retention effects to rather build up over time. Intuitively, both types of investment are positively correlated with the timing of robot adoption, impacting the estimated coefficients downwards. Reassuringly, after this initial phase, the positive retention effects become larger, but also more imprecise.

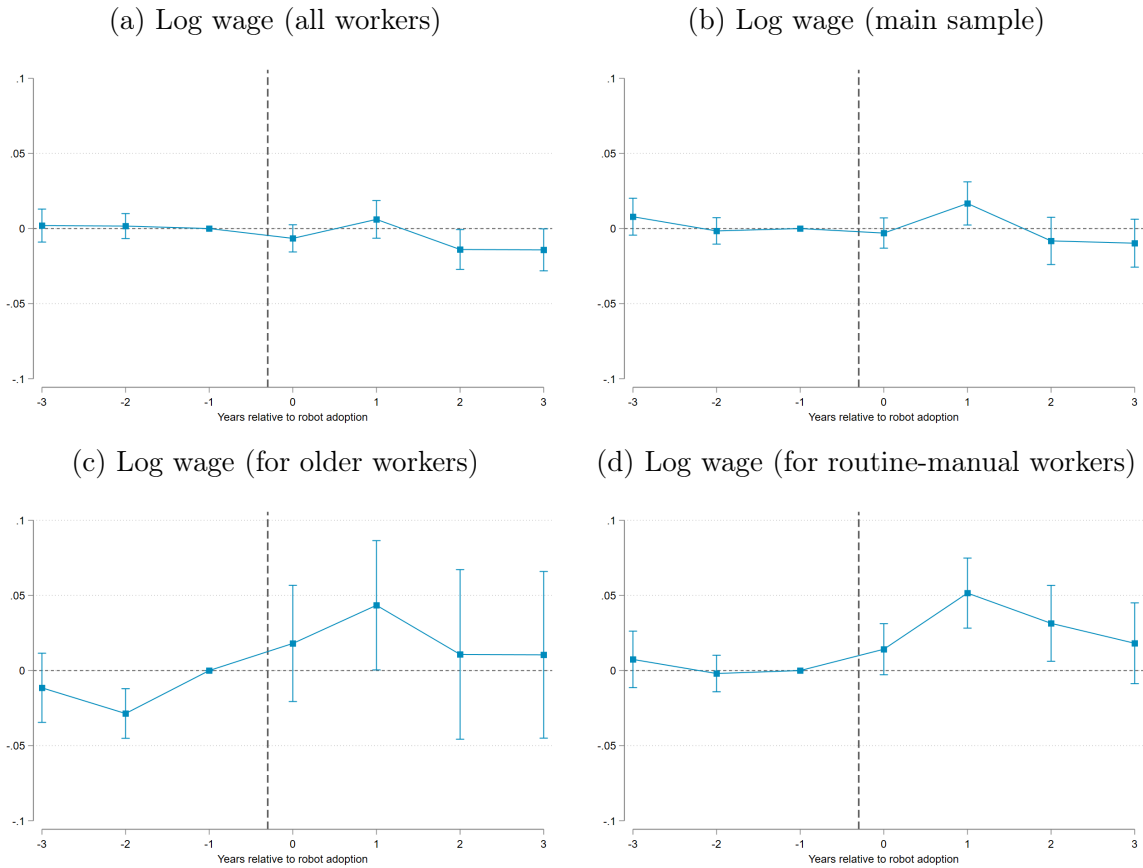
5 Wage Effects

Works councils lack direct wage bargaining rights. However, besides job security – analyzed in the previous section – works councils can influence wages indirectly through negotiations over classifications of workers into pay grades (Mohrenweiser, 2022) and, particularly, prevent workers’ demotion to lower pay grades. In the context of automation, we therefore test the hypothesis of whether works councils protect wages and prevent cuts, especially for vulnerable subgroups. This analysis complements the rent-sharing literature surveyed in Card et al. (2018), which typically examines positive firm shocks like innovation events that are expected to increase (or at least maintain) workers’ wages.

We estimate the same set of event studies as in the previous section – and described in Equations (1) and (2) in Section 3 – using log wages as the outcome variable. To be able to capture the evolution of wages within establishments, we have to concentrate on observations of workers who are employed at their initial plant. Thus, our results should be interpreted as the wage effect of works councils during automation for those workers who remain at their initial establishment, at least for some time.

First, Figure 4 shows the results from the triple difference specification (Equation (2)) for the entire workforce (Panel (a)) as well as for all production workers (Panel (b)). For both groups, we find no wage effects of works councils during automation. This

Figure 4: Wage effect of works councils by workers' characteristics



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption on workers' log daily wage in the initial establishment. Panel (a) to (d) display the triple differences (DiDiD) estimates obtained from Equation 2, separately estimated across worker groups. For Panel (a), the sample of workers is restricted to individuals aged 25 to 60 in the year of adoption and being employed at least two years prior to robot adoption in the plant. In all other panels, the sample is restricted to production workers (i.e. Blossfeld occupations 2-5). For Panel (c), the sample is additionally restricted to workers aged 55 and above in the year of robot adoption. For Panel (d), the restriction is based on whether workers' occupation are characterized to be routine-manual (RMW, referring to Blossfeld occupation 2) in the year prior to the event. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

is in line with the literature, which, although not differentiating by firm-level policies, has found only limited effects of automation on wages at adopting firms for all workers on average (e.g., Koch et al., 2021; Dixon et al., 2021). Next, we investigate the wage effects separately for older production workers (Panel (c)) and workers in routine-manual occupations (Panel (d)), as in the previous section. For both groups, there is evidence that works councils increase wages during automation events, although the estimates suffer from imprecision for older workers. For workers with routine-manual task profiles, this effect is both economically and statistically significant, with a relative increase in wages of around 3.5%.

Figure 5: Wage effect of automation in plants with vs. without works council – by age and occupation



Notes: This figure shows the effect of robot adoption in plants on workers' log daily wage in the initial establishment. All sub-figures display the difference-in-differences estimates obtained from Equation 1, separately for workers in plants with and without works council. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). For Panel (a) and (b), workers are divided into groups based on whether they are below 55 or between 55 and 60 in the year of robot adoption. For Panel (c) and (d), this division is based on whether workers' occupation is characterized to be routine-manual (RMW, referring to Blossfeld occupation 2) in the year prior to the event. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

To see if these positive wage effects of works councils for older and routine-manual workers are driven by excess positive wage growth or rather an avoidance of wage cuts, we estimate the difference-in-differences specification as in Equation 1 separately for each of the two vulnerable subgroups. Panel (b) and (d) in Figure 5 clarify that the positive wage effect of works councils is driven by the prevention of wage cuts. Both older and routine-manual workers face a significant reduction of wages in plants without works councils, while wages for both groups remain relatively constant in plants with works councils. In contrast, wages of younger workers (Panel (a)) and workers in non-routine-manual occupations (Panel (c)) are not affected by robot adoption overall, neither in plants with

or without works councils. This leads us to conclude that works councils play a decisive role in the earnings prospects of vulnerable subgroups, not only through increased job security – as highlighted in Section 4 – but also by preventing wage cuts.

6 Mechanisms: Technology Direction and Training

In this section, we investigate the mechanisms through which organized labor interacts with automation decisions at the plant level. We do so by comparing features of the plants regarding the direction of technology adoption, productivity, and training of similar first-time robot-adopting plants that differ in their works council status. While this comparison of plants by works council status does not identify causal effects it reveals meaningful descriptive patterns.

Data and Estimation. In addition to the information on the presence of works councils and the number of robots used between 2014 and 2018, the IAB Establishment Panel contains variables that capture the direction of technology adoption, as well as value-added measures (as a proxy for productivity) and worker training. For the year 2018, we have information on the robot density and the type of robot installed. Plants were asked about the number of robots (i) with a price below 50,000 Euro, which we refer to as ‘cheap robots,’ and (ii) that are separated from the workforce with a fence, which we call ‘cage robots.’ Cage robots are large, versatile, and highly productive – but need to be separated from the workforce to prevent hazards (Taesi et al., 2023). Cheap robots, by contrast, are more likely to be collaborative robots, or ‘cobots,’ which demand a high degree of human-machine interaction (Gerbert et al., 2015; Plümpe and Stegmaier, 2023) and are constructed with a focus on human safety (Gurgul, 2018). By linking administrative data, we observe changes in employment, productivity, provision of training, and the skill structure of the plant over time and use information on size, industry, and organization as control variables.

For every first-time robot-adopting plant j in industry i we estimate the following regression across different outcomes:

$$Y_j = \beta_{WC} WC_j + I_{i(j)} + X_j + e_j, \quad (3)$$

The estimand of interest is β_{WC} and captures the difference in features of the automation process for similar first-time robot adopters depending on their works council

status. $I_{i(j)}$ are industry fixed effects, which ensure that β_{WC} is identified only from comparing plants within the same aggregate industry, controlling for potentially confounding industry-specific trends. With X_j we further control for plant size (10 groups), the share of high-skilled workers, and the year of foundation. For outcomes that we observe repeatedly over time, we run this regression in pooled cross-sections around the year of robot adoption, which are either prior to adoption ($\tau < -1$), during adoption ($\tau = -1, 0$), or after adoption ($\tau > 0$).¹⁴

Results. We start by relating the presence of co-determination to the type of automation technology adoption. Panel A in Table 3 contains the estimates for β_{WC} from the cross-sectional regression. Column 1 shows, that there is no difference in robot density, measured by the number of robots per worker in 2018. In columns 2 and 3, we distinguish different types of robots. Plants with works councils have a 17 percentage points higher share of cage robots, which are commonly associated with higher productivity (Gurgul, 2018).

Consistent with a mechanism in which automating firms with works councils employ not more but higher quality robots, they appear to have fewer cheap (collaborative) robots installed, although this result is not statistically significant. In a standard model of automation decisions, as the one by Acemoglu and Restrepo (2018), the profitability of investments is the main concern of firms. Internalizing part of the displacement costs and weighting the welfare of incumbent workers could create a wedge into the decision, shifting up the threshold for automation investments with a positive return. Thus, conditional on robot adoption, plants with works councils should have higher productivity gains.

We test this in Panel B, where we study the log of value added per worker as a measure of labor productivity, leveraging also the time dimension of the data. Already prior to robot adoption, plants with works councils are more productive, in line with previous findings (Addison et al., 2001; Mueller and Neuschaeffer, 2021), although the difference is statistically insignificant. This difference increases in the aftermath of robot adoption and becomes highly statistically and economically significant. The estimates imply that the difference in productivity increases from around 0.11 log points to more than 0.26, or 30%. Overall, these findings align with the idea that works councils drive up the requirement for

¹⁴Pooling years ensures a sufficient number of observations as many plants have missings across years. We account for plants having multiple observations per year pool by clustering robust standard errors at the plant level. Additionally, we only include plants that have at least one observation in every year pool. We use up to 4 years before and up to 2 years after adoption.

Table 3: Firm-level mechanisms

<i>Panel A: Equipment</i>	Robots/worker	Share cage	Share cheap
Works council	0.006 (0.008)	16.716** (8.427)	-9.250 (8.887)
Mean of Y	0.08	67.81	37.05
SD of Y	0.17	45.93	46.53
R-squared	0.66	0.40	0.22
Observations	187	187	187
<i>Panel B: Log value added</i>	$\tau_{<-1}$	$\tau_{-1,0}$	$\tau_{>0}$
Works council	0.114 (0.155)	0.151 (0.103)	0.261*** (0.089)
Mean of Y	10.97	10.97	10.93
SD of Y	0.71	0.72	0.69
R-squared	0.37	0.43	0.42
Observations	203	191	171
<i>Panel C: Training</i>	$\tau_{<-1}$	$\tau_{-1,0}$	$\tau_{>0}$
Works council	4.922 (5.906)	14.536*** (5.440)	6.960 (5.539)
Mean of Y	29.16	30.30	26.62
SD of Y	32.19	31.51	32.41
R-squared	0.14	0.15	0.12
Observations	264	266	249

Notes: This table shows results from regressions of various outcome variables on an indicator whether a plant has a works council. Panel A shows the results for robot density (robots per worker), the share of cheap (price below 50,000 Euro) and cage robots (separated through a fence) from the total number of installed robots. All outcome variables in Panel A refer to the year 2018. In Panel B and C, the outcome variable is the log value added per worker and the share of trained workers around robot adoption. Columns $\tau_{<-1}/\tau_{-1,0}/\tau_{>0}$ report results from a pooled regression prior/during/after the event. In each regression, we control for 10 plant size dummies, the share of highly qualified workers, the year of foundation, as well as industry fixed effects. Further, we restrict the sample to first-time adopters that have non-missing values in any of the outcome variables in 2018 (for Panel A) or at least one observation in all year pools $b \in \{-4, -2; -1, 0; 1, 2\}$ (for Panel B and C). Standard errors are robust and clustered at the plant level.

the returns to automation to be sufficiently high to offset the higher costs of displacement. However, we do not find evidence that works councils change the purpose or even hinder the adoption of robots, as we do not observe systematic differences in the type (process vs. product improvement) and the probability of using or newly adopting robots (see

Appendix Table E.2).

In Section 4, we showed how works councils reduce turnover during automation events. Previous papers in the literature, such as Acemoglu et al. (2023), have documented that robot adopting firms tend to grow. An interesting open question, therefore, is how labor representation might affect the total firm employment effects of automation and the relative balance between increased retention and new hires. In the Appendix (Table E.3), we investigate this issue with three main findings: First, there is no difference in the total employment effects among adopters by works council status (Panel A). Second, the share of plants with unfilled vacancies does not differ by works council status over the course of automation (Panel B). Third, automating firms with worker representation engage in significantly fewer new hires around the event of adoption compared to automating firms without representation (Panel C). In combination with our findings regarding retention, these patterns indicate that works councils lead plants to rely more on their existing workforce rather than new hires when adopting new technologies – without hindering overall firm growth.

Finally, we shed light on whether works councils accompany the increased retention during automation events as documented in Section 4 by increasing re-training efforts among their workforce. Once again, we leverage the time dimension, using periods pre- and post-adoption as well as contemporaneous to adoption. Column 1 in Panel C of Table 3 demonstrates that firms with works councils have a 5 percentage point (statistically insignificant) higher share of workers who receive training in a given year, consistent with previous findings on plants with worker representation generally providing more training (Stegmaier, 2012; Mohrenweiser, 2022).¹⁵ Notably, during automation events, the gap in the share of trained workers increases sharply to 14.5 percentage points and returns approximately to previous levels thereafter. When distinguishing between low- and high-skilled workers participating in training (Appendix Table E.4), we find that the spike in training provision in plants with works council is, in absolute terms, similar across worker groups. This increase is economically meaningful given that, pre-adoption and across all plants, only 18% of low-skilled workers receive training. This pattern might not only reflect the increased willingness of firms to supply training, but also a higher propensity of workers to take up firm-specific training, as the increased job security makes investments into firm-specific skills more worthwhile (Freeman and Lazear, 1995). These findings are

¹⁵Relatedly, Dustmann and Schönberg (2012) demonstrate that unions can increase training by compressing the wage distribution.

also related to the work by Battisti et al. (2023), who show that technological and organizational change at the establishment level is, on average, associated with higher retraining efforts.

7 Conclusion

We find that work councils moderate adverse effects from automation events on incumbent workers by reducing separations. Older workers, who have limited adjustment possibilities, benefit the most in terms of employment. When replacing workers is costly for firms – as reflected by high plant-specific labor market tightness – separation and retention effects of automation in plants with and without work councils converge as the objectives of works councils and the management are more closely aligned. We also observe that works councils prevent wage cuts for workers with a high share of routine-manual tasks, which are prone to automation.

Robot adoption is associated with larger productivity growth and increased training efforts in the presence of worker representation. The fact that rising productivity goes hand in hand with retaining and retraining incumbent workers supports the idea that works councils facilitate cooperative solutions when the interests of capital owners and workers diverge (Müller-Jentsch, 1995). Understanding how other labor market institutions might alter the direction and consequences of new (automation) technologies remains an important topic for future research.

References

- Abadie, Alberto and Jann Spiess (2022). “Robust Post-Matching Inference”. In: *Journal of the American Statistical Association* 117(538), pp. 983–995.
- Abowd, John M., Francis Kramarz, and David N. Margolis (1999). “High Wage Workers and High Wage Firms”. In: *Econometrica* 67(2), pp. 251–333.
- Acemoglu, Daron and Simon Johnson (2023). *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*. Hachette UK.
- Acemoglu, Daron and Simon Johnson (2024). *Learning from Ricardo and Thompson: Machinery and Labor in the Early Industrial Revolution, and in the Age of AI*. NBER Working Paper 32416.
- Acemoglu, Daron, Hans RA Koster, and Ceren Ozgen (2023). *Robots and Workers: Evidence from the Netherlands*. NBER Working Paper 31009.
- Acemoglu, Daron and Pascual Restrepo (2018). “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment”. In: *American economic review* 108(6), pp. 1488–1542.
- Acemoglu, Daron and Pascual Restrepo (2020). “Robots and Jobs: Evidence from US Labor Markets”. In: *Journal of Political Economy* 128(6), pp. 2188–2244.
- Acemoglu, Daron and Pascual Restrepo (2022). “Tasks, Automation, and the Rise in U.S. Wage Inequality”. In: *Econometrica* 90(5), pp. 1973–2016.
- Adachi, Daisuke, Daiji Kawaguchi, and Yukiko U. Saito (2024). “Robots and Employment: Evidence from Japan, 1978–2017”. In: *Journal of Labor Economics* 42(2), pp. 591–634.
- Addison, John T. (2009). *The Economics of Codetermination*. Palgrave Macmillan US: New York.
- Addison, John T., Claus Schnabel, and Joachim Wagner (2001). “Works Councils in Germany: Their Effects on Establishment Performance”. In: *Oxford Economic Papers* 53(4), pp. 659–694.
- Aubert, Patrick, Eve Caroli, and Muriel Roger (2006). “New technologies, organisation and age: firm-level evidence”. In: *The Economic Journal* 116(509), F73–F93.
- Autor, David (2024). *Applying AI to Rebuild Middle Class Jobs*. NBER Working Paper 32140.
- Battisti, Michele, Christian Dustmann, and Uta Schönberg (2023). “Technological and organizational change and the careers of workers”. In: *Journal of the European Economic Association* 21(4), pp. 1551–1594.

- Behaghel, Luc, Eve Caroli, and Muriel Roger (2014). “Age-biased technical and organizational change, training and employment prospects of older workers”. In: *Economica* 81(322), pp. 368–389.
- Bellmann, Lisa, Benjamin Lochner, Stefan Seth, and Stefanie Wolter (2020). *AKM Effects for German Labour Market Data*. FDZ-Methodenreport 01/2023.
- Bellmann, Lutz, Laura Brunner, Peter Ellguth, Philipp Grunau, Christian Hohendanner, Susanne Kohaut, Ute Leber, Iris Möller, Barbara Schwengler, Jens Stegmaier, Alexandra Schmucker, Hannes Walz, Stefanie Wolter, and Viktoria Nußbeck (2021). *IAB-Betriebspanel (IAB BP) – Version 9320 v1*. Forschungsdatenzentrum der Bundesagentur für Arbeit (BA) im Institut für Arbeitsmarkt- und Berufsforschung (IAB).
- Beraja, Martin and Nathan Zorzi (2024). “Inefficient Automation”. In: *Review of Economic Studies* forthcoming.
- Bertheau, Antoine, Edoardo Maria Acabbi, Cristina Barceló, Andreas Gulyas, Stefano Lombardi, and Raffaele Saggio (2023). “The Unequal Consequences of Job Loss across Countries”. In: *American Economic Review: Insights* 5(3), pp. 393–408.
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge (2023). “What Happens to Workers at Firms That Automate?” In: *The Review of Economics and Statistics*, pp. 1–45.
- Blien, Uwe, Wolfgang Dauth, and Duncan H.W. Roth (2021). “Occupational Routine Intensity and the Costs of Job Loss: Evidence from Mass Layoffs”. In: *Labour Economics* 68, p. 101953.
- Blossfeld, Hans-Peter (1987). “Labor-Market Entry and the Sexual Segregation of Careers in the Federal Republic of Germany”. In: *American Journal of Sociology* 93(1), pp. 89–118.
- Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia (2022). *Robot Imports and Firm-Level Outcomes*. CEPR Discussion Paper No. 14593.
- Bosler, Mario and Martin Popp (2024). *Labor Demand on a Tight Leash*. Working Paper, arXiv:2203.05593 [econ.GN].
- Budde, Julian, Thomas Dohmen, Simon Jäger, and Simon Trenkle (2024). *Worker Representatives*. NBER Working Paper 32740.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018). “Firms and Labor Market Inequality: Evidence and Some Theory”. In: *Journal of Labor Economics* 36(S1), S13–S70.
- Card, David, Jörg Heining, and Patrick Kline (2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality”. In: *The Quarterly Journal of Economics* 128(3), pp. 967–1015.

- Cortes, Guido Matias (2016). “Where have the middle-wage workers gone? A study of polarization using panel data”. In: *Journal of Labor Economics* 34(1), pp. 63–105.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner (2021). “The Adjustment of Labor Markets to Robots”. In: *Journal of the European Economic Association* 19(6), pp. 3104–3153.
- Deng, Liuchun, Steffen Müller, Verena Plümpe, and Jens Stegmaier (2023). *Robots, Occupations, and Worker Age: A Production-Unit Analysis of Employment*. IZA Discussion Paper No. 16128.
- Deng, Liuchun, Verena Plümpe, and Jens Stegmaier (2024). “Robot Adoption at German Plants”. In: *Jahrbücher für Nationalökonomie und Statistik* 244(3), pp. 201–235.
- Dixon, Jay, Bryan Hong, and Lynn Wu (2021). “The Robot Revolution: Managerial and Employment Consequences for Firms”. In: *Management Science* 67(9), pp. 5586–5605.
- Dobbelaere, Sabien, Boris Hirsch, Steffen Müller, and Georg Neuschäffer (2024). “Organised Labour, Labour Market Imperfections, and Employer Wage Premia”. In: *ILR Review* forthcoming.
- Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener (2014). “From sick man of Europe to economic superstar: Germany’s resurgent economy”. In: *Journal of economic perspectives* 28(1), pp. 167–188.
- Dustmann, Christian and Uta Schönberg (2012). “What makes firm-based vocational training schemes successful? The role of commitment”. In: *American Economic Journal: Applied Economics* 4(2), pp. 36–61.
- Ellguth, Peter and Susanne Kohaut (2015). “Tarifbindung und betriebliche Interessenvertretung. Ergebnisse aus dem IAB-Betriebspanel 2014”. In: *WSI-Mitteilungen* 68(4), pp. 290–297.
- Freeman, Richard B and Edward P Lazear (1995). “An economic analysis of works councils”. In: *Works councils: Consultation, representation, and cooperation in industrial relations*. University of Chicago Press, pp. 27–52.
- Furman, Jason and Robert Seamans (2019). “AI and the Economy”. In: *Innovation Policy and the Economy* 19, pp. 161–191.
- Gerbert, Philipp, Markus Lorenz, Michael Rießmann, Manuela Waldner, Jan Justus, Pascal Engel, and Michael Harnisch (2015). “Industry 4.0: The future of productivity and growth in manufacturing industries”. In: *Boston Consulting Group* 9(1), pp. 54–89.
- Goodman-Bacon, Andrew (2021). “Difference-in-differences with variation in treatment timing”. In: *Journal of Econometrics* 225(2). Themed Issue: Treatment Effect 1, pp. 254–277.

- Gordon, Robert J. (2016). “Perspectives on The Rise and Fall of American Growth”. In: *American Economic Review* 106(5), pp. 72–76.
- Graetz, Georg and Guy Michaels (2018). “Robots at Work”. In: *The Review of Economics and Statistics* 100(5), pp. 753–768.
- Gurgul, Michał (2018). *Industrial robots and cobots: Everything you need to know about your future co-worker*. Michał Gurgul.
- Hirsch, Boris, Thorsten Schank, and Claus Schnabel (2010). “Works Councils and Separations: Voice, Monopoly, and Insurance Effects”. In: *Industrial Relations: A Journal of Economy and Society* 49(4), pp. 566–592.
- Hirvonen, Johannes, Aapo Stenhammar, and Joonas Tuhkuri (2023). *New Evidence on the Effect of Technology on Employment and Skill Demand*. Working Paper.
- Illing, Hannah, Johannes Schmieder, and Simon Trenkle (2024). “The Gender Gap in Earnings Losses after Job Displacement”. In: *Journal of the European Economic Association* forthcoming.
- Jäger, Simon and Jörg Heining (2022). *How substitutable are workers? Evidence from worker deaths*. NBER Working Paper 30629.
- Jäger, Simon, Shakked Noy, and Benjamin Schoefer (2022). “What Does Codetermination Do?” In: *ILR Review* 75(4), pp. 857–890.
- Jäger, Simon, Benjamin Schoefer, and Jörg Heining (2021). “Labor in the Boardroom”. In: *The Quarterly Journal of Economics* 136(2), pp. 669–725.
- Jäger, Simon, Shakked Noy, and Benjamin Schoefer (2022). “The German Model of Industrial Relations: Balancing Flexibility and Collective Action”. In: *Journal of Economic Perspectives* 36(4), 53–80.
- Keynes, John Maynard (1930). “Economic Possibilities for Our Grandchildren”. In: *Essays in Persuasion*. Ed. by John Maynard Keynes. Palgrave Macmillan UK: London, pp. 321–332.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar (2019). “Who Profits from Patents? Rent-Sharing at Innovative Firms”. In: *The Quarterly Journal of Economics* 134(3), pp. 1343–1404.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka (2021). “Robots and Firms”. In: *The Economic Journal* 131(638), pp. 2553–2584.
- Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury (2020). “Sources of Displaced Workers’ Long-Term Earnings Losses”. In: *American Economic Review* 110(10), pp. 3231–3266.
- Mohrenweiser, Jens (2022). *Works Councils*. GLO Discussion Paper 1103.

- Mueller, Steffen and Georg Neuschaeffer (2021). “Worker Participation in Decision-making, Worker Sorting, and Firm Performance”. In: *Industrial Relations: A Journal of Economy and Society* 60(4), pp. 436–478.
- Müller-Jentsch, Walther (1995). “Germany: From collective voice to co-management”. In: *Works councils: Consultation, representation, and cooperation in industrial relations*. University of Chicago Press, pp. 53–78.
- Olden, Andreas and Jarle Møen (2022). “The Triple Difference Estimator”. In: *The Econometrics Journal* 25(3), pp. 531–553.
- Plümpe, Verena and Jens Stegmaier (2023). “Micro Data on Robots from the IAB Establishment Panel”. In: *Jahrbücher für Nationalökonomie und Statistik* 243(3-4), pp. 397–413.
- Schmieder, Johannes F., Till von Wachter, and Jörg Heining (2023). “The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany”. In: *American Economic Review* 113(5), pp. 1208–1254.
- Spitz-Oener, Alexandra (2006). “Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure”. In: *Journal of Labor Economics* 24(2), pp. 235–270.
- Stansbury, Anna and Lawrence H. Summers (2020). *The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy*. NBER Working Paper 27193.
- Stegmaier, Jens (2012). “Effects of Works Councils on Firm-Provided Further Training in Germany”. In: *British Journal of Industrial Relations* 50(4), pp. 667–689.
- Stuart, Elizabeth A. and Donald B. Rubin (2008). “Best Practices in Quasi-Experimental Designs”. In: *Best practices in quantitative methods*, pp. 155–176.
- Taesi, Claudio, Francesco Aggogeri, and Nicola Pellegrini (2023). “COBOT applications—recent advances and challenges”. In: *Robotics* 12(3), p. 79.

Appendix

Organized Labor Versus Robots? Evidence from Micro Data

Sebastian Findeisen, Wolfgang Dauth, Oliver Schlenker

A Plant Statistics

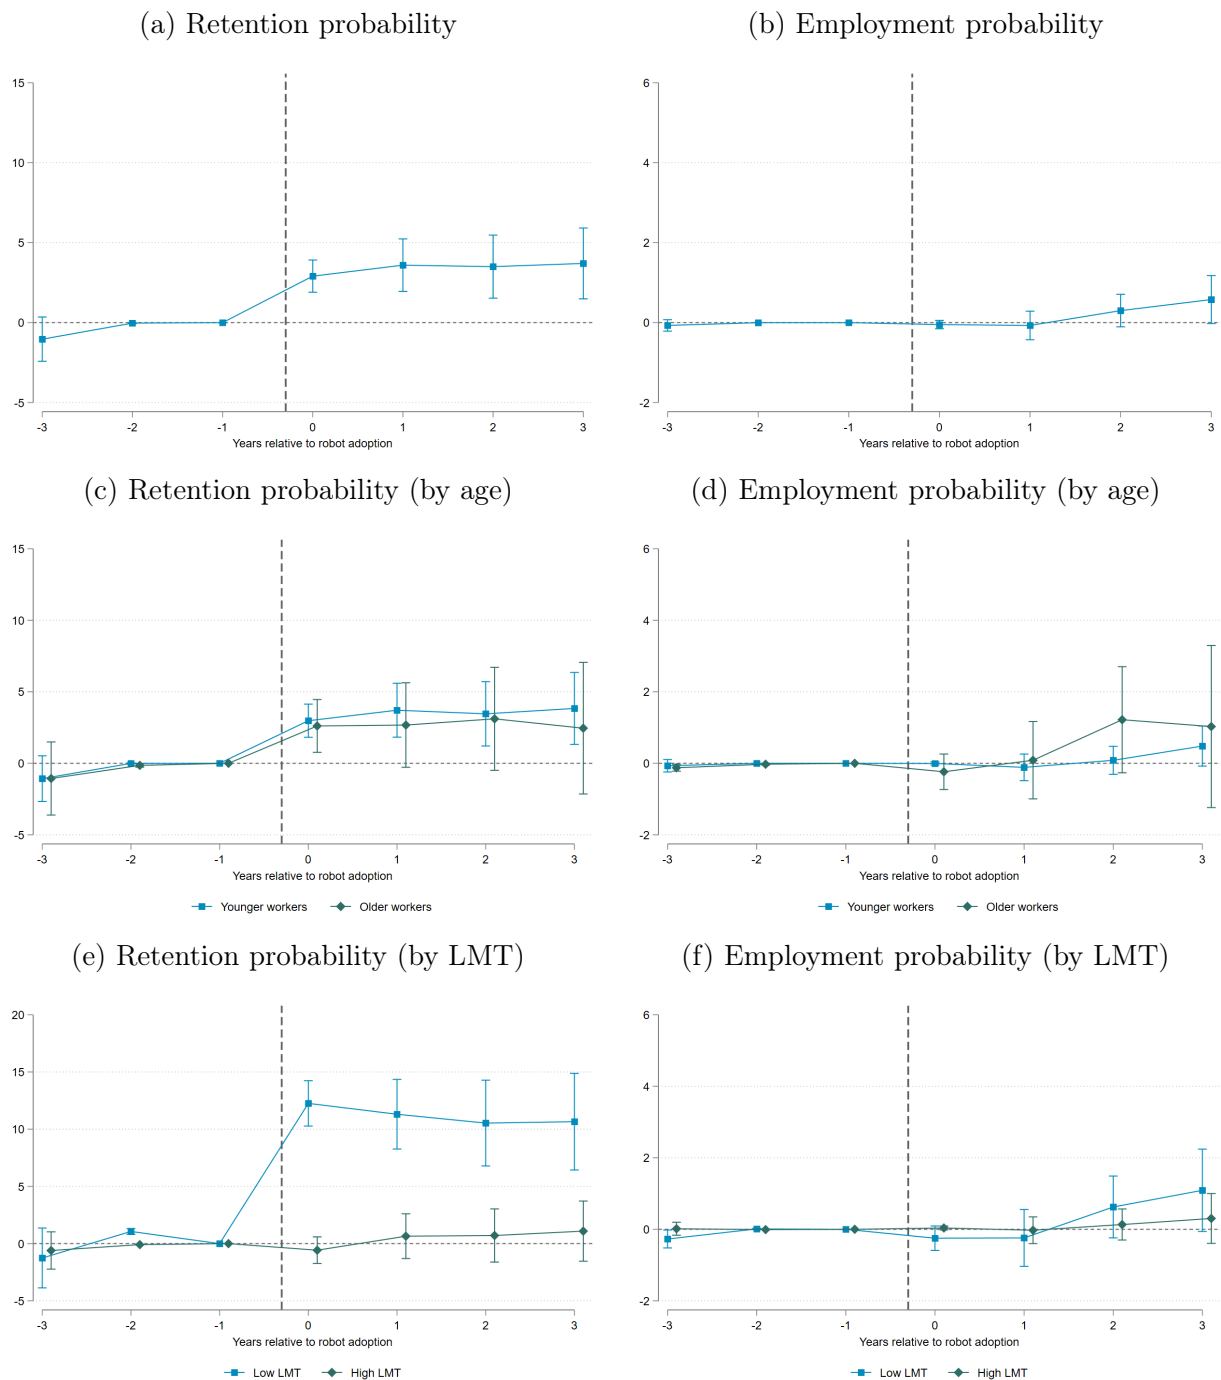
Table A.1: Summary statistics by adopter status

	Early adopters	New adopters	Never users (manufacturing only)
Log employment	4.65 (1.68)	4.45 (1.63)	3.19 (1.53)
Share high qualified workers	11.44 (11.64)	11.50 (12.33)	10.95 (14.98)
Robots/worker	0.15 (0.66)	0.08 (0.17)	. (.)
Robots	44.03 (384.78)	3.53 (10.09)	. (.)
Log total investment	13.46 (2.26)	13.02 (2.11)	11.73 (2.08)
Log VA per worker	11.14 (0.63)	10.98 (0.78)	10.80 (0.73)
Founding year	1990.78 (13.28)	1993.80 (14.24)	1996.03 (12.40)
Vacancies to fill	0.54 (0.50)	0.51 (0.50)	0.44 (0.50)
Manufacturing (excl. auto.)	0.67 (0.47)	0.62 (0.49)	0.96 (0.20)
Automotive	0.10 (0.30)	0.07 (0.26)	0.04 (0.20)
Non-manufacturing	0.17 (0.37)	0.27 (0.45)	0.00 (0.00)
Works council	0.52 (0.50)	0.43 (0.50)	0.24 (0.43)
Observations	280	215	2221

Notes: This table presents mean values and standard deviations for plants in 2018, categorized by their robot adoption status. Early adopters are plants that reported using robots in 2014. New adopters are plants that did not use robots in 2014 but reported using them at least once in the subsequent years (2015 to 2018). Never adopters are plants that did not use robots in any year between 2014 and 2018. Vacancies refers to the proportion of plants reporting open positions that could not be filled. All variables below are also expressed as shares, reflecting either the sector composition (classified according to WZ08) or the proportion of plants with a works council. Sector shares do not sum to 1 because of missing industry classifications.

B Results for the Entire Workforce (Including Non-Production)

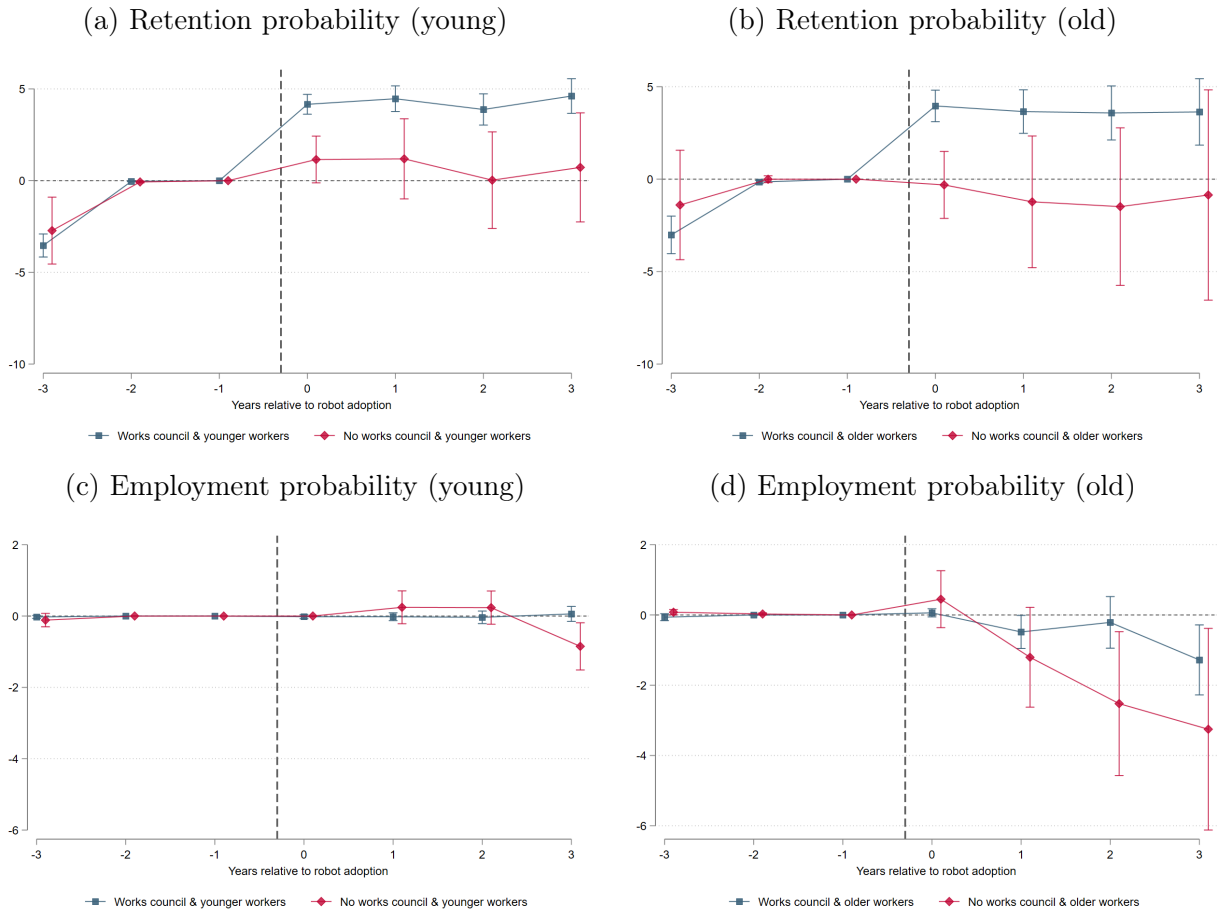
Figure B.1: Employment effect of automation in plants with vs. without works council – entire workforce



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes.

C Difference-in-Difference Results


Figure C.2: Employment effect of automation in plants with vs. without works council – by age



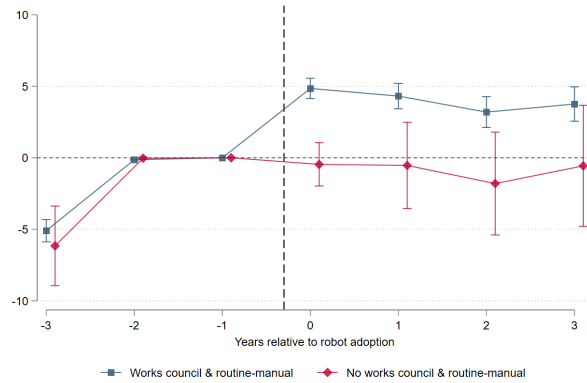
Notes: This figure shows the effect of robot adoption in plants on workers' employment, either at the initial plant (in Panel (a) and (b)) or anywhere (in Panel (c) and (d)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences estimates obtained from Equation 1, separately for workers in plants with and without works council. Workers are divided into groups based on whether they are below 55 (Panel (a) and (c)) or between 55 and 60 (Panel (b) and (d)) in the year of robot adoption. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

Figure C.3: Employment effect of automation in plants with vs. without works council – by occupation


(a) Retention probability (not routine-manual)




(b) Retention probability (routine-manual)



(c) Employment probability (not routine-manual)

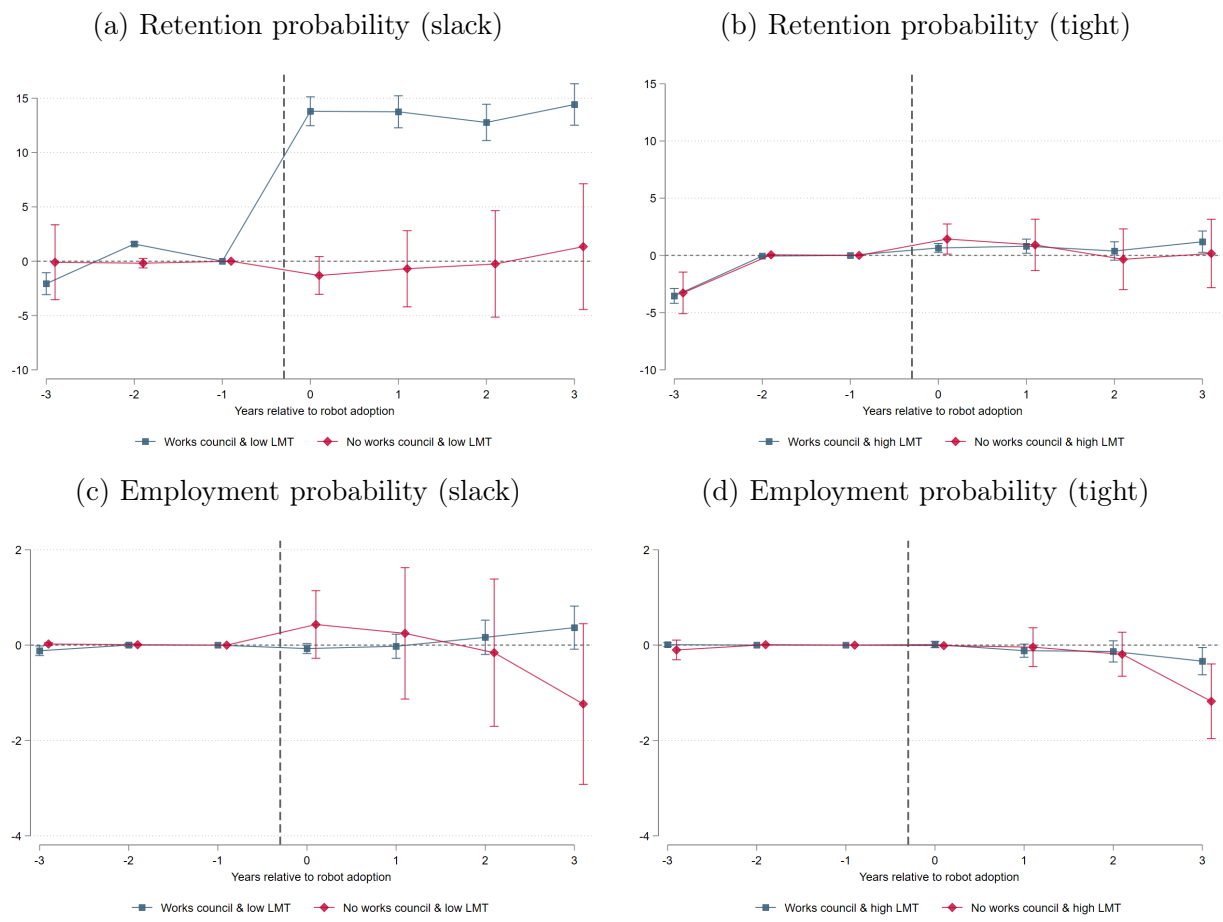


(d) Employment probability (routine-manual)



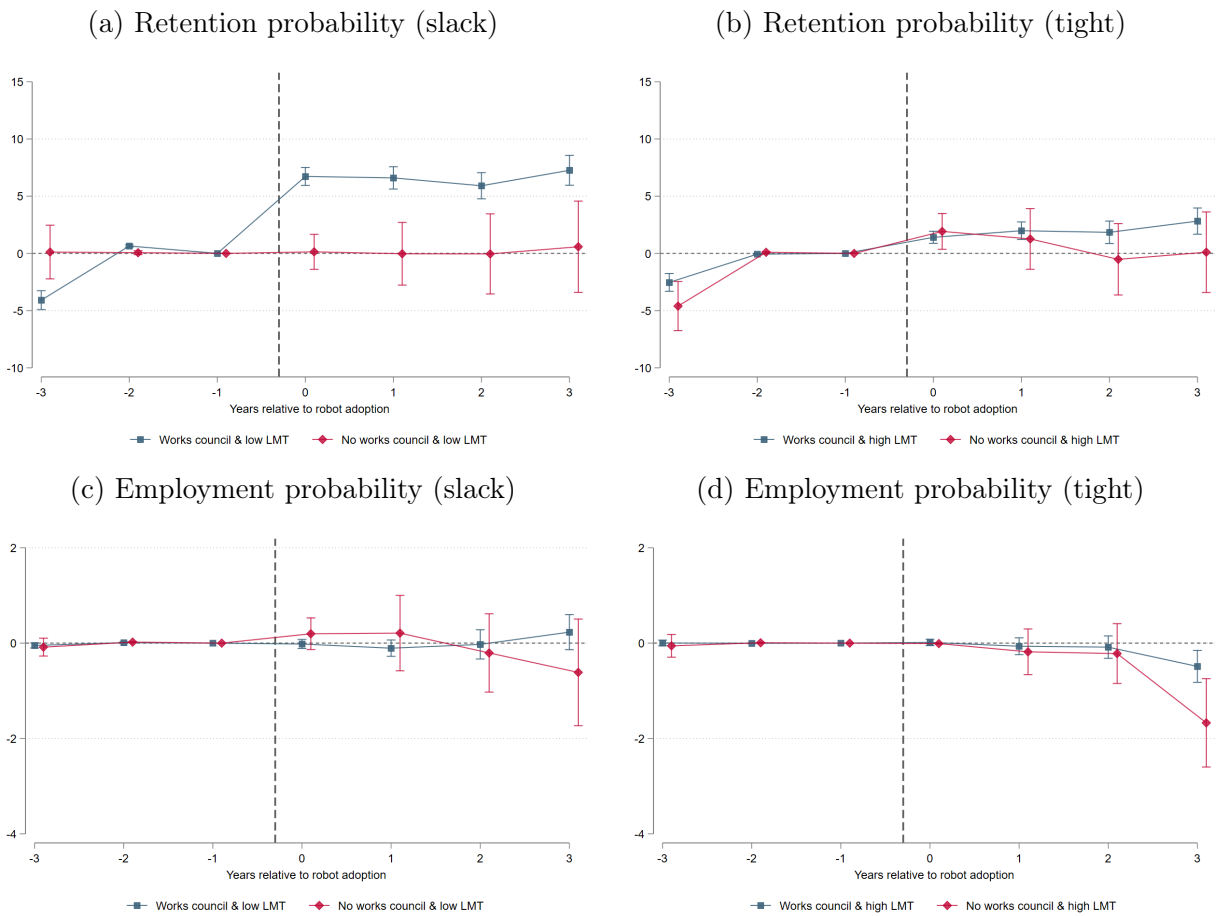
Notes: This figure shows the effect of robot adoption in plants on workers' employment, either at the initial plant (in Panel (a) and (b)) or anywhere (in Panel (c) and (d)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences estimates obtained from Equation 1, separately for workers in plants with and without works council. Workers are divided into groups based on whether they hold a routine-manual occupation (RMW), thus belong to Blossfeld occupation 2, (Panel (a) and (c)) or not (Panel (b) and (d)) in the year prior to robot adoption. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

Figure C.4: Employment effect of automation in plants with vs. without works council – by labor market tightness (25th percentile)



Notes: This figure shows the effect of robot adoption in plants on workers' employment, either at the initial plant (in Panel (a) and (b)) or anywhere (in Panel (c) and (d)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences estimates obtained from Equation 1, separately for workers in plants with and without works council. Workers are divided into groups based on plants' local labor market tightness (LMT) in the year prior to adoption (cutoff is the 25th percentile). This measure is obtained from Bossler and Popp (2024) and defined as the ratio of the number of vacancies to job seekers at the occupation-region level, weighted by plants' employment shares. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

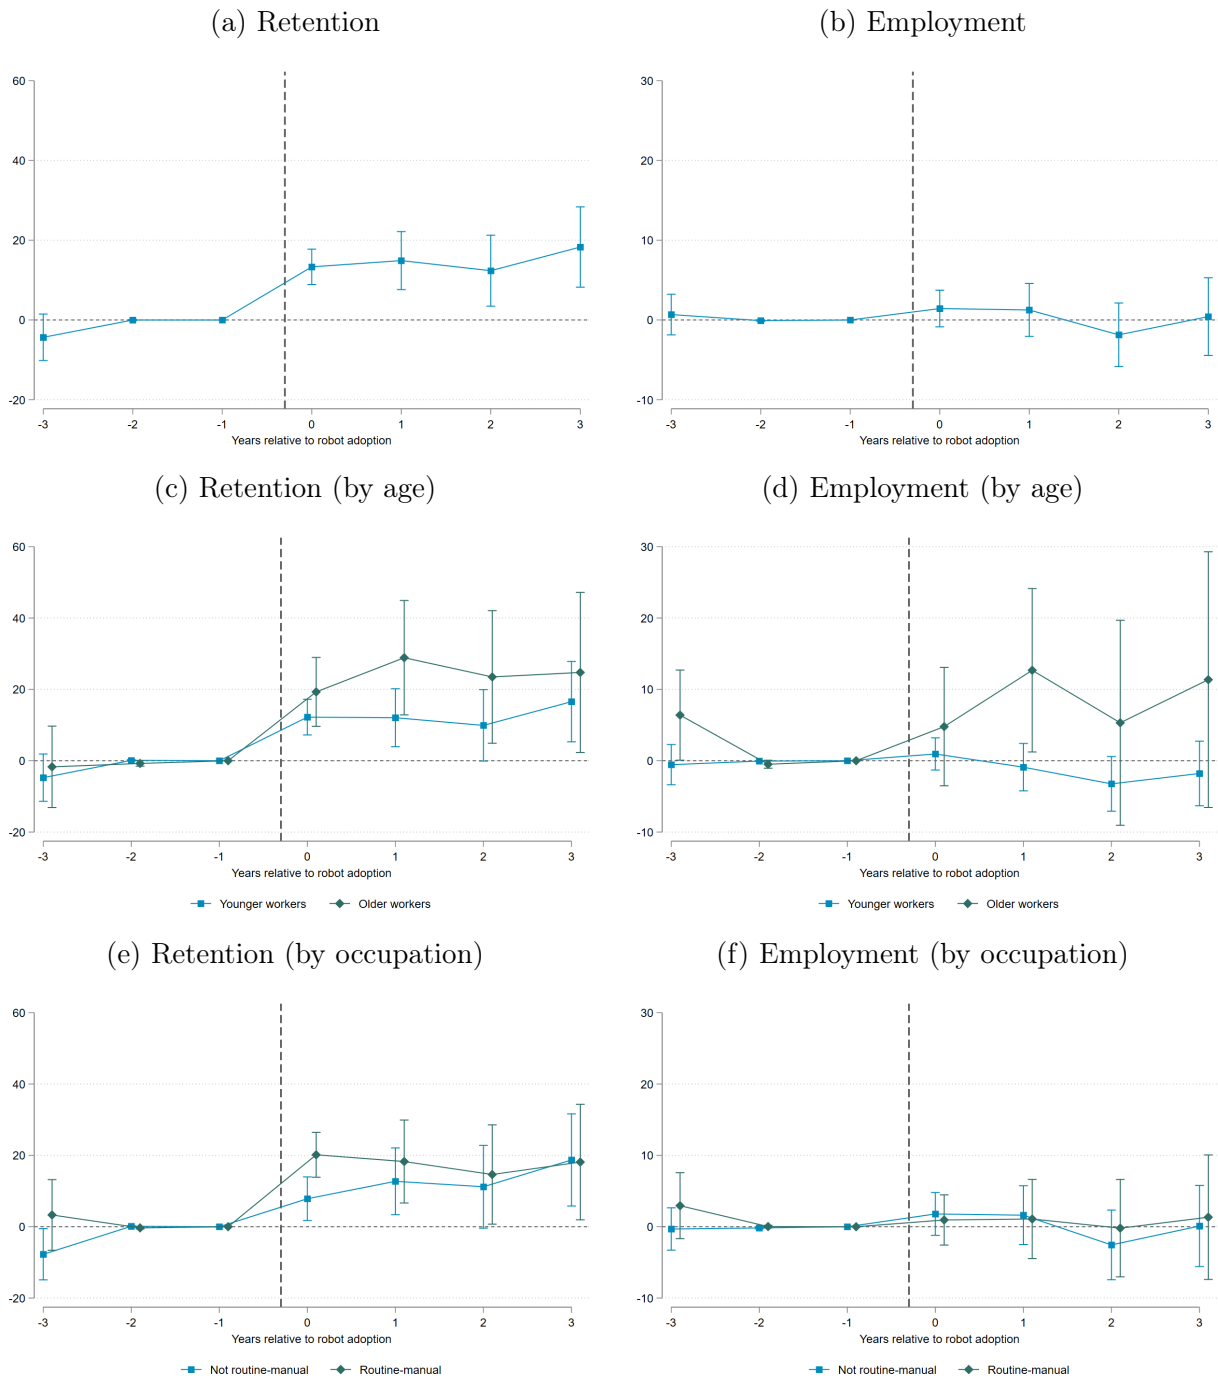
Figure C.5: Employment effect of automation in plants with vs. without works council – by labor market tightness (50th percentile)



Notes: This figure shows the effect of robot adoption in plants on workers’ employment, either at the initial plant (in Panel (a) and (b)) or anywhere (in Panel (c) and (d)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences estimates obtained from Equation 1, separately for workers in plants with and without works council. Workers are divided into groups based on plants’ local labor market tightness (LMT) in the year prior to adoption (cutoff is the 50th percentile). This measure is obtained from Bossler and Popp (2024) and defined as the ratio of the number of vacancies to job seekers at the occupation-region level, weighted by plants’ employment shares. Vertical bars indicate 90% confidence intervals based on robust standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

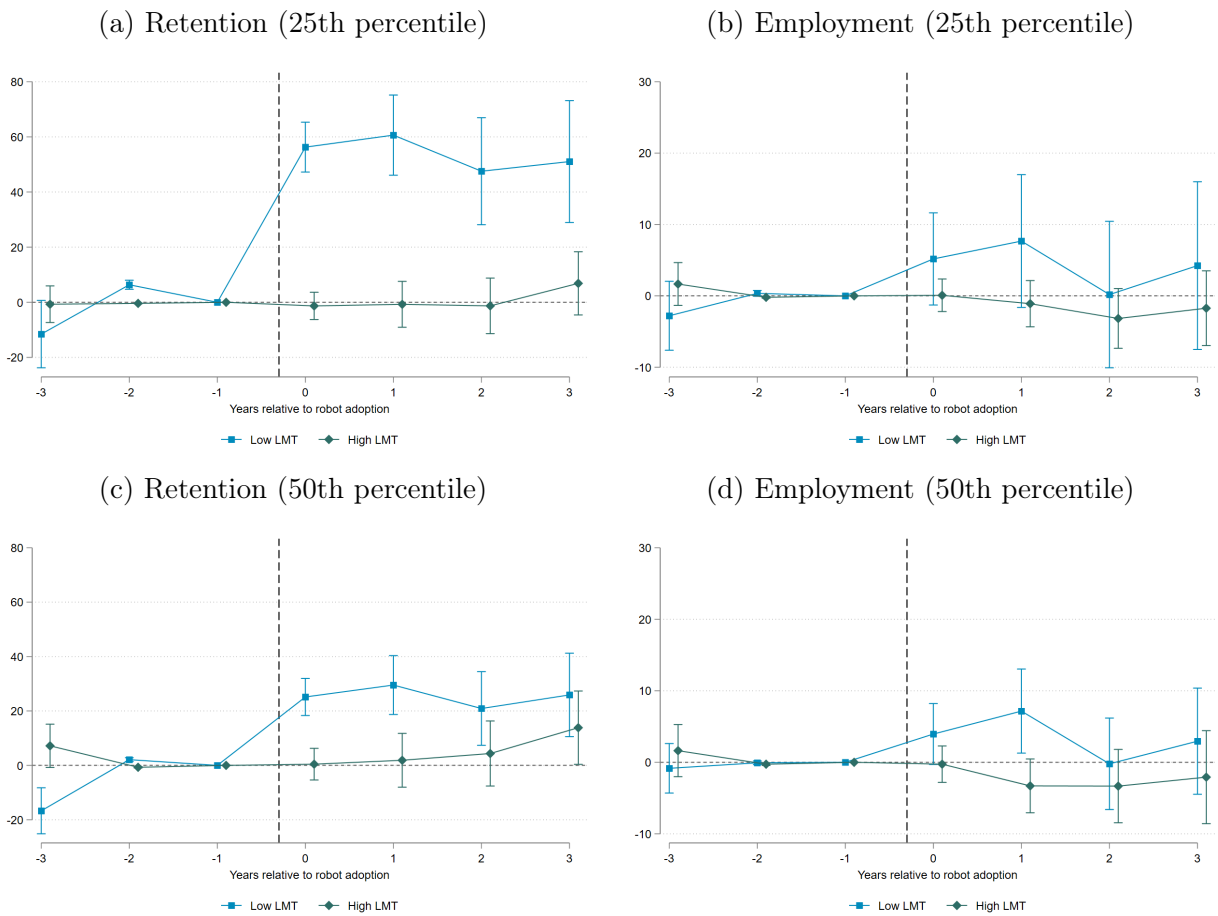
D Main Sensitivity and Robustness Checks

Figure D.6: Employment effect of works council by workers' characteristics – sensitivity using days employed



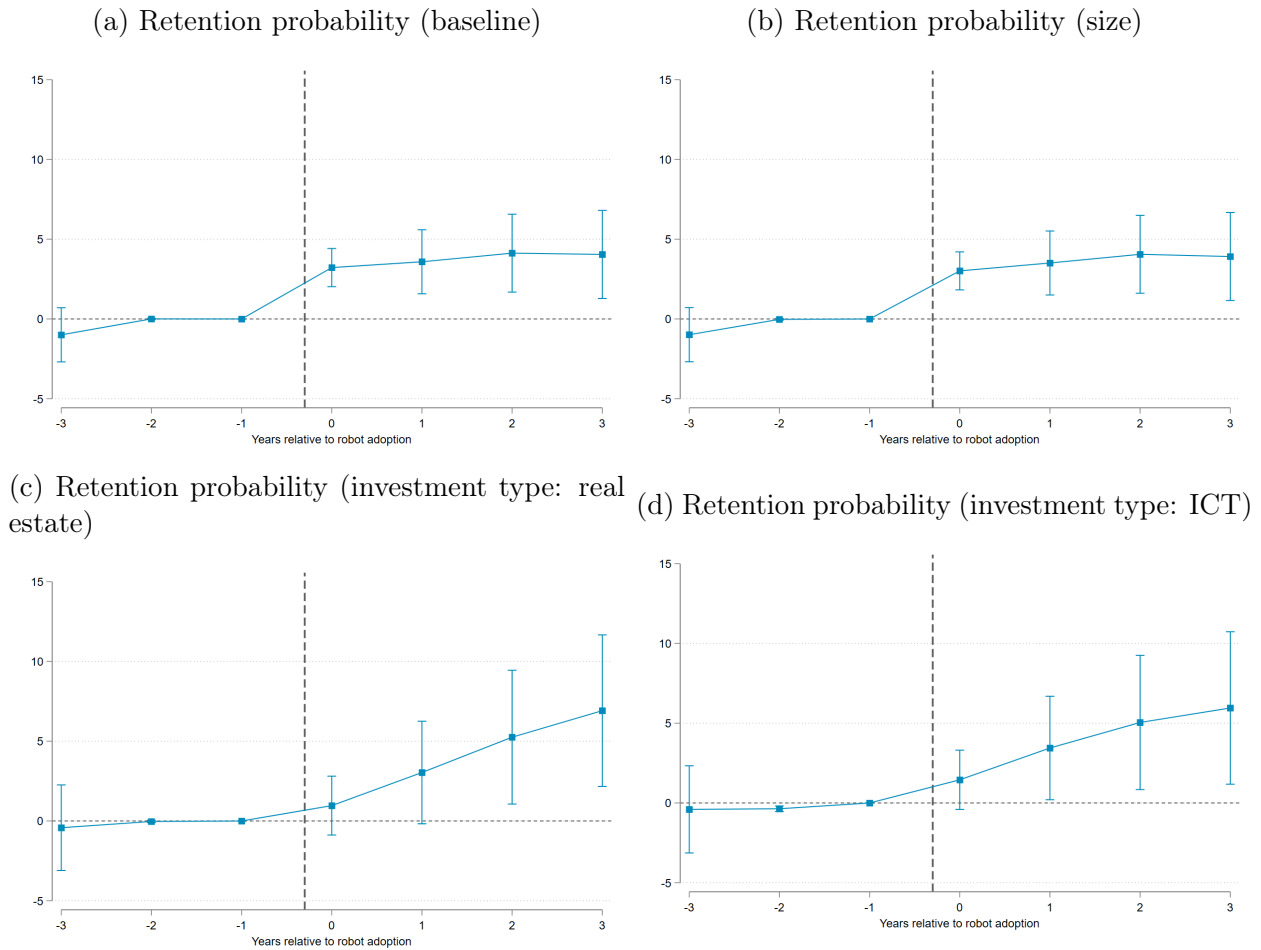
Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the number of days employed per calendar year, either at the initial plant (in Panel (a), (c), and (e)) or anywhere (in Panel (b), (d), and (f)). All sub-figures display the triple differences (DiDiD) estimates obtained from Equation 2, either for the whole sample or separately estimated across worker groups. In Panel (c) and (d), workers are divided into groups based on whether they are below 55 or between 55 and 60 in the year of robot adoption. In Panel (e) and (f), the division is based on workers' occupation in the year prior to adoption. Routine manual workers (RMW) are defined as being in Blossfeld occupation 2.

Figure D.7: Employment effect of works council by labor market tightness – sensitivity using days employed



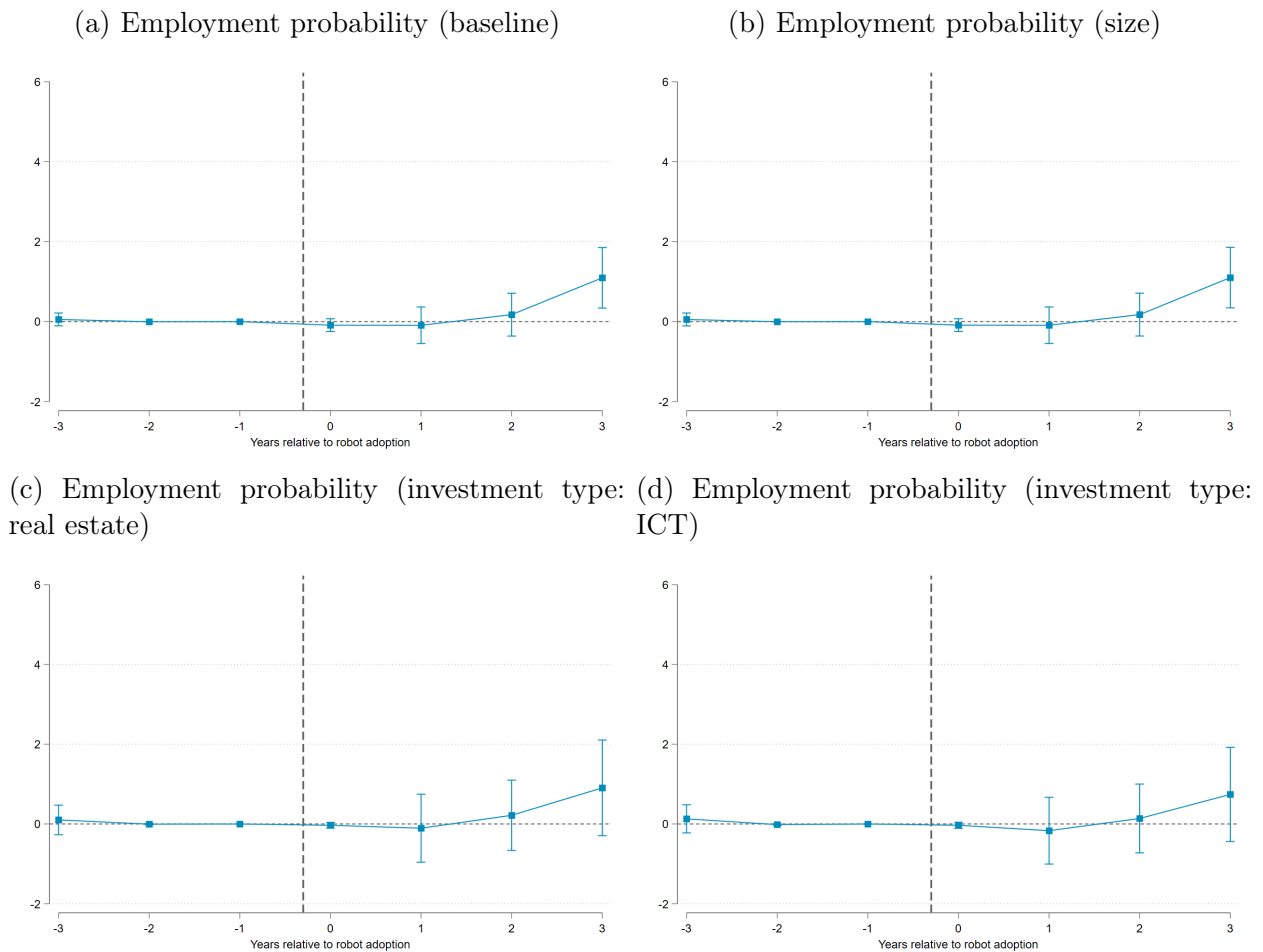
Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the number of days employed per calendar year, either at the initial plant (in Panel (a) and (c)) or anywhere (in Panel (b) and (d)). All sub-figures display the triple differences (DiDiD) estimates obtained from Equation 2, either for the whole sample or separately estimated across worker groups. Workers are divided into groups based on plants' local labor market tightness (LMT) in the year prior to adoption (cutoff is the 25th percentile in Panel (a) and (b), 50th percentile in Panel (c) and (d)). This measure is obtained from Bossler and Popp (2024) and defined as the ratio of the number of vacancies to job seekers at the occupation-region level, weighted by plants' employment shares. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

Figure D.8: Retention effect of automation in plants with vs. without works council – robustness



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' retention probability. Retention is measured as the probability of being employed at least one day per calendar year at the initial plant. Panel (a) displays the triple differences (DiDiD) estimates obtained from Equation 2. Panel (b) augments the regression equation by including a dummy variable that equals one if a worker is among the top 25th percentile employed at the largest plants regarding employment size, which is then interacted with event-year indicators. In Panels (c) and (d), the dummy variable equals one if a plant reports having additionally invested in real estate or information and communication technologies (ICT) in the year of robot adoption. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

Figure D.9: Employment effect of automation in plants with vs. without works council – robustness



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment probability. Employment is measured as the probability of being employed at least one day per calendar year. Panel (a) displays the triple differences (DiDiD) estimates obtained from Equation 2. Panel (b) augments the regression equation by including a dummy variable that equals one if a worker is among the top 25th percentile employed at the largest plants regarding employment size, which is then interacted with event-year indicators. In Panels (c) and (d), the dummy variable equals one if a plant reports having additionally invested in real estate or information and communication technologies (ICT) in the year of robot adoption. The dashed vertical line marks the event of robot adoption. Vertical bars indicate 90% confidence intervals based on standard errors clustered at the matching pair level. The sample of workers is restricted to individuals aged 25 to 60 in the year of adoption, being employed at least two years prior to robot adoption in the plant, and working in production (i.e. Blossfeld occupations 2-5). To ensure comparability, workers in robot-adopting plants are matched to similar counterparts in non-adopting plants. The matching process is described in detail in Section 3.

E Further Outcomes and Mechanisms

Table E.2: Works councils and likelihood of robot adoption and investment

	Robot adopter (New adopters vs. never users)	Robot user (All plants)	Log total investment (New adopters)
Works council	-0.409 (0.407)	-0.093 (0.579)	0.337 (0.469)
Mean of Y	1.60	3.60	13.26
SD of Y	12.54	18.64	2.00
R-squared	0.05	0.12	0.61
Observations	11,888	12,150	118

Notes: This table shows results from regressions of various outcome variables on an indicator whether a plant has a works council. Columns 1 and 2 show the results for the probability of newly adopting robots (between 2015 and 2018, Column 1) and having robots installed at any point in time (Column 2). The outcome in column 3 is the log of total investment in the year of robot adoption. In each regression, we control for 10 plant-size dummies, the share of highly qualified workers, the year of foundation, as well as industry fixed effects. For columns 1 and 2, we restrict the sample to all firms in the 2019 wave of the Establishment Panel with non-missing information, and for column 3 to all first-time robot-adopting plants. Standard errors are robust and clustered at the plant level.

Table E.3: Works councils, robot adoption and employment changes

<i>Panel A: Log employment</i>	$< \tau_{-1}$	$\tau_{-1,0}$	$> \tau_0$
Works council	0.010 (0.053)	-0.016 (0.053)	-0.056 (0.051)
Mean of Y	4.48	4.43	4.35
SD of Y	1.48	1.58	1.56
R-squared	0.97	0.97	0.97
Observations	288	271	258
<i>Panel B: Vacancies</i>	$< \tau_{-1}$	$\tau_{-1,0}$	$> \tau_0$
Works council	-11.559 (9.748)	-6.629 (9.048)	-5.767 (8.801)
Mean of Y	47.37	55.63	47.06
SD of Y	50.01	49.77	50.01
R-squared	0.11	0.11	0.18
Observations	287	271	258
<i>Panel C: Hires</i>	$< \tau_{-1}$	$\tau_{-1,0}$	$> \tau_0$
Works council	-3.413 (2.344)	-7.228*** (2.473)	-3.697 (3.081)
Mean of Y	10.83	13.96	11.52
SD of Y	21.43	28.72	21.36
R-squared	0.53	0.61	0.56
Observations	215	207	185

Notes: This table shows results from regressions of log employment (Panel A), the reporting of unfilled vacancies (Panel B), and the number of new hires in the first half of the calendar year (Panel C) on an indicator whether a plant has a works council. Columns $\tau_{<-1}/\tau_{-1,0}/\tau_{>0}$ report results from a pooled regression prior/during/after the event of robot adoption. In each regression, we control for 10 plant-size dummies, the share of highly qualified workers, the year of foundation, as well as industry fixed effects. Further, we restrict the sample to first-time adopters that have at least one observation in all year pools $b \in \{-4, -2; -1, 0; 1, 2\}$. Standard errors are robust and clustered at the plant level.

Table E.4: Works councils, robot adoption and the provision of training

<i>Panel A: Before adoption</i>	All workers	In simple tasks	In qualified tasks
Works council	4.922 (5.906)	12.165 (7.623)	8.070 (6.365)
Mean of Y	29.16	18.27	37.43
SD of Y	32.19	36.46	41.92
R-squared	0.14	0.19	0.17
Observations	264	198	257
<i>Panel B: During adoption</i>	All workers	In simple tasks	In qualified tasks
Works council	14.536*** (5.440)	15.190** (6.355)	13.450** (6.074)
Mean of Y	30.30	20.59	37.18
SD of Y	31.51	37.62	36.34
R-squared	0.15	0.17	0.19
Observations	266	199	250
<i>Panel C: After adoption</i>	All workers	In simple tasks	In qualified tasks
Works council	6.960 (5.539)	1.970 (6.390)	9.472 (6.340)
Mean of Y	26.62	22.79	30.57
SD of Y	32.41	39.20	39.60
R-squared	0.12	0.17	0.13
Observations	249	185	238

Notes: This table shows results from regressions of the share of trained workers (all workers, workers performing simple/qualified tasks) on an indicator whether a plant has a works council. Simple (qualified) tasks refer to the requirement of workers performing them having no (at least a) vocational degree. The panels report results from a pooled regression prior ($\tau_{<-1}$), during ($\tau_{-1,0}$), and after ($\tau_{>0}$) the event of robot adoption. In each regression, we control for 10 plant-size dummies, the share of highly qualified workers, the year of foundation, as well as industry fixed effects. Further, we restrict the sample to first-time adopters that have at least one observation in all year pools $b \in \{-4, -2; -1, 0; 1, 2\}$. Standard errors are robust and clustered at the plant level.