

The Adjustment of Labor Markets to Robots*

Wolfgang Dauth[†] Sebastian Findeisen[‡] Jens Suedekum[§] Nicole Woessner[¶]

December 3, 2019

Abstract

We use detailed administrative data to study the adjustment mechanisms of labor markets to industrial robots in Germany, employing local labor market shift-share designs. Robot exposure reduces the labor share, increases productivity, and leads to displacement effects in manufacturing which are offset by new jobs in services. The majority of the incidence falls on young workers just entering the labor force. Automation causes more stable employment within firms for incumbents and this is driven by workers taking on new jobs (measured by occupations) in their original plants. Several measures indicate the new jobs are of higher quality than the old ones for incumbents. Young workers substitute away from vocational training towards colleges and universities in their education choices. Finally, we find that automation creates substantial earnings inequality. Industrial robots have benefited workers in occupations with complementary tasks like, e.g. managers or technical scientists.

JEL-Classification: J24, O33, F16, R11

Keywords: Automation, Labor Market Institutions, Skill Upgrading

*We thank Daron Acemoglu for suggestions and very helpful feedback. We are grateful to David Autor, Uwe Blien, Georg Graetz, Xavier Jaravel, Michal Kolesár, Attila Lindner, Guy Michaels, Gianmarco Ottaviano, Pascual Restrepo, Daniel Sturm, Danny Yagan, Fabrizio Zilibotti and many seminar audiences for helpful comments. We thank Hans Ludsteck for answering our data questions. We received financial support from the DFG-priority program 1764 “The German Labour Market in a Globalised World - Challenges through Trade, Technology, and Demographics”.

[†]University of Würzburg and IAB. wolfgang.dauth@uni-wuerzburg.de

[‡]University of Konstanz. sebastian.findeisen@uni-konstanz.de

[§]DICE Heinrich-Heine-Universität Düsseldorf. suedekum@dice.hhu.de

[¶]DICE Heinrich-Heine-Universität Düsseldorf.

1 Introduction

How have new automation and robotics technologies transformed the labor market? Theoretical work on this question has identified two competing high-level effects for employment and wages (Acemoglu and Restrepo, 2018b; Acemoglu and Restrepo, 2019). At first, the adoption of automation technologies will cause a displacement effect, as robots take over in jobs or tasks performed by labor. Sooner or later, however, productivity gains will lead to new jobs elsewhere in the economy. Careful empirical work is now needed to provide evidence on which effect dominates. Furthermore, understanding and examining the mechanisms behind the displacement and productivity effects is crucial for a wide range of questions high on the agenda for economists and policymakers alike. First, displacement can trigger painful adjustment processes and large earnings losses (Jacobson, LaLonde, and Sullivan, 1993), which might imply a bigger role for policy in targeting those displaced by automation technologies. Second, incumbent workers and also young labor market entrants might have to be re-trained or acquire more education in order to be able to transition to new tasks and jobs created by the productivity effect. Finally, a reasonable prior is that different labor market institutions might mediate the displacement/productivity effects very differently, providing potential lessons how to maximize the positive impacts of automation.

In this paper, we examine those mechanisms and how firms and individual workers adjust to automation exposure. The labor replacing technology we focus on are industrial robots, primarily used in the manufacturing sector. Following significant technical advances, robotic capabilities have made great strides in limiting the need for human intervention while autonomously operating production processes. According to the International Federation of Robotics (2016), the stock of industrial robots rose by a factor of five between 1993 and 2015 in North America, Europe, and Asia. An estimated 1.5 million industrial robots are currently used. A large number of industries have thus already undergone dramatic changes in the organization of production in the last two decades.

We use Germany as our "laboratory" and make use of local labor market variation as our main source. It is clear that Germany provides an important benchmark case when it comes to the equilibrium effects of robots and how labor markets adjust to increasing automation. Figure 1 shows the penetration of robots, dividing the stock of robots by the number of workers in different regions of the world between 1994 and 2014. Korea (the world leader), Japan, and Germany are technologically much more advanced in robotics than other countries in Europe and the United States. In addition, to get a solid understanding of the adjustment process and to grasp the incidence of automation, one needs high-quality longitudinal data, following workers over time across firms, occupations, and sectors. For Germany, we can leverage that matched employer-employee data extracted from social security records exist.

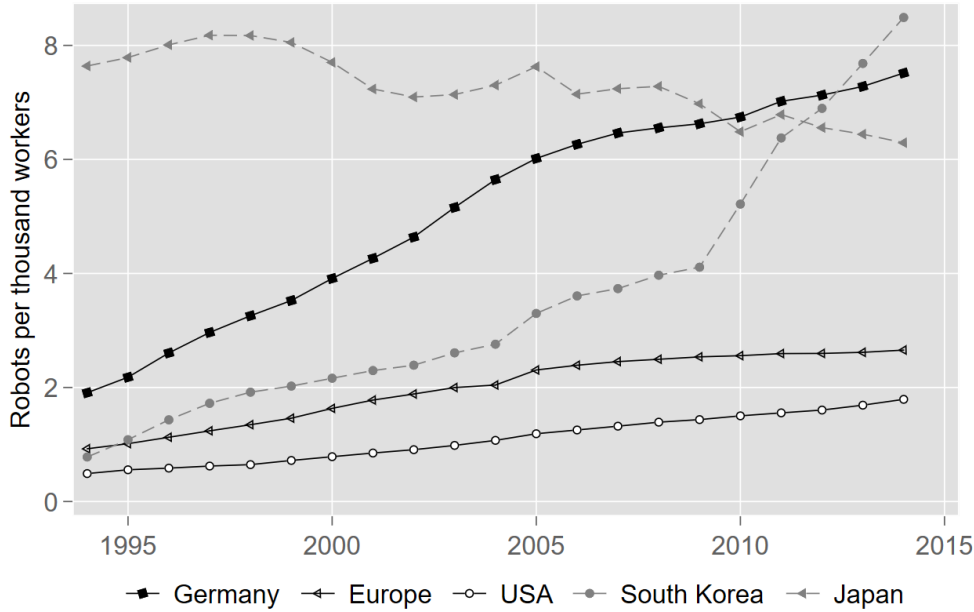


Figure 1: Robot penetration, 1994-2014

Notes: *Europe* = Germany, France, Italy, Spain, Finland, Sweden, Norway, UK. Robot penetration is the robot stock relative to the dependent employment in full-time equivalents (FTE). Employment data from the IAB for Germany and from OECD.Stat for the remaining countries. Dependent employment in Korea and Japan was imputed from total employment and the ratio of dependent to total employment in the European countries, where data on both dependent and total employment is available. Source: IFR, OECD, and BEH V10.01.00, own calculations.

The first part of the paper replicates the strategy of Acemoglu and Restrepo (2019), who have found alarmingly negative impacts on labor demand in the US. We find no such negative effects on total employment but show that this masks the presence of considerable *displacement* and *reallocation* effects. Within manufacturing, robot exposure leads to fewer jobs, but, new labor demand in the service sector – in particular local services used by other businesses – leads to an offsetting force. We present new pieces of evidence using data on productivity and the labor share at the local labor market level. They corroborate the presence of substantial productivity/reallocation effects and displacement effects, manifesting themselves in a lower labor share.¹ We then extend the literature in four ways, which we describe now.

The second main contribution is a complete characterization of the incidence of the displacement and reallocation effects. The main finding is that the majority of the incidence falls on young workers, just entering the labor force. They face lower labor demand in automating industries and adjust by taking jobs in the expanding service sector. Incumbent workers, maybe paradoxically at first glance, actually see an increase in their plant tenure in response to automation.

Our third main contribution shows that this last effect – i.e. automation causing more stable employment within firms – is driven by workers taking on new jobs (measured by

¹Autor and Salomons (2018) show in a cross-country and cross-industry design that TFP growth is associated with a reduction in the labor share, consistent with these findings.

occupations) in their original plants. Displacement from old tasks, hence, takes place, but is (immediately) offset by taking on new tasks for the same employer. There are three measures we use, which indicate the new jobs are of higher quality than the old ones for incumbents: the new occupations pay higher wages, are characterized by a higher share of abstract instead of routine tasks, and the college share is higher. Young workers in local labor markets with more exposure to automation substitute away from vocational training towards colleges and universities.

Fourth, we divide local labor markets, using the vote share of the German Social Democratic Party as a proxy for the strength of labor market institutions favoring labor. We work under the hypothesis that stronger labor market protections for workers should manifest themselves in smaller displacement effects, but also reduced productivity gains from automation. Results from sample splits along the (imperfect) measure support this hypothesis. Robots reduce the labor share and increase productivity in regions with weak worker protections.

In the fifth and last contribution, we shift the focus from local labor market adjustments to the individual worker. This complements the previous models because it allows to directly study the effects of automation on earnings and wages using a more compelling design. Comparing wage or earnings growth across local labor markets, in contrast, can lead to biased results because automation changes the composition of employed workers. At the individual level, we can follow the same workers, who start competing with industrial robots, over time and across all possible margins of adjustments (plants, occupations, sectors). One result is that average earnings are hardly affected. However, we find that automation creates substantial earnings inequality: workers who are retained by their plants, experienced positive earnings effects. Workers which are forced to switch plants, industries, or leave manufacturing see significant earnings losses. Finally, we show how industrial robots have benefited workers in occupations with complementary tasks like, e.g. managers or technical and natural scientists, while hurting those in routine intensive tasks like machine operators.

The characterizing attribute of automation and robotics is the high substitutability with human labor in some tasks. The theoretical implications of automation for wages, employment, productivity, and other outcomes have been studied by Acemoglu and Restrepo (2018b), Acemoglu and Restrepo (2019), and Moll, Restrepo, and Rachel (2019).² The important empirical paper by Acemoglu and Restrepo (2019) has documented strong negative effects for wages and employment across US commuting zones, implying a strong displacement force and productivity effects, which do not seem to lead to more employment or wage increases elsewhere. Our point of departure is the adoption of the empirical strategy developed by Acemoglu and Restrepo (2019). Quantitatively, we also find significant displacement effects, although around 50% smaller on average. The key difference here is that we additionally

²This active literature builds on older papers, which highlighted the usefulness of the task framework to explain a variety of empirical phenomena concerning the distribution of wages and employment – see Acemoglu and Autor (2011) for an exhaustive survey.

identify significant and offsetting reallocation effects. Concerning the displacement effect, we only find a negative impact of robot exposure on the labor share and labor demand in local labor markets with weaker protections for labor (as measured by the strength of workers' rights parties). This hints at the potential importance of labor market institutions in explaining differences in the strength of displacement also across countries.

The use of data assembled by the International Federation of Robotics (IFR) originates with the innovative paper by Graetz and Michaels (2018). Consistent with our results, they uncover positive productivity effects and zero effects on total employment, using variation in robot usage across industries in different countries. However, as our analysis shows, the zero employment effect can mask substantial displacement and reallocation effects. We complement Graetz and Michaels (2018) (and also Acemoglu and Restrepo, 2019), by leveraging administrative labor market data as the first study. We can, therefore, study the mechanisms how automation affects labor markets; in particular, if workers separate from firms, how the set of tasks carried out by exposed workers evolves in response to automation, and what role the transitions of individual workers across industries and sectors play.

An important part of the adjustment process to automation and robots is the skill upgrading process, as our evidence shows. Changes in the demand for high-skilled workers also feature prominently role in the polarization literature (Michaels, Natraj, and van Reenen, 2014; Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014). We document direct and indirect evidence for two margins of human capital adjustments to automation: within the plant for workers who are retained in the face of automation and for young labor market entrants. The first channel of within firm upgrading is consistent with the famous plant-level study by Bartel, Ichniowski, and Shaw (2007) on American valve-makers. They chronicle how the adoption of new IT-enhanced capital equipment leads to increases in the skill requirements of machine operators and a transition from routine to abstract/cognitive tasks.³ Finally, our analysis reveals that the reallocation effect is driven by increased employment in the business service sector, showing that the spillovers seem to operate locally through firms expanding their demand for complementary tasks. Relatedly, Helm (2019) also finds positive local spillovers of export shocks across German labor markets, consistent with agglomeration economies.⁴

Section 2 describes our empirical approach and the data. Section 3 studies the impact of robots on equilibrium employment, the labor share, and productivity across local labor markets. Sections 4 and 5 investigate adjustment mechanisms. Section 6 studies the adjustment process of individual workers. Section 7 concludes.

³Notably, the plants in the study accompanied the transition process with the adoption of new human resource practices to support these skills.

⁴While automation creates productivity growth – see our evidence on this in Section 3 – it also coincides with displacement. This presumably counters agglomeration forces and explains why one should expect different effects on agglomeration of automation than for positive trade shocks.

2 Data and Methodology

2.1 Administrative Labor Market Data

Our main source is administrative German labor market data provided by the Institute for Employment Research (IAB) at the German Federal Employment Agency. Specifically, we use data from the Employee History (Beschäftigtenhistorik – BEH, Version V10.01.00). The raw version of the BEH is a spell-dataset of the complete job histories of the universe of private workers from 1978 to 2014, excluding the self employed and sworn civil servants. Eastern Germany enters the data in 1992. We use a simplified version of this dataset that contains only one observation for each individual and year, pertaining to the spell of the highest paid job that stretches over June 30th of a given year. We also drop observations on marginal jobs since those are only included in the data from 1999 onwards. The individual level information contains information on gender, year of birth, educational attainment, a unique plant-id, as well as codes for industries workplace locations, and occupations.⁵ This allows us to aggregate the dataset to the county level and obtain a precise picture on the size, the industry composition, and the workforce characteristics of local labor markets. Moreover, the worker level panel structure of the dataset allows us to observe the mobility patterns of individuals as they enter the labor market, move between jobs, firms, industries, and regions and finally exit the labor market. As described below, we mainly work at the local labor market level. Our main outcome is the percentage change in a county’s employment. We construct this from the aggregate worker counts on June 30 of the start year 1994 and end year 2014, where part-time workers are weighted by 0.5 to get a measure for full-time equivalent employment. The information on the industry of the workplace plant allows us to construct this variable separately for the manufacturing and non-manufacturing sectors. The advantage of using percentage changes rather than the log-difference is that this growth rate can be additively decomposed into the contributions of various groups defined by worker mobility, such as workers who enter the labor market, who stay with their original plant, who move to a different plant in the same industry, etc.

Our second outcome variable is the log change in average wages. To construct this variable, we first impute the individual wages, which are censored at the social security contribution ceiling, using a procedure suggested by Card, Heining, and Kline (2013). We then compute the average daily wage for full-time workers on June 30 of the start and end year for demographic cells defined by gender, three age groups (below 30, 30 to 44, 45 or more years old), and three education groups (no or unknown degree, vocational training degree, university degree). A further dependent variable is the total yearly wagebill, which is 365 times the individual daily wage, aggregated to the county level. To compare our results to the findings of Acemoglu and Restrepo (2019), we also construct the change in the employment to population ratio as a

⁵We distinguish between 102 2/3 digit NACE Rev. 2 industries, 402 counties, and 54 occupational fields.

further outcome variable. The employment numbers again stem from the aggregate BEH data, while population counts stem from the German Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR). Two final outcome variables stem from the German Federal Statistical Office and relate to the productivity of the regional economy. These are the log change in GDP per worker and the percentage point change in the total employee compensation (Arbeitnehmerentgelte) relative to total GDP. Since this data is only available from 2000 onwards, these two measures are constructed as changes between 2000 and 2014.

2.2 Robot Usage

We combine our administrative labor market data set with data on the stock of robots for 25 industries in 50 countries over the period from 1994 to 2014 from the International Federation of Robotics (IFR). This data set has been used before by Graetz and Michaels (2018) in a cross-country study at the industry level and by Acemoglu and Restrepo (2019) for the US. A *robot* in this data is defined as an “automatically controlled, re-programmable, and multipurpose machine”. As explained in more detail in International Federation of Robotics (2016), this means that robots are “fully autonomous machines that do not need a human operator and that can be programmed to perform several manual tasks such as welding, painting, assembling, handling materials, or packaging.” Single-purpose machines such as elevators or transportation bands are, by contrast, not robots in this definition, as they cannot be reprogrammed to perform other tasks, require a human operator, or both. These data are based on yearly surveys of robot suppliers and capture around 90 % of the world market. The information is broken down at the industry level.⁶ The industry classification of this data conforms to 2-digit ISIC Rev. 4 codes, where 3-digit information is available for manufacturing of electronic devices, electrical equipment, and motor vehicles. Since our administrative data has time-consistent NACE Rev. 2 industry codes, which correspond to the ISIC Rev. 4 codes at the 2/3-digit level, both datasets can be matched without using any further crosswalk.⁷

The 25 industries consist of 20 manufacturing industries, agriculture, mining, supply, construction, and education. Appendix Figure A.1 illustrates the change in the number of robots per thousand workers in all 25 industries, and as an interesting comparison, we also present the US numbers here. By far the strongest increase can be observed in the different branches of the automobile industry (motor vehicles, auto bodies and parts). Here, more than 100 additional robots were installed per thousand workers in 2014 compared to 1994. Other industries

⁶But data availability differs across countries but coverage is comprehensive for Germany. As Graetz and Michaels (2018), we do not use the IFR industries *all other manufacturing*, *all other non-manufacturing*, and *unspecified*. Those categories cover less than 5% of the total robot stock in Germany.

⁷Data used for a previous version of the paper (Dauth, Findeisen, Suedekum, and Wößner, 2018) only had time consistent NACE Rev. 1 codes. This required us to construct a crosswalk from the IFR classification to the classification of the labor market data, where we apportioned ambiguous cases according to employment shares. The results were qualitatively very similar.

that became vastly more robot-intensive include rubber and plastic products, electronic components, and domestic appliances. On the other side of the spectrum we find cases where robot usage has hardly changed, and sometimes (e.g., in manufacturing of instruments for measuring) it even decreased over time. In non-manufacturing industries, robots are used much less than in manufacturing.

2.3 Local Labor Market Approach

In this paper, we are interested in how the labor market and its main actors, firms and workers, adjust to increasing automation possibilities. Our research design is motivated by the important paper on the impact of robots on the US labor market by Acemoglu and Restrepo (2019), who also provide a theoretical micro foundation. This empirical approach bases on the fact that a country’s local labor markets differ markedly with respect to their industry composition. These differences create differences in the exposure to technological change in the form of industrial robots.⁸ The regional perspective allows us to observe equilibrium adjustments and spillovers from directly affected to indirectly affected industries.⁹

Following Acemoglu and Restrepo (2019), our main variable of interest, the change in robot exposure in region r , is constructed as:

$$\Delta\text{robots}_r = \sum_{j=1}^J \left(\frac{\text{emp}_{jr}}{\text{emp}_r} \times \frac{\Delta\text{robots}_j}{\text{emp}_j} \right) \quad \text{with } J = 25. \quad (1)$$

The term $\Delta\text{robots}_j = \frac{\Delta\text{robots}_j}{\text{emp}_j}$ measures the national industry robot adoption as the increase in the robot count in industry j relative to its workforce size in the base year 1994. We allocate this industry-level exposure to regions according to their shares of national industry employment by multiplying Δrobots_j with emp_{jr} , which is the initial employment in industry-region cell jr . For each local labor market r , we sum the exposures of all local industries and scale it by the region’s total employment emp_r , also measured in the base year 1994.

This measure is a typical shift-share variable where an industry-level shock is apportioned across regions. In a recent paper, Adão, Kolesár, and Morales (2019) point out that such an explanatory variable can cause problems with statistical inference: Regions with similar industry structures are likely to have correlated error terms, which means that conventional standard errors may be underestimated. Adão, Kolesár, and Morales (2019) propose to account for this by calculating standard errors in a fashion similar to cluster-robust standard errors, where the correlation structure of the error terms is represented by a matrix of regional

⁸Faber (2019) extends this approach and regresses employment changes in Mexican labor markets on an adjusted measure of exposure to robots adopted in other countries, US robots in his study.

⁹As is widely discussed in the literature, regional difference-in-difference designs have well-known limitations when it comes to gauging absolute or national impacts. But, relative to other structural approaches, the design offers more transparent and clearer identification. The results from various strands of literature show that many equilibrium adjustments take indeed place at the local labor market level (Moretti, 2011).

industry shares rather than by discrete clusters. We adopt their construction of robust standard errors and also apply their adjustment for small industry numbers by imposing the null hypothesis of the true coefficient being zero.¹⁰

The identification of the effects of robots on the labor market builds on the assumption that differences in robot exposure across industries are generated because robots have become better utilizable in some industries than in others. However, the pattern of robot exposure in Germany may be the result of domestic industry-specific demand shocks in Germany. To address this endogeneity concern, we also apply the instrumental variable strategy proposed by Acemoglu and Restrepo (2019). In this approach, we employ robot adoptions across industries in other high-income countries as an instrument for German robot exposure.¹¹ More specifically, we construct the instrumental variables analogously to equation 1 but use the increases in the robot count in the same set of industries in each country and use employment counts from 1984 for normalization and apportioning across regions.¹² Figure 2 summarizes our empirical approach. The horizontal axis shows the variation of the regional robot exposure, conditional on regional employment shares in nine broad industry groups and federal state dummies. The highest ranking regions are Wolfsburg, Dingolfing-Landau, and Ingolstadt, which are heavily concentrated in the automotive industry (Volkswagen, Audi, and BMW produce there, respectively). In our empirical analysis we will pay attention to the special role of the automobile industry in robustness checks. But also aside from those extremes, the variation across regions is substantial. There is no positive relation with employment growth. In our empirical analysis in Section 3, we discuss this result in more detail.

2.4 Descriptive Overview

Table 1 provides a descriptive overview over the data we use for the region level regressions. The average region saw a slight decline in employment. When weighting by the number of full-time equivalent jobs in 1994, this decline becomes sharper, which demonstrates that larger regions declined more strongly.¹³ The overall decline stems mostly from the declining manufacturing sector, which has not been compensated by growth of non-manufacturing industries.

¹⁰The exact procedure is laid out in Remarks 5 and 6 in Adão, Kolesár, and Morales (2019). We thank Michal Kolesár for very valuable advice how to adapt their standard error adjustment for the overidentified IV case.

¹¹See Autor, Dorn, and Hanson (2013) and Bloom, Draca, and van Reenen (2016) for similar approaches to study the effects of Chinese import competition. The validity of this approach hinges on the assumption that the industry pattern of robot adoption is an exogenous shock, while the allocation of industries across regions may be endogenous (see Borusyak, Hull, and Jaravel, 2018, for technical details). One way to verify this is to rule out that there is a correlation of pre-trends in the outcome variables and current robot adoption, which we do in Appendix Table A.2.

¹²We construct one instrument for each country $k = (\text{Spain, Finland, France, Italy, Norway, Sweden, and UK})$ and estimate an over-identified model. In a further robustness check, we also aggregate the robot exposures of all k countries to build a single instrument in a just identified 2SLS model. Notice that it is not possible to use time lags for East German regions; here we are confined to use 1994 in the deflator.

¹³Note that this picture changes when part-time jobs are not weighted down. In this case, the growth rate of the total number of jobs is positive.

ment, respectively, to construct two further shift-share variables, which both have positive averages.¹⁴

Finally, we report the means and deciles of the measure of robot exposure in Panel C. In the average region, the number of robots has increased by around 4.6 robots per 1000 workers. However, as shown in Figure 2, the distribution is skewed to the right, with a handful of very large values.

Table 1: Summary statistics, region level 1994-2014.

observations	unweighted		weighted	
	402		23,884,076	
	mean	(sd)	mean	(sd)
[A] Outcomes				
% change in total employment	-0.145	(18.181)	-1.860	(16.080)
% change in manuf. employment	-9.312	(25.517)	-16.408	(23.744)
% change in non-manuf. employment	5.872	(22.738)	5.031	(21.041)
100 x ln-change in average wage	32.640	(10.022)	32.751	(9.468)
100 x ln-change in average wage, manuf.	40.041	(15.728)	40.512	(14.234)
100 x ln-change in average wage, non-manuf.	28.896	(11.715)	29.516	(11.134)
100 x ln-change in total wagebill	37.968	(18.856)	37.310	(16.921)
100 x ln-change in total wagebill, manuf.	33.462	(32.570)	26.726	(31.537)
100 x ln-change in total wagebill, non-manuf.	39.232	(20.952)	40.286	(19.783)
%-point change in emp/pop-ratio	-0.116	(3.656)	-0.783	(3.537)
%-point change in emp/pop-ratio, manuf.	-0.816	(2.332)	-1.378	(2.284)
%-point change in emp/pop-ratio, non-manuf.	0.700	(3.310)	0.595	(3.320)
100 x ln-change in GDP per worker	46.529	(21.149)	43.455	(19.419)
%-point change in labor share	-12.754	(6.893)	-11.155	(6.560)
[B] Control variables				
% female	34.715	(4.674)	35.155	(4.706)
% foreign	6.981	(4.782)	8.071	(5.147)
% age \geq 50 years	20.101	(2.366)	21.192	(2.450)
% unskilled	11.063	(4.435)	10.794	(4.218)
% vocational training	80.296	(4.117)	78.220	(4.851)
% university degree	7.956	(3.965)	10.154	(4.592)
% manufacturing	30.473	(12.559)	27.772	(12.881)
% food products	3.443	(2.076)	2.814	(1.752)
% consumer goods	4.609	(4.012)	3.876	(3.494)
% industrial goods	11.846	(7.516)	10.491	(7.725)

¹⁴For the measurement of trade exposure, we closely follow Dauth, Findeisen, and Suedekum (2017) and Dauth, Findeisen, and Suedekum (2019), who compute the increase in German net exports vis-à-vis China and 21 Eastern European countries over the period 1994-2014 for every manufacturing industry j using UN Comtrade data, normalized by the initial wage bill to account for industry size. For ICT, we exploit information about installed equipment at the industry level as provided in the EU KLEMS database. It is defined as the change in real gross fixed capital formation volume per worker for computing and communications equipment from 1994 to 2014.

Table 1: Summary statistics, region level (continued).

	unweighted		weighted	
	mean	(sd)	mean	(sd)
% capital goods	11.048	(8.733)	11.068	(8.315)
% construction	13.562	(4.717)	12.514	(4.773)
% maintenance; hotel & catering	19.231	(4.469)	19.594	(4.193)
% services	14.186	(5.271)	17.908	(7.864)
% public sector	19.913	(6.397)	19.963	(6.312)
dummy, 1=north	0.159	(0.366)	0.149	(0.357)
dummy, 1=south	0.348	(0.477)	0.282	(0.451)
dummy, 1=east	0.192	(0.394)	0.230	(0.421)
Δ net exports in 1000 € per worker	0.956	(3.146)	1.002	(2.758)
Δ ICT equipment in € per worker	661.942	(157.081)	733.603	(185.298)
[C] Exposure to robots				
Δ robots per 1000 workers	4.617	(8.028)	4.642	(7.808)
p10-p90 interval	[0.982 ; 8.527]		[0.982 ; 8.527]	
p25-p75 interval	[1.438 ; 4.540]		[1.394 ; 5.108]	

Notes: Summary statistics of region level variables. In Columns 2 and 3, the data is weighted by full-time-equivalent number of jobs in 1994. Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

2.5 Regressions Models

In Sections 3 to 5 we estimate models of the following form at the local labor market level:

$$\Delta Y_r = \alpha \cdot \mathbf{x}'_r + \beta_1 \cdot \Delta \text{robots}_r + \beta_2 \cdot \Delta \text{trade}_r + \beta_3 \cdot \Delta \text{ICT}_r + \phi_{REG(r)} + \epsilon_r. \quad (2)$$

We regress a change – sometimes a percentage change – of an outcome variable, such as total employment, manufacturing employment, the employment-to-population ratio, or the labor share, over the period 1994-2014, on the change in the number of robots per worker (i.e., on Δrobots_r as defined in (1)). In the vector \mathbf{x}'_r we control for detailed demographic characteristics of the local workforce (such as age, gender, and qualification) in levels, aggregated up from the universe of individual social security records. To avoid contamination by the endogenous adjustment of the local labor force after the shock, we use levels before the start of the periods and not changes. We also include controls for the employment shares of nine broadly defined industry groups. Moreover, we add four broad region dummies, and we add the local exposures to net exports and ICT specifications.

As discussed above in Section 2.3, for inference we apply the method proposed by Adão, Kolesár, and Morales (2019). In the tables, we label them shift-share standard errors. We

additionally present standard errors which are calculated in a conventional way, using 50 clusters which represent a higher geographical aggregation of local labor markets.¹⁵ On average the shift-share standard errors are larger and make the inference more conservative.

3 Employment Effects, the Labor Share, and Productivity

At the very start of our analysis, we begin with results for employment, following Acemoglu and Restrepo (2019). We then separate the effects on employment, wage, wage bills, and employment-to-population ratios across sectors. The end of the section investigates their robustness.

3.1 Employment Effects

We first look at employment changes in percentage terms in Table 2 using OLS regressions in Panel A.¹⁶ We include a separate row for the shift-share standard errors using the method proposed by Adão, Kolesár, and Morales (2019). Standard errors calculated in the more conventional way allowing for 50 regional clusters are presented in parentheses below the estimates.

Column 1 presents a parsimonious specification (with the initial manufacturing and regional dummies as the only additional control variables). The estimated effect is positive but very small and statistically insignificant. Quantitatively, comparing a local labor market at the 75th percentile of exposure to a local labor market at the 25th percentile, the magnitudes imply that the 75th percentile labor market experiences 0.155% points ($[4.540 - 1.438] \times 0.05 = 0.155$) higher employment growth or around 100 more (full-time equivalent) jobs for an average region.

The estimates stay small and statistically insignificant as we enrich the specifications. First, as there may be more fine-grained industry trends within the manufacturing sector, which are correlated with employment outcomes and robot installations, we next include the initial employment shares of nine industry groups instead of the overall manufacturing share. The coefficient in column 2 changes signs but importantly stays close to zero.

Column 3 adds the trade exposure of local labor markets, using exports and imports with Eastern Europe and China, as described in Section 2.4.¹⁷ Column 4 additionally includes exposure to ICT investments. The inclusion of both variables clearly has a visible effect on

¹⁵These 50 clusters are highly aggregated labor market regions defined for use in German regional policy. Most economic interactions should be confined to those areas.

¹⁶Using changes in log employment yields very similar results. We prefer the percentage changes since they allow for a clean additive decomposition into various channels. This will be the focus of Section 4. As an additional employment measure, we will also use the employment-to-population ratio in the next subsection.

¹⁷As is well known, Germany is a very export-oriented economy. If export intensive industries also rely more heavily on robots, this might alleviate possible job losses from technological change. Conversely, robots might have lowered production costs and thus spurred demand for German products.

Table 2: Robot Exposure and Employment.

	Dependent variable:			
	% change in total employment between 1994 and 2014			
	(1)	(2)	(3)	(4)
[A] OLS				
Δ robots per 1000 workers	0.0530 (0.106) [0.089]	-0.0396 (0.125) [0.121]	-0.0665 (0.121) [0.124]	0.0858 (0.122) [0.167]
% manufacturing	-0.1107 (0.173)			
% food products		2.5142 (0.480)	2.4864 (0.480)	2.4233 (0.458)
% consumer goods		0.4912 (0.329)	0.5671 (0.335)	0.5389 (0.321)
% industrial goods		0.5823 (0.292)	0.5529 (0.299)	0.5457 (0.280)
% capital goods		0.9398 (0.287)	0.9059 (0.296)	0.9126 (0.278)
% construction		1.0598 (0.312)	1.0437 (0.317)	1.0624 (0.297)
% consumer services		1.5169 (0.372)	1.5114 (0.377)	1.6461 (0.365)
% business services		0.4782 (0.306)	0.4730 (0.308)	0.8484 (0.280)
% public sector		0.9281 (0.283)	0.9206 (0.286)	1.1057 (0.273)
Δ net exports in 1000 € per worker			0.3735 (0.230)	0.3593 (0.227)
Δ ICT equip. in 1000 € per worker				-0.0251 (0.007)
R ²	0.506	0.570	0.573	0.586
[B] 2SLS				
Δ robots per 1000 workers	0.0660 (0.105) [0.085]	-0.0559 (0.132) [0.139]	-0.0811 (0.129) [0.141]	0.0675 (0.137) [0.181]
Kleibergen-Paap weak ID test	562.668	391.407	383.098	378.041
Hansen J p-value	0.432	0.223	0.217	0.203

Notes: $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*). Regressions of total employment growth (in %) on the change in robot exposure between 1994 and 2014. All specifications include a constant, broad region dummies indicating if the region is located in the north, west, south, or east of Germany and demographic control variables, measured in the base year 1994. The demographic control variables are the employment shares of female, foreign, age ≥ 50 , medium skilled (with completed apprenticeship), and high skilled (with a university-degree) workers relative to total employment (reference category: unskilled workers and with unknown education). In column 1, we control for the manufacturing share in total employment. In columns 2-4, we instead include broad industry shares to control better for regional industry patterns. Industry shares cover the percentage of workers in nine broad industry groups (agriculture (reference); food products; consumer goods; industrial goods; capital goods; construction; consumer related services; business related services; public sector) in the base year 1994. Columns 3 and 4 successively take into account the change in German net exports vis-à-vis China and 21 Eastern European countries (in 1000 € per worker), and the change in ICT equipment (in € per worker), both between 1994 and 2014. Panel B reports results of a two-stage least squares (2SLS) IV approach where German robot exposure is instrumented with robot installations across industries in other high-income countries. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, Comtrade, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

the main coefficient, moving its magnitude by around 0.03 and 0.12 points. However, the main results remain unaffected and the coefficient estimates imply only small employment effects of automation.

Panel B shows the results when the regressions are estimated with 2SLS. First, across the different specifications, the 2SLS estimates are close to their OLS counterparts. The Kleibergen and Paap (2006) statistic indicates there is no problem of weak instruments and the Hansen test values imply no rejection of the null hypothesis of valid instruments. Given, however, that the OLS and 2SLS estimates are so similar because of the closely correlated patterns of robot adoption across countries, we will put our focus for the rest of the paper on the simpler OLS strategy and present the instrumental variable estimates in the appendix for the most part.

3.2 Displacement versus Reallocation

We next study the displacement and reallocation/productivity effects of automation separately. To analyze decomposition effects, we opt for (arguably) the most transparent cut of the data. In particular, the displacement of workers should occur within the robot adopting manufacturing sector. At the same time, the demand for labor in all other local industries increases when industries are gross complements in the production of a final consumption good. A plausible hypothesis is that industries in the service sector should see an increase in labor demand.

Panel A in Table 3 presents the results for employment changes.¹⁸ Column 1 repeats the main estimate from Table 2 column 4, which was the fully specified model with the most control variables. The models in columns 2 to 4 use the (percentage) change in manufacturing as the outcome variable. Column 2 has the same control variables as column 2 of Table 3, namely broad industry employment shares and regional dummies. The next columns add trade and ICT exposure, respectively. The estimates in all three columns show a negative coefficient and, importantly, the effect size is around one order of magnitude larger than the effects on total employment. Columns 5 to 7 investigate the impact on employment in the service sector. The positive coefficients reveal the presence of substantial reallocation forces, offsetting the adverse impact of the displacement effects. Approximately, displacement and reallocation effects tend to be of similar magnitudes, which explains the robust finding of a zero total employment effect.

We re-estimate the models in Panel B with the change in the employment-to-population ratio as the dependent variable.¹⁹ Column 1 shows an effect close to zero. But again this

¹⁸Here we focus on the OLS results for brevity. The very similar 2SLS estimates can be found in Appendix Table A.1.

¹⁹We measure employment by all jobs in Germany subject to social security. This yields small E/POP ratios between 0.25 and 0.5 in our sample since we have excluded civil servants and self-employed workers. Including civil servants and self-employed workers in the E/POP with data from the German Federal Statistical Office

Table 3: Composition Effects

	Total		Manufacturing		Non-manufacturing		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
[A] Employment: % change in total employment between 1994 and 2014							
Δ robots per 1000 workers	0.0858 (0.122) [0.167]	-0.5690 (0.177) [0.304]	-0.6055 (0.173) [0.304]	-0.4467 (0.160) [0.281]	0.5599 (0.275) [0.382]	0.5412 (0.272) [0.390]	0.7051 (0.275) [0.442]
[B] E/Pop: 100 x Δ in employment/population between 1994 and 2014							
Δ robots per 1000 workers	0.0091 (0.060) [0.031]	-0.0449 (0.028) [0.033]	-0.0501 (0.029) [0.033]	-0.0423 (0.030) [0.029]	0.0405 (0.039) [0.036]	0.0399 (0.039) [0.037]	0.0514 (0.038) [0.043]
Effect of 1 robot	0.3	-1.6	-1.8	-1.5	1.4	1.4	1.8
[C] Wages: 100 x Log- Δ in average wage between 1994 and 2014							
Δ robots per 1000 workers	-0.0345 (0.049) [0.029]	-0.1367 (0.049) [0.072]	-0.1450 (0.050) [0.073]	-0.0960 (0.060) [0.070]	0.0894 (0.037) [0.061]	0.0809 (0.038) [0.060]	0.0904 (0.037) [0.063]
[D] Wagebill: 100 x Log- Δ in total wagebill on June 30							
Δ robots per 1000 workers	0.0851 (0.125) [0.197]	-0.6017 (0.177) [0.330]	-0.6568 (0.175) [0.331]	-0.4509 (0.182) [0.328]	0.3763 (0.212) [0.287]	0.3524 (0.208) [0.293]	0.5180 (0.211) [0.354]
Δ net exports in 1000 € per worker	Yes	No	Yes	Yes	No	Yes	Yes
Δ ICT equipment in € per worker	Yes	No	No	Yes	No	No	Yes

Notes: In all regressions, the variable of interest is the change in robot exposure between 1994 and 2014. The estimates in panels A, B, and D are based $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*), while the unit of observation in the wage estimates in panel (C) are $N = 7,236$ region x demographic cells. Demographic cells are defined by gender, three age groups, and three education groups. We only include cells containing at least 10 observations, and perform the regressions at the region x demographic cell level including fixed effects for demographic cells. The dependent variable in Panel D is the log-difference total amount of gross salaries paid to employees subject to social security on June 30 in 1994 and 2014. All specifications include a constant, broad region dummies, demographic control variables, and employment shares of nine aggregate industry groups, measured in the base year 1994. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, Comtrade, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

hides significant displacement in columns 2 to 4 and strong reallocation effects in columns 5 to 7. Since the sum of the employment-to-population ratios in the two sectors equals the total employment-to-population ratio in a region, the coefficients of the fully specified models in 4 and 7 sum up to 1. We can translate these numbers into head counts.²⁰ This makes the does not affect our results. See also column 6 of Appendix Table A.4, which shows no effect of robots on public employment.

²⁰If we have two time periods, E_t is job head counts in t , R installed robots, and Pop population, then:

$$\frac{E_2}{Pop_2} - \frac{E_1}{Pop_1} = \beta \left(\frac{R_2 - R_1}{E_1} \right) \times 1000.$$

If we assume a constant population, we get:

$$E_2 - E_1 = \beta \left(\frac{R_2 - R_1}{E_1/Pop_1} \right).$$

estimates directly comparable to Acemoglu and Restrepo (2019) for the US, since our E/POP ratio is calculated differently here (see footnote 19). The numbers are shown in the second to last row of Panel B. The preferred estimate from column 4 implies a displacement effect of -1.5 workers per new robot.²¹

In Panel C, we repeat the analysis using the change in local average daily log wages as the outcome variable. We note that the wage estimates must be interpreted with some caution. Robot exposure displaces workers at least in the manufacturing sector, which creates selection since wage outcomes are only available for employed workers.²² We circumvent these selection issues when we look at labor earnings directly for exposed individual workers in section 6. The results by and large mirror the employment effects. Column 1 shows a small, negative, and insignificant impact of robot exposure on wage growth. Consistent with the employment results, however, we see negative effects within manufacturing in columns 2 to 4 and positive effects in the service sector in columns 5 to 7. The results strongly support the hypothesis of decreased manufacturing labor demand in regions with high robot exposure and an offsetting increase in labor demand for local services.

Panel D combines the wage and employment information by calculating sectoral total wage bills (based on the universe of social security records). The results in columns 1 to 7 strongly support the interpretation of reduced manufacturing labor demand in regions strongly exposed to automation but increasing labor demand in local service industries.

The results represent strong evidence that the adoption of robots has led to *positive employment spillovers* on other local industries in non-manufacturing.²³ Our data allow us to further look at this channel. Table A.4 in the appendix presents estimates when we split up the non-manufacturing sector into several subsectors. We differentiate business services, consumer services, construction, and public government services. The first category includes employment in establishments that render their services to other businesses on a contract or fee basis. This includes services related to information and communication technology, cleaning, or security. The second category, consumer services, contains employment in hotel and restaurant services, as well as beauty services such as haircutting.

Finally, normalizing to one additional robot per 1,000 workers, and using a ratio of the number of jobs covered by social security relative to the population of 0.28, which is the average value across regions in 1994, we get the numbers from Table 3.

²¹Using the model from Acemoglu and Restrepo (2019), one can calibrate the impact assuming dampening general equilibrium effects to arrive at numbers which are around 10-15% smaller. For comparison, the number of jobs lost per robot – both from direct and indirect – effects is calculated as around 3 for the US by Acemoglu and Restrepo (2019).

²²We conduct our analysis at the demographic group-region cell level, as in Acemoglu and Restrepo (2019) to deal with the changing observables of employed workers. Using residualized wages from Mincer regressions gives us very similar effects.

²³Any negative spillover effects from the displacement forces of automation appear, hence, to be dominated by new labor demand, at least for service industries. Gathmann, Helm, and Schönberg (2019) consider the regional effects of mass layoffs and detect significant negative spillovers. While the displacement effects we document are economically significant, industrial robots did not trigger mass layoff episodes in Germany, which limits the scope for negative spillovers.

By far the largest employment effect is on business services with a coefficient of 0.59. The consumer service coefficient, in contrast, is only estimated to have a value of 0.041. The other coefficients on construction employment and public sector employment are close to zero. Positive employment spillovers are hence driven by spending from local firms on local services. This result is consistent with the model by Acemoglu and Restrepo (2019) where increased robot adoption raises demand for complementarity inputs by producers. Relatedly, Goldschmidt and Schmieder (2017) show that task outsourcing has increased within Germany. It is conceivable that increased automation may be related to changing boundaries of the firm, and may accelerate these processes. This would be consistent with a positive effect of automation on business service employment. We leave a further empirical investigation of this very interesting issue for further research.

The appendix contains important robustness checks to our findings (Table A.2). First, we check for the presence of pre-trends by regressing lagged outcome variables on future exposure.²⁴ Second, we restrict the time window for the analysis to stop before the global great recession in 2007. Third, we conduct various checks concerning the regional dimensions.²⁵ Finally, we pay special attention to the car industry. We split up the treatment variable into exposure to robots in automobile production and robots in other industries. The displacement effect is relatively homogenous across sectors. Reallocation is driven by the exposure to robots in automobile production, in contrast. This suggests that the productivity effects were particularly large in this sector.²⁶ An alternative way to look at the automotive sector is to distinguish between automotive and other manufacturing when constructing the outcome variables, as we show in Panel H. While the effect of robots on other manufacturing industries is quantitatively similar to the overall effect, we find an exorbitant but also very imprecisely estimated negative coefficient for car manufacturing. We conclude that our main results are not exclusively driven by this sector but are rather representative for manufacturing as a whole.

3.3 Productivity and the Labor Share

The reallocation of labor should be directly tied to productivity gains. Table 4 presents direct evidence on this, using data on labor productivity which is measured as output per worker. Because of constraints on the available data, which comes from the German Federal Statistical Office, the period of analysis is the long difference between the years 2000 and 2014. As can be seen in Panel C of Table A.2, this restriction does not alter our previous results on employment

²⁴The results here imply that labor demand in manufacturing and services was trending in the opposite direction, so that higher future robots exposure was correlated with higher manufacturing employment growth.

²⁵Leaving out East Germany leaves the results unaffected. The results remain also very similar when we include time trends at the level of 16 federal states, and changing the local labor market definition, making labor markets broader (reducing the number of units from 402 to 258).

²⁶It also points to an interaction of automation with trade, since Germany is a large exporter of cars, and one would expect that productivity effects are increasing in market size.

Table 4: Labor Share and Productivity

	Dependent variable: % change between 2000 and 2014					
	Labor productivity			Labor share		
	(1)	(2)	(3)	(4)	(5)	(6)
[A] OLS						
Δ robots per 1000 workers	0.7844 (0.399) [0.401]	0.7320 (0.371) [0.345]	0.7515 (0.351) [0.338]	-0.2819 (0.141) [0.141]	-0.3207 (0.134) [0.145]	-0.3095 (0.123) [0.134]
[B] 2SLS						
Δ robots per 1000 workers	0.8030 (0.384) [0.431]	0.7533 (0.360) [0.369]	0.7774 (0.351) [0.361]	-0.2910 (0.136) [0.149]	-0.3345 (0.126) [0.153]	-0.3276 (0.119) [0.141]
Δ net exports in 1000 € per worker	No	Yes	Yes	No	Yes	Yes
Δ ICT equipment in € per worker	No	No	Yes	No	No	Yes

Notes: $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*). The period of analysis is 2000-2014 due to availability of regional GDP and employee compensation data. The dependent variable in columns 1-3 is the log change in output per worker $\times 100$ and in column 4-6 the percentage point change in the total employee compensation (*Arbeitnehmerentgelte*) over total GDP $\times 100$. All specifications include a constant, broad region dummies, demographic control variables, and employment shares of nine aggregate industry groups, measured in the base year 2000. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, Comtrade, EU KLEMS, German Federal Statistical Office, and BHP 7514 v1, own calculations.

in manufacturing and non-manufacturing. The columns are ordered as in Table 3, increasing the set of control variables. The estimated magnitudes from Panel A imply that a labor market at the 75th percentile compared to a 25th percentile labor market of the exposure distribution has accumulated around 2.3% higher productivity.

An important and common prediction made by recent task-based models of automation (e.g. Acemoglu and Restrepo, 2018b), is that the labor share is expected to decrease with the advancement of labor replacing technologies such as industrial robots.²⁷ We can leverage the local labor market design to credibly test this prediction. Panel B uses changes in the labor share again for the time period from 2000 to 2014 as the outcome of a 2SLS IV regression. Using the interquartile spread again to benchmark the magnitude, the estimates imply a negative cumulated impact of around 1% point of the labor share.

²⁷In contrast to, for example, a multi sector growth model featuring a unitary elasticity of substitution between capital and labor in production (Ngai and Pissarides, 2007).

4 Adjust Mechanism I: Reduced Creation of New Jobs for Young Workers

We have documented the presence of substantial and roughly offsetting displacement and reallocation effects of automation. This analysis relied on a local labor market approach as in Acemoglu and Restrepo (2019). In this section, we leverage the availability of detailed administrative panel data to dig deeper and understand which kind of workers are actually displaced and reallocated in response to automation. This sheds novel light on how labor markets adjust in response to automation. One of the main results will highlight that a large portion of the incidence of displacement and reallocation is borne by young workers, who face reduced (increased) job creation in the manufacturing (service) sector.

We analyze the adjustment process by decomposing the employment variables from Section 3 into mutually exclusive channels. The decomposition is additive and, hence, easy to interpret. We start by characterizing the displacement effect. Conceptually, we distinguish between workers who were working in the exposed manufacturing sector at the start of the period in 1994 and non-incumbents who were not working in manufacturing in 1994.

The set of different channels for the displacement effect to materialize are listed in the seven columns of Panel A of Table 5. Columns 1 to 3 summarize the outcomes for incumbent manufacturing workers, defined as the set of workers employed in manufacturing in the year 1994 at the start of the period analyzed. They include employment at the same plant,²⁸ within the original industry but at a different plant, and finally employment in the non-manufacturing sector. Columns 4 to 7 encompass all margins related to workers not in the manufacturing sector at the start of the period in 1994. They comprise workers who had not entered the labor market yet in 1994, workers who were already in the same local labor market but not in the manufacturing sector, workers who were employed in a different region, and temporarily non-employed workers in 1994. The coefficients from columns 1 to 7 add up to the coefficient from column 8, which is the full effect on manufacturing employment from column 4 of Table 3 and re-stated here to facilitate reading and interpretation.

Column 1 in Panel A starts with a – perhaps – surprising finding. Exposure to automation increases employment at one’s original employer. The effect is sizable and around a third of the total displacement effect from column 8. We will devote parts of the next section of the paper to explain the mechanisms and document how workers relocate within firms across tasks and occupations. So while incumbent workers face a lower layoff risk, this is offset – almost one-to-one – by decreased employment in other firms in the same industry, as evidenced by the estimate in column 2.²⁹

²⁸In our data, we only observe plants but not firms. On a few occasions, we use these term interchangeably.

²⁹These results are in line with Koch, Manuylov, and Smolka (2019), who find that Spanish firms create jobs after investing in robots. By contrast, Bessen, Goos, Salomons, and Berge (2019) find that in particular older workers are more likely to leave firms that invested in automation technologies in a broader sense. This

These two findings are consistent with the following interpretation(s). Labor market institutions in the form of firing costs make it costly to lay off workers even though the tasks previously performed by those workers are now carried out by industrial robots. At the same time, productivity effects are plausibly occurring mostly within the same firms adopting robots, which allows the re-shuffling of workers from automated tasks to other tasks, since new demand for non-automated task arises in those firms. These two forces explain why robot adoption actually increases employment with the original plant. In Section 5, we document how automation indeed causes the re-shuffling of workers across tasks within plants. In addition, below, we present (indirect) evidence on how variation in labor market institutions influences the "retainment" effect from column 1. However, the estimate in column 2 shows that – conditional on a separation – workers have a harder time regaining employment in the same industry, consistent with general reduced labor demand in robot adopting industries. Moreover, column 3 confirms that the displacement effect for incumbents also leads to reduced employment in the manufacturing sector as a whole. The overall impact on incumbent workers can then be benchmarked as around 20% of the total displacement effect, obtained by summing up columns 1 to 3 and dividing by the estimate in 8.

On the flipside, the remaining 80% of the displacement incidence falls on non-incumbent workers. The negative coefficient in column 4 reveals that the largest burden is on young workers, who had not entered the labor market in 1994 (and subsequently entered in some year between 1995 and 2014). Consistently, automation reduces flows from the service into the manufacturing sector and lowers entry from unemployment, as evidenced by columns 5 and 7. The effect sizes, however, are much smaller compared to the entrants margin. Reduced in-migration, as measured by column 6, plays no role in explaining the displacement force.

Panel B provides the same decomposition for the non-manufacturing sector to study the reallocation effect. By construction the sum of columns 1 to 7 equals the estimate from column 8 (and column 7 from Table 3). We expect zero or only very small impacts for non-manufacturing workers, since their task set is not exposed to automation. This is confirmed in columns 1 to 3. An important open question is, if the manufacturing displacement experienced by entering labor market cohorts leads to offsetting gains for young workers in services. The estimate in column 4 provides the answer and implies gains for young workers, partly offsetting the adverse impacts from displacement. If productivity effects also spillover into the service sector – something which should be expected given that tasks in this sector are q-complements to automated tasks – robot exposure should presumably also increase labor demand in services at other margins. There is indeed a positive effect – shown in column 6 of Panel B – on pulling in workers into an expanding service sector from other regions.

suggests that the concrete firm-level effects of robots might vary across countries, depending on their labor market institutions.

Table 5: Adjustment

Dependent variable: 100 x Number of workers in 2014 / total employment in 1994								
	Incumbent workers			Entrants			total	
	yes	no	no	entered	same region,	in diff.	not	
	yes	yes	no	labor mkt.	diff. sector	region	emp.	
	yes	yes	yes	after 1994	in 1994	in 1994	in 1994	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
[A] Manufacturing								
Δ robots per 1000 workers	0.1579 (0.052) [0.074]	-0.1524 (0.051) [0.089]	-0.0895 (0.014) [0.039]	-0.2315 (0.091) [0.137]	-0.0608 (0.026) [0.043]	0.0003 (0.036) [0.058]	-0.0706 (0.027) [0.039]	-0.4467 (0.160) [0.281]
[B] Non-Manufacturing								
Δ robots per 1000 workers	-0.0423 (0.014) [0.028]	-0.0251 (0.012) [0.013]	-0.0153 (0.014) [0.015]	0.5582 (0.194) [0.337]	-0.0121 (0.012) [0.006]	0.2002 (0.053) [0.106]	0.0415 (0.039) [0.037]	0.7051 (0.275) [0.442]

Notes: $N = 402$ In this table, the employment growth rate is additively split up into the contributions of different groups of incumbent workers or workers that enter the region's manufacturing (Panel A) or non-manufacturing sector (Panel B) between 1994 and 2014. The coefficients of columns 1-7 sum up to the coefficient in column 8. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Given that the incidence of the reallocation effect falls primarily on young workers, one should expect that the age structure in the manufacturing sector evolves differently than in the service sector. In the appendix in Table A.5, we find that automation reduces the average age of workers in the service sector and increases the average age of manufacturing workers (although the latter effect is small and imprecisely estimated). Our results are consistent with a two-way interaction between automation and aging. Acemoglu and Restrepo (2018a) investigate the effect of an older population on more automation. We find that more automation causes an increase in the average age of the working population in regions more affected by automation. These effects could reinforce each other.

Heterogeneity. In this subsection, we present additional results for the displacement, reallocation, and productivity effects, splitting labor markets into how strong the Social Democratic Party of Germany (SPD) polled during the 1980, 1983 and 1987 federal elections. The idea is that the strength of the German social-democratic labor party can act as a proxy for different labor market institutions strengthening incumbent workers' rights, like, for example, the power of unions, works councils, or how labor courts rule. Spier (2017) argues that for most of post-war German history, the SPD and the major trade unions were very closely aligned.³⁰ Relatedly, Berger and Neugart (2011) show that labor court verdicts on dismissal cases systematically vary with the political leaning of the government that has appointed judges across

³⁰Under the direction of chancellor Gerhard Schröder in the early 2000s, these relations started to weaken. To partly circumvent this, we use the historical vote shares preceding this period.

German states. This is a direct way in which the relative strength of left-leaning versus conservative parties introduces variation in institutions, in particular firing costs of incumbent workers.

In this exercise, we focus on West Germany since we measure the strength of the SPD according to their average vote shares during the three federal elections in the 1980ies. Unlike the two elections in 1990 and 1994, which were conducted in all of Germany, we argue that the earlier elections were not affected by the regional distribution of robot installments starting in the 1990ies and the very special events due to the German Reunification.³¹ While our measure is clearly imperfect and potentially imprecise, we would expect that exposure to automation impacts labor markets with different political leanings quite differently. Displacements and therefore productivity effects should be weaker where worker protections are stronger.³²

First, we again analyze the effects of robots on productivity and the labor share in Table 6. The earlier results from Table 4 appear to be driven by regions with a below median share of votes for the SPD. In those regions, both the positive effect of robots on labor productivity and the negative effect of robots on the labor share are particularly strong. By contrast, in regions with an above median share of SPD-votes, both effects are small and insignificant. This suggests that adjustment was more difficult for firms in regions with strong labor market institutions that favor incumbent employees.

Above in this section, we presented a worker retention result: workers in more exposed local labor markets are more likely to stay with their original plant. Are firms retaining their workers voluntarily at higher rates in the wake of automation because they value their firm-specific human capital? Or does this finding capture high firing costs? We examine this in Table 7. Comparing the coefficients in column 1 from panels A and B reveals that the retention effect is much stronger in areas with higher worker protection. At least part of the retention therefore seems to reflect institutional constraints on firms to adjust to technological change. By contrast, conditional on leaving the original plant, workers are not protected by these institutions any more and hence columns 2 and 3 show that the effects of robots on mobility to other plants within the manufacturing sector do not differ between regions with higher and lower job protection. In columns 4 to 7, we report the effects on entrants into the local manufacturing sector. Aside from the lower retention of incumbent workers, the manufacturing sector in low job protection regions also attracts fewer young entrants, formerly unemployed, and workers changing between sectors. In total, column 8 shows displacement measured by manufacturing employment was much stronger in the low worker protection regions.

³¹We show in Appendix Tables A.6 and A.7 that the results also hold when we use the vote shares of the SPD and the PDS, the successor of the socialist single political party of the DDR, in the 1994 elections.

³²Data on election outcomes are available from the Federal Returning Officer (Bundeswahlleiter) at the level of electoral districts. Since those are not congruent with counties, we aggregate the election data to administrative districts (NUTS-2 regions) before merging them to our main data.

Table 6: Labor Share and Productivity - by vote shares of the social democratic party (SPD)

	Dependent variable: % change between 2000 and 2014	
	Labor productivity	Labor share
	(1)	(2)
[A] Above median share of SPD votes in 1980, 1983, 1987 federal elections		
Δ robots per 1000 workers	0.1648 (0.207) [0.271]	-0.1358 (0.136) [0.127]
[B] Below median share of SPD votes in 1980, 1983, 1987 federal elections		
Δ robots per 1000 workers	1.5593 (0.176) [0.765]	-0.6003 (0.088) [0.254]

Notes: $N = 158$ (Panel A) and 167 (Panel B). The period of analysis is 2000-2014 due to availability of regional GDP and employee compensation data. The dependent variable in column 1 is the log change in GDP per worker $\times 100$ and in column 2 the percentage point change in the total employee compensation (Arbeitnehmerentgelte) over total GDP $\times 100$. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, BHP 7514 v1, and Federal Returning Officer, own calculations.

5 Adjust Mechanism II: Skill Upgrading

In this section, we turn our attention to different mechanisms of adjustment: the re-assignment of workers to new tasks and the upgrading of skills. The analysis will establish four new results how automation through robots influences the transition across tasks and human capital acquisition of the labor force. First, a majority of workers who are retained by their firms despite increases in automation, are re-assigned to new occupations. Second, their new occupations feature a more abstract and less routine intensive task content. Third, they are higher up in the wage ladder and are characterized by a higher college share of their workforce. Finally, the skill (college) share among labor market entrants increases significantly, the apprentice share goes down, and (relatedly) the jobs held by labor market entrants become more abstract and less routine intensive.

Table 8 presents the results from models which analyze the adjustment process for incumbent manufacturing workers. All models follow the specification with the most control variables as in Sections 3 and 4. Our linked employer-employee data allows us to observe the workplace of every worker at all points in time. We also observe 3-digit observation codes, which we aggregate to 54 economically more meaningful occupational fields according to the German Federal Institute for Vocational Education and Training (Tiemann, Schade, Helmrich, Hall, Braun, and Bott, 2008). We measure the quality of occupations according to four dimensions: the median wage of all full-time employees, the share of workers with a college degree, and the intensity in abstract and routine tasks. For the latter two, we follow Spitz-

Table 7: Manufacturing Adjustment - by vote shares of the social democratic party (SPD)

Dependent variable: 100 x Number of workers in 2014 / total employment in 1994								
	Incumbent workers			Entrants				total
Same plant as in 1994	yes	no	no	entered	same region,	in diff.	not	
Same industry as in 1994	yes	yes	no	labor mkt.	diff. sector	region	emp.	
Same sector as in 1994	yes	yes	yes	after 1994	in 1994	in 1994	in 1994	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
[A] Above median share of SPD votes in 1980, 1983, 1987 federal elections								
Δ robots per 1000 workers	0.2379 (0.070) [0.107]	-0.1524 (0.118) [0.095]	-0.0973 (0.020) [0.047]	-0.0455 (0.148) [0.131]	-0.0049 (0.038) [0.043]	-0.0655 (0.050) [0.055]	-0.0051 (0.046) [0.043]	-0.1328 (0.308) [0.269]
[B] Below median share of SPD votes in 1980, 1983, 1987 federal elections								
Δ robots per 1000 workers	0.0746 (0.070) [0.074]	-0.1475 (0.037) [0.109]	-0.0818 (0.023) [0.040]	-0.2716 (0.131) [0.196]	-0.0679 (0.020) [0.038]	0.0939 (0.047) [0.082]	-0.0856 (0.039) [0.047]	-0.4858 (0.253) [0.341]

Notes: $N = 158$ (Panel A) and 167 (Panel B). In this table, the employment growth rate is additively split up into the contributions of different groups of incumbent workers or workers that enter the region's manufacturing sector between 1994 and 2014. The coefficients of columns 1-7 sum up to the coefficient in column 8. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Oener (2006) and construct task intensities as the average shares of abstract or routine tasks in all tasks performed by around 20,000 workers surveyed in the 1991 BIBB/IAB Employment Survey.³³

Panel A starts with a decomposition of the "retainment" effect, shedding light on the question, how plants keep workers around in the wake of automation. Column 3 repeats this retainment effect from column 1 of Table 5. In columns 1 and 2, the coefficient is additively decomposed into the contribution of days employed *in the same plant* in a worker's origin occupation in 1994 and days employed in other occupations by defining the dependent variables in this way. The magnitudes imply that 77% (0.1214/0.1579) of the total effect stem from days worked in a different occupation.

It is not clear to what extent workers profit from these occupational transitions. To address this, the next set of models in Panel B investigate several dimensions of the occupational quality of jobs. The first measure is the change in median occupational wages. Concretely, we measure the quality of an occupation at any point in time as the median wage of all workers in this occupation in 1994.³⁴ The outcome variable is the log difference of the median wage of the occupation a worker held in 2014 vs. the median wage of the occupation the same worker held in 1994. In column 1, this variable is constructed only from workers who stayed in their

³³A third task category is manual tasks, which we omit here as it is mostly relevant for individual-related services.

³⁴Using median wages from earlier years as measure leaves the results unaffected.

Table 8: Occupational Upgrading Within and Across Firms

	(1)	(2)	(3)	(4)
[A] Occupational adjustment				
	Dependent variable: 100 x Number of workers in 2014 / total employment in 1994			
Same plant as in 1994	yes	yes	yes	
Same occupation as in 1994	yes	no	(total)	
Δ robots per 1000 workers	0.0366 (0.027) [0.022]	0.1214 (0.030) [0.058]	0.1579 (0.052) [0.074]	
[B] Occupational upgrading: Wages and skills				
	Dependent variables: Δ log median wage in € 100 x Δ college share			
Same plant as in 1994	yes	no	yes	no
Δ robots per 1000 workers	0.0620 (0.025) [0.045]	0.0411 (0.029) [0.033]	0.0687 (0.020) [0.037]	0.0256 (0.019) [0.015]
[C] Occupational upgrading: Tasks				
	Dependent variables: 100 x Δ abstract task intensity 100 x Δ routine task intensity			
Same plant as in 1994	yes	no	yes	no
Δ robots per 1000 workers	0.0808 (0.024) [0.044]	-0.0082 (0.022) [0.017]	-0.1276 (0.031) [0.074]	-0.0466 (0.024) [0.030]

Notes: $N = 402$. In this table, we analyze the effect of robots on the occupation dimension of exposed workers. In Panel A, the dependent variables are 100x the number of workers who stay in the manufacturing sector of their original region but show different kinds of job mobility, relative to total employment in 1994. The coefficients of Panel A, columns 1 and 2 add up to the coefficient in column 1 of Panel A, Table 5 (also reported in column 3). In Panels B and C, we focus on the occupational quality of workers who stay in the manufacturing sector of their original region but possibly switch into a different occupation. The dependent variable in columns 1 and 2 of Panel B is the difference of the median wage, measured in 1994, of the occupation of workers staying in the same plant in 2014 versus the occupation in 1994. The dependent variable in columns 3 and 4 of Panel B is the difference of the percentage of people with a college degree, measured in 1994, of the occupation of workers staying in the same plant in 2014 versus the occupation in 1994. The dependent variable in Panel C is the difference of the abstract (columns 1 and 2) and routine (columns 3 and 4) task intensities, measured in 1994, of the occupation of workers staying in the same plant in 2014 versus the occupation in 1994. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets. Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

initial plant, while in column 2, the outcome is analogously defined for workers who switched between plants. Positive coefficients would indicate that robots exposure leads to occupational upgrading. Column 1 displays a positive coefficient, around twice the size of the coefficient in column 2. So, on average, higher robots exposure causes occupational mobility up the wage ladder and the effect is much stronger within plants, so for workers who are retained by their original employer.

Table 9: Robots and skill share of people younger than 30

	Dependent variable:			
	100 x Δ Share of workers with		Task intensity	
	university degree (1)	apprenticeship degree (2)	abstract (3)	routine (4)
Δ robots per 1000 workers	0.0932 (0.039) [0.049]	-0.0684 (0.033) [0.042]	0.0686 (0.031) [0.040]	-0.0604 (0.020) [0.038]

Notes: In this table, we analyze the effect of robots on occupational quality of younger workers. The estimates are based on $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*). The dependent variables is the change in various measures for occupation quality of workers 30 years old or less between 1994 and 2014: Share of workers with university degree (column 1), share of workers with apprenticeship degree (2), average abstract task intensity (3), and average routine task intensity (4). The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

The dependent variable in column 3 (4) of Panel B is the difference of the percentage of people with a college degree, measured in 1994, of the occupations of workers working in the same plant (a different plant) in 2014 versus the college percentage of their occupation in 1994.³⁵ The results imply a positive effect of automation on occupational quality and the effect is again much larger for firm stayers.³⁶

Finally, Panel C studies the re-assignment of tasks for exposed workers. The dependent variable in Panel C is the difference of the abstract (columns 1 and 2) and routine (columns 3 and 4) task intensities, measured in 1994, of the occupation of workers staying in the same plant in 2014 versus their occupation in 1994. Columns 1 and 3 present evidence that automation causes a shift in the careers of workers away from routine to abstract tasks within plants. The coefficients in columns 2 and 4 show much smaller effects across plants.

We next turn to human capital adjustments of young cohorts. The first dependent variable of Table 9 is the change in the share of college educated workers of age 30 and lower, which hold at least a degree which requires 3 years or more of tertiary education. The positive coefficient indicates that young people adjust to local automation by increasing their level of education. Column 2 shows that is counteracted by a significant reduction in the apprenticeship share. Importantly, the table also shows that the adjustment efforts of young workers extends beyond education choices and into occupational choices. In columns 3 and 4 we measure the effect of robots on changes in the task contents of jobs held by people below 30. One sees a reallocation in highly exposed regions from routine towards abstract tasks. These results are robust to

³⁵Again, using lagged college share produces almost the same results, since skill share remained fairly stable during this time period.

³⁶Comparing plant stayers and switchers may be difficult if automation changes the composition of workers who stay/leave their original firm. We found no significant differences in the difference between stayers and leavers in highly exposed versus weakly exposed regions in terms of observables.

using different age cutoffs than 30; in Appendix Table A.3 we present the results with a cutoff of 40 as an example.

6 Individual Level Results

We now shift the focus from local labor market adjustments to the individual worker. This complements the previous models because it allows to directly study the effects of automation on earnings and wages using a more compelling design. Comparing wage or earnings growth across local labor markets, in contrast, can lead to biased results because automation changes the composition of employed workers. By following the same workers, we can circumvent those selection issues.

6.1 Earnings and Employment

We use an exposure to automation design which compares the outcomes of workers which were employed in a manufacturing industry in 1994.³⁷ We follow the standard practice in the literature and focus on workers with sufficiently high labor force attachment. This means that we restrict the sample to workers who were i) between 22 and 44 years old, ii) earned more than the marginal-job threshold, and iii) had job tenure for at least two years in the base year 1994.³⁸ Finally, we keep only workers in manufacturing industries that can be matched to the IFR data. The specification is:

$$Y_{ij} = \alpha \cdot \mathbf{x}'_{ij} + \beta \cdot \Delta\text{robots}_j + \gamma \cdot \mathbf{z}'_j + \epsilon_{ij}.$$

Y_{ij} represents the cumulated number of days spent in employment – irrespective if employed in a manufacturing or a different sector – over the 1995-2014 period in the first set of regressions. In the vector \mathbf{x}'_{ij} we include worker-level controls, measured in the base year 1994: dummies for gender, foreign nationality, three skill categories, and three tenure categories. In addition we include a full set of age dummies, federal state dummies (there are 16 federal states), and dummies for six plant size groups. We also control for the log of yearly earnings of a worker at the start of the period, i.e. 1994.

The term Δrobots_j is the change in robot adoption per worker – with the number of workers fixed at the starting level in 1994 – in industry j . As described in Section 2, the IFR classification allows to distinguish 20 manufacturing industries. To account for this, we cluster standard errors at the levels of the IFR classification with 20 clusters. \mathbf{z}'_j is a vector of industry controls with dummies for broad industry groups to control for broad trends.³⁹ It

³⁷This approach has also been used by Autor, Dorn, Hanson, and Song (2014) to study the worker-level impacts of trade shocks. We follow their method here.

³⁸Results are very similar, however, when including also workers with lower attachment.

³⁹The categories are, as in Section 3, food products, consumer products, capital goods, and industrial goods.

also contains changes in trade exposure at the 3-digit level and ICT exposure at the 2-digit level.⁴⁰

As of the data in this section, we use a 30 percent sample of the Integrated Labor Market Biographies (IEB V12.00.00) of the Institute for Employment Research. This data is similar to the administrative data introduced in Section 2.1 but covers the complete employment biographies with daily precision and not only the main observation on June 30. Since East Germany saw very strong wage growth up until 1995, related to other factors besides automation, we drop workers who were employed there in 1994 in a robustness check. Our results are unaffected, consistent with the analogue robustness checks at the regional level. Table 10 reports descriptive statistics of the variables used in the worker level analysis. The average manufacturing worker in our sample has experienced an exposure equal to $\Delta\text{robots}_j = 24.4$ (see panel C). Notice the large variation across individuals. The worker at the 75th percentile has seen an increase in exposure that is almost five times larger than for the worker at the 25th percentile (26.1 versus 5.5 additional robots per thousand workers), and the comparison between the 90th and the 10th percentiles is even more dramatic (104.3 versus -2.7). This reflects the extremely skewed distribution of robot installation across industries that is illustrated in Figure A.1. The average worker in our sample is employed for 5,980 days during the 20 years after 1994, which amounts to 82% of the duration of this period (7,305 days). In terms of earnings, workers have on average almost exactly retained their base year earnings over the 20-year period.

Table 11 shows how workers have adjusted in response to the emergence of industrial robots. In both panels, the coefficients listed in columns 2 to 5 sum up to the total effect in 1. Column 1 shows a small, positive impact on employment. From column 2 of Panel A, it becomes clear that this positive effect is driven by increased employment at one's original plant, echoing the local labor market results from Table 5. The economic magnitude of this effect is large and around eleven times the size of the total employment effect. Quantitatively, it translates into an increase of 197 days of employment (over 20 years) in one's original plant for a worker starting out in the industry at the 75th percentile relative to a worker from 25th percentile exposed manufacturing industry. This number grows to 1,024 days roughly equal to three years comparing the 90th and the 10th percentile.

Column 3 shows reduced transitions into other firms within the same industry.⁴¹ This is consistent with our interpretation that workers are institutionally protected from displacement at one's own firm but have a hard time to find other gainful employment within the same industry in the face of automation. Movements to other industries are reduced, as shown by columns 4 and 5.

⁴⁰See the data part in Section 2 for a description. See Dauth, Findeisen, and Suedekum (2019) for details on the trade variables.

⁴¹Industry mobility is classified according to the 20 IFR industries, so at the level of robot adoption variation.

Table 10: Summary statistics, worker level.

Observations	720,562	
	mean	(sd)
[A] Outcomes, cumulated over years following base year		
Days employed	5,980	(1,986)
Average daily wage	121.3	(71.2)
100 x earnings / base year earnings	1,949.8	(1,000.3)
[B] Control variables, measured in base year		
Base year earnings	38,683	(20,599)
Base year average wage	106.55	(55.14)
Dummy, 1=female	0.211	(0.408)
Dummy, 1=foreign	0.110	(0.313)
Birth year	1960	(6)
Dummy, 1=low skilled	0.160	(0.366)
Dummy, 1=medium skilled	0.751	(0.432)
Dummy, 1=high skilled	0.089	(0.285)
Dummy, 1=tenure 2-4 yrs	0.397	(0.489)
Dummy, 1=tenure 5-9 yrs	0.317	(0.465)
Dummy, 1=tenure ≥ 10 yrs	0.247	(0.431)
Dummy, 1=plant size ≤ 9	0.054	(0.225)
Dummy, 1=plant size 10-99	0.224	(0.417)
Dummy, 1=plant size 100-499	0.289	(0.453)
Dummy, 1=plant size 500-999	0.122	(0.328)
Dummy, 1=plant size 1000-9999	0.225	(0.418)
Dummy, 1=plant size ≥ 10000	0.084	(0.277)
Dummy, 1=food products	0.094	(0.291)
Dummy, 1=consumer goods	0.033	(0.178)
Dummy, 1=industrial goods	0.393	(0.488)
Dummy, 1=capital goods	0.465	(0.499)
Δ net exports / wagebill in %	14.413	(57.996)
Δ ICT equipment in € per worker	254.9	(271.7)
[C] Exposure to robots		
Δ robots per 1000 workers	24.400	(40.119)
p10-p90 interval	[-2.721 ; 104.258]	
p25-p75 interval	[5.547 ; 26.052]	

Sources: IFR, COMTRADE, EU KLEMS, and IEB V12.00.00, own calculations.

Panel B extends the analysis to individual adjustments across occupations, using the same classification of 54 occupational fields as in Section 5. Again, of high interest here is how adjustments within firms take place, given displacement by robots. Columns 2 and 3 examine this by splitting employment within spells at the original plant into time worked in the base year occupation and other occupations – consequently the two estimates sum up to the coefficient in column 2 of Panel A. Approximately two thirds of the employment at the original plant effect are driven by employment in a different occupation. Both coefficients are statistically and economically significant. The decompositions can also be used to get a total occupational mobility effect across all firms. We can add columns 2 and 4 to obtain the effect of robot exposure on time spent in one’s original occupation and compare it to the sum of column 3 and 5, which encompasses time spent in a different occupation. This gives

Table 11: Individual Adjustment to Robot Exposure (Employment)

[A] Industry mobility	(1)	(2)	(3)	(4)	(5)
	all employers	same sector		Service Sector	
Same industry		yes	yes	no	no
Same employer		yes	no	no	no
Δ robots per 1000 workers	0.8659 (1.290)	9.5719 (1.559)	-3.5587 (2.689)	-2.1185 (1.508)	-3.0288 (1.247)

[B] Occupational mobility	(1)	(2)	(3)	(4)	(5)
	all jobs	same employer	other employer		
Same occupational field		yes	no	yes	no
Δ robots per 1000 workers	0.8659 (1.290)	2.9067 (1.002)	6.6652 (1.417)	-6.1868 (1.562)	-2.5193 (0.978)

Notes: Based on 720,562 workers. OLS results for period 1994-2014. The outcome variables are cumulated days of employment. For column 1, employment days are cumulated over all employment spells in the 20 years following the base year. Panel A: For column 2 employment days are cumulated only when they occurred at the original workplace. For the other columns, employment days are cumulated only when they occurred at a different plant in the same industry 3, at a plant in a different manufacturing industry 4, and outside the manufacturing sector 5, respectively. Panel B: Employment days are cumulated only when they occurred in the original occupation and workplace column 2, in a different occupation but at the original workplace column 3, in the original occupation but at a different workplace column 4, and in a different occupation and workplace, respectively. Control variables are log base year earnings and indicator variables for gender, foreign nationality, birth year, educational degree (3 categories), tenure (3 categories), plant size (6 categories), broad manufacturing industry groups (4 categories), and 16 federal states. Standard errors are clustered by 20 ISIC Rev.4 industries in parentheses.

Sources: IFR, Comtrade, EU KLEMS, and IEB V12.00.00, own calculations.

$2.9067 - 6.1868 = -3.2801$ versus $6.6652 - 2.5193 = 4.1459$: in sum, automation has increased occupational mobility.

A popular narrative in the public debate is that affected workers will have to be flexible and mobile across tasks and occupations to be "one step ahead" of labor displacing technologies. Those sets of results first imply that workers in Germany already responded by switching tasks to the rise of industrial robots. Second, the reassignment of workers to new tasks happen frequently within a worker's original firm.

Table 12 extends the analysis to earnings. The models are an important complement to the employment regressions since they help to paint a more complete picture about workers' welfare than looking at employment outcomes alone. Following Autor, Dorn, Hanson, and Song (2014), to create the outcome variable, we cumulate earnings over the whole period and divide them by average earnings in 1994. The regressions can hence be interpreted as differences-in-differences designs.

We begin in Panel A by studying the effect on earnings from all sources. In contrast to the employment effects, one obtains a negative albeit very small and insignificant point estimate of -0.19. To interpret the coefficient, we calculate the quartile spread again, comparing an industry at the 75th percentile to an industry at the 25th percentile. The implied reduction in earnings (over the whole 20 year period and *not* per year) would be 3.9% of annual initial

Table 12: Individual Adjustment to Robot Exposure (Earnings)

[A] Industry mobility	(1)	(2)	(3)	(4)	(5)
	all employers	same sector			Service Sector
Same industry		yes	yes	no	no
Same employer		yes	no	no	no
Δ robots per 1000 workers	-0.1852 (0.869)	3.1138 (0.561)	-1.3986 (1.018)	-0.8332 (0.629)	-1.0671 (0.388)

[B] Occupational mobility	(1)	(2)	(3)	(4)	(5)
	all jobs	same employer		other employer	
Same occupational field		yes	no	yes	no
Δ robots per 1000 workers	-0.1852 (0.869)	0.7997 (0.313)	2.3141 (0.525)	-2.3107 (0.631)	-0.9883 (0.404)

Notes: Based on 720,562 workers. OLS results for period 1994-2014. The outcome variables are 100 x earnings (normalized by earnings in the base year) cumulated over the 20 years following the base year. For column 1, earnings are cumulated over all employment spells in the 20 years following the base year. Panel A: For column 2 earnings are cumulated only when they occurred at the original workplace. For the other columns, earnings are cumulated only when they occurred at a different plant in the same industry 3, at a plant in a different manufacturing industry 4, and outside the manufacturing sector 5, respectively. Panel B: Earnings are cumulated only when they occurred in the original occupation and workplace column 2, in a different occupation but at the original workplace column 3, in the original occupation but at a different workplace column 4, and in a different occupation and workplace, respectively. Control variables are log base year earnings and indicator variables for gender, foreign nationality, birth year, educational degree (3 categories), tenure (3 categories), plant size (6 categories), broad manufacturing industry groups (4 categories), and 16 federal states. Standard errors clustered by 20 ISIC Rev.4 industries in parentheses.

Sources: IFR, Comtrade, EU KLEMS, and IEB V12.00.00, own calculations.

earnings. Or about 1,500 Euro in absolute terms for a worker with average earnings. The effect increases by an order of magnitude to 3.11, estimated to be statistically significant, for earnings at the original plant. This is offset, approximately equally across the different channels, by reduced earnings in other plants, industries, and the service sector, however. In other words, robots create inequality between workers who separate and plant stayers.

To measure the role of occupational adjustments, Panel B examines the effects of earnings across occupations. Did occupational switching help workers to respond to automation? Of particular interest are the coefficients in columns 2 and 3, which decompose the original plant earnings effect into impacts for the starting versus other occupations. The split is very close to 75%. Occupational (and presumably task) transitions within firms play, hence, a large role for the labor earnings impacts of automation. Columns 4 and 5 complete this picture. While earnings at other firms decrease in all occupations, the decrease is much more pronounced for a worker's original occupation.

6.2 Skill or Task Bias?

In the final step of our analysis, we explore heterogeneous impacts across occupations and skill groups. A very influential literature has investigated the *skill bias* of technological change (Katz and Murphy, 1992). A newer literature has instead emphasized the *task bias* of techno-

logical developments.⁴² This section presents new evidence how the advancements of industrial robot technology have affected inequality across different occupation and task groups.

The results are contained in Figure 3, where we show the point estimates of interaction terms of the increase in robot exposure and 95% confidence intervals, based on clustered standard errors across the 20 IFR industries, for different groups of workers. The regression models are the same as in the last section for earnings. So we include controls for skill categories, tenure categories, age, plant size categories, initial industry, and region – and the dependent variable is cumulated labor earnings.⁴³ Panel A differentiates six broad occupational categories that can be found among the individual manufacturing workers in our sample. Panel B distinguishes three skill categories.

In Panel A, for two occupation groups, the estimated impact is economically meaningful and positive (only the point estimate for the former group is statistically significant at the conventional 5% level). These are managers and legal specialists, as well as occupations in the fields of technical science and natural science. This group encompasses, for example, all kinds of engineers, as well as chemists. Automation through robots has benefited these cognitive-task heavy and, in general, very skilled occupations.


In the middle of the spectrum, with estimated coefficients slightly positive but very close to zero, one finds the point estimates for clerical/sales workers and a bundle of occupations, encompassing e.g. security and transportation workers. The common theme here is that the task set of those occupations is mostly non-routine and, hence, at least during the period we study, technically harder to automate. Interestingly, the rents from robots are seemingly passed on at higher rates to the set of skilled, technical occupations discussed in the preceding paragraph.

The next lines present the results for a set of occupations, which are suspected to be strongly exposed to replacement. Indeed, we find significant earnings losses mainly for machine operators. Industrial robots – by definition – do not require a human operator anymore but have the potential of conducting many production steps autonomously. Robots therefore directly substitute the task sets of this group. The point estimate here implies that a manufacturing worker at the third quartile of exposure sees an earnings reduction of around 25% of initial annual earnings, relative to a worker at the first quartile of exposure. A qualitatively similar finding is obtained for workers in processing and maintaining jobs but the effect size here is only a third of the effect for machine operators.


A second natural way to cut the data is to consider impacts across education groups, following an enormous literature investigating how technological change affects relative skill demand. In the German context, because of the prevalence of the apprenticeship system, it

⁴²See Acemoglu and Autor (2011) for a survey of both literatures and Autor and Dorn (2013) or Goos, Manning, and Salomons (2014) for prominent empirical applications.

⁴³We obtain similar effects for wages but prefer the earnings models since they avoid the classical selection problem that wages are not observed for non-employed people.



(a) Occupation: Heterogenous Impacts



(b) Education: Heterogenous Impacts

Notes: The figures report the coefficients of interaction terms of Δ robots per 1000 workers and dummies indicating the respective worker group. The outcome variables are 100 x earnings (normalized by earnings in the base year) cumulated over the 20 years following the base year. In panel A, occupations base on the definition of aggregate occupational fields by the German Federal Institute for Vocational Education and Training (BIBB) with the following modifications: Sales and clerical occupations are combined and agriculture, mining, and construction (that would have a point estimate of zero with a huge standard error) are omitted. In panel B, high skilled is defined as having a degree from a university or university of applied sciences, and medium skilled is defined as having a vocational training degree. All other educational levels are subsumed as low skilled. All regressions include the same full set of control variables as in Table 12. The confidence intervals are constructed from standard errors clustered by 20 ISIC Rev.4 industries.

Figure 3: Heterogeneous earnings effects

makes sense to split the population not just into two but three skill groups. In Panel B, high skilled is defined as having a degree from a university or college, and medium skilled is defined as having a vocational training degree. All other educational levels are subsumed as low skilled (i.e., high school graduates and high school dropouts). Completed apprenticeship is the typical profile for manufacturing workers in Germany, accounting for almost 75% of all individuals in the sample. 16% are low-skilled and 9% high-skilled according to the classification.

The general take-away here is that occupations represent a much more powerful cut of the data. Although for each of the three skill groups, sample sizes are much larger than for the occupations split, confidence bands are much wider. The figure shows approximately equal negative point estimates for low- and medium-skilled workers. In contrast, college-educated workers see earnings increases.

7 Conclusion

Many people foresee a further rise of robots, artificial intelligence, and other automation technologies, which can potentially disrupt labor markets. The still small but growing empirical literature on this topic, most importantly Acemoglu and Restrepo (2019) and Graetz and Michaels (2018), have documented the (negative) effects of industrial robots on employment and wages and (positive) impacts on productivity. Nevertheless, there has been little work on studying the adjustment processes of labor markets and its main actors (workers and firms) in response to new automation technologies. German administrative labor market data provides us with a very rare longitudinal perspective over an extended period of time to examine how workers and firms respond to automation in a context with very high adoption rates.

The results paint a nuanced picture how the strong rise in automation affected workers. They also point to a strong interaction with labor market institutions. Relatively strong protections for incumbent workers shift the incidence of job displacement on young workers and labor market entrants. In order to retain workers whose task set is automated, one expects transition to new occupations and tasks within employers. We find several pieces of evidence that these transitions are going on and that they contribute significantly to soften the blow of automation. Encouragingly, the data suggests that skill upgrading goes hand-in-hand with those transitions. Such skill upgrading is also observed for young workers and labor market entrants.

Labor market institutions are an important mediator of the effects of technological advance on the labor market. How the next generation of advances in AI, machine learning, and new manufacturing technologies will impact workers, will also depend on the design of these institutions. We believe these questions should be investigated with, on the one hand, more empirical evidence on the interaction, but also theoretical work incorporating institutional aspects and the frictions inherent in labor markets.


References

- Acemoglu, Daron and David Autor (2011). “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: *Handbook of Labor Economics* 4, pp. 1043–1171.
- Acemoglu, Daron and Pascual Restrepo (2018a). *Demographics and Automation*. NBER Working Paper No. 24421.
- Acemoglu, Daron and Pascual Restrepo (2018b). “The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment”. In: *American Economic Review* 108(6), pp. 1488–1542.
- Acemoglu, Daron and Pascual Restrepo (2019). “Robots and Jobs: Evidence from US Labor Markets”. In: *Journal of Political Economy*.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales (2019). “Shift-Share Designs: Theory and Inference”. In: *Quarterly Journal of Economics* 134(4), pp. 1949–2010.
- Autor, David H. and David Dorn (2013). “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”. In: *American Economic Review* 103(5), pp. 1553–1597.
- Autor, David H., David Dorn, and Gordon H. Hanson (2013). “The China Syndrome: Local Labor Market Effects of Import Competition in the United States”. In: *American Economic Review* 103(4), pp. 2121–2168.
- Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song (2014). “Trade Adjustment: Worker Level Evidence”. In: *Quarterly Journal of Economics* 129(4), pp. 1799–1860.
- Autor, David and Anna Salomons (2018). “Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share”. In: *Brookings Papers on Economic Activity* 2018(1), pp. 1–87.
- Bartel, Ann, Casey Ichniowski, and Kathryn Shaw (2007). “How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills”. In: *Quarterly Journal of Economics* 122(4), pp. 1721–1758.
- Berger, Helge and Michael Neugart (2011). “Labor courts, nomination bias, and unemployment in Germany”. In: *European Journal of Political Economy* 27(4), pp. 659–673.

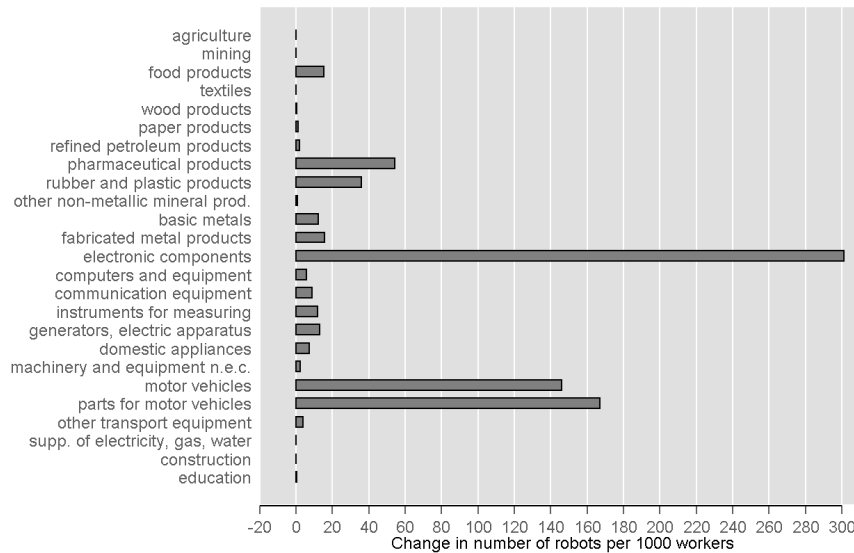
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge (2019). *Automatic Reaction – What Happens to Workers at Firms that Automate?* Boston University Law & Economics Series Paper No. 19-2.
- Bloom, Nicholas, Mirco Draca, and John van Reenen (2016). “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity”. In: *Review of Economic Studies* 83(1), pp. 87–117.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2018). *Quasi-experimental Shift-Share Research Designs*. NBER Working Paper 24997.
- Card, David, Joerg Heining, and Patrick Kline (2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality”. In: *Quarterly Journal of Economics* 128(3), pp. 967–1015.
- Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2017). “Trade and Manufacturing Jobs in Germany”. In: *American Economic Review Papers & Proceedings* 107(5), pp. 337–342.
- Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2019). “Adjusting to Globalization in Germany”. In: *Journal of Labor Economics*.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Wößner (2018). *German robots - the impact of industrial robots on workers*. CEPR Discussion Paper 12306.
- Faber, Marius (2019). *Robots and reshoring: Evidence from Mexican labor markets*. Working Paper.
- Gathmann, Christina, Ines Helm, and Uta Schönberg (2019). “Spillover Effects of Mass Lay-offs”. In: *Journal of the European Economic Association*.
- Goldschmidt, Deborah and Johannes F. Schmieder (2017). “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure”. In: *The Quarterly Journal of Economics* 132(3), pp. 1165–1217.
- Goos, Maarten, Alan Manning, and Anna Salomons (2014). “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring”. In: *American Economic Review* 104(8), pp. 2509–2526.
- Graetz, Georg and Guy Michaels (2018). “Robots at Work”. In: *Review of Economics and Statistics* 100(5), pp. 753–768.

- Helm, Ines (2019). “National Industry Trade Shocks, Local Labor Markets, and Agglomeration Spillovers”. In: *Review of Economic Studies*.
- International Federation of Robotics (2016). *World Robotics Industrial Robots 2016*.
- Jacobson, Louis S., Robert LaLonde, and Daniel Sullivan (1993). “Earnings Losses of Displaced Workers”. In: *American Economic Review* 83(4), pp. 685–709.
- Katz, Lawrence F. and Kevin M. Murphy (1992). “Changes in Relative Wages, 1963-1987: Supply and Demand Factors”. In: *Quarterly Journal of Economics* 107(1), pp. 35–78.
- Kleibergen, Frank and Richard Paap (2006). “Generalized Reduced Rank Tests using the Singular Value Decomposition”. In: *Journal of Econometrics* 133(1), pp. 97–126.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka (2019). *Robots and firms*. CESifo Working Paper No. 7608.
- Michaels, Guy, Ashwini Natraj, and John van Reenen (2014). “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years”. In: *Review of Economics and Statistics* 96(1), pp. 60–77.
- Moll, Benjamin, Pasqual Restrepo, and Lukasz Rachel (2019). *Uneven Growth: Automation’s Impact on Income and Wealth Inequality*. Working Paper.
- Moretti, Enrico (2011). “Local Labor Markets”. In: *Handbook of Labor Economics* 4, pp. 1237–1313.
- Ngai, L. Rachel and Christopher A. Pissarides (2007). “Structural Change in a Multisector Model of Growth”. In: *American Economic Review* 91(1), pp. 429–443.
- Spier, Tim (2017). “Growing Apart?” In: *Left-of-Centre Parties and Trade Unions in the Twenty-First Century*. Ed. by Elin Haugsgjerd Allern and Tim Bale. Oxford Scholarship Online. Chap. 7.
- 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.
- Tiemann, Michael, Hans-Joachim Schade, Robert Helmrich, Anja Hall, Uta Braun, and Peter Bott (2008). *Berufsfeld-Definitionen des BIBB auf Basis der Klassifikation der Berufe 1992*. Wissenschaftliche Diskussionspapiere des Bundesinstituts für Berufsbildung (BIBB), Heft 105.

Appendix



(a) German robots.



(b) US robots.

Figure A.1: Industry-level distribution of increase in number of robots

Notes: The figure displays the change in the number of robots per thousand workers by ISIC Rev.4 industries (German Classification of Economic Activities, Edition 2008), for the period 1994-2014. Increase in the number of US robots in panel (b) is also normalized by German industry-level employment.

Source: International Federation of Robotics (IFR) and BHP 7514 v1, own calculations.

Table A.1: Composition Effects - 2SLS

	Total	Manufacturing		Non-manufacturing			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
[A] Employment: % change in total employment between 1994 and 2014							
Δ robots per 1000 workers	0.0675 (0.137) [0.181]	-0.6029 (0.167) [0.319]	-0.6302 (0.160) [0.315]	-0.4745 (0.162) [0.297]	0.5869 (0.325) [0.400]	0.5670 (0.322) [0.408]	0.7299 (0.329) [0.466]
Kleibergen-Paap weak ID test	378.041	391.407	383.098	378.041	391.407	383.098	378.041
Hansen J p-value	0.203	0.036	0.022	0.037	0.552	0.547	0.568
[B] E/Pop: 100 x Δ in employment/population between 1994 and 2014							
Δ robots per 1000 workers	0.0073 (0.062) [0.032]	-0.0525 (0.025) [0.035]	-0.0570 (0.025) [0.033]	-0.0495 (0.027) [0.030]	0.0461 (0.046) [0.039]	0.0451 (0.047) [0.040]	0.0568 (0.046) [0.046]
Effect of 1 robot	0.3	-1.9	-2.0	-1.7	1.6	1.6	2.0
Kleibergen-Paap weak ID test	378.041	391.407	383.098	378.041	391.407	383.098	378.041
Hansen J p-value	0.671	0.014	0.014	0.017	0.097	0.092	0.109
[C] Wages: 100 x Log- Δ in average wage between 1994 and 2014							
Δ robots per 1000 workers	-0.0402 (0.045) [0.031]	-0.1498 (0.052) [0.083]	-0.1574 (0.054) [0.085]	-0.1107 (0.068) [0.080]	0.1022 (0.040) [0.065]	0.0940 (0.040) [0.064]	0.1050 (0.041) [0.067]
Kleibergen-Paap weak ID test	401.095	395.706	389.132	386.957	412.681	404.781	400.810
Hansen J p-value	0.126	0.053	0.061	0.055	0.105	0.079	0.089
[D] Wagebill: 100 x Log- Δ in total wagebill on June 30							
Δ robots per 1000 workers	0.0589 (0.153) [0.210]	-0.6783 (0.168) [0.368]	-0.7243 (0.163) [0.366]	-0.5250 (0.195) [0.358]	0.4177 (0.248) [0.319]	0.3930 (0.244) [0.325]	0.5595 (0.251) [0.388]
Kleibergen-Paap weak ID test	378.041	391.407	383.098	378.041	391.407	383.098	378.041
Hansen J p-value	8.269	13.221	14.463	12.888	5.913	5.760	5.212
Δ net exports in 1000 € per worker	Yes	No	Yes	Yes	No	Yes	Yes
Δ ICT equipment in € per worker	Yes	No	No	Yes	No	No	Yes

Notes: In all regressions, the variable of interest is the change in robot exposure between 1994 and 2014. The estimates in panels A, B, and D are based $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*), while the unit of observation in the wage estimates in panel (C) are $N = 7,236$ region x demographic cells. Demographic cells are defined by gender, three age groups, and three education groups. We only include cells containing at least 10 observations, and perform the regressions at the region x demographic cell level including fixed effects for demographic cells. The dependent variable in Panel D is the log-difference total amount of gross salaries paid to employees subject to social security on June 30 in 1994 and 2014. All specifications include a constant, broad region dummies, demographic control variables, and employment shares of nine aggregate industry groups, measured in the base year 1994. The regressions are estimated by applying the 2SLS IV approach where German robot exposure is instrumented with robot installations in other high-income countries. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, Comtrade, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Table A.2: Robustness checks.

	Employment			Average Wages		
	(1) Total	(2) Manuf.	(3) Non-Manuf.	(4) Total	(5) Manuf.	(6) Non-Manuf.
[A] Pre-Trends						
Δ robots per 1000 workers	0.1076 (0.159) [0.115]	0.4021 (0.206) [0.202]	-0.0074 (0.121) [0.097]	0.0172 (0.021) [0.023]	-0.0012 (0.024) [0.018]	0.0251 (0.027) [0.036]
N	326	326	326	5846	5242	5828
[B] 1994-2007						
Δ robots per 1000 workers	-0.1194 (0.087) [0.114]	-0.5543 (0.185) [0.307]	0.1925 (0.172) [0.171]	-0.0883 (0.048) [0.039]	-0.2220 (0.079) [0.101]	0.0905 (0.053) [0.091]
N	402	402	402	7235	6898	7231
[C] 2000-2014						
Δ robots per 1000 workers	-0.2600 (0.177) [0.159]	-0.8499 (0.232) [0.384]	0.3789 (0.275) [0.275]	-0.0892 (0.063) [0.051]	-0.1239 (0.079) [0.081]	0.0807 (0.053) [0.051]
N	402	402	402	7236	6956	7236
[D] West Germany, OLS						
Δ robots per 1000 workers	0.0075 (0.138) [0.126]	-0.4493 (0.175) [0.249]	0.6486 (0.276) [0.400]	-0.0403 (0.049) [0.028]	-0.1437 (0.062) [0.066]	0.1003 (0.036) [0.065]
N	325	325	325	5849	5545	5845
[E] Federal state dummies, OLS						
Δ robots per 1000 workers	0.0629 (0.126) [0.166]	-0.4326 (0.168) [0.262]	0.6796 (0.274) [0.414]	-0.0418 (0.051) [0.029]	-0.1297 (0.061) [0.077]	0.0911 (0.038) [0.059]
N	402	402	402	7235	6896	7231
[F] 258 Local labor markets, OLS						
Δ robots per 1000 workers	-0.1155 (0.154) [0.170]	-0.6475 (0.295) [0.444]	0.5161 (0.216) [0.293]	-0.0431 (0.064) [0.036]	-0.0979 (0.067) [0.092]	0.1053 (0.052) [0.071]
N	258	258	258	4643	4489	4643
[G] Split automotive and other manufacturing in treatment variables						
Δ robots per 1000 workers <i>automobile industry</i>	0.0882 (0.113) [0.158]	-0.4454 (0.157) [0.288]	0.7125 (0.244) [0.393]	-0.0338 (0.048) [0.029]	-0.0958 (0.060) [0.070]	0.0924 (0.034) [0.046]
Δ robots per 1000 workers <i>other industries</i>	-0.1316 (0.205) [0.204]	-0.5616 (0.416) [0.524]	0.0512 (0.257) [0.274]	-0.0951 (0.075) [0.068]	-0.1140 (0.113) [0.112]	-0.0833 (0.045) [0.066]
N	402	402	402	7235	6896	7231
[H] Split automotive and other manufacturing in outcome variables						
	total manuf.	car manuf.	other manuf.	total manuf.	car manuf.	other manuf.
Δ robots per 1000 workers	-0.4467 (0.160) [0.281]	-5.2065 (21.951) [18.740]	-0.5326 (0.242) [0.323]	-0.0960 (0.060) [0.070]	-0.1207 (0.126) [0.131]	-0.1785 (0.106) [0.079]
N	402	382	402	6896	2830	6866

Notes: This table presents modifications the baseline specifications for employment and average wages as of columns 1, 4 and 7 of Table 3. The dependent variables are employment growth rates (column 1-3) and log-differences in average wages (column(4-6)). Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, and BHP 7514 v1, own calculations.

Table A.3: Robots and skill share of people younger than 40

	Dependent variable:			
	100 x Δ Share of workers with		Task intensity	
	university degree	apprenticeship degree	abstract	routine
	(1)	(2)	(3)	(4)
Δ robots per 1000 workers	0.0913 (0.045) [0.052]	-0.0912 (0.035) [0.060]	0.0660 (0.030) [0.041]	-0.0581 (0.018) [0.038]

Notes: In this table, we analyze the effect of robots on occupational quality of younger workers. The estimates are based on $N = 402$ local labor market regions (*Landkreise und kreisfreie Staedte*). The dependent variables is the change in various measures for occupation quality of workers 40 years old or less between 1994 and 2014: Share of workers with university degree (column 1), share of workers with apprenticeship degree (2), average abstract task intensity (3), and average routine task intensity (4). The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Table A.4: Disaggregating the Service Sector

	Dependent variable:					
	100 x 2014 employment in industry / total non-manuf. employment in 1994					
	(1)	(2)	(3)	(4)	(5)	(6)
[A] Broad industry groups						
	Non-Manuf.	Agg/Mining	Constr.	Cons. serv.	Business serv.	Public sect.
Δ robots per 1000 workers	0.7051 (0.275) [0.442]	0.0223 (0.021) [0.023]	-0.0026 (0.027) [0.027]	0.0410 (0.064) [0.051]	0.5934 (0.213) [0.352]	0.0356 (0.036) [0.053]

Notes: $N = 402$. In this table, the employment growth rate in the non-manufacturing sector is the contributions of different industries. The dependent variables are constructed as 100x the number of employees in 2014 in each industry relative to total non-manufacturing employment in 1994. Consequently, the coefficients in each panel sum up to the coefficient in column 7 of panel A, Table 3. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Table A.5: Change in average age

	Dependent variable:	
	change in average age between 1994 and 2014	
	Manufacturing	Non-manufacturing
	(1)	(2)
Δ robots per 1000 workers	0.7049 (0.740) [0.867]	-2.7519 (1.063) [1.722]

Notes: $N = 402$. The dependent variable is the change in the average age of workers in 1994 vs. 2014. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, own calculations.

Table A.6: Labor Share and Productivity - by vote shares of left parties

	Dependent variable: % change between 2000 and 2014	
	Labor productivity	Labor share
	(1)	(2)
[A] Above median share of left party votes in 1992 federal elections		
△ robots per 1000 workers	0.0254 (0.265) [0.296]	-0.0446 (0.144) [0.124]
[B] Below median share of left party votes in 1992 federal elections		
△ robots per 1000 workers	1.3572 (0.297) [0.657]	-0.5333 (0.115) [0.227]

Notes: $N = 198$ (Panel A) and 204 (Panel B). The period of analysis is 2000-2014 due to availability of regional GDP and employee compensation data. The dependent variable in columns (1) is the log change in GDP per worker x 100 and in column (2) the percentage point change in the total employee compensation (Arbeitnehmerentgelte) over total GDP x 100. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, BHP 7514 v1, and Federal Returning Officer, own calculations.

Table A.7: Manufacturing Adjustment - by vote shares of left parties

	Dependent variable: 100 x Number of workers in 2014 / total employment in 1994							
	Incumbent workers			Entrants				total
	yes	no	no	entered	same region,	in diff.	not	
	yes	yes	no	labor mkt.	diff. sector	region	emp.	
	yes	yes	yes	after 1994	in 1994	in 1994	in 1994	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
[A] Above median share of left party votes in 1994 federal elections								
△ robots per 1000 workers	0.3295 (0.072) [0.139]	-0.2132 (0.104) [0.107]	-0.0913 (0.026) [0.048]	0.0062 (0.115) [0.222]	-0.0285 (0.038) [0.076]	-0.0352 (0.060) [0.103]	0.0015 (0.037) [0.071]	-0.0308 (0.233) [0.476]
[B] Below median share of left party votes in 1994 federal elections								
△ robots per 1000 workers	0.0528 (0.059) [0.058]	-0.1117 (0.036) [0.092]	-0.0885 (0.017) [0.038]	-0.3396 (0.110) [0.201]	-0.0658 (0.020) [0.036]	0.0390 (0.051) [0.054]	-0.1112 (0.029) [0.056]	-0.6251 (0.189) [0.366]

Notes: $N = 198$ (Panel A) and 204 (Panel B). In this table, the employment growth rate is additively split up into the contributions of different groups of incumbent workers or workers that enter the region's manufacturing sector between 1994 and 2014. The coefficients of columns 1-7 sum up to the coefficient in column 8. The regressions include the full set of control variables as in column 4 of Table 2. Standard errors clustered at the level of 50 aggregate labor market regions in parentheses. Shift-share standard errors in brackets.

Sources: IFR, COMTRADE, EU KLEMS, BEH V10.01.00, and BHP 7514 v1, and Federal Returning Officer, own calculations.