

Automatic Reaction – What Happens to Workers at Firms that Automate? *

James Bessen[†]
Boston University

Maarten Goos[‡]
Utrecht University

Anna Salomons[§]
Utrecht University

Wiljan van den Berge[¶]
CPB

January 2019

Abstract

We provide the first estimate of the impacts of automation on individual workers by combining Dutch micro-data with a direct measure of automation expenditures covering firms in all private non-financial industries over 2000-2016. Using an event study differences-in-differences design, we find that automation at the firm increases the probability of workers separating from their employers and decreases days worked, leading to a 5-year cumulative wage income loss of about 8% of one year's earnings for incumbent workers. We find little change in wage rates. Further, lost wage earnings are only partially offset by various benefits systems and are disproportionately borne by older workers and workers with longer firm tenure. Compared to findings from a literature on mass layoffs, the effects of automation are more gradual and automation displaces far fewer workers, both at the individual firms and in the workforce overall.

Keywords: Automation, Technological Change, Displacement

JEL: J23, J31, J62, J63, O33

*Helpful comments by David Autor, Guy Michaels, and participants of workshops at CPB, ROA Maastricht, and Utrecht School of Economics are gratefully acknowledged. James Bessen thanks Google.org for financial support.

[†]Boston University Technology & Policy Research Initiative, jbessen@bu.edu

[‡]Utrecht University School of Economics, m.goos@uu.nl

[§]Utrecht University School of Economics, a.m.salomons@uu.nl

[¶]CPB Netherlands Bureau for Economic Policy Analysis and Erasmus University Rotterdam, w.van.den.berge@cpb.nl

1 Introduction

Advancing technologies are increasingly able to fully or partially automate a broad range of job tasks, even ones that had previously been considered non-routine, such as interpreting X-rays to diagnose disease, selecting job applicants for interviewing, picking orders in a warehouse, or driving a car. These technologies range from robotics to speech recognition and other applications of artificial intelligence, and are being applied across a broad range of sectors of the economy. Some commentators warn that the pace of automation may be accelerating and the range of jobs affected is widening, threatening displacement across large shares of jobs in the near future (Frey and Osborne 2016; Ford 2015). A large literature on worker displacement¹ suggests that the effects of such developments could indeed be dire: individual workers subject to plant closings and mass layoffs experience reduced employment probabilities and wage scarring, leading to long-term earnings losses, as well as reductions in consumption and worse health outcomes. These discussions have led to a call for new policies to support workers displaced by automation, such as a universal basic income.

The potential for automation to displace workers is being taken seriously in recent labor market models where technology changes the comparative advantage of workers across job tasks (Autor et al. 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2018a,d,b; Benzell et al. 2016; Susskind 2017). In these theories, worker displacement is a possible outcome of automation as machines take over tasks previously performed by humans. Automation can also lead to worker displacement if the new technologies require workers with different skills than those of the incumbent workforce.

However, direct empirical evidence on the worker-level impacts of automation is lacking: this paper provides the first estimate of the economic impacts on workers when their firm invests in automation technology. We link an annual firm survey on automation costs to Dutch administrative firm and worker databases, allowing us to consider automation across all private non-financial economic sectors. The data are provided by Statistics Netherlands and cover years 2000-2016: we observe 36,085 firms with at least three years of automation cost data, employing close to 5 million unique workers per year on average.

This paper makes several contributions. Firstly, we directly measure automation at the firm level and can therefore study the worker impacts of automation where they originate: at the

¹Starting with Ruhm (1991) and Jacobson et al. (1993).

automating firms. Secondly, we develop and implement an empirical methodology combining event study with difference-in-differences analysis, leveraging the timing of firm-level automation events for identifying causal effects. Thirdly, we consider automation events as they occur across all private non-financial sectors of the economy rather than considering a specific automation technology in isolation. Fourthly, we measure a rich array of outcomes for individual workers for the years surrounding the automation event. These outcomes include wage earnings, daily wages, firm separation, days spent in non-employment, self-employment, early retirement, and unemployment insurance and welfare receipts. We also look separately at outcomes for incumbent workers employed three or more years at the firm prior to the automation event, and for the firm's more recent hires, and consider how impacts differ across worker characteristics. Lastly, we compare the impacts of displacement from automation to those arising from mass lay-offs and firm closures, both within firms and, using a back-of-the-envelope calculation, in the economy as a whole.

We find that automation at the firm increases the probability of workers separating from their employers. For incumbent workers (those with at least three years of firm tenure), this firm separation is followed by a decrease in annual days worked, leading to a 5-year cumulative wage income loss of about 8% of one year's earnings. Wage rates, however, are not much affected. Further, wage income losses are disproportionately borne by older workers and are only partially offset by various benefits systems. Compared to findings from a literature on mass layoffs, however, the effects of automation are more gradual and automation displaces far fewer workers, both at the individual firms and in the workforce overall.

While the related economic literature on technological change is large, little of it specifically addresses questions about automation. Indeed, empirical work on automation has so far mostly focused on robotics – a prime example of an automation technology, albeit one that has penetrated only a limited number of sectors– and on aggregate outcomes.² The evidence is mixed: Graetz and Michaels (2018) find that industrial robots have had positive wage effects and no employment effects across a panel of countries and industries, whereas Acemoglu and Restrepo (2018c) find that wages and employment have decreased in US regions most exposed to automation by robots. Applying Acemoglu and Restrepo (2018c)'s empirical design to German regions, Dauth et al. (2017) find evidence of positive wage effects, and no changes in total employment.

²Other papers have looked at cross-sectional features of automation in manufacturing including Doms et al. (1997) and Dinlersoz and Wolf (2018).

A broader empirical literature indeed makes clear that large-scale automation need not bring about labor displacement in aggregate, but rather, leads to labor reallocation. For one, even within the affected industry, automation can increase employment if industry demand is sufficiently elastic (Acemoglu and Restrepo (2018a,d); Bessen (2018)). Moreover, there is evidence that productivity gains generate employment increases in other industries through input-output linkages as well as final demand effects, offsetting any employment losses in automating industries (Autor and Salomons (2018); Gregory et al. (2018)). It should be noted that our analyses do not consider these countervailing forces: this implies our findings do not inform on the macroeconomic impacts of automation. However, to understand how automation affects work, it is critical to also study its effects on individual workers. After all, the absence of displacement in aggregate need not imply the absence of losses for individual workers directly affected by automation. These micro-level impacts are also of first-order importance for policymakers aiming to assuage adverse impacts out of distributional concerns.

This paper is structured as follows. We first introduce our data source, Dutch matched employer-employee data which we link to a firm survey containing a direct measure of automation expenditures. Section 3 contains our empirical approach, outlining a definition of automation events and the resulting estimation framework using a combination of event study and differences-in-differences. Our results are divided into total impacts on workers' wage income (section 4), which we decompose into firm separation and employment impacts (section 5), and daily wage impacts conditional on employment (section 6). We next consider to what extent wage income losses are compensated by various benefit schemes (section 7), and how these losses differ across worker types (section 8). Lastly, in section 9 we consider the worker costs of automation conditional on displacement and compare these to income losses arising from mass lay-offs and firm closures in section 10. The final section concludes.

2 Data

We use Dutch data provided by Statistics Netherlands. In particular, we link an annual firm survey to administrative firm and worker databases covering the universe of firms and workers in the Netherlands. The firm survey is called 'Production statistics' ("Productiestatistieken") and includes a direct question on automation costs— it covers all non-financial private firms with

more than 50 employees, and samples a subset of smaller non-financial private firms.³ This survey can be matched to administrative company (“Algemeen Bedrijfsregister”) and worker records (“GBA” and “BAAN” files).

Our data cover years 2000-2016, and we retain 36,085 unique firms with at least 3 years of automation cost data – together, these firms employ 4,962,682 unique workers annually on average. We remove firms where Statistics Netherlands indicate that the data are (partly) imputed. We further remove workers enrolled in full-time studies earning either less than 5,000 euros per year or less than 10 euros per day, as well as workers earning more than half a million euros per year or more than 2,000 euros on average per day.⁴ For workers observed in multiple jobs simultaneously, we only retain the one providing the main source of income in each year. We use their total earnings in all jobs as the main measure of wage income.

At the worker level, we observe gross wage income as well as days worked – since we do not observe hours worked, we use daily wages as a measure of wage rates. We further observe workers’ gender, age, and nationality⁵. A downside of these data is that we do not observe workers’ occupations nor their education level: the former is unavailable entirely, whereas the latter is only defined for a small and selected subset of workers (with availability skewed towards the high-educated). We further match worker-level data to administrative records on receipts from unemployment, welfare, disability, and retirement benefits. We can track workers across firms on a daily basis, allowing us to construct indicators for firm separation and days spent in non-employment.

The main advantage of the dataset we construct is the availability of a direct measure of automation at the firm level. In particular, “Automation costs” is an official bookkeeping term defined as costs of third-party automation services. This also includes non-activated purchases of custom software and costs of new software releases, but excludes prepackaged software licensing costs. These automation costs are not restricted to manufacturing firms and could include expenditures related to automation technologies such as self-service check-outs, warehouse and storage systems, or automated customer service. Another example are robotics integrator services highlighted (and used as an instrument for robotic technology adoption) in Acemoglu and

³Firms are legally obliged to respond to the survey when sampled. However, the sampling design implies our data underrepresent smaller firms: we will examine effect heterogeneity across firm size classes to consider how this sample selection affects our overall findings.

⁴In Appendix A.3 we perform robustness checks from several other sample restrictions, including removing firms with outlier employment changes and those undergoing events such as mergers and acquisitions.

⁵In these data, individuals are classified as “Dutch” if they themselves and both of their parents have been born in the Netherlands.

Table 1: Automation cost share distribution

	All observations		Automation costs >0	
	<i>Cost level</i>	<i>Cost share</i>	<i>Cost level</i>	<i>Cost share</i>
p5	0	0.00%	2,026	0.03%
p10	0	0.00%	3,663	0.06%
p25	0	0.00%	9,775	0.14%
median	10,988	0.15%	27,912	0.31%
p75	49,446	0.47%	87,115	0.67%
p90	179,641	1.05%	283,703	1.36%
p95	422,000	1.69%	658,695	2.11%
mean	196,351	0.45%	284,692	0.65%
% of N with 0 costs	31.03%	31.03%	0.00%	0.00%
N firms × years	241,292	241,292	166,418	166,418

Notes: Automation cost level in 2010 euros, automation cost share as a percentage of total costs, excluding automation costs. The number of observations is the number of firms times the number of years.

Restrepo (2018d). While the downside of this measure is that we do not know the exact automation technology being used by the firm, it does capture all such technologies rather than focusing on a single one, and we measure it at the level of the firm rather than industry.

Table 1 shows summary statistics on annual automation costs for firms, both in levels and as a percentage of total costs (excluding automation costs). This highlights several things. Firstly, we observe zero automation expenditures for almost one-third of firm-year observations. Further, the average automation cost share is 0.45 percent, corresponding to an outlay of around 200,000 euros annually, but this distribution is highly right-skewed as the median is only 0.15 percent – this is true even when removing observations with zero automation costs.

Table 2 further shows how these automation costs and cost shares differ by broad sector. Our comprehensive measure of automation technologies indicates that all sectors have automation expenditures, though there is substantial variation at the firm level within each of these broad sectors. The highest mean automation cost share is observed in Professional, scientific, and technical activities, followed by Information and communication (which also has the highest level), and Administrative and support activities; conversely, Construction has the lowest share of automation expenditures in total costs. Most variance in automation cost expenditures can be observed in wholesale and retail trade: this may of course be because this broad sector covers firms engaged in very different activities which require different automation technologies. While we do not use either this sectoral or between-firm variation in our empirical identification strategy, we will consider effect heterogeneity across sectors since the nature of automation

Table 2: Automation costs by sector

Sector	Cost level	Cost share		Nr of observations	
	<i>Mean</i>	<i>Mean</i>	<i>Std dev</i>	<i>Firms</i>	<i>Firms × yrs</i>
Manufacturing	395,310	0.36%	0.63%	5,605	45,308
Construction	72,400	0.20%	0.36%	4,529	28,927
Wholesale and retail trade	107,600	0.33%	7.51%	11,086	76,283
Transportation and storage	255,270	0.40%	0.98%	3,194	21,779
Accommodation and food serving	53,550	0.30%	0.49%	1,176	6,355
Information and communication	427,080	0.87%	3.10%	2,651	16,804
Professional, scientific, and technical activities	142,720	1.05%	1.80%	3940	22,826
Administrative and support activities	123,030	0.50%	1.24%	3,904	23,010

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,085; Total N firms × years is 241,292.

Table 3: Automation costs by firm size class

Firm size class	Cost level	Cost share		Nr of observations	
	<i>Mean</i>	<i>Mean</i>	<i>Std dev</i>	<i>Firms</i>	<i>Firms × yrs</i>
0-9 employees	22,670	0.35%	1.05%	9,241	48,869
10-19 employees	24,890	0.47%	9.34%	8,036	49,112
20-49 employees	48,750	0.44%	1.61%	9,771	68,872
50-99 employees	105,130	0.44%	1.04%	4,464	34,758
100-199 employees	265,650	0.49%	1.14%	2,344	19,825
200-499 employees	712,990	0.59%	1.37%	1,446	12,620
≥500 employees	3,285,220	0.70%	1.78%	783	7,236


Notes: Automation costs as a percentage of total costs, excluding automation costs. Total N firms is 36,085; Total N firms × years is 241,292.

technologies may be sector-specific.

Table 3 reports the same statistics but separately by firm size class, grouped into 7 classes used by Statistics Netherlands: the smallest firms have fewer than 10 employees whereas the largest have more than 500. Unsurprisingly, automation cost levels rise with firm size: firms with fewer than 20 employees spend around 25,000 euros annually on automation services, whereas the largest firms spend close to 3.3 million. Less obviously, this table also reveals that automation cost shares increase with firm size: the very smallest firms have average automation cost shares of 0.35%, whereas firms with between 10 to 200 employees have a cost share of around 0.48%. This increases to 0.59% for firms between 200 and 500 workers, and 0.70% for firms with more than 500 workers. Again, however, there is substantial variation within these size classes.

Figure 1 further shows how the distribution of automation cost shares changes over time: this shows that the mean automation cost share is rising in the Netherlands, from 0.27% to 0.57% relative to total costs over 2000-2016. All else equal, this implies workers' exposure to automation is also rising. Furthermore, besides an increase in the average, there is a fanning

Figure 1: Firm-level automation cost shares over time



out of the distribution with automation cost shares rising faster for higher percentiles.⁶

Lastly, we find that automation costs are not correlated with computer investment⁷, highlighting that this is a different type of technology. In a future version of this paper, we will consider how the impact of automation differs from that of computerization, a more commonly considered type of technological change.

3 Empirical approach

3.1 Defining automation cost spikes

The main challenge for empirically identifying the worker-level impacts of automation lies in finding a group of workers who can be used as a control group. A further challenge is to distinguish automation events at the firm level, especially when using survey data. Our approach for both of these challenges is to use what we term *automation spikes*. In particular, we hypothesize that spikes in automation cost shares at the firm level signal changes in work processes related to automation.

⁶In a future version of this paper, we will additionally consider to what extent this evolution of the automation cost share distribution is driven by changes within and between firms.

⁷Computer investment is available in another firm-level survey: the coverage of this survey partially overlaps with the sample of firms where we observe automation costs.

We define automation cost spikes as follows. Firm j has an automation cost spike in year τ if its real automation costs $AC_{j\tau}$ relative to real total operating costs (excluding automation costs) averaged across all other years $t \neq \tau$ ($\overline{TC}_{j,t \neq \tau}$) are at least thrice the average firm-level cost share (again excluding year τ):

$$spike_{j\tau} = \mathbb{1} \left\{ \frac{AC_{j\tau}}{\overline{TC}_{j,t \neq \tau}} \geq 3 \times \frac{\overline{AC}_{j,t \neq \tau}}{\overline{TC}_{j,t \neq \tau}} \right\}, \quad (1)$$

where $\mathbb{1}\{\dots\}$ denotes the indicator function. As such, a firm that has automation costs around one percent of all other operating costs for year $t \neq \tau$ will be classified as having an automation spike in $t = \tau$ if its automation costs in τ exceed three percent of average operating costs over years $t \neq \tau$.

Note that this is a firm-specific measure, intended to identify automation events that are large for the firm, independent of that firm's initial automation expenditure level. As such, this indicator does not mechanically correlate with firm characteristics such as firm size, sector, or capital-intensity. Although we could possibly exploit the size of the automation spike, this is not our baseline specification for a number of reasons. Firstly, there may be measurement error in the survey variable making it more difficult to measure the exact size of a spike. Secondly, we use the automation costs survey variable to flag automation events, but other (indirect) costs may be incurred which are not directly surveyed: as such, our baseline approach identifies automation events without taking a strong stance on their exact size.⁸

3.2 Summary statistics on automation cost spikes

We now document the existence and frequency of automation spikes by firm and sector. In order to identify spikes, we need at least three years of automation cost data at the firm level: this is the sample of 36,085 firms described above.

Table 4 shows that close to two-thirds of firms never spike, whereas the remaining third spike at least once over the 16 years of observation. Note that non-spiking firms do not necessarily have zero automation costs: it is just that their automation expenditures do not fluctuate much as a percentage of total costs, implying they do not undergo large automation events as we define them. Out of the firms that do have such an event, the large majority spikes only once

⁸Preliminary results varying the spike size threshold do suggest impacts are larger when using more stringent spike definitions. We therefore also report results scaled by the absolute size of the average automation spike, below.

over 2000-2016, although some spike twice and up to four times at most.

Figure 2 shows what automation spikes look like on average across firms where spikes are observed. This is constructed by defining time relative to the spike year, such that all spikes line up in $t = 0$. When firms spike multiple times, we only include the largest spike. Note that this is not a balanced panel of firms: rather, all 12,278 spiking firms are observed in $t = 0$, and the number of observations for other years depends on when the spike took place (if in the first year of data, there are no observations for $t < 0$; if in the last year, there are no observations for $t > 0$), and on how often the firm enters in the automation survey. Nevertheless, we see a clear spike pattern.

Figure 3 restricts the sample to treated firms and our estimation window: these are the spikes we can actually use for identification of treatment. This shows that these automation events are quite cleanly identified – automation events are not preceded by a substantial lead-up of automation spending relative to total costs, nor is there evidence of much slow tapering off afterwards. Rather, automation spike years stand out as years when the firm made a large (relative to its normal automation expenditure share) investment in automation. Figure 4 plots this same graph but for the implied level of automation expenditure, showing that the average firm-level automation spike amounts to an investment of close to 200K euros, compared to a usual level of around 70K euros in years close to the spike, as shown in Figure 4. Of course, larger firms have higher absolute levels of expenditure: if we weight these numbers by firms' number of workers in $t = -1$, we find that the average spike size is 4.0 million euros, compared to 1.2 million euros on average in other years.⁹

The existence of these automation cost spikes is consistent with a literature on lumpy investment (Haltiwanger et al. 1999; Doms and Dunne 1998). Indeed, such spikes occur when the investment is irreversible and there are important indivisibilities. Under uncertainty, irreversibility creates an option value to waiting (Pindyck 1991; Nilsen and Schiantarelli 2003); whereas indivisibilities can arise from fixed adjustment costs (Rothschild 1971) – together, this implies investment occurs in relatively infrequent episodes of disproportionately large quantities. It is plausible that investments in automation meet these two criteria: major automation investments likely include both substantial irreversible investments, for example in terms of worker training or from developing custom software; as well as introduce fixed adjustment costs from


⁹In Figures 15 and 16 in the Appendix, we show that the same patterns hold when considering an entirely balanced sample of treatment firms where we observe automation cost share information in every single year.

Table 4: Firm-level automation spike frequency

Nr of spikes	Frequency	Percent
0	23,807	66.0%
1	9,572	26.5%
2	2,297	6.4%
3	359	1.0%
4	46	0.1%

Notes: Total N firms is 36,085.

Figure 2: All firm-level automation cost share spikes




Notes: N=12,278 in $t = 0$.

reorganizing production processes.

3.3 An event study differences-in-differences design


We now outline our empirical design to leverage the observed automation cost spikes for identification. Three measures help identification. Firstly, as mentioned above, automation spikes are by definition a big event for the firm – by construction, this is not “run-of-the-mill” automation, aiding identification. The idea is that such major events help distinguish the effects from other factors that might affect worker outcomes by making the signal strong relative to the noise. This is especially important since our automation cost measure – albeit an official bookkeeping term– derives from a survey.

Figure 3: Automation cost share spikes for treatment firms



Notes: N=2,753 in $t = 0$.

Figure 4: Automation cost spike level for treatment firms



Notes: N=2,753 in $t = 0$.

Secondly, we assume that the timing of automation cost spikes is essentially random conditional on observables. This means that for each treatment group of workers experiencing a spike in a given year, we can select a control group of workers at firms that spike in a later year. The workers and firms in the control group should have similar observed and unobserved characteristics to the treatment group. Third, we match the control group workers to the treatment group workers.

Our use of timing differences across firms is in the spirit of a recent literature exploiting event timing differences in other contexts, e.g. see Duggan et al. (2016); Fadlon and Nielsen (2017); Miller (2017); Lafortune et al. (2018). In the context of automation, our identification relies in part on the nature of major automation events. Indeed, as argued above, because these investments typically involve both uncertainty about the payoff and irreversible investments, they can create substantial option value to waiting to invest. This means that small differences in the payoffs to automating can generate substantial differences in the timing of investment.¹⁰ This sensitivity implies that small, idiosyncratic differences can change the exact timing of automation events across firms. Consequently, cohorts of firms that spike a few years apart should have similar characteristics, their workers should also have similar characteristics on average, and can thus serve as a counterfactual. For each treatment cohort we define a control group of workers from firms that spike 5 or more years after the treatment cohort.

Given this assumption, we can use a differences-in-differences (DiD) model where treatment is based on automation spike timing. Our worker-level event study DiD specification is

$$y_{ijt} = \alpha + \beta \text{treat}_i + \sum_{t \neq -1; t=-3}^4 \gamma_t \times I_t + \sum_{t \neq -1; t=-3}^4 \delta_t \times I_t \times \text{treat}_i + \lambda X_{ijt} + \varepsilon_{ijt}, \quad (2)$$

where i indexes workers, j firms, and t time, henceforth measured as years relative to the automation spike in period τ , i.e. $t \equiv \text{year} - \tau$. y_{ijt} is the outcome variable (such as total wage earnings, annual days in non-employment, wages conditional on working, and firm separation), and treat_i is a treatment indicator, equal to 1 for worker i if their firm is experiencing an automation spike at time $t = 0$, and 0 otherwise (i.e. if their firm spikes in $t = 5$ or later). Further, I_t are indicators for time relative to the spike year, where we choose $t \in \{-3, 4\}$, with $t = -1$ as the reference category. Lastly, X_{ijt} are controls: these are a set of worker characteristics (age and age squared, gender, and nationality); sector and size class of the spiking

¹⁰For example, Bessen (1999) finds that a 6 percent payoff difference generated a decade difference in when firms chose to switch from mule-spinning to ring-spinning in the British textile industry.

firm; as well as fixed effects for years. We cluster standard errors at the worker level.

Workers at a firm are treated in year t if that firm undergoes an automation cost share spike (or its largest automation cost spike if multiple spikes are observed) in year t , and is not observed to spike again during the estimation window ($t = -3$ to $t = 4$). Control group workers are those employed at firms who spike at $t + 5$ or later, but not during the estimation window. We exclude both treatment and control group firms with multiple spikes in the estimation window such that estimates of pretrends and treatment lags are not contaminated.¹¹

In the event study differences-in-differences (DiD) literature, never-treated units are sometimes used for identification as well (e.g. see Miller 2017; Lafortune et al. 2018). Our specification, however, only considers workers who are employed in firms that spike at some point over 2000-2016: we exclude workers at firms that never spike. As such, our specification strictly exploits differences in event timing rather than also using event incidence for identification. We make this choice since including firms that never experience an automation spike may be on different growth trajectories, leading to different outcomes for their workers.¹²

Further, we divide treated workers into two groups, incumbent workers and recent hires. We define incumbent workers as those who have been at the same firm for at least three years prior to the treatment year. On average across firms, 64 percent of workers are incumbents (where the median is 70 percent) in $t - 1$. Incumbent workers at a firm are treated in year t if that firm undergoes an automation cost share spike (or its first automation cost spike if multiple spikes are observed) in year t . Incumbent workers at firms who spike at $t + 5$ or later are used as controls. These incumbent workers are similar to definitions used in the mass lay-off literature (e.g. see Jacobson et al. 1993). Dutch labor law typically ensures that temporary contracts are of a maximum duration of two years, implying that workers with three years of tenure are very likely to have a stable working relationship with the firm.

The second group of treated workers we identify are those with less than three years of firm tenure prior to the automation event: compared to incumbent workers, these workers have been hired relatively recently – we therefore refer to them as recent hires. This worker group is more likely to hold temporary contracts, which could imply different treatment effects.

¹¹Although this sample restriction conceptually provides for cleaner identification, our results are similar when not imposing this restriction.

¹²An additional consideration is that we do not observe firms with fewer than 50 employees in the survey in all years: firms where we do not observe at least three automation cost observations may still experience an (unobserved) automation event yet be assigned to the control group when including never-spikers, leading to possible contamination of the treatment effect estimate.

Moreover, causal identification of the treatment effect for recent hires could prove more difficult as they may have been hired in anticipation of the automation event. We therefore analyze them separately, and generally put more stock in our results for incumbent workers.

In equation 2, the parameters of interest are δ_t : these estimate period t treatment effect ($t > 0$) relative to pre-treatment period $t = -1$. As with all DiD models, identification requires parallel trends in the absence of treatment. Our timing strategy provides some support for the parallel trends assumption. Assuming that later-spiking firms have very similar worker and firm characteristics to the treatment firms, we expect outcome variables to follow similar trends in a counterfactual without treatment. We can strengthen identification further by matching on worker and firm observables to ensure that $\delta_t = 0$ for all $t < 0$, making the assumption of parallel trends more plausible (Azoulay et al. 2010).

In our baseline specification, we match treated and control group workers on pre-treatment annual real wage income, separately by sector and calendar year. While the match is exact for calendar year and sector, we use coarsened exact matching (CEM, see Iacus et al. 2012; Blackwell et al. 2009) for pre-treatment income. To this end, we construct separate strata for each 10 percentiles of real wage income, as well as separate bins for the 99th and 99.9th percentiles, and a bin for zero income for recent hires, in each of the three pre-treatment years $t = -3, -2, -1$. We then match treated workers to control group workers for each of these income bins, while additionally requiring them to be observed in the same calendar year, and work in the same sector one year prior to treatment. We include calendar year and sector matching to ensure we are not capturing sector-specific business cycle effects, or other unobserved time-varying shocks affecting workers based on their original sector of employment. As such, each treated worker is matched to a set of controls from the same calendar and sector and belongs to the same pre-treatment earnings percentile bin. For recent hires, we additionally require them to experience similar trends in non-employment duration before treatment¹³ – for incumbents, this dimension is less relevant since incumbents are employed in the pre-treatment period by definition. This procedure results in 30,247 strata for incumbents and 82,942 strata for recent hires, and in doing so can match 98 percent of treated incumbents (using 93 percent of control group incumbents) and 95 percent of treated recent hires (using 65 percent of control group recent hires).¹⁴ While

¹³In particular, we estimate a linear trend in non-employment duration for individual recent hires before treatment, and match treat and control group recent hires using four bins of this trend: up to the 10th percentile, the 10th percentile to the median, the median to the 90th percentile and higher than the 90th percentile.

¹⁴Note that some strata may not contain any treated workers, in which case they are irrelevant for estimation.

we could refine the matching procedure further by increasing the number of strata, this comes at the cost of external validity, as a lower share of treated individuals would be matched.¹⁵ After matching, our sample contains 1,165,244 unique incumbent workers and 473,190 unique recent hires: given our observation window of 8 years ($t = -3$ through $t = 4$) this results in 9,321,952 observations for incumbents and 3,785,520 for recent hires.

This matching procedure ensures that outcome variables in most of the analyses below exhibit no significant pre-trends from years $t = -3$ through $t = -1$. Further support for the parallel trends assumption comes from Table 10 in Appendix A.1. This table compares observables across the treatment and control groups using matching weights. The two groups are closely matched on a wide range of variables for both firms and workers. In Appendix A.3 we check our results by additionally matching on pre-treatment employment growth at the firm level: this produces similar results, supporting the robustness of our identification.

More generally, Table 9 in Appendix A.1 shows summary statistics for incumbent and recently hired workers, showing averages and standard deviations across the balanced panel of workers and years. On average, recent hires earn lower wages and spend a higher number of days in non-employment compared to incumbents (76 versus 29 days annually). They also have higher benefit receipts, and are more likely to be female, and younger. Compared to incumbents, recent hires are overrepresented in larger firms, and are most commonly employed in firms in administrative and support activities, whereas incumbent workers are most often observed in the manufacturing sector.

Our estimation sample for identifying treated and control group workers contains 6,915 unique firms, all of which experience an automation spike at some point over the period. Workers employed at 2,753 of these firms are treated (that is, experience an automation spike early in the observation window), and workers at 5,301 firms serve as controls (that is, are experiencing an automation spike later). Workers at 1,139 firms are treated while its workers also serve as controls for workers at another treated firm either earlier or later in the sample period. In total, we observe 1,165,244 unique incumbent workers (115,698 of whom are treated) and 473,190 unique recent hires (94,757 of whom are treated)¹⁶ – all of these workers are employed at a firm in $t = -1$ that spikes at some point over 2000-2016. Note that incumbents and recent hires from both treated and control groups may move to never-spiking firms after $t = -1$.

¹⁵Further refinement of these strata does not change our results, although it leads to a smaller percentage of treated workers matched.

¹⁶This corresponds to around 20% of the total number of employed persons in the Netherlands over this period.

4 Impacts on wage income

In this section, we consider how workers' wage income evolves as the result of an automation event. We consider impacts on both incumbent workers and recent hires. Subsequent sections will decompose these income impacts into employment (section 5) and wage (section 6) effects, and consider to what extent income impacts are offset by various benefit payments (section 7).

Figure 5 shows the impact of automation events on annual real wage income in euros, separately for incumbent workers (3+ years of firm tenure) and recent hires. We scale each individual worker's wage income by their real wage income level one year before the automation spike, to obtain relative impacts.¹⁷ That is, it shows estimates of equation 2, where the outcome variable is annual real wage income (in 2010 euros) over $t = -1$ real income. Firstly, there are only negligible pre-trends in wage earnings for either incumbents and recent hires.

The estimates for incumbents, shown in black, highlight that these workers lose income as the result of an automation event. Indeed, in the automation year, incumbent workers lose some 0.7 percent in wage income and this effect increases over time, cumulating to 8.2 percent in total after five years. Estimates are statistically significant for all treatment and post-treatment years. Given that annual earnings grow by 1.6 percent annually on average, this reflects a non-negligible loss compared to usual earnings trajectories. In levels, this corresponds to approximately 257 euros lost for the average worker in the treatment year (0.7 percent of the average pre-treatment income of 36,706 euros, see Table 10); and 3,010 euros after 5 years in total ($0.082 \times 36,706$). This is evidence in favor of automation leading to labor-displacement for workers: compared to workers employed at firms who automate later, workers employed at currently automating firms experience income losses. The accumulation of effects over time is consistent with incumbent workers likely having open-ended contracts, making it costly to fire them. These gradual changes could in part also result from a time delay in the effective implementation of automation technologies relative to the cost outlay, or because it takes time for workers and firms to learn about changes to their match quality under the new technology.

For recent hires, shown in gray in Figure 5, we find no such income losses. Relative to recent hires in the control group, these workers actually have a 4.4 percent higher income in total over five years following an automation spike (corresponding to around 959 euros from a pre-treatment average income of 21,791 euros, see Table 10). However, this five-year effect

¹⁷This is preferable to log income impacts since this would eliminate zeros: this approach is also taken in e.g. Autor et al. (2014). In appendix A.2 we additionally show impacts in levels.

is only just statistically significant, and no significant impact on wage income is found in the automation year. This could be the case because recent hires have built up less firm-specific human capital, and therefore are more able to adapt to new job tasks either within the same firm or when moving to a new employer. This is consistent with Carneiro et al. (2015); Lefranc (2003), who finds that income losses following displacement result mostly from accumulated firm-specific human capital. Note that it may also be the case that recent hires do not lose income because these workers are in part hired in anticipation of the automation event – in this case their outcomes are endogenous to the event. However, as we will show below, recent hires also experience increased separation from their firms, suggesting endogeneity cannot account for the entire observed difference with incumbents.


All in all, the results in Figure 5 highlight that automation negatively affects the wage incomes of workers with longer firm tenures employed at the automating firm. To further scale these costs, we calculate the incumbent worker income losses resulting from one hundred euros spent on additional automation during an automation event. For incumbent workers on average, 100 euros spent on automation during an automation event results in the loss of 0.5 euros of income for the firm’s incumbent workers in that same year, and 6 euros after five years.¹⁸ The same calculation for recent hires reveals that 100 euros spent during an automation spike leads to an income gain of 0.6 euros of income in total after five years – this effect is an order of magnitude smaller than the 5-year loss for incumbents both because recent hires experience smaller income gains from automation than incumbents’ losses, and because the former on average are employed in firms with larger automation spike sizes.¹⁹

In Appendix A.3, we show that these results are robust to a number of alternative model specifications. In particular, we remove outlier firms in terms of employment changes (those experiencing an employment change exceeding 90 percent in any one year); and eliminate firms that experience a range of firm-level events (such as take-overs, mergers, and acquisitions; restructuring; and firm splits) either in the pre-treatment period or over the entire event estimation window. This alleviates the concern that unobserved firm-level events correlated with automa-

¹⁸For the average incumbent worker, firm-level automation costs during a spike year are 2.12 millions of euros. Since there are 2,753 treated firms, total automation event expenditures in our sample amount to 5,836 millions of euros. At the same time, 257 euros are lost per incumbent worker in the year of the automation event on average (3,010 euros in total over five years)– and the total number of (matched) unique incumbent workers in these firms is 115,698. Hence, 30 million euros of labor income is lost to incumbent workers in the year of the automation event (348 million euros over five years).

¹⁹There are 94,757 (matched) unique recent hires in treated firms, and the automation cost spike they experience on average is 5.36 million euros.

Figure 5: Annual real wage income, relative to $t = -1$



Notes: N=9,321,952 for incumbents; 3,785,520 for recent hires. Whiskers represent 95 percent confidence intervals.

tion spikes are driving our results. Further, we additionally match workers across firms with similar pre-treatment employment growth trajectories: this eliminates the concern that while individual treated and control group workers are on similar earnings and employment trajectories, firms need not be. Throughout, we find substantial income losses for incumbents, and much less evidence of such effects for recent hires (with estimates either positive but small, or statistically insignificant).

We now consider to what extent workers are effectively displaced from their firms, and then analyze to what extent the average annual income effects found here are due to changes in annual days worked after firm separation, or from changes in daily wage trajectories conditional on employment.

5 Impacts on firm separation and employment

The annual wage impacts found above may be driven by changes in days worked, changes in daily wages conditional on being employed, or a combination of both. Here we consider the first of these adjustment margins: do workers separate from their firms as a result of automation,

and does this lead to non-employment spells?²⁰

To answer this question, we first consider displacement effects in their most literal sense: does automation result in pre-existing workers leaving the firm? Figure 6 considers the impact of automation on firm separation by presenting estimates from equation 2, where the dependent variable is a dummy for the worker separating from the pre-treatment employer. All coefficients have been multiplied by 100, such that the effects are in percentage points. As before, we estimate the model separately for incumbents and recent hires. This highlights that workers' probability of separating from their employer following an automation spike is rising over time compared to control group workers, for both incumbents and recent hires.

Indeed, in the automation year, the separation probability for incumbent workers is 2.1 percentage points higher, where the (matched) control group incumbent separation probability is 13 percent, such that the automation-year effect corresponds to a 16 percent rise in firm separation for these workers. Cumulatively after five years, incumbents have a 8.6 percentage point higher chance of firm separation: this is a 24 percent rise compared to the average five-year (cumulative) chance of firm separation among control group workers of 36 percent).

Recent hires experience a 4.1 percentage point increase in the chance of firm separation in the year of the automation event. Although higher in absolute terms compared to the effect for incumbents, this is lower relative to the average separation chance of 35 percent among the (matched) control group for that year. Further, the cumulative treatment effect on separation probability for recent hires grows somewhat less over time – these different dynamics compared to incumbents could be the result of a higher prevalence of temporary contracts, allowing recent hires to be laid off more quickly. In total after five years, recent hires have a 7.2 percentage point higher chance to have separated from the automating firm, an increase of 10.3 percent over the 70 percent separation probability for the control group over that same period.

All in all, this suggests both incumbents and recent hires are more likely to separate from their employer as a result of automation. For incumbents, the resulting increase in separation is substantially higher relative to what they experience in the absence of an automation event – this effect is less sizable for recent hires, who have a higher separation probability at baseline. Although this increase in firm separation implies automation leads to displacement for the firm's

²⁰Any adjustment in days worked can in principle come from either the intensive or extensive margin: that is, workers may work fewer days with their current employer, or separate from their employer and experience a non-employment spell before finding re-employment. However, we do not find any evidence of intensive margin changes in non-employment – as such, any change found here reflects adjustments along the extensive margin.

pre-existing workers, it need not translate to income losses if these workers find re-employment quickly (and at similar wage rates): we now first consider impacts on non-employment.


Results are shown in Figure 7, where we define the dependent variable in equation 2 as the annual number of days spent in non-employment. Note that incumbents are by definition employed in the three years prior to the automation event – although they need not work full-time, their number of days in non-employment does not evolve much prior to $t = 0$. Starting in the year in which the automation spike takes place, however, their days worked gradually decreases. In particular, non-employment increases by 1.2 day in the automation event year, and this increases to around 5.2 days annually four years after the automation event, with a total cumulative increase in non-employment of 15.7 days compared to the control group. In the event year, matched control group incumbents spend around 22 days in non-employment on average, suggesting automation produces an increase of 5.4 percent in non-employment in the automation year itself. The cumulative five-year impact corresponds to 71 percent of this average treatment-year non-employment duration; and a 9.1 percent increase relative to the five-year cumulative non-employment duration (172 days) experienced by control group incumbents.

Unlike incumbents, new hires do not experience any statistically significant changes in non-employment as a result of automation, despite their increased firm separation rates. Indeed, the cumulated five-year impact is only an increase of 1.2 days in non-employment, which besides not being statistically significant is a negligible number compared to the control group’s cumulated five-year non-employment duration of 357 days.

These findings suggest that while both incumbents and recent hires separate from the firm at a higher rate as a result of automation, the former experience a non-negligible increase in non-employment whereas the latter move on to new jobs quite seamlessly. These different adjustment costs to displacement may be the result of incumbent workers having accumulated more (firm-specific) human capital which is depreciating in the face of automation. Also note that besides having higher firm-tenure, incumbents are on average around 5 years older than recent hires (see Table 10 in the Appendix). A further explanation could therefore be that compared to recent hires, incumbents had achieved a better firm-match prior to the automation event. If recent hires were not that well-matched with the firm in the first place (as also indicated by their higher baseline leave probability, and consistent with evidence in Fredriksson et al. (2018)), finding a similar or better match after separation may be easier.²¹

²¹This hypothesis may be investigated more directly using the methodology pioneered by Abowd et al. (1999)


Figure 6: Firm separation probability



Notes: N=9,321,952 for incumbents; 3,785,520 for recent hires. Whiskers represent 95 percent confidence intervals. Effects at $t = -2$ and $t = -3$ are zero by definition since incumbents are at the firm for three years before $t = 0$, and recent hires are in the sample conditional on being observed in the firm in $t = -1$.


in estimating worker-firm match quality.

Figure 7: Annual number of days in non-employment



Notes: N=9,321,952 for incumbents; 3,785,520 for recent hires. Whiskers represent 95 percent confidence intervals.

Figure 8: Log daily wages conditional on employment



Notes: N=9,321,952 for incumbents; 3,785,520 for recent hires. Whiskers represent 95 percent confidence intervals.

6 Impacts on daily wages


We now turn to the wage impacts of automation. In Figure 8, we consider the effect of automation on log daily wages (conditional on employment), for both incumbent workers and recent hires. Recall that we do not observe daily hours worked in our data: changes in daily wages can therefore result from changes in hourly wages and/or changes in daily hours worked. We find statistically significant but economically very small negative daily wage losses of at most 0.8 percent four years after the automation event for incumbents and 1.4 percent for recent hires.²² The absence of substantial daily wage effects implies that the income losses incumbent workers experience are almost entirely accounted for by firm separation followed by non-employment spells.

Of course, these wage effects combine effects across job leavers and job stayers, which may cancel out on average – we therefore also estimate our daily wage models separately for these two groups, relative to control group workers of job leavers and job stayers, respectively.

Figures 9 and 10 consider how daily wages evolve, in log points, for treated versus control

²²Together with the effect on days in non-employment, these estimates on daily wages provide a decomposition of the total wage income effect shown in section 4.


Figure 9: Log daily wages conditional on employment, firm leavers



Notes: N=2,591,038 for incumbents; 2,197,693 for recent hires. Whiskers represent 95 percent confidence intervals.

group workers who leave or stay at their firm, respectively. This reveals only very limited heterogeneity across these two groups: we find very small negative wage impacts from automation for both incumbents and recent hires who leave their firm, and no statistically significant effects for those who remain. This suggests automation does not lead to substantial wage losses for these workers.

Figure 10: Log daily wages conditional on employment, firm stayers



Notes: N=6,410,742 for incumbents; 1,115,624 for recent hires. Whiskers represent 95 percent confidence intervals.

7 Impacts on alternative income sources


So far, we have shown that incumbent workers experience income losses as a result of automation in their firm, cumulating to around 3,010 euros on average per incumbent worker over the five-year post-treatment window. These losses are almost entirely driven by higher firm-separation probabilities accompanied by non-employment spells, raising the question to what extent various benefit systems are absorbing these impacts. Figure 11 considers total benefit receipts (in annual real 2010 euros), comprised of unemployment benefits, disability benefits, and welfare payments, as well as the separate contributions from these three different sources.²³

We find that incumbent workers do receive additional benefit income following an automation event: however, after five years, the cumulative amount is around 446 euros on average, implying that only 15 percent of the negative wage income impact is offset.²⁴ This finding is comparable to that in other worker displacement events, where typically only a small part of the negative impact is compensated by social security (Hardoy and Schøne 2014). Figure 11 further shows

²³Since we do not find any non-employment increases for recent hires, we only consider these impacts for incumbent workers.

²⁴We find similar percentages of income loss offset when using different spike threshold cutoffs and/or longer observation windows.

Figure 11: Annual real benefit income



Notes: N=9,321,952 (incumbents only). Benefits are the sum of received unemployment and disability benefits, and welfare. Whiskers represent 95 percent confidence intervals.


that all of the benefit payments for incumbents arise from unemployment insurance: this is expected, as unemployment benefit eligibility is very high among workers with at least three years of firm tenure.²⁵ Consistent with high unemployment benefit eligibility, we do not see any contribution from welfare payments for incumbents. Lastly, disability benefits²⁶ are actually slightly decreasing over time. Since all incumbents were previously employed, this implies some were working part-time and receiving benefits for a partial disability prior to the automation event. Further analysis shows that the decline in disability insurance receipts is driven by incumbents who find re-employment – that is, they no longer receive these benefits with their new employer.

Besides benefits and welfare payments, displaced workers could also adjust by entering self-employment: since self-employment income is not observed in our data, we may be overestimating the income losses workers experience. Indeed, Figure 12 shows that treated incumbents are somewhat more likely to enter self-employment following an automation event, although

²⁵For most of the observation period, eligible workers in the Netherlands are entitled to up to 38 weeks of unemployment benefits following job loss.

²⁶Disability benefits in the Netherlands cover impairment whether full or partial, and whether temporary or permanent, and replace up to 70-75 percent of workers' past wages. Benefits are financed by employers without worker contributions, and there is a long history of the use of these schemes in the Netherlands as hidden unemployment (Koning and Lindeboom 2015).

Figure 12: Probability of entering self-employment or early retirement



Notes: N=9,321,952 (incumbents only). Whiskers represent 95 percent confidence intervals.

the effect is very small indeed – 0.3 percentage points cumulated over five years (relative to a five-year probability of 5.2 percent among the control group), suggesting this cannot be an important compensating income source. Lastly, we find evidence that automation also leads to impacts on early retirement (defined as the receipt of retirement benefits prior to reaching the legal retirement age of 65), as shown in Figure 12. These effects are larger: four years after the automation event, treated incumbent workers are 0.4 percentage points more likely to be observed in early retirement, whereas the average five-year probability of early retirement among control-group incumbents is around 1.7 percent. As such, the treatment effect represents a 24 percent increase in the incidence of early retirement.

Taken together, these results show that the documented income losses experienced by automation-impacted workers are only partially offset by benefit systems, implying automation-affected workers largely bear the cost themselves.

8 Effect heterogeneity

So far, we have shown income losses for incumbent workers driven by non-employment spells, only partially compensated by benefit payments. For recent hires, on the other hand, automation

does not lead to substantial income losses. Here, we investigate further how different types of workers are affected: in particular, we consider workers by age, gender, and wage quartile. In Appendix A.2 we additionally consider differences by workers' initial sector of employment and firm size, as well as two alternative skill measures.

For each of the groups considered here, we contrast the effect against the same group at the control firm by using an interaction term – this results in a decomposition of the mean effects found in previous sections. In particular, we estimate the following model:

$$y_{ijt} = \alpha + \beta \text{treat}_i + \gamma \text{post}_{it} + \delta_0 \times \text{treat}_i \times \text{post}_{it} + \sum_k [\delta_k \times \text{treat}_i \times \text{post}_{it} \times z_{ki}] + \lambda X_{ijt} + \varepsilon_{ijt}, \quad (3)$$

where, as before, i indexes workers, j firms, and t time relative to the automation spike. For succinctness, we estimate the average annual effect over the entire post-treatment period rather than reporting the year-by-year coefficients. As such, post_{it} is a dummy variable indicating the post-treatment period (i.e. $t \geq 0$). Further, z_{ki} is a dimension of worker heterogeneity, such as gender, age in the year before automation, or age-specific wage rank, containing k categories. In addition to the controls included in equation 2, X_{ijt} also contains z_{ki} as well as the interaction terms $z_{ki} \times \text{treat}_i$ and $z_{ki} \times \text{post}_{it}$. In equation 3, δ_0 gives the estimated treatment effect for the reference group, and δ_k the deviation from that effect for category k of worker characteristic z_i .

Throughout this section, we will show effects on relative earnings, firm separation, non-employment duration, and log daily wages (conditional on employment). While the impacts on relative earnings are a good summary measure, any given income loss may come about through different channels (e.g. wage changes versus non-employment changes) for different workers. Lastly, we also show results for the probability of early retirement.

8.1 Worker age

We first consider how automation affects workers of different age groups: those younger than 30 years old, between the ages of 30 and 39, 40 and 49, and 50 years and older: results are shown in Table 5. In this table, coefficients are shown for the reference category, workers below the age of 30, and the remaining estimates represent deviations from this reference group. We consider impacts on relative earnings, the firm separation probability, non-employment duration, log daily wages (conditional on employment), and the probability of early retirement – coefficients in columns 1, 2, 4, and 5 have been multiplied by 100 to represent percentage point impacts for

relative earnings, firm separation, and early retirement, and log point impacts for daily wages. Panel A presents estimates for incumbent workers, and panel B for recent hires.

The oldest incumbent workers experience the largest income losses: 2.3 percent annually on average over 5 years. Qualitatively, we see a similar pattern among recent hires, with the difference that these losses (of -0.9 percent annually) are not statistically different from zero.

However, the patterns of displacement from the firm by age category do differ between incumbents and recent hires. In particular, for recent hires, the probability of displacement increases monotonically with age, with the oldest recent hires 6.7 percentage points more likely to leave the firm than the youngest workers. This suggests younger recent hires may be better able to deal with any automation-driven workplace changes. Among incumbents, on the other hand, the youngest workers are most likely to leave the firm, but differences with older workers are smaller (these differences range between 0.5 to 1.6 percentage points and are not always statistically significant). This may be related to “last-in first-out” policies among high-tenure workers, but also possibly reflects younger workers’ higher flexibility in terms of job transitions.

These age-dependent costs of job transitions are also reflected in non-employment durations, which are monotonically increasing in age for incumbents, despite older incumbents not having a higher separation probability. Indeed, the youngest incumbents do not experience any statistically significant increase in non-employment duration, although workers of all other age groups do. Among recent hires, we do not find any statistically significant increases in non-employment duration for any group, even for older workers.

Log daily wage effects (which are only defined for workers with non-zero wages, i.e. those in employment) are economically very small for incumbents of all ages. For recent hires, however, we do find one group which loses a somewhat larger amount: workers over the age of 50 have 2.3% lower wages annually. These effects are not visible in the average (reported above) since older workers are only a small fraction of the group of recent hires, and because the youngest workers gain 1.1% annually.

Lastly, and unsurprisingly, the early retirement effects found above are entirely driven by workers over the age of 50, particularly incumbents.

Taken together, these results show that older workers generally bear more of the cost of automation-driven displacement: they either experience longer non-employment durations or enter early retirement (as is the case for incumbents), or have higher firm separation probabilities and lower daily wages (as is the case for recent hires).

Table 5: Effect heterogeneity by worker age

	(1)	(2)	(3)	(4)	(5)
	Relative wage income	Firm separation probability	Non-employm. duration	Log daily wage	Early retirem. probability
Panel A. Incumbents					
Age <30 (ref)	-0.73 (0.56)	5.79*** (0.37)	0.66 (0.53)	-0.64** (0.23)	-0.00 (0.01)
<i>Deviations from reference group for:</i>					
Age 30-39	-1.19 (0.62)	-1.58*** (0.44)	2.23*** (0.63)	-0.17 (0.27)	0.01 (0.01)
Age 40-49	-0.83 (0.61)	-0.46 (0.43)	2.60*** (0.62)	0.60* (0.27)	0.01 (0.01)
Age 50+	-1.56* (0.64)	-1.02* (0.44)	4.33*** (0.72)	0.65* (0.29)	0.83*** (0.11)
Observations	9,319,736	9,319,736	9,319,736	8,999,617	9,319,736
Panel B. Recent hires					
Age <30 (ref)	2.26* (0.92)	3.08*** (0.30)	0.63 (0.71)	1.05*** (0.31)	0.01 (0.01)
<i>Deviations from reference group for:</i>					
Age 30-39	-2.24 (1.29)	1.96*** (0.44)	1.03 (1.02)	-1.77*** (0.44)	-0.02 (0.02)
Age 40-49	1.12 (1.48)	5.19*** (0.50)	-0.82 (1.16)	-1.12* (0.49)	-0.06 (0.03)
Age 50+	-3.13 (1.91)	6.74*** (0.64)	1.11 (1.52)	-3.37*** (0.71)	0.49** (0.18)
Observations	3,785,520	3,785,520	3,785,520	3,313,317	3,785,520

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Workers under the age of 30 are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

8.2 Worker gender

There are only modest differences in how male and female workers are affected by automation events at their firm, as shown in Table 6. Indeed, there are virtually no gender differences in most outcomes – wage earnings, firm separation rates, and wages conditional on employment – for either incumbents or recent hires. Female workers do experience a somewhat lower non-employment duration than male workers, for both incumbents and recent hires – but these differences are not statistically significant. The only margin where gender differences are detected is early retirement: virtually the entire observed increase in early retirement is driven by male workers, both among incumbents and among recent hires.

Table 6: Effect heterogeneity by worker gender

	(1) Relative wage income	(2) Firm separation probability	(3) Non-employm. duration	(4) Log daily wage	(5) Early retirem. probability
Panel A. Incumbents					
Male (ref)	-1.69*** (0.17)	4.92*** (0.15)	3.36*** (0.23)	-0.37*** (0.08)	0.32*** (0.03)
<i>Deviations from reference group for:</i>					
Female	0.11 (0.37)	-0.02 (0.29)	-0.56 (0.45)	0.12 (0.20)	-0.27*** (0.05)
Observations	9,319,736	9,319,736	9,319,736	8,999,617	9,319,736
Panel B. Recent hires					
Male (ref)	0.94 (0.66)	5.76*** (0.24)	1.81*** (0.54)	-0.03 (0.21)	0.10** (0.03)
<i>Deviations from reference group for:</i>					
Female	0.07 (1.11)	-0.64 (0.38)	-1.57 (0.87)	-0.52 (0.39)	-0.11* (0.05)
Observations	3,785,520	3,785,520	3,785,520	3,313,317	3,785,520

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Male workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

8.3 Worker wage quartile

Since we only observe education levels for a small and selected subsample of workers, we obtain a measure of workers' skill level by calculating each worker's wage rank by age in $t = -1$, calculated separately for incumbents and recent hires. We then group workers into quartiles based on this rank. For example, the highest-quartile workers in this measure are those who earn in the top 25 percent of earnings across the sample for workers of their age in the year before the automation event.²⁷

Results are reported in Table 7: workers in the lowest residual wage quartile are used as the reference category throughout. This shows that for incumbents, income losses are increasing monotonically in wage quartiles: the highest-paid workers experience the largest losses, percentage-wise. Among recent hires, workers in the bottom quartile experience the strongest increase in income, with no statistically significant effect for workers in higher wage quartiles. This shows that, across the economy, automation is not biased against lower-wage workers.

Indeed, the firm separation probability is *increasing* in wage quartile for both incumbents

²⁷As an alternative skill measure, we additionally calculate wage quartiles based on a residual wage regression: results (reported in Table 11 in the Appendix) are very similar.

and recent hires. This suggests that “lower-skilled” workers are actually less likely to be affected by automation events compared to “higher-skilled” workers, where skill is defined by age-specific wage ranks as described above. However, despite a higher change of firm separation, these higher-skilled incumbent workers do not experience statistically significantly higher average increases in non-employment duration, implying they are able to find re-employment more quickly. For recent hires, we do find some higher non-employment duration increases for the higher-skilled workers, matching their higher separation rates. However, non-employment duration effects are lower for recent hires as a group compared to incumbents.

Similar to the average treatment effect, we do not see much action along the daily wage margin conditional on employment when breaking out effects by age-specific wage quartile. However, higher-quartile workers do experience slightly stronger daily wage decreases than do lower-quartile workers. Lastly, early retirement take-up occurs for incumbents across all but the lowest wage quartile, and for recent hires for the top quartile workers, only.

We should be careful about drawing strong conclusions from these results since they may be capturing other factors than pure worker skill, such as the quality of the worker-firm match. However, Table 7 does provide an important insight that is counter to a common intuition: workers lower down the wage rank for their age (“lower-skilled” workers) are not displaced more often by automation events: higher-skilled incumbents are actually more likely to be affected and experience displacement. These findings are broadly confirmed when we estimate our models for the subset of data where education level data is available, as shown in Table 13 in the Appendix.

These higher losses for higher-skilled workers may of course be the result of differences in the firms where automation spikes occur: specifically, if larger firms both have bigger automation spikes and also pay higher wages. While this matters for the average worker’s exposure to displacement from automation, we are also interested in which workers are displaced within firms. Therefore, Table 7 reports estimates by workers’ age-specific *within-firm* wage quartile. That is, the bottom quartile reflects incumbents (panel A) or recent hires (panel B) who are in the lowest 25% of their firm’s wage distribution for their age.²⁸

When considering workers by their within-firm wage rank, automation is remarkably neutral across the “skill” distribution for incumbents: there are no statistically significant differences in wage losses across the four quartiles, and these workers have very similar probabilities of

²⁸Note that these quartiles cannot be calculated for the smallest firms: however, all our previous findings are very similar in this subsample, suggesting that this is not driving the results.

Table 7: Effect heterogeneity by age-specific wage quartile

	(1) Relative wage income	(2) Firm separation probability	(3) Non-employm. duration	(4) Log daily wage	(5) Early retirem. probability
Panel A. Incumbents					
Bottom quartile (ref)	-0.94* (0.39)	4.25*** (0.27)	2.75*** (0.43)	0.31 (0.18)	0.05 (0.04)
<i>Deviations from reference group for:</i>					
Second quartile	-0.63 (0.47)	0.99** (0.36)	0.63 (0.57)	-0.68** (0.21)	0.27*** (0.07)
Third quartile	-0.78 (0.47)	0.97** (0.36)	0.63 (0.57)	-0.48* (0.22)	0.25*** (0.07)
Top quartile	-1.76*** (0.49)	1.33*** (0.37)	1.04 (0.59)	-1.00*** (0.24)	0.22** (0.07)
Observations	9,321,952	9,321,952	9,321,952	9,001,780	9,321,952
Panel B. Recent hires					
Bottom quartile (ref)	5.20*** (1.31)	2.33*** (0.33)	-2.24** (0.87)	0.58 (0.44)	-0.01 (0.05)
<i>Deviations from reference group for:</i>					
Second quartile	-4.93** (1.64)	2.29*** (0.48)	3.80** (1.22)	-0.51 (0.55)	0.03 (0.06)
Third quartile	-4.46** (1.56)	7.60*** (0.52)	4.20*** (1.20)	-0.23 (0.53)	0.08 (0.06)
Top quartile	-5.71*** (1.55)	5.66*** (0.52)	4.19*** (1.14)	-1.46** (0.52)	0.17* (0.08)
Observations	3,785,520	3,785,520	3,785,520	3,313,317	3,785,520

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Quartile 1 workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

leaving the firm. The adjustment mechanisms do differ slightly, with the higher-paid workers displaced from the firm experiencing larger daily wage losses but lower non-employment increases compared to workers displaced from firms' lower wage ranks. Also for recent hires, displacement effects are neutral across the "skill" groups. However, we see a polarized pattern of wage impacts, with the bottom and top quartiles experiencing income gains, and no impacts for the middle two quartiles.

All in all, the results in this section show that the average income losses found earlier are borne disproportionately by older workers, and workers in higher wage quartiles across the economy. Within firms, however, automation displaces workers from all ranks of the "skill" distribution at similar rates. Lastly, early retirement effects are entirely driven by (older) males.

In Appendix A.2 we also consider how worker displacement effects from automation differ by the worker's firm size and sector. Here, the main takeaway is that although there are some

Table 8: Effect heterogeneity by age-specific within-firm wage quartile

	(1)	(2)	(3)	(4)	(5)
	Relative wage income	Firm separation probability	Non-employm. duration	Log daily wage	Early retirem. probability
Panel A. Incumbents					
Bottom quartile (ref)	-1.86*** (0.40)	5.33*** (0.27)	4.31*** (0.44)	-0.01 (0.18)	0.24*** (0.06)
<i>Deviations from reference group for:</i>					
Second quartile	0.14 (0.50)	-0.99** (0.37)	-1.49* (0.60)	-0.53* (0.23)	0.05 (0.08)
Third quartile	0.17 (0.48)	-0.51 (0.37)	-1.14 (0.59)	-0.44 (0.23)	0.09 (0.08)
Top quartile	0.40 (0.50)	0.11 (0.38)	-2.13*** (0.60)	-0.70** (0.25)	-0.11 (0.08)
Observations	9,205,656	9,205,656	9,205,656	8,890,803	9,205,656
Panel B. Recent hires					
Bottom quartile (ref)	3.95** (1.33)	6.54*** (0.36)	-1.82* (0.85)	-0.50 (0.43)	0.04 (0.05)
<i>Deviations from reference group for:</i>					
Second quartile	-4.27* (1.70)	-0.00 (0.52)	3.15* (1.23)	-0.05 (0.56)	-0.03 (0.07)
Third quartile	-3.40* (1.64)	-0.75 (0.52)	3.87** (1.22)	0.47 (0.54)	-0.03 (0.06)
Top quartile	-1.70 (1.65)	0.52 (0.53)	2.78* (1.21)	1.25* (0.55)	0.10 (0.08)
Observations	3,664,008	3,664,008	3,664,008	3,203,250	3,664,008

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Quartile 1 workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

differences in effect size, our findings for incumbents are not driven by workers in a small subset of sectors or firm sizes. Income losses are larger for incumbents and recent hires displaced from smaller firms, suggesting we would find larger losses if our data were more representative in terms of firm size. Further, while incumbents initially employed in manufacturing sectors do experience some of the larger firm separation and non-employment increases, the automation-driven displacement effects found for incumbents are pervasive across many different sectors of the economy. Lastly, overall differences found between incumbents and recent hires are not only due to their different sectoral or firm size affiliations: manufacturing workers and workers at the largest firms are displaced from their firms at similar rates whether they are incumbents or recent hires, but only incumbents experience increases in non-employment and early retirement.

9 Impacts conditional on firm separation

The above sections document the effects of automation on wage income, finding losses for incumbent workers at automating firms. These estimates are calculated for all incumbent workers and all recently hired workers who are at the firm the year prior to the major automation event. However, most of the earnings losses are likely borne just by those workers who leave the firm.


In this section, we consider the worker-level income loss from automation *conditional on* separating from the firm by comparing the evolution of wage income for workers who leave the automating firm (“leavers”) to that of all control group workers (i.e. both leavers and stayers). This mimics the empirical specification from the displaced worker literature, where earnings trajectories for workers displaced from firms with annual employment losses exceeding a certain threshold (usually, 30 percent) are compared to those from control group firms which did not experience such a displacement event.²⁹

Unsurprisingly, relative losses in wage income are larger among the group of firm leavers than for the entire group of treated workers: Figure 13 indicates incumbents lose 4.7 percent in the year of the automation event, and these losses increase to 17.2 percent after five years (and 57.4 percent cumulatively). Recent hires also lose income, but much less: 0.9 percent in the automation year and 3.6 percent four years after the event (13.5 percent cumulatively).

As before, these earnings losses are largely driven by non-employment. Incumbent workers who leave their firm work around 11 fewer days during the year of the automation event and 44 fewer days after 5 years, as seen in Figure 14 – the cumulative effect over 5 years is 149 days. Non-employment effects are again smaller for recently hired workers: 30 days cumulatively after 5 years. As above (section 6), effects on daily wages are small for both groups (not shown here).


²⁹However, note that our model also differs from the displaced worker literature in that *all* firms in our sample experience the event, and we exploit only its timing.

Figure 13: Annual real wage income, relative to $t = -1$ and conditional on firm separation



Notes: N=8,786,432 for incumbents; 3,647,448 for recent hires. Whiskers represent 95 percent confidence intervals.

Figure 14: Annual number of days in non-employment, conditional on firm separation



Notes: N=8,786,432 for incumbents; 3,647,448 for recent hires. Whiskers represent 95 percent confidence intervals.

10 Impact comparison to mass layoffs and plant closings

Our evidence confirms that some workers bear a considerable loss of earnings and work following major automation investments by their employers. But how significant are these losses for social policy? To gauge the overall magnitude of these losses for society, it helps to have a relevant metric to compare them to. A natural point of comparison is the displaced worker literature: not only is there a substantial literature on the costs to workers of mass layoffs and plant closings, but there is also a history of policy interventions aimed at countering these losses.

To compare the overall impact on the workforce in terms of displacement, we need to compare the shares of workers separating from their firm each year in these two different types of events. Firstly, Abbring et al. (2002) find that about 4 percent of workers in both the US and the Netherlands are displaced each year because of their employers' adverse economic conditions. For the US, the estimate is 4.4 percent; for the Netherlands, the estimates range from 3.5 to 7.2 percent. To construct an analogous number for automation events, we need to consider the probability that any one worker experiences the event in a given year ("event exposure"), and the probability of firm separation conditional on the event occurring ("event impact"). As for the former, in our sample, some 9 percent of workers are employed in firms that have an automation event in any year on average. This is an upper bound for the total workforce population since our sample is weighted toward larger firms who automate more frequently. As for event impact, the difference between our automation results and the worker displacement literature is that by comparison, far fewer workers separate from their firms each year as a result of automation at the affected firms. The mass layoff studies typically include workplaces where 30 percent or more of the employees are separated, and even larger portions of the workforce are laid off in studies of plant closings. In contrast, only 2 percent of incumbent workers leave their employer in the year of an automation event relative to the control group. Even after 5 years, automation only accounts for the separation of 8.6 percent of incumbent workers (7.2 percent for recent hires). Putting event exposure and impact together, the annual separation hazard from automation is around 0.77 percent (0.086×0.09), compared to 4 percent for mass lay-offs. As such, the overall risk that a worker will separate from their firm in a mass layoff or plant closing in any year is almost an order of magnitude greater than the risk of firm separation due to automation.

This discussion highlights another difference from findings on mass lay-offs and plant closings: the impact of automation occurs relatively slowly while the mass layoff literature reports an

immediate impact followed by a period of adjustment. The slow response may reveal something about the nature of automation. If automation were a simple story of machines replacing humans, then most of the job losses would occur the year of the automation event or, perhaps, the year after. The slow response suggests that automation appears to involve substantial adjustment costs, perhaps reflecting significant complementarities with workers. For example, slow adjustment might reflect a learning and selection process: over time, workers learn what work is like using the new technology and some choose to leave as a result; similarly, employers learn which workers have the skills to adapt to the new technology and they lay off some workers or encourage them to leave based on this new knowledge. Further research, including a consideration of workers hired after the automation event, and a comparison between leaving workers' pre- and post-treatment firms may shed some light on the changing nature of employer-employee matches.

The different time patterns affect comparisons of earnings losses. While workers separated from their firms by automation experience similar earnings losses to displaced workers in the fourth year after the event, their initial earnings losses are much lower. Incumbent workers who leave their employer after an automation event experience an average loss of earnings of about 17 percent in the fifth year after the event.³⁰ By comparison, Jacobson et al. (1993) find that long-tenured workers in Pennsylvania experienced a loss of earnings of 25 percent six years after a mass layoff. Couch and Placzek (2010) find losses of 13 to 15 percent six years after a mass layoff in another sample.³¹ However, Jacobson et al. (1993) report an average earnings loss of over 40 percent immediately following the layoff and Couch and Placzek (2010) report an immediate loss of 33 percent. The earnings loss during the first year of automation is less than 4 percent by comparison. Thus, the workers displaced from automation have somewhat smaller cumulative earnings losses than workers experiencing mass layoffs.

A final difference with the worker displacement literature is that automation seems to have significantly smaller effects on wages. The worker displacement literature reports differing results regarding the relative importance of wage declines and non-employment in explaining the observed loss of earnings (see discussion in Hijzen et al. 2010). However, studies that carefully

³⁰The loss of earnings for recently hired workers averages just 3.6 percent four years after the automation event. We cite the result for incumbent workers because most of the mass layoff literature studies impacts on long-tenured workers. Note also that the 17 percent figure might overstate the long-term loss of earnings. Because workers leave over time, the average losses in the fifth year reflect workers who have left their employer relatively recently; longer term earnings losses may be smaller.

³¹Couch and Placzek (2010) also review other studies of displaced workers that report long-term average earnings losses ranging from 7 to 16 percent.

measure wage effects do find significant long-term declines in wages, sometimes called “wage scarring”. In the Netherlands specifically, Mooi-Reci and Ganzeboom (2015) find wage declines following unemployment of 10 percent for men and 17 percent for women. Deelen et al. (2018) study the impact of firm closures and find immediate wage declines of 4 and 6 percent for prime-age and older workers, respectively; and 3 percent wage declines 6 years after the closure. In contrast, our study finds wages declines of at most 1 percent for both incumbent and recently hired workers who leave following automation, both immediately after the event and in subsequent years.³² Our evidence suggests that automation induces relatively little wage scarring, even for workers with relatively long firm tenure.³³

To summarize, compared to mass layoffs and plant closings, automation events cause far fewer firm separations, both at the affected firms individually and for the workforce overall. Further, these layoffs occur more gradually, perhaps facilitating adjustment. Indeed, there is little effect on daily wages, and cumulative earnings losses are somewhat smaller even conditional on displacement. Overall, the displacement impact of automation seems mild in comparison to the effects arising from changing economic conditions more generally.

11 Conclusion

Using evidence on firm-level automation expenditures across all non-financial private sectors in the Netherlands over 2000-2016, we provide the first estimate of the impacts of automation on individual workers. Using an event study differences-in-differences design, we find that automation at the firm increases the probability of workers separating from their employers. Incumbent workers are around 25 percent more likely to separate from their firm, followed by a decrease in days worked. This leads to a five-year cumulative wage income loss of about 8% of one year’s earnings for incumbent workers. We find little change in wage rates. Further, lost wage earnings are only partially offset by various benefits systems, and disproportionately borne by older workers – the rate of early retirement increases by 24% as a result of automation.

Further, we compare the effects of automation events to what happens when firms conduct mass layoffs or close down. A back-of-the-envelope calculation shows that far fewer workers are

³²This number compares workers who leave their employers in the treatment group to workers who leave their employers in the control group.

³³This could of course in part be because the workers affected by automation spikes are different from those affected by firm closures, in ways that allow them to better deal with job transitions. Indeed, to the extent that workers can anticipate firm closures more than the timing of automation events, workers affected by closures may be negatively selected.

displaced from their firms following automation events, both at the individual firms and in the workforce overall. Results suggest that the risk of displacement from automation is an order of magnitude smaller than the risk of losing work due to adverse economic conditions.

It should be noted that although we have shown displacement for automating firms' pre-existing workforce, this of course need not imply that jobs are destroyed on net in the economy or even at the firm as a result of automation. As a related macro literature has shown, there are various countervailing mechanisms which our specification does not inform on, including effects operating through firms' input-output linkages and changes in final demand. However, our results contribute to understanding the distributional impacts of automation across individual workers.

References

- Abbring, J., van den Berg, G., Gautier, P., van Lomwel, A., van Ours, J., and Ruhm, C. (2002). Displaced Workers in The United States and The Netherlands. In Kuhn, P., editor, *Losing Work, Moving on: International Perspectives on Worker Displacement*. W.E. Upjohn Institute.
- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.
- Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, 4:1043–1171.
- Acemoglu, D. and Restrepo, P. (2018a). Artificial Intelligence, Automation and Work. In Agrawal, A. K., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence*. University of Chicago Press.
- Acemoglu, D. and Restrepo, P. (2018b). Modeling automation. *AEA Papers and Proceedings*, 108:48–53.
- Acemoglu, D. and Restrepo, P. (2018c). Robots and Jobs: Evidence from US Labor Markets. MIT.
- Acemoglu, D. and Restrepo, P. (2018d). The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review*, 108(6):1488–1542.
- Autor, D. and Salomons, A. (2018). Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. *Brookings Papers on Economic Activity*, Spring.
- Autor, D. H., Dorn, D., Hanson, G. H., and Song, J. (2014). Trade Adjustment: Worker-Level Evidence *. *The Quarterly Journal of Economics*, 129(4):1799–1860.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content Of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar Extinction. *The Quarterly Journal of Economics*, 125(2):549–589.

- Benzell, S. G., Kotlikoff, L. J., LaGarda, G., and Sachs, J. D. (2016). Robots Are Us: Some Economics of Human Replacement. NBER Working Paper 20941, National Bureau of Economic Research, Inc.
- Bessen, J. (1999). Real Options and the Adoption of New Technologies. Online working paper at <http://www.researchoninnovation.org/>.
- Bessen, J. (2018). AI and Jobs: The Role of Demand. In *The Economics of Artificial Intelligence: An Agenda*, NBER Chapters. National Bureau of Economic Research, Inc.
- Blackwell, M., Iacus, S., King, G., and Porro, G. (2009). CEM: Coarsened Exact Matching in Stata. *Stata Journal*, 9(4):524–546.
- Carneiro, A., Portugal, P., and Raposo, P. (2015). Decomposing the Wage Losses of Displaced Workers: The Role of the Reallocation of Workers into Firms and Job Titles. IZA Discussion Papers 9220, Institute for the Study of Labor (IZA).
- Couch, K. A. and Placzek, D. W. (2010). Earnings Losses of Displaced Workers Revisited. *American Economic Review*, 100(1):572–89.
- Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2017). German Robots – The Impact of Industrial Robots on Workers. Technical report, Institute for Employment Research, Nuremberg, Germany.
- Deelen, A., de Graaf-Zijl, M., and van den Berge, W. (2018). Labour market effects of job displacement for prime-age and older workers. *IZA Journal of Labor Economics*, 7(1):3.
- Dinlersoz, E. and Wolf, Z. (2018). Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing. Working Papers 18-39, Center for Economic Studies, U.S. Census Bureau.
- Doms, M., Dunne, T., and Troske, K. R. (1997). Workers, Wages, and Technology. *The Quarterly Journal of Economics*, 112(1):253–290.
- Doms, M. E. and Dunne, T. (1998). Capital Adjustment Patterns in Manufacturing Plants. *Review of Economic Dynamics*, 1(2):409–429.

- Duggan, M., Garthwaite, C., and Goyal, A. (2016). The Market Impacts of Pharmaceutical Product Patents in Developing Countries: Evidence from India. *American Economic Review*, 106(1):99–135.
- Fadlon, I. and Nielsen, T. H. (2017). Family Health Behaviors. Working Paper 24042, National Bureau of Economic Research.
- Ford, M. (2015). *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.
- Fredriksson, P., Hensvik, L., and Skans, O. N. (2018). Mismatch of talent: Evidence on match quality, entry wages, and job mobility. *American Economic Review*, 108(11):3303–38.
- Frey, C. B. and Osborne, M. (2016). The Future of Employment: How Susceptible are Jobs to Computerisation? *Technological Forecasting and Social Change*, 114:254–280.
- Graetz, G. and Michaels, G. (2018). Robots at Work. *Review of Economics and Statistics*, forthcoming.
- Gregory, T., Salomons, A., and Zierahn, U. (2018). Racing With or Against the Machine? Evidence from Europe. CESifo Working Paper Series 7247, CESifo Group Munich.
- Haltiwanger, J., Cooper, R., and Power, L. (1999). Machine Replacement and the Business Cycle: Lumps and Bumps. *American Economic Review*, 89(4):921–946.
- Hardoy, I. and Schøne, P. (2014). Displacement and household adaptation: insured by the spouse or the state? *Journal of Population Economics*, 27(3):683–703.
- Hijzen, A., Upward, R., and Wright, P. (2010). The Income Losses of Displaced Workers. *Journal of Human Resources*, 45(1).
- Iacus, S. M., King, G., and Porro, G. (2012). Causal Inference Without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, pages 685–709.
- Koning, P. and Lindeboom, M. (2015). The Rise and Fall of Disability Insurance Enrollment in the Netherlands. *Journal of Economic Perspectives*, 29(2):151–72.

- Lafortune, J., Rothstein, J., and Schanzenbach, D. W. (2018). School Finance Reform and the Distribution of Student Achievement. *American Economic Journal: Applied Economics*, 10(2):1–26.
- Lefranc, A. (2003). Labor Market Dynamics and Wage Losses of Displaced Workers in France and the United States. *SSRN Electronic Journal*.
- Miller, C. (2017). The Persistent Effect of Temporary Affirmative Action. *American Economic Journal: Applied Economics*, 9(3):152–90.
- Mooi-Reci, I. and Ganzeboom, H. B. (2015). Unemployment Scarring by Gender: Human Capital Depreciation or Stigmatization? Longitudinal Evidence from the Netherlands, 1980–2000. *Social Science Research*, 52:642–658.
- Nilsen, O. A. and Schiantarelli, F. (2003). Zeros and Lumps in Investment: Empirical Evidence on Irreversibilities and Nonconvexities. *The Review of Economics and Statistics*, 85(4):1021–1037.
- Pindyck, R. (1991). Irreversibility, Uncertainty, and Investment. *Journal of Economic Literature*, 29(3):1110–48.
- Rothschild, M. (1971). On the cost of adjustment. *The Quarterly Journal of Economics*, 85(4):605–622.
- Ruhm, C. J. (1991). Displacement Induced Joblessness. *The Review of Economics and Statistics*, 73(3):517–522.
- Susskind, D. (2017). A Model of Technological Unemployment. Economics Series Working Papers 819, University of Oxford, Department of Economics.

A Appendix

This supplemental appendix contains additional summary statistics and results, as well as several robustness checks on our baseline specifications.


A.1 Additional summary statistics

Here we present some additional summary statistics.

Firstly, Figures 15 and 16 show automation cost spikes in both shares and levels for a balanced sample of firms – that is, firms which are observed in the Production Statistics survey every single year over 2000-2016.

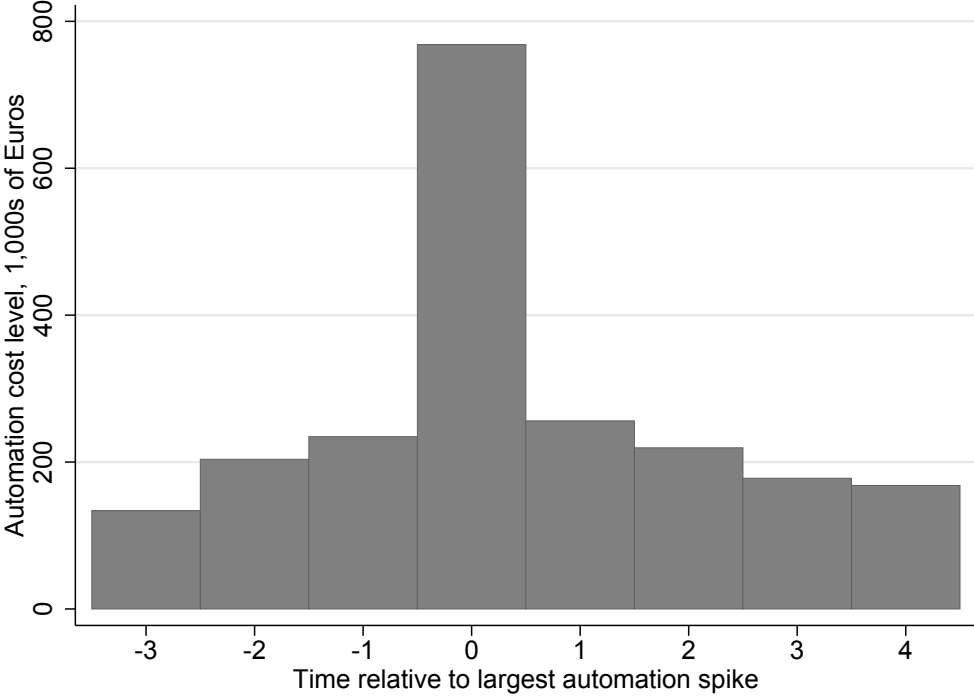
Next, Table 9 provides summary statistics on our sample of workers across all years, before matching. Next, Table 10 provides summary statistics on our matched sample of workers. For both incumbents and recent hires, we show the averages and standard deviations for the dependent as well as independent variables used in our models, separately for the treated and control group. Note that we have $115,698 + 1,049,546 = 1,165,244$ observations for incumbents and $94,757 + 378,433 = 473,190$ for recent hires: given our observation window of 8 years ($t = -3$ through $t = 4$) this adds up to the $1,179,584 \times 8 = 9,321,952$ observations for incumbents and $473,190 \times 8 = 3,785,520$ for recent hires used in our regressions.

Figure 15: Automation cost share spikes for treatment firms, balanced sample



Notes: N=330 in all years.

Figure 16: Automation cost spike level for treatment firms, balanced sample



Notes: N=330 in all years.

Table 9: Descriptives for all workers

	Incumbents (1)	Recent hires (2)
Annual wage earnings	36826.28 (25158.64)	23893.32 (21772.39)
Daily wage if employed	156.11 (95.18)	122.34 (80.50)
Annual non-employment duration (in days)	29.10 (62.47)	75.82 (95.53)
Probability of leaving the spiking firm	0.13 (0.33)	0.35 (0.48)
Year-round employment status	0.96 (0.19)	0.88 (0.33)
Total benefits	404.67 (2715.01)	1606.56 (4411.76)
Unemployment benefits	193.24 (1859.36)	819.11 (3109.38)
Disability benefits	199.52 (1943.60)	403.96 (2513.62)
Welfare	11.91 (353.24)	383.49 (2059.87)
Probability of entering early retirement	0.01 (0.09)	0.00 (0.05)
Probability of becoming self-employed	0.03 (0.18)	0.04 (0.20)
Share female	0.26 (0.44)	0.38 (0.49)
Foreign born or foreign-born parents	0.16 (0.36)	0.29 (0.45)
Age	42.68 (10.26)	36.74 (10.15)
Calendar year	2006.91 (3.38)	2006.97 (3.44)
Manufacturing	0.36 (0.48)	0.14 (0.35)
Construction	0.11 (0.32)	0.07 (0.25)
Wholesale and retail trade	0.19 (0.39)	0.16 (0.37)
Transportation and storage	0.09 (0.29)	0.08 (0.26)
Accommodation and food serving	0.02 (0.13)	0.03 (0.16)
Information and communication	0.05 (0.22)	0.05 (0.22)
Professional, scientific, and technical activities	0.08 (0.27)	0.07 (0.26)
Admin and support activities	0.09 (0.29)	0.41 (0.49)
0-9 employees	0.00 (0.06)	0.00 (0.07)
10-19 employees	0.03 (0.18)	0.03 (0.17)
20-49 employees	0.13 (0.34)	0.10 (0.30)
50-99 employees	0.11 (0.32)	0.09 (0.29)
100-199 employees	0.12 (0.33)	0.10 (0.30)
200-499 employees	0.15 (0.35)	0.11 (0.31)
≥500 employees	0.45 (0.50)	0.56 (0.50)
Observations	9,969,232	5,453,584


Notes: Unweighted means for all worker-year observations. Standard deviations in parentheses.

Table 10: Descriptives on matched worker samples

	Incumbents		Recent hires	
	<i>Treated</i>	<i>Control</i>	<i>Treated</i>	<i>Control</i>
	(1)	(2)	(3)	(4)
Annual wage earnings	36706.58 (24227.08)	36708.00 (24426.67)	21790.80 (17998.26)	21773.95 (17985.85)
Daily wage if employed	149.92 (93.35)	149.23 (93.34)	108.09 (66.38)	105.52 (66.86)
Annual non-employment duration (in days)	20.96 (42.75)	20.04 (41.47)	72.84 (77.47)	68.32 (76.51)
Total benefits	0.00 (0.00)	0.00 (0.00)	1268.14 (3471.95)	1215.47 (3413.38)
Unemployment benefits	0.00 (0.00)	0.00 (0.00)	705.79 (2461.18)	647.49 (2385.00)
Disability benefits	0.00 (0.00)	0.00 (0.00)	254.10 (1895.58)	247.61 (1848.60)
Welfare receipts	0.00 (0.00)	0.00 (0.00)	308.25 (1660.73)	320.36 (1725.35)
Probability of entering early retirement	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Probability of becoming self-employed	0.03 (0.16)	0.03 (0.17)	0.03 (0.18)	0.03 (0.17)
Share female	0.33 (0.47)	0.31 (0.46)	0.44 (0.50)	0.41 (0.49)
Foreign born or foreign-born parents	0.18 (0.39)	0.18 (0.38)	0.34 (0.47)	0.37 (0.48)
Age	40.93 (10.10)	40.98 (9.99)	35.98 (10.14)	35.68 (10.11)
Calendar year	2006.35 (2.42)	2006.35 (2.42)	2007.16 (1.93)	2007.16 (1.93)
Manufacturing	0.25 (0.43)	0.25 (0.43)	0.07 (0.26)	0.07 (0.26)
Construction	0.08 (0.28)	0.08 (0.28)	0.03 (0.18)	0.03 (0.18)
Wholesale and retail trade	0.25 (0.44)	0.25 (0.44)	0.12 (0.32)	0.12 (0.32)
Transportation and storage	0.08 (0.28)	0.08 (0.28)	0.05 (0.22)	0.05 (0.22)
Accommodation and food serving	0.02 (0.14)	0.02 (0.14)	0.02 (0.13)	0.02 (0.13)
Information and communication	0.05 (0.22)	0.05 (0.22)	0.04 (0.20)	0.04 (0.20)
Professional, scientific, and technical activities	0.10 (0.30)	0.10 (0.30)	0.06 (0.23)	0.06 (0.23)
Admin and support activities	0.16 (0.36)	0.16 (0.36)	0.61 (0.49)	0.61 (0.49)
0-9 employees	0.00 (0.03)	0.00 (0.07)	0.00 (0.02)	0.00 (0.06)
10-19 employees	0.04 (0.20)	0.04 (0.19)	0.02 (0.13)	0.02 (0.15)
20-49 employees	0.15 (0.36)	0.14 (0.34)	0.08 (0.26)	0.07 (0.26)
50-99 employees	0.12 (0.32)	0.12 (0.32)	0.07 (0.25)	0.07 (0.26)
100-199 employees	0.12 (0.33)	0.13 (0.33)	0.08 (0.27)	0.09 (0.29)
200-499 employees	0.15 (0.36)	0.15 (0.36)	0.07 (0.26)	0.10 (0.30)
≥500 employees	0.41 (0.49)	0.42 (0.49)	0.69 (0.46)	0.64 (0.48)
Observations	115,698	1,049,546	94,757	378,433

Notes: Weighted means for the full regression sample at $t = -1$, where weights are obtained from coarsened exact matching as described in section 3.3. Standard deviations in parentheses.

Figure 17: Annual real wage income in levels



Notes: N=9,321,952 for incumbents; 3,785,520 for recent hires. Whiskers represent 95 percent confidence intervals.

A.2 Additional results

Here we present some additional results: in particular, we show estimates on income levels, and discuss how our findings differ by workers' initial sector of employment and firm size.

A.2.1 Effects in levels

update this Figure 17 shows estimates for our main specification in levels. This shows that incumbents lost around 3,800 euros over five years, in total, which is comparable to the 3,010 euros lost when estimating impacts on relative income (shown in Figure 5). Recent hires lose around 937 euros in total over the same period: note that the relative income effects for recent hires are positive (see Figure 5), highlighting that results for recent hires are much more dependent on the exact model specification

A.2.2 Effect heterogeneity by residual wage quartile

We obtain an alternative measure of skill by calculating each worker's skill level based on a residual wage regression. That is, we regress worker wages in $t = -1$ onto a set of observables and their interactions: gender, age and age squared, nationality, and sector, as well as year

Table 11: Effect heterogeneity by residual wage quartile

	(1) Relative wage income	(2) Firm separation probability	(3) Non-employm. duration	(4) Log daily wage	(5) Early retirem. probability
Panel A. Incumbents					
Bottom quartile (ref)	-1.44*** (0.38)	3.99*** (0.26)	3.69*** (0.42)	0.07 (0.17)	0.20*** (0.05)
<i>Deviations from reference group for:</i>					
Second quartile	-0.02 (0.47)	0.77* (0.36)	-0.68 (0.57)	-0.45* (0.21)	0.02 (0.07)
Third quartile	0.46 (0.47)	1.49*** (0.37)	-1.65** (0.57)	-0.18 (0.21)	-0.06 (0.07)
Top quartile	-1.34** (0.49)	1.83*** (0.36)	0.50 (0.59)	-0.52* (0.23)	0.13 (0.08)
Observations	9,319,736	9,319,736	9,319,736	8,999,617	9,319,736
Panel B. Recent hires					
Bottom quartile (ref)	4.76*** (1.30)	2.82*** (0.36)	-1.53 (0.87)	0.21 (0.42)	-0.02 (0.05)
<i>Deviations from reference group for:</i>					
Second quartile	-3.27* (1.65)	1.07* (0.51)	1.43 (1.22)	0.41 (0.53)	0.11 (0.06)
Third quartile	-5.38*** (1.61)	4.07*** (0.52)	4.01*** (1.22)	-0.24 (0.52)	0.00 (0.06)
Top quartile	-3.46* (1.61)	6.46*** (0.53)	2.93* (1.20)	-0.14 (0.53)	0.14 (0.08)
Observations	3,783,024	3,783,024	3,783,024	3,311,071	3,783,024

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Quartile 1 workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

fixed effects. We perform this regression separately for incumbents and recent hires, and then construct four skill quartiles for each of these two worker groups based on the residuals of this model. In this measure, higher quartiles correspond to higher skill levels. This measure has the advantage that it also captures unobserved components of skill, while it of course has the downside of potentially also capturing other factors impacting wage levels.³⁴

Results are reported in Table 11: workers in the lowest residual wage quartile (quartile 1) are used as the reference category throughout. Consistent with findings reported in the main text, we conclude that “higher-skilled” incumbents are most negatively affected by automation.

³⁴Workers’ residual skill quartile is indeed positively correlated with education levels, where this information is available – see Table 12.

Table 12: Distribution of education groups across residual wage quartiles

Education level	Residual wage quartile at $t = -1$				Total
	Bottom	Q2	Q3	Top	
Low	37%	29%	22%	12%	100%
Middle	19%	26%	29%	26%	100%
High	6%	14%	28%	52%	100%

Notes: Percentage of workers of 3 education levels categorized into 4 residual wage quartiles, at $t = -1$. Low education are workers with a 4-year high-school degree or less (*seq*“MBO”), middle-skilled workers are those with 5+ year high-school degrees and/or a vocational diploma (“HAVO, VWO, MBO2-MBO4”), and high-educated are those with tertiary degrees.

A.2.3 Effect heterogeneity by education level

While we only observe education level for a much smaller and selected subset of workers, we here report estimates to corroborate our findings from alternative skill measures. In particular, education level is available as three categories – low, middle, and high, which are described in the footnote to Table 12.

Table 12 also provides the distribution of workers of three different education levels across residual wage quartiles. Since the availability of skill data is limited and biased towards higher-educated workers, we ignore any differences in the shares of each of the education groups in constructing this table. As such, if workers of different education types were evenly distributed across the four residual wage quartiles, all numbers in this table should be 25%. Instead, we find that 37% of low-educated workers are found in the lowest wage quartile, compared to only 19% and 6% for middle- and high-educated workers, respectively. Further, only 12% of the low-educated are observed in the fourth, and highest, residual wage quartile, compared to 26% and 52% of middle- and high-educated workers. This confirms that these residual wage quantiles broadly correlate with workers’ education level, even if this need not be the only factor they correlate with.

Table 13 shows estimates by education level. This confirms that the firm separation probability increases with workers’ skill level, for both incumbents and recent hires. However, income losses are not statistically significantly different across education levels, suggesting higher-skilled workers are better able to adjust to the automation shock. This is mostly because higher-skilled incumbents and recent hires experience smaller increases in non-employment.

Table 13: Effect heterogeneity by education level

	(1)	(2)	(3)	(4)	(5)
	Relative wage income	Firm separation probability	Non-employm. duration	Log daily wage	Early retirem. probability
Panel A. Incumbents					
Low education (ref)	-1.98** (0.64)	4.07*** (0.48)	4.36*** (0.83)	0.13 (0.48)	0.05 (0.08)
<i>Deviations from reference group for:</i>					
Middle education	0.24 (0.76)	2.74*** (0.58)	-1.92 (0.99)	-1.43* (0.56)	0.13 (0.09)
High education	0.48 (0.83)	3.54*** (0.60)	-2.62** (0.96)	-1.20* (0.53)	0.04 (0.09)
Observations	2,749,439	2,749,439	2,749,439	2,662,211	2,749,439
Panel B. Recent hires					
Low education (ref)	-0.38 (1.35)	4.24*** (0.40)	3.18** (1.13)	-2.69*** (0.55)	0.03 (0.05)
<i>Deviations from reference group for:</i>					
Middle education	2.57 (1.71)	1.24* (0.52)	-4.14** (1.40)	2.17** (0.66)	-0.06 (0.07)
High education	3.52 (1.90)	3.04*** (0.61)	-5.16*** (1.47)	2.75*** (0.69)	-0.03 (0.06)
Observations	1,763,766	1,763,766	1,763,766	1,584,999	1,763,766

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Quartile 1 workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.2.4 Effect heterogeneity by sector

We now consider to what extent the impacts of automation differ depending on which sector the worker's firm belongs to: that is, do automation events originating in different sectors have different impacts on workers? Note that such sectoral differences may exist for various reasons. Firstly, automation technologies may be sector-specific, and differ in terms of how much they displace labor. For example, it is possible that industrial or warehouse robots are more labor-replacing than automated check-out systems. Secondly, the workers employed in these different industries may have different characteristics (such as average age), making the impacts differ by sector. Thirdly, to the extent that skills are industry-specific, sectoral labor market conditions matter: displacement would be costlier in sectors with an excess supply of workers. While we cannot distinguish between these different explanations, it is still important to consider whether our results are driven by displacement effects in a subset of sectors, or whether the found impacts are pervasive. As shown in Table 9, the main sectors of employment for incumbents are Manufacturing and Wholesale and retail as over 50 percent of these workers are initially

employed there; followed by Construction, and Administrative and support activities, which employ another 21 percent of incumbent workers. For recent hires, 41 percent is employed in Administrative and support activities, followed by around 15 percent in Manufacturing and Wholesale and retail trade, each.

Estimates interacting our treatment effects with workers' sector of employment in $t = -1$ are shown in Table 14. For each of these models, manufacturing is the reference category. A first important finding is that firm separation effects are found across a wide range of sectors. Indeed, for incumbents, substantial automation-driven displacement from the firm is found for all sectors except Accommodation and food serving – though firm separation rates (weakly) increase for incumbents from all sectors. Among recent hires, firm separation rates also increase everywhere, with the exception of Accommodation and food serving, where they actually decline (although only 3 percent of recent hires are employed in this sector). Despite the effects being pervasive across sectors, the increase in worker separations is highest among incumbents employed in Manufacturing, Transportation and storage, and Administrative and support activities.

Increases in non-employment duration for incumbents are seen across most sectors as well (though not for recent hires, as found on average across sectors as well), with two exceptions. Firstly, no increase (and rather, a decrease) is seen for Accommodation and food serving, where no increase in firm separation for incumbents occurred in the first place. Secondly, incumbents initially employed in Administrative and support activities do not experience increases in non-employment after the automation event despite increased firm separation rates, suggesting these workers face more favorable labor market conditions or are otherwise more able to adjust.

Effects on wage rates, while occasionally statistically significant, are small across the board: this implies our overall finding is not concealing substantial sectoral heterogeneity. Lastly, early retirement effects for incumbents are largely driven by manufacturing workers.

All in all, these results show that while incumbents initially employed in manufacturing sectors do experience some of the larger firm separation and non-employment increases, the automation-driven displacement effects found for incumbents are pervasive across many different sectors of the economy. Further, overall differences found between incumbents and recent hires are not only due to their different sectoral affiliations: for example, manufacturing workers are displaced from their firms at similar rates whether they are incumbents or recent hires, but only incumbents experience increases in non-employment and early retirement.

A.2.5 Effect heterogeneity by firm size

We consider how the worker effects of automation effects differ by firm size. For our purposes, this is also important because our automation cost survey overrepresents large firms – while these of course employ the majority of workers, it could still bias the found worker-level effect of automation events by including too low a number of workers experiencing such events in small firms.³⁵

Estimates are reported in Table 15, where the largest firms (those with 500 employees or more) are the reference category. Note that this effect only pertains the workers' firm size in the year before the automation event: firm size may change over time, and workers may move to firms of different sizes. This highlights several findings.

Firstly, effects on firm separation and non-employment duration for incumbents are robust across different firm sizes: workers in firms of all sizes experience an increase in the probability of firm separation and days spent in non-employment four years after an automation event. The separation effect is biggest for the largest firms but increases in non-employment and daily wage losses are largest for the very smallest firms. All in all, this indicates that while displacement occurs more frequently as a result of an automation event in the largest firms, finding (similarly paid) re-employment is easier, also. As a result, income losses for incumbents occur everywhere but are decreasing in firm size: and workers displaced from the very largest firms do not experience wage losses at all.

Secondly, results for recent hires are qualitatively similar to those for incumbents: workers at larger firms are better off. Indeed, those employed at the largest firms experience substantial wage gains of 3.5% annually on average (and no increase in non-employment duration), whereas those at smaller firms lose income (and increased non-employment). Since recent hires are overrepresented among the largest firms, the average effect is dominated by the largest firms, where losses are absent. This highlights that a firm-size effect does matter in understanding the different outcomes of incumbents and recent hires. However, recent hires do also differ in terms of their ability to adjust to firm separation, as discussed in the main text: workers at the largest firms are displaced from their firms at similar rates whether they are incumbents or recent hires, but only incumbents experience increases in non-employment and early retirement.

Lastly, increases in early retirement are predominantly observed among incumbents initially

³⁵Note that we do not include these workers in our regression models as the control group, either, as explained in section 3.3.

Table 14: Effect heterogeneity by sector

	(1) Relative wage income	(2) Firm separation probability	(3) Non-employm. duration	(4) Log daily wage	(5) Early retirem. probability
Panel A. Incumbents					
Manufacturing (ref)	-1.52*** (0.21)	5.34*** (0.20)	3.89*** (0.35)	0.07 (0.13)	0.35*** (0.06)
<i>Deviations from reference group for:</i>					
Construction	-0.68 (0.50)	-4.03*** (0.41)	-0.06 (0.74)	-0.34 (0.24)	-0.10 (0.15)
Wholesale and retail trade	-1.83*** (0.36)	-1.34*** (0.30)	2.09*** (0.50)	-0.86*** (0.20)	-0.16* (0.07)
Transportation and storage	0.96 (0.54)	1.70*** (0.46)	-1.06 (0.76)	0.25 (0.29)	-0.40** (0.13)
Accommodation and food serv	4.07** (1.33)	-4.59*** (0.95)	-8.63*** (1.45)	0.14 (0.59)	-0.72*** (0.14)
Information and comm	-3.10*** (0.68)	-0.95 (0.57)	-0.82 (0.94)	-2.58*** (0.38)	-0.33** (0.11)
Prof scientific techn act	-0.95 (0.50)	-0.15 (0.44)	-1.22 (0.70)	-1.40*** (0.27)	0.00 (0.09)
Admin and support act	2.72*** (0.66)	2.35*** (0.48)	-4.40*** (0.76)	0.80* (0.34)	-0.07 (0.08)
Observations	9,321,952	9,321,952	9,321,952	9,001,780	9,321,952
Panel B. Recent hires					
Manufacturing (ref)	0.80 (0.93)	4.80*** (0.56)	0.02 (1.02)	-0.08 (0.39)	-0.11 (0.06)
<i>Deviations from reference group for:</i>					
Construction	-5.61*** (1.45)	3.54*** (1.02)	5.29** (1.71)	-1.47* (0.66)	0.39* (0.16)
Wholesale and retail trade	-1.93 (1.26)	2.72*** (0.72)	4.27** (1.32)	0.24 (0.51)	0.20** (0.07)
Transportation and storage	3.37 (1.96)	-2.69** (0.96)	-1.76 (1.86)	-1.23 (0.75)	0.11 (0.12)
Accommodation and food serv	5.94 (3.67)	-9.40*** (1.41)	-5.57 (2.91)	2.00 (1.25)	-0.10 (0.15)
Information and comm	-6.17*** (1.80)	-1.06 (0.99)	3.78* (1.84)	-1.01 (0.71)	0.16* (0.08)
Prof scientific techn act	-4.34** (1.59)	-1.35 (0.90)	3.60* (1.66)	-1.30* (0.64)	0.29** (0.09)
Admin and support act	1.79 (1.24)	1.23* (0.61)	0.08 (1.20)	0.28 (0.48)	0.15* (0.07)
Observations	3,785,520	3,785,520	3,785,520	3,313,317	3,785,520

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Workers employed in manufacturing firms at the time of treatment are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

employed in the largest and smallest firms.

All in all, the results in this section suggest we would probably find higher average wage losses from automation for both incumbents and recent hires if our data were more representative in terms of firm size.

Table 15: Effect heterogeneity by firm size

	(1)	(2)	(3)	(4)	(5)
	Relative wage income	Firm separation probability	Non-employ. duration	Log daily wage	Early retirem. probability
Panel A. Incumbents					
≥500 employees (ref)	-0.01 (0.28)	6.81*** (0.22)	3.38*** (0.33)	0.71*** (0.14)	0.41*** (0.04)
<i>Deviations from reference group for:</i>					
<10-19 employees	-4.55*** (0.75)	-2.98*** (0.61)	2.72** (1.04)	-3.46*** (0.42)	-0.06 (0.13)
20-49 employees	-4.22*** (0.46)	-3.26*** (0.37)	0.93 (0.59)	-2.31*** (0.23)	-0.18* (0.08)
50-99 employees	-2.23*** (0.50)	-2.88*** (0.41)	-1.11 (0.64)	-1.59*** (0.25)	-0.45*** (0.08)
100-199 employees	-2.20*** (0.49)	-2.52*** (0.41)	-0.72 (0.64)	-1.80*** (0.26)	-0.34*** (0.08)
200-499 employees	-2.14*** (0.45)	-3.07*** (0.37)	-0.70 (0.60)	-0.76** (0.24)	-0.43*** (0.08)
Observations	9,321,952	9,321,952	9,321,952	9,001,780	9,321,952
Panel B. Recent hires					
≥500 employees (ref)	3.50*** (0.74)	6.48*** (0.22)	-1.07 (0.57)	0.08 (0.25)	0.05 (0.03)
<i>Deviations from reference group for:</i>					
<10-19 employees	-11.10*** (2.31)	-3.85** (1.25)	5.93* (2.38)	-2.87** (0.93)	-0.06 (0.09)
20-49 employees	-6.46*** (1.44)	-1.60* (0.67)	3.95** (1.33)	-1.15* (0.52)	-0.09 (0.06)
50-99 employees	-5.97*** (1.61)	-3.69*** (0.73)	7.29*** (1.44)	-0.65 (0.57)	-0.09 (0.07)
100-199 employees	-8.97*** (1.72)	-4.88*** (0.68)	8.57*** (1.44)	-1.38* (0.62)	0.09 (0.09)
200-499 employees	-3.58* (1.67)	-9.88*** (0.68)	4.59** (1.40)	1.97*** (0.59)	-0.05 (0.07)
Observations	3,785,520	3,785,520	3,785,520	3,313,317	3,785,520

Notes: All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$), and coefficients in columns 1, 3, 4, and 5 have been multiplied by 100. Male workers are the reference group. Log daily wages are observed conditional on employment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.3 Robustness checks

In this section, we present several robustness checks on our results. In particular, one may be concerned that our findings for automation spikes are the result of other firm-level events which impact the firm’s workers. We consider robustness to this in three main ways.

Firstly, our data include administrative information on some of these events, such as mergers, take-overs, acquisitions, firm splits, and restructuring. We additionally observe firm births and deaths, but these are already excluded since we consider a balanced sample of firms over the observation window: we do allow firm births in the first year of observation, however. We therefore compare the estimated effects in our baseline model to those when we remove all events (including firm births in $t = -3$) occurring in the entire observation window.


Beyond these observable events, there may be other firm-level changes correlating with worker-level outcomes. Note that in our baseline specification, we do not see any pretrends at the worker level, but the parallel trends assumption may of course still fail in the post-treatment period if such unobserved events coincide with the automation spike. In order to deal with this possibility, we remove firms that experience at least one very large employment change at any time over the observation window – where we define this as an employment change exceeding 90 percent (in absolute terms). This is our second robustness check, and the idea is that such large changes may signal other (unobserved) events.

Lastly, as a third robustness check, we additionally match workers across firms with similar pre-treatment employment growth trajectories. This implies we now compare treated workers with control workers not only matched on individual pre-treatment income levels, calendar year, and sector of employment, but also ensure the treated and control workers are employed at firms where pre-treatment employment growth is in the same quartile of the between-firm pre-treatment employment growth distribution. Also here, the idea is to control for unmeasured events which impact labor demand at the firm level.

Figures 18 and 19 summarize the results for all three robustness checks (along with baseline model estimates), for incumbents and recent hires, respectively. Results for incumbent workers are very robust: incumbents lose wage income as a result of automation, cumulating to around 10 percent of average annual earnings. For recent hires, estimates vary but are smaller than the found effects for incumbents, and not statistically significantly different from zero in some models.


It is worth noting that we do still find an increased firm separation probability in all models for both incumbents (Figure 20) *and* recent hires (Figure 21). This reinforces our conclusion that while there is evidence that automation displaces both incumbent workers and recent hires, only the former suffer income losses.

Figure 18: Wage income effects for incumbent workers: robustness checks




Notes: Whiskers represent 95 percent confidence intervals.

Figure 19: Wage income effects for recent hires: robustness checks




Notes: Whiskers represent 95 percent confidence intervals.

Figure 20: Firm separation probability for incumbent workers: robustness checks



Notes: Whiskers represent 95 percent confidence intervals.

Figure 21: Firm separation probability for recent hires: robustness checks



Notes: Whiskers represent 95 percent confidence intervals.