



Robots and Workers: Evidence from the Netherlands

Daron Acemoglu
Hans Koster
Ceren Ozgen

March 2023

NBER WORKING PAPER SERIES

ROBOTS AND WORKERS:
EVIDENCE FROM THE NETHERLANDS

Daron Acemoglu
Hans R. A. Koster
Ceren Ozgen

Working Paper 31009
<http://www.nber.org/papers/w31009>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2023

We thank Matt Cole, Francesco Esposito, Maarten Goos, Ted Pinchbeck, James Rockey, Anna Salamons, Benedikt Vogt, Ulrich Zierahn, and participants at seminars in Amsterdam, The Hague, and Utrecht for valuable feedback. Daron Acemoglu acknowledges funding from the Hewlett Foundation and National Science Foundation. Ceren Ozgen gratefully acknowledges Birmingham Business School research funding and University of Birmingham Covid support fund. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w31009>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Daron Acemoglu, Hans R. A. Koster, and Ceren Ozgen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Robots and Workers: Evidence from the Netherlands
Daron Acemoglu, Hans R. A. Koster, and Ceren Ozgen
NBER Working Paper No. 31009
March 2023
JEL No. D63,E22,E23,E24,J24,O33

ABSTRACT

We estimate the effects of robot adoption on firm-level and worker-level outcomes in the Netherlands using a large employer-employee panel dataset spanning 2009-2020. Our firm-level results confirm previous findings, with positive effects on value added and hours worked for robot-adopting firms and negative outcomes on competitors in the same industry. Our worker-level results show that directly-affected workers (e.g., blue-collar workers performing routine or replaceable tasks) face lower earnings and employment rates, while other workers indirectly gain from robot adoption. We also find that the negative effects from competitors' robot adoption load on directly-affected workers, while other workers benefit from this industry-level robot adoption. Overall, our results highlight the uneven effects of automation on the workforce.

Daron Acemoglu
Department of Economics, E52-446
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
daron@mit.edu

Ceren Ozgen
Department of Economics
University House
University of Birmingham
Edgbaston
Birmingham B15 2TT
United Kingdom
c.ozgen@bham.ac.uk

Hans R. A. Koster
Vrije Universiteit Amsterdam
Department of Spatial Economics
De Boelelaan 1105
1081HV Amsterdam
the Netherlands
and Tinbergen Institute and CEPR
h.koster@vu.nl

1 Introduction

Industrial robots have spread rapidly in most advanced economies as well as in some emerging ones. Their adoption has been particularly pronounced in economies facing labor shortages and aging workforces (Acemoglu and Restrepo, 2022a). In the US, which was initially slow in introducing robotics into the production process, the number of robots per 10,000 industrial workers has increased from 35 in 1993 to 149 in 2014 to 255 in 2020. The same numbers in the Netherlands are, respectively, 12, 68 and 209 per 10,000 industrial workers. Although industrial robots have automated a variety of production tasks from painting to welding, sorting and assembling, and in many cases boosted productivity, their effects on workers are debated.

Firm-level studies on the effects of robot adoption paint a mixed picture. Most of these studies find that robot-adopting firms not only increase their productivity but also expand their employment (see, for example; Acemoglu et al., 2020 for France; Koch et al., 2021 for Spain; Dixon et al., 2021 for Canada; Humlum, 2019 for Denmark; Acemoglu et al., 2022 for the US). These firm-level outcomes reflect several forces, however. First, robot-adopting firms are typically more productive and often on a different trend than non-adopters (*e.g.* Koch et al., 2021; Acemoglu and Restrepo, 2022b). Second, adopters may be expanding at the expense of rivals in the same industry (Acemoglu et al., 2020). Because of this equilibrium effect of robots, overall industry or nation-wide employment could decline as non-adopting competitors significantly reduce employment. This is the pattern found by Acemoglu et al. (2020) and Koch et al. (2021) as well as by Bessen et al. (2020) for the Netherlands and Bonfiglioli et al. (2020) for France.¹

Studies focusing on equilibrium (industry-level) implications of robots typically find negative effects on employment and wages. For example, Acemoglu and Restrepo (2020) estimate negative impacts on workers — especially low- and mid-skill workers and those in manufacturing and in bluecollar occupations — in US local labor markets who are more exposed to the spread of industrial robots. Dauth et al. (2021) estimate similar negative wage and employment impacts in manufacturing in Germany, but the negative employment effects are smaller compared to those in the US and are compensated by local expansion of non-manufacturing employment. Acemoglu and Restrepo (2022b) estimate negative effects on wages and employment on demographic groups most exposed to automation, driven by robots and specialized software. Graetz and Michaels (2018), Acemoglu and Restrepo (2020) and Acemoglu et al. (2022) also report negative effects on the labor share at the industry level.

Nevertheless, we are far from a consensus on what types of workers are affected by robot adoption and what the impact of robotization is on *individual* workers. Using aggregate data on workers, Acemoglu and Restrepo (2020), Acemoglu et al. (2020), Dauth et al. (2021) and Humlum (2019) estimate negative effects on production workers and Bonfiglioli et al. (2020) and Barth et al. (2020) estimate negative impacts on low-skilled workers. In contrast, Aghion et al. (2021) estimate positive employment effects, even for unskilled production workers in France, while Hirvonen et al. (2022) do not find negative effects for low-skilled workers in Finland.

We contribute to this emerging literature in two ways. First, we confirm several of the important

¹Aghion et al. (2021), on the other hand, find positive firm-level and industry-level effects, focusing on various proxies of equipment investment (rather than direct measures of robot adoption).

firm-level and industry-level findings of the literature using high-quality Dutch employer-employee panel dataset on robots, firms and workers.² Robot-adopting firms increase output by about 14.9%, increase employment (hours worked) by 4.3% and reduce the labor share by 4.6 percentage points, relative to comparable non-adopting firms. The quantitative magnitudes of these estimates are very similar to those from France and Spain. As in these countries, we find negative effects on non-adopting rivals in the same industry. For example, a non-adopting firm experiences a 6.2% decline in hours worked when competitor robot adoption — that is, the share of sales by robot adopters in the same four-digit industry — increases by one standard deviation.

Second, our major contribution is to investigate the impact of robot adoption on workers, utilizing a large panel dataset of workers. In addition, to analyzing the overall impact of robot adoption on individual workers, we shed light on the heterogeneous effects of robot adoption on different types of workers. We distinguish “*directly-affected*” workers from those that are “*indirectly-affected*”. To motivate this distinction, recall that, as emphasized in [Acemoglu and Restrepo \(2020\)](#), robot adoption creates a negative *displacement effect* on workers whose tasks are being replaced. Simultaneously, it produces a positive *productivity effect*, as non-automated tasks expand, and it is reasonable to expect that indirectly-affected workers will be the main beneficiaries of this productivity effect. We construct three alternative, though complementary, measures of directly-affected workers. The first is bluecollar workers employed in routine tasks (constructed using the routine task intensity index developed in [Autor and Dorn, 2013](#) and [Koster and Ozgen, 2021](#)). Previous work has documented that these workers are more likely to perform tasks that can be more easily automated (see *e.g.* [Autor and Dorn, 2013](#); [Oesch, 2013](#)) and have tended to be more adversely affected by the adoption of automation technologies at the aggregate level (see *e.g.* [Acemoglu and Restrepo, 2020](#); [Bonfiglioli et al., 2020](#); [Barth et al., 2020](#)). The second measure is based on the replaceability index of [Graetz and Michaels \(2018\)](#) and, similarly, captures workers in occupations that can be more easily replaced by automation. The third measure simply focuses on the highest completed level of education by a worker. We further motivate the choice of these three measures in [Section 2.2.2](#).

Using all three measures, we find that workers who do not perform routine production tasks indirectly gain from robot adoption, while routine production workers, workers in replaceable occupations, or low-education workers lose out. These patterns are similar when we look at the effects of robot adoption on non-adopting rivals. The negative effects of robot adoption on workers employed in routine production work and replaceable occupations is predominantly through lower wages. The much smaller impacts on employment are broadly consistent with the idea that rigidities may be leading to slower or even muted quantity adjustments in the Dutch labor market.

Our discussion so far has already placed our work in the context of the recent literature. Here we only add that our paper is distinguished by the use of high-quality, longitudinal data on robot adoption matched to a panel of employer-employee administrative data and by the length of the period covered. We build our comprehensive measure of firm-level robot adoption and worker-level outcomes by linking *International Trade Register* data to firm-level *Production Statistics* and to

²Our definition of robots corresponds to the code 8479500 in the international trade codes of commodities, which is defined as *industrial robots*, not elsewhere specified or included. According to the International Standards Organisations, an industrial robot is an actuated mechanism programmable in two or more axes, with a degree of autonomy, moving within its environment, to perform intended tasks.

the worker-level *Tax register*. In the Dutch context we are able to do this for the period covering 2009-2020, which gives us a longer sample than in Acemoglu et al. (2020) and, more importantly, we are able to study *worker*-level outcomes. The use of actual, longitudinal robot data also distinguishes our paper from Aghion et al. (2021) for France and Bessen et al. (2020) for the Netherlands, which use proxies for automation; from Acemoglu et al. (2022) who use cross-sectional data on automation technologies and robots for the US; and from Hirvonen et al. (2022) for Finland, who focuses on a variety of advanced equipment, which includes other automation and non-automation technologies as well as robots.

This paper continues as follows. In Section 2 we outline the data construction and introduce our summary measures based on the task content of occupations. Section 3 analyses firm-level outcomes, followed by worker-level outcomes in Section 4. Section 5 concludes.

2 Data

2.1 Data description

This study benefits from a number of administrative datasets provided by the *Statistics Netherlands*. We combine a number of datasets namely: *Production Statistics*, *Tax Registers*, the *International Trade Register*, *Labor Force Surveys*, *Investment Statistics* and the *Firm Register*.

2.1.1 Firm-level data

Production Statistics constitute the core of our analysis on firms. They include very detailed firm-level information on firms' production input/outputs such as number of employees, value added, sales, total costs, personnel costs and total wage bill. The dataset contains all firms that have 50 employees and above, and a representative sample of firms smaller than 50 employees per year for the 2000-2020 period. We observe around 55 thousand unique firms per year. We focus on manufacturing firms yet we use a broader definition of the manufacturing industry that includes manufacturing, energy, water and waste, construction, mining, and transportation.

We link *Production Statistics* to the *Tax Registers*, which is based on the employers' tax declarations. It includes employees that are employed by formally registered firms. Hence, self-employed that do not work at formally registered firms are not included. We observe the monthly wages and hours worked of around 10 million employees per year. By linking *Production Statistics* to *Tax Registers*, we construct a near universe employer-employee dataset (LEED) dataset on active firms in manufacturing industry and their employees over time.

Following the literature we calculate the labor share as the total wage costs over gross value added (GVA). We set the labor share to missing if it is larger than one.³

Tax Registers include two main job related measures that are annual earnings before tax and hours worked in a year. From this we calculate hourly wages. In the analysis, to ensure comparability we

³This holds for about 2% of the cases. These firms are slightly less productive, older and export less in real terms. However, in terms of robot adoption they are similar to the firms in our sample, so removing these firms is unlikely to introduce any selection issues.

drop firms from the LEED dataset where at least one of the following variables; GVA, labor share, sales or total hours worked are coded as missing. Moreover, firms with more than 25 thousand workers are also dropped from the dataset. These selections decrease the number of observations by about 30%. This reduction is mainly caused by the limited coverage of *Production Statistics* of the small firms with less than 50 employees. However, this should not be a major problem as we will show that essentially only large firms are robot adopters.

The *International Trade Register (ITR)* includes all trade transactions in the Netherlands with other countries at the firm level from 2009 onwards, yet makes a distinction between within-EU trade and non-EU trade. With respect to trade with non-EU countries, the information is gathered from the customs data. With respect to trade within the EU countries, Statistics Netherlands runs their own survey called *Intrastat*. Enterprises that import and/or export goods to the EU in *total* in excess of € 1.2 million in a year are required to specify the exact commodity code of the goods they traded and with which member state. Overall, the *ITR*, including the *Intrastat* survey, roughly contains 80% of total Dutch imports and exports (in value) that can be attributed to a firm.

We can trace robot importing firms based on the specific commodity code, 847950, in line with the international trade codes of commodities. We define robot-adopting firms as firms that have cumulative imports of robots exceeding the median value of robot imports in our dataset, which is € 2,500. Imports below this value are unlikely to be referring to significant capacity of industrial robots that may be influential enough to change the course of production.⁴ An important concern is if due to the threshold value of € 1.2 million we are missing out significant number of robot imports from the EU countries. One advantage of our data is that the threshold value applies only to *total* imports value of a firm in a year from an EU country, meaning that the commodity code registration is not exclusively linked to the value of a single item imported. In other words, when a firm imports from *e.g.* France, for each item we would know the commodity code even as small as € 200 unless firm's total number of imports from France remains under less than € 1.2 million in a year. This would make it very unlikely that we will be missing out major robot imports from within the EU, as our dataset mostly consists of 50 employees or more, which easily trade more than € 1.2 million a year with EU countries. However, for example for trade with *e.g.* Japan, we would observe every single item imported and its respective purchase value. Moreover, robot production in the Netherlands is negligible, therefore we are not likely to miss out significant robot adopters in the country by focusing on robot imports. Similarly, if there are large firms that would import and sell robots in the domestic market, they are unlikely to be listed as a manufacturing firm, but rather as a wholesale firm.⁵ Finally, we will show later that our results are robust to excluding firms that re-export robots.

In Appendix A.1 we discuss two other datasets that we link to *Production Statistics* data, one on investments and another on the age of firms. By combining these seven datasets, we create a thorough picture of robot-adopting firms between 2009 and 2020.

⁴We have also tested different cut-off values based on the robot import value distribution and we have used a continuous measure of robot adoption in Appendix B.1.

⁵There are some firms in the Netherlands that are robot-related service providers (rather than producers). They offer help in assembling robots or advising on setting up robotic processing infrastructure.

2.1.2 Worker-level data

After constructing the LEED data, we link it to the *Demographic* register that contains the universe of population in the Netherlands, hence information on workers' age, gender, and whether a worker is born in the Netherlands or not. The resulting worker-level dataset contains almost the universe of employees in all sectors, though in our analysis we focus on the broader manufacturing sector where robot adoption is most prevalent. We then keep the working population by dropping workers that are younger than 18 or older than 67. We further drop all observations for an employee who earned more than half a million euros; worked more than 4,380 hours; earned less than € 2.5 and more than € 500 hourly wages per annum. These selections correspond to around two standard deviations from the mean of each indicator. We further focus on workers that had a job at one employer in a given year. Our final data is a balanced panel of 333 thousand unique workers that have been employed in manufacturing sector at least once between 2009 and 2020.

For each worker we have longitudinal information on hourly wages and hours worked when the worker is employed in a certain firm. This means we know whether a worker is employed in a certain year. If a worker is not in employment in a certain year, we cannot associate her with firm characteristics. To be able to analyze the impact of robot adoption on the probability to be in employment for this worker, we assign firm characteristics of the last firm the worker has been employed.

The *Tax Registers* do not include workers who are self-employed. Hence, instead of being unemployed, workers may for example have set up their own firm. In order to obtain information on whether a worker is in fact unemployed or not, we merge in the so-called *Personal Income* data. These data include information on the total income and disaggregated income resources, such as rental income, of the universe of the population as well as the employment status. By combining our data with the *Personal Income* dataset, we are able to distinguish unemployed from those who are not participating in work due to other reasons such as retirement or study. Furthermore, we obtain the type of household, such as whether a worker lives with her partner or lives together with multiple adults on the same address.

2.1.3 Task content, education and the most affected workers

The displacement and productivity effects suggest that the impact of automation should be uneven across workers. In order to investigate the heterogeneous effects of robot adoption, we define groups of workers that are potentially directly-affected by robots. Because in the Netherlands, education and occupation levels of the employees can only be observed from the *Labor Force Surveys (LFS)*, we link our worker-level LEED dataset with observations in the 10 years prior to and including the year of observation. This implies that, although we will not have the universe of employees in our data set, we have access to a large number of workers matched to our firm-level data.

We construct three measures of directly-affected workers based on education level and task content of a job in a worker's occupation. We label workers as *directly-affected*, denoted by a_{it} , when they belong to one of these groups. The remaining workers are referred to as *indirectly-affected*, since the impact of robot adoption on them will be mostly through indirect channels, such as productivity increases, reorganization, or reallocation to new tasks.

Bluecollar-routine workers. First, we focus on the effect of robot adoption on firm and worker outcomes when bluecollar workers perform highly routine tasks. These workers are likely most impacted by robot adoption, since current robots are designed to perform routine tasks. Using *O*NET Online* and occupational codes from the LFS, we compute a routine task intensity index (RTI), following Autor and Dorn (2013)’s construct and Koster and Ozgen (2021)’s application of it to Dutch ISCO. The exact definition of RTI is in Appendix A.2.

Given the RTI index, bluecollar-routine workers are defined as follows:

$$a_{it} = \max_{\tau=-9, \dots, t} (I(\mathcal{B}_{i o \tau} = 1) \times I(\mathcal{RTI}_{i o \tau} > 1)), \quad (1)$$

where $\mathcal{B}_{i o \tau}$ is an indicator variable whether worker i is in a bluecollar occupation o in year τ . Similarly, $I(\mathcal{RTI}_{i o \tau} > 1)$ is an indicator function that equals one when the routine-task-intensity index exceeds 1 in τ . We use a 10-year window prior to the year of observation to match the workers to an *LFS* wave to obtain information on whether these workers are in bluecollar-routine occupations.

By adopting a 10-year window we assume that workers do not change occupations frequently. This is a plausible assumption as Visser et al. (2018) show that occupational mobility in the Netherlands is uncommon, particularly for groups that are likely affected by robots. Occupational mobility is more likely to be observed among the 18-25 year-old workers, transitioning from education to employment. There is a significant path-dependency in terms of job changes, and this trend is even stronger for occupational changes. Moreover, in the study period almost 70% of the employees have not experienced an earnings transition in consecutive years, even independently of occupational mobility (Bachmann et al., 2020).⁶ Although we believe that potential occupational mobility is not likely to affect our results, to ascertain the robustness of our worker level results, in Appendix C.1 we show similar results when we narrow the window down to one year.

According to this definition, about 11% of the workers in the Dutch broader manufacturing industry during the study period are bluecollar-routine workers.

Replaceable workers. Not all routine-bluecollar workers are equally susceptible to robot adoption. Although some occupations require the performance of highly routine tasks, they still need to be complemented by non-routine tasks which may require assessment and discretion, *e.g.* a call center agent, metal working machinist, wood cutting operator and metal driller.

To account for these differences, we construct a worker-level replaceability index at the 4-digit ISCO level. Our replaceability index is based on the description of robot applications by the *International Federation of Robots (IFR)* and occupational classifications in the US Censuses. *IFR* distinguishes the applications that can be executed by robots on the basis of tasks such as welding, assembling and painting. If an occupational title includes one of these keywords we assign the value of 1 to that occupation to indicate that workers in that occupation is replaceable by robots, as in Graetz and Michaels (2018). To apply this measure to the Dutch occupational classification, we use a crosswalk to concord the occupations from SOC to ISCO. Similar to the definition of bluecollar-routine workers

⁶An earnings transition is defined as a switch from one decile of the country- and year-specific earnings distribution to another decile

we look at a worker’s occupation within a 10 year window. Hence,

$$a_{it} = \max_{\tau=-9,\dots,t} (I(\mathcal{V}_{i\sigma\tau} = 1)), \quad (2)$$

where $\mathcal{V}_{i\sigma\tau}$ is an indicator variable whether a worker performs a replaceable job in year τ .

Low-education workers. The final measure of workers likely to be adversely affected by robot adoption is based on education. We generate a measure of low-education workers by using the educational classification in the Dutch *LFS*. For this, we assign workers to have a low education when the highest level of educational degree corresponds to secondary education. Hence, these workers would have in total a maximum of 10 years of primary and secondary education. Our measure is then:

$$a_{it} = \max_{\tau=-9,\dots,t} (I(\mathcal{E}_{i\sigma\tau} = 1)), \quad (3)$$

where $\mathcal{E}_{i\sigma\tau}$ denotes the educational classification.

2.2 Descriptive statistics

2.2.1 Firm-level data

Our yearly unbalanced panel data spans 12 years and includes 162,220 firm-year observations and 46,914 thousand unique firms. We observe 218 unique robot-adopting firms (0.5%). Although only a small fraction of firms are adopting robots, they tend to be larger, and thus 6.8% of the workers in our sample are employed in a firm that adopt robots at some point during our time window.

Robot adoption primarily concentrates in the (narrowly-defined) manufacturing sector (2.1%). Other sectors with substantial robot adoption are mining (3.8%), energy (1.0%), and transport and logistics (0.6%). There is a positive secular trend in robot adoption over the 12 years both at the sector and at the firm level. For instance, the correlation between firms’ import value of robots between t and $t - 1$ is 0.76.

Because it may take time to observe the effects of robot adoption, especially in a highly-regulated labor market like the Netherlands where laying off workers is costly and time-consuming, the effects at the firm level may take place with long and variable lags, and thus we also look at long-differences models, focusing on the years 2009 and 2020 compare it to the 12-waves firm panel. The long-differences sample now includes 3,989 unique firms, 1.1% of which have adopted robots. Table A1 in Appendix A.3 reports descriptive statistics for the 2-wave balanced panel, indicating very similar values to those in Table 1.

Table 1 presents descriptives of the main variables of interest for the unbalanced panel of firms between 2009 and 2020. The histograms of main variables of interest are shown in Appendix Figure A2. This will be our main dataset throughout the paper from which we will make further selections depending on the type of the analysis. In Panel A, we present descriptives for firms that adopt robots sometime between 2009 and 2020. A comparison of robot-adopting firms with non-adopting firms in Panel B indicates that, as expected, the former are, on average, larger, produce much higher value added, pay higher wages and have a larger workforce. It also shows that robot-adopting firms generate more than 10 times as much GVA than non-adopting firm. The average number of

TABLE 1 – SUMMARY STATISTICS OF 12-WAVE UNBALANCED PANEL 2009-2020

	mean	std. dev.	5 th perc.	Median	95 th perc.	N
PANEL A: Robot-adopting firms	(1)	(2)	(3)	(4)	(5)	(6)
Gross value added in market prices (<i>in 1000 €</i>)	85,844	309,884	1,405	12,935	368,738	1,712
Hours worked	912,129	1,788,000	30,360	279,157	3,924,000	1,712
Number of workers	528.1	1,019	24	164	2,200	1,712
Labor share	0.528	0.183	0.210	0.535	0.819	1,712
Total wage bill (<i>in 1000 €</i>)	30,927	76,213	765	6,588	137,633	1,712
Mean hourly wage (<i>in €</i>)	28.01	25.30	14.74	23.68	44.15	1,696
Robot adopter	0.591	0.492	0	1	1	1,712
Competition by robot adopters	0.117	0.239	0	0.00694	0.823	1,712
PANEL B: Non-adopters	(1)	(2)	(3)	(4)	(5)	(6)
Gross value added in market prices (<i>in 1000 €</i>)	7,866	65,214	121	1,737	23,463	160,508
Hours worked	137,441	526,061	2,082	45,985	453,288	160,508
Number of workers	85.93	321.1	2	30	279	160,508
Labor share	0.553	0.189	0.198	0.575	0.841	160,508
Total wage bill (<i>in 1000 €</i>)	3,397	16,664	42	950	11,496	160,508
Mean hourly wage (<i>in €</i>)	24.35	31.19	10.40	19.62	40.50	156,753
Robot adopter	0	0	0	0	0	160,508
Competition by robot adopters	0.0248	0.0963	0	0	0.146	159,982

Notes: Panel A reports summary statistics for robot-adopting firms in manufacturing sector. Panel B reports summary statistics for non-adopters in manufacturing sector. For confidentiality reasons, the min and max values cannot be reported. Competition by robot adopters refers to the share of sales by robot adopting firms within the same 4-digit industry.

workers of robot-adopter firms is more than 6 times the size of the workforce of the non-adopting firms. Interestingly, we do not find large differences in the average labor share, but we will see in our regression analysis that the labor share declines after robot adoption. In addition, almost all robot-adopting firms are exporters.

Figure 1 displays the cumulative value of robot imports versus the number of firms adopting robots over the period 2009-2020, Although the cumulative trend in all indicators is towards a steady increase over time, the imports value fluctuates significantly annually.

We explore the determinants of robot adoption more in Table 2, where we estimate simple exploratory regressions using data from 2009. The dependent variable is a dummy whether firms will adopt robots in the future. We first show the individual correlations between robot adoption and main firm level indicators, where we control for 4-digit industry and location fixed effects. We subsequently augment the model by including all firm-level variables together. In column 6 we show that robot adoption increases with GVA, while other measures of a firm’s productivity are not statistically significant determinants of the robot adoption decision. This means that essentially only firm size is determining robot adoption.

Figure 2 shows that in 2020, more than 35% of all robots were adopted by firms in the top 2.5% of the distribution in terms of value added, confirming that mostly large firms are adopting robots.




FIGURE 1 – CUMULATIVE ROBOT ADOPTION

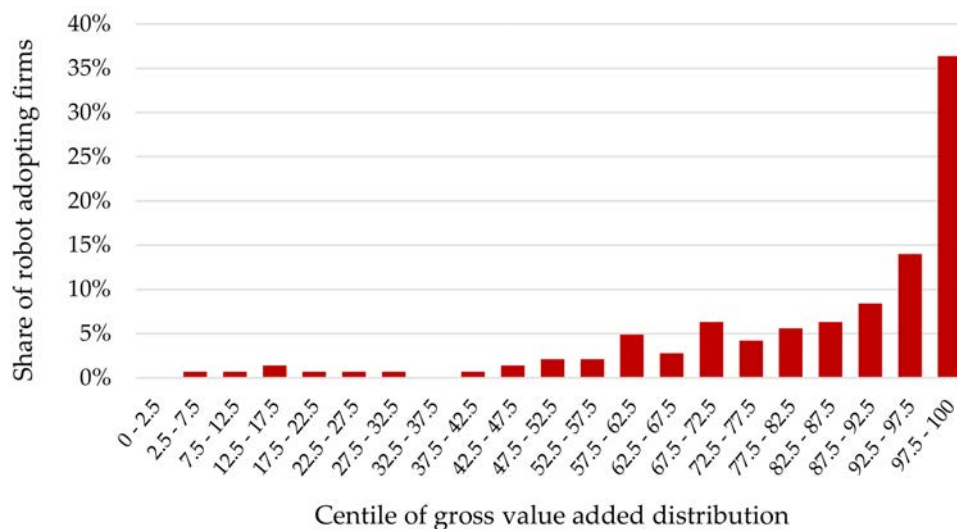


FIGURE 2 – ROBOT ADOPTION BY VALUE ADDED PERCENTILES

2.2.2 Worker-level data

The descriptive statistics reported in Table 3 are for the matched firm-worker data. We keep workers that appear at least once in an *LFS*-wave during our study period.

Overall, 6.1% of the employees work in a robot-adopting firm. The mean hourly wage and annual earnings of employees in robot adopters are € 32 and € 65,841, that are 30% higher than those in non-adopters. The employee characteristics are in general similar between robot-adopting and non-adopting firms, except for lower share of low-education workers (about 50%); lower share of replaceable workers (about 25%) and lower share of immigrant workers (about 15%) in robot

TABLE 2 – ROBOT ADOPTION BY FIRMS

<i>Dependent variable:</i>	<i>Robot adopter</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
Gross value added (<i>log</i>)	0.008*** (0.001)					0.015*** (0.005)
Gross value added per hour (<i>log</i>)		-0.000 (0.001)				
Labor share			-0.000 (0.006)			0.007 (0.017)
Hourly wage (<i>log</i>)				0.004** (0.002)		-0.009 (0.006)
Hours worked (<i>log</i>)					0.006*** (0.001)	-0.006 (0.005)
4-digit industry fixed effects	✓	✓	✓	✓	✓	✓
Municipality fixed effects	✓	✓	✓	✓	✓	✓
Number of observations	6,784	6,784	6,784	6,593	6,784	6,593
R^2	0.207	0.195	0.195	0.200	0.205	0.213

Notes: The dependent variable takes the value of 1 when a firm adopts robots any time between 2009-2020. The regressions are estimated only for the year 2009. All regressions include 4-digit industry and municipality fixed effects. We exclude gross value added per hour in column 6 as to avoid collinearity. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

adopters.

About 10% of the workers in the full data are classified as bluecollar-routine workers (*i.e.*, those who are in an occupation with an RTI value exceeding 1 and in a bluecollar occupation). The share of this type of workers is not very different between robot adopters and other firms (9.2% versus 10.5%). The share of replaceable workers follows a different pattern. 10.4% of the workers are replaceable in non-adopting firms, while this value is 7.8% in non-adopters. Similarly, low-education workers represent 34.5% of the workforce, but this share is only 18.8% in robot adopters. The summary statistics is consistent with the idea that robot-adopting firms have more skilled workforces. This skill differential is also part of the explanation for why hourly wages are higher among robot adopters. Finally, more than 80% of manufacturing workers are male, regardless of whether a firm adopts robots or not.


Following workers over time within the study period, on average $1 - 0.953 = 4.7\%$ of the workers become unemployed. This rate is 3.1% that is about 25% lower for workers who were previously employed in robot-adopting firms.

We now further motivate our three definitions of directly-affected workers, documenting that these workers are indeed more likely to be adversely impacted by robot adoption. We plot trends in the *total* hours worked and hourly wage of bluecollar-routine workers, replaceable workers, low-education workers, and all workers in Figure 3 by tapping into data from LFS linked to Tax Registers from 2001 onwards. In Panel A we depict the share of workers by worker type in the last 20 years. There is clearly a substantial decrease in the share of all directly-affected worker types, with the share of replaceable workers and bluecollar-routine workers declining by about 45% by 2020.

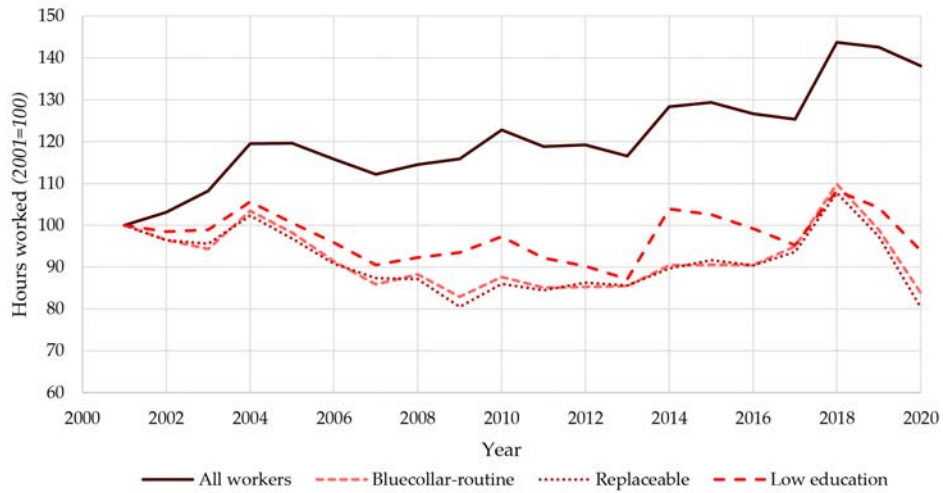
TABLE 3 – SUMMARY STATISTICS OF MATCHED LFS SAMPLE OF WORKERS 2009-2020

	<i>mean</i>	<i>std. dev.</i>	<i>5th perc.</i>	<i>median</i>	<i>95th perc.</i>	<i>N</i>
PANEL A: Workers in robot-adopting firms	(1)	(2)	(3)	(4)	(5)	(6)
Mean hourly wage (<i>in €</i>)	32.09	19.94	14.01	27.66	62.92	122,439
Hours worked	1,858	499.3	590	2,076	2,179	126,169
Employed	0.969	0.174	1	1	1	11,023
Personal income (<i>in €</i>)	65,841	42,924	25,055	56,898	131,139	120,885
Robot adopter	0.615	0.487	0	1	1	126,169
Competition by robot adopters	0.0906	0.175	0	0.000336	0.453	113,743
Bluecollar-routine worker	0.0923	0.289	0	0	1	74,943
Replaceable worker	0.0788	0.269	0	0	1	82,532
Low-education worker	0.188	0.390	0	0	1	88,789
Male	0.831	0.374	0	1	1	126,169
Age	46.03	10.64	27	47	62	126,169
Migrant	0.111	0.314	0	0	1	126,169
2 nd generation migrant	0.166	0.372	0	0	1	126,169
Household type – single	0.174	0.379	0	0	1	120,885
Household type – couple	0.822	0.382	0	1	1	120,885
Household type – other	0.00342	0.0584	0	0	0	120,885
PANEL B: Workers in non-adopters	(1)	(2)	(3)	(4)	(5)	(6)
Mean hourly wage (<i>in €</i>)	23.62	15.05	11.07	20.10	45.89	1,661,936
Hours worked	1,773	623.5	109.3	2,043	2,303	1,734,836
Employed	0.953	0.211	1	1	1	1,566,314
Personal income (<i>in €</i>)	48,205	31,224	16,662	42,514	94,951	1,616,555
Robot adopter	0	0	0	0	0	1,734,836
Competition by robot adopters	0.0318	0.109	0	0	0.205	1,582,445
Bluecollar-routine worker	0.105	0.306	0	0	1	1,102,731
Replaceable worker	0.104	0.305	0	0	1	1,143,449
Low-skilled worker	0.357	0.479	0	0	1	1,223,219
Male	0.813	0.390	0	1	1	1,734,836
Age	45.78	11.64	25	47	63	1,734,836
Migrant	0.0853	0.279	0	0	1	1,734,836
2 nd generation migrant	0.139	0.346	0	0	1	1,734,836
Household type – single	0.195	0.396	0	0	1	1,616,555
Household type – couple	0.799	0.401	0	1	1	1,616,555
Household type – other	0.00603	0.0774	0	0	0	1,616,555

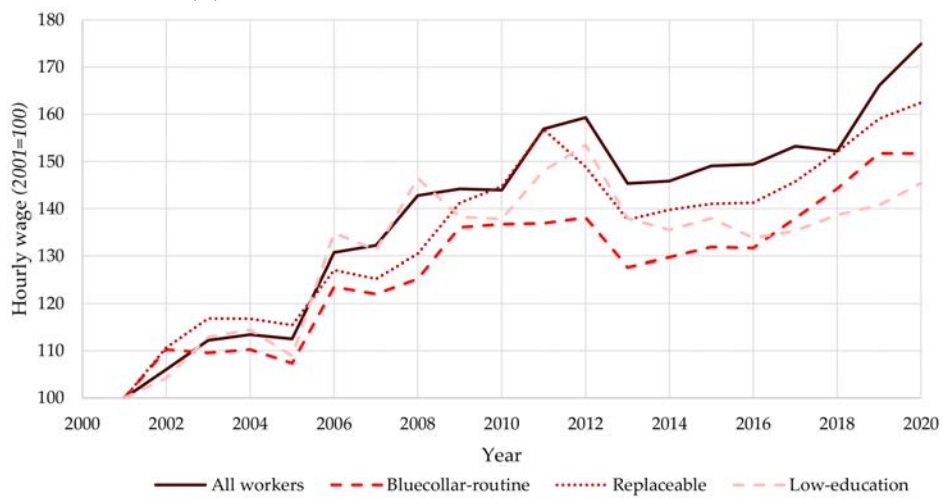
Notes: The data include workers that are in manufacturing sector and appear in an *LFS* wave at least once in a 10-year window including the year of observation. Panel A reports summary statistics for workers in robot-adopting firms in manufacturing sector. Panel B reports summary statistics for workers in non-adopters in manufacturing sector. Competition by robot adopters refers to the share of sales by robot adopters within the same 4-digit industry. For confidentiality reasons, the min and max values cannot be reported.



(A) SHARE OF WORKERS BY WORKER TYPE



(B) TOTAL HOURS WORKED BY WORKER TYPE



(C) MEAN HOURLY WAGE BY WORKER TYPE

FIGURE 3 – TRENDS IN HOURS WORKED AND HOURLY WAGE BY WORKER TYPE

Figure 3b shows that there is an overall increase, of about 25%, in total hours worked since 2001, while hours worked by workers performing routine tasks in bluecollar occupations decreased by 40%. The same pattern can be seen for for replaceable and low-education workers.

Figure 3c reports the trends for (nominal) hourly wages. It indicates that average hourly wages have grown relatively fast, by about 75% between 2001 and 2020. This average masks significant heterogeneity, with slower growth for bluecollar-routine and replaceable workers than the rest. Wage growth is even slower for low-education workers. Our subsequent analysis sheds light on whether robot adoption has been a contributing factor to this slower wage growth in the Dutch economy.

3 Firm-level evidence on the effects of robot adoption

This section presents our baseline firm-level results. We focus on the effects of robot adoption on (gross) value added, the labor share, the hourly wage and hours worked both for robot-adopting firms and their competitors. Section 3.1 outlines our econometric framework, Section 3.2 reports our main estimates for robot-adopting firms, followed by a discussion in Section 3.3 on robustness of the results to relaxing various assumptions. Section 3.4 turns to the effects of robot adoption on competitors.

3.1 Econometric framework

We present both long-differences regressions, focusing on 11-year changes, and panel data (fixed effects) estimates using annual data. As in Graetz and Michaels (2018) and Acemoglu et al. (2020), the advantage of the long-differences specification is that it focuses on a time horizon during which most of the (potentially slow-acting) effects of robot adoption may be realized. By contrast, the fixed effects estimates use all of the available data and thus exploit all of the yearly variation in the sample. Hence, we find it useful to look at both sets of estimates.

Let y_{fmt} denote one of our four dependent variables (gross value added, the labor share, the hourly wage and and hours worked) for firm f located in municipality m in year t . Then, our long-differences estimation equation is:

$$\Delta y_{fmt} = \beta \Delta r_{fmt} + \zeta x_{fmt} + \lambda_{f \in s} + \mu_m + \epsilon_{fmt}, \quad (4)$$

where Δ denotes the change between \underline{t} and t , spanning the years of 2009 and 2020 and r_{fmt} indicates whether a firm is a robot adopter, as defined in Section 2.1. In addition, the x_{fmt} 's are firm-level control variables in the first year of observation \underline{t} , including the log of number of workers and the log of value added per worker. $\lambda_{f \in s}$ are 4-digit industry fixed effects, and μ_m capture location fixed effects. Using the same notation, our panel data specification is:

$$y_{fmt} = \beta r_{fmt} + \zeta_t x_{fmt} + \kappa_f + \lambda_{f \in s, t} + \mu_{mt} + \epsilon_{fmt}, \quad (5)$$

where x_{fmt} again denote beginning-of-sample control variables (which are not time-varying but we estimate time-varying coefficients ζ_t to allow for trends in x_{fmt}), κ_f are firm fixed effects, $\lambda_{f \in s, t}$ are sector-by-year fixed effects and μ_{mt} are municipality-by-year fixed effects. Note that there are about 330 municipalities and 500 4-digit sectors in the SBI sector classification in the Netherlands.

We include these controls and fixed effects to mitigate the issue that firms that adopt robots have underlying characteristics that are different and may therefore be on different trends.⁷

The competition variable is defined on the basis of the share of sales in a given 4-digit industry accounted for by robot adopters (leaving out the sales of the own firm in question). Specifically, we define *robot adoption by competitors* as

$$r_{ft}^c = \frac{\left(\sum_{f \in s} q_{ft} r_{ft}\right) - q_{ft} r_{ft}}{\left(\sum_{f \in s} q_{ft}\right) - q_{ft}}, \quad (6)$$

where r_{ft} is our usual robot adoption measure at the firm level and q_{ft} denotes firm sales.⁸ Using this variable, we estimate analogues of equations (4) and (5), except with r_{ft}^c on the right-hand side and focusing on non-adapter firms. As in this case the identifying variation comes from the differences in competition *between* 4-digit industries we cannot include 4-digit industry-by-year fixed effects, and only include 2-digit industry-by-year fixed effects.

Additionally, one may be concerned that industry robot adoption can be endogenous, for example, because it is correlated with other technological investments in the same industry.⁹ There also may be attenuation in the estimates of the effects of competitors' robot adoption, since we do not have product-level sales information. Motivated by these concerns, we next report instrumental-variables (IV) estimates.

We follow the strategy in [Acemoglu and Restrepo \(2020\)](#) and exploit the variation coming from a five-year lag of industry-level robot adoption in South Korea and Taiwan. These two countries are further ahead than the Netherlands in terms of adoption and are not directly competing with Dutch firms. Specifically, using *IFR* data, we construct the following exposure variable as instrument:

$$r_{\tilde{s}t}^{\mathcal{E}} = \frac{\mathcal{R}_{\tilde{s},t-5} - \mathcal{R}_{\tilde{s},t-5}}{n_{\tilde{s},t-5}}, \quad (7)$$

where \tilde{s} refers to the *IFR* sector, $\mathcal{R}_{\tilde{s},t}$ is the total number of robots in Korea and Taiwan in sector \tilde{s} in year t , and $n_{\tilde{s},t}$ is the total employment in sector \tilde{s} in the Netherlands in t . Because $r_{\tilde{s}t}^{\mathcal{E}}$ has some extreme outliers we cap the instrument at its 99th percentile value.¹⁰ We additionally control for 2-digit industry-by-year fixed effects, intended to purge any differential industry productivity trends. There may still be productivity trend differences *within* 2-digit industries, but in this case, to the extent that productivity is positively associated with robot adoption, such residual correlation would bias our estimates upwards. Instead, our estimates point to sizable negative effects from

⁷We also perform several sensitivity analyses to investigate whether omitted variable bias is an issue. These analyses include the inclusion of leads and lags of robot adoption to equation (5). We also obtain [Oster's \(2019\)](#) bias-adjusted estimates.

⁸Note that we do not have detailed information on the product composition of different firms, and hence are using coarser information to construct the competition variable than in [Acemoglu et al. \(2020\)](#).

⁹As discussed above, the adoption of other technologies does not appear to be correlated with robot adoption at the firm level, though there may be other technological or organizational changes at the industry level that may still confound the effects of robots.

¹⁰The *IFR* data are from 2004-2014. Hence, for 2020 we would need data for 2015. We predict robots in 2015 by the linear trend of robot adoption in each sector in each country between 2010 and 2014. In [Appendix B.7](#) we also provide estimations with non-extrapolated 2015 wave. Our results remain robust to baseline predictions.

TABLE 4 – FIRM-LEVEL EVIDENCE FOR THE EFFECTS OF ROBOT ADOPTION

<i>Dependent variable:</i>	ΔGVA (log)	$\Delta Labor$ share	$\Delta Hourly$ wage (log)	$\Delta Hours$ worked (log)
PANEL A: Long-differences				
Robot adopter	0.257*** (0.085)	-0.065** (0.031)	-0.002 (0.049)	0.076 (0.062)
Firm-level control variables	✓	✓	✓	✓
4-digit industry fixed effects	✓	✓	✓	✓
Municipality fixed effects	✓	✓	✓	✓
Number of observations	4,298	4,298	4,227	4,298
R^2	0.513	0.503	0.444	0.0403
<i>Dependent variable:</i>	GVA (log)	$Labor$ share	$Hourly$ wage (log)	$Hours$ worked (log)
PANEL B: Fixed effects				
Robot adopter	0.139*** (0.031)	-0.046*** (0.010)	0.011 (0.018)	0.042** (0.021)
Firm-level control variables	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930
R^2	0.984	0.848	0.838	0.985

Notes: We weight all regressions by total hours worked in the firm in 2009. In Panel A reports the estimates based on 2009 & 2020 waves. We add the log of number of workers in t as well as the log value added per worker in t as controls. Panel B reports the estimates based on 2009-2020 and includes year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

competitors' robot adoption.

3.2 Firm-level impact of robot adoption

Table 4 presents our baseline firm-level results for manufacturing in the Netherlands.¹¹ Panel A reports results for our long-differences specification from equation (4) on a balanced panel of firms. All regressions are weighted by total hours worked in the firm in 2009, and include the log of number of workers at time t as well as the log value added per worker at time t as controls.

The results are comparable to those in Acemoglu et al. (2020). Compared to France, we report somewhat larger effects of robot adoption on value added ($\exp(0.257) - 1 \times 100\% = 29\%$ vs 9% in France), on the labor share (-6.5 percentage points vs -2.7 percentage points), and on total hours worked (7.9% vs 5.5%). The decrease in labor share by 6.5 percentage points from robot adoption is

¹¹Our results are based on firms that are within the broader manufacturing sectors. The results are very similar when we focus only on manufacturing sector firms. Moreover, we do not find sizable heterogeneity in firm outcomes between different sectors (*i.e.*, construction, energy, manufacturing, mining, transport, water and waste).

also in line with [Kehrig and Vincent \(2020\)](#) who show a 5 percentage points decline of labor share in US manufacturing firms and with [Koch et al. \(2021\)](#) who show a 7 percentage points decline of labor share in Spanish manufacturing firms. We do not detect any effects on wages, which may be because Dutch manufacturing firms are able to expand employment without putting upward pressure on wages (and we return to this issue in worker level analysis).

In Panel B in [Table 4](#) we report estimates from equation [\(5\)](#) using our unbalanced annual panel. Because we control for the number of workers and gross value added per worker in 2009, we only include firms that we observe in 2009. As a result, the number of observations vary from those in summary statistics table. In these panel data regressions, value added increases by $(\exp(0.139) - 1) \times 100\% = 15\%$, while the labor share decreases by 4.6 percentage points and hours worked increase by 4.3%. As in [Koch et al. \(2021\)](#) and [Acemoglu et al. \(2020\)](#), the effects on wages are positive but insignificant. Since these estimates are close to the long-difference results, but are considerably more precise, in the rest of the paper we focus on these panel results.¹²

3.3 Firm-level impact of robot adoption — robustness

A central question is whether our estimates, and other similar ones in the literature, are due to robot adoption or other technologies that may be introduced at the same time as robots (see [Bessen et al., 2023](#)). Our *Investments* data, described in [Appendix A.1](#), allow us to separate investments in IT and other technologies and explore this issue. In [Table 5](#) we estimate analogous specifications to those reported in Panel B of [Table 4](#), augmented with controls for investments in computers and machinery. Similar to our robot adoption dummy, we define these investments by a dummy variables that equals one when a firm’s investments in computers or machines in the past years were at least once among the top 5% of all investments in that year. These controls do not affect our estimates for the impact of robot adoption, though they tend to increase value added and employment.¹³ These findings are consistent with [Koch et al. \(2021\)](#) who also find industrial robots are the only advanced technology that reduce the labor share.

We also checked the robustness of our estimates to several modifications of our baseline specification in [Appendix B](#).

We performed several robustness checks to make sure that the robot effects can be interpreted as causal effects. First, we provided a sensitivity analysis of the definition of our robot adoption dummy by substantially changing the threshold regarding the cumulative value of robot imports. We also used a continuous measure of robot value by using the cumulative value of robot imports.

Second, we undertook event studies showing that there are no pre-trends in the variables of interest. We further re-estimate the baseline specification by including leads and lags of robot adoption. These leads and lags are mostly statistically insignificant and do not materially influence the coefficients

¹²The similarity between the long-differences and panel estimates suggests that most of the effects of robot adoption are realized rather quickly. This conclusion is also backed up by the fact that leads and lags of robot adoption do not appear to be significant in panel regressions and the event-studies also suggest an immediate effect (see [Appendix B.2.1](#)).

¹³As an additional robustness check we also estimated regressions where we control for the cumulative investments in computers and machines, which did not materially influence the results presented here. These results are presented in [Table 5](#).

TABLE 5 – CONTROLLING FOR INVESTMENTS
IN COMPUTERS AND MACHINES

<i>Dependent variable:</i>	<i>GVA</i>	<i>Labor</i>	<i>Hourly</i>	<i>Hours</i>
	<i>(log)</i>	<i>share</i>	<i>wage (log)</i>	<i>worked (log)</i>
	(1)	(2)	(3)	(4)
Robot adopter	0.151*** (0.031)	-0.046*** (0.010)	0.010 (0.018)	0.053*** (0.020)
Computer investment	0.067*** (0.010)	0.000 (0.003)	-0.011* (0.006)	0.072*** (0.007)
Machine investment	0.122*** (0.012)	-0.005 (0.0042)	0.003 (0.007)	0.100*** (0.010)
Firm-level control variables	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930
R^2	0.985	0.848	0.838	0.985

Notes: The computer/machine investment dummy equals one when a firm’s investments in past years were at least once among the top 5% investments in computers/machines in that year. We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

of interest.

Third, we considered a ‘placebo’ treatment where we randomized the timing of robot adoption for firms that will adopt robots in the future. We find zero effects of these placebo-treatments.

Fourth, as an omnibus measure against various omitted variable bias issues, we present bias-adjusted estimates, following Oster (2019), with and without detailed fixed effects.¹⁴ These results do not alter our baseline predictions.

Fourth, we address concerns related to a possible violation of the stable unit treatment variance assumption (SUTVA) in our econometric framework. In particular, non-robot adapters may be indirectly impacted by the introduction of this technology, and we will explore this issue explicitly in the next subsection. As an alternative strategy, we also confirm that the results are very similar when all non-adopting firms in the same 3-digit industry are excluded from the sample.

Fifth, we also checked against issues of ‘negative weights’, which can arise in models with two-way fixed effects models. We verify that the results are very similar when we use a weighted least squares estimator that includes interactions of the fixed effects with the year of adoption, so that there is no staggered treatment within groups.

¹⁴Specifically, this approach corrects for any possible effects of unobservables not included in a regression. Building on Altonji et al. (2005), it looks at the relationship between the set of covariates included in regression and the coefficient estimate of interest, and presumes that adding covariates tends to reduce the degree of omitted variable bias. In addition, Oster’s approach looks at not just how much the coefficient of interest moves, but also the change in the variance explained after adding the controls.

TABLE 6 – FIRM-LEVEL EVIDENCE FOR THE EFFECTS
OF ROBOT COMPETITION, 2SLS ESTIMATES

<i>Dependent variable:</i>	ΔGVA <i>(log)</i>	$\Delta Labor$ <i>share</i>	$\Delta Hourly$ <i>wage (log)</i>	$\Delta Hours$ <i>worked (log)</i>
	(1)	(2)	(3)	(4)
Competition by robot adopters	-0.464 (0.389)	-0.139 (0.181)	0.0329 (0.211)	-0.623*** (0.187)
Firm-level control variables	✓	✓	✓	✓
2-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	70,313	70,313	69,701	70,292
Kleibergen-Paap <i>F</i> -statistic	15.06	15.06	15.09	15.06

Notes: These estimations exclude the robot-adopting firms. Competition by robot adopters refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). We weight all regressions by total hours worked in the firm in 2009. We include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses and clustered at the IFR-industry×year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Sixth, we verified that unweighted results are very similar to our baseline estimates that are weighted by firm size in 2009. We also obtained similar results when including non-manufacturing firms (meaning all firms in all sectors) in our sample. We further extended the sample back to 2004, assuming zero robot adoption before 2009. The baseline predictions survive these checks.

Seventh, we investigate the robustness of our results to the € 1.2 million threshold value that allows Dutch firms trading within the EU not to report commodity codes for the imported goods. In particular, we excluded imports from all countries that are below this threshold and verified that the threshold value has no effect on our results.

Eighth, we checked the robustness of our results to the possible re-exporting of robots by domestic firms. Finally, we estimated the effects of robot adoption separately for large and small firms. None of these checks materially change our findings.

3.4 Effects of robot adoption on competitors

In this subsection, we study the effects of robot adoption on competitors. As explained above, we limit the sample to non-adopting firms and look at the effects of robot adoption in their 4-digit industry, and we instrument this variable with exposure to robots in the same industries in Taiwan and South-Korea five years earlier. The first-stages, reported in Appendix B.5, are precisely estimated, and the Kleibergen-Paap *F*-statistic is about 15 in all specifications.

Our main results, reported in Table 6, indicate sizable negative effects from robot adoption on competitors, though these estimates are sometimes imprecise. In column 1, for example, a one standard deviation increase in robot adoption in the firm’s 4-digit industry reduces value added by $(\exp(-0.464 \times 0.0993) - 1) \times 100\% = 4.5\%$. The imprecision may be partly due to the fact that our measures of competition are coarser than those in Acemoglu et al. (2020).

Column 2 shows a negative but small and insignificant effect on the competitors’ labor share (a one standard deviation increase in competitors’ robot adoption reduces labor share by 1.38 percentage points). Column 3 does not detect statistically significant effects on wages, although the standard error is too large to draw strong conclusions. Finally, we find more precise negative impacts on hours worked: a one standard deviation increase in competitors’ robot adoption reduces hours worked by 6%.

Positive effects on adopting firms and negative effects on competitors combined imply that industry-level implications of robots are ambiguous in general. If we focus on the more precise estimates in column 4, we find that the overall effects are slightly negative, because the negative impacts on competitors are larger, and thus overall hours worked in the industry declined by about 2.7%. These negative effects are broadly consistent with past work, such as [Graetz and Michaels \(2018\)](#), [Koch et al. \(2021\)](#), [Acemoglu and Restrepo \(2020\)](#) and [Acemoglu et al. \(2020\)](#).

In [Appendix B.7](#) we confirm the robustness of these results to variations of the instrument. Specifically, we show that the results are robust to lagging robot values by six years instead of five, as to avoid extrapolation of the data to 2015 (in this case we have to exclude 2009 because the *IFR* data are only available from 2004 onwards). Further, we show that results are robust to using values of the instrument at time t , in which case we extrapolate the data to 2020. Finally, we show that the results are essentially unaffected if we also include *IFR* data from Hong Kong and Singapore to construct our instrument.

4 Worker-level analysis

We next turn to our main focus: the effects of robot adoption on workers. In addition to confirming the main outlines of our and other authors’ firm-level results, our high-quality employer-employee data enable us to investigate which types of workers are negatively impacted by robot adoption. [Section 4.1](#) outlines the econometric framework we use for investigating worker-level effects. [Section 4.2](#) turns to heterogeneous effects of robot adoption on different types of workers, while [Section 4.4](#) studies the heterogeneous effects of competitors’ robot adoption.

4.1 Econometric framework

Let w_{ift} and h_{ift} denote, respectively, hourly wage and total hours for employee i working at firm f in year t . Then the main relationships of interest are:

$$\{\log w_{ift}, \log h_{ift}\} = \beta r_{ft} + \zeta z_{it} + \kappa_f + \lambda_{t,f \in s} + \mu_{f \in m,t} + \nu_i + \epsilon_{fmt}, \quad (8)$$

where z_{it} are worker characteristics such as age and immigration background. We further include firm fixed effects κ_f to control for the fact that more productive workers may be more likely to be employed in high-productivity firms, which are in turn more likely to adopt robots. As in the firm-level analysis, equation (8) also includes 4-digit industry-by-year and municipality-by-year fixed effects to address the issue that more productive firms may be more likely to adopt robots.

Finally, given the nature of our data we can follow workers over time and include worker fixed effects, ν_i . Specifications that include worker fixed effects focus on the impact of robots adoption

on the same worker and are particularly useful, since Table 2 provided suggestive evidence of endogenous sorting of workers across robot-adopting and non-adopting firms (with adopting firms having better-paid workers on average). In all specifications standard errors are clustered at the firm-year and worker levels.

Our main interest in this section, however, is not the overall impact of robot adoption on workers, captured by the parameter β , but heterogeneous effects. In particular, as explained above, we are interested in the differences between directly-affected workers (who are subject to the direct displacement effects of robot adoption) and indirectly-affected workers (who should generally benefit from the indirect productivity effects, which induce additional hiring and non-automated tasks). We will use the three measures of directly-affected workers (based on workers performing blue-collar-routine tasks, performing replaceable tasks and having low education), as defined in equations (1), (2) and (3). The econometric specification in this case can be written as

$$\{\log w_{ift}, \log h_{ift}\} = \beta_1 r_{ft} a_{ift} + \beta_2 r_{ft} (1 - a_{ift}) + \delta a_{ift} + \zeta z_{it} + \kappa_f + \lambda_{t,f \in s} + \mu_{f \in m,t} + \nu_i + \epsilon_{fmt}, \quad (9)$$

where a_{it} is an indicator for whether the worker is directly affected. We control for the direct effects of a_{it} , firm fixed effects, industry-year fixed effects, and worker fixed effects. To further assuage concerns related to endogeneity, we estimate a version of (9) where we include firm-year fixed effects, which enables us to control for all direct effects of robot adoption on firms and focus on differential impacts on directly-affected workers *within* firms.

Beyond hours worked and wages, robot adoption may also change the employment status of these workers. To study the effects of robots on the probability of employment, we estimate the relationship between being employed, denoted by the dummy variable, e_{ift} , and firm-level robot adoption. Specifically, we estimate the following linear probability model:

$$e_{ift} = \beta_1 r_{ft} a_{it} + \beta_2 r_{ft} (1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \kappa_f + \lambda_{t,f \in s} + \mu_{f \in m,t} + \nu_i + \epsilon_{fmt}. \quad (10)$$

As usual, the specification excludes the retired, self-employed and students. We also assign the last employer's characteristics to workers who are currently unemployed but were previously employed.

We use analogous models to study the impacts of competitors' robot adoption on directly-affected and indirectly-affected workers. Specifically, we estimate models of the following form:

$$\{\log w_{ift}, \log h_{ift}, e_{ift}\} = \gamma_1 r_{ft}^C a_{it} + \gamma_2 r_{ft}^C (1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \kappa_f + \lambda_{t,f \in s} + \mu_{f \in m,t} + \nu_i + \epsilon_{fmt}, \quad (11)$$

where r_{ft}^C captures robot adoption in the same 4-digit industry. Let $\lambda_{t,f \in s}$ now denote 2-digit sector-by-year fixed effects. In line with the firm-level results, we instrument for r_{ft}^C using the robots exposure instrument as defined in equation (7). To address the issue that the instrument varies only at the IFR-industry level, we cluster our standard errors at the IFR-industry-by-year and worker levels.

4.2 The effects of robot adoption on workers

In Table 7 we first report the average effects of robot adoption on hourly wages, the employment probability and hours worked. Our main estimations focus on workers matched with the *LFS*

surveys, hence provide occupation information of directly-affected workers. For consistency, we focus on the same set of workers in Table 7 as well, although we do not yet use the information on workers' occupations.

In columns 1-3 of Table 7 we estimate that robot adoption is associated with an average increase in hourly wages of 2.5%. Notably, this estimate is essentially the same regardless of whether a battery of worker characteristics (in particular, age, gender, migrant background, household type) are included, as we do in columns 2 and 3. When we additionally include worker fixed effects, the impact is smaller (hourly wages increase by 1.6%). We interpret this smaller effect to be indicative of the endogenous sorting of workers — whereby more productive workers tend to work for more productive firms that are more likely to adopt robots — generating a slight overestimate of the positive effect of robot adoption on hourly wages.

Columns 4-6 turn to the impact of robot adoption on employment. Here, we do not detect consistent significant effects, which may again reflect the rigidities in the Dutch labor market, where laying workers off can be difficult and slow.¹⁵ In columns 7-9, we look at hours worked, where we detect negative and fairly stable estimates. For example, without worker-level controls, hours worked declined by about 1.7%, and when detailed worker-level controls are added, this negative effect is about 1.3% (see column 8). With worker fixed effects in column 9, the impact is in the same ballpark but larger, -2.1% . This negative wage impact at the worker level is in line with the aggregate negative employment effects in the literature. However, controlling for worker fixed effects halves the initial estimates in the hourly wage regressions.

The negative implications for hours worked, combined with positive wage impacts, already suggests that there may be heterogeneous effects from robot adoption — some workers getting pay increases, while others have their hours cut.

We next turn to the heart of our worker-level analysis by allowing for differential effects on directly-affected and indirectly-affected workers. In Table 8 we focus on hourly wages. Different columns focus on different controls and our three measures of directly-affected workers. The pattern is fairly clear and confirms our conjecture about the juxtaposition of negative hours effects and positive wage effects. For example, in columns 1-3, using the definition based on workers performing bluecollar-routine tasks, we find precisely-estimated and sizable positive impacts on indirectly-affected workers, and negative and equally precisely-estimated impacts on directly affected workers. Quantitatively, the estimate in column 1 implies that robot adoption increases hourly wages of indirectly-affected workers by $(\exp(0.034) - 1) \cdot 100\% = 3.5\%$. At the same time, directly-affected workers suffer from hourly wage declines of about 5.5%. The patterns are similar in column 2 when we include worker fixed effects.

In column 3, when we include firm-year fixed effects, we can only estimate the differential impact on directly-affected workers, which is estimated to be about 2.3% for directly-affected workers compared

¹⁵Dutch firms operate under strict rules of dismissal. For example, grounds for dismissal include: (i) when an employee's performance is not satisfactory, then a firm has to prove that it has informed the employee about this and has given the employee sufficient opportunities to improve their performance; (ii) if an employee has serious conscientious objections to the business activities and the firm is unable to offer alternative work; (iii) if an employee is disabled for 2+ years etc., for more details, see <https://business.gov.nl/running-your-business/staff/dismissing-staff/grounds-for-dismissal/>.

TABLE 7 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>			<i>Employed</i>			<i>Hours worked (log)</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter	0.025*** (0.005)	0.023*** (0.005)	0.016*** (0.004)	-0.005 (0.007)	-0.005 (0.006)	0.000 (0.003)	-0.017*** (0.006)	-0.013** (0.006)	-0.021*** (0.006)
Worker-level variables		✓	✓		✓	✓		✓	✓
Worker fixed effects			✓			✓			✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
4-digit industry×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	1,778,509	1,655,095	1,601,266	1,679,993	1,679,993	1,636,397	1,778,509	1,655,095	1,601,266
R^2	0.384	0.457	0.918	0.145	0.165	0.614	0.215	0.287	0.697

Notes: The table reports results from a regression of worker-level effects of firm-level robot adoption. Worker-level variables include age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults. Standard errors are clustered at the firm×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE 8 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON HOURLY WAGE – HETEROGENEITY

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>								
	<i>Bluecollar-routine workers</i>			<i>Replaceable workers</i>			<i>Low-education workers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter× indirectly-affected worker	0.034*** (0.007)	0.015*** (0.005)		0.032*** (0.007)	0.014*** (0.005)		0.049*** (0.007)	0.018*** (0.005)	
Robot adopter× directly-affected worker	-0.054*** (0.015)	-0.008 (0.008)	-0.023*** (0.008)	-0.061*** (0.014)	-0.014 (0.009)	-0.026*** (0.009)	-0.069*** (0.012)	-0.013* (0.007)	-0.034*** (0.008)
Directly-affected worker	-0.175*** (0.003)	-0.001 (0.007)	0.002 (0.006)	-0.178*** (0.002)	-0.007 (0.006)	-0.005 (0.006)	-0.189*** (0.002)	-0.015*** (0.005)	-0.013*** (0.005)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm×year fixed effects			✓			✓			✓
Number of observations	771,681	731,290	702,987	806,390	764,276	736,273	793,550	753,123	724,914
R^2	0.476	0.937	0.948	0.477	0.937	0.948	0.498	0.937	0.948

Notes: The table reports the worker-level heterogeneous hourly wage effects of firm-level robot adoption. Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the firm×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

to other employees of the firm at the same time.

The results are quite similar in columns 4-6 when we use the replaceable worker definition of [Graetz and Michaels \(2018\)](#), with the main exception being that in column 5, when we include worker fixed effects, the negative impact on the directly-affected workers becomes imprecise and is no longer statistically significant at the 5% level. The results are also broadly similar in columns 7-9, when we focus on low-education workers. In this case, too, there are precisely-estimated positive impacts for indirectly-affected workers, and significant and again fairly precisely-estimated negative implications for directly affected workers.

Table 9 turns to the implications of robot adoption for employment and hours worked. In Panel A, we find small positive employment impacts on indirectly-affected workers and negative effects on directly-affected workers. For example, the estimates that control for worker fixed effects with the replaceable worker and low-education worker measures (columns 5 and 8) are positive for indirectly-affected workers, and of larger magnitude, though less precisely estimated for the directly-affected workers. We interpret the general imprecision of the results for employment to be again related to rigidities in the Dutch labor market, which tend to slow down or prevent worker layoffs and also discourage or slow down hiring. With all three measures, when we include firm-year fixed effects (columns 3, 6 and 9) we estimate statistically significant differential impacts for directly-affected workers.

In Panel B of Table 9 we turn to effects on hours worked. In this case, the results are less clear-cut. In some specifications, we estimate negative impacts on both directly-affected and indirectly-affected workers, though in specifications with worker fixed effects, the magnitudes are larger for directly-affected workers. One reason for the less clear-cut nature of these results may be that our measures of who is directly affected may not fully capture which workers will be reallocated towards tasks with lower hours.

4.3 The effects of robot adoption on workers — robustness

First, in the previous, we match workers in LEED-data to workers in the past 10 LFS waves to obtain their occupational information. This matching increases the number of observations but may exacerbate measurement error, *e.g.* because workers may have switched occupations in the meantime. To address this issue, in Appendix C.1, we reduce this matching window down to one year, which reduces the number of observations but largely addresses the issue of measurement error. The results show that the effect sizes are somewhat larger, with hourly wages of directly affected workers decreasing by 3-6% in relative terms once a firm adopts robots. However, the standard errors become at least twice as high, and the baseline estimates are not significantly smaller than the new results.

Second, we offer a more detailed analysis of the impact of robotization by worker skill groups and confirm in Appendix C.2 that, when we distinguish between medium and low-education workers, the negative effects of robot adoption are more pronounced for the lowest-education category.

Finally, in Appendix C.3 we investigate the effects of robot adoption on a fourth measure, *personal income*, which combines information from all three of our measures (wage, extensive margin of

TABLE 9 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON EMPLOYMENT – HETEROGENEITY

Dependent variable:	Employment								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL A: Employment									
Robot adopter × indirectly-affected worker	0.003 (0.004)	0.004 (0.003)		0.002 (0.004)	0.005* (0.003)		0.003 (0.004)	0.007** (0.003)	
Robot adopter × directly-affected worker	-0.005 (0.008)	-0.010 (0.008)	-0.014* (0.008)	-0.003 (0.008)	-0.010 (0.008)	-0.016** (0.008)	-0.006 (0.007)	-0.009 (0.006)	-0.016*** (0.006)
Directly-affected worker	-0.005*** (0.001)	0.006 (0.008)	0.003 (0.006)	-0.004*** (0.001)	0.004 (0.006)	-0.004 (0.006)	-0.007*** (0.001)	-0.002 (0.004)	-0.000 (0.004)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm × year fixed effects			✓			✓			✓
Number of observations	771,016	736,034	708,116	805,725	769,235	741,566	792,772	757,842	730,012
R ²	0.161	0.596	0.629	0.158	0.595	0.627	0.159	0.596	0.628
Dependent variable:	Hours worked (log)								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL B: Hours worked									
Robot adopter × indirectly-affected worker	-0.024*** (0.008)	-0.015** (0.007)		-0.022*** (0.007)	-0.012* (0.007)		-0.019** (0.007)	-0.009 (0.007)	
Robot adopter × directly-affected worker	-0.011 (0.011)	-0.027* (0.014)	-0.019 (0.015)	-0.021* (0.012)	-0.034* (0.017)	-0.026 (0.018)	-0.037*** (0.011)	-0.032** (0.013)	-0.023* (0.014)
Directly-affected worker	0.003 (0.003)	-0.010 (0.012)	-0.020* (0.012)	0.000 (0.003)	-0.005 (0.012)	-0.013 (0.012)	0.002 (0.002)	0.007 (0.008)	0.004 (0.008)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm × year fixed effects			✓			✓			✓
Number of observations	771,681	731,290	702,987	806,390	764,276	736,273	793,550	753,123	724,914
R ²	0.295	0.731	0.766	0.291	0.728	0.763	0.293	0.729	0.764

Notes: The table reports the worker-level heterogeneous employment effects of firm-level robot adoption. Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry × year fixed effects and municipality × year fixed effects). Standard errors are clustered at the firm × year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

employment captured by our employment dummy, and the intensive margin represented by hours worked). The results in this case confirm the pattern shown so far: robot adoption increases the personal income of indirectly-affected workers (by about 1.5% in our preferred specification with worker fixed effects) and reduces the income of directly-affected workers (by about 1.5% with worker fixed effects).

4.4 The effects of competitors' robot adoption on workers

Our results in the previous section suggest that the most negative effects of robot adoption may be on competitors, and hence at the worker level we may expect these results to fall on directly-affected workers employed in non-adopting firms whose competitors are intensively investing in robots. In this subsection, we provide evidence consistent with this expectation.

Table 10 reports estimates of the overall effects of robot adoption by competitors on hourly wage, employment status and hours worked of workers. Once again, we instrument for competitors' robot adoption, as in equation (7). The relevant first stages are reported in Appendix C.4 and continue to show a strong relationship, with Kleibergen-Paap F -statistics exceeding 10 in all specifications.

Going back to Table 10, the overall effect of competitors' robot adoption is quite similar to the effects of robot adoption by one's own firm. There are positive impacts on hourly wage, no effects on employment, and imprecise, though typically negative effects on hours worked. The increase in the hourly wage is somewhat unexpected. One possible explanation is that the increase in the hourly wage among adopting firms, shown in Table 7, puts upward pressure on the wages of the employees at non-adopting firms. Whether this is the case or not can be more easily understood once we look at heterogeneous effects, which we turn to next.

Table 11 explores heterogeneous effects. The patterns are consistent with our overall interpretation, though in some specifications somewhat imprecise. In sum, we find positive hourly wage impacts from competitors' robot adoption on indirectly-affected workers, and negative impacts for directly-affected workers. For example, in column 1, where we look at the measure based on bluecollar-routine work, we find that one standard deviation increase in competitors' robot adoption increases hourly wages by 3% for indirectly-affected workers. By contrast, the impact on the hourly wages of directly-affected workers is negative, even if imprecisely estimated. In column 3, when we include firm-year fixed effects, we estimate a fairly precise negative differential impact on the directly-affected workers. The pattern using the other two measures of who is directly affected are quite similar.

Overall, these patterns are in line with the interpretation we offered for the results in Table 10: firms whose competitors are investing in robots find themselves in a double squeeze. The demand for workers employed in non-automated tasks goes up among their competitors, forcing them to increase wages, while they are also experiencing lower demand for their products, as their competitors expand at their expense.

Table 12 turns to the effects of competitors' robot adoption on employment and hours worked. These results are less precisely estimated. Almost in all of our specifications, the differential impact on directly-affected workers are negative, though never significant at conventional levels in this table.

TABLE 10 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>			<i>Employed</i>			<i>Hours worked (log)</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Competition by robot adopters	0.211*** (0.074)	0.276*** (0.073)	0.285*** (0.076)	0.051 (0.102)	0.071 (0.099)	0.030 (0.050)	-0.045 (0.087)	0.001 (0.084)	-0.145 (0.093)
Worker-level variables		✓	✓		✓	✓		✓	✓
Worker fixed effects			✓			✓			✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	1,504,477	1,399,702	1,347,437	1,426,511	1,426,511	1,384,170	1,504,477	1,399,702	1,347,437
Kleibergen-Paap <i>F</i> -statistic	67.03	66.35	57.11	68.14	68.14	56.04	67.03	66.35	57.11

Notes: Competition by robot adoption refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). Worker-level variables include age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults. Standard errors are clustered at the IFR-industry×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE 11 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION ON HOURLY WAGE – HETEROGENEITY

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>								
	<i>Bluecollar-routine workers</i>			<i>Replaceable workers</i>			<i>Low-education workers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Competition by robot adopters× indirectly-affected worker	0.258** (0.117)	0.572*** (0.175)		0.286** (0.116)	0.536*** (0.170)		0.396*** (0.120)	0.666*** (0.184)	
Competition by robot adopters× directly-affected worker	-0.109 (0.136)	0.165 (0.185)	-0.353*** (0.097)	-0.113 (0.142)	0.058 (0.202)	-0.433*** (0.101)	-0.176 (0.136)	0.149 (0.185)	-0.449*** (0.098)
Directly-affected worker	-0.161*** (0.004)	0.015** (0.007)	0.014** (0.007)	-0.162*** (0.004)	0.010 (0.008)	0.011 (0.007)	-0.170*** (0.003)	-0.003 (0.006)	-0.004 (0.005)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	668,181	628,689	603,115	694,639	653,510	628,200	683,591	644,062	618,596
Kleibergen-Paap <i>F</i> -statistic	29.04	12.25	62.29	28.26	11.42	54.63	28.31	11.10	83.88

Notes: These estimations exclude the robot-adopting firms. Competition by robot adoption refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the IFR-industry×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE 12 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION ON EMPLOYMENT – HETEROGENEITY

Dependent variable:	Employment								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL A: Employment									
Competition by robot adopters× indirectly-affected worker	0.107 (0.092)	0.015 (0.080)		0.101 (0.091)	-0.007 (0.082)		0.076 (0.090)	0.003 (0.084)	
Competition by robot adopters× directly-affected worker	0.160 (0.100)	-0.122 (0.112)	-0.092 (0.072)	0.114 (0.102)	-0.155 (0.126)	-0.089 (0.071)	0.109 (0.092)	-0.135 (0.101)	-0.055 (0.051)
Directly-affected worker	-0.006*** (0.002)	0.013 (0.009)	0.008 (0.008)	-0.005** (0.002)	0.011 (0.008)	0.004 (0.008)	-0.009*** (0.001)	0.002 (0.005)	0.001 (0.005)
Control variables variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	668,448	634,200	608,988	694,976	659,299	634,309	683,804	649,606	624,500
Kleibergen-Paap F -statistic	28.33	11.72	55.49	27.86	10.88	48.82	27.83	10.60	76.22
Dependent variable:	Hours worked (log)								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL B: Hours worked									
Competition by robot adopters× indirectly-affected worker	-0.057 (0.118)	-0.251 (0.214)		-0.029 (0.116)	-0.267 (0.212)		-0.025 (0.120)	-0.306 (0.220)	
Competition by robot adopters× directly-affected worker	0.130 (0.137)	-0.345 (0.247)	-0.003 (0.167)	0.090 (0.139)	-0.340 (0.269)	-0.084 (0.162)	-0.052 (0.125)	-0.287 (0.251)	0.122 (0.168)
Directly-affected worker	-0.004 (0.004)	-0.002 (0.013)	-0.014 (0.013)	-0.0031 (0.004)	0.005 (0.013)	0.001 (0.014)	0.002 (0.003)	0.007 (0.009)	0.002 (0.009)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	668,181	628,689	603,115	694,639	653,510	628,200	683,591	644,062	618,596
Kleibergen-Paap F -statistic	29.04	12.25	62.29	28.26	11.42	54.63	28.31	11.10	83.88

Notes: These estimations exclude the robot-adopting firms. Competition by robot adoption refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the IFR-industry×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Finally, in Appendix C.3 we show that the results for personal income are very similar to our hourly wage results—positive impacts on indirectly-affected workers and negative effects for directly-affected workers. The differential effects are quite sizable. For example, a one standard deviation increase in competitors’ robot adoption has a differential negative impact on personal incomes of directly-affected workers of about 4.6%.

5 Conclusions

Despite the rapid spread of robots in most industrialized nations and some emerging economies, there is still much controversy about their effects. Previous work has focused on either market- or industry-level outcomes, or on firm-level outcomes. Much of this work finds negative market-level effects from robots on employment and wages, but positive firm-level effects. Robot-adopting firms benefit, in part, from the ability to expand their business at the expense of their competitors. The latter is consistent with negative industry-level effects. However, positive effects of robot adoption may also reflect pre-existing differential trends between adopting and non-adopting firms. This literature has not focused on worker-level outcomes, and particularly, on which types of workers are positively or negatively impacted by robot adoption.

This paper, for the first time investigates the *worker-level* implications of robot adoption using high-quality data on robot imports, spanning a longer time period than most other studies. We combine these with detailed linked employer-employee data from the Dutch manufacturing sector. The Dutch economy provides an interesting context, since it has invested in automation technologies rapidly, but at the same time is subject to various labor market regulations and rigidities that may protect workers in the face of automation.

We first confirm that the firm-level effects of robot adoption are very similar in the Netherlands to those we observe in other industrialized economies. In particular, robot-adopting firms increase their value added and employment, and reduce their labor share. This overall pattern and the quantitative magnitudes of our estimates are very similar to those presented in [Acemoglu et al. \(2020\)](#) for France and [Koch et al. \(2021\)](#) for Spain. Moreover, as in French and Spanish manufacturing, these positive effects on adopting-firms are associated with negative impacts on competitors. Similarly to the French case, our estimates suggest that the negative effects are somewhat larger than the positive ones, so overall industry employment declines following robot adoption.

The main contribution of the paper is to estimate the effects of robot adoption on worker outcomes. Our detailed data enable us to construct several measures concerning which workers are likely to be more negatively impacted by robot adoption. Specifically, task-based frameworks imply that workers performing tasks that will be replaced by robots will suffer from adoption, while workers employed in complementary tasks may benefit, as higher productivity translates into greater demand for skills associated with these tasks. We use three measures of which types of workers are going to be more directly affected and thus likely to suffer the negative consequences of robot adoption ([Acemoglu and Restrepo, 2020](#); [Acemoglu et al., 2022](#)). These are: workers employed in bluecollar-routine tasks, those employed in replaceable tasks (as defined in [Graetz and Michaels, 2018](#)), and workers with low education levels. Consistent with theoretical expectations, using all three measures we find that robot adoption either by own employer or by competitors has more negative effects on

directly-affected workers. For example, robot adoption by one's own employer leads to higher hourly wages for indirectly-affected workers, but to lower hourly wages for directly-affected workers.

Several questions and areas call for future inquiry. One important set of issues relates to the role of labor market institutions. Although our estimates are similar to those from other countries, the Dutch labor market is more rigid than those of many other industrialized nations and restricts firms' ability to adjust both employment and wages. Investigating the role of labor market institutions in mediating the effects of automation technologies is an important and interesting area for future work. Secondly, more granular data on market structure and competition patterns would be very useful for understanding how the adoption of automation technologies (and more broadly other new technologies) affects employees currently working for competitors. Third, although our paper shows that robot adoption may increase labor market inequality, future research may further delve into the inequality implications of robots and other automation technologies. Recent work by [Acemoglu and Restrepo \(2022b\)](#) documents substantial inequality impacts from the adoption of automation technologies in the US labor market. It would be interesting to investigate how these effects may or may not be different under more rigid labor market institutions. Finally, an open area of inquiry is whether there are other technologies, such as those creating new tasks, which firms can adopt simultaneously with robots that might have more favorable implications for workers.

References

- Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P. and Zolas, N. (2022), ‘Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey’, *Working Paper Number CES-22-12* .
- Acemoglu, D. and Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, in O. Ashenfelter and D. Card, eds, ‘Handbook of Economics, vol. 4’.
- Acemoglu, D., Lelarge, C. and Restrepo, P. (2020), ‘Competing with Robots: Firm-level Evidence from France’, *American Economic Review: Papers and Proceedings* **110**, 383–388.
- Acemoglu, D. and Restrepo, P. (2020), ‘Robots and Jobs: Evidence from US labor markets’, *Journal of Political Economy* **128**(6), 2188–2244.
- Acemoglu, D. and Restrepo, P. (2022a), ‘Demographics and Automation’, *The Review of Economic Studies* **89**, 1–44.
- Acemoglu, D. and Restrepo, P. (2022b), ‘Tasks, automation, and the rise in u.s. wage inequality’, *Econometrica* **90**(5), 1973–2016.
- Adermon, A. and Gustavsson, M. (2015), ‘Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975–2005’, *Scandinavian Journal of Economics* **117**(3), 878–917.
- Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. (2021), ‘What Are the Labor and Product Market Effects of Automation? New Evidence from France’.
- Altonji, J. G., Elder, T. E. and Taber, C. R. (2005), ‘Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools’, *Journal of Political Economy* **113**(1), 151–184.
- Autor, D. and Dorn, D. (2013), ‘The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market’, *American Economic Review* **103**(5), 1533–1597.
- Autor, D. H., Dorn, D. and Hanson, H. G. (2015), ‘Untangling trade and technology: Evidence from local labor markets’, *Economic Journal* **125**(584), 621–646.
- Bachmann, R., Bechara, P. and Vonnahme, C. (2020), ‘Occupational Mobility in Europe: Extent, Determinants and Consequences’, *De Economist* **168**(1), 79–108.
- Barth, E., Røed, M., Schøne, P. and Umblijs, J. (2020), ‘How Robots Change Within-Firm Wage Inequality’, *IZA Discussion Paper No. 13605* .
- Bessen, J., Goos, M., Salomons, A. and den Berge, W. V. (2023), ‘What happens to workers at firms that automate?’, *Review of Economics and Statistics* .
- Bessen, J., Goos, M., Salomons, A. and van den Berge, W. (2020), ‘Firm-Level Automation: Evidence from the Netherlands’, *AEA Papers and Proceedings* **110**, 389–393.
- Bonfiglioli, A., Crinò, R., Fadinger, H. and Gancia, G. (2020), ‘Robot Imports and Firm-Level Outcomes’, *CESifo Working Paper no. 8741* .
- Callaway, B. and Sant’Anna, P. H. C. (2021), ‘Difference-in-differences with multiple time periods’, *Journal of Econometrics* **225**, 200–230.
- Chaisemartin, C. D. and D’Haultfœuille, X. (2020a), ‘Difference-in-differences estimators of intertemporal treatment effects’, *Available at SSRN 3731856* .
- Chaisemartin, C. D. and D’Haultfœuille, X. (2020b), ‘Two-way fixed effects estimators with hetero-

- geneous treatment effects’, *American Economic Review* **110**, 2964–2996.
- Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. (2021), ‘The adjustment of labor markets to robots’, *Journal of the European Economic Association* **19**(6), 3104–3153.
- Dixon, J., Hong, B. and Wu, L. (2021), ‘The robot revolution: Managerial and employment consequences for firms’, *Management Science* **67**, 5586–5605.
- Goos, M., Manning, A. and Salomon, A. (2014), ‘Explaining Job Polarization: Routine-Biased Technological Change and Offshoring’, *American Economic Review* **104**(8), 2509–2526.
- Graetz, G. and Michaels, G. (2018), ‘Robots at Work’, *Review of Economics and Statistics* **100**(5), 753–768.
- Hirvonen, J., Stenhammar, A. and Tuhkuri, J. (2022), ‘New Evidence on the Effect of Technology on Employment and Skill Demand’, *ETLA Working Papers No.93* **93**.
- Humlum, A. (2019), ‘Robot Adoption and Labor Market Dynamics’, *Mimeo, Princeton University*.
- Kehrig, M. and Vincent, N. (2020), *The Micro-Level Anatomy of the Labor Share Decline*.
- Koch, M., Manuylov, I. and Smolka, M. (2021), ‘Robots and firms’, *The Economic Journal* **131**(638), 2553–2584.
- Koster, H. R. A. and Ozgen, C. (2021), ‘Cities and Tasks’, *Journal of Urban Economics* **126**(103386).
- Oesch, D. (2013), *Occupational Change in Europe*, Oxford University Press.
- Oster, E. (2019), ‘Unobservable Selection and Coefficient Stability: Theory and Evidence’, *Journal of Business and Economic Statistics* **37**(2), 187–204.
- Visser, M., Gesthuizen, M., Kraaykamp, G. and Wolbers, M. H. J. (2018), ‘Labor Market Vulnerability of Older Workers in the Netherlands and Its Impact on Downward Mobility and Reduction of Working Hours’, *Work, Aging and Retirement* **4**(3), 289–299.

Online Appendix A Data

A.1 Other datasets

We further link *Investments* data to the *Production Statistics*. This dataset is produced in the same way as *Production Statistics* and includes the population of firms for firm size 50+ workers, and a representative survey of firms for those with firm size lower than 50. This dataset allows us to observe firms' tangible and intangible investments from at least 2003 onwards. These investments are detailed by for example investments on machinery (installations, machines and devices), as well as computers and hardware (computer, data processing electronic equipment etc.).

Another dataset we use is the *Firm Register*, which is a register of the universe of firms in the Netherlands to identify the location and age of the firms at a very refined spatial scale corresponding to more or less street level; that is the 6-digit postal code.

A.2 The routine task intensity index

We construct a measure of routine task intensity (RTI) that concords Autor and Dorn (2013)'s SOC level RTI to Dutch occupations as in Koster and Ozgen (2021), at the highest possible resolution, which is 4-digit ISCO (ISCO'08). The RTI informs us on the task content of occupations workers perform, which varies within educational levels. We gather data from *LFSs* from 1996-2020. The mapping of SOC level to ISCO subdivisions enables us identifying routinization level of occupations at the lowest level of breakdown. We construct five categories of task groups based on their degree of routineness, namely: routine cognitive (\mathcal{RC}), routine manual (\mathcal{RM}), non-routine manual (\mathcal{NRM}), non-routine analytic (\mathcal{NRA}) and non-routine interactive (\mathcal{NRI}). Following Autor and Dorn (2013), let \mathcal{RTI}_{ot} be the routine task intensity of an occupation o in year t :

$$\mathcal{RTI}_{ot} = \mathcal{RC}_{ot} + \mathcal{RM}_{ot} - \mathcal{NRM}_{ot} - \mathcal{NRA}_{ot} - \mathcal{NRI}_{ot}, \quad (\text{A.1})$$

\mathcal{RTI}_{ot} is normalized to have mean zero and unit standard deviation.

A.3 Firm-level descriptives

Table A1 presents the same statistics for the 2-wave balanced panel of firms. These are the firms in our dataset that could be observed over the 12 years. The descriptive statistics shows that the summary statistics of the variables in the long-differences panel are very similar to that of the year-to-year panel.

Figure A1 shows the distributions of robot imports for firms that are above and below the EU threshold of € 1.2 million. We find mild differences in the distribution of robot imports. In any case we will show robustness of the main results to excluding firms that are below the threshold in Appendix B.1.

In Figure A2 we report descriptive statistics of the main variables of interest. The distributions of value added, labor share, and wages are essentially normally distributed. The distribution of hours worked is almost log-normally distributed, there is a spike at about 1,600 hours, which is the full-time equivalent of one worker. This means there are some firms in our dataset with one

TABLE A1 – SUMMARY STATISTICS OF 2-WAVE BALANCED PANEL 2009 AND 2020

	mean	std. dev.	5 th perc.	Median	95 th perc.	N
PANEL A: Robot-adopting firms	(1)	(2)	(3)	(4)	(5)	(6)
Gross value added in market prices (<i>in 1000 €</i>)	77,705	135,021	1,859	20,069	358,898	174
Hours worked	1,089,046	1,580,383	67,650	468,242	4,483,371	174
Number of workers	613.7	885.3	41	270.5	2,431	174
Labor share	0.557	0.184	0.236	0.569	0.874	174
Total wage bill (<i>in 1000s</i>)	33,334	54,380	1,340	10,755	143,993	174
Mean hourly wage (<i>in €</i>)	26.24	10.06	15.51	24.16	42.46	174
Robot adopter	0.569	0.497	0	1	1	174
Competition by robot adopters	0.089	0.203	0	0	0.594	174
PANEL B: Non-adopters	(1)	(2)	(3)	(4)	(5)	(6)
Gross value added in market prices (<i>in 1000s</i>)	12,579	66,718	566	3,897	41,732	8,554
Hours worked	221,003	539,126	15,991	101,526	732,507	8,554
Number of workers	132.2	315.2	11	63	428	8,554
Labor share	0.555	0.168	0.255	0.569	0.813	8,554
Total wage bill (<i>in 1000 €</i>)	5,375	14,319	303	2,130	18,739	8,554
Mean hourly wage (<i>in €</i>)	22.66	15.40	13.20	20.83	35.17	8,482
Robot adopter	0	0	0	0	0	8,554
Competition by robot adopters	0.0237	0.0833	0	0	0.126	8,535

Notes: Panel A reports summary statistics for robot-adopting firms in manufacturing sector in 2009 & 2020. Panel B reports summary statistics for non-adopters in manufacturing sector in 2009 & 2020. For confidentiality reasons, the min and max values cannot be reported. Competition by robot adopters refers to the share of sales by robot adopting firms within the same 4-digit industry.



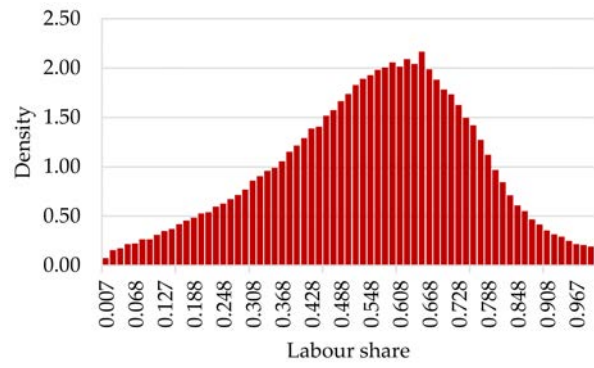


FIGURE A1 – HISTOGRAMS – ROBOT IMPORTS AND THE EU THRESHOLD

full-time employee only. However note that the regressions are weighted with firms' number of workers in 2009, so these observations are contributing very little to the estimations. Because firms often hire more than one full-time worker, the mean hours worked is obviously considerably larger as indicated in the summary statistics.




(A) GROSS VALUE ADDED (*log*)



(B) LABOR SHARE



(C) HOURLY WAGE (*log*)



(D) HOURS WORKED (*log*)

FIGURE A2 – HISTOGRAMS OF KEY VARIABLES

Online Appendix B Robustness of firm-level results

B.1 Definition and measurement of robot adoption

Our definition of robot adoption may raise some concerns. Recall that we define robot adoption such that the cumulative robot imports exceed € 2,500. One may be concerned that our results are sensitive to the precise threshold.

To address this issue, we first alter the threshold value significantly. In Table B1 we present the baseline results where the robot adoption indicator is now based on (approximately) the 75th percentile value of the cumulative robot imports of robot-adopting firms in the *ITR* data, that is €50,000. In columns 1-4 we repeat the baseline specification to estimate the impact of robot adoption on our four dependent variables. The results are entirely in line with the baseline results both in terms of magnitudes and signs of the coefficients.

Second, we use the continuous cumulative robot imports at the firm level. Since most of firms do not adopt robots, the value of robot adoption is strongly right-skewed. Therefore we use an inverse-hyperbolic-sine transformation of the cumulative value of robot imports. Columns 5-8 presents these results. Again, these results are robust to baseline estimates in terms of signs and statistical significance. However, quantitatively, the estimates are not directly comparable to the robot adoption dummy. For example, if we would interpret the coefficients related to the cumulative robot value as a percentage effect, doubling the cumulative robot value increases value added by about 1%. However, since robot imports have increased by so much (more than 10 times over the last 10 years, see Figure 1), one may question what is a meaningful increase in the cumulative value of robot imports. Hence, we prefer to stick to the robot adoption dummy as used in the main analyses.

There can also be other concerns that directly relate to the measurement of the value of imports. For example, the threshold value within EU trade to be recorded, which could introduce a selection bias such that firms trading above the €1.2 million threshold within the EU may have different characteristics, hence are differently impacted by robot adoption than those trading with non-EU countries. Therefore, in our analysis, we exclude all firms that have total import values below the EU threshold regardless of the origin of the trade partners. We show in columns 1-4 of Table B2 that using EU threshold values do not materially change the robot adoption effect on the dependent variables.

Another issue is selection among firms. One may be concerned that the comparison between firms that do not import and firms that do import is not correct. The reason is that firms, according to the data do not import, may still import goods, but in practice the value may be so low that it is not recorded, or they may import goods within the Netherlands. To circumvent this issue we exclude all firms that do not import in a certain year and are below the EU threshold of €1.2 million. This reduces the number of observations by about 75%. We show in columns 5-9 of Table B2 that the estimations excluding firms that do not import display remarkably similar results for all dependent variables, despite the number of observations being reduced to just 17 thousand.

Finally, re-exporting firms, which import robots and export robots to other countries can bias our

TABLE B1 – ROBUSTNESS FOR THE EFFECTS OF ROBOT ADOPTION ON FIRMS: MEASUREMENT OF ROBOT ADOPTION

<i>Dependent variable:</i>	<i>Higher threshold (>€50,000)</i>				<i>Continuous values</i>			
	<i>GVA</i>	<i>Labor</i>	<i>Hourly</i>	<i>Hours</i>	<i>GVA</i>	<i>Labor</i>	<i>Hourly</i>	<i>Hours</i>
	<i>(log)</i>	<i>share</i>	<i>wage (log)</i>	<i>worked (log)</i>	<i>(log)</i>	<i>share</i>	<i>wage (log)</i>	<i>worked (log)</i>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot adopter	0.204*** (0.032)	-0.060*** (0.013)	0.046*** (0.019)	0.041* (0.023)				
Cumulative robot value (<i>ih</i> s)					0.012*** (0.002)	-0.003*** (0.001)	0.001 (0.001)	0.006*** (0.002)
Firm-level control variables	✓	✓	✓	✓	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930	71,953	71,953	71,297	71,930
R^2	0.984	0.848	0.838	0.985	0.984	0.848	0.838	0.985

Notes: Columns 1-4 uses an alternative threshold value of the cumulative imports of robots to €50,000 to define robot adoption. Columns 5-8 include the inverse-hyperbolic-sine transformation (*ih*s) of the cumulative investments in robots instead. We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE B2 – ROBUSTNESS FOR THE EFFECTS OF ROBOT ADOPTION ON FIRMS: MEASUREMENT OF ROBOT ADOPTION

Dependent variable:	Remove values below 1.2 million threshold in ITR data				Remove firms that do not import				Remove re-exporters of robots			
	GVA	Labor	Hourly	Hours	GVA	Labor	Hourly	Hours	GVA	Labor	Hourly	Hours
	(log)	share	wage (log)	worked (log)	(log)	share	wage (log)	worked (log)	(log)	share	wage (log)	worked (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Robot adopter	0.143*** (0.044)	-0.053*** (0.013)	-0.022 (0.020)	0.077*** (0.024)	0.143*** (0.048)	-0.054*** (0.014)	-0.033 (0.021)	0.097*** (0.026)	0.132*** (0.038)	-0.060*** (0.012)	-0.013 (0.022)	0.033 (0.025)
Firm-level control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	32,875	32,875	32,568	32,865	17,390	17,390	17,333	17,389	70,998	70,998	70,343	70,975
R^2	0.886	0.886	0.993	0.985	0.991	0.899	0.906	0.994	0.983	0.842	0.822	0.984

Notes: Column 1-4 tests whether the minimum threshold requirement introduced by the Dutch government for the registration of imported goods creates a bias in our baseline estimates. To assess whether the importers are different than non-importers, or due to the threshold firms importing goods below the 1.2M threshold value are assigned as non-importers, column 5-8 remove the firms that do not import. Column 9-12 remove the re-exporting robot-adopting firms. We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

estimates. Those firms may not necessarily be robot adopters themselves, but just intermediaries. We think this unlikely to be a main issue as we are not aware of any large firms re-selling industrial robots, but if there were, they would be part of the wholesale sector, which is not included in our main analyses. To be safe, we exclude the firms that ever (re-)exported robots in our sample period, which applies to about 25% of the robot adopters. The results reported in columns 9-12 in Table B2 are virtually the same.

B.2 Identification issues

To further ascertain whether the robot adoption dummy captures a causal effect of robot adoption we undertake a series of checks to investigate whether our results are sensitive to various identifying assumptions.

B.2.1 Event studies

Although the estimations so far inform us on the relationship between robot adoption and firm level productivity, one may be concerned that robot adoption is endogenous and correlated to unobservable firm traits so that firm adoption captures some other shock to firm’s productivity. One way to investigate this further is to undertake event studies, which track the evolution of the dependent variables of interest before the adoption of robots.

Given that our panel covers just 12 years, event studies substantially restrict the number of observations (because, say, if we aim to analyze the robot adoption effects in 6 years we cannot observe whether a firm in 2015 will adopt robots in 2021 because our sample runs until 2020). In other words the time span we have with our data is just long enough to observe pre-adoption period changes with respect to robot adoption. However the time span is too short to take into account also the post-robot adoption effects. Therefore, we estimate the following panel data regression where dependent variables of interest are regressed on robot adoption before and after *first-time* the cumulative investments in robots exceeded the threshold of € 2,500.¹⁶ We estimate the following panel data regression where dependent variables of interest are regressed on the years before and after robot adoption, denoted by τ :

$$y_{fmt} = \sum_{\tau=-5}^0 \beta_{\tau} r_{fmt,\tau} + \gamma(x_{f\bar{m}\bar{t}} \times t) + \lambda_{ti \in s} + \mu_{mt} + \epsilon_{fmt}, \quad (\text{B.1})$$

where β_{τ} indicates the treatment effect in year τ , while robot adoption occurs in year $\tau = 0$. For example, $\tau = -1$ indicates whether the firm’s cumulative robot imports exceed the threshold of € 2,500 in the following year. The control variables included in these regressions are identical to those equation (5).

Our results reported in Figure B1 indicate that robot adoption leads to an immediate and significant productivity increase of about 11% in productivity; a reduction of 2 percentage point in the labor

¹⁶In our dataset 42% of the firms imported robots only once in the 11 year period. Only 16% of them imported robots two times. For all robot-buying firms, the first robot purchase was by far the highest value invested among almost all purchases. For example even among two or more times buyers for 67% of the firms’ first purchase was higher than their median robot investment value.

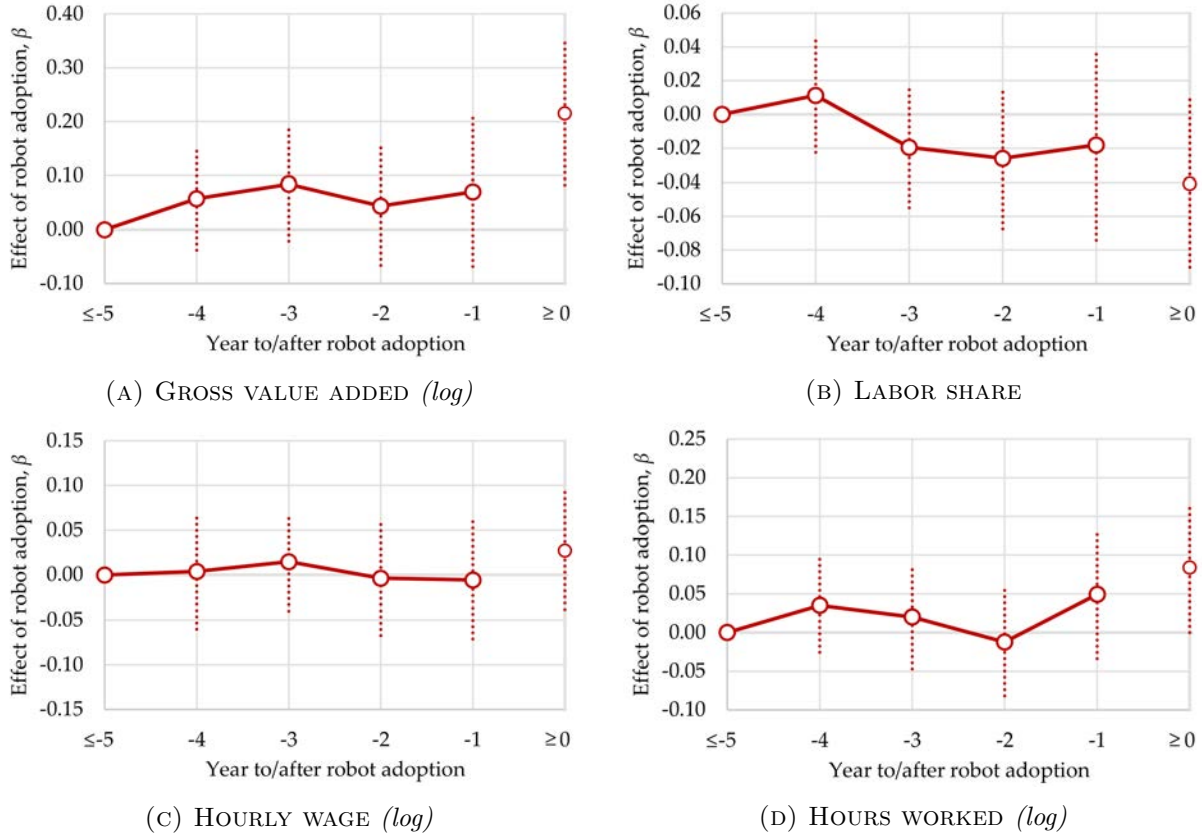


FIGURE B1 – PRE-TRENDS

Notes: The dotted lines denote 95% confidence bands based on robust standard errors.

share; a small positive effect on wages; and about a 3% increase in employment in robot-adopting firms. Our estimation for value added seems to be in line with Humlum (2019), who reports a value added increase of 13% two years after robot adoption. The trends in these variables prior to the firm’s robot adoption are largely flat. Hence, this event-study type of analysis seems to confirm that pre-trends are not a major issue.

In addition to the above, we also take a simpler approach and add three variables: whether the firm will adopt robots in $t + 1$, whether the firm has adopted robots in t , and whether the firm has adopted robots in $t - 1$. We show in Table B3 that the main effects are hardly affected. The coefficients of the leads and lags of robot adoption are mostly statistically insignificant and considerably smaller than the main effect. This suggests that robot adoption indeed implies a shock to value added and hours worked. Having said that, we find some evidence that the labor share already adjusts one year before robot adoption, although the coefficient of robot adoption in $t - 1$ is smaller than the main effect.

B.2.2 A ‘placebo’ test for robot adoption

We further explore whether observed changes in firms’ productivity is indeed due to robot adoption. In order to do this, we construct a placebo experiment. We reconstruct our firm sample such that we drop all firm-year observations for the year of robot adoption and after. This selection leaves us

TABLE B3 – ROBUSTNESS FOR THE EFFECTS OF
ROBOT ADOPTION ON FIRMS: LEADS AND LAGS

<i>Dependent variable:</i>	<i>GVA (log)</i>	<i>Labor share</i>	<i>Hourly wage (log)</i>	<i>Hours worked (log)</i>
	(1)	(2)	(3)	(4)
Robot adopter	0.152*** (0.046)	-0.048*** (0.015)	0.024 (0.024)	0.058** (0.028)
Robot adopter, $t - 1$	0.047 (0.050)	-0.034** (0.015)	0.011 (0.027)	-0.018 (0.030)
Robot adopter, t	-0.067 (0.041)	-0.006 (0.012)	-0.029 (0.023)	-0.049 (0.030)
Robot adopter, $t + 1$	-0.052 (0.037)	-0.004 (0.013)	-0.034 (0.022)	-0.017 (0.025)
Firm-level control variables	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
2-digit industry×year fixed effects				
4-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	67,020	67,020	66,377	66,997
R^2	0.985	0.854	0.841	0.986

Notes: We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

with all firm-year observations prior to robot adoption as well as those of the non-robot adopters. We then randomly assign a placebo treatment to 218 firms that will adopt robots in the future. We choose 218 as we observe 218 robot-adopting firms in 2020. This procedure is meaningful because if, say, robot-adopting firms are on increasing trends in terms of productivity we would expect to find a positive effects of the placebo-treatment. We report our findings in Table B4. We show that the ‘placebo’ robot adoption is statistically and economically insignificant in all specifications.

B.2.3 Omitted variable bias

Here we investigate the concern that our estimates can be biased due to omitted variables that correlate to robot adoption. We apply Oster’s (2019) bias-adjusted estimator. Essentially, the idea is to use coefficient movements together with changes in the R^2 after the inclusion of control variables to investigate whether omitted variable bias is important. Hence, coefficient movements alone are not a sufficient statistic to calculate the bias due to omitted variables, but the explained variance of the added control variables also matters. Oster (2019) then derives a GMM estimator to correct estimates for omitted variable bias, given the assumption that the relationship between the variables of interest and unobservables can be established from the relationship between the variables of interest and the observable control variables. In the current context, this makes sense because control variables that are added likely bear some relationship to unobservables. More specifically, we add the following firm characteristics: year-specific effects for the log of numbers of workers in 2009, the log of value added per worker in 2009, the log of cumulative investments in computers,

TABLE B4 – ROBUSTNESS FOR THE EFFECTS OF ROBOT
ADOPTION ON FIRMS: PLACEBO TREATMENT

<i>Dependent variable:</i>	<i>GVA (log)</i>	<i>Labor share</i>	<i>Hourly wage (log)</i>	<i>Hours worked (log)</i>
	(5)	(6)	(7)	(8)
Robot adopter, placebo treatment	-0.040 (0.087)	0.005 (0.036)	0.042 (0.041)	-0.085 (0.087)
Firm-level control variables	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	71,208	71,208	70,553	71,185
R^2	0.983	0.843	0.824	0.984

Notes: We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. We further exclude all observations of firms after and including the year of robot adoption. We then assign a ‘placebo’ treatment to firms that will adopt robots in the future but have not done so yet at the time of the assignment. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

machines, transport, and real estate and land, as well as the log of firm age.

There are two key input parameters that must be determined. First, a parameter must be chosen that determines the relative degree of selection on observed and unobserved variables, which is denoted by Π . Despite this parameter being fundamentally unknown, Altonji et al. (2005) and Oster (2019) show that $\Pi = 1$ is a reasonable upper-bound value. Second, there is the maximum R^2 from a hypothetical regression of the dependent variable of interest on robot adoption, and all observable and unobservable controls. $R_{\max}^2 = 1$ is then again a reasonable upper bound value, assuming that measurement error is zero. Therefore, we expect that the bias-adjusted coefficients to be similar to our baseline estimates if the added controls in our regressions are correlated with the unobservables. Given these values, we display the results in Table B5.

In columns 1 and 2 of Table B5 we show two specifications, one without 4-digit industry×year and municipality×year fixed effects and one with these fixed effects included. Since Oster (2019)’s method critically depends on the number and type of controls included, we think it is important to show parsimonious specifications as well as specifications with a more elaborate set of controls.

The coefficient in column 1 is very similar to the comparable baseline coefficient in column 1 in Panel B, Table 4, but even somewhat higher because we do not include the detailed industry-year and municipality-year fixed effects. The estimate in column 2 is essentially the same, suggesting that omitted variables do not cause a major (upward) bias in the estimated effect.

We repeat the same set of regressions but now take labor share as the dependent variable. We find negative effects in the same order of magnitude as the baseline regressions, but the coefficients are somewhat imprecisely estimated, so the bias-adjusted estimates are not significantly lower than the baseline estimates. Since the bias-adjusted estimator is derived through GMM, estimates are less

TABLE B5 – ROBUSTNESS FOR THE EFFECTS OF ROBOT ADOPTION ON FIRMS:
BIAS-ADJUSTED ESTIMATES

<i>Dependent variable:</i>	<i>GVA (log)</i>		<i>Labor share</i>		<i>Hourly wage (log)</i>		<i>Hours worked (log)</i>	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot adopter	0.186*** (0.046)	0.127*** (0.037)	-0.044 (0.030)	-0.017 (0.022)	-0.087 (0.058)	0.043 (0.060)	0.083** (0.040)	0.042 (0.035)
Firm-level control variables	✓	✓	✓	✓	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
4-digit industry×year fixed effects		✓		✓		✓		✓
Municipality×year fixed effects		✓		✓		✓		✓
Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930	71,953	71,953	71,297	71,930
R_{\max}^2	1	1	1	1	1	1	1	1
δ	1	1	1	1	1	1	1	1

Notes: We weight all regressions by total hours worked in the firm in 2009. The controls are year-specific effects for the log of numbers of workers in 2009, the log of value added per worker in 2009, the log of cumulative investments in respectively computers, machines, transport, and real estate and land, as well as the log of firm age. Bootstrapped standard errors (100 replications) are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

efficient than OLS. In the baseline regressions we did not find any statistically significant effects on hourly wages, which is confirmed by the bias-adjusted estimates in columns 5 and 6.

For hours worked we find a strong and significant positive effect on hours worked of about 8.7% (see column 7). The estimate is about halved when we control for industry-year and municipality-year fixed effects in column 8. The point estimate is essentially the same as the baseline estimate. Unfortunately, the standard errors are somewhat larger.

All in all, these results do not give the impression that omitted variable bias is a serious problem, which is given the crucial assumption that the included control variables bear some relationship with the unobservables.

B.2.4 The SUTVA assumption and negative weights

There are two more issues that deserve attention. First, one may be concerned that in the baseline regressions in Section 3.2 – where we test for the impact of robot adoption on value added, the labor share, wages and hours worked – the stable unit treatment value assumption (SUTVA) may not hold. More specifically, as shown in Section 3.4, firms that are in the same industry are affected negatively when a firm within the same industry adopts robots. Hence, the coefficients reported in Table 4 may capture both a direct effect on the treated firms as well as an indirect competition effect on the non-treated firms. However, we do not think this is a major issue because the competition effect caused by one robot-adopting firm on other firms is likely an order of magnitude smaller as compared to the direct effect of robot adoption.

Yet to investigate this further, we estimate regressions where we exclude firms that are likely to be indirectly affected by the robot-adopting firms. More specifically, we exclude all non-adopting firms that are in the same 3-digit industry as robot-adopting firms.

Unsurprisingly, this selection strongly reduces the number of observations by about 65%. Because of a lower number of observations we have to include less detailed fixed effects. More specifically, by construction we cannot include 4-digit industry \times year trends because we will not have any firms left to compare robot-adopting firms with. We therefore rather include firm and year fixed effects. The results shown in Panel A of Table B6 are remarkably similar to the baseline results. Hence, because firms in industries without any robot-adopting firms are unlikely to be strongly affected by robot adoption in other industries, these results suggest that the effects we report in Table 4 are most likely capturing direct effects, rather than indirect effects through competition.

Another issue is that of negative weights in difference-in-difference design (Chaisemartin and D’Haultfœuille, 2020a,b; Callaway and Sant’Anna, 2021). To determine treatment effects in two-way fixed effects designs like ours, a weighted sum of several difference-in-differences (DID) calculations is used to compare the changes in productivity or wage between consecutive time periods across pairs of groups. However, in some instances, the ‘control group’ may have already have been treated in the previous time period, leading to the treatment effect being differenced-out in the later period, resulting in negative weights. Negative weights are particularly an acute issue if a large share of observations in the control group are already treated and when there is substantial treatment heterogeneity. As there are only few robot adopters, the lion’s share of observations are never-treated firms. We therefore do not expect this issue to influence our results.

TABLE B6 – FIRM-LEVEL EVIDENCE FOR THE EFFECTS OF ROBOT ADOPTION:
SUTVA AND STAGGERED TREATMENT

	<i>Dependent variable:</i>	<i>GVA (log)</i>	<i>Labor share</i>	<i>Hourly wage (log)</i>	<i>Hours worked (log)</i>
PANEL A: Addressing SUTVA issues					
Robot adopter		0.134*** (0.038)	-0.031** (0.016)	0.023 (0.016)	0.055* (0.032)
Firm fixed effects		✓	✓	✓	✓
Year fixed effects		✓	✓	✓	✓
Number of observations		26,554	26,554	26,285	26,544
R^2		0.978	0.675	0.763	0.983
PANEL B: Staggered treatment in DID					
Robot adopter		0.140*** (0.030)	-0.046*** (0.010)	0.011 (0.017)	0.043** (0.020)
Firm-level control variables		✓	✓	✓	✓
Firm fixed effects		✓	✓	✓	✓
4-digit industry×year fixed effects		✓	✓	✓	✓
Municipality×year fixed effects		✓	✓	✓	✓
Number of observations		71,953	71,953	71,297	71,930
R^2		0.984	0.848	0.838	0.985

Notes: We weight all regressions by total hours worked in the firm in 2009. In Panel A reports regressions where we exclude non-robot adopting firms in the same 3-digit industries as robot-adopting firms. Panel B reports the estimates based on regressions by year of treatment. This involves separate regressions in which we keep the never-treated firms and compare them to robots firms that have adopted robots in each year. Because the control observations are duplicated, we re-eight the estimates accordingly. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

In any case, we investigate this further in Panel B of Table B6, where we present the estimates obtained from regressions conducted *by year of treatment*. For example, we estimate a separate regression including all never-treated observations as well as firms that adopt robots in 2010. Hence, this methodology involves running separate regressions for each year-of-adoption, where never-treated firms are compared to firms that have adopted robots. For a given year-of-adoption, there is no staggered treatment and the issue of negative weights is addressed. To account for the duplicated never-treated observations, the reported coefficients are re-weighted accordingly. We show that the results are essentially unaffected, which confirms that negative weights are not a pressing issue in our research design.

B.3 Sample selection and weighting

In columns 1-4 of Table B7 we report results where we do not weight the estimations with firm size. The estimates are barely affected. The effect of robot adoption on hours worked in column 4 is slightly stronger than the baseline estimate reported in Table 4, which further implies that the

TABLE B7 – ROBUSTNESS FOR THE EFFECTS OF ROBOT ADOPTION ON FIRMS: SAMPLE SELECTION AND WEIGHTING

Dependent variable:	Unweighted results				All firms				Manufacturing firms since 2004			
	GVA	Labor	Hourly	Hours	GVA	Labor	Hourly	Hours	GVA	Labor	Hourly	Hours
	(log)	share	wage (log)	worked (log)	(log)	share	wage (log)	worked (log)	(log)	share	wage (log)	worked (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Robot adopter	0.136*** (0.026)	-0.029*** (0.010)	-0.024 (0.015)	0.097*** (0.023)	0.201*** (0.026)	-0.039*** (0.089)	0.019 (0.015)	0.107*** (0.021)	0.099 (0.076)	-0.045** (0.022)	-0.005 (0.031)	0.062 (0.052)
Firm-level controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930	177,483	177,483	174,370	177,430	15,408	15,408	15,393	15,408
R^2	0.951	0.717	0.680	0.948	0.978	0.823	0.825	0.979	0.986	0.845	0.852	0.992

Notes: We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. For robustness checks, however, Columns 1-4 presents unweighted results. Columns 5-8 are weighted regressions and include all firms in all sectors. Columns 9-12 are weighted regressions and include the active firms since 2004 to extend the time period. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

effect on labor share is somewhat smaller.

Column 5-8 include all firms above and beyond manufacturing firms. This means that the number of observations increase by 250%. However, as most of the variation in robot adoption comes from manufacturing firms, the estimates are very comparable. We find somewhat stronger effects on value added and hours worked, but the results are qualitatively similar.

In columns 9-12 we investigate whether the choice of a relatively short time period of 12 years, due to data availability, affects the outcomes. For example, if the effects of robot adoption takes time, we may find underestimates of the impact of robot adoption. Therefore, in Table B7 we extend the panel of manufacturing firms to 2004 and assume that there are no robot adopters before 2009 (recall that the *ITR* data on trade transactions of robots is not available before 2009). We now weight by the number of hours worked in 2004 and control for value added per worker and number of workers in 2004. Because fewer firms are observed since 2004, the number of observations is reduced to just 15 thousand. Looking at the point estimates, they are very similar to the baseline estimates. However, because of the strong reduction in the number of observations, the coefficients capturing the impact of robots on value added and hours worked cease to be statistically significant.

B.4 Large and small firms

Here we distinguish between the effects of robot adoption on large and small firms so to assure that the effect we find is not just an effect that entirely applies to large firms. We define the large firms to be in the top 1% of firms with the largest workforce in 2009 and we define the rest of the firms as ‘small’. Table B8 reports the results.

In column 1 we show that robot adoption increases value added of large firms and small firms, although the effect is almost twice as strong for large firms (20% versus 10%). The reduction in the labor share (see column 2) is also stronger for large firms, as the reduction is 5.5 percentage points for large firms, while it is only -3.7 percentage points for small firms.

The effects on hourly wage are also very different between large and small firms. While the average effect was close to zero and statistically insignificant, we find a strong positive effects of robot adoption on hourly wages for large firms (*i.e.* 7.4%), while it is negative and sizable for small firms (*i.e.* -4.9%). One possible interpretation is that within large firms, after robot adoption, the match between new technologies and workers’ skills is likely to be better. Therefore, large firms do not need to hire new workers, but just reallocate existing workers to the new tasks, which results in higher wages. Smaller firms will need to hire workers that will be able to handle robots in production.

In column 4 we investigate the effects of robot adoption on hours worked. Here we only find statistically significant effects for small firms, which would be in line with the previous suggestion that reallocating replaceable workers within a firm is harder in smaller firms, therefore they seem to hire new workers.

In sum, these results confirm that robot adoption is not a phenomenon that only impact large firms, but also can affect the productivity of smaller firms. However, this suggestive evidence indicates that the robotization effects operate differently for these two groups of firms.

TABLE B8 – ROBUSTNESS FOR THE EFFECTS OF ROBOT ADOPTION
ON FIRMS: LARGE AND SMALL FIRMS

<i>Dependent variable:</i>	<i>GVA (log)</i>	<i>Labor share</i>	<i>Hourly wage (log)</i>	<i>Hours worked (log)</i>
	(1)	(2)	(3)	(4)
Robot adopter×large firm	0.183*** (0.040)	-0.055*** (0.014)	0.071*** (0.025)	0.002 (0.030)
Robot adopter×small firm	0.095*** (0.042)	-0.037*** (0.022)	-0.050** (0.024)	0.083*** (0.029)
Firm-level control variables	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	71,953	71,953	71,297	71,930
R^2	0.984	0.848	0.838	0.985

Notes: We define large firms to be in the top 1% of firms with the largest workforce in 2009. We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.5 First-stage results for robot competition

In Table B9 we report first-stage estimates. We regress our measure of competition by robot-adopting firms on the robot exposure instrument as defined in equation (7). We find a strong and statistically significantly positive effect of industry-level exposure to robots on competition. The coefficient indicates that a standard deviation increase in robot exposure increases the robot competition by $0.0943 \times 0.270 \times 100 = 2.6$ percentage points, which is about a quarter of a standard deviation. Hence, the effect is sizable.

B.6 WLS results for robot competition

We report the WLS competition results *without* instrumenting for robot competition in Table B10. We first show the main average effects for competition on firm level outputs in Panel A. The coefficient capturing competition effects has the opposite sign compared to robot adoption. The coefficients, however, are small and statistically insignificant.

As stressed earlier, one may be concerned that the effect of competition due to robot adoption is likely to be endogenous. For example, industrial sectors that are more productive are also more likely to adopt robots and invest in other automation technologies at the same time. Note that although this could explain why hours worked decrease, but this could not explain why we estimate a negative sign value added, though the estimation is imprecise (see columns 1 and 4). In any case, we will therefore instrument for robot competition in the specifications reported in Section 3.4.

TABLE B9 – FIRM-LEVEL EVIDENCE FOR THE EFFECTS OF ROBOT COMPETITION: FIRST-STAGE RESULTS

	(1)
<i>Dependent variable:</i>	<i>Competition by robot adopters</i>
Robots exposure	0.270*** (0.065)
Firm-level control variables	✓
Firm fixed effects	✓
4-digit industry×year fixed effects	✓
Municipality×year fixed effects	✓
Number of observations	25,816
R^2	0.858

Notes: These estimations exclude the robot-adopting firms. We weight all regressions by total hours worked in the firm in 2009 and include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE B10 – FIRM-LEVEL EVIDENCE FOR THE EFFECTS OF ROBOT COMPETITION, WLS ESTIMATES

<i>Dependent variable:</i>	ΔGVA (log)	$\Delta Labor$ share	$\Delta Hourly$ wage (log)	$\Delta Hours$ worked (log)
	(1)	(2)	(3)	(4)
Competition by robot adopters	-0.075 (0.051)	0.009 (0.017)	0.019 (0.028)	-0.056 (0.040)
Firm-level control variables	✓	✓	✓	✓
2-digit industry×year fixed effects	✓	✓	✓	✓
Municipality fixed effects	✓	✓	✓	✓
Number of observations	71,145	71,145	70,503	71,124
R^2	0.981	0.810	0.800	0.982

Notes: These estimations exclude the robot-adopting firms. Competition by robot adopters refers to the share of sales by robot adopting firms within the same 4-digit industry. We weight all regressions by total hours worked in the firm in 2009. We include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.7 Results for robot competition with alternative instruments

Here we provide additional robustness checks with respect to the instruments. For the instruments, we use data on the number of robots in each industry in South Korea and Taiwan between 2004 and 2014. Because we use the number of robots lagged by five years (see equation 7), we use extrapolated values in 2015 based on previous years. Alternatively, we lag the values by 6 years and exclude 2009. We show the results in Panel A in Table B11. We find similar outcomes as reported in Table 6. Hence, we think extrapolation of the *IFR* data with one year is not likely to change our findings.

TABLE B11 – THE EFFECTS OF ROBOT COMPETITION,
ADJUSTING THE INSTRUMENT

<i>Dependent variable:</i>	ΔGVA <i>(log)</i>	$\Delta Labor$ <i>share</i>	$\Delta Hourly$ <i>wage (log)</i>	$\Delta Hours$ <i>worked (log)</i>
PANEL A: Instrument based on $t - 6$				
Competition by robot adopters	-0.701* (0.412)	-0.020 (0.135)	-0.160 (0.216)	-0.484** (0.229)
Firm-level control variables	✓	✓	✓	✓
2-digit industry×year fixed effects	✓	✓	✓	✓
Municipality fixed effects	✓	✓	✓	✓
Number of observations	59,317	59,317	58,935	59,304
Kleibergen-Paap F -statistic	11.74	11.74	11.79	11.74
PANEL B: Instrument based on t				
Competition by robot adopters	-0.684 (0.428)	-0.149 (0.173)	-0.116 (0.191)	-0.711*** (0.256)
Firm-level control variables	✓	✓	✓	✓
2-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	70,313	70,313	69,701	70,292
Kleibergen-Paap F -statistic	11.74	11.74	11.91	11.74
PANEL C: Add Hong Kong and Singapore				
Competition by robot adopters	-0.480 (0.398)	-0.146 (0.185)	0.032 (0.218)	-0.631*** (0.192)
Firm-level control variables	✓	✓	✓	✓
2-digit industry×year fixed effects	✓	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓	✓
Number of observations	70,313	70,313	69,701	70,292
Kleibergen-Paap F -statistic	15.39	15.39	15.42	15.39

Notes: These estimations exclude the robot-adopting firms. Competition by robot adopters refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). We weight all regressions by total hours worked in the firm in 2009. We include year-specific coefficients for the log of numbers of workers in 2009 and the log of value added per worker in 2009. Robust standard errors are in parentheses and clustered at the IFR-industry×year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Conversely, in Panel B we extrapolate the number of robots until 2020 and use the robot exposure based on the current year, t , instead of $t - 5$. The results again confirm that our results are robust to this alternative instrument. The effect on hours worked is now somewhat stronger as compared to the estimate in Panel A, but still very close to the baseline estimate.

In Panel C in Table B11 we add the number of robots in Hong Kong and Singapore, which are arguably other good candidates for countries are ahead in term of technological development,

while at the same time are not important trading partners of the Netherlands. We show that the results are very similar.

Online Appendix C Robustness of worker-level results

C.1 LEED and LFS: changing the matching window

In the main analyses, we match each worker in the LEED-data to workers in *LFS* waves in the past 10 years to obtain their occupational info. This matching strongly increases the number of observations, but also has the downside that it may increase measurement error in the proxies for directly affected workers. More specifically, although occupational mobility of working age population in the Netherlands is not a major concern, some workers may still have changed their occupations, possibly as a response to robot adoption. If indeed occupational mobility is important, we would expect our estimates to be biased towards the average effect of robot adoption on workers.

To investigate this issue further, we narrow the matching window down to just 1 year where we link workers in the LEED data to workers in the *LFS* in the current or past year. This will largely address the issue of measurement error in our construct of worker types variables. Expectedly, this strongly reduces the number of observations by about two-thirds. The results for hourly wages are reported in Table C1.

Compared to the corresponding Table 8 we find qualitatively the same results. The main difference is that effect sizes are somewhat larger. For example, once including worker fixed effects, hourly wages of directly affected workers decrease by 3-6% in relative terms once a firm adopts robots. However, the standard errors become at least twice as high, therefore our initial estimates are not significantly smaller than the results reported here.

When turning to effects on employment (see Table C2) the issue of imprecise results becomes more acute. Although point estimates in Panel B on hours worked are sizable, the standard errors are too large to draw strong conclusions. Still, we re-iterate that the qualitative conclusions do not change when decreasing the window to only 1 year.

C.2 Robot adoption and the impact of low and medium-education workers

In line with the literature, *e.g.* Acemoglu and Restrepo (2020), Bonfiglioli et al. (2020) and Barth et al. (2020) who estimate negative impacts of robotization on low-skilled workers, one way to define ‘directly-affected’ workers is to focus on low-education workers. We define these workers as those whose highest degree of education is primary or secondary education. One may argue, however, that *medium-skilled* workers should be the ones who are more likely to be impacted by the broader automation effects (see Acemoglu and Autor, 2011; Oesch, 2013; Goos et al., 2014; Adermon and Gustavsson, 2015; Autor et al., 2015). While this may true in general, in the Dutch context this may not be entirely applicable given that low-education workers are more likely to be in replaceable occupations. Similarly, most bluecollar-routine workers do have a low, rather than a medium, degree. More specifically, 56% of the replaceable workers have a low degree, 41% are medium-education workers, while 3% has a high degree. These numbers are approximately the same for bluecollar-routine workers.

In any case, we test for differential effects between the low and medium-education workers in Table C3 to ascertain our choice of low-education workers as being directly affected by robot adoption.

TABLE C1 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON HOURLY WAGE:
REDUCING THE MATCHING WINDOW TO ONE YEAR

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>								
	<i>Bluecollar-routine workers</i>			<i>Replaceable workers</i>			<i>Low-education workers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter × indirectly-affected worker	0.028* (0.016)	-0.019 (0.012)		0.0266* (0.015)	-0.016 (0.011)		0.030** (0.015)	-0.014 (0.011)	
Robot adopter × directly-affected worker	-0.106*** (0.024)	-0.048*** (0.017)	-0.031** (0.015)	-0.094*** (0.024)	-0.063*** (0.018)	-0.044*** (0.016)	-0.084*** (0.022)	-0.046* (0.027)	-0.065*** (0.024)
Directly-affected worker	-0.193*** (0.004)	0.000 (0.013)	0.003 (0.015)	-0.193*** (0.004)	-0.007 (0.011)	-0.016 (0.012)	-0.198*** (0.003)	0.002 (0.009)	0.022** (0.010)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm × year fixed effects			✓			✓			✓
Number of observations	173,865	143,115	115,272	183,294	151,189	123,744	180,471	148,879	121,594
R^2	0.505	0.971	0.978	0.504	0.972	0.979	0.521	0.972	0.979

Notes: The table reports the worker-level heterogeneous hourly wage effects of firm-level robot adoption. Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry × year fixed effects and municipality × year fixed effects). Standard errors are clustered at the firm × year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE C2 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON EMPLOYMENT:
REDUCING THE MATCHING WINDOW TO ONE YEAR

Dependent variable:	Employment								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL A: Employment									
Robot adopter× indirectly-affected worker	0.003 (0.004)	-0.000 (0.004)		0.000 (0.003)	-0.001 (0.004)		0.001 (0.004)	-0.000 (0.004)	
Robot adopter× directly-affected worker	0.004 (0.007)	-0.001 (0.003)	-0.003 (0.002)	0.002 (0.007)	-0.000 (0.003)	0.001 (0.002)	-0.007 (0.006)	-0.004 (0.006)	0.001 (0.002)
Directly-affected worker	-0.002 (0.001)	0.013 (0.012)	0.023 (0.014)	-0.001 (0.001)	-0.005 (0.014)	-0.012 (0.021)	-0.003*** (0.001)	-0.009 (0.009)	-0.018 (0.012)
Control variables variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	170,470	142,007	115,139	179,786	150,028	123,069	117,015	147,731	120,919
R ²	0.180	0.570	0.583	0.176	0.572	0.584	0.177	0.571	0.583
Dependent variable:	Hours worked (log)								
	Bluecollar-routine workers			Replaceable workers			Low-education workers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PANEL B: Hours worked									
Robot adopter× indirectly-affected worker	-0.029* (0.016)	-0.010 (0.023)		-0.026 (0.016)	-0.009 (0.022)		-0.026 (0.016)	0.001 (0.020)	
Robot adopter× directly-affected worker	-0.010 (0.028)	-0.048 (0.062)	-0.086 (0.066)	-0.019 (0.027)	-0.072 (0.071)	-0.102 (0.077)	-0.038* (0.022)	-0.092 (0.058)	-0.094 (0.059)
Directly-affected worker	-0.004 (0.005)	-0.026 (0.033)	-0.018 (0.044)	-0.005 (0.005)	-0.005 (0.039)	0.020 (0.049)	0.000 (0.003)	0.013 (0.018)	-0.018 (0.023)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	173,865	143,115	115,272	183,294	151,189	123,744	180,471	148,879	121,594
R ²	0.369	0.850	0.878	0.364	0.848	0.876	0.366	0.849	0.877

Notes: The table reports the worker-level heterogeneous employment effects of firm-level robot adoption. Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the firm×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

We find results that are in line with our previous findings. In column 1 we find that, without worker fixed effects, only high-education workers benefit from the robot adoption in terms of hourly wage, while low-education workers seem to lose. When we include the worker fixed effects in column 2 the results are comparable. Column 3 includes firm-year fixed effects, which implies that we identify the effects of robots *within* firms. Relative to the high-education workers, medium and low-education workers lose from robot adoption.

Columns 4-6 are less clear cut, in line with the results reported in Table 9. Columns 5 and 6 seem to confirm that the effects of robot adoption are increasing in education, as the probability to be employed within a firm is the most reduced for low-education workers, followed by medium-education workers.

The results in columns 7-9 study employment effects at the intensive margin. Here again, we seem to find the strongest reductions in hours worked for the low-education workers, followed by the medium-education workers. Hence, all these results seem to point towards the idea that the low-education workers are indeed more affected by robot adoption.

C.3 Robot adoption and competition on personal income

From the *Personal Income* dataset we also obtained the pre-tax annual income. Personal income is a summarizing measure, which includes effects on hourly wages, hours worked and the probability to be employed. Hence, we think it is relevant to repeat our analysis, but take annual (personal) income as dependent variable. We start by studying the effects of robot adoption on annual income.

In columns 1-3 we consider directly-affected workers to be bluecollar-routine workers. We find that indirectly-affected workers gain from robot adoption, while directly-affected workers lose. In column 1, where we do not include worker fixed effects, we find that annual income increases by 2% for indirectly-affected workers when robots are adopted, while they decrease by -6.4% for directly-affected workers. Part of this effect is due to worker sorting, because when we control for worker fixed effects in column 2 we find smaller effects, but, still, we find that directly-affected workers experience negative income effects of robot adoption, while indirectly-affected workers gain. These results are confirmed, and even slightly more convincing, when we focus on the two other definitions of directly-affected workers: replaceability and whether a worker has a low education. Columns 3, 6 and 9, where we include firm-year fixed effects, confirm that directly-affected workers are much worse off in relative terms compared to indirectly-affected workers when robot competition intensifies.

We further investigate the effects of robot competition in Table C5. In line with Tables 11 and 12, we do not find effects of robot competition when we do not include worker fixed effects. Still, the signs are in line with expectations with indirectly-affected workers benefiting, while directly-affected workers are losing (see columns 1, 4 and 7). When we include worker fixed effects, the effects become more pronounced. For instance, in column 2, a standard deviation increase in competition increases incomes of indirectly-affected workers by $\exp(0.0993 \times 0.459) \times 100\% = 4.4\%$, which is a sizable effect.

In sum, the results reported in Tables 11 and 12 show that increases in hourly wages apparently dominates the reductions in employment.

TABLE C3 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION – LOW AND MEDIUM-EDUCATION WORKERS

<i>Dependent variable:</i>	<i>Hourly wage (log)</i>			<i>Employed</i>			<i>Hours worked (log)</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter × low-education worker	-0.017* (0.010)	-0.013* (0.007)	-0.061*** (0.010)	-0.006 (0.007)	-0.009 (0.006)	-0.021*** (0.007)	-0.036*** (0.011)	-0.032** (0.013)	-0.027 (0.018)
Robot adopter × medium-education worker	0.007 (0.008)	-0.004 (0.005)	-0.047*** (0.007)	-0.002 (0.004)	0.004 (0.003)	-0.010** (0.004)	-0.025*** (0.008)	-0.011 (0.009)	-0.007 (0.012)
Robot adopter × high-education worker	0.046*** (0.010)	0.043*** (0.006)		0.007 (0.005)	0.010*** (0.004)		-0.013 (0.008)	-0.006 (0.010)	
Low-education worker	-0.435*** (0.003)	-0.038*** (0.010)	-0.042*** (0.009)	-0.008*** (0.001)	-0.005 (0.006)	-0.005 (0.006)	-0.001 (0.003)	-0.012 (0.04)	-0.011 (0.014)
Medium-education worker	-0.327*** (0.003)	-0.024*** (0.009)	-0.031*** (0.008)	-0.001* (0.001)	-0.004 (0.004)	-0.004 (0.004)	-0.005** (0.002)	-0.022* (0.012)	-0.017 (0.012)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm × year fixed effects			✓			✓			✓
Number of observations	793,550	753,123	724,914	792,772	757,842	730,012	793,550	753,123	724,914
R^2	0.569	0.938	0.948	0.160	0.596	0.628	0.293	0.729	0.764

Notes: Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry × year fixed effects and municipality × year fixed effects). Standard errors are clustered at the firm × year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE C4 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON INCOME – HETEROGENEITY

<i>Dependent variable:</i>	<i>Annual income (log)</i>								
	<i>Bluecollar-routine workers</i>			<i>Replaceable workers</i>			<i>Low-education workers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter× indirectly-affected worker	0.020** (0.008)	0.012** (0.005)		0.021*** (0.008)	0.012*** (0.005)		0.037*** (0.008)	0.018*** (0.005)	
Robot adopter× directly-affected worker	-0.065*** (0.019)	-0.014 (0.009)	-0.022*** (0.009)	-0.070*** (0.016)	-0.013 (0.009)	-0.022** (0.009)	-0.087*** (0.014)	-0.020** (0.008)	-0.040*** (0.009)
Directly-affected worker	-0.173*** (0.003)	0.005 (0.008)	0.002 (0.009)	-0.176*** (0.003)	-0.000 (0.008)	-0.008 (0.009)	-0.191*** (0.002)	-0.009 (0.007)	-0.010 (0.007)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm×year fixed effects			✓			✓			✓
Number of observations	829,365	794,477	767,359	866,208	829,811	802,991	852,534	817,741	790,731
R^2	0.481	0.888	0.895	0.481	0.888	0.895	0.495	0.889	0.895

Notes: Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the firm×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE C5 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION ON INCOME – HETEROGENEITY

<i>Dependent variable:</i>	<i>Annual personal income (log)</i>								
	<i>Bluecollar-routine workers</i>			<i>Replaceable workers</i>			<i>Low-education workers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Competition by robot adopters× indirectly-affected worker	0.122 (0.183)	0.459** (0.207)		0.193 (0.182)	0.409** (0.191)		0.260 (0.187)	0.468** (0.208)	
Competition by robot adopters× directly-affected worker	-0.129 (0.202)	0.029 (0.200)	-0.410*** (0.114)	-0.230 (0.215)	-0.072 (0.211)	-0.506*** (0.123)	-0.285 (0.198)	0.141 (0.205)	-0.211** (0.102)
Directly-affected worker	-0.162*** (0.005)	0.023** (0.011)	0.018 (0.012)	-0.159*** (0.006)	0.018 (0.011)	0.013 (0.012)	-0.174*** (0.007)	-0.000 (0.008)	-0.006 (0.008)
Control variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	716,330	682,294	658,054	744,424	708,965	684,990	732,686	698,741	674,628
Kleibergen-Paap <i>F</i> -statistic	6.923	3.002	18.62	6.865	2.895	18.34	6.867	2.825	21.24

Notes: These estimations exclude the robot-adopting firms. Competition by robot adoption refers to the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (7). Control variables include worker-level variables (*i.e.* age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults) and firm-level variables (*i.e.* 4-digit industry×year fixed effects and municipality×year fixed effects). Standard errors are clustered at the IFR-industry×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

TABLE C6 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION:
FIRST-STAGE RESULTS

<i>Dependent variable:</i>	<i>Competition by robot adopters</i>		
	(1)	(2)	(3)
Robot exposure	0.213*** (0.051)	0.212*** (0.052)	0.184*** (0.051)
Worker-level variables		✓	✓
Worker fixed effects			✓
Firm fixed effects	✓	✓	✓
4-digit industry×year fixed effects	✓	✓	✓
Municipality×year fixed effects	✓	✓	✓
Number of observations	1,694,145	1,559,308	1,516,999
R^2	0.872	0.871	0.894

Notes: These estimations exclude the robot-adopting firms. Competition by robot adoption refers to the share of sales by robot adopting firms within the same 4-digit industry. Worker-level variables include age dummies, as well as indicators for whether the worker is male, has a migrant background, and whether the worker is part of a couple or a household with multiple adults. Standard errors are clustered at the firm×year and worker levels and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.4 First-stage results for robot competition

In Table C6 we report first-stage results of the robot competition regressions for the worker-level analysis. Hence, we regress robot competition on robot exposure, like we did in Appendix B.5, yet the unit of analysis is now the worker.

In column 1 we only include firm, industry-year and municipality-year fixed effects. We find a somewhat stronger coefficient of robot exposure than at the firm level. One standard deviation increase in robot exposure increases competition by robot adopters by $0.0993 \times 0.213 \times 100 = 2.12$ percentage points. The coefficient is essentially the same once we include worker control variables in column 2 and is slightly lower when including worker fixed effects in column 3.