

Findeisen, Sebastian; Dauth, Wolfgang; Schlenker, Oliver

Working Paper

Organized labor versus robots? Evidence from micro data

Working Paper Series, No. 25

Provided in Cooperation with:

University of Konstanz, Cluster of Excellence "The Politics of Inequality. Perceptions, Participation and Policies"

Suggested Citation: Findeisen, Sebastian; Dauth, Wolfgang; Schlenker, Oliver (2024) : Organized labor versus robots? Evidence from micro data, Working Paper Series, No. 25, University of Konstanz, Cluster of Excellence "The Politics of Inequality", Konstanz,
<https://doi.org/10.48787/kops/352-2-pkkgn822nr6u9>

This Version is available at:

<https://hdl.handle.net/10419/300230>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

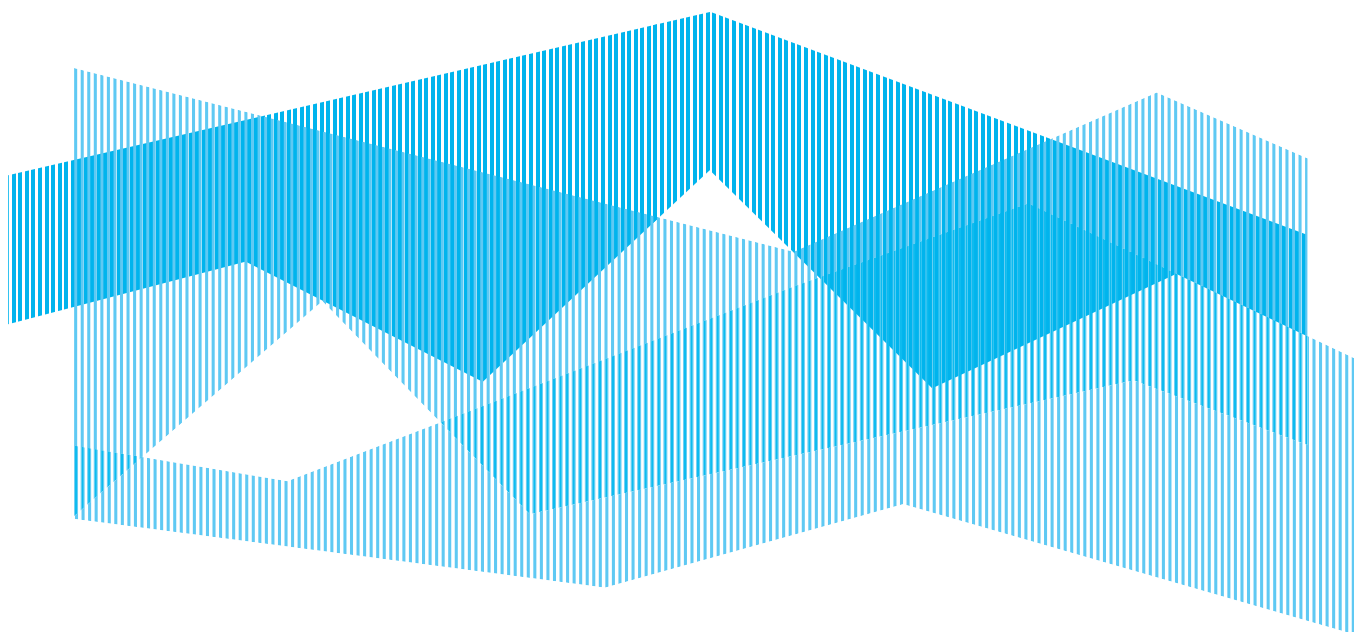
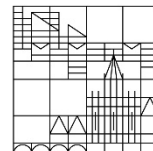
Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Organized Labor Versus Robots? Evidence from Micro Data

Sebastian Findeisen, University of Konstanz, sebastian.findeisen@uni-konstanz.de

Wolfgang Dauth, IAB Nuremberg and University of Bamberg, wolfgang.dauth@iab.de

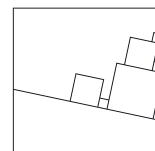
Oliver Schlenker, University of Konstanz and ZEW Mannheim, oliver.schlenker@uni-konstanz.de

Working Paper Series of the Cluster “The Politics of Inequality”:

→ <http://inequality.uni.kn/working-papers>

Working papers of the Cluster of Excellence “The Politics of Inequality” serve to disseminate the research results of work in progress prior to publication. Inclusion of a paper in the working paper series does not constitute publication and should not limit publication in any other outlet. The working papers published by the Cluster represent the views of the respective author(s) and not of the Cluster as a whole.

Cluster of Excellence
The Politics of Inequality



About the authors

Sebastian Findeisen is a Full Professor in the Department of Economics at the University of Konstanz since 2019 and a Principal Investigator in the Cluster of Excellence “The Politics of Inequality”. Prior to his current position, he served as Assistant and Associate Professor in the Department of Economics at the University of Mannheim. His research focuses on labor markets and economic policies. He is also affiliated with the CEPR, ZEW, Ceslfo, and IZA.

Wolfgang Dauth is head of the Research Department "Regional Labour Markets" and Professor of Regional Labor Economics at the Otto Friedrich University in Bamberg. His research lies at the cross-section of labor economics and the fields of Urban-/Regional Economics and International Economics. He is a Research Fellow at the Institute of Labor Economics (IZA).

Oliver Schlenker is a PhD student at University of Konstanz under the supervision of Sebastian Findeisen and Nicolas Ziebarth. He is also a member of the Cluster of Excellence “The Politics of Inequality” and a researcher at ZEW Mannheim. His research focuses on labor economics and health economics, especially on how technological change affects employees.

Organized Labor Versus Robots?

Evidence from Micro Data*

Sebastian Findeisen[†](r) Wolfgang Dauth[‡](r) Oliver Schlenker[§]

July 1, 2024.

Abstract

New technologies drive productivity growth but the distribution of gains might be unequal and is mediated by labor market institutions. We study the role that organized labor plays in shielding incumbent workers from the potential negative consequences of automation. Combining German individual-level administrative records with information on plant-level robot adoption and the presence of works councils, a form of shop-floor worker representation, we find positive moderating effects of works councils on retention for incumbent workers during automation events. Separations for workers with replaceable task profiles are significantly reduced. When labor markets are tight and replacement costs are high for firms, incumbent workers become more valuable and the effects of works councils during automation events start to disappear. Older workers, who find it more challenging to reallocate to new employers, benefit the most from organized labor in terms of wages employment. Concerning mechanisms we find that robot-adopting plants with works councils employ not more but higher quality robots. They also provide more training during robot adoption and have higher productivity growth thereafter.

Keywords: automation, organized labor, work councils, labor market tightness, worker re-training

JEL classification: J20, J30, J53, O33

*We thank Fabrizio Colella, Bernd Fitzenberger, Georg Grätz, Simon Jäger, Philipp Kircher, Peter Nilsson, Torsten Persson, Pascual Restrepo, David Seim, Jens Stegmaier, Simon Trenkle, Joonas Tuhkuri, and many 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 Politics of Inequality, CEPR

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

[§]University of Konstanz - Department of Economics & Cluster of Excellence Politics of Inequality, 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 conflict between workers and capital owners over the distribution of benefits and costs associated with new technologies (Acemoglu and Johnson, 2023).¹ At present, the recent rise of automation through robotics and artificial intelligence (AI) has sparked fresh debates about strategies for workers, employers, and governments to navigate labor market disruptions moving forward (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 of 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 no 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 and how policy responses could be deployed in light of ongoing technological disruption. We focus on codetermination, 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 about 40% of the workforce and have potent means to protect workers in terms of employment and working conditions. Their power ranges from veto rights against dismissals (that can only be overruled by labor courts) to co-decision rights in matters that concern, e.g., working hours, pay schemes, or workplace monitoring (Jäger et al., 2022).² In the literature, the common view is that firms automate as long as it is profitable, not internalizing the consequences for displaced workers.³ 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,

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).

²Work councils have been associated with reduced separations of (blue collar) workers (Hirsch et al., 2010; Budde et al., 2024) and a shift of bargaining power towards employees (Dobbelaere et al., 2024).

³Beraja 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.

2018). As works councils protect the interests of the incumbent workforce, their presence and rights to be involved in changes in work procedures caused by technology adoption can change 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), and account for the presence of works councils (across plants).⁴ We focus on robot adoption events at the plant level and estimate the effects on incumbent workers using event studies. Our approach is akin to computing the difference between two distinct 'difference-in-differences' estimators, where each model is estimated separately by works council status. 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.

Our main finding is that automation events increase retention of incumbent blue-collar workers by decreasing separation probabilities – but only in plants with active works councils. This finding is consistent with works councils, as a form of shop-floor representation, acting in the interest of incumbent workers. We find that the positive retention effects are very similar for older workers (above 55) versus younger ones, as separation probabilities decline by around 3.5%-points (around a 30% reduction) in both groups. However, it is older workers who are the main beneficiaries of these policies as employment rates increase by 2%-points in the following years. In the wake of automation events, young displaced workers, on average, adjust successfully and transition smoothly into new employment, while older workers are much more likely to remain unemployed after an automation-induced separation consistent with increasing adjustment costs over the life-cycle.⁵

Next, we investigate firms' incentives to shield incumbent workers from layoffs. In particular, in frictional labor markets, the value of retaining incumbent workers in automating plants increases in labor market tightness, as replacement and recruitment costs are higher (Kline et al., 2019; Jäger and Heining, 2022). Difficulties in hiring, thus, 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 is accompanied by similar retention effects in plants with and without works councils.

⁴The German economy has one of the highest robot densities in the world, providing rich variation in adoption events across plants.

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

Works councils might also engage in bargaining and limiting wage cuts for vulnerable workers who have competing task profiles with automation technologies or who face worse outside options.⁶ We test this by studying wages for two groups of workers, namely older employees, for whom we have documented difficulties in the adjustment process indicating scarce outside options, and workers with a routine-manual task profile who are confronted with high automation risk. We show that works councils have a sizeable positive wage effect of around 5% for these vulnerable workers. Compared to not directly affected production workers and white-collar workers, whose wages are not affected by automation, routine-manual and older workers experience significant wage reductions – but only in plants without works councils, highlighting the ability of shop-floor representation to prevent wage cuts.

In the last part of the paper, we investigate by which means works councils dampen the negative effects of automation for incumbent workers. When automation-related investment decisions internalize (part of) the cost of displacement, this raises the threshold for the marginal investment decision to become profitable. All else being equal, this diminishes the incentives to implement ‘so-so technologies’ Acemoglu and Restrepo (2018), which are investments in automation equipment that displace workers but come only with modest productivity improvements. Consequently, automation events in firms with worker representation should go along with larger productivity gains, as those must counterbalance the internalized displacement costs. In line with this argumentation, we find that robot-adopting plants with works councils experience greater productivity growth after robot adoption compared to their counterparts without entrenched co-decision-making rights for workers.

We use proxies for robot quality and find evidence that this 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 are plausibly directly related to increased retention. Using the panel dimension of our data, we can detect ‘training spikes’ around adoption, solidifying this interpretation. In the years before automation events, firms with works councils have a larger share of workers who participate in training compared to their counterparts. This gap increases from around 5%-points to around 15%-points during automation events while in the years after, it reverts back to around 5%-points. Since co-determination rights increase job security in the face of automation, this should make workers more willing to invest in firm-specific skills and participate in training, as argued in Freeman and Lazear (1995).

Our paper contributes to the literature using data on firms and worker outcomes

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.

to understand the effects of new automation technologies on the labor market. Robot-adopting plants typically expand employment (Koch et al., 2021; Hirvonen et al., 2023), often at the expense of competing firms. On average, directly affected workers with replaceable task profiles lose out (Bessen et al., 2023), while other workers might see gains (Acemoglu et al., 2023) with differential effects across age groups (Deng et al., 2023a). 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. Dauth et al. (2021) identify a suggestive interaction between the displacement effect at the local labor market level and local unionization rates. In contrast, in this paper we are able to leverage detailed micro data on firm adoption and worker trajectories, utilizing event studies to demonstrate the dynamic interaction between institutions and automation. In addition, our baseline empirical strategy, using event studies in combination with matched control group workers, is related to papers in the literature studying the cost of job loss (Bertheau et al. (2023) and Illing et al. (2024)).

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) and Mueller (2012) highlight a positive correlation between plant-level productivity and the presence of works councils. 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 of shared governance by being able to identify an interaction effect of work councils with technology adoption.

The paper proceeds as follows. In the next section, we describe the data and briefly discuss the institutional background. In Section 3, we present the empirical models and strategy. Section 4 contains the main results. In Section 5, we show evidence on mechanisms with a focus on technology adoption and worker training. Section 6 concludes.

2 Data

Administrative Labor Market Data. For our main analysis in Section 4, 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 IAB Establishment Panel (Bellmann et al., 2021), an 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 codetermination, as well as information on the plant’s technical endowment.

We link this survey to the administrative records of all individuals who were ever employed in one of the plants of the 2019 wave. A detailed description of the data sources can be found in the Online Appendix. 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 merge further information on job characteristics. First, we classify occupations according to the popular classification by Blossfeld (1987), which permits us to separate blue-collar production jobs from others. Second, we merge data from the 1991 BIBB/IAB Employment Survey and follow Spitz-Oener (2006) to obtain the share of routine-manual tasks performed in each occupation. We use this additional information on the task profile to identify workers who are most prone to being directly affected by automation, as their job contains a high share of potentially replaceable routine tasks (Acemoglu and Restrepo, 2018; Acemoglu et al., 2023). 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.

In addition to the worker-level analyses in Section 4, we study the mechanisms of how codetermination interacts with technology adoption, further training, and increasing productivity in Section 5. For those analyses, we use the plant-level data of the IAB establishment survey to construct a sample of first-time robot adopter plants between the years 2015 and 2018.

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.⁷ We first distinguish between robot users and plants that have never used robots up until 2020. 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

⁷Plümpe 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). Deng et al. (2023b) report that robot adopters are positively selected among firms, e.g. in terms of size and productivity, as in other periods and countries (e.g. Koch et al., 2021).

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.

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; 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 Stegmaier (2012).

3 Event-Study Models with Double and Triple Differences

Our empirical strategy aims to identify the consequences of an event where firms automate parts of their routine-manual tasks on directly affected incumbent workers – and to analyze 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 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 matching pairs to have the same sex, nationality, contract, Blossfeld occupation, aggregate industry, missing vs. non-missing AKM plant fixed effect, and works council status (see Appendix Table A.1 for additional details regarding the variables used). The restrictions ensure that matched workers operate under similar market/working conditions, are exposed to similar production technologies, and account for common trends and sorting. Matching is repeated twice to increase the sample size, whereby the caliper ensures that pairs cannot be too different from one another regarding their propensity score. Thus the final dataset consists of treated workers who experienced robot adoption in a year between 2015 and 2018 and up to two matched never-treated workers who, in the same year, were equally likely to experience this event.

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 a matched plant at least two years prior to robot adoption. Additionally, we restrict our sample to directly affected production workers (as defined in Blossfeld (1987)) who are aged between 25 and 60 in the year of robot adoption. This leaves us with yearly observations of 26,047 individuals in 962 plants. Appendix Table A.2 shows descriptive statistics among matched workers using a two-sided t-test. Reassuringly, both groups differ only very little, even for characteristics that were not part of the matching algorithm.

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:

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

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 ($g = NWC$), one could run separate regressions and compare

the results in a further step. To directly quantify the differences between plants with and without works councils, we instead use 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. The 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 all standard errors at the level of matched worker pairs.

4 Results

4.1 Works Councils and Retention During Automation Events

We begin by showing how automation – induced by the adoption of robots at the plant level – affects incumbent workers’ employment prospects. 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 3.5%-points more likely to remain in their initial establishment relative to their matched counterparts in non-automating plants with a works council. Thus, it seems that works

councils actively prevent the workforce from leaving when plants adopt robots.

Potentially, this has the aim to ease the consequences of automation, as indeed, the higher retention translates into a small positive employment effect (Panel (b)). While the adaptation of robots does not worsen the employment probability of workers in plants with works councils, it decreases by around 0.5%-points for workers from plants without works councils.

4.2 Effects By Age and Employer-Specific Tightness

Next, we turn to the triple difference (DiDiD) design of Equation 2 to study the heterogeneity of the mitigating effect of works councils with respect to worker and firm characteristics. Panels (a) and (b) in Figure 2 are equivalent to the previous results. The coefficients now reflect the difference between the separately regressed DiD models.

We start by focusing on older workers, who typically adjust to new technologies less easily. Employers may prefer to substitute those workers with young ones. Thus, works councils may focus their efforts on this particular group because older workers also face stronger barriers to finding new employment. 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 about 3.5%-points more likely to remain in their initial plant, indicating that works councils have no age bias. However, unlike young workers, who have low adjustment costs of displacement, works councils increase the probability of being employed for older workers by 2%-points. These effects are both statistically and economically significant given that older workers are, on average, non-employed with a probability of around 2% three years after the adoption of robots.⁸

In our main specifications, we focus on production workers (i.e. Blossfeld occupations 2-5). When we instead consider all workers, we find a similar retention effect, see Appendix Figure A.1. The same is true if we split our sample into routine-manual and all other production workers, see Appendix Figure A.2. However, the overall employment effect gets closer to zero as non-production workers are less affected by automation and have lower adjustment costs of switching across employers and industries. This highlights, that although works councils matter differently across worker groups, they do not seem to be biased when negotiating with the management about retaining incumbent workers.

In contrast to the characteristics of the workforce, we expect external circumstances to lead to differences in the retention effect. For the management, the value of retaining workers increases as replacement and recruitment costs become higher (Kline et al., 2019; Jäger and Heining, 2022) – and by that aligning the incentives of the management and the incumbent workforce regarding retention. To test whether works councils have a higher

⁸These effects are not driven by early retirement as we only consider spells of individuals subject to social security contributions.

retention effect when labor is scarce, we use the novel measure of plant-specific labor market tightness from Bossler and Popp (2024). We use this measure to categorize plants based on their specific local labor market tightness, a proxy for replacement costs, in the year before robot adoption. Indeed, Panel (e) reveals that works councils only have a positive retention effect during automation when labor markets are not tight, i.e., when plant-specific replacement costs are low. For the subgroup of workers in plants with the lowest labor market tightness, we find that works councils increase retention by around 18%-points, translating into a positive employment effect (see Panel (f)). In the Appendix in Figures A.3 and A.4, we show the corresponding DiD results of this section.

4.3 Wage Effects

In order to see whether works councils not only benefit workers by increasing their employment prospects (at their initial plant) but also allow them to partake in the gains of automation, we examine workers' wages. To this end, we restrict the sample to spells of individuals who are continuously employed at the initial establishment between $t-2$ and $t+3$, which allows us to capture the evolution of wages within the plants.

In contrast to labor-augmenting technological change, firm-level automation has been found to have only limited effects on wage (e.g., Koch et al. (2021) and Dixon et al. (2021)). However, these average effects mask substantial heterogeneity as low-skilled and routine-manual workers may see their human capital being depreciated by automation and, therefore, experience wage losses.

Although works councils do not have the right to enter into wage negotiations, they can exert indirect pressure through their right of codetermination in other areas, thus redistributing rents from the employer to the employees (Freeman and Lazear, 1995). In addition, they can affect wages by negotiations over classifications of workers into pay grades (Mohrenweiser, 2022) and also prevent workers from slipping into lower pay grades.

In the context of automation, we expect works councils to have a positive rent-sharing impact not by increasing overall wages but by protecting the wages of vulnerable subgroups from falling. One vulnerable subgroup is production workers with a routine-manual intensive task profile.⁹ Another group is older employees. Automation not only replaces but also generates new tasks, for instance, robot maintenance. Older workers usually have more difficulties learning new (digital) tasks and, therefore, may face downward pressure on their wages.

Looking at the wages for all incumbent production workers (Panel (a)) and the entire workforce (Panel (b)) in Figure 3, we find no wage effects of works councils during automation. However, when differentiating by task profile (Panel (c)) and age (Panel (d)),

⁹We classify workers as routine-manual intensive based on whether they are in Blossfeld occupation group 2 in the year prior to robot adoption. These occupations are characterized by a high share of routine-manual tasks which are easily automatable.

we find that the positive wage effect of works councils during automation is especially present for workers aged 55 and above and in routine-manual occupations. For them, works councils increase wages by around 5%. This positive effect is driven by a stark reduction of wages for routine-manual and older workers in plants without works council (see Figure 3 Panel (c) and (d)). In plants with works councils, however, the wage level for these (and all other workers) does not change as a result of robot adoption, suggesting that works councils use their power to protect automation-exposed workers from wage cuts.

5 Mechanisms: Technology Direction and Worker 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. This comparison of plants by works council status does not reveal causal effects but reveals meaningful patterns that distinguish those two groups of plants.

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 ($\tau > 0$).¹⁰

Results. We start by relating the presence of codetermination to the type of automation technology adoption. Panel A in Table 1 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%-points higher share of cage robots, which are commonly associated with higher productivity (Gurgul, 2018) but also higher investment costs.

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 are 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, which is consistent with the literature on codetermination (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 the returns to automation to be sufficiently high to offset the higher costs of displacement. However, as we show in the Appendix, we do not find evidence that works councils change the purpose or even hinder

¹⁰Pooling 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.

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 A.3). We also do not detect differences between plants with and without works council around adoption regarding employment levels or vacancies (see Table A.4 in the Appendix).

Finally, we want to shed light on whether works councils do not only increase retention during automation events as documented in Section 4 but also increase re-training efforts. 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 1 demonstrates that firms with works council have a 5%-points (statistically insignificant) higher share of workers who receive training in a given year, conditional on the set of controls. This is consistent with previous findings (Stegmaier, 2012; Mohrenweiser, 2022). However, during automation events, the gap in the share of trained workers increases sharply to 14.5%-points and returns approximately to previous levels thereafter. When distinguishing between low- and high-skilled workers participating in training in Appendix Table A.5, we find that the spike in training provision in plants with works council is similar across worker groups. This increase is meaningful given that, pre-adoption and across all plants, only 18% of low-skilled workers receive training. The patterns are in line with systematic investment in training during automation events in plants with works councils. It might not only reflect the willingness of firms to supply training, but also an elevated 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).

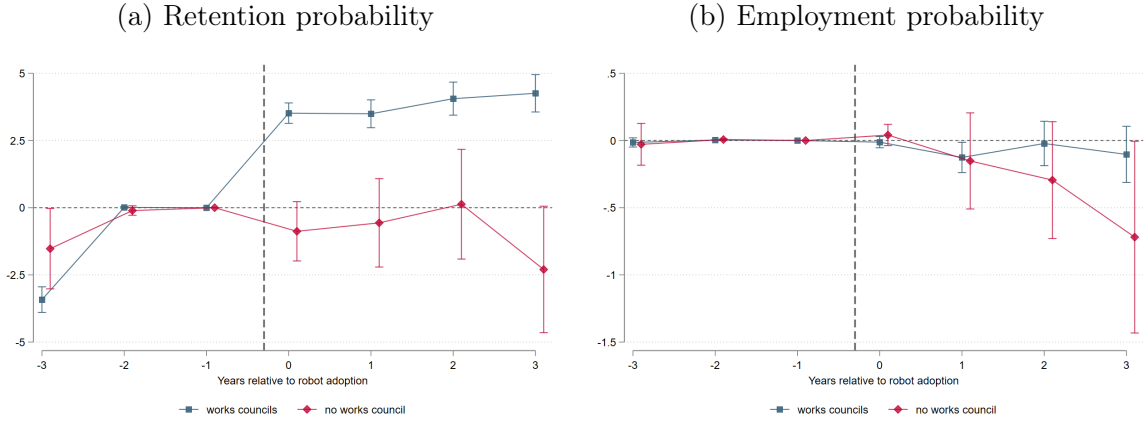
6 Conclusion

We find that work councils moderate adverse effects from automation events on incumbent workers by reducing separations. Older workers, with limited adjustment possibilities, benefit the most. When replacing workers is costly for firms, as measured by high plant-specific labor market tightness, separation and retention effects in automating firms with and without work councils converge. A higher value of incumbent workers for firms, hence, aligns the objectives of work councils and the management. We further find that works councils prevent negative wage effects from automation for vulnerable workers.

We observe that robot adoption is associated with larger productivity growth and increased training efforts in the presence of work councils. Our observation that increasing productivity goes hand in hand with retaining and retraining incumbent workers is consistent with the view that work councils facilitate cooperative solutions in the wake of conflicting interests between capital owners and workers (Müller-Jentsch, 1995). Understanding how other types of labor market institutions might alter the direction and consequences of new (automation) technologies is an important topic for future research.

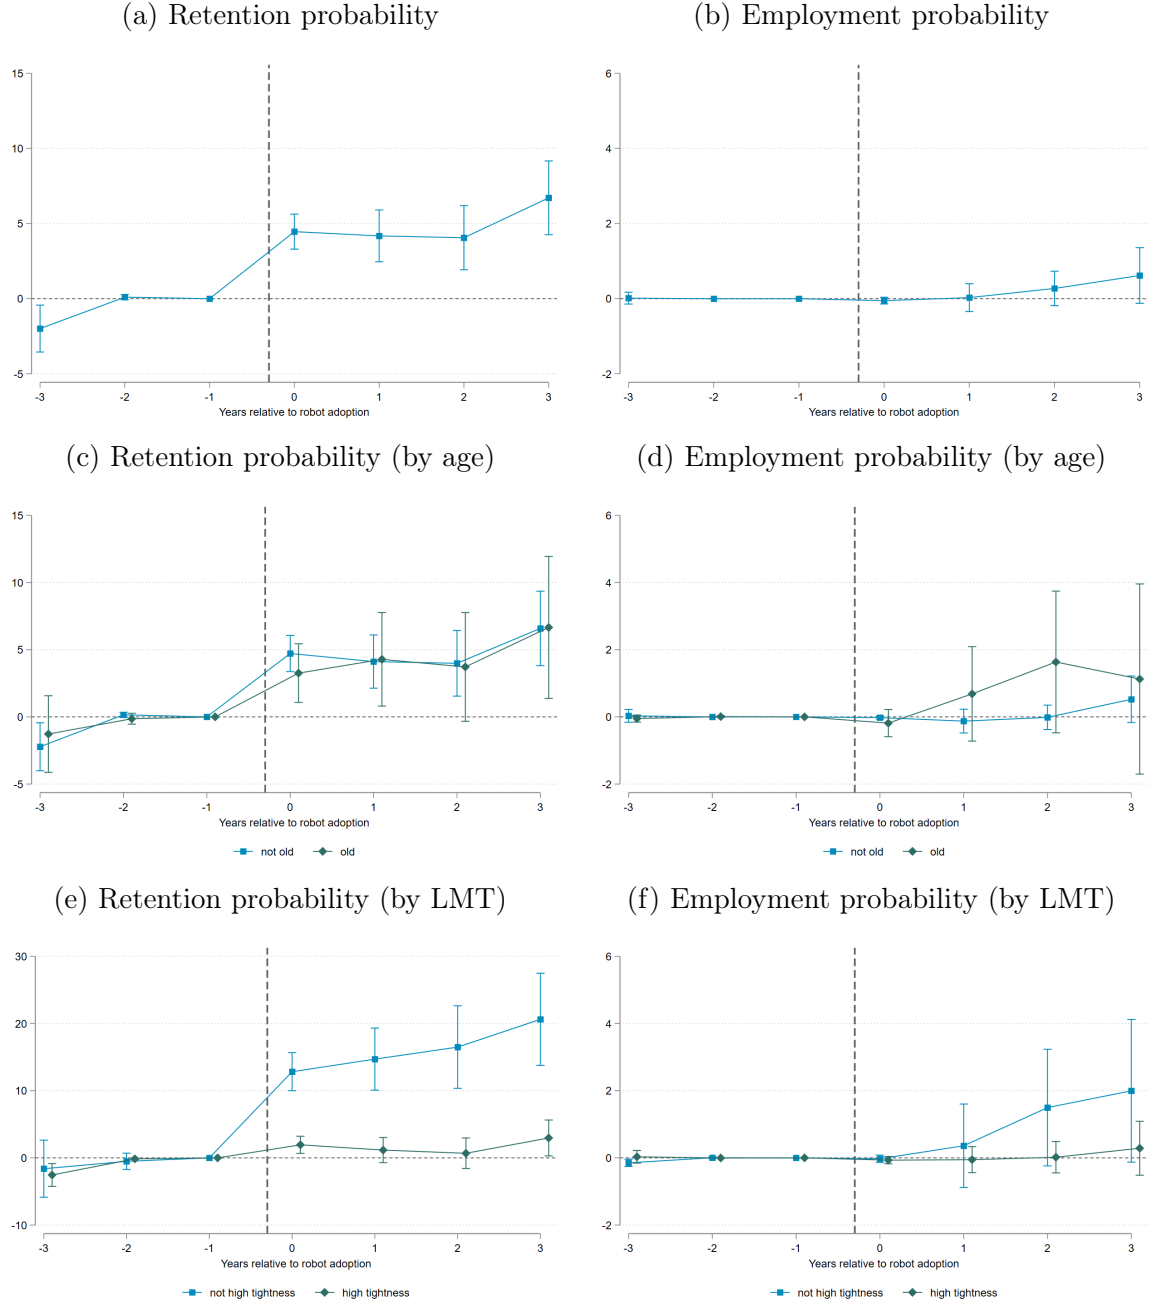
Exhibits

Figure 1: Employment effect of robotization by works council status



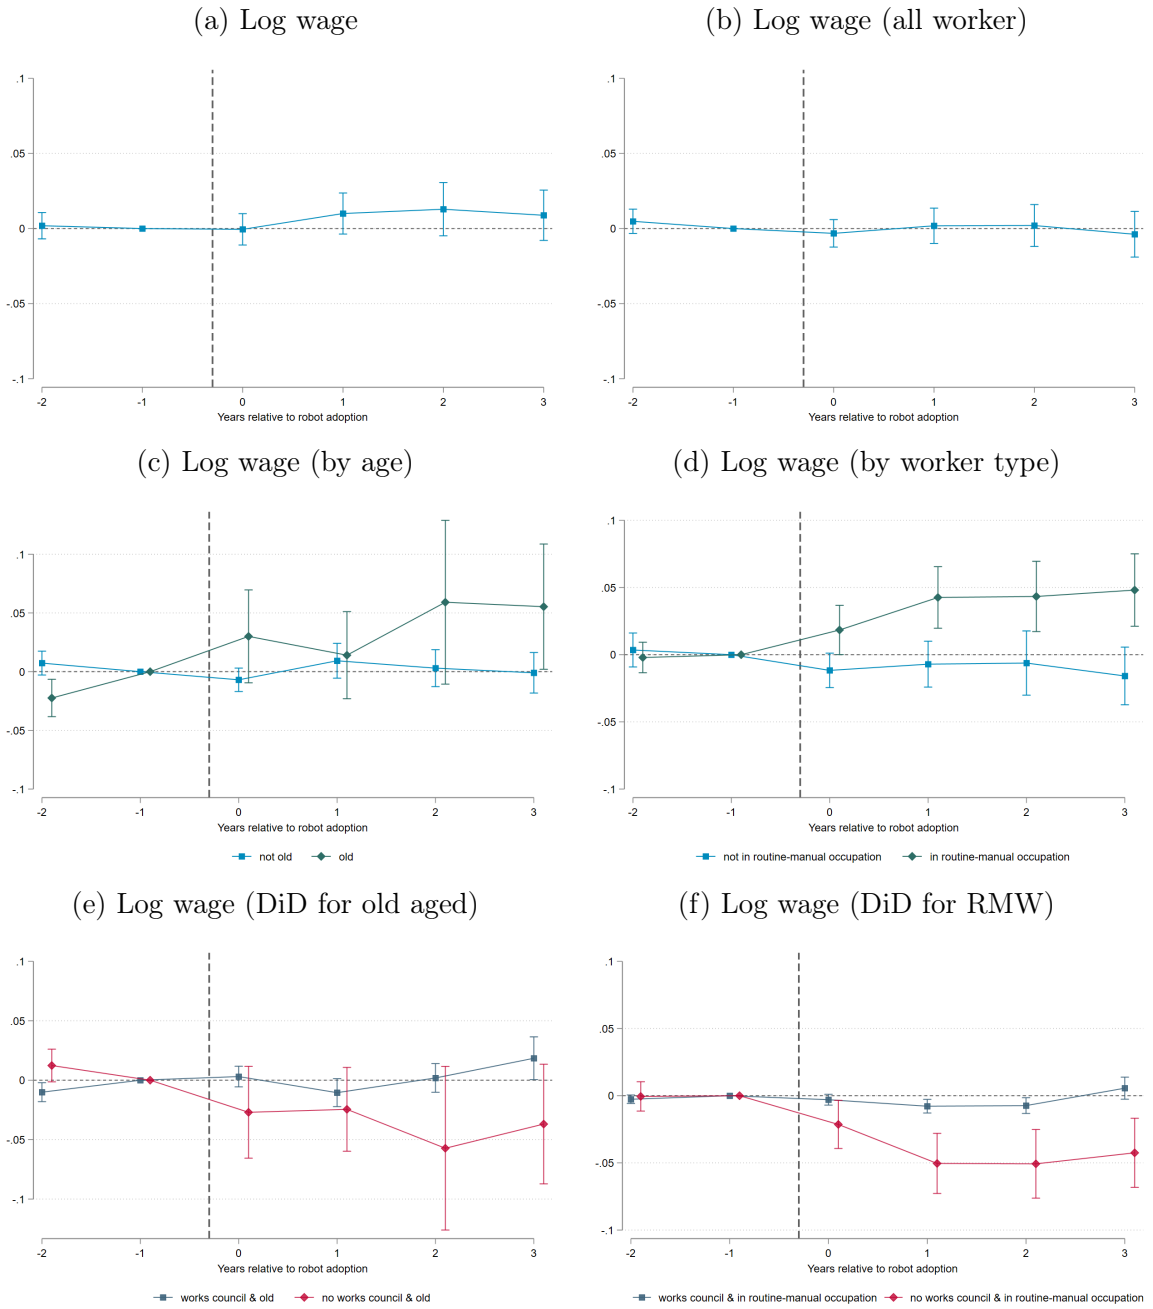
Notes: This figure shows the effect of robot adoption in plants on workers' employment, either at the initial plant (in Panel (a)) or anywhere (in Panel (b)). Employment is measured as the probability of being employed at least one day per calendar year. All sub-figures display the difference-in-differences (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.

Figure 2: Employment effect of works councils during automation



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 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. 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 A.3, and A.4 in the Appendix.

Figure 3: Wage effect of works councils during automation



Notes: This figure shows the effect of the presence of a works council during the event of robot adoption (in Panel (a) to (d)), and the effect of robot adoption on workers' log daily wage (Panel (e) and (f)). Panel (a) to (d) display the triple differences (DiDiD) estimates obtained from Equation 2, separately estimated across worker groups. Panel (e) and (f) display difference-in-differences (DiD) estimates obtained from Equation 1 for a specific worker group. For Panel (b), 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 and three years after 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), 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 (e), 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. For Panel (e) and (f), the sample is restricted to workers aged 55 and above and in routine-manual occupations 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.

Table 1: 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 $b \in \{-4, 1; -1, 0; 1, 2\}$ (for Panel B and C). Standard errors are robust and clustered at the plant level.

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.
- Autor, David (2024). *Applying AI to Rebuild Middle Class Jobs*. NBER Working Paper 32140.
- 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.
- 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.
- Bossler, 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*. Working Paper.
- 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.
- 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 (2023a). *Robots, Occupations, and Worker Age: A Production-Unit Analysis of Employment*. IZA Discussion Paper No. 16128.
- Deng, Liuchun, Verena Plümpe, and Jens Stegmaier (2023b). “Robot Adoption at German Plants”. In: *Jahrbücher für Nationalökonomie und Statistik*.
- 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.
- 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 (2012). “Works Councils and Establishment Productivity”. In: *ILR Review* 65(4), pp. 880–898.
- 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.

Supplementary Online Appendix

Organized Labor Versus Robots? Evidence from Micro Data

Sebastian Findeisen, Wolfgang Dauth, Oliver Schlenker

Appendix A: Additional Tables and Figures

Appendix Table A.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

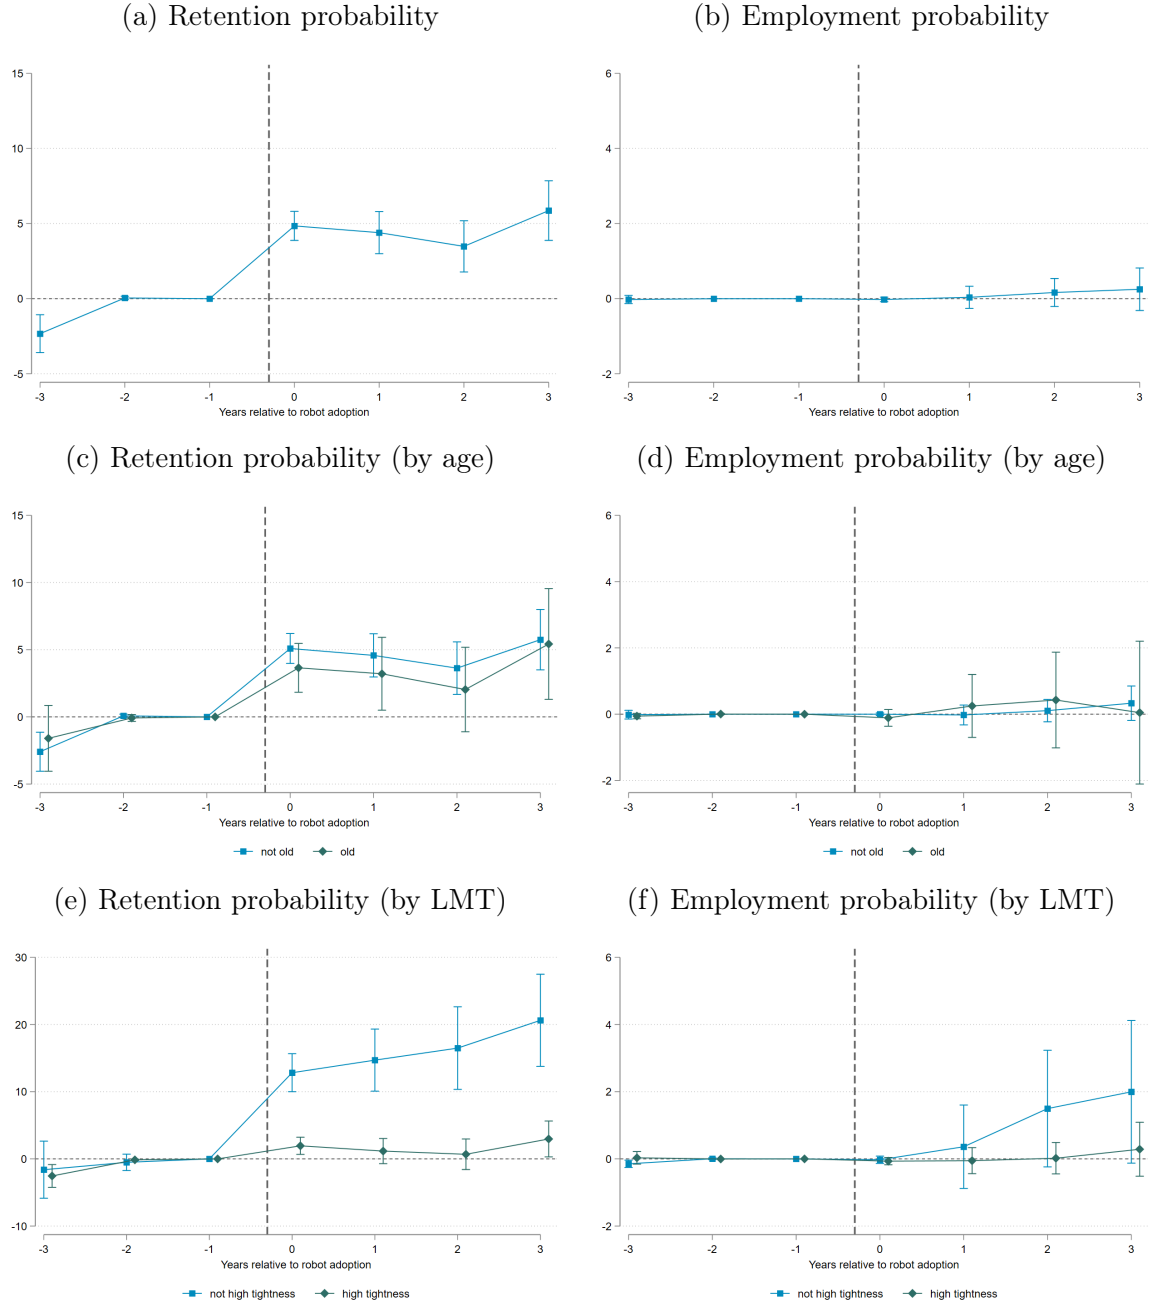
Notes: This table contains variables used for the matching approach described in Section 3. The matching type refers to whether the variable is used to calculate the propensity score for the 1-nearest neighbor matching or the subsequent restriction of matches having identical characteristics, for instance,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

Appendix Table A.2: Summary statistics of matched workers

	Adopters	Non-Adopters	Difference
Age	44.16 (9.86)	44.86 (10.10)	0.70 (0.13)
Female	0.20 (0.40)	0.11 (0.32)	-0.08 (0.00)
No degree	0.06 (0.24)	0.06 (0.23)	-0.01 (0.00)
Vocational degree	0.85 (0.36)	0.85 (0.35)	0.01 (0.00)
College degree	0.08 (0.28)	0.08 (0.28)	0.00 (0.00)
Simple manual	0.49 (0.50)	0.46 (0.50)	-0.03 (0.01)
Qualified manual	0.33 (0.47)	0.35 (0.48)	0.02 (0.01)
High-skilled manual	0.18 (0.38)	0.19 (0.39)	0.01 (0.00)
Share routine tasks	0.53 (0.22)	0.50 (0.20)	-0.03 (0.00)
Share routine manual tasks	0.49 (0.25)	0.45 (0.23)	-0.04 (0.00)
Tenure	13.58 (8.37)	15.00 (8.57)	1.42 (0.11)
Log daily wage	4.67 (0.42)	4.73 (0.45)	0.06 (0.01)
Works council	0.89 (0.32)	0.87 (0.33)	-0.01 (0.00)
Number of employees	535.03 (504.12)	392.69 (499.35)	-142.35 (6.48)
Share of production workers	0.76 (0.11)	0.76 (0.14)	0.00 (0.00)
Observations	9443	16604	26047

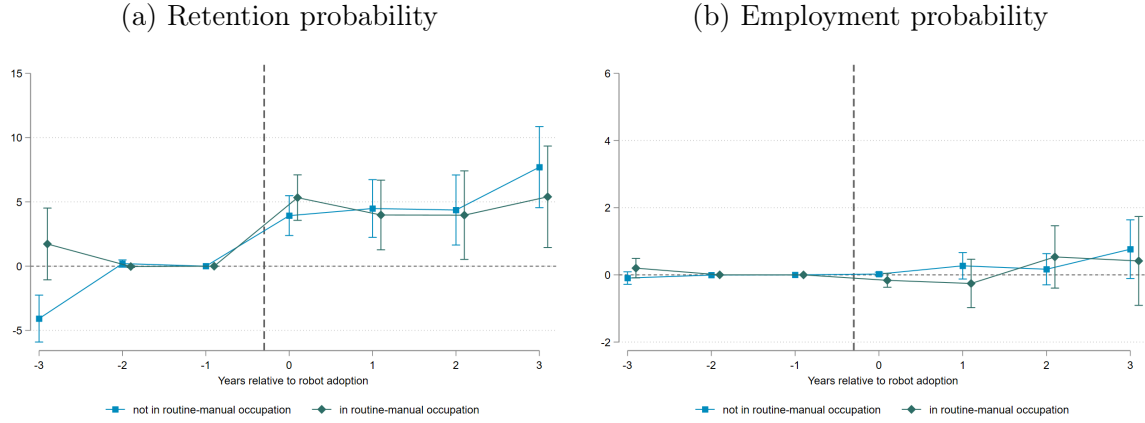
Notes: This table displays means and standard deviations for matched workers in robot adopting and non-robot adopting plants in the year prior to adoption. The last column shows the difference in means and the standard error from a two-sided t-test accounting for unequal variances between the two groups. Production workers are defined as being in Blossfeld occupation 2-5. 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.

Appendix Figure A.1: Employment effect of works councils during automation, entire workforce



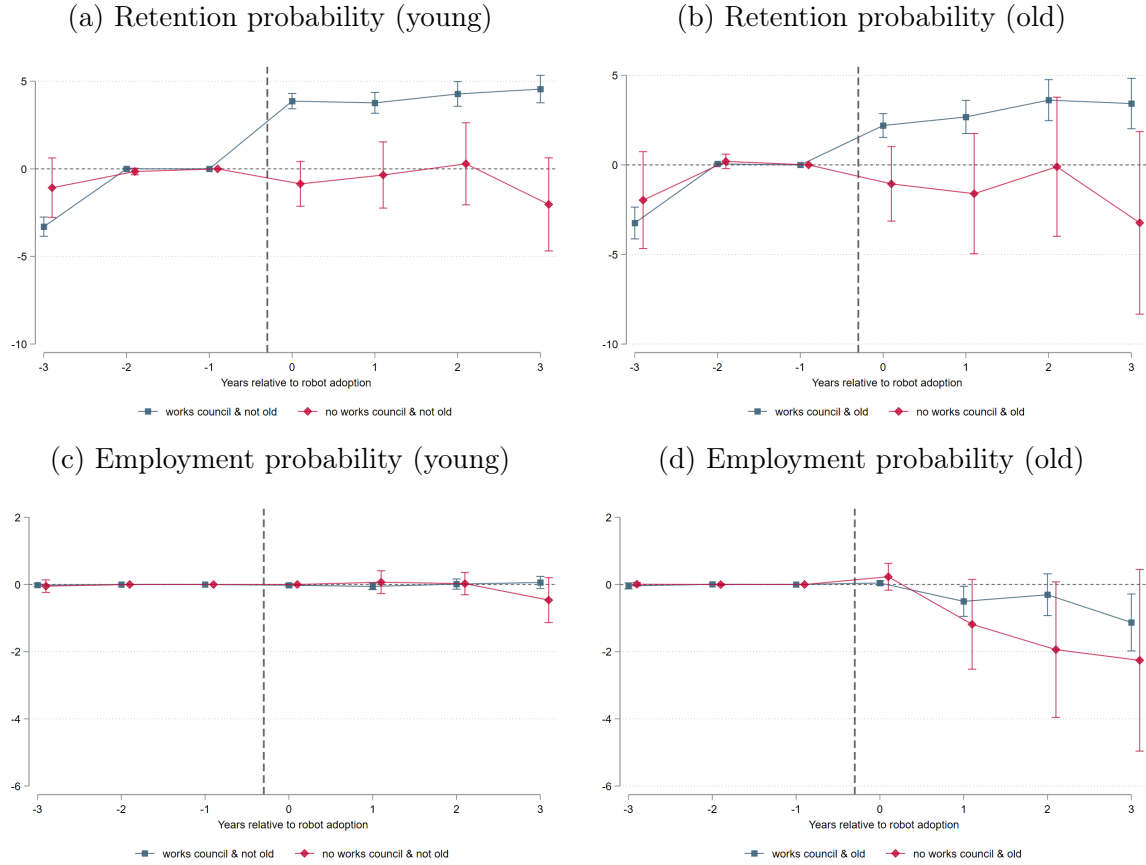
Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the probability of being employed at least one day per calendar year, either at the initial plant (in Panel (a), (c), and (e)) 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 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. 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. 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.

Appendix Figure A.2: Employment effect of works councils during automation (routine-manual vs. other production workers)



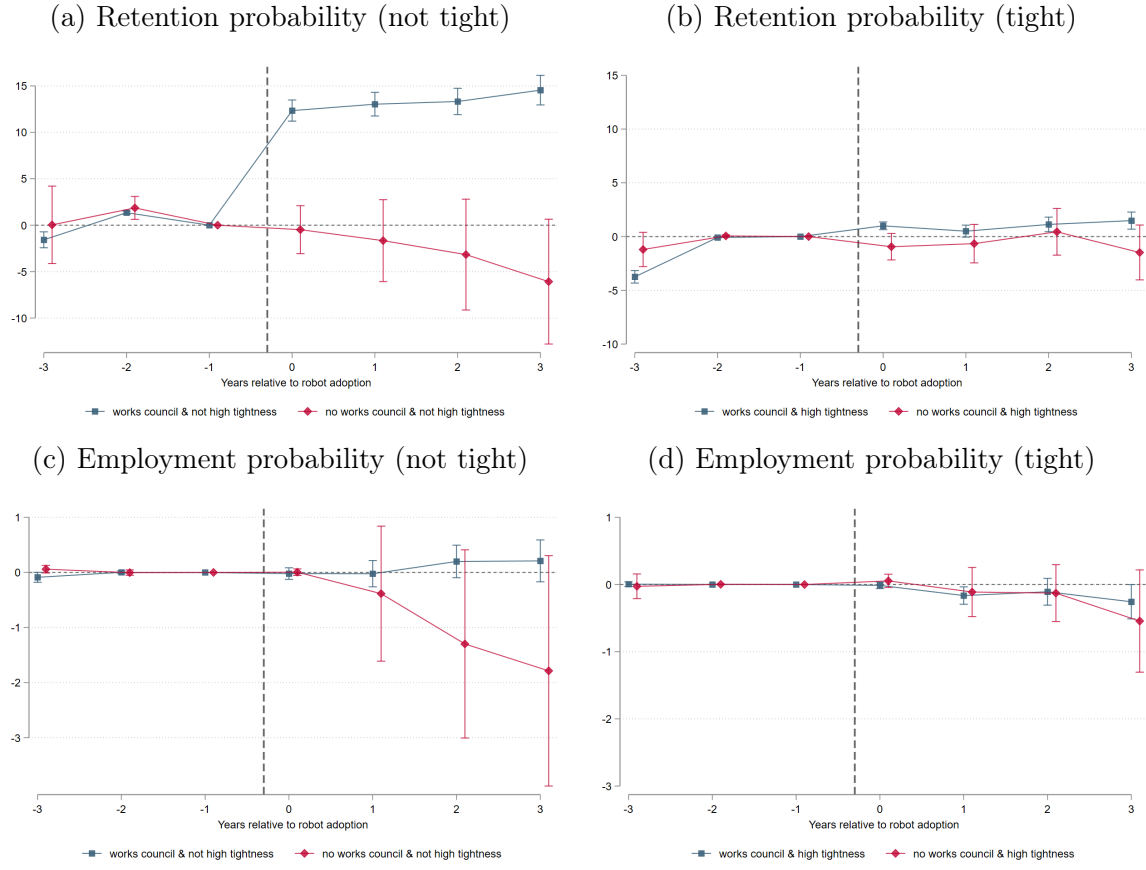
Notes: This figure shows the effect of the presence of a works council during the event of robot adoption in plants on workers' employment outcomes. Employment is measured as the probability of being employed at least one day per calendar year, either at the initial plant (in Panel (a)) or anywhere (in Panel (b)). 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 whether they are in a routine-manual occupation (Blossfeld occupation 2) or not 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.

Appendix Figure A.3: Employment effect of robotization by works council status 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.

Appendix Figure A.4: Employment effect of robotization by works council status by LMT



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.

Appendix Table A.3: Works councils and likelihood and purpose of robot adoption

<i>Panel A: Adoption</i>	Robot adopter	Robot user		
Works council	-0.409 (0.407)	-0.093 (0.579)		
Mean of Y	1.60	3.60		
SD of Y	12.54	18.64		
R-squared	0.05	0.12		
Observations	11888	12150		
<i>Panel B: Purpose</i>	Product improvement	New offering	New product	Process improvement
Works council	-1.819 (9.573)	-2.591 (11.009)	-8.951 (10.303)	-14.721 (11.557)
Mean of Y	67.10	34.19	17.42	49.68
SD of Y	47.14	47.59	38.05	50.16
R-squared	0.23	0.12	0.21	0.16
Observations	148	148	148	148

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 the probability of newly adopting robots (between 2015 and 2018, Column 1) and having robots installed at any point in time (Column 2). The outcomes in Panel B refer to the question of whether the plant introduced a certain type of improvement in the year of robot adoption (answers are not mutually exclusive). 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 Panel A, we restrict the sample to all firms in the 2019 wave of the Establishment Panel with non-missing information, and for Panel B to all first-time robot-adopting plants. Standard errors are robust and clustered at the plant level.

Appendix Table A.4: 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	-0.116 (0.097)	-0.066 (0.090)	-0.058 (0.088)
Mean of Y	0.47	0.56	0.47
SD of Y	0.50	0.50	0.50
R-squared	0.11	0.11	0.18
Observations	287	271	258

Notes: This table shows results from regressions of log employment (Panel A) and the reporting of unfilled vacancies (Panel B) 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 $b \in \{-4, 1; -1, 0; 1, 2\}$. Standard errors are robust and clustered at the plant level.

Appendix Table A.5: 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 $b \in \{-4, 1; -1, 0; 1, 2\}$. Standard errors are robust and clustered at the plant level.

Appendix B: Data Description

Since the IAB Establishment Panel (Bellmann et al., 2021) of the Institute for Employment Research 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.¹¹ Specifically, we use the Integrated Employment Biographies (IEB v16.01.00) prepared by the Institute for Employment Research use¹², which comprises the full universe of all individuals who have held a job subject to social security contributions, marginal employment, or received unemployment benefits. From this data, we identify all individuals who were ever employed in one of the surveyed plants in the 2019 wave of the Establishment Panel. This dataset is similar to the longitudinal model of the Linked Employer-Employee Dataset of the IAB (LIAB LM, DOI: 10.5164/IAB.LIABLM7519.de.en.v1). However, due to size restrictions, the standard LIAB dataset contains only a subset of all plants from the original employment panel, the so-called panel cases 2009-2016, which would further limit the number of observations in the subsequent analyses. 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 the pension insurance. We employ the procedure introduced by Card et al. (2013) to impute those censored wages.

¹¹This 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.

¹²Access to this data is regulated by Section 75 of the German Social Code (Book X).