STONE CENTER | INEQUALITY & SHAPING THE FUTURE OF WORK

Robots and Workers

Daron Acemoglu
Hans R.A. Koster
Ceren Ozgen

December 2025

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the James M. and Cathleen D. Stone Center on Inequality and Shaping the Future of Work, the Massachusetts Institute of Technology, or any affiliated organizations. Stone Center working papers are circulated to stimulate discussion and invite feedback. They have not been peer-reviewed or subject to formal review processes that accompany official Stone Center publications.

Stone Center on Inequality and Shaping the Future of Work
70 Memorial Drive, Building E51-355 - Cambridge, MA 02142

shapingwork.mit.edu

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Robots and Workers*

Daron Acemoglu† Hans R.A. Koster‡ Ceren Ozgen§

December 24, 2025

Abstract—We estimate the effects of robot adoption on workers in the Netherlands using a large employer-employee panel dataset spanning 2009-2020. Our results show that *directly-affected* workers (*e.g.*, bluecollar workers performing routine or replaceable tasks) face lower earnings and employment rates, while *indirectly-affected workers* gain from robot adoption. About half of these effects are due to worker heterogeneity: robot-adopting firms hire workers with greater earnings potential, even before they adopt robots, and do so particularly in occupations that are not directly impacted by robots. We also find that robot adoption reduces employment but increases wages in competing firms.

Keywords—robots, workers, technology, productivity, the Netherlands

JEL codes—D63, E22, E23, E24, J24, O33

*We thank David Autor, Matt Cole, Francesco Esposito, Maarten Goos, Ted Pinchbeck, James Rockey, Anna Salamons, Benedikt Vogt, Ulrich Zierahn, the editor and three anonymous referees, and participants at seminars in Amsterdam, Essex, Gyumri (Armenia), London, Prague, The Hague, Sussex and Utrecht for valuable feedback. Daron Acemoglu acknowledges funding from the Hewlett Foundation and the Sloan Foundation. Ceren Ozgen gratefully acknowledges Birmingham Business School research funding and University of Birmingham Covid support fund. Earlier version of this paper also appeared as NBER Working Paper 31009.

†Department of Economics, MIT, email: daron@mit.edu.

‡Department of Spatial Economics, Vrije Universiteit Amsterdam, email: h.koster@vu.nl; Tinbergen Institute and CEPR.

§Corresponding author. Department of Economics, University of Birmingham, University House 1105, B15 2TT Birmingham, UK; IZA and GLO. e-mail: c.ozgen@bham.ac.uk.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

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 & Restrepo 2022*a*). In the US, which was initially slow in robotics investments, the number of robots per 10,000 industrial workers increased from 35 in 1993 to 149 in 2014, and then accelerated to 255 in 2020. The Dutch labor market has seen a similar increase from 12 robots per 10,000 industrial workers in 1993 to 68 in 2014 and to 209 in 2020. Although industrial robots have automated a variety of production tasks from painting to welding, sorting and assembly, their effects on workers are still debated.

Most firm-level studies find that robot-adopting firms experience faster productivity growth and also expand 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. First, robot-adopting firms are typically more productive and often on a different trend than non-adopters (*e.g.* Koch et al. 2021, Acemoglu et al. 2022). Second, adopters may be expanding at the expense of their rivals in the same industry, as documented in Acemoglu et al. (2020) and Koch et al. (2021), as well as in Bessen et al. (2020) for the Netherlands and Bonfiglioli et al. (2020) for France. Because of this equilibrium impact of robots, overall industry or nationwide employment could decline if non-adopting competitors significantly reduce employment. Indeed, studies focusing on industry-level implications of robots typically find negative effects on employment and wages. For example, Acemoglu & Restrepo (2020) estimate negative impacts in more exposed US local labor markets—especially for low- and mid-skill workers and those in manufacturing and in bluecollar occupations. 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 & Restrepo (2022*b*) estimate negative effects on wages and employment on demographic groups most exposed to automation, driven by robots and specialized software. Graetz & Michaels

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

(2018), Acemoglu & Restrepo (2020) and Acemoglu et al. (2022) also report negative effects on the labor share at the industry level.

Much less is known about impacts on workers. Because robots reduce employment in some firms and tasks and expand it in others, they may systematically benefit some employees while harming the labor market prospects of other groups of workers. Consistent with this expectation, studies using aggregate data, such as Acemoglu & Restrepo (2020), Dauth et al. (2021) and Humlum (2019), find negative effects on production workers, and Bonfiglioli et al. (2020) and Barth et al. (2020) estimate negative impacts on lower-education workers. In contrast, Aghion et al. (2021) estimate positive employment effects, even for unskilled production workers in France, using proxies of equipment investment (which involves some automated and non-automated capital), and Hirvonen et al. (2022) find that investment in advanced equipment has not had a negative effect on low-skilled workers in Finland.

Our main contribution in this paper is to estimate the impacts of robots, as well as robot adoption by competitors, on individual worker outcomes, especially distinguishing workers employed in tasks that can be automated by robots and those performing other tasks.[1]

We first confirm several of the important firm-level and industry-level findings of the literature using our Dutch firm panel dataset on robots and firms.[2] Specifically, we find that firms adopting robots experience a roughly 15% boost in output, a 4% increase in employment, and a decrease in the labor share of about 5 percentage points compared to similar non-adopting firms. The quantitative magnitudes of these estimates are very similar to previous studies from France and Spain.

Turning to our main contribution, we investigate the impact of robot adoption on individual workers, utilizing a large matched employer-employee dataset. Our starting point is the theoretical prediction that the effects of robots should be different on *directly-affected* workers employed in

[1] Table A1 in Appendix A.1 presents an overview of the existing literature and places our paper within it.

[2] Our definition of robots corresponds to code 8479500 in the international trade classification of commodities, which is for *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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

tasks that can be automated with this new technology from the impact on *indirectly-affected* workers, typically performing complementary tasks. While robot adoption should generate a negative *displacement effect* on directly-affected workers (Acemoglu & Restrepo 2020), it simultaneously produces a positive *productivity effect*, as non-automated tasks expand, and we would expect that indirectly-affected workers be the main beneficiaries of this expansion (for example because these are the workers that have comparative advantage in non-automated tasks, see Acemoglu & Restrepo 2022*b*).

To explore these issues systematically, we construct three alternative, though complementary, measures of directly-affected workers. The first one identifies blue-collar workers engaged in routine tasks (constructed using the routine task intensity index from Autor & Dorn 2013 and Koster & 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 & 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 & Restrepo 2020, Bonfiglioli et al. 2020, Barth et al. 2020). The second measure is based on the replaceability index of Graetz & 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.

Using all three measures, we find negative impacts on directly-affected workers, and positive impacts for indirectly-affected workers. For example, our baseline estimates indicate that real wages for directly-affected workers decline by about 5.4%, while those of indirectly-affected workers increase by about 3.4% following robot adoption. We also estimate a lower likelihood of continued employment and fewer hours worked for directly-affected workers, but generally positive effects for indirectly-affected workers.

Our baseline estimates could be influenced by the sorting of heterogeneous workers into different firms after robot adoption or by initial selection of workers (because firms that adopt robots in the future may already employ workers with different earning potentials). Including worker fixed

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

effects in the wage regressions reduces the impact of robots on both types of workers by about a half: 1.5% for indirectly-affected workers and −1% for directly-affected workers. We further explore the role of sorting and initial selection, and find that robot adoption does not significantly alter the composition of the workforce, but firms adopting robots were already employing workers with higher earning potential in non-automated tasks prior to the introduction of robots.

We also explore the effects of robot adoption by competing firms. While Acemoglu et al. (2020) and Koch et al. (2021) estimate negative impacts on firms whose competitors are adopting robots, we are able to look at the effects on workers, and find that there are negative impacts on directly-affected workers, but positive impacts on indirectly-affected workers' wages, presumably because the industry-level demand for these workers increases following robot adoption.

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. 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 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 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 automated and non-automated technologies as well as robots. Most importantly, our study uses worker-level outcomes, such as hourly wages and hours worked, and detailed measures of task content in workers' occupations. Differently from other studies in the literature, our data further enable us to separately estimate the impacts of robot adoption on directly-affected and indirectly-affected workers.

The rest of the paper is organized as follows. Section 2 introduces our various data sources and

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

describes how we construct the measures of directly-affected and indirectly-affected workers. Section 3 presents our main results on the effects of robot adoption on worker-level outcomes. Section 4 concludes, while Appendices A and B provide information on our data and additional results at the firm-level and Appendix C presents further robustness checks.

2 Data

Our study leverages a number of administrative datasets provided by *Statistics Netherlands*. We combine the following datasets: *Production Statistics*, *Tax Registers*, the *International Trade Register*, *Labor Force Surveys*, *Investment Statistics*, and the *Firm Register*. We now describe various aspects of our data.

2.1 Linked employer-employee data and robot adoption

Firm-level *Production Statistics* constitute the starting point of our sample. The data include detailed information on firms' production inputs/outputs, such as the number of employees, value-added, sales, total costs, personnel costs, and total wage bill. The dataset covers all firms with 50 employees and above, and is a representative sample of firms smaller than 50 employees per year for the 2000-2020 period. We observe around $55,000$ unique firms per year. We focus on firms in the industrial sector, covering manufacturing, energy, water and waste, construction, mining, and transportation sectors.

We link *Production Statistics* to *Tax Registers*, which are based on the employers' tax declarations. This dataset includes all individuals employed by formally registered firms, but excludes the self-employed. In total, we observe the monthly wages and hours worked of around 10 million employees per year. By linking *Production Statistics* to *Tax Registers*, we are able to construct our matched employer-employee dataset (LEED), which covers the near-universe of firms in the manufacturing industry and their employees over time.[3]

From *Tax Registers*, we use before-tax annual earnings and hours worked. To ensure comparability,

[3] We calculate the labor share as the total wage costs over value-added. We set the labor share to missing if this quantity is larger than one, which applies to about 2% of the firms. While firms with missing values are slightly less productive, are older, and export less, they have similar robot adoption behavior.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

we drop firms from the LEED where at least one of the following variables is missing: value-added, labor share, sales, or total hours worked. Firms with more than 25,000 workers are also dropped because they are very likely to refer to employment agencies. In total, our data covers about 70% of workers in the Netherlands (the remaining 30% are missing mostly due to the fact that we only have a subsample of firms with fewer than 50 employees).

Additionally, the *International Trade Register* (*ITR*) includes all trade transactions in the Netherlands with other countries at the firm level from 2009 onwards. In the case of trade with non-EU countries, the information is gathered from customs data. For trade within the EU, *Statistics Netherlands* runs its 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 trade, as well as the identity of the trading partner. In total, the *ITR* and the *Intrastat* survey contain about 80% of total Dutch imports and exports (in value) that can be attributed to a firm. From these data, we obtain robot imports (corresponding to commodity code 847950). 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 about € 2,500.[4]

In Appendix B.3, we present evidence that, although robot imports are positively correlated with computer and machinery investments, they are distinct, and our results remain largely unchanged when we control for investments in computers and other machinery. Imports below this value are unlikely to refer to significant robotics investments and may also contain measurement error.[5] Note also that the threshold of € 1.2 million for within-EU transactions refers to total imports of a firm from the same country. Hence, we only miss robot imports of a firm if it imports robots from some country, say Germany, but also engages in less than € 1.2 million of imports from that country that year. We also verify the robustness of our results by excluding firms that re-export

[4]Total costs exceed € 2,500, because they typically include expenses for installation, programming, and integration. In practice, they can be as high as $50,000-100,000 (see *e.g.* Robots.com 2013). Because these costs are not included in the robot import values, they cannot be compared to estimates of total automation costs.

[5]We checked robustness to different cut-off values based on the robot import value distribution and obtained very similar results. Appendix B.3 also presents results with a continuous measure.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

robots. Overall, our data provides a fairly reliable source of information on robot adoption and is not subject to the type of measurement error and recall bias involved in self-reported measures of robotics investments.

In summary, robot adoption in our dataset is defined as follows:

$$r_{ft} = I\left(\sum_{\tau=2009}^{t} \mathcal{RI}_{f\tau} \geq \overline{\mathcal{RI}}\right), \tag{1}$$

where $\mathcal{RI}_{f\tau}$ are robot imports of firm f in year $\tau \leq t$. The robot adoption dummy equals one (and stays at one) when the cumulative robot imports since 2009 exceed the minimum threshold $\overline{\mathcal{RI}}$.

We link *Production Statistics* to investments data and to the firm register to obtain firm age. This dataset further allows us to observe firms' investments by type (*e.g.*, computers, hardware, machinery, etc.) from 2003 onwards. Finally, *Firm Register*, which is a register of the universe of firms in the Netherlands, helps us identify the location and age of each firm. By combining these seven datasets, we create a thorough picture of robot-adopting firms between 2009 and 2020.

We also estimate the effects of competitors' robot adoption. Following Acemoglu et al. (2020), the competition variable is defined based on the share of sales in a given 4-digit industry accounted for by robot-adopters (leaving out the sales of the firm in question). Specifically, we define *robot adoption by competitors* as

$$r^{\mathcal{C}}_{ft} = \frac{\left(\sum_{\tilde{f}\in s} q_{\tilde{f}t} r_{\tilde{f}t}\right) - q_{ft} r_{ft}}{\left(\sum_{\tilde{f}\in s} q_{\tilde{f}t}\right) - q_{ft}}, \tag{2}$$

where $r_{\tilde{f}t}$ captures robot adoption in another firm, $\tilde{f}$, and $q_{\tilde{f}t}$ denotes firm sales.[6]

2.2 Worker-level data

We link the LEED to the *Demographic* register, which contains the universe of the Dutch population, and obtain information on workers' age, gender, and whether they were born in the

[6] Because we do not have detailed information on the product composition of different firms, we are using coarser information to construct the competition variable than in Acemoglu et al. (2020).

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Netherlands. The resulting worker-level dataset contains the near-universe of employees in all sectors, though in our analysis, we focus on the industrial sector where robot adoption is most prevalent. We drop individuals younger than 18 or older than 67, focusing on the working-age population. We also drop observations with outlying values—specifically those with more than 4,380 hours worked per year; less than € 2.5 and more than € 500 hourly wages per annum; and more than € 500,000 income a year. Our final data is a balanced panel of 333,000 unique workers who have been employed in the industrial 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 by a certain firm. In our analysis of how robot adoption affects the employment status of workers, we assign the characteristics of the last firm that employed the worker.

As noted above, *Tax Registers* do not include workers who are self-employed. To obtain information on whether a worker is unemployed or self-employed, we combine our data with the *Personal Income* data, which enables us to distinguish different sources of income and determine employment status for a larger sample. From these data, we can also exclude individuals who do not participate in the labor force due to study or retirement.

2.3 Defining directly-affected workers

As discussed in the Introduction, the displacement and productivity effects suggest that the impact of automation should be uneven across workers. 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 employees can only be observed from the *Labor Force Surveys* (*LFS*), we link the LEED with observations from 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 dataset, 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 the task

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

content of a job in a worker's occupation. We label workers as *directly-affected*, denoted by $a_{it} = 1$, when they belong to one of these groups. The remaining workers are referred to as *indirectly-affected* ($a_{it} = 0$), since the impact of robot adoption on them will be mostly through indirect, equilibrium channels, such as productivity increases, reorganization, or reallocation to new tasks.

Bluecollar routine workers. Our first measure of directly-affected workers is those employed in bluecollar occupations performing routine tasks. These workers are likely most impacted by robot adoption, since current robotics technology focuses on automating routine manual tasks. Using *O*NET Online* and occupational codes from the LFS, we compute a routine task intensity index (RTI), following Autor & Dorn (2013) and Koster & Ozgen (2021) in the Dutch context. The exact definition of RTI is in Appendix A.2.

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

$$a_{it} = \max_{\tau=-9,\dots,t} \left(I(\mathcal{B}_{io\tau} = 1) \times I(\ RTI_{io\tau} > 1) \right), \tag{3}$$

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

By adopting a 10-year window, we are assuming 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 among 18- to 25-year-old workers transitioning from education to employment. Although we believe that potential occupational mobility is not a major issue, in Appendix C.4, we confirm the robustness of our worker-level results by narrowing this window to one year.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

According to our definition, about 11% of the workers in the Dutch industrial sector are bluecollar-routine workers during the study period.

Replaceable workers. Not all bluecollar routine workers are equally susceptible to robot adoption. Although some occupations require the performance of certain routine tasks, they still need to be complemented by non-routine tasks that may require assessment and discretion (consider, for example, call center agents, metalworking machinists, wood-cutting operators, and metal drillers).

To account for these differences, our second measure leverages a worker-level replaceability index, constructed at the 4-digit ISCO level, following Graetz & Michaels (2018). This replaceability index is based on the description of robot applications by the *International Federation of Robots* (*IFR*) and occupational classifications in the US Censuses. The *IFR* distinguishes the tasks that can be executed by robots, such as welding, assembling, and painting. If an occupational title includes one of these keywords, we assign a value of 1 to that occupation to indicate that workers in that occupation are replaceable by robots, as in Graetz & 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, we look at a worker's occupation within a 10-year window. Hence,

$$a_{it} = \max_{\tau=-9,\ldots,t} \left(I(\mathcal{V}_{io\tau} = 1)\right), \tag{4}$$

where $\mathcal{V}_{io\tau}$ is an indicator variable for 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 the low-education group when their highest degree corresponds to secondary schooling. Hence, these workers would

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

have in total a maximum of 10 years of primary and secondary education. Our measure is then:

$$a_{it} = \max_{\tau=-9,\ldots,t} \left(I(\mathcal{E}_{io\tau} = 1) \right), \tag{5}$$

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

2.4 Descriptive statistics for the firm-level data

Our firm-level panel data spans 12 years and includes 162,220 firm-year observations and 46,914 unique firms. We observe 218 unique robot-adopting firms (0.5%). Although only a small fraction of firms are robot-adopters, they tend to be larger, and thus 6.8% of the workers in our sample are employed in a firm that adopts robots at some point during our time window.[7] Appendix A.3 reports detailed descriptive statistics.

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 a firm's import value of robots between t and $t-1$ is 0.76.

Figure 1 displays the cumulative value of robot imports versus the number of firms adopting robots over the period 2009-2020. The cumulative trend in both indicators shows a steady increase over time.

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 it is mostly larger firms that are adopting robots.

[7] Based on IFR data, the number of robots in Dutch manufacturing in 2020 was approximately 14,000. This implies that there are on average 65 installed robots per robot-adopting firm in our sample, which is reasonable, since adopters tend to be larger firms.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

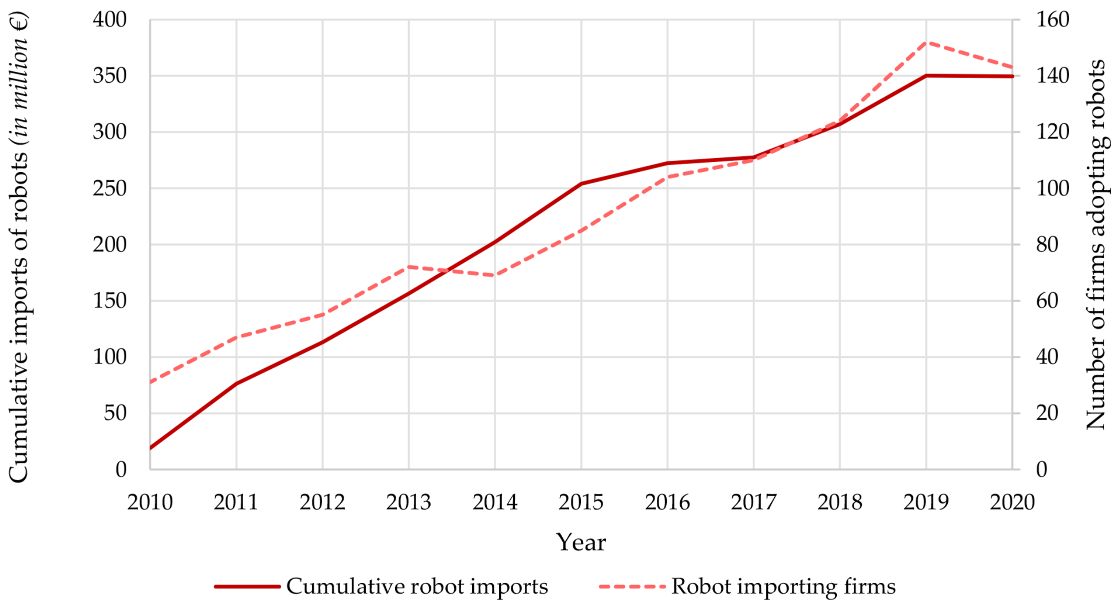

Figure 1 – Cumulative robot adoption

Notes: We present the cumulative value of robot imports, in million *Euros*, over the study period (2009–2020), alongside the number of robot-adopting firms. The data come from the Dutch International Trade Register and the robot import value refers to products imported under the specific commodity code 847950.

2.5 Descriptive statistics of the worker-level data

Descriptive statistics for the matched worker data are reported in Table 1. 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, respectively, €32 and €65,841, which are 30% higher than those in non-adopters. The employee characteristics are generally similar between robot-adopting and non-adopting firms, except for a lower share of low-education workers (about 50%); a lower share of replaceable workers (about 25%); and a lower share of immigrant workers (about 15%) among robot-adopters.

About 10% of the workers in the full data are classified as bluecollar routine workers (which, recall, is defined as those who are in an occupation with an RTI value exceeding 1 and in a bluecollar occupation). The share of this type of worker 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 employees are replaceable in non-adopting firms, while this value is 7.8% in adopters. Similarly, low-education workers represent 34.5% of the workforce, but this share is only 18.8% in robot-adopters. The summary statistics are consistent with the idea that robot-adopting firms

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

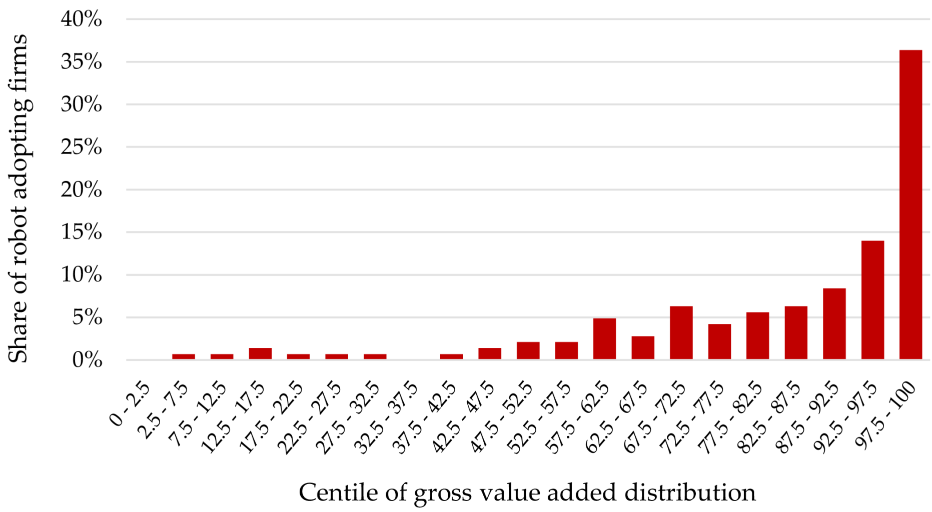

Figure 2 – Robot adoption by value-added percentiles

Notes: This figure shows the share of robot-adopting firms across the centiles of value-added at the firm level during the study period (2009–2020). Value-added for each firm is from the Dutch Production Statistics.

have more skilled workforces. This skill differential also partly explains why hourly wages are higher among robot-adopters. Finally, more than 80% of industrial workers are male, regardless of whether a firm adopts robots or not.

During the study period, $1 - 0.953 = 4.7\%$ of the workers became unemployed the following year. This rate is 3.1%, about 25% lower for workers who were previously employed in robot-adopting firms.

We 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. In Panel A, we depict the share of workers by worker type over 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.

Figure 3b shows that there is an overall increase of about 40% in total hours worked since 2001, while hours worked by workers performing routine tasks in bluecollar occupations decreased by 15%. The same pattern can be seen for replaceable and low-education workers.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Table 1 – Summary statistics of matched LFS sample of workers 2009-2020

	mean	*std. dev.*	*5^{th} perc.*	*median*	*95^{th} perc.*	*N*
Panel A: Workers in robot-adopting firms	(1)	(2)	(3)	(4)	(5)	(6)
Mean hourly wage *(in €)*	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 *(in €)*	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 *(in €)*	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 *(in €)*	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 the manufacturing sector. Panel B reports summary statistics for workers in non-adopters in the 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

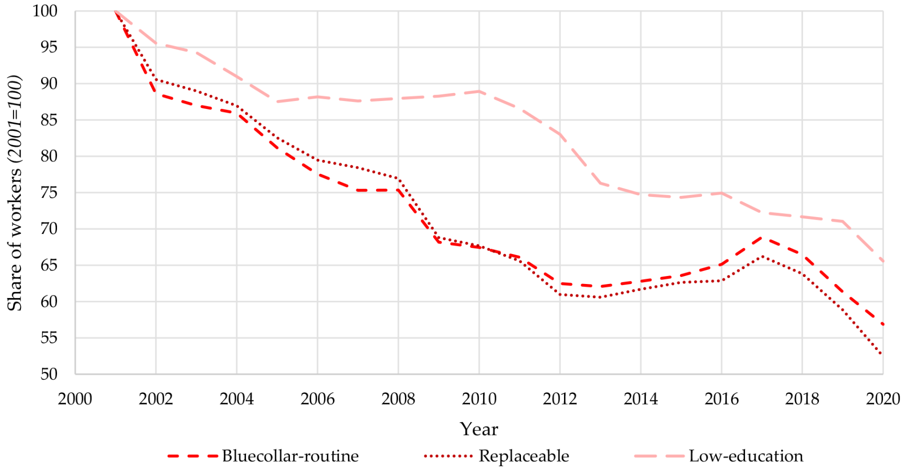

(A) Share of workers by worker type

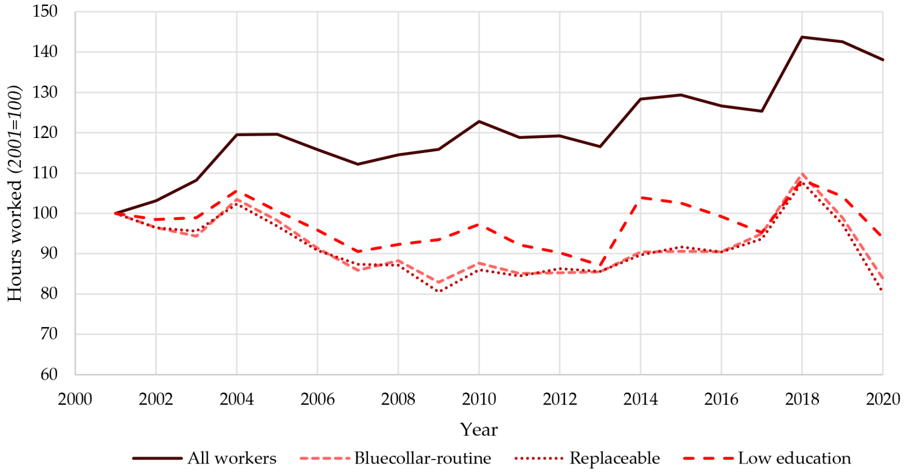

(B) Total hours worked by worker type

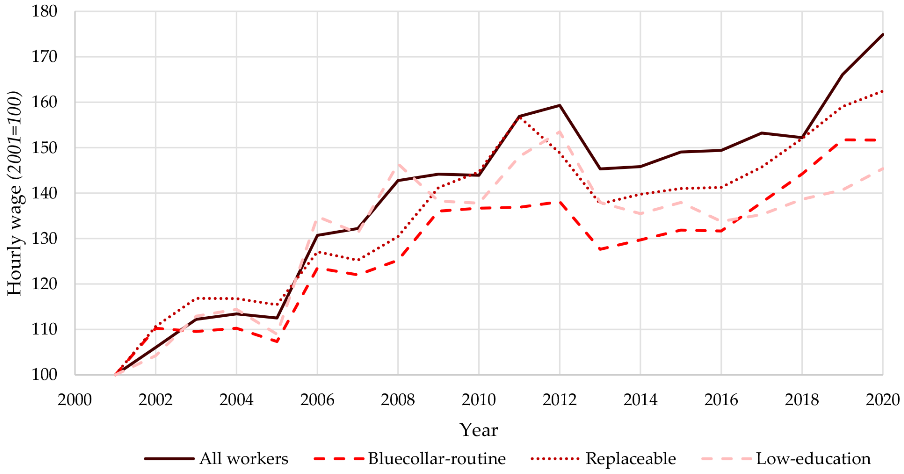

(C) Mean hourly wage by worker type

Figure 3 – Trends in hours worked and hourly wage by worker type

Notes: We present trends in: (A) the share of worker types, (B) hours worked by worker type, and (C) hourly wages by worker type. All values in the corresponding figures are normalized relative to their respective 2001 values, so the changes represent percentage changes for each indicator relative to 2001. The data are sourced from the Labor Force Surveys for 2001–2020.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

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, however, that is slower growth for bluecollar routine and replaceable workers than for 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 Netherlands.

3 Main results

Before turning to our main focus, which is the worker-level analysis, we first briefly summarize the impact of robot adoption on firms in the next subsection. We then report baseline effects of robot adoption on workers in Section 3.2, and our most important results, distinguishing the impacts of robots on directly-affected and indirectly-affected workers, are presented in Section 3.3. Section 3.4 provides a range of robustness checks. Lastly, Section 3.5 studies the effects of competitors' robot adoption on wages.

3.1 Firm-level effects of robots

We start with firm-level effects of robots. These are useful both for comparison with the previous literature and to show how overall employment at the firm level and the labor share are impacted by robots (which cannot be seen from the worker-level analysis). Since they are not our main focus, these results are presented in Appendix B, and here we provide a brief summary. We look at both long-differences and panel data estimates for value-added, labor share, hourly wage, and hours worked. The long-differences regressions focus on changes over a 12-year period, while the panel data estimates use annual data. All our firm-level specifications are weighted by total hours worked in the firm in 2009 and allow for differential trends by baseline characteristics (in particular, log number of workers in 2009 and log value-added per worker in 2009, interacted with a full set of time dummies) and also control for industry fixed effects and location fixed effects.

Our results point to significant associations between robot adoption and firm value-added, labor

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

share, average hours worked, and average hourly wage. Specifically, 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 qualitative and quantitative impacts are very similar to previous studies using firm-level data (see particularly Acemoglu et al. 2020, Koch et al. 2021, for evidence on France and Spain, respectively).

We subject these results to several robustness checks. These include complementary event-study estimates, placebo treatments, various corrections for omitted variable bias, checks against possible violation of the stable unit treatment variance assumption (SUTVA), verifying the sensitivity of our estimates to the exclusion of firms re-exporting robots, and finally, estimates using various alternative measures of robot adoption. The details of these robustness checks, which confirm the results summarized here, are reported in Appendices B.3, B.4, and B.5.

Additionally, Appendix B.6 shows that wage inequality within firms rises significantly following robot adoption, which is consistent with our main results below, indicating opposite-signed impacts on directly- and indirectly-affected workers.

In Appendix B.7, we further study the effects of robot adoption on competitors. We focus on firms that do not adopt robots and look at the share of sales by robot-adopters within 4-digit industries in which these firms operate (leaving out the sales of the firm in question). We also instrument the competition variable (see equation (2) above), as in Acemoglu et al. (2020). The results indicate that robot adoption has a negative impact on competitors: a one standard deviation increase in competitors' robot adoption reduces value-added by 4.4%, the labor share by 1.4 percentage points, and hours worked by 6%. These estimates are also robust.

3.2 The baseline effects of robots on workers

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:

$$\{\log w_{ift}, \log h_{ift}\} = \beta r_{ft} + \zeta z_{it} + \kappa_f + \lambda_{s_f,t} + \mu_{m_f,t} + \nu_i + \epsilon_{ift}, \tag{6}$$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

where z_{it} are worker characteristics such as age and immigration background. We 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. Equation (6) also includes 4-digit industry-by-year ($\lambda_{s_f,t}$) and municipality-by-year ($\mu_{m_f,t}$) fixed effects to control for any industry or location-specific trends in robot adoption and labor market outcomes (where s_f and m_f, respectively, denote the sector and the municipality to which firm f belongs).

Finally, we can follow workers over time and include worker fixed effects, ν_i. Specifications that include worker fixed effects focus on the impact of robot adoption on the same worker and are particularly useful to control for sorting of workers with different earning potentials in robot-adopting firms, as well as the initial selection of workers with different earning potentials across firms that will and will not adopt robots in the future. In all specifications, standard errors are clustered at the firm-year and worker levels.

We also performed event studies to explore whether there are no pre-trends in hourly wages prior to the adoption of industrial robots by firms (see for a similar application Bessen et al. 2020):

$$\log w_{ift} = \sum_{\tau=-4}^{1} \beta_\tau r_{fmt,\tau} + \zeta z_{it} + \kappa_f + \lambda_{s_f,t} + \mu_{m_f,t} + \nu_i + \epsilon_{ift}, \tag{7}$$

where β_τ indicates the treatment effect in year τ, with the robot adoption year designated as $\tau = 0$. The effect for $\tau = -2$ is normalized to zero, and the specifications always include firm and industry-year fixed effects, as in our baseline estimates. Because we only have 12 years, our data do not permit us to estimate long-term impacts. We therefore focus on effects on hourly wages from 4 years prior to robot adoption to 1 year after the adoption.

Our main focus is not just the average impact of robot adoption on workers, but also the heterogeneous effects across worker types. 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 the expansion of non-automated

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

tasks). We will use the three measures of directly-affected workers (based on workers performing bluecollar routine tasks, performing replaceable tasks, and having low education), as defined in equations (3), (4), and (5). The econometric specification in this case can be written as

$$\{\log w_{ift}, \log h_{ift}\} = \beta_1 r_{ft} a_{it} + \beta_2 r_{ft}(1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \kappa_f + \lambda_{s_f,t} + \mu_{m_f,t} + \nu_i + \epsilon_{ift}, \quad (8)$$

where a_{it} is an indicator for whether the worker is directly-affected. We also control for the direct effects of a_{it}, firm fixed effects, industry-year fixed effects, and worker fixed effects. To further address concerns related to endogeneity, we estimate a version of equation (8) 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 impact employment. 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, specified as follows:

$$e_{ift} = \beta_1 r_{ft} a_{it} + \beta_2 r_{ft}(1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \kappa_f + \lambda_{s_f,t} + \mu_{m_f,t} + \nu_i + \epsilon_{ift}. \quad (9)$$

Recall that the specification excludes the retired, the self-employed, and students. In our analysis, we follow the same workers from 2009 to 2020. When constructing the employment variable, we retain workers in the dataset for one year after they become unemployed. This approach allows us to include both currently and previously employed workers in our analysis. To account for firm effects, we assign the characteristics of the last employer to workers who are currently unemployed but were previously employed. We reiterate that regressions using this employment measure will not account for effects on new workers joining a firm.

In Table 2, we begin by reporting the average effects of robot adoption on hourly wages, employment status, and hours worked. However, our primary estimations focus on workers matched with the *LFS* surveys, which provide occupational information on directly- and indirectly-affected workers. For consistency with the subsequent tables, we restrict our analysis

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

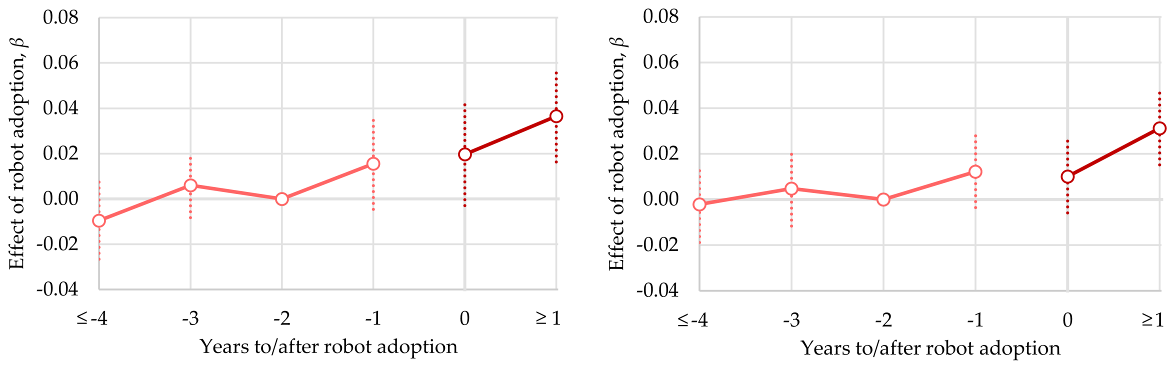

(A) Baseline with controls (B) With worker fixed effects

Figure 4 – Event studies for workers' hourly wages

Notes: This figure presents event-study estimates from equation (7) using data from 2010-2016. The dotted lines denote 95% confidence bands based on robust standard errors.

in Table 2 to this same set of workers, although we do not yet distinguish between worker types in these estimations.

In columns 1-3 of Table 2, we estimate that robot adoption is associated with an average increase in hourly wages of 2.5%. This estimate remains essentially the same regardless of whether a battery of worker characteristics (age, gender, migrant background, household type) are included, as we do in column 2. When we additionally include worker fixed effects in column 3, the impact is slightly smaller (hourly wages increase by 1.6%), which may suggest the possibility that workers with different earning potentials are differentially sorted into robot-adopting firms, either before or after adoption. We discuss this issue in greater detail below.

Figure 4 depicts the results for event studies for hourly wages. In the left panel, we depict estimates from a specification that corresponds to column 1 in Table 2, where we control for firm fixed effects, industry-year fixed effects, and municipality-year fixed effects. We do not find evidence of any pre-trends or anticipation effects two years or earlier prior to robot adoption. However, we see that the adjustment of wages starts one year before robot adoption, which may be due to planning for automation and/or robotics-related reorganization before the actual installation of robots.

In the right panel of Figure 4, we additionally control for worker fixed effects, and the overall pattern remains very similar. The evidence confirms that there are no pre-trends in hourly

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

TABLE 2 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION

Dependent variable:	*Hourly wage (log)*			*Employment status of existing workers*			*Hours worked (log)*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter	0.025	0.023	0.016	-0.003	-0.004	-0.000	-0.017	-0.013	-0.022
	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.0023)	(0.006)	(0.006)	(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	1,644,662	1,644,662	1,600,922	1,778,509	1,655,095	1,601,266
R^2	0.384	0.457	0.918	0.099	0.112	0.501	0.215	0.287	0.697

Notes: This table reports results from a regression of worker wages on firm-level robot adoption. In columns 4-6, we include employed workers and unemployed workers in their first year of unemployment. We attribute the characteristics of the last employer to workers who are currently unemployed. 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

wages, while wages seem to adjust in the year before adoption.[8]

Columns 4-6 turn to the impact of robot adoption on the employment status of existing workers and workers who became unemployed in the past year. Here, we do not find significant effects, but the point estimates are generally negative.[9] In columns 7-9, we look at hours worked, where we see negative and fairly stable estimates. For example, without worker-level controls, hours worked decline 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 range but larger, at −2.1%. The negative implications for hours worked, combined with positive wage impacts, suggest that there may be heterogeneous effects from robot adoption across worker types. Some workers receiving pay increases while others have their hours reduced. This is what we investigate next.

3.3 Effects of robots on different types of workers

We next turn to the heart of our analysis, where we allow for differential effects of robot adoption on directly-affected and indirectly-affected workers. In Table 3, we examine hourly wages. Different columns focus on different controls and our three measures of directly-affected workers. The main results confirm our conjecture about the juxtaposition of negative hours effects and positive wage effects, on average, masking major heterogeneous effects. For example, using the definition based on workers performing bluecollar-routine tasks, we find positive impacts on indirectly-affected workers and negative, equally precisely-estimated impacts on directly-affected workers. In column 1, the estimate for directly-affected workers is −0.054, while the impact on indirectly-affected workers is 0.034. Quantitatively, this implies that robot adoption increases

[8] Appendix C.1 presents longer post-event trends in hourly wages, but we prefer to focus on shorter post-event windows. Although the time span covered by our sample is long by the standards of other studies on robots, it is still shorter than in typical event-study settings. As a result, longer post-event windows create systematic changes in sample selection, because we do not know whether a firm adopted robots two years or more before the date in question, and we do not know whether a firm towards the end of our sample will adopt robots two years from now.

[9] Part of the reason we do not find employment effects on existing workers may be due to labor market rigidities in the Netherlands. Dutch firms operate under strict rules for dismissal. For example, grounds for dismissal are restricted to: *(i)* when an employee's performance is not satisfactory, the firm must prove that it has informed the employee and given them sufficient opportunities to improve; *(ii)* if an employee has serious conscientious objections to the business activities and the firm is unable to offer alternative work; and *(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/`.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

hourly wages of indirectly-affected workers by $(\exp(0.034) - 1) \cdot 100\% = 3.5\%$, while reducing it by about 5.3% for directly-affected workers.

The patterns are similar in column 2 when we include worker fixed effects, but the estimates are smaller. This difference could arise due to differential sorting, where robot adoption leads to changes in the composition of the workforce, or because robot-adopting firms were already employing workers with high earning potential (initial selection), meaning these firms were already different in their organizational or technological structures (for example, with different vertical task specialization structures as in Faia et al. 2023). We investigate this issue further in Table 4 below.

In column 3 of Table 3, we include firm-year fixed effects, which means we can only estimate the differential impact on directly-affected workers, and that is about 2.3% lower for directly-affected workers compared to other employees in the firm.

The results are quite similar in columns 4-6 when we use the replaceable worker definition of Graetz & Michaels (2018). 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, fairly precisely-estimated negative implications for directly-affected workers.

Table 4 explores the issue of differential sorting. We build on the literature studying the role of worker and firm fixed effects in wages and worker sorting across firms (see Abowd et al. 1999, Card et al. 2018, Bonhomme et al. 2019). We regress the worker fixed effects estimated in equation (8) on the same right-hand side variables, but also crucially in our context, on interactions of a dummy for a firm ever adopting a robot, interacted with directly- and indirectly-affected worker dummies. This "ever adopter" variable captures the possibility that robot-adopting firms systematically hire different types of workers into different positions, even before robot adoption.

Column 1 starts with the specification without firm fixed effects and relates worker fixed effects to these variables. It shows that worker fixed effects are systematically higher for indirectly-affected

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

TABLE 3 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON HOURLY WAGE – BY WORKER TYPE

Dependent variable:	*Hourly wage (log)*								
	Bluecollar-routine workers			*Replaceable workers*			*Low-education workers*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter×	0.034	0.015		0.032	0.014		0.049	0.018	
indirectly-affected worker	(0.007)	(0.005)		(0.007)	(0.005)		(0.007)	(0.005)	
Robot adopter×	-0.054	-0.008	-0.023	-0.061	-0.014	-0.026	-0.069	-0.013	-0.034
directly-affected worker	(0.015)	(0.008)	(0.008)	(0.014)	(0.009)	(0.009)	(0.012)	(0.007)	(0.008)
Directly-affected worker	-0.175	-0.001	0.002	-0.178	-0.007	-0.005	-0.189	-0.015	-0.013
	(0.003)	(0.007)	(0.006)	(0.002)	(0.006)	(0.006)	(0.002)	(0.005)	(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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Table 4 – Robot adoption – Worker sorting and initial selection

Dependent variable:	*The worker fixed effect, $\hat{\nu}_i$*				
	Bluecollar-routine workers				
	All workers	*All workers*	*Before robot adoption*	*New hires*	*Leavers*
	(1)	(2)	(3)	(4)	(5)
Robot adopter× indirectly-affected worker	0.027	0.009		0.009	-0.039
	(0.014)	(0.006)		(0.026)	(0.031)
Robot adopter× directly-affected worker	-0.002	0.014		0.024	-0.040
	(0.021)	(0.018)		(0.044)	(0.060)
Firm ever adopted robots× indirectly-affected worker	0.065	0.075	0.075	0.040	0.104
	(0.012)	(0.016)	(0.017)	(0.029)	(0.046)
Firm ever adopted robots× directly-affected worker	0.011				
	(0.018)				
Directly-affected worker	-0.176	-0.173	-0.173	-0.179	-0.184
	(0.003)	(0.003)	(0.003)	(0.006)	(0.007)
Control variables	✓	✓	✓	✓	✓
Firm fixed effects		✓	✓	✓	✓
Number of observations	731,288	731,288	700,716	73,548	62,402
R^2	0.218	0.431	0.431	0.437	0.451

Notes: This table reports regressions of worker fixed effects, $\hat{\nu}_i$, estimated from hourly wage regressions reported in Table 3 on hourly wages. The regression equation here is: $\hat{\nu}_i = \beta_1 r_{ft} a_{it} + \beta_2 r_{ft}(1 - a_{it}) + \beta_3 \bar{r}_f a_{it} + \beta_4 \bar{r}_f (1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \lambda_{t,f\in s} + \mu_{f\in m,t} + \epsilon$. In the above specification, β_1 and β_2 capture sorting effects (compositional changes) as a result of robot adoption. $\bar{r}_f$ is a time-invariant variable indicating whether a firm adopted robots any time during the sample period so β_3 and β_4 capture initial selection effects. New hires are defined as workers who were not employed in the firm in $t-1$ and leavers as workers who are not at the firm in $t+1$. 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.

workers at firms that adopt robots at some point in the sample. There is a much smaller and marginally significant additional impact after robot adoption, but this is generally not robust. In column 2, we include firm fixed effects, so we can only estimate the differential influence of robot adoption for indirectly-affected workers (relative to directly-affected workers). We see the same pattern as in column 1, with a large and significant effect from ever adopting a robot, although the additional impact of robot adoption is now indistinguishable from zero. Column 3 confirms the same results when focusing only on observations before robot adoption and shows an estimated differential selection of indirectly-affected workers to robot-adopters that is very similar to the previous two columns.

Finally, columns 4 and 5 look at the worker fixed effects of new hires (those not employed by the firm in $t-1$) and employees who leave the firm (those no longer with the firm in $t+1$) and show

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

no major compositional effects. The only significant coefficient is in column 5, but it goes in the opposite direction, implying that indirectly-affected workers who leave firms ever adopting a robot have higher earning potential. Moreover, this coefficient is not significantly different from the coefficient on new hires. Overall, there is little evidence for systematic compositional changes following robot adoption. We therefore conclude that estimates that include worker fixed effects are about 50% smaller than those that do not, mainly because firms that will later adopt robots already employ workers with higher earning potential in non-automated tasks, and these workers experience faster wage growth.

Table 5 turns to the implications of robot adoption for employment and hours worked. In Panel A, we find small positive and sometimes statistically significant impacts on the employment status of indirectly-affected workers, and negative, albeit typically insignificant, effects on directly-affected workers. For example, the estimates that control for worker fixed effects using the replaceable worker and low-education worker measures (columns 5 and 8) are positive and significant for indirectly-affected workers, and the evidence in columns 3, 6 and 9, from specifications that include firm-year fixed effects, indicates that there is a differential negative impact on directly-affected workers. We again reiterate that the somewhat weaker employment results compared to the wage estimates may be due to the rigidities in the Dutch labor market, which presumably slow down or even prevent worker layoffs and may also discourage or delay hiring.

In Panel B of Table 5, we turn to the effects on hours worked. Without worker fixed effects, there is evidence of negative hours effects for both types of workers, though these estimates are typically not significant. When we include worker fixed effects, the estimates become more negative and statistically significant. In any case, consistent with our theoretical expectations, the impact on directly-affected workers is larger (in absolute value).[10]

[10]The firm-level results in Table B1 in Appendix B.2 demonstrate a positive and statistically significant effect of robot adoption on hours worked. However, in the worker-level results, we can only include currently and previously employed workers (in order to include firm and worker fixed effects). This limitation explains the discrepancy between the firm-level and worker-level results. To estimate the overall effect of robot adoption on employment, it is necessary to consider all employees of robot-adopting firms, including new hires.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

TABLE 5 – WORKER-LEVEL EFFECTS OF ROBOT ADOPTION ON EMPLOYMENT – BY WORKER TYPE

Dependent variable:	*Employment status of existing workers*								
PANEL A: Employment	*Bluecollar-routine workers*			*Replaceable workers*			*Low-education workers*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter×	-0.000	0.003		0.000	0.004		0.001	0.005	
indirectly-affected worker	(0.002)	(0.002)		(0.003)	(0.002)		(0.003)	(0.002)	
Robot adopter×	-0.004	-0.008	-0.012	-0.002	-0.008	-0.012	-0.004	-0.007	-0.012
directly-affected worker	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)	(0.004)	(0.005)	(0.005)
Directly-affected worker	-0.003	0.006	0.004	-0.002	0.004	0.001	-0.004	-0.002	-0.001
	(0.001)	(0.006)	(0.005)	(0.001)	(0.005)	(0.005)	(0.001)	(0.003)	(0.003)
Control variables variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Worker fixed effects		✓	✓		✓	✓		✓	✓
Firm fixed effects	✓	✓		✓	✓		✓	✓	
Firm×year fixed effects			✓			✓			✓
Number of observations	760,896	725,962	698,020	795,251	758,810	731,120	782,502	747,611	719,759
R^2									

Dependent variable:	*Hours worked (log)*								
PANEL B: Hours worked	*Bluecollar-routine workers*			*Replaceable workers*			*Low-education workers*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot adopter×	-0.024	-0.015		-0.022	-0.012		-0.019	-0.009	
indirectly-affected worker	(0.008)	(0.007)		(0.007)	(0.007)		(0.007)	(0.007)	
Robot adopter×	-0.011	-0.027	-0.019	-0.021	-0.034	-0.026	-0.037	-0.032	-0.023
directly-affected worker	(0.011)	(0.014)	(0.015)	(0.012)	(0.017)	(0.018)	(0.011)	(0.013)	(0.014)
Directly-affected worker	0.003	-0.010	-0.020	0.000	-0.005	-0.013	0.002	0.007	0.004
	(0.003)	(0.012)	(0.012)	(0.003)	(0.012)	(0.012)	(0.002)	(0.008)	(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. In Panel A, for the employment status variable, we include employed workers and unemployed workers in their first year of unemployment. We attribute the characteristics of the last employer to workers who are currently unemployed. 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

3.4 Robustness checks

First, as indicated before, we report extended event-study estimates for the changes in hourly wages of workers in response to the future robot adoption of firms. As detailed in Appendix C.1, there appear to be no significant effects on hourly wages up to two years prior to the adoption of robots. In the year before adoption, hourly wages start adjusting, completing most of their adjustment during this year, likely due to the anticipated effects of oncoming robot adoption.[11]

Second, although the event studies indicate a rather immediate wage increase following robot adoption, one may still worry about anticipation and adjustment issues. In Appendix C.2, we replicate our baseline analysis from Table 2 using data only from the first and last years. This yields a selective sample, leveraging variation only from firms existing in 2009 and 2020 and workers active in the labor market during these years. Nevertheless, the results are very similar to the baseline results reported earlier.

Third, we examine the robustness of our results to allowing the effects of robot adoption to differ between workers with different characteristics. In Appendix C.3, we find that older workers (≥ 55) and women are more negatively impacted by robot adoption. However, when we control for the interaction of robot adoption with demographics, the key finding that directly-affected workers are negatively affected by robot adoption is robust.

Fourth, we currently match workers in the LEED to workers in the past 10 LFS waves to obtain their occupational information (recall equations (3), (4), and (5)). This matching procedure increases the number of observations but may exacerbate measurement error, *e.g.*, if workers have switched occupations in the meantime. To address this issue, in Appendix C.4 we limit this matching window to one year, which reduces the number of observations but alleviates 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% once a firm adopts robots. However, due to the increase in standard errors, these estimates are not statistically different from our baseline

[11] Event-study estimates that separately look at directly-affected and indirectly-affected workers are noisier, and, to save space, we do not present them here.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

effects.

Fifth, the allocation of workers to a specific worker type—directly- and indirectly-affected—is largely fixed throughout the study period, with very few exceptions. Nonetheless, in Appendix C.5, we examine whether the wage changes observed in the baseline results can be attributed to workers transitioning from indirectly- to directly-affected or vice versa. To address this concern, we focus exclusively on workers whose occupations were recorded prior to the adoption of robots. The results are unaffected by this selection, suggesting that the observed wage changes cannot be explained by worker type evolutions.

Sixth, Appendix C.6 further examines the two sources of identifying variation driving our baseline results with worker fixed effects: within-firm wage changes over time for existing workers and wage changes for workers moving between firms. The results for existing workers closely mirror our baseline findings. This suggests that the observed effects are primarily driven by within-firm wage changes.

Seventh, in Appendix C.7, we study job mobility between firms as a response to robot adoption by worker type. We find weak evidence that job mobility is higher for directly-affected workers after robot adoption compared to indirectly-affected workers.

Eighth, we offer a more detailed analysis of the impact of robot adoption by worker skill groups and confirm in Appendix C.8 that, when we distinguish between medium and low-education workers, the negative effects of robot adoption are more pronounced for the lowest-education category.

Last, in Appendix C.9, we investigate the effects of robot adoption on a fourth outcome measure, *personal income*, which combines information from all three of our measures (wage, the extensive margin of employment captured by our employment dummy, and the intensive margin represented by hours worked). The results 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

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

worker fixed effects).

3.5 Effects of competitors' robot adoption on workers

In this subsection, we focus on the effect of competitors' robot adoption, $r^{\mathcal{C}}_{ft}$, on workers' outcomes, following the specification in equation (2). We use analogous models as before to study the impacts of competitors' robot adoption on directly-affected and indirectly-affected workers, except that our sample now only includes workers from non-adopting firms. As the identifying variation comes from the differences in competition *between* 4-digit industries, we cannot control for 4-digit industry-by-year fixed effects and instead include 2-digit industry-by-year fixed effects. Specifically, we estimate models of the following form:

$$\log w_{ift} = \gamma_1 r^{\mathcal{C}}_{ft} a_{it} + \gamma_2 r^{\mathcal{C}}_{ft}(1 - a_{it}) + \delta a_{it} + \zeta z_{it} + \kappa_f + \lambda_{s_f,t} + \mu_{m_f,t} + \nu_i + \epsilon_{fmt}, \tag{10}$$

where $r^{\mathcal{C}}_{ft}$ is robot adoption in the same 4-digit industry, and $\lambda_{s_f,t}$ denotes 2-digit sector-by-year fixed effects.

Industry robot adoption may be endogenous, for example, because it is correlated with other technological investments in the industry.[12] There may also be some attenuation in the estimates of the effects of competitors' robot adoption, especially since we do not have product-level sales information. Motivated by these concerns, we report instrumental-variables (IV) estimates. In particular, we follow the strategy in Acemoglu & 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 ahead of 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 an instrument:

$$r^{\mathcal{E}}_{\tilde{s}t} = \frac{\mathcal{R}_{\tilde{s},t-5} - \mathcal{R}_{\tilde{s},\underline{t}-5}}{n_{\tilde{s},t-5}}, \tag{11}$$

where $\tilde{s}$ refers to the *IFR* sector, $\mathcal{R}_{\tilde{s},t}$ is the total number of robots in Korea and Taiwan in

[12]As 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Table 6 – Worker-level effects of robot competition on hourly wages

Dependent variable:	*Hourly wage (log)*		
	(1)	(2)	(3)
Competition by robot adopters	0.211	0.276	0.285
	(0.074)	(0.073)	(0.076)
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
Kleibergen-Paap F-statistic	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 (11). 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.

sector $\tilde{s}$ in year t, and $n_{\tilde{s},t-5}$ is the total employment in sector $\tilde{s}$ in the Netherlands in year $t-5$. Because $r^{\mathcal{E}}_{\tilde{s}t}$ has some extreme outliers, we cap the instrument at its 99$^{\text{th}}$ percentile value.[13] Because the instrument varies only at the IFR-industry level, we cluster our standard errors at the IFR-industry-by-year and worker levels.[14]

Table 6 reports estimates of the overall effects of robot adoption by competitors on hourly wages. The relevant first stages are reported in Appendix C.10 and show a strong relationship, with Kleibergen-Paap F-statistics exceeding 10 in all specifications.

In Table 6, there are positive impacts on hourly wages, but as we will see, this result is a combination of both negative and positive effects on different types of workers.

Table 7 explores heterogeneous effects by worker type. We find positive impacts on hourly wages from competitors' robot adoption for indirectly-affected workers and negative impacts for

[13] The *IFR* data are from 2004-2014. Hence, for 2020 we would need data for 2015. We predict robots in 2015 by a linear trend of robot adoption in each sector in each country between 2010 and 2014. In Appendix B.7 we also provide estimates based on a lag of 6 years so that this extrapolation is not necessary. Our results remain robust to this variation.

[14] Our firm-level results in Appendix B.7 suggest that there are sizable negative effects of robot adoption 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

TABLE 7 – WORKER-LEVEL EFFECTS OF ROBOT COMPETITION ON HOURLY WAGE – BY WORKER TYPE

Dependent variable:	*Hourly wage (log)*								
	Bluecollar-routine workers			*Replaceable workers*			*Low-education workers*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Competition by robot adopters× indirectly-affected worker	0.258	0.572		0.286	0.536		0.396	0.666	
	(0.117)	(0.175)		(0.116)	(0.170)		(0.120)	(0.184)	
Competition by robot adopters× directly-affected worker	-0.109	0.165	-0.353	-0.113	0.058	-0.433	-0.176	0.149	-0.449
	(0.136)	(0.185)	(0.097)	(0.142)	(0.202)	(0.101)	(0.136)	(0.185)	(0.098)
Directly-affected worker	-0.161	0.015	0.014	-0.162	0.010	0.011	-0.170	-0.003	-0.004
	(0.004)	(0.007)	(0.007)	(0.004)	(0.008)	(0.007)	(0.003)	(0.006)	(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 F-statistic	29.04	12.25	62.29	28.26	11.42	54.63	28.31	11.10	83.88

Notes: This table reports heterogeneous wage effects from growth adoption by competitors. The sample includes only workers from non-adopting firms. Competition by robot adoption is computed as the share of sales by robot adopting firms within the same 4-digit industry. Competition is instrumented by robot exposure as defined by equation (11). 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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

directly-affected workers. For instance, in column 1, which looks at bluecollar-routine workers, a one standard deviation increase in competitors' robot adoption raises hourly wages by 3% for indirectly-affected workers. In contrast, the impact on directly-affected workers is negative, though imprecisely estimated. When firm-year fixed effects are included in column 3, the negative impact on directly-affected workers becomes more precise. The pattern holds across other measures of directly-affected workers.

Overall, these results suggest that workers performing tasks later automated by robots face negative impacts from both their own firm's and competitors' robot adoption, while indirectly-affected workers benefit from both. This likely occurs because more robot-intensive production increases demand for their services. Even non-adopting firms end up paying more to indirectly-affected workers, as robot-adopting firms increase their hiring of such workers, pushing up wages across the board.

In Appendix C.11, we examine the effects of competitors' robot adoption on employment and hours worked. These results are less precisely estimated, with overall effects close to zero, but the differential impacts on directly-affected workers are negative, though not statistically significant. We also find that the results for personal income mirror the hourly wage results, showing positive effects for indirectly-affected workers and negative effects for directly-affected workers. The differential effects are substantial, with a one standard deviation increase in competitors' robot adoption reducing the personal incomes of directly-affected workers by approximately 4.6%.

4 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 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, which aligns with negative industry-level effects. This literature has not focused on worker-level

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

outcomes, particularly on which types of workers are positively or negatively impacted by robot adoption.

This paper investigates the worker-level implications of robot adoption using high-quality data on robot imports, spanning a comparatively long time period. We combine our robot adoption data with detailed linked employer-employee data from the Dutch manufacturing sector. The Dutch economy provides an interesting context, since it has invested substantially in automation technologies, 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 quite similar in the Netherlands to those observed 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.

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, the task-based framework implies that workers performing tasks that can 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 & Restrepo 2020, Acemoglu et al. 2022). These are: workers employed in bluecollar-routine tasks, those employed in replaceable tasks (as defined in Graetz & Michaels 2018), and workers with low education levels. Consistent with theoretical expectations, using all three measures we find that robot adoption either by one's 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. We also show that controlling for worker

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

fixed effects is important because robot-adopting firms tend to hire workers with higher earning potential in non-automated tasks (even before they adopt robots).

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, while our paper shows that robot adoption generates different effects on different types of workers, future research may further delve into the inequality implications of robots and other automation technologies. Recent work by Acemoglu & Restrepo (2022*b*) 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.

5 Data availability

The replication package provides the full code to reproduce all tables and figures in this article, together with detailed documentation on data construction and instructions for applying for access to the underlying confidential administrative microdata from Statistics Netherlands. Due to confidentiality restrictions, the data cannot be shared publicly; instead, the package includes do-files, variable documentation, and guidance for authorised access. The replication materials are available as Acemoglu et al. (2025), Harvard Dataverse, `https://doi.org/10.7910/DVN/OWGYIS`.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

References

Abowd, J. M., Kramarz, F. & Margolis, D. N. (1999), 'High Wage Workers and High Wage Firms', *Econometrica* **67**(2), 251–333.

Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P. & 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., Koster, H. R. A. & Ozgen, C. (2025), 'Replication Data for: Robots and Workers'. **URL:** *https://doi.org/10.7910/DVN/OWGYIS*

Acemoglu, D., Lelarge, C. & Restrepo, P. (2020), 'Competing with Robots: Firm-level Evidence from France', *American Economic Review: Papers and Proceedings* **110**, 383–388.

Acemoglu, D. & Restrepo, P. (2020), 'Robots and Jobs: Evidence from US labor markets', *Journal of Political Economy* **128**(6), 2188–2244.

Acemoglu, D. & Restrepo, P. (2022*a*), 'Demographics and Automation', *The Review of Economic Studies* **89**, 1–44.

Acemoglu, D. & Restrepo, P. (2022*b*), 'Tasks, automation, and the Rise in U.S. Wage inequality', *Econometrica* **90**(5), 1973–2016.

Aghion, P., Antonin, C., Bunel, S. & Jaravel, X. (2021), 'What Are the Labor and Product Market Effects of Automation? New Evidence from France'.

Autor, D. & 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.

Barth, E., Røed, M., Schøne, P. & Umblijs, J. (2020), 'How Robots Change Within-Firm Wage Inequality', *IZA Discussion Paper No. 13605* .

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Bessen, J., Goos, M., Salomons, A. & 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. & Gancia, G. (2020), 'Robot Imports and Firm-Level Outcomes', *CESifo Working Paper no. 8741* .

Bonhomme, S., Lamadon, T. & Manresa, E. (2019), 'A Distributional Framework for Matched Employer Employee Data', *Econometrica* **87**(3), 699–739.

Card, D., Cardoso, A. R., Heining, J. & Kline, P. (2018), 'Firms and Labor Market Inequality: Evidence and Some Theory', *Journal of Labor Economics* **36**(S1), S13–S70.

Dauth, W., Findeisen, S., Suedekum, J. & Woessner, N. (2021), 'The adjustment of labor markets to robots', *Journal of the European Economic Association* **19**(6), 3104–3153.

Dixon, J., Hong, B. & Wu, L. (2021), 'The robot revolution: Managerial and employment consequences for firms', *Management Science* **67**, 5586–5605.

Faia, E., Ottaviano, G. I. P. & Spinella, S. (2023), 'Robot adoption, worker-firm sorting and wage inequality: Evidence from administrative panel data', *Centre for Economic Performance, Discussion Paper* (1902).

Graetz, G. & Michaels, G. (2018), 'Robots at Work', *Review of Economics and Statistics* **100**(5), 753–768.

Hirvonen, J., Stenhammar, A. & 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* .

Koch, M., Manuylov, I. & Smolka, M. (2021), 'Robots and firms', *The Economic Journal* **131**(638), 2553–2584.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/740218 in an upcoming issue of Journal of Labor Economics,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Koster, H. R. A. & Ozgen, C. (2021), 'Cities and Tasks', *Journal of Urban Economics* **126**(103386).

Oesch, D. (2013), *Occupational Change in Europe*, Oxford University Press.

Robots.com (2013), 'What is the cost of a robot?', *www.robots.com* .

Visser, M., Gesthuizen, M., Kraaykamp, G. & 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.