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*

September 2019

Abstract

We provide the first estimates of automation's impacts 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 a differences-in-differences design exploiting automation event timing, we find that firm-level automation increases the probability of workers separating from their employers and decreases days worked, leading to a 5-year cumulative wage income loss of 11 percent of one year's earnings. These losses are only partially offset by benefits systems, and quite pervasive across worker types, firm sizes and sectors. Further, no such losses are found for computerization events.

Keywords: Automation, Technological Change, Displacement

JEL: J23, J31, J62, J63, O33

*Helpful comments by David Autor, Egbert Jongen, Guy Michaels, Pascual Restrepo, Bas van der Klaauw, and participants of seminars and conferences at CPB, the European Central Bank, Erasmus University Rotterdam, Groningen University, IZA, NBER, ROA Maastricht, TASKS V Bonn, Utrecht School of Economics, and University of Sussex 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 job tasks. These technologies range from robotics to machine learning and other forms of artificial intelligence, and are being adopted across many sectors of the economy. Applications range from selecting job applicants for interviewing, picking orders in a warehouse, interpreting X-rays to diagnose disease, and automated customer service. These developments have raised concern that workers are being displaced by advancing automation technology. Indeed, opinion surveys from the US and Europe highlight that a majority of individuals are worried about the future of work and expect worsening employment prospects, even as they foresee a positive impact on the economy and on society more generally (Eurobarometer 2017; Pew 2017).

This potential for automation to displace workers is studied 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 at the micro level plays a central role, as machines take over tasks previously performed by humans. Under certain conditions, such displacement is a possible outcome of automation even in aggregate.

Empirical work on automation has so far mostly focused on robotics – a prime example of automation technology, albeit one that has penetrated only a limited number of sectors – and on more aggregate outcomes.[1] The macro-economic 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. Further, Koch et al. (2019) show that firms that adopt robots experience net employment growth compared to firms that do not, and Dixon et al. (2019) find that a firm's employment growth rises with their robot stock.

Besides macro-economic and firm-level impacts, it is critical to also study automation's 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. Any such adjustment

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

costs are also of first-order importance for policymakers aiming to diminish adverse impacts out of distributional concerns.

So far, direct empirical evidence on the worker-level impacts of automation is lacking. Existing studies on worker adjustments have used more aggregate sources of variation and do not always focus on causal effects. In particular, Dauth et al. (2018) correlate regional variation in robot exposure with worker outcomes; Cortés (2016) finds that workers switching out of routine-intense occupations experience faster wage growth relative to those who stay; while Edin et al. (2019) show that workers have worse labor market outcomes when their occupation is experiencing long-term decline. To our knowledge, our paper provides the first estimate of the economic impacts on workers when their firm invests in automation technology.

This study makes several contributions. First, we directly measure automation at the firm level and can therefore analyze the worker impacts of automation where they originate: at the automating firms. We do so by linking 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 36,490 firms with at least three years of automation cost data, employing close to 5 million unique workers per year on average. Second, we develop and implement a differences-in-differences methodology leveraging the timing of firm-level automation events for identifying causal effects. Third, 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, complementing the literature focused on robotics. Fourth, we measure a rich array of outcomes for individual workers for the years surrounding the automation event: this provides insight in how any adjustment costs come about. These outcomes include annual 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 (those 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. Finally, we directly compare the current worker-level impacts of automation to those of computerization.

We find that automation at the firm increases the probability of incumbent 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 11 percent of one year's earnings. On the other hand, wage rates are not much affected: that is, we do not see wage scarring for workers impacted by automation. This is in contrast to displacement from mass lay-offs or firm closures, which have been studied in another literature (see Jacobson et al. 1993; Couch and Placzek 2010; Davis and Von Wachter 2011). However, lost wage earnings from non-employment spells are only partially offset by various benefits systems, and older workers are more likely to enter early retirement. Further, automation's displacement effects are found to be quite pervasive across different incumbent worker types as well as firm sizes and sectors. In contrast, we do not find evidence for such displacement from investments in computer technology. This suggests that for incumbent workers, automation is a more labor-displacing force.

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 differences-in-differences estimation framework. Our main results are reported in section 4: subsections consider impacts on workers' wage income and its components; additional adjustment mechanisms; robustness checks; and effect heterogeneity. In section 5 we directly compare the worker-level impacts of automation to those of computerization. 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.[2] This survey can be matched to administrative company ("*Algemeen Bedrijfsregister*") and worker records ("GBA" and "BAAN" files).

Our data cover the years 2000-2016, and we retain 36,490 unique firms with at least 3 years of automation cost data – together, these firms employ around 5 million unique workers annually

[2] Firms 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.

on average. We remove firms where Statistics Netherlands indicate that the data are (partly) imputed.[3] We further remove workers enrolled in full-time studies, and those 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. 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 nationality.[4] A downside to these data is that we neither observe workers' occupations nor their level of education: the former is unavailable entirely, whereas the latter is only defined for a small and selected subset of workers (with availability skewed toward younger and high-educated workers). 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.[5] While the disadvantage of this measure is that we do not know the exact automation technology being used by the firm, it does capture all automation technologies rather than focusing on a single one, and we measure it at the level of the firm rather than the industry, and across all private non-financial sectors. From discussions with company representatives and automation services providers, we know that these expenditures are related to automation technologies such as self-service check-outs, warehouse and storage systems, data-driven decision making, or automated customer service. Another example are robotics integrator services highlighted (and used as an instrument for robotic technology adoption) in Acemoglu and Restrepo (2018d).

Table 1 shows summary statistics on annual automation costs for firms, both in levels, per worker, and as a percentage of total costs (excluding automation costs). This highlights several

[3]In 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.

[4]In these data, individuals are classified as "Dutch" if they themselves and both of their parents have been born in the Netherlands.

[5]This also includes non-activated purchases of custom software and costs of new software releases, but excludes prepackaged software licensing costs.

Table 1. Automation cost share distribution

	All observations			Automation costs >0		
	Cost level	Cost per worker	Cost share (%)	Cost level	Cost per worker	Cost share (%)
p5	0	0	0	2,026	54	0.04
p10	0	0	0	3,652	92	0.06
p25	0	0	0	9,537	234	0.14
p50	10,508	257	0.15	27,390	587	0.32
p75	48,000	899	0.47	85,597	1,322	0.68
p90	175,035	2,058	1.05	278,213	2,697	1.37
p95	412,945	3,305	1.69	650,966	4,200	2.13
mean	*192,391*	*953*	*0.44*	*280,713*	*1,391*	*0.64*
N firms × years		240,320			164,707	
N with 0 costs		31%			0%	

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

things. First, almost one-third of firm-year observations has zero automation expenditures. Second, the average automation cost share is 0.44 percent, corresponding to an outlay of around 200K euros annually, or 953 euros per worker. Third, this distribution is highly right-skewed as the median is only 0.15 percent – this skewness persists even when removing observations with zero automation costs.

Table 2 further shows how these automation costs and cost shares differ by broad (one-digit) sector. Our comprehensive measure of automation technologies indicates that all sectors have automation expenditures, though there is substantial variation at the firm level both between and within each of these sectors. Average expenditures at the sectoral level range from 220 to 1,636 euros per worker. The highest mean automation expenditures per worker are observed in Professional, scientific, and technical activities, followed by Information and communication, Wholesale and retail, and Manufacturing. Conversely, Accommodation and food serving has the lowest expenditure per worker, followed by Construction, Administrative and support activities, and Transportation and storage. However, there is much variation between firms in the same sector, as shown by the standard deviations of the automation cost share in total (other) costs. 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 technologies may be sector-specific.

Table 3 reports the same statistics but separately by firm size class, grouped into 6 classes used by Statistics Netherlands: the smallest firms have up to 19 employees whereas the largest

Table 2. Automation costs by sector

Sector	Mean cost level (€) Total	Per worker	Cost share (%) *Mean*	*SD*	Nr of obs *Firms*	*Firms × yrs*
Manufacturing	391,214	986	0.36	0.58	5,655	44,636
Construction	71,150	414	0.2	0.36	4,688	28,757
Wholesale & retail trade	106,259	1,075	0.31	0.80	11,041	75,421
Transportation & storage	257,057	834	0.42	1.07	3,122	21,235
Accommodation & food serving	49,475	220	0.29	0.50	1,292	6,761
Information & communication	409,511	1,636	0.85	2.92	2,655	16,854
Prof'l, scientific, & technical activities	136,437	1,174	1.02	1.76	4,074	23,692
Administrative & support activities	121,301	761	0.49	1.18	3,963	22,964

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,490; Total N firms × years is 240,320.

Table 3. Automation costs by firm size class

Firm size class	Total cost Mean	Cost per worker Mean	SD	Cost share (%) Mean	SD	Nr of obs *Firms*	*Firms × yrs*
1-19 employees	11,135	836	13,255	0.40	1.29	9,850	48,758
20-49 employees	25,287	815	4,152	0.42	1.34	13,777	87,188
50-99 employees	56,336	873	3,975	0.42	0.96	6,291	47,209
100-199 employees	132,573	1,038	5,318	0.44	0.94	3,471	28,748
200-499 employees	372,095	1,440	19,498	0.51	1.11	1,969	17,897
≥500 employees	2,885,712	1,937	13,082	0.76	1.60	1,132	10,520

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,490; Total N firms × years is 240,320.

have more than 500. Unsurprisingly, automation cost levels rise with firm size: firms with fewer than 20 employees spend around 11K euros annually on automation services, whereas the largest firms spend close to 2.9 million. Less obviously, this table also reveals that automation cost shares increase with firm size, particularly at the very top. The smallest firms have average automation cost shares of 0.40 percent[6], whereas firms with between 20 to 200 employees have a cost share of around 0.44 percent. This increases to 0.51 percent for firms between 200 and 500 workers, and 0.76 percent for firms with more than 500 workers. There is substantial variation within size classes, also.

Figures 1 and 2 further show how the distribution of automation cost shares and expenditures per worker change over time. Mean automation cost share and outlays per worker are rising in the Netherlands over 2000-2016, from 0.28 to 0.57 percent relative to total other costs, and from 744 to 1,103 euros per worker. All else being equal, this implies that workers' exposure to automation is also rising. Furthermore, besides an increase in the average, there is a fanning

[6] The relatively high expenditure per worker for the smallest firm size class is driven by a small number of one-person firms with high automation expenditures – when we eliminate the top 1 percent of observations in terms of automation cost per worker, outlays per worker are monotonically rising in firm size as reported in Table 14 in the Appendix.

Figure 1. Firm-level automation cost shares over time

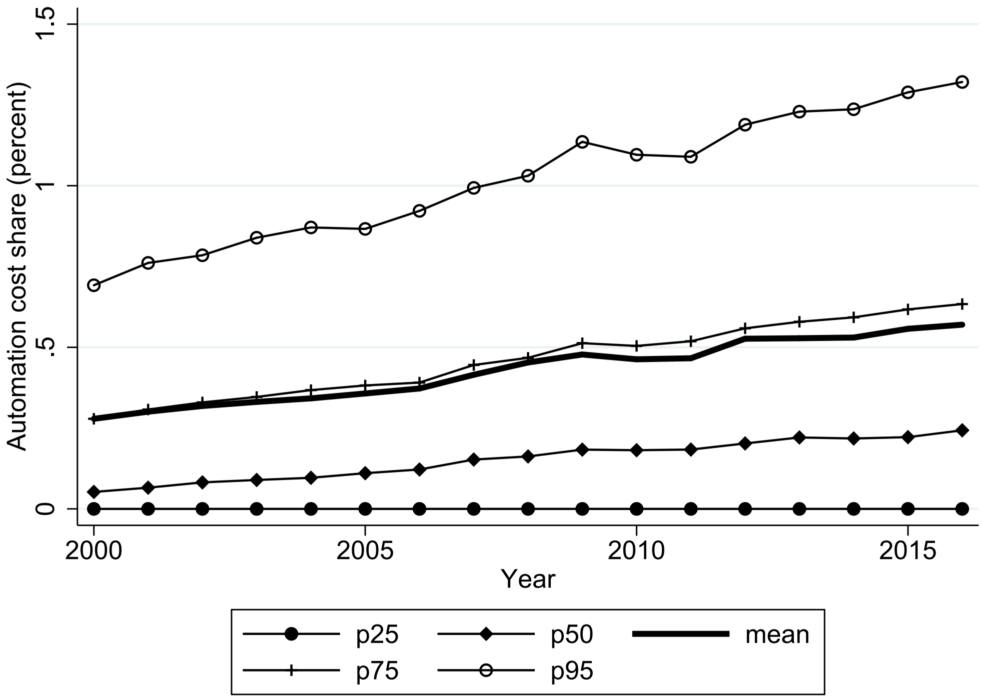

Figure 2. Firm-level automation cost per worker over time

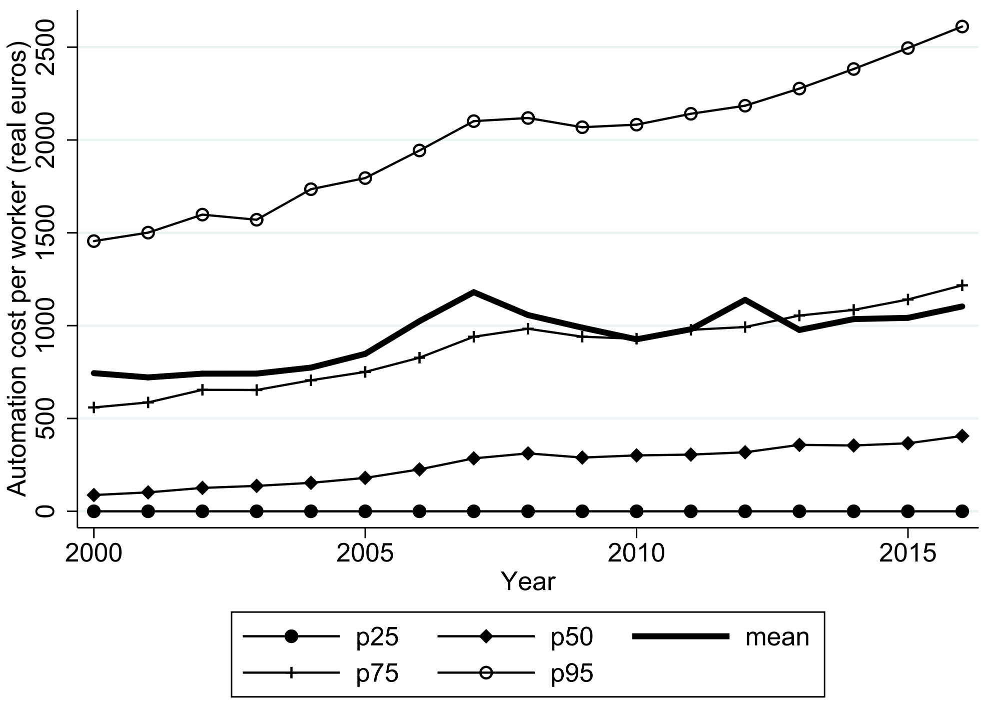

out of the distribution with automation cost shares rising faster for higher percentiles.

Lastly, we find that automation expenditures are somewhat correlated with computer investments: these are available from a different, and partially overlapping, firm-level survey. In section 5, below, we consider the robustness of our results to excluding firms that have investment events in both technology types within the estimation window, as well as study how the worker impact of computer investment events differs from that of automation.

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 novel approach for both of these challenges is to use what we term *automation spikes*. In particular, we assume that spikes in automation cost shares at the firm level signal changes in work processes related to automation.

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 years t, $\overline{TC}_j$, are at least thrice the average firm-level cost share excluding year τ:

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

where $\mathbb{1}\{\ldots\}$ 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.

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 specification for a number of reasons. First, there may be measurement error in the survey variable making it more difficult to measure the exact size of a spike. Second, 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. In Appendix A.3.5, we report several robustness checks, including changing the automation spike definition and varying the spike threshold.

The existence of these automation cost spikes would be consistent with a literature on lumpy investment (Haltiwanger et al. 1999; Doms and Dunne 1998). In fact, 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 both include substantial irreversible investments (for example in terms of worker training or from developing custom software) as well as involve fixed adjustment costs from reorganizing production processes.

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 36,490 firms described above.

Table 4 shows that around 70 percent of firms never spike, whereas the remaining 30 percent spike at least once over the 17 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 over 2000-2016, although some spike twice and up to five times at most. Automation spikes are observed across all sectors, as Table 5 highlights. However, a higher share of firms in Information and communication experience such an event compared to firms in Construction or in Accommodation and food serving.

Figure 3 shows what automation spikes look like on average across firms where spikes are observed. This is constructed by redefining time t as the number of years relative to the spike in period τ, i.e. $t \equiv year - \tau$, such that all spikes line up in $t = 0$. When firms spike multiple times, we only include the largest spike. Figure 3 is not using a balanced panel of firms: rather, all 10,476 spiking firms are observed in $t = 0$, and the number of observations for other years

Table 4. Firm-level automation spike frequency

Spike frequency	**N firms**	**% of N firms**
0	26,014	71.3
1	8,411	23.1
2	1,764	4.8
3	267	0.7
4	30	0.1
5	4	0.0
Total	36,490	100.0

Notes: Spike frequency is defined as the total number of spikes occurring over 2000-2016. The total number of firms is 36,490.

depends on when the spike took place[7], and on how often the firm enters in the automation survey. Nevertheless, we see a clear spike pattern.

Figure 4 restricts the sample of firms with spikes in $t = 0$ to those firms that are observed in all years $t \in [-3, 4]$, as these are the treatment group firms we will actually use in the empirical design explained below.[8] Figure 4 shows that automation events are quite cleanly identified: these 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 5 repeats Figure 4 but for the implied level of automation expenditure per worker, showing that the average firm-level automation spike amounts to an investment of close to 1,900 euros per worker, compared to a usual level of around 440 euros in years close to the spike. Figures 4 and 5 are both weighted by firms' employment size: as such, they reflect the exposure to automation for the average treated worker in our sample.[9]

3.3 How do automating firms differ?

A potential control group for workers in automating firms are workers in firms that are not automating. However, here we show that these groups are not comparable. Table 6 first considers how the average automation expenditures compare across these two groups. This reveals that firms with automation events have higher average levels of automation expenditures, whether

[7] For example, if the spike occurred in the first calendar year of data, there are no observations for $t < 0$; if it took place in the last calendar year, there are no observations for $t > 0$.

[8] See Appendix A.1 for details on sample construction.

[9] In Figures 21 and 22 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 5. Automation spike frequency by sector

Sector	**N firms**	**N firms with spike**	**Spike frequency (%)**
Manufacturing	5,655	1,606	28.4
Construction	4,688	1,143	24.4
Wholesale & retail trade	11,041	3,004	27.2
Transportation & storage	3,122	937	30.0
Accommodation & food serving	1,292	329	25.5
Information & communication	2,655	1,023	38.5
Prof'l, scientific, & technical activities	4,074	1,293	31.7
Administrative & support activities	3,963	1,141	28.8

Notes: A spiking firm has at least once automation spike over 2000-2016. The total number of firms is 36,490, the total number of spiking firms is 10,476. Spike frequency is the ratio of spiking firms over total firms by sector.

Figure 3. Automation cost share spikes

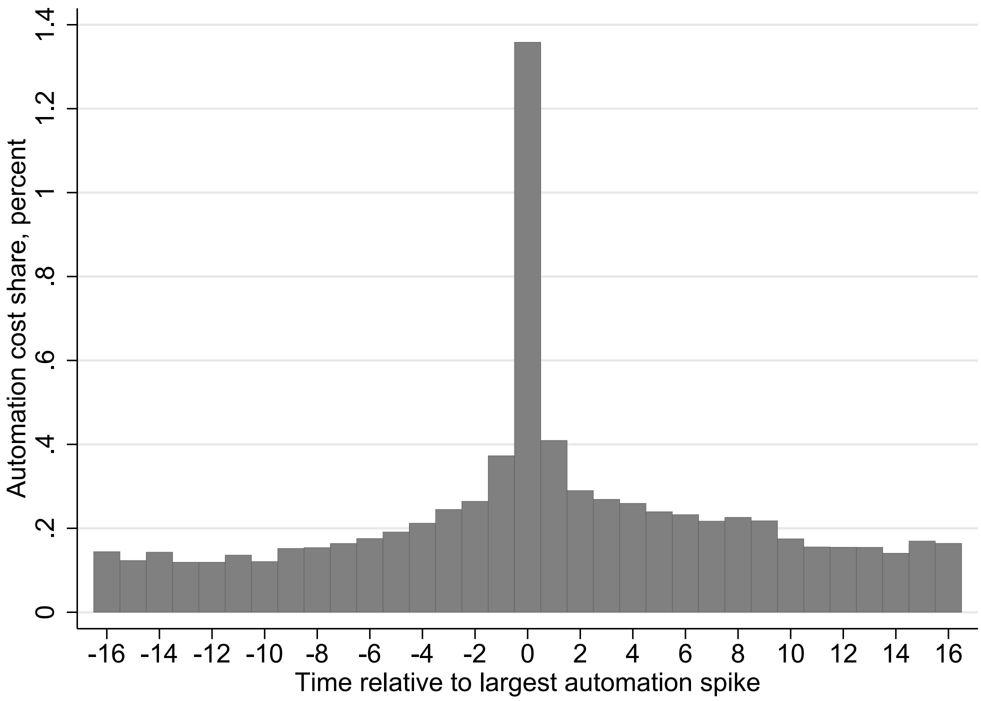

Notes: Unbalanced panel of firms, N=10,476 in $t = 0$.

expressed in absolute terms, or relative to the number of workers, or as a share in total costs. These differences are considerable: firms with automation events spend around twice as much on automation per worker or relative to total operating costs.[10]

Importantly, firms that make large automation investments have faster employment growth compared to firms that do not have automation spikes. This is shown in Figures 6 and 7 which respectively plot firm-level log employment and wagebill trajectories, for a balanced sample of

[10]Further, Table 15 in the Appendix shows how spiking firms' time-invariant characteristics differ from non-spiking ones: the main finding is that firms that experience automation spikes are larger.

Figure 4. Automation cost share spikes for treated firms

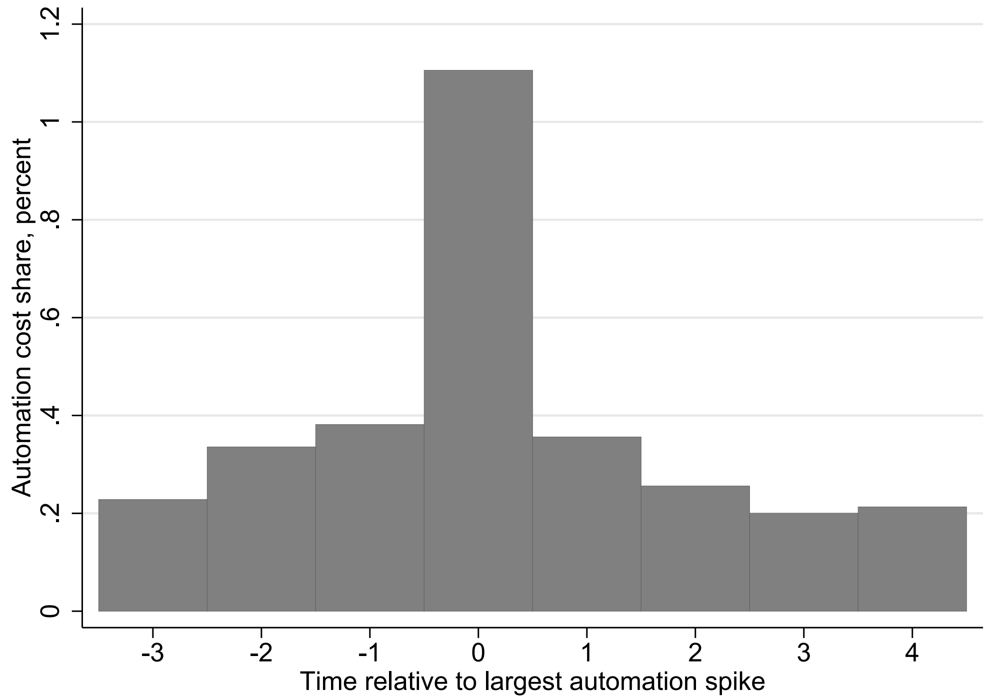

Notes: N=2,446 in $t = 0$.

Figure 5. Automation cost level per worker for treated firms

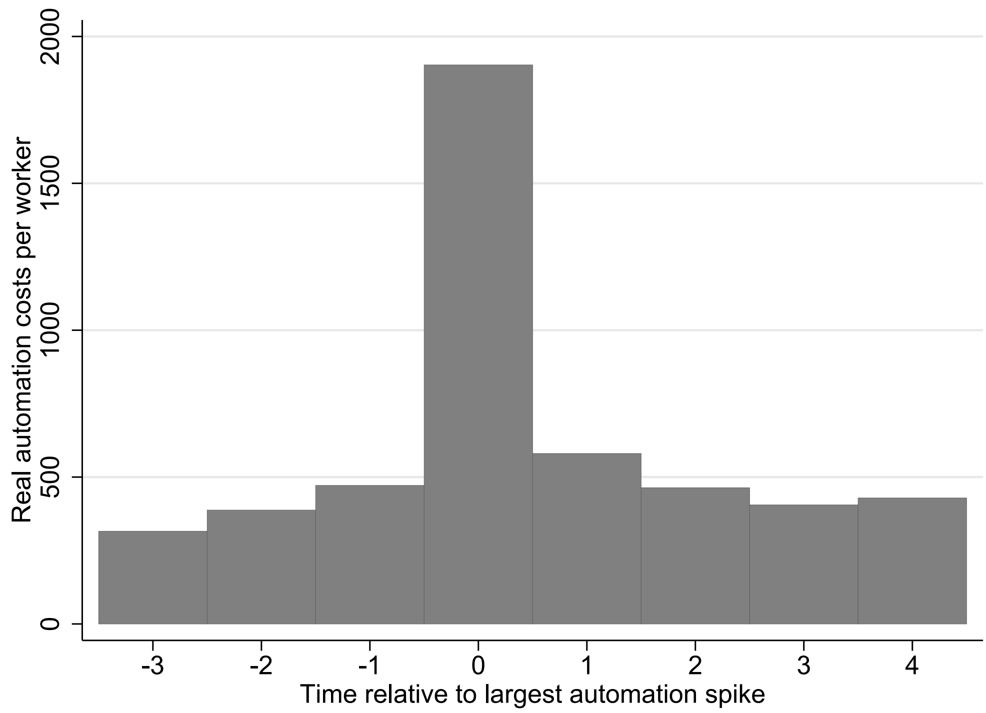

Notes: N=2,446 in $t = 0$.

Table 6. Automation expenditures by firm type

	Mean automation cost:		
Firm type	level	per worker	share (%)
No automation spike	245,066	1,389	0.62
≥1 Automation spike	359,797	2,547	1.29

Notes: Total N firms is 36,490.

Figure 6. Log employment for firms with and without automation events

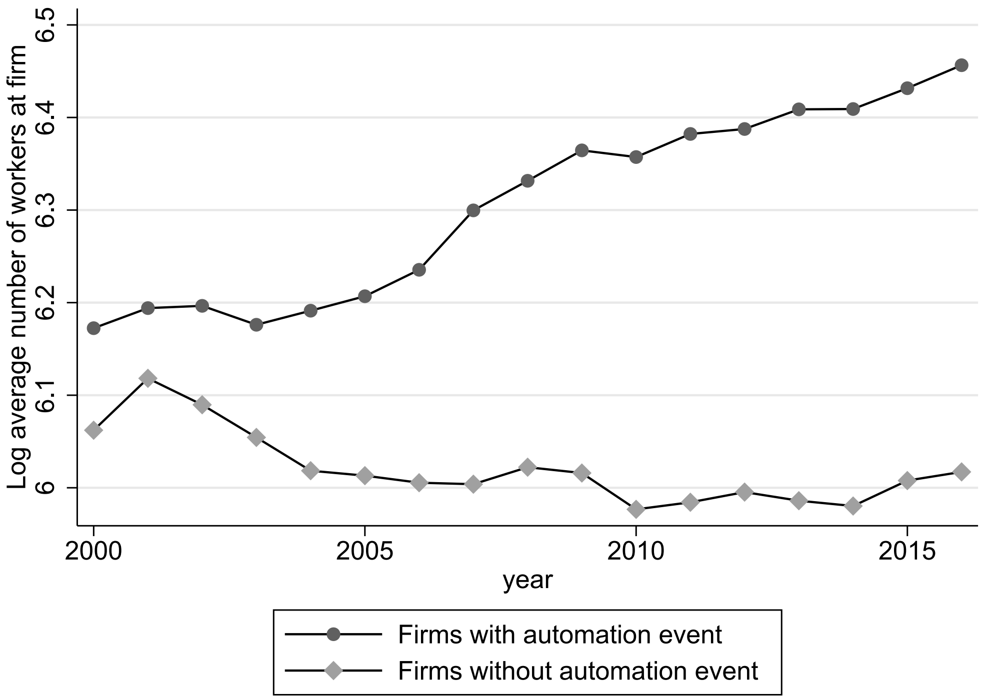

Notes: Balanced sample of 399 firms with and 623 firms without an automation event.

firms existing over the entire 17-year period. These stark descriptive differences in trajectories between automating and non-automating firms are consistent with findings in Koch et al. (2019), and in part motivate our empirical design, outlined in the next section.

3.4 Empirical design

We now outline our empirical design to leverage the observed automation cost spikes for identification. Our specification only considers incumbent workers who are employed in firms that spike at some point over 2000-2016. We define incumbent workers as workers with at least 3 years of firm tenure. This by and large captures workers with permanent contracts and hence workers who have a stable working relation with the firm.[11] This is important because identification

[11]Dutch labor law during almost our entire data period ensures that temporary contracts are of a maximum duration of 3 years, implying that workers with 3 years of tenure are very likely to have permanent contracts. On

Figure 7. Log wage bill for firms with and without automation events

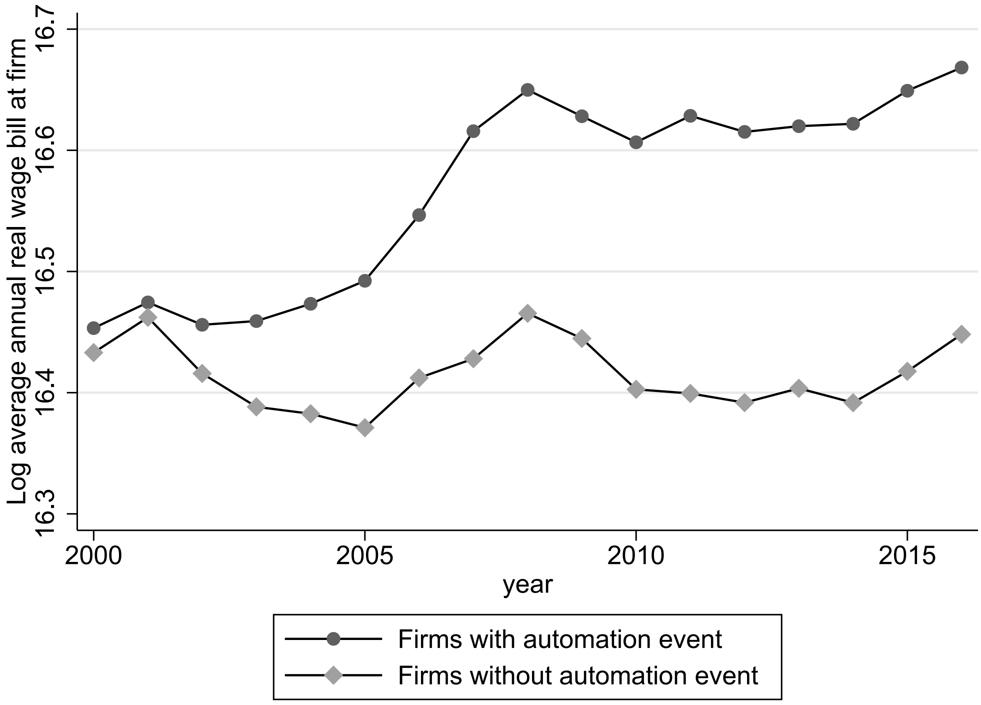

Notes: Balanced sample of 399 firms with and 623 firms without an automation event.

requires that workers are not self-selected into the firm in anticipation of an automation event occurring in the near future.[12] This reasoning is similar to the focus on incumbent workers in the mass lay-off literature (e.g. see Jacobson et al. 1993; Couch and Placzek 2010; Davis and Von Wachter 2011).

We define the group of treated workers as those with 3 or more years of firm tenure at $t-1$ in treatment group firms, i.e. firms that spike in $t = 0$ and are observed in all years $t \in [-3, 4]$. Treated workers are further divided into cohorts by the calendar year in which their firm spikes. Specifically, given that our sample covers calendar years 2000 to 2016, the earliest cohort of treated workers are those employed between 2000 and 2002 at a firm that spikes in 2003. Similarly, the last cohort of treated workers are those employed between 2008 and 2010 in firms that spike in 2011.[13]

For each cohort of treated workers, we then define a control group of workers with at least 3 years of firm tenure at $t-1$ and who are, at $t-1$, employed in firms that spike in $t+5$ or

average across firms in our data, 64 percent of workers are incumbents (where the median is 70 percent).

[12]In section 4.4, we also estimate impacts for the group of workers with less than three years of firm tenure prior to the automation event. Causal identification of the treatment effect for this group is 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.

[13]See Appendix A.1 for more details on sample construction.

later.[14] For example, the control group for the earliest cohort of treated workers are workers employed between 2000 and 2002 at a firm that spikes in 2008 or later. Similarly, the control group workers for the last cohort of treated workers are those employed between 2008 and 2010 at the same firm that spikes in 2016. Finally, we exclude both treatment and control group firms with multiple spikes in the estimation window such that estimates of pre-trends and treatment lags are not contaminated, but our results are similar when not imposing this restriction.

By defining treatment and control group workers from firms that spike at least once (i.e. excluding workers from firms that never spike in our control group), our specification strictly exploits differences in event timing rather than also using event incidence for identification. As such, we assume that from the perspective of incumbent workers, the timing of automation cost spikes is essentially random conditional on observables.[15] Another way to think about our approach is that we match workers on the firm-level outcome of making large investments in automation technology at some point in time. Only exploiting spike timing (rather than also spike incidence) across firms is important since in section 3.3 we showed that firms with automation events are on very different employment and wagebill trajectories: as such, the employment trajectories for workers employed at firms without these events are not an appropriate counterfactual.

Our use of timing differences across firms is in the spirit of a recent literature exploiting event timing differences in other contexts (see e.g. 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.[16] This sensitivity implies that small, idiosyncratic differences can change the exact timing of automation events across firms. Consequently, workers employed at cohorts of firms that spike a few years apart should be on similar trends, and can thus serve as a counterfactual.

[14] We only require control group workers to be at a firm j that spikes at $t+5$ or later to stay at firm j from $t=-3$ until $t=-1$. Hence, they do not have to be employed at firm j when firm j actually spikes in year $t+5$ or later.

[15] In Appendix A.2, we use a k-fold cross-validation prediction to show that spike timing is difficult to predict based on observables, increasing our confidence that event timing is plausibly random from the perspective of a firm's incumbent workers.

[16] For 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.

We use a Differences-in-Differences (DiD) specification for each cohort of treatment and control group workers, with the data stacked across cohorts:

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

where i indexes workers, j firms, and $t \in [-3, 4]$ the number of years relative to the timing of the automation spike.[17] 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. Further, I_t are indicators for time relative to the spike year, 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 firm; as well as fixed effects for years. In our baseline specification, we replace $\beta treat_i$ with individual fixed effects[18] – this also absorbs non-time varying controls (gender, nationality, firm size and sector). We cluster standard errors at the level of the firm where treatment occurs: that is, all workers employed at the same firm in $t - 1$ are one cluster.

In equation 2, the parameters of interest are δ_t: these estimate period t treatment effect for $t \geq 0$, relative to pre-treatment period $t = -1$. As with all DiD models, identification requires parallel trends in the absence of treatment, or that $\delta_t = 0$ for all $t < 0$. Our event timing strategy is intended to support the assumption that worker outcome variables would have followed similar trends in the absence of treatment.

We can further strengthen the assumption of parallel trends by matching on worker and firm observables to ensure that $\delta_t = 0$ for all $t < 0$ (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, 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

[17] Our results are robust to changes in the number of estimated post-treatment periods, which in our setting also changes the set of control group firms.

[18] Except when we estimate individual workers' hazard of leaving the firm.

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. This procedure results in 30,247 strata for incumbent workers[19], and in doing so can match 98 percent of treated incumbents (using 93 percent of control group incumbents).[20]

After matching, our sample contains $1,046,995$ distinct incumbent workers in treatment and control groups. Of those incumbent workers, $102,599$ are treated. Given we observe each of these individuals for 8 years, this results in $8,375,960$ observations. Our estimation sample of firms for identifying these treated and control group workers contains $5,970$ unique firms, all of which experience an automation spike at some point over the period and are observed for at least 8 consecutive years. Workers employed at $2,429$ of these firms are treated, and workers at $4,543$ firms serve as controls at least once.[21]

4 The impact of automation on incumbent workers

Here, we consider how incumbent workers are impacted by an automation event at their firm. In the first section, we study impacts on wage income and its components: changes in firm separation and non-employment on the one hand, and changes in daily wages on the other. The next subsection then considers adjustment margins for displaced workers (in particular, sectoral switching, early retirement, and self-employment) and to what extent income impacts are offset by various benefit payments. The third subsection performs a range of important robustness checks on our main results, and the final subsection considers effect heterogeneity in these impacts.

[19] Note that some strata may not contain any treated workers, in which case they are irrelevant for estimation. Further refinement of these strata does not change our results, although it leads to a smaller percentage of treated workers being matched.

[20] Further support for the parallel trends assumption is given in Table 18 in the Appendix. 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.3.3, we also show that our results are robust to additionally matching on pre-treatment employment growth at the firm level, incumbent worker firm tenure, and firm size.

[21] Firms can serve as control group firm more than once for different treatment group cohorts. For example, a firm spiking in 2016 might be a control group firm for a firm spiking in 2005 and for another firm spiking in 2008. In addition, some firms can serve as control group firm and as treatment group firm. For example, a firm spiking in 2010 could be a control group firm for a firm spiking in 2005, while also serving as a treatment group firm in 2010.

4.1 Impacts on wage income, firm separation, and non-employment

Figure 8 shows the impact of automation events on annual real wage income for incumbent workers, that is, workers with at least three years of firm tenure. We scale each individual worker's real wage income by their real wage income level one year before the automation spike, to obtain relative impacts.[22] That is, it shows estimates of equation 2, where the outcome variable is annual real wage income over $t = -1$ annual real wage income.

Figure 8 shows there are no pre-trends in wage earnings. The estimates highlight that incumbent workers lose income as the result of an automation event. Indeed, in the automation year, incumbent workers lose some 0.9 percent in wage income and this effect increases over time, cumulating to 10.7 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 323 euros lost for the average worker in the treatment year (0.9 percent of the average pre-treatment income of 35,885 euros)[23]; and 3,839 euros after 5 years in total (0.107 $\times$ 35,885). This suggests automation leads to displacement for workers: compared to workers employed at firms who automate later, workers employed at currently automating firms experience income losses.

To further scale these costs, we can calculate the incumbent worker income losses resulting per euro of automation expenditure per worker during an automation event. These estimates should be interpreted as an upper bound, since we do not necessarily observe all outlays associated with the automation event we identify. As reported above, during an automation event, automation expenditures are on average 1,912 euros per worker. Since the average worker loses 323 euros in the automation year, and 3,839 euros after 5 years in total, 1 euro invested per worker leads to a loss of around 0.30 euros of per incumbent worker in the automation year, and 2 euros after 5 years.

These annual wage impacts may be driven by changes in days worked following firm separation, changes in daily wages conditional on being employed, or a combination of both. We now 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?[24]

[22]This is preferable to log income impacts since this would eliminate zeros: this approach is also taken in e.g. Autor et al. (2014). In Figure 23 in the Appendix we additionally show impacts in levels.

[23]See Table 18 in the Appendix.

[24]Any adjustment in days worked can in principle come from either the intensive or extensive margin: that

Figure 8. Annual real wage income, relative to $t = -1$

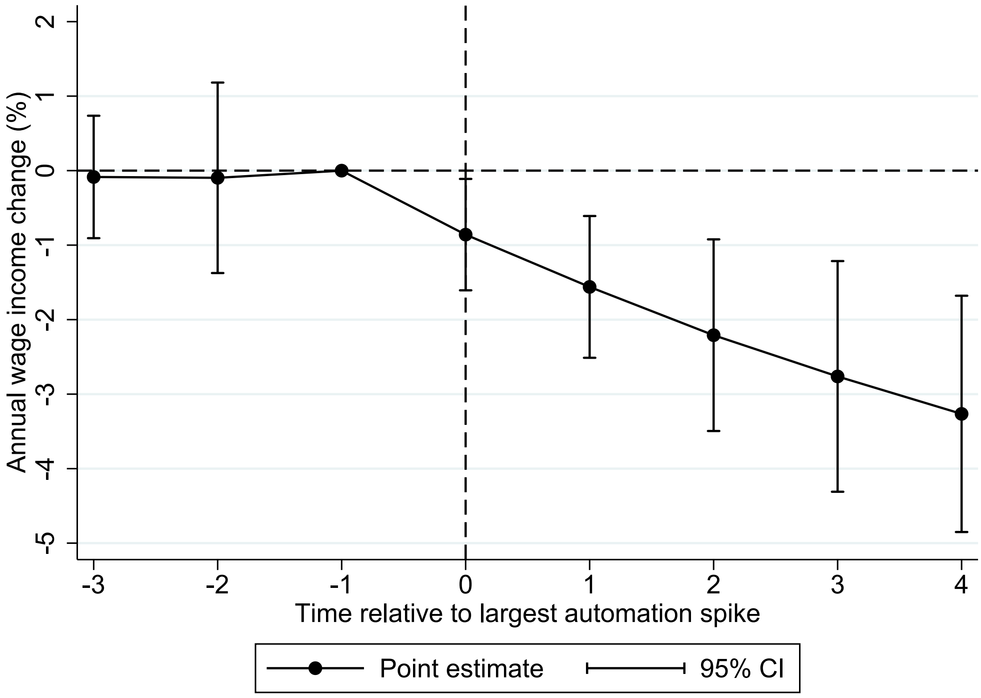

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

Figure 9 considers the impact of automation on displacement in its most literal sense: does automation result in incumbent workers leaving the firm? This figure presents estimates from equation 2, where the dependent variable is worker-level hazard of separating from their pre-treatment employer.[25] All coefficients have been multiplied by 100, such that the effects are in percentage points. This highlights that workers' probability of separating from their employers following an automation spike is rising over time compared to control group workers.

Specifically, in the automation year, the separation hazard for incumbent workers is 2.1 percentage points higher (though this estimate is statistically insignificant), where the (matched) control group incumbent separation probability is 9.6 percent. After five years, incumbents have a statistically significant 3.6 percentage point higher firm separation hazard: this is a substantial 41 percent rise compared to the average corresponding hazard among control group workers of 8.8 percent.

It is noteworthy that worker displacement does not occur instantly: rather, displacement

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.

[25]Because the dependent variable is a hazard rate, this model does not include worker fixed effects (unlike estimates for all other dependent variables).

effects arise over time. There are various (and non-mutually exclusive) possible explanations for this. For one, these patterns are consistent with incumbent workers having open-ended contracts, making it costly to fire them. Further, 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. Gradual displacement is in contrast to the patterns seen during mass lay-off events, where at least 30 percent of the firm's incumbent workforce is laid off at once (see Davis and Von Wachter (2011) for an overview). These gradual changes in firm separation are also what underlie the increasing effects for total wage income seen in Figure 8: estimates for any one post-event year reflect the impacts of all worker cohorts having left up until that point.[26]

All in all, this shows incumbents are more likely to separate from their employer as a result of automation. Indeed, the resulting increase in separation is substantially higher relative to what they experience in the absence of an automation event. Although this increase in firm separation implies automation leads to displacement for the firm's incumbent workers, it need not translate to income losses if these workers find re-employment quickly (and at similar wage rates): we now turn to impacts on non-employment.

Results are shown in Figure 10, 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 and may change their annual days worked, 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 1.3 days in the automation event year, and this increases to around 5.7 days annually five years after the automation event, with a total cumulative increase in non-employment of 18 days compared to the control group. By comparison, in the event year, matched control group incumbents spend around 5.7 days in non-employment on average, suggesting automation produces an increase of 22 percent in non-employment in the automation year itself. The cumulative five-year impact also corresponds to a 22 percent increase relative to the five-year cumulative non-employment duration (82 days) experienced by control group incumbents.

[26]Put differently, our estimates combine time and cohort effects: since workers are still leaving the automating firm at increased rates several years after the event, and assuming that displaced workers do not all adjust within the span of a year, the average treatment effect consists of cohort effects that cumulate over time.

Figure 9. Firm separation hazard

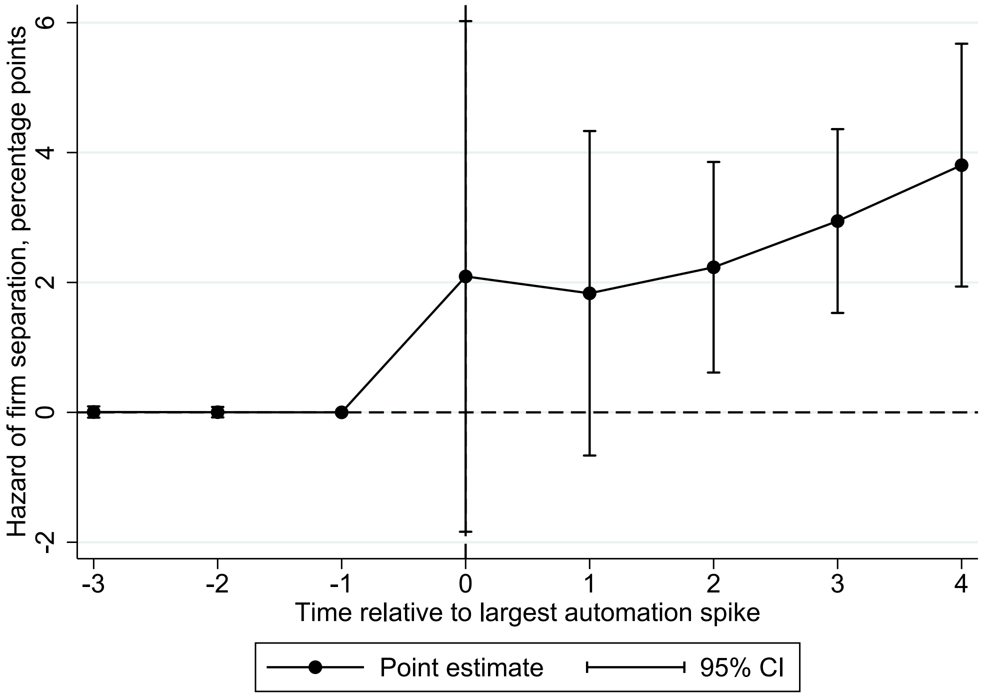

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals. Effects at $t = -2$ and $t = -3$ are zero by definition since incumbents are at the firm for three years before $t = 0$.

Figure 10. Annual number of days in non-employment

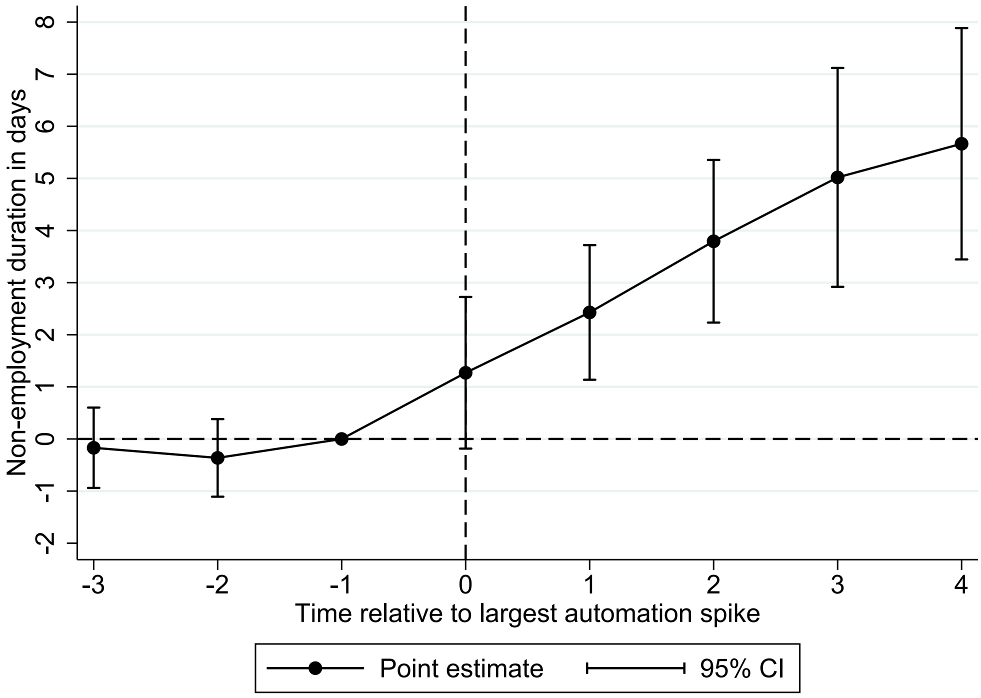

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

Figure 11. Log daily wages

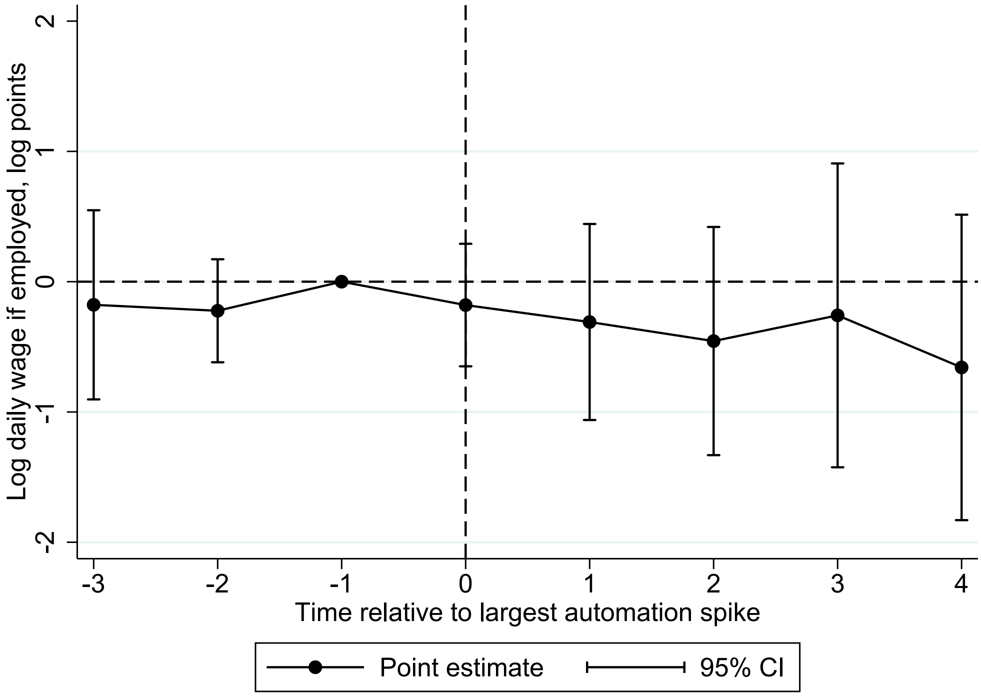

Notes: N=8,094,856. Whiskers represent 95 percent confidence intervals.

We now turn to the daily wage impacts of automation. In Figure 11, we consider the effect of automation on log daily wages (conditional on employment) of incumbents. 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 do not find any statistically significant effects, and the point estimates are also economically very small.[27] These findings are in contrast to long-term wage scarring found in mass lay-off studies: also in the Dutch context (and with the same administrative records we are using), such wage scarring has been found (see Deelen et al. 2018; Mooi-Reci and Ganzeboom 2015). This suggests the adjustment costs arising from automation events are more transitory, even for workers with relatively long firm tenure. This need of course not be because of the nature of the automation event, but could in part be because the workers affected by automation are different from those affected by firm closures, in ways that allow them to better deal with job transitions. The absence of daily wage effects also implies that the income losses incumbent workers experience are entirely accounted for by firm separation followed by non-employment spells.

[27]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. Also for these two groups separately, we find only economically small and statistically insignificant impacts.

Overall, the results in this section show that automation leads to displacement for individual incumbent workers. This of course does not imply automating firms are necessarily displacing workers on net. In Section 3.3, we already documented that automating firms expand employment more rapidly than non-automating firms. While our empirical design cannot claim causal identification of automation's employment effects at the firm level (as the automation event and its timing are clearly endogenous from the firm's perspective), in a forthcoming paper we show descriptively that firms do appear to be labor-saving on net around automation events (Bessen et al. 2020). This suggests that labor-saving automation is one way firms may gain a competitive advantage that allows them to expand the size of their operations – including employment– in the longer run. However, these events are accompanied by adjustment costs borne by incumbent workers.

4.2 Where do automation-affected workers go?

So far, we have shown that incumbent workers experience income losses as a result of automation in their firm, cumulating to 3,839 euros on average per incumbent worker in total over the five-year post-treatment window. These losses are entirely driven by higher firm-separation probabilities accompanied by non-employment spells. This raises the question where workers affected by automation events go: in this section, we consider a range of outcomes and adjustment mechanisms. First, we will consider to what extent workers impacted by automation are switching to firms in other industries, and to firms of different sizes and different average wages. Next, we study to what extent various benefit systems are compensating for lost wage income. Lastly, we consider whether treated workers are more likely to be observed in early retirement or self-employment compared to the control group.

Figure 12 estimates our empirical specification with the probability of switching two-digit industries as the dependent variable. This shows that workers impacted by an automation event are 5 percentage points more likely to switch industries after five years. While the estimates are relatively noisy, this is a 17 percent increase compared to the 30 percent probability of switching industry for control group workers over the same time period.[28] That is, the additional displacement produced by automation is translating to some increased industry switching.

Besides industry switches, we do not find economically sizable or statistically significant

[28]As shown in the figure results are very similar for one-digit industries, though not statistically significant at the 5 percent level.

Figure 12. Probability of switching industries

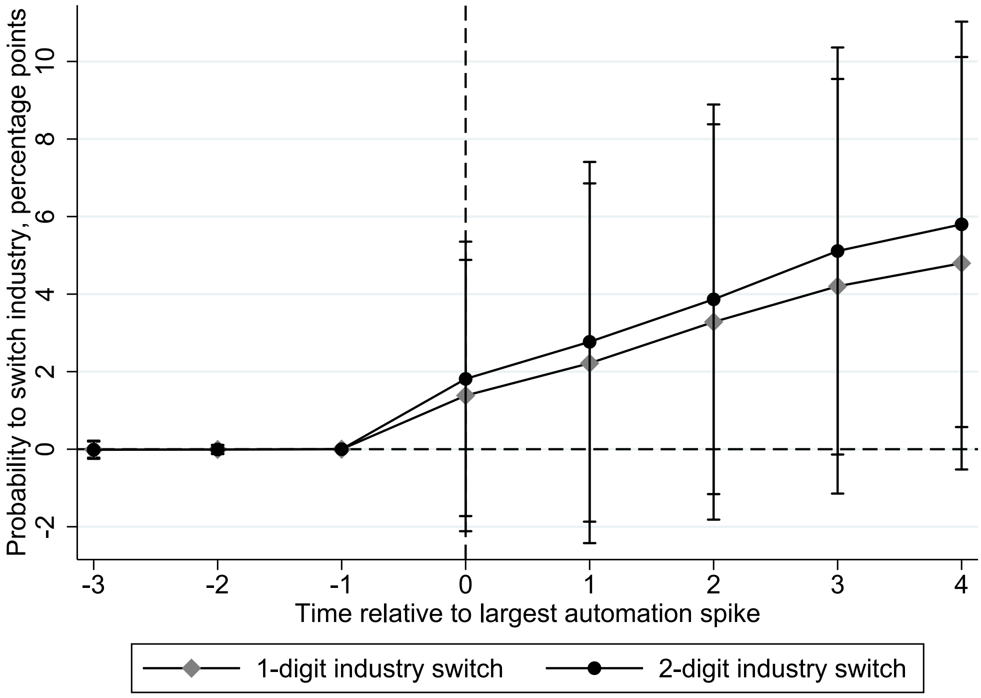

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

changes in terms of workers' average or median firm wage, firm size, or firm automation expenditure.[29] This implies that automation-affected workers are not structurally moving to firms that pay different wages, are of different sizes, or are differently automation-intense.

Figure 13 considers the impact of automation on workers' total benefit receipts (in annual real euros), comprised of unemployment benefits, disability benefits, and welfare payments, as well as the separate contributions from these three different sources.

We find that incumbent workers receive additional benefit income following an automation event: after five years, the cumulative amount received is around 514 euros on average, implying that only 13 percent of the negative wage income impact is offset. 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).[30]

Figure 13 further shows that all of the benefit payments for incumbents arise from unemploy-

[29]For each of these (firm wage, firm size, and firm automation expenditure), we measure the dependent variable in $t = -1$, to keep overall changes in firm characteristics from impacting the estimate. As such, for workers remaining with their pre-event firm, or those switching to a firm that had the same size, wage, or automation expenditure in the pre-treatment year, the change in the dependent variable is zero.

[30]In part, this is by law: unemployment benefits in the Netherlands have a replacement rate of 75 percent in the first two months of unemployment, which then decreases to 70 percent. Further, there is a maximum ceiling, such that workers with higher wages earn lower replacement rates than the 70 or 75 percent maximum.

Figure 13. Annual real benefit income

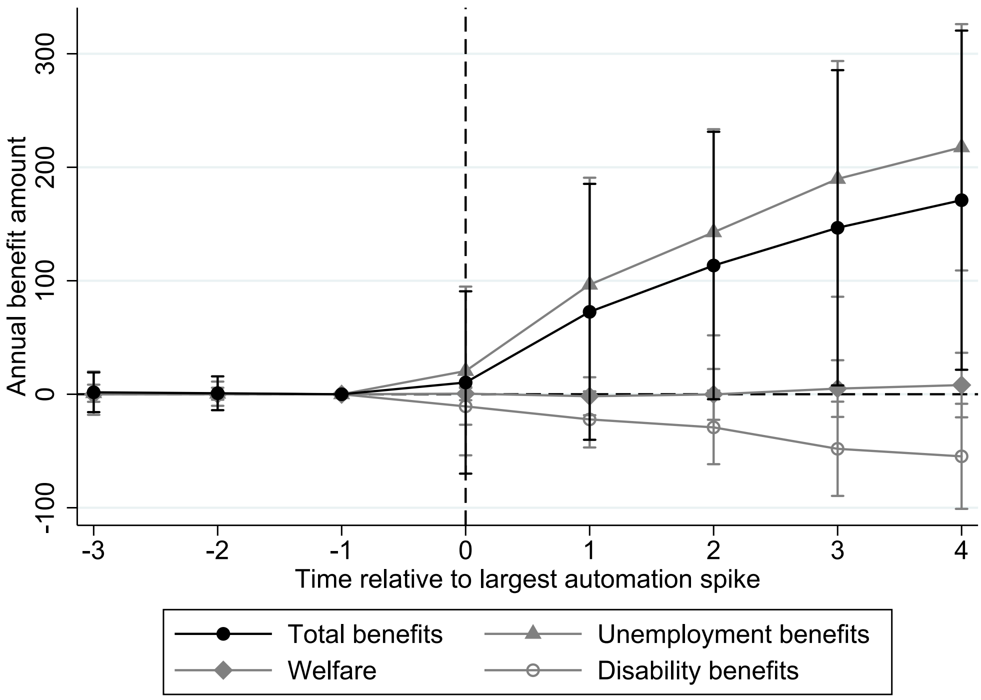

Notes: N=8,375,960. Total benefits are the sum of received unemployment and disability benefits, and welfare. Whiskers represent 95 percent confidence intervals.

ment insurance: this is expected, as unemployment benefit eligibility is very high among workers with at least three years of firm tenure.[31] Consistent with high unemployment benefit eligibility, we do not see any contribution from welfare payments for incumbents. Lastly, disability benefits[32] 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. 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 14 shows that treated incumbents are somewhat more likely to enter self-employment following an automation event, although the effect is very small – 0.3 percentage points cumulated over five years (a 6 percent increase relative

[31]For most of the observation period, eligible workers in the Netherlands are entitled to up to 38 months of unemployment benefits following job loss. In the last year of observation (2016), this has been decreasing: currently, eligibility is 24 months.

[32]Disability 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 14. Cumulative probability of entering self-employment or early retirement

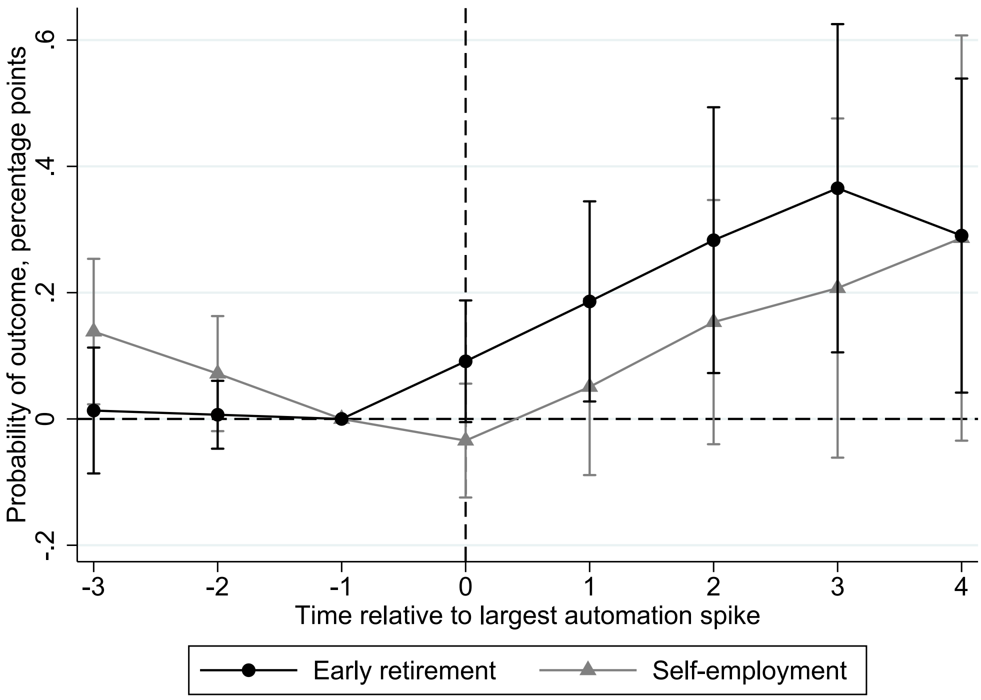

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

to a five-year probability of 5.2 percent among the control group). This means self-employment is unlikely to 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, as shown in Figure 14. In particular, five years after the automation event, treated incumbent workers are 0.3 percentage points more likely to be observed in early retirement. While these effects are similar in absolute size to those for self-employment, the average five-year probability of early retirement among control-group incumbents is much lower, around 1.7 percent. As such, the treatment effect represents an 18 percent increase in the incidence of early retirement.[33]

Taken together, these results show that automation-impacted workers are more likely to switch industries, but do not move to firms with different characteristics in terms of size, wages, or automation intensity. However, the documented income losses experienced by these workers are only partially offset by benefit systems and other income sources, implying automation-affected workers largely bear the adjustment cost themselves.

[33]However, early retirement does not account for all wage income losses: when we estimate our models separately for workers under the age of 55, we still find statistically significant relative wage income losses of the same size as in our full sample.

4.3 Robustness checks

Having laid out our main results, we show that these findings are robust to a number of alternative model specifications and other checks.

First, we subject our results to a randomization test as first introduced by Fisher (1935).[34] To do this, we take our sample of 36,490 firms, randomly draw firms with replacement, and then for each of these firms randomly assign a year to have a placebo automation event.[35] We then construct treated and control firms based on these placebo events. We repeat this procedure 100 times, where each permutation sample contains the same number of treated and control firms we have in our actual estimation sample.

Results are shown in Figure 15: each gray line presents a set of placebo (dynamic) treatment estimates, whereas the black line presents our actual treatment estimates. The graph also shows probability values calculated using the rank of the absolute value of our estimated coefficient among the 100 permutated estimates.[36] Something at least as extreme as our treatment estimate is unlikely to occur by chance: from the first post-event year onwards, the probability is very close to zero. Permutation estimates for two of our other three outcome variables also reject the null hypothesis: the hazard of leaving the firm and days in non-employment.[37] For the log daily wage impact, on the other hand, the randomization test shows that our point estimate is likely to occur when assigning automation spikes at random – this is as expected, since we did not find a statistically significant impact of the automation event on daily wages conditional on employment. All in all, this increases confidence that our estimates are not a statistical false positive.

However, even if our results are not occurring by chance, they may be driven by other real firm-level events that correlate with automation. Such events may impact labor demand at the firm level and thereby affect individual incumbent workers. Note that in our baseline specification, we do not see any pre-trends at the worker level, but the parallel trends assumption may of course still fail in the post-treatment period if such events coincide with the automation spike. We address this concern in three main ways.

First, we match additionally workers on their firms' pre-treatment employment trends. This

[34] Also see Kennedy (1995) for an overview and Young (2018) for a recent application and evaluation of the value of these tests.

[35] Note that this permutates both the assignment of treatment to firms, and their timing across years, since both are part of our empirical procedure.

[36] Results are very similar when using t-statistics rather than coefficient estimates to calculate probability values.

[37] See Appendix A.3.

implies we now ensure that treated and control workers are not only employed at firms that experience an automation event at some point in time, but where pre-treatment employment growth is similar. Second, our data include administrative information on some of these events, namely mergers, take-overs, acquisitions, firm splits, and restructuring.[38] As a second robustness check, we eliminate firms that experience such events anywhere in the estimation window. Third, we remove outlier firms in terms of employment changes (those experiencing an employment change exceeding 90 percent in any one year), both in the estimation window and outside of it.[39] The removal of these outliers is intended to capture any firm-level events which are not formally documented in our administrative records. Last, we remove firms where there was a new worker among the firm's top-decile annual wage income earners[40] in the three years prior to the automation event. This is intended to capture automation events coinciding with managerial change, which may bring changes in personnel policy unrelated to automation.

Figure 16 summarizes the results for all three robustness checks pertaining to firm-level events (along with baseline model estimates). Estimates are very similar across the board, though effects are somewhat smaller when eliminating firms with (suspected) management change: this suggests that automation may sometimes be the result of a new manager changing business practices. Overall, however, our findings are very robust, showing that firm-level events other than automation are unlikely to be the driving force behind the worker impacts we find.

We have performed further robustness checks.[41] Specifically, we change our empirical specification by removing individual fixed effects (and replacing them with worker- and firm-level control variables), or additionally matching incumbent workers on firm size and on firm tenure (i.e. within the three year minimum firm incumbency requirement). We also consider various alternative spike definitions – including spikes in automation per worker rather than in total costs, and when measuring average costs only in the pre-event period. This all leads to very similar results. Lastly, we show that results are robust to varying the spike threshold from two to four times the average automation costs (our baseline is thrice the average automation costs). Estimated effect sizes are somewhat larger for higher compared to lower thresholds, as expected, but these differences are not statistically significant. This shows that our results are not driven by the specific spike size cut-off we employ in our baseline estimates.

[38]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.

[39]Results are similar when only removing firms with outliers inside the estimation window.

[40]Conditional on this worker earning at least 150 real euros a day, i.e. 40.000 euros a year.

[41]These are reported in Appendix A.3.

Figure 15. A randomization test for relative wage income estimates

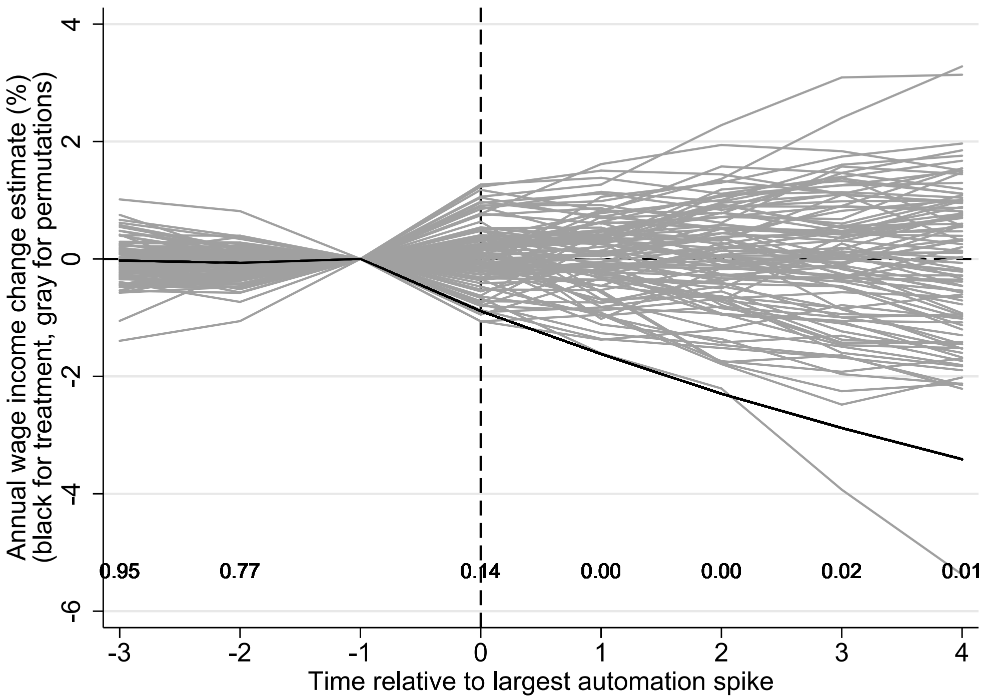

Notes: 100 permutations. The numbers printed at the bottom of the graph are probability values for the treatment estimates, based on the randomization test.

Figure 16. Robustness to removing other firm-level events

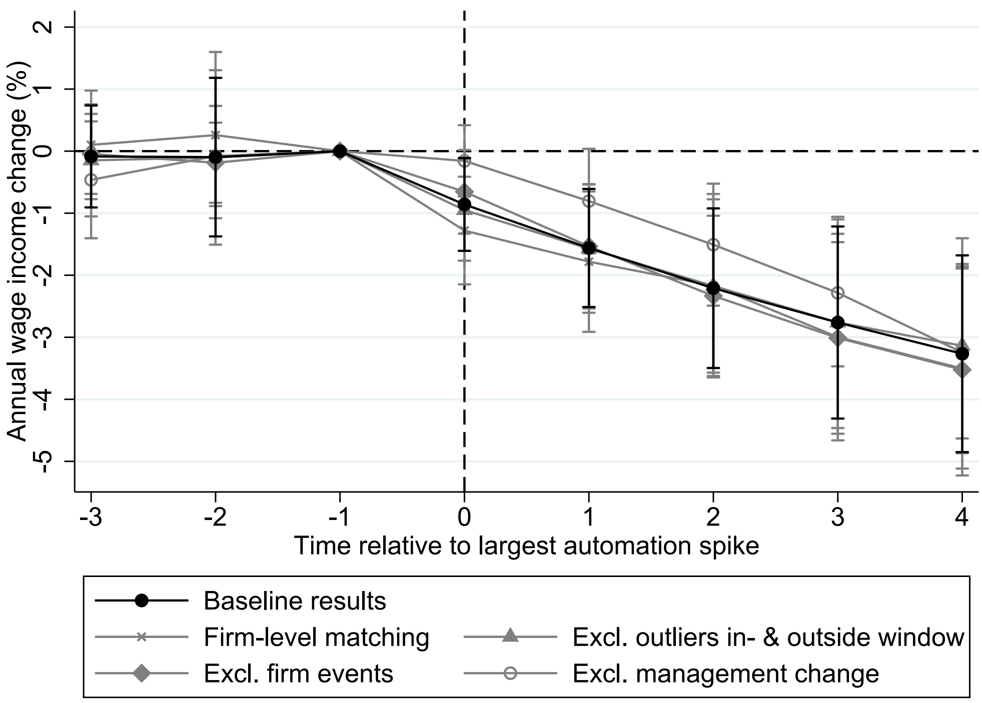

4.4 Effect heterogeneity

So far, we have shown income losses for incumbent workers driven by non-employment spells, only partially compensated by benefit payments. Here, we investigate further how different types of workers are affected: in particular, we consider effect hetereogeneity by incumbent age, gender, firm size, and sector of employment; and also separately study workers with less firm tenure than incumbents. Although our data lack a direct skill measure, we consider how impacts differ by age-specific wage quartiles, both overall and within firms. For succinctness, we will only show estimates for relative annual wage income, as this is the summary measure capturing all other impacts. Any noteworthy differences in results for firm separation, non-employment duration, log daily wages, and early retirement are described where relevant.

4.4.1 Incumbent worker characteristics

Here, we consider how incumbent workers with different characteristics fare after an automation event. 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 section 4.1. In particular, we estimate the following model:

$$y_{ijt} = \alpha + \beta D_i + \gamma post_t + \delta_0 \times treat_i \times post_{it} + \sum_k \left[\delta_k \times treat_i \times post_t \times z_{ki}\right] + \lambda X_{ijt} + \varepsilon_{ijt}, \quad (3)$$

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, $post_t$ 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 $k+1$ categories. In addition to the controls included in equation 2, X_{ijt} also contains z_{ki} as well as the interaction terms $z_{ki} \times treat_i$ and $z_{ki} \times post_i$. 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. βD_i capture worker fixed effects, and standard errors are clustered at the treatment level as before.

Table 7 summarizes how average post-treatment effects for annual wage income differ across workers of different ages, gender, and their (initial) firms' sector and employment size. First, we find that workers over the age of 50 are most negatively affected by automation events: while differences with younger age groups are not always statistically significant, the point estimates suggest all other groups experience somewhat smaller income losses. This is not because older workers leave the automating firm at higher rates, but rather, because they experience larger increases in non-employment duration. Unsurprisingly (and not reported here), the early retirement effects we found are entirely driven by the oldest workers. Taken together, older workers appear to face higher adjustment costs from automation than do younger ones.

We do not find any statistically significant differences in impact by gender and firm size. However, we may be underpowered in detecting differences across these groups. It should also be noted that effect heterogeneity across firm size (reported in column 3) is important for our purposes 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. For this reason, it is reassuring that displacement effects are found across the firm size distribution. If anything, losses are somewhat higher (though not statistically significantly so) for workers employed at smaller firms, which implies we would probably find somewhat higher average wage losses from automation if our data were more representative in terms of firm size. Lastly, although not reported here, we find that firm separation increases as a result of automation across all firm sizes, but most strongly so for the largest firms – the fact that this does not translate to larger wage losses for these workers suggests they have better outside options.

In column 4, we consider to what extent the impacts of automation differ depending on which sector the worker's firm belongs to: that is, our treatment effect is interacted with workers' sector of employment in $t = -1$. For this model, Manufacturing is the reference category. Note that sectoral differences may exist for various reasons. First, 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. Second, the workers employed in these different industries may have different characteristics (including unobservable ones), making the impacts differ. Third, to the extent that skills are industry-specific, sectoral labor market conditions matter: displacement would be more costly 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.

Our finding here is that automation leads to wage income losses that are quite pervasive

Table 7. Relative wage income effects by incumbents' characteristics

(1) Age		(3) Gender	
Age 50+ (ref)	-3.04***	Male (ref)	-1.52***
	(1.15)		(0.57)
Deviations from reference group for:		*Deviations from reference group for:*	
Age <30	1.20	Female	-1.39
	(3.94)		(0.97)
Age 30-39	0.96	(4) Sector	
	(0.93)	Manufacturing (ref)	-1.98**
Age 40-49	1.61*		(0.99)
	(0.92)	*Deviations from reference group for:*	
(2) Firm size		Construction	1.05
500+ employees (ref)	-1.53		(1.73)
	(1.35)	Wholesale & retail trade	-2.23
Deviations from reference group for:			(1.51)
200-499 employees	1.21	Transportation & storage	0.71
	(1.77)		(1.79)
100-199 employees	-2.19	Accommodation & food serving	4.57**
	(1.77)		(2.32)
50-99 employees	0.17	Information and communication	-0.25
	(1.57)		(1.76)
20-49 employees	-2.18	Prof'l, scientific, & techn'l act's	-0.24
	(1.46)		(1.80)
1-19 employees	-2.06	Administrative & support act's	1.55
	(1.52)		(2.01)

Notes: Estimates from four separate models, N=8,375,960 for each model. All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$); coefficients have been multiplied by 100. *p<0.10, **p<0.05, ***p<0.01.

across sectors: this highlights that robotics is likely not the only automation technology displacing workers from their jobs. The exception is Accommodation and food serving, where no income losses (nor increases in firm separation) are detected. However, Accommodation and food serving is also a sector with one of the lowest automation expenditures per worker (and the lowest number of automation events, see Table 5), as well as contributing only 2 percent of the sample of incumbent workers.[42] On the other hand, incumbent workers in Wholesale and retail and Manufacturing do experience earnings losses – together, these two sectors employ almost half of all incumbents in our sample (26 and 23 percent, respectively). We find that automation leads to increased firm separation rates for all sectors except Accommodation and food serving and Construction.[43] All in all, we find that automation events originating in different sectors have qualitatively similar impacts on workers.

Unfortunately, our data do not contain any occupation information, and only contain edu-

[42]See Table 18 in the Appendix.

[43]In a separate analysis, we find that the overall differences found between incumbents and recent hires in the next section are not only due to their different sectoral or firm size affiliations.

cation information for a small and selected subsample of workers. Instead, we obtain a measure of workers' skill level by calculating each worker's wage rank by age in $t = -1$. We then group workers into quartiles based on this rank. For example, the top-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.[44]

Results are reported in the first column of Table 8: workers in the highest age-specific wage quartile are used as the reference category. We do not detect any statistically significant differences: that is, workers across all wage quartiles experience displacement from automation. Indeed, the point estimates for deviations from effects for the top quartile are quite small.

The similarity of losses across the wage distribution may of course be partially driven by differences in the firms where automation spikes occur: lower losses for one "skill" group may be offset by higher exposure to automation events in our sample. While the estimates in column 1 matter for the average worker's exposure to displacement from automation, we are also interested in which workers are displaced within firms. Therefore, the second column in Table 8 reports estimates by workers' age-specific *within-firm* wage quartile. That is, the bottom quartile reflects incumbents who are in the lowest 25 percent of their firm's wage distribution for their age.[45] If anything, this reveals that the highest-paid workers by age *within* firms appear to lose more wage income than do lower quartiles, although these differences are not statistically significant.

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 8 does provide an important insight that is counter to a common intuition: there is no evidence that workers lower down the wage rank for their age ("lower-skilled" workers) are displaced more often by automation events.[46] Since our approach captures a wide range of automation technology, this could be consistent with Webb (2019), who uses patent data to show that while low-skilled workers are most exposed to robotics, other automation technologies such as software and artificial intelligence impact more on work performed by medium- and high-skilled workers.

[44]As an alternative skill measure, we calculate residual wage quartiles (by first regressing worker wages in $t = -1$ onto a set of observables and their interactions): results (not reported here) are very similar.

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

[46]These findings are confirmed when we estimate our models for the (small and selected) subset of observations where education level data is available.

Table 8. Relative wage income effects by incumbents' wage quartile

(1) Overall age-specific wage quartile		(2) Within-firm age-specific wage quartile	
Top quartile (ref)	-2.17**	Top quartile (ref)	-2.43*
	(1.06)		(1.27)
Deviations from reference group for:		*Deviations from reference group for:*	
Third quartile	0.39	Third quartile	0.79
	(0.84)		(0.86)
Second quartile	0.09	Second quartile	0.33
	(1.07)		(1.10)
Bottom quartile	-0.09	Bottom quartile	1.69
	(1.65)		(2.20)

Notes: The two models are estimated separately. 8,375,960 observations for column (1); 5,894,240 observations for column (2). All coefficients are average annual effects over the post-treatment period ($t = 0$ through $t = 4$); coefficients have been multiplied by 100. *p<0.10, **p<0.05, ***p<0.01.

4.4.2 Worker tenure: incumbents versus recent hires

Our identification strategy for the impacts of automation is to consider individual workers who have a pre-existing working relationship with the firm, as evidenced by at least three years of firm tenure. We now turn to estimating our models for a second group of workers: 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. However, 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.

On average, recent hires earn lower wages and spend a higher number of days in non-employment compared to incumbents.[47] 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.

We now estimate equation 2 for recent hires in the same way we have for incumbents, while additionally creating a zero income bin when matching on pre-event income, and matching individual workers on pre-event trends in non-employment duration.[48] After matching, our

[47] Table 17 in the Appendix shows summary statistics both incumbent and recently hired workers, showing averages and standard deviations across the balanced panel of workers and years.

[48] In particular, we estimate a linear trend in non-employment duration for individual recent hires before treatment, and match treated and control group recent hires using four bins of this trend: up to the 10th percentile,

sample contains $404,796$ unique recent hires ($78,282$ of whom are treated): given our observation window of 8 years ($t = -3$ through $t = 4$) this results in $3,238,368$ observations.

Unlike for incumbent workers, we find no income losses from automation for recent hires, as shown in Figure 17. Relative to recent hires in the control group, point estimates are positive[49] but never statistically significant – hence, recent hires do not have different annual wage earnings as a result of automation. 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) and Lefranc (2003), who find that income decreases following displacement result mostly from the loss of returns to accumulated firm-specific human capital. However, 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. Consistent with new hires being better matched (or able to adjust) to their firms' new technologies, we do not find any statistically significant increase in firm separation for these workers, and differences in non-employment duration with the control group are very close to zero over the entire pre- and post-treatment period.

Taken together, these results in this section show that although we detect some effect heterogeneity, our findings for incumbents are not driven by workers in a small subset of sectors, firm sizes, or age groups. Further, the income losses found for incumbent workers are not seen among recent hires, and automation affects incumbent workers from all ranks of the "skill" distribution.

the 10th percentile to the median, the median to the 90th percentile and higher than the 90th percentile. Together with the other matching variables, we obtain 82,942 strata for recent hires, and can match 95 percent of treated recent hires (using 65 percent of control group recent hires).

[49]The point estimate suggests recent hires gain 4.5 percent of an annual income in total over five years following an automation event (corresponding to around 1,032 euros from a pre-treatment average income of 22,944 euros, see Table 18 in the Appendix).

Figure 17. Relative annual wage income effects for incumbents versus recent hires

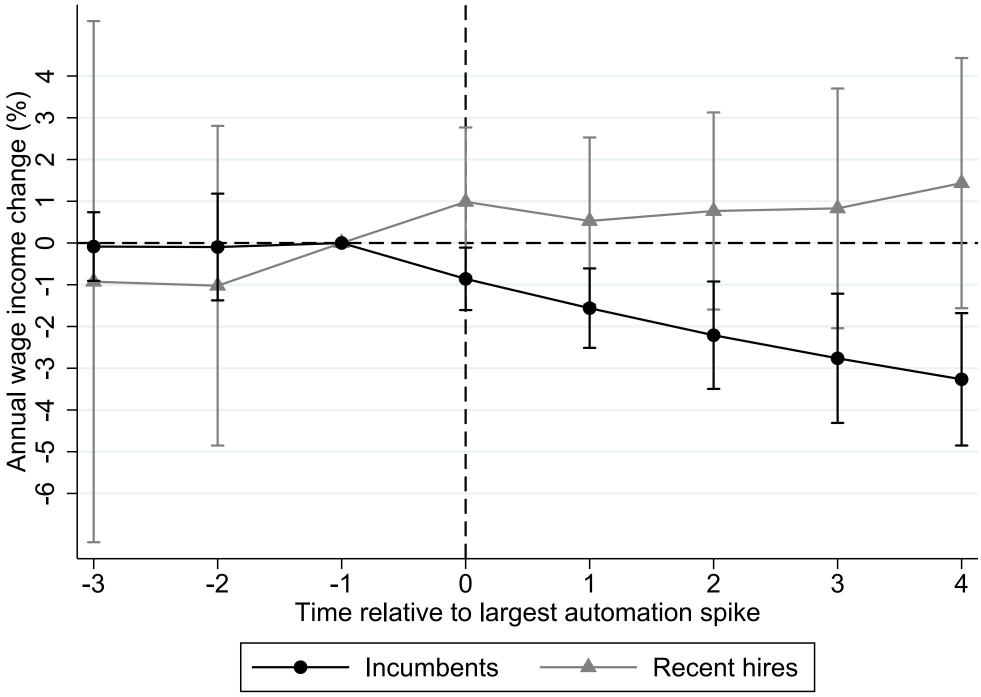

Notes: N=3,238,368 for recent hires and N=8,375,960 for incumbents. Whiskers represent 95 percent confidence intervals.

5 Comparison to computerization

We have found that automation displaces incumbent workers: this raises the question whether this effect is specific to automation technologies or occurs with investment in new technology more generally. This question is also relevant from the perspective of recent theoretical frameworks which distinguish labor-replacing technologies from labor-augmenting ones (e.g. see Acemoglu and Autor 2011; Acemoglu and Restrepo 2018d): automation is typically modeled as being labor-replacing in nature.

While we do not have a specific technology inventory at the firm level, Statistics Netherlands conducts a separate and partially overlapping firm survey on investments, including computer investments.[50] This item is called ‘computers’ or ‘computers and other hardware’ (depending on the year) and consistently defined as follows: “All data-processing electronic equipment insofar as they can be freely programmed by the user, including all supporting appliances. Do not include software.”. All investment within the company counts towards the expenditures, also if the equipment is second-hand, or leased or rented, or produced within the company. It excludes

[50]Investments in software and in communication equipment are only measured from 2012 onwards, so we only consider computer investments. In 2012, software investments are of a similar magnitude as computer investments.

investments in plants that are located abroad or resulting from take-overs of other organizations whose operations are continued without change.

In this section, we analyze the effects of computer investments in a similar way to that of automation investment, and directly contrast it to the impacts of automation in the part of the sample where we have overlapping data. This serves two purposes. First, as outlined above, we can consider to what extent spikes in automation costs have different effects on workers than do spikes in computer investment. Second, since automation cost and computer investments are somewhat correlated at the firm level, we can remove firms which have computer investment spikes within our estimation window to rule out that our automation event is partially capturing investment in computers. Conversely, we will also estimate the effects of computer investments in isolation, that is, excluding any events where automation spikes occur within the estimation window.

We first show some summary statistics on computer investments (section 5.1) before turning to the comparison between automation and computerization (section 5.2). Throughout, we consider the overlapping sample of firms where we observe both automation cost and computer investment data. This means our dataset has a smaller number of observations, and is more skewed towards larger firms as these are most likely to be sampled in both surveys. However, our results are qualitatively identical in the full samples for both automation and computerization.

5.1 Summary statistics

Here, we show summary statistics on both computer investments and automation costs for the overlapping sample of firms where we observe both: this allows for the most direct comparison.

Table 9 informs on the distribution of automation costs and computer investment across firms and years. Automation costs are higher than computer investments across the distribution, both in total and per worker. Of course, it should be noted that both can come with other unmeasured correlated costs, such as software for computers, and machinery for automation.

Further, Tables 10 and 11 compare automation and computer investments across firms of different sizes and sectors. In both tables, we also report the ratio of observed automation to computer expenditures per worker. As expected, Information and communication has the highest computer investment per worker, followed by Professional, scientific, and technical activities. Accomodation and food serving and Construction have the lowest computer investment per worker. When considering the relative importance of automation and computer technology,

Table 9. Computer and automation cost share distributions

	Automation cost		**Computer investment**	
	level	*per worker*	*level*	*per worker*
p5	0	0	0	0
p10	0	0	0	0
p25	0	0	0	0
p50	16,747	297	5,554	99
p75	69,617	957	31,042	447
p90	241,274	2,175	112,889	1,126
p95	568,915	3,518	250,652	1,868
mean	249,275	1,032	99,666	559
mean excl. zeros	346,396	1,434	155,619	873
N firms × yrs	171,549		171,549	
N firms × yrs with 0 costs	48,098		61,680	

Notes: All numbers are in 2010 euros. The number of observations is the number of firms times the number of years.

Table 10. Automation costs and computer investments by sector

Sector	**Autom. cost per worker**	**Comp. inv. per worker**	**Ratio autom. to comp.**	**Nr of obs** *Firms*	*Firms × yrs*
Manufacturing	998	369	2.7	5,153	40,773
Construction	497	215	2.3	2,821	18,319
Wholesale & retail trade	1,152	544	2.1	7,220	50,381
Transportation & storage	917	456	2.0	2,279	15,834
Accommodation & food serving	256	151	1.7	742	4,462
Information & communication	2,030	2,420	0.8	1,562	9,756
Prof'l, scientific, & techn'l activities	1,272	772	1.6	2,345	14,708
Admin & support activities	863	388	2.2	2,914	17,316

Notes: Overlapping sample.

Manufacturing is the most automation-intense compared to other sectors, whereas Information and communication is the most computer-intense. These patterns are reassuring. Like for automation, we generally see higher computer investment per worker for larger than smaller firms, but the pattern is less dramatic: this is reflected by the ratio of automation to computer expenditures rising with firm size.

Lastly, Figure 18 plots quantiles of the distribution of computer investments per worker over time. Unlike for automation costs, investments per worker are declining initially (perhaps reflecting the aftermath of the dot-com bubble) and relatively flat thereafter. Of course, effective computing power per worker is likely to have grown: computer investments have been deflated by the overall price index, which is unlikely to capture quality improvements in computing equipment.

Table 11. Automation costs and computer investments by firm size

Firm size	Autom. cost per worker	Comp. inv. per worker	Ratio autom. to comp.	Nr of obs *Firms*	*Firms × yrs*
1-19 employees	2,233	1,091	2.0	2,267	11,326
20-49 employees	851	543	1.6	10,451	66,339
50-99 employees	838	456	1.8	5,804	41,460
100-199 employees	944	500	1.9	3,418	26,466
200-499 employees	1,204	569	2.1	1,929	16,202
≥500 employees	1,640	637	2.6	1,167	9,756

Notes: Overlapping sample.

Figure 18. Firm-level computer investment per worker over time

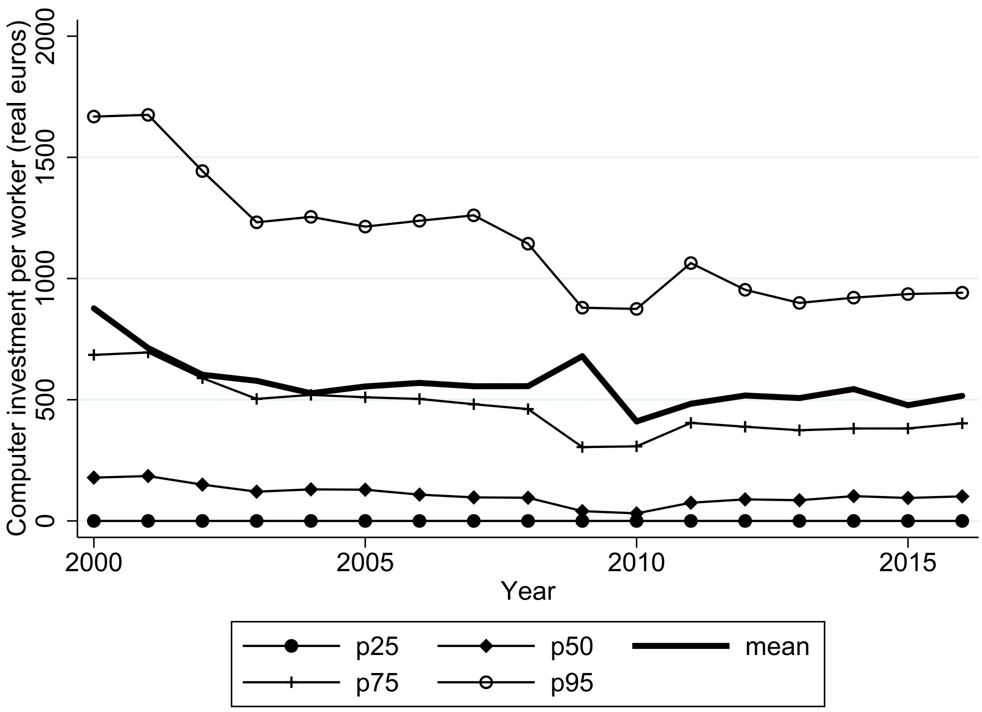

5.2 Automation versus computerization

In order to compare automation to computerization, we construct computer investment events in the same way we have for automation, but using computer investment per worker.[51] We use the same threshold, assigning firms a computer investment spike if their computer investment per worker exceeds three times their usual level.

The resulting distribution of computerization events is reported in Table 12. Compared to automation events, computerization events are more frequent. However, Figure 19 shows that

[51]This is necessary because, unlike automation expenditures, computer investments are not part of total costs; and because total investments are inconsistently defined over our sample period. In Appendix A.3.5 we also define automation spikes based on outlays per worker (rather than based on cost shares), and find very similarly sized effects.

Table 12. Automation costs and computer investments by firm size

	Percentage of firms with event type:	
Nr of events	Automation	Computerization
0	71.8	47.9
1	22.5	41.9
2	4.8	9.1
3	0.7	1.1
4	0.1	0.1

Notes: Overlapping sample, N=25,036 firms.

Figure 19. Computer investment per worker for treated firms

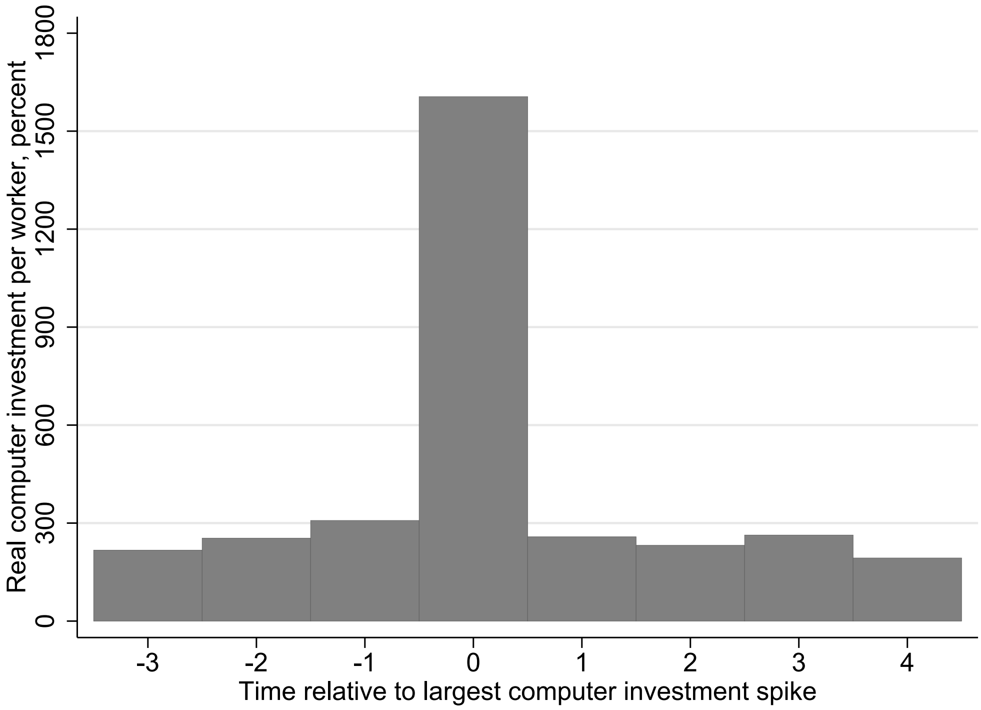

Notes: Overlapping sample, N=2,745.

computerization events are also clearly visible in the estimation sample: in the event year, treated firms spend around 1,605 euros per worker, compared to around 246 euros in the years before and after. This is similar to the 1,912 euros treated firms spend per worker during automation events.

Armed with our overlapping sample and both types of events at the firm level, we now construct four different datasets. First, we consider automation events and computer events in isolation: that is, we identify treated and control group workers for one type of event while ignoring the other. This allows us to estimate our DiD model for automation and computerization separately. However, these two events are correlated across firms over time: that is, firms

Figure 20. Relative annual wage income effects of automation and computerization

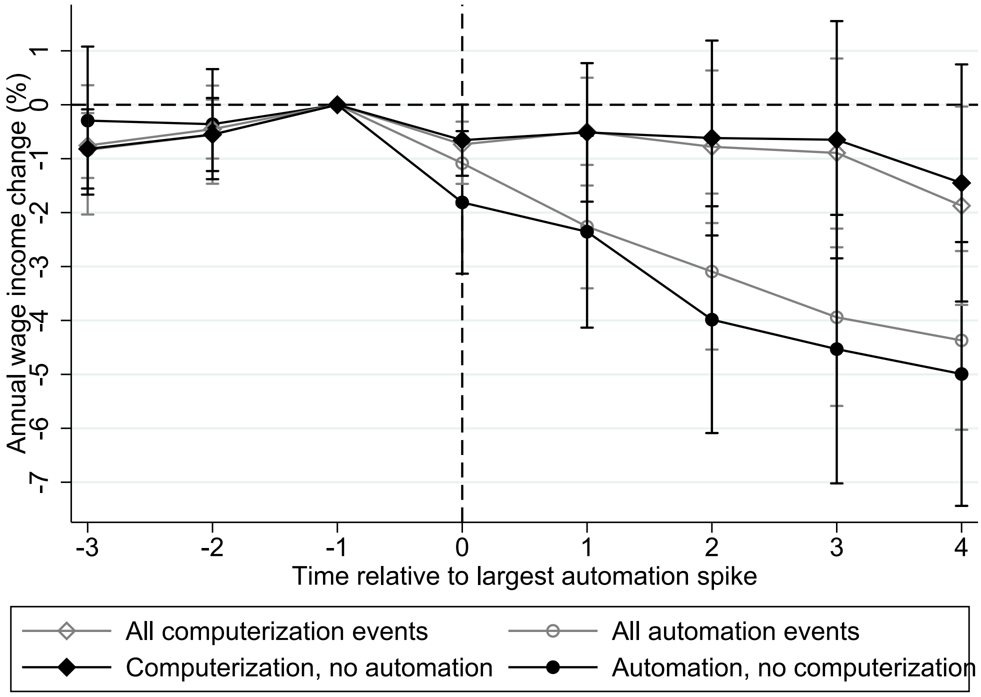

Notes: All estimates are for the overlapping sample where we observe data on both automation costs and computer investments. N=10,217,088 for all computerization; N=7,783,929 for all computerization excluding automation; N=8,110,456 for all automation; and N=4,632,880 for automation excluding computerization.

that have recently had one type of event are more likely to also experience the other sometime soon – sometimes even in the same year. This implies any estimated impact of automation may be contaminated by computerization, and vice versa. We therefore construct two additional samples of events which occur in isolation: that is, we only retain those automation (computerization) events where there is no computerization (automation) event occurring in the estimation window for either treated or control group firms. For each of the four samples, we then estimate equation 2 and report results in Figure 20.

This comparison leads to several findings. First and foremost, computerization does not lead to income losses for incumbent workers: estimates are small and never statistically significant. This is in contrast to automation, which does lead to income losses. Further, the income losses of automation are larger when removing concurrent computerization, and the effects of computerization on income are (slightly) smaller when removing concurrent automation. Consistent with these results, we do not find any increase in firm separation or non-employment duration for workers impacted by computerization. This may be because automation is generally more labor-displacing in nature than computerization, but also because computer technologies are already further along the adoption curve (as also evidenced by the higher frequency of computer investment spikes). This could imply that during computer spikes, it is mostly older vintages of computer capital that are being replaced, while any displacement of workers by computers has already played out in the past. All in all, irrespective of the reason, automation is (currently) a more labor-displacing force than computerization from the perspective of a firm's incumbent workers.

6 Conclusion

We provide the first estimate of the impacts of automation on individual workers, using firm-level data on automation expenditures across all non-financial private sectors in the Netherlands over 2000-2016. Leveraging a novel differences-in-differences design exploiting automation event timing, we show that automation at the firm significantly increases incumbent workers' hazard of separating from their employers. This firm separation is followed by a decrease in days worked, leading to a five-year cumulative wage income loss of some 11 percent of one year's earnings. Wage income losses are only partially offset by various benefits systems, and older workers are more likely to enter early retirement.

In contrast to displacement from mass lay-offs, however, workers do not experience daily wage scarring as a result of automation events, nor do we find evidence that automation displaces the firm's more recent hires. Further, automation events affect relatively few workers and these effects occur more gradually: two percent of incumbent workers separate from their employers in the first year, followed by a trickle of ongoing separations.

Our findings are robust to a range of specification and falsification tests, including controlling for other firm-level events such as mergers and acquisitions, firm restructurings, and suspected management change. Furthermore, effects are quite pervasive across different incumbent worker types, as well as firm sizes and sectors. Put simply, where work is automated, incumbent workers face a higher risk of displacement, resulting in adjustment costs. These adjustment costs may partly explain anxiety about new workplace technology expressed in opinion surveys, despite ample macro-economic evidence on technology's economic benefits. In contrast, we do not find evidence that workers face adjustment costs from firms' investments in computer technology. This suggests that, from the perspective of incumbent workers, automation is (currently) a more labor-displacing force.

Our findings of course do not imply that automation destroys jobs on net in the economy. As a related macro literature has shown, there are various countervailing mechanisms which our models do not inform about, including effects operating through firms' input-output linkages and changes in final demand. However, by focusing on workers directly impacted by automation events, our results contribute to understanding the adjustment costs of automation, which matter for science and policy alike.

References

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. 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., Goos, M., Salomons, A., and van den Berge, W. (2020). Evidence on Firms' Automation Activities. *American Economic Association Papers & Proceedings*, Forthcoming.

Blackwell, M., Iacus, S., King, G., and Porro, G. (2009). CEM: Coarsened Exact Matching in Stata. *Stata Journal*, 9(4):524–546.

Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78:1–3.

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).

Cortés, G. (2016). Where Have the Middle-Wage Workers Gone? A Study of Polarization using Panel Data". *Journal of Labor Economics*, forthcoming.

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.

Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2018). Adjusting to Robots: Worker Level Evidence . Working paper.

Davis, S. J. and Von Wachter, T. (2011). Recessions and the Costs of Job Loss. *Brookings Papers on Economic Activity*, 42(2):1–71.

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.

Dixon, J., Hong, B., and Wu, L. (2019). The Employment Consequences of Robots: Firm-Level Evidence. Ssrn discussion paper, SSRN.

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.

Edin, P.-A., Evans, T., Graetz, G., Hernnäs, S., and Michaels, G. (2019). Individual Consequences of Occupational Decline. Working paper.

Eurobarometer (2017). Attitudes Towards the Impact of Digitisation and Automation on Daily Life. Technical report, European Commission, Directorate-General for Communication, Special Eurobarometer 460.

Fadlon, I. and Nielsen, T. H. (2017). Family Health Behaviors. Working Paper 24042, National Bureau of Economic Research.

Fisher, S. R. A. (1935). *The Design of Experiments*. Macmillan.

Graetz, G. and Michaels, G. (2018). Robots at Work. *Review of Economics and Statistics*, forthcoming.

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.

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.

Kennedy, P. E. (1995). Randomization Tests in Econometrics. *Journal of Business & Economic Statistics*, 13(1):85–94.

Koch, M., Manuylov, I., and Smolka, M. (2019). Robots and Firms. Working paper.

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.

Pew (2017). Automation in Everyday Life. Technical report, Pew Research Center.

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.

Susskind, D. (2017). A Model of Technological Unemployment. Economics Series Working Papers 819, University of Oxford, Department of Economics.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. Working paper, Stanford University.

Young, A. (2018). Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results. *The Quarterly Journal of Economics*, 134(2):557–598.

A Appendix

This supplemental appendix contains details on sample construction, additional summary statistics, and additional robustness checks on our baseline specifications.

A.1 Sample construction

We follow several steps to construct our sample. We start with 36,490 firms for which we observe at least 3 years of automation costs data in the Production Statistics (PS) survey for 2000-2016. Then for each firm we determine if they have at least one spike, with spikes defined as explained in section 3.1. We keep all firms with at least one spike observed over 2000-2016. This leaves us with 10,476 firms. We lose an additional 2 firms because we cannot merge them to administrative worker records, so we continue with 10,474 firms. Then, for each calendar year y we define a set of potential treatment and control group automation events as follows.

Potential treatment events for y are defined as a firm having its largest spike in y. y has to lie between 2003 and 2011, so that for each event we at least have a window of three years before and five years after the event. Events are excluded if the firm also has another spike in the $t = [-3, 4]$ window around the event. This gives us 2,446 potential treatment group events. Note that we do not require that firms are observed in the PS survey in all 8 years around the potential event, but they do have to exist in each year in the window. Figures 4 and 5 are based on this sample. Figures 6 and 7 furthermore require firms to be observed in the PS survey in all 17 years, that is from 2000 to 2016.

Potential control events for y are defined as firms that have their largest spike in year $y + 5$ or later. Hence, these spikes have to occur between 2008 and 2016. Furthermore, events are excluded if a firm has another spike in the $t = [-8, -1]$ window around the event. Again, we do not require that firms are observed in the PS survey in all the years surrounding the event, but the firm does have to exist in this period. This gives us 21,575 potential control events.

Columns (1) and (2) in Table 13 show the number of potential treatment and control events per calendar year. Note that our procedure implies that multiple control group events can involve the same firm, but for different years y. It is also possible that one treatment group event and one or more control group events involve the same firm in different years y. For example, a firm that has its largest spike in 2010 can be a potential treatment event in 2010, but also serve as a potential control event for treatment events in 2003, 2004, or 2005. Similarly, a firm having its largest spike in 2011 can serve as a control group event for treatment events in 2003, 2004, 2005, or 2006. For our 21,575 potential control events, 20,838 involve a firm that is involved in more than one potential control event, while 737 events involve a firm that is involved in only one potential control event. Firms with potential control events are on average involved in 6.3 potential control events, with a maximum of 9 events. For our 2,446 potential treated events, 1,021 involve a firm that is also involved in at least one potential control event in another year and 1,425 involve a firm that is not involved in a potential control event.

We then merge our firm-level data to worker data and keep only events for which we can find at least one incumbent worker who is between 18 and 65 years old at $t = -1$. This leaves us with 2,439 potential treatment events merged to 124,225 incumbent workers and 21,399 potential control events merged to 1,157,536 incumbent workers.

We then take these samples of potential treatment and control group events and incumbent workers and apply our matching procedure. The details of our matching procedure are discussed in the main text in section 3.4. During matching we also apply some basic sample selection procedures to remove outliers. In particular, we remove students, people with total wage earnings above 0.5 million euros in a year or 2,000 euros per day and people earning less than 1/4 of the fulltime minimum wage (5,000 euros per year) or less than 10 euros per day on average. We also drop people whose earnings or daily wages increase more than tenfold since the year before treatment. After these sample selection procedures and matching, we are left with 102,599 treated incumbents in 2,429 events and 944,396 control incumbents in 21,175 events. Columns (3) and (4) in Table 13 show the number of events in each calendar year after matching. This is our main balanced panel of 1,046,995 workers, observed for 8 years, which means we have 8,375,960 observations in total.

A.2 Additional descriptives

Here we present some additional summary statistics for firms' automation expenditures and events; and for the main estimation sample of workers.

A.2.1 Summary statistics on automation expenditures and events

First, Table 14 shows automation costs per worker across firms of different sizes, when outliers in automation costs have been removed – that is, observations with automation costs per worker in the 99th percentile or above are dropped. This removes fewer than 50 firms, predominantly

Table 13. Number of treatment and control events at the firm level by calendar year

	Potential events		Events after matching	
Calendar year	Control	Treatment	Control	Treatment
2003	3,492	199	3,411	199
2004	3,271	200	3,208	198
2005	2,965	200	2,922	198
2006	2,725	223	2,688	223
2007	2,452	321	2,392	320
2008	2,205	311	2,161	308
2009	1,926	363	1,894	361
2010	1,544	315	1,521	312
2011	995	314	978	310
Total	**21,575**	**2,446**	**21,175**	**2,429**
Unique firms involved	4,588	2,446	4,543	2,429
Unique firms only used once	737	2,446	751	2,429

Notes: Table show the number of potential treatment and control events, and the number of events remaining after matching, for each calendar year.

one-person firms with high automation outlays per worker.

Second, Figures 21 and 22 show automation cost spikes in both shares and levels per worker for a balanced sample of firms – that is, firms which are observed in the Production Statistics survey every single year over 2000-2016.

Third, Table 15 shows how spiking and non-spiking firms differ in terms of observables. In particular, it estimates a firm-level linear probability model where the dependent variable is a dummy for the firm having at least one automation spike over 2000-2016: this model is estimated for the sample of firms where we observe at least 3 years of automation cost data. This highlights that firms that we observe having automation spikes are different from those where we do not observe a spike. In particular, the smallest firms are less likely to experience an automation event, and there are some sectoral differences, with automation events least likely to be observed in Accommodation and food serving, and most likely in Information and communication.

Lastly, firm-level spike timing is not easy to predict based on observables. Specifically, a predictive model with observables performs only marginally better than a random prediction where we uniformly distribute spikes across years where the firms are observed. This is reflected in the Brier (1950) skill scores for ten k-folded samples reported in Table 16. These are constructed as follows. We draw a 10 percent random sample without replacement from the main sample of spiking firms, and do this ten times: these are the test samples. The remaining 90 percent of observations for each of these test samples constitute the ten training samples. We then

Table 14. Automation costs by firm size class after removing outliers

Firm size class	Total cost Mean	Cost per worker Mean	Cost per worker SD	Cost share (%) Mean	Cost share (%) SD	Nr of obs *Firms*	Nr of obs *Firms × yrs*
1-19 employees	9,028	613	1,074	0.35	0.75	9,836	48,378
20-49 employees	20,393	639	1,068	0.37	0.74	13,755	86,523
50-99 employees	48,418	690	1,099	0.39	0.73	6,282	46,756
100-199 employees	105,035	754	1,172	0.41	0.70	3,470	28,472
200-499 employees	287,917	930	1,421	0.45	0.75	1,966	17,600
≥500 employees	1,584,027	968	1,522	0.62	0.99	1,134	10,188

Notes: Automation cost level in 2010 euros, automation cost shares as a percentage of total costs, excluding automation costs. Total N firms is 36,443; Total N firms × years is 237,917.

estimate a logit model with firm fixed effects and time-varying observables[52] on each training sample and predict the probability of having a spike in a year for each test sample, assuming that each firm will have exactly one spike. We also calculate the spike probability by year per firm from random prediction, simply as one over the number of years the firm is observed. For the model-based and random predictions in each of the ten test samples, we calculate the Brier score, defined as the mean squared difference between the prediction and the actual outcome. Lastly, we obtain the Brier skill score as $1 - \frac{Brier_{model}}{Brier_{random}}$, reflecting the percent prediction improvement of the model relative to random prediction. This improvement is 2.9 to 4.2 percent, confirming that spike timing is hard to predict.

A.2.2 Summary statistics for workers

Table 17 provides summary statistics on our sample of workers across all years, before matching. Table 18 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 $102,599 + 944,396 = 1,046,995$ observations for incumbents and $78,282 + 326,514 = 404,796$ for recent hires: given our observation window of 8 years ($t = -3$ through $t = 4$) this adds up to the $1,179,584 \times 8 = 8,375,960$ observations for incumbents and $404,796 \times 8 = 3,238,368$ for recent hires used in our regressions.

[52]These observables are: firm average log yearly and daily wages, log total wage bill, log number of workers, log average worker age, log average worker tenure at the firm, share female and a full set of interactions. Results are similar when we do not include these interactions, or additionally include lagged observables in the model.

Figure 21. Automation cost share spikes for treated firms, balanced sample

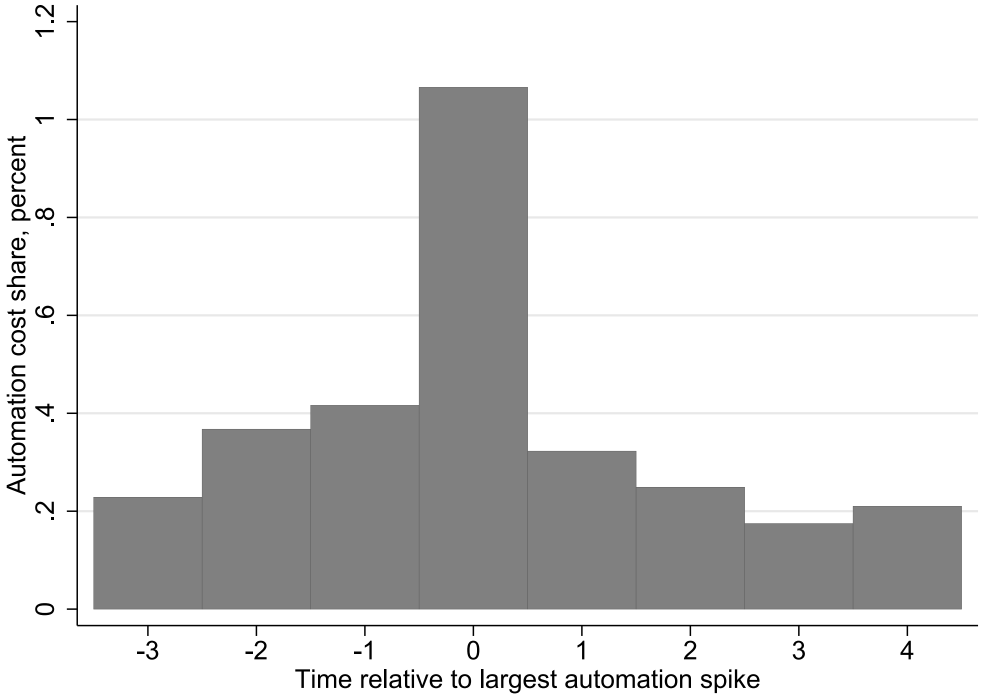

Notes: N=315 in all years.

Figure 22. Automation cost level per worker for treated firms, balanced sample

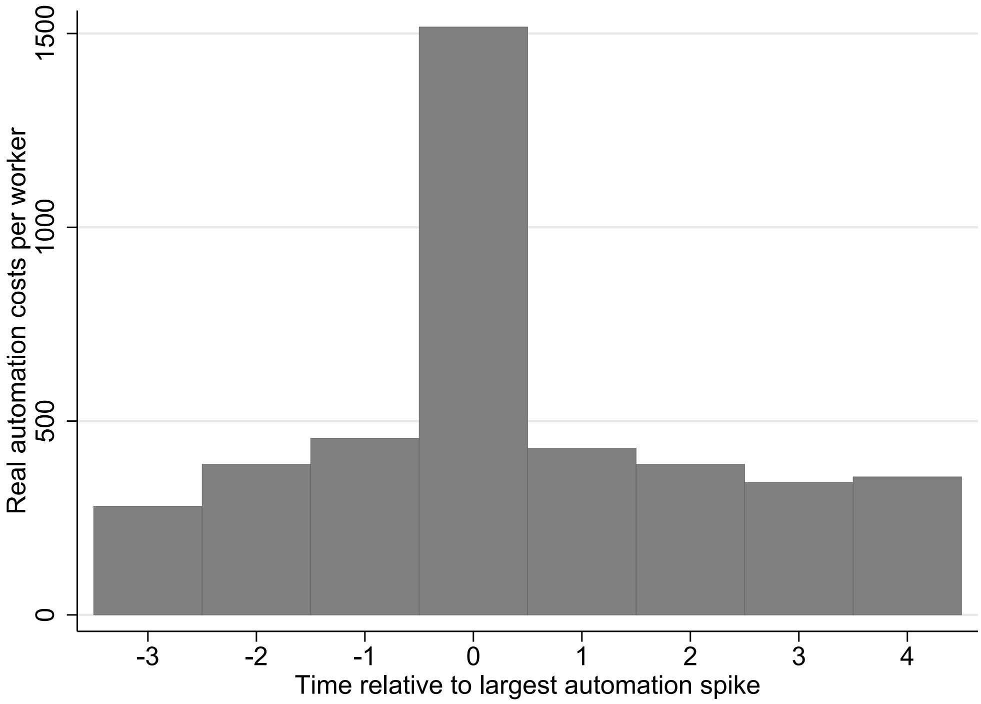

Notes: N=315 in all years.

Table 15. Correlates of a firm ever having an automation spike

Mean worker age	-0.0027***	Manufacturing	*reference*
	(0.0005)		
Share of women	-0.0075	Construction	-0.0294***
	(0.0126)		(0.0093)
Mean real annual wage / 1,000	0.0009***	Wholesale & retail trade	-0.0057
	(0.0002)		(0.0080)
1-19 employees	*reference*	Transportation & storage	0.0268***
			(0.0102)
20-49 employees	0.0347***	Accommodation & food serving	-0.0333**
	(0.0060)		(0.0151)
50-99 employees	0.0495***	Information & communication	0.0977***
	(0.0075)		(0.0114)
100-199 employees	0.0291***	Prof'l, scientific, & techn'l act's	0.0598***
	(0.0091)		(0.0102)
200-499 employees	0.0534***	Admin & support act	0.0089
	(0.0114)		(0.0099)
≥500 employees	0.0250*	Constant	0.3304***
	(0.0142)		(0.0214)

Notes: 36,489 observations, each observation is a unique firm. The dependent variable is having an automation spike at any point in the sample. Standard errors in parentheses * p<0.1 ** p<0.05, *** p<0.01

Table 16. Brier skill scores for predicting automation spikes

Sample	**N**	**Brier skill score**
1	127,378	0.040
2	127,232	0.037
3	126,485	0.038
4	127,886	0.033
5	126,812	0.042
6	126,440	0.029
7	127,599	0.033
8	126,954	0.032
9	126,497	0.028
10	127,616	0.036

Table 17. Descriptives for all workers

	Incumbents (1)	Recent hires (2)
Annual wage income	37329.65	25872.14
	(25159.74)	(22092.42)
Daily wage if employed	156.91	126.32
	(94.43)	(80.91)
Annual non-employment duration (in days)	28.59	66.32
	(61.96)	(90.25)
Hazard of leaving the firm	0.04	0.12
	(0.21)	(0.32)
Total benefits	403.45	1610.27
	(2724.93)	(4449.22)
Probability of entering early retirement	0.01	0.00
	(0.09)	(0.05)
Probability of becoming self-employed	0.03	0.04
	(0.18)	(0.21)
Share female	0.25	0.35
	(0.43)	(0.48)
Foreign born or foreign-born parents	0.16	0.27
	(0.36)	(0.44)
Age	42.60	36.60
	(10.24)	(10.06)
Calendar year	2006.90	2006.86
	(3.37)	(3.43)
Manufacturing	0.37	0.15
	(0.48)	(0.36)
Construction	0.11	0.07
	(0.32)	(0.26)
Wholesale and retail trade	0.19	0.16
	(0.39)	(0.36)
Transportation and storage	0.09	0.07
	(0.28)	(0.26)
Accommodation and food serving	0.02	0.02
	(0.13)	(0.15)
Information and communication	0.06	0.05
	(0.23)	(0.23)
Professional, scientific, and technical activities	0.08	0.08
	(0.28)	(0.28)
Administrative and support activities	0.09	0.38
	(0.29)	(0.49)
0-19 employees	0.06	0.06
	(0.23)	(0.23)
20-49 employees	0.14	0.12
	(0.35)	(0.32)
50-99 employees	0.12	0.10
	(0.32)	(0.30)
100-199 employees	0.12	0.10
	(0.33)	(0.30)
200-499 employees	0.15	0.11
	(0.36)	(0.31)
$\geq$500 employees	0.42	0.52
	(0.49)	(0.50)
Observations	9,017,448	4,392,416

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

Table 18. Descriptives on matched worker samples

	Incumbents		Recent hires	
	Treated	*Control*	*Treated*	*Control*
	(1)	(2)	(3)	(4)
Annual wage income	35885.04	35950.43	22944.41	22963.90
	(23717.08)	(23973.53)	(17594.43)	(17658.74)
Daily wage if employed	145.77	146.13	109.72	109.10
	(90.79)	(90.11)	(64.91)	(64.29)
Annual non-employment duration (in days)	21.51	22.48	64.09	63.84
	(43.22)	(44.25)	(71.02)	(71.60)
Total benefits	0.00	0.00	1401.83	1406.47
	(0.00)	(0.00)	(3648.65)	(3688.33)
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	0.03	0.04	0.03
	(0.16)	(0.17)	(0.18)	(0.18)
Share female	0.33	0.31	0.43	0.38
	(0.47)	(0.46)	(0.49)	(0.48)
Foreign born or foreign-born parents	0.18	0.17	0.31	0.33
	(0.38)	(0.38)	(0.46)	(0.47)
Age	40.78	40.65	36.01	35.27
	(10.18)	(10.05)	(10.15)	(10.03)
Calendar year	2006.41	2006.41	2007.09	2007.09
	(2.39)	(2.39)	(1.99)	(1.99)
Manufacturing	0.23	0.23	0.07	0.07
	(0.42)	(0.42)	(0.25)	(0.25)
Construction	0.08	0.08	0.04	0.04
	(0.28)	(0.28)	(0.19)	(0.19)
Wholesale and retail trade	0.26	0.26	0.12	0.12
	(0.44)	(0.44)	(0.32)	(0.32)
Transportation and storage	0.09	0.09	0.05	0.05
	(0.28)	(0.28)	(0.22)	(0.22)
Accommodation and food serving	0.02	0.02	0.02	0.02
	(0.15)	(0.15)	(0.12)	(0.12)
Information and communication	0.05	0.05	0.04	0.04
	(0.21)	(0.21)	(0.19)	(0.19)
Professional, scientific, and technical activities	0.11	0.11	0.07	0.07
	(0.31)	(0.31)	(0.25)	(0.25)
Administrative and support activities	0.17	0.17	0.60	0.60
	(0.38)	(0.38)	(0.49)	(0.49)
0-19 employees	0.07	0.06	0.04	0.04
	(0.25)	(0.24)	(0.19)	(0.20)
20-49 employees	0.16	0.16	0.09	0.10
	(0.37)	(0.36)	(0.29)	(0.30)
50-99 employees	0.12	0.13	0.07	0.09
	(0.33)	(0.33)	(0.25)	(0.29)
100-199 employees	0.12	0.13	0.07	0.09
	(0.32)	(0.34)	(0.26)	(0.28)
200-499 employees	0.14	0.15	0.07	0.10
	(0.35)	(0.36)	(0.25)	(0.30)
$\geq$500 employees	0.39	0.37	0.67	0.58
	(0.49)	(0.48)	(0.47)	(0.49)
Observations	102,599	944,396	78,282	326,514

Notes: Weighted means for the full regression sample at $t = -1$, where weights are obtained from coarsened exact matching as described in section 3.4. Standard deviations in parentheses.

A.3 Additional robustness checks

In this section, we present several additional robustness checks on our results. First, we present estimates for wage income in levels. Second, we show permutation estimates for each of our other dependent variables (the hazard of firm separation, days in non-employment, and log daily wages). Next, we change our empirical specification in a number of ways. Lastly, we consider how changing the spike definition changes our results.

A.3.1 Effects in levels

Figure 23 shows estimates for our main specification in levels. This shows that incumbents lost around 3,842 euros over five years in total, which is very similar to the 3,839 euros lost when estimating impacts on relative income (shown in Figure 8).

Figure 23. Annual real wage income in levels

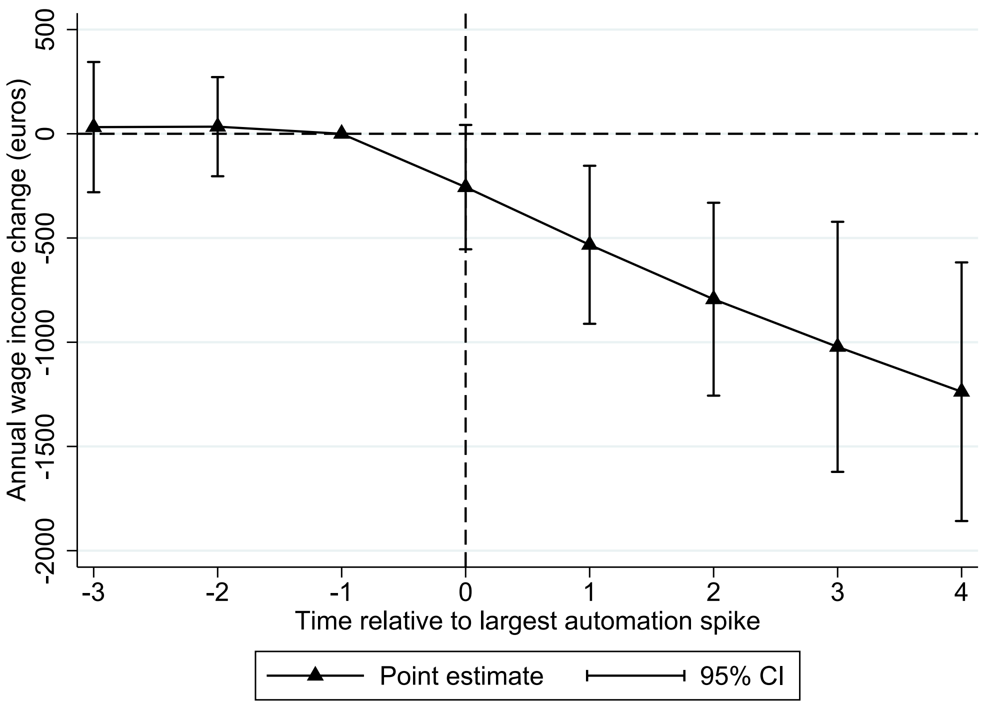

Notes: N=8,375,960. Whiskers represent 95 percent confidence intervals.

A.3.2 Randomization tests for other outcome variables

Figure 24 shows permutation estimates for the hazard of firm separation (panel a), days in non-employment (panel b), and log daily wages (panel c). Two-sided probability values are reported below each estimate: see the main text for a discussion of these results.

Figure 24. Additional randomization tests

(a) Probability of firm separation

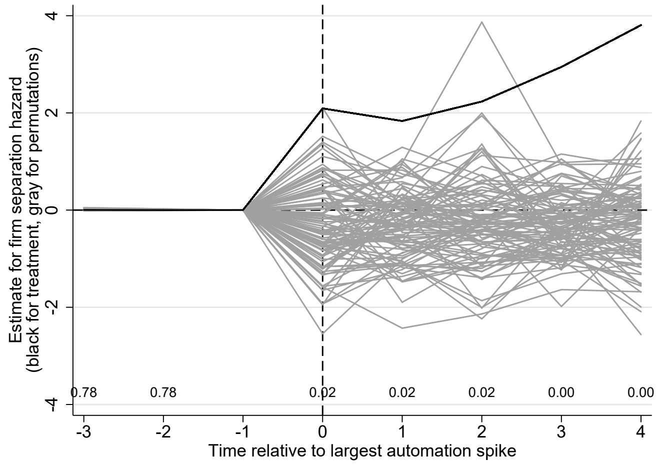

(b) Days in non-employment

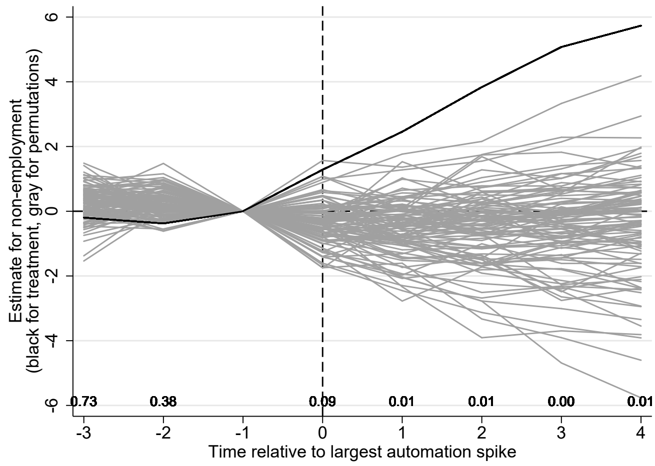

(c) Log daily wage

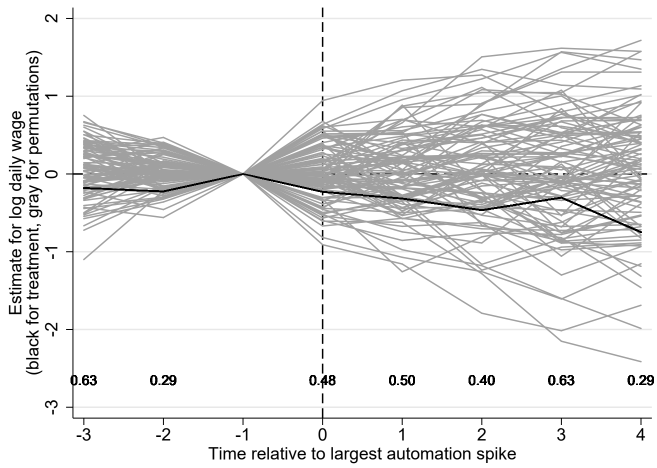

Notes: 100 permutations. The numbers printed at the bottom of the graph are probability values for the treatment estimates, based on the randomization test.

A.3.3 Changes in model specification

Here, we change our model specification in a number of ways. In particular, compared to our baseline estimates, Figure 25 shows results when additionally matching workers on their firm tenure in years (that is, beyond the three years of firm tenure that all treated and control group workers have); additionally matching workers on firm size; and when removing individual fixed effects from the model (these are then replaced by dummies for worker gender and nationality, as well as fixed effects for firm size categories, and for firm sector). Although estimates without individual fixed effects are a little less precise, results are extremely robust to these changes in specification.

Figure 25. Robustness to changes in model specification

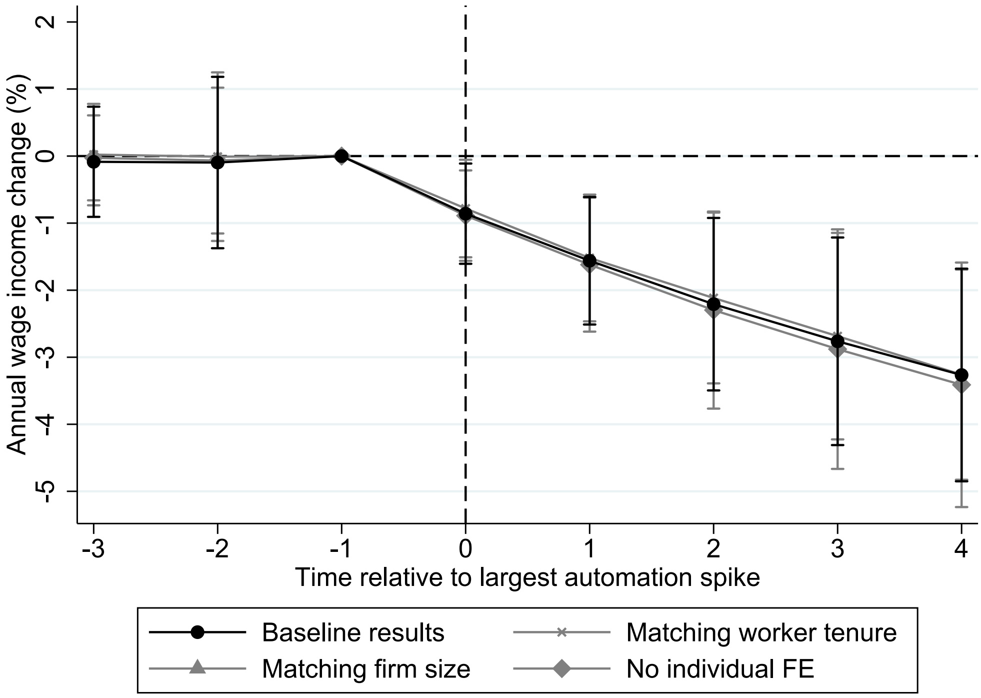

A.3.4 Removing other firm events

Figure 26 shows that our results for firm separation are also robust to excluding other firm-level events: see section 4.3 in the main text for a description.

Figure 26. Robustness to removing other firm events

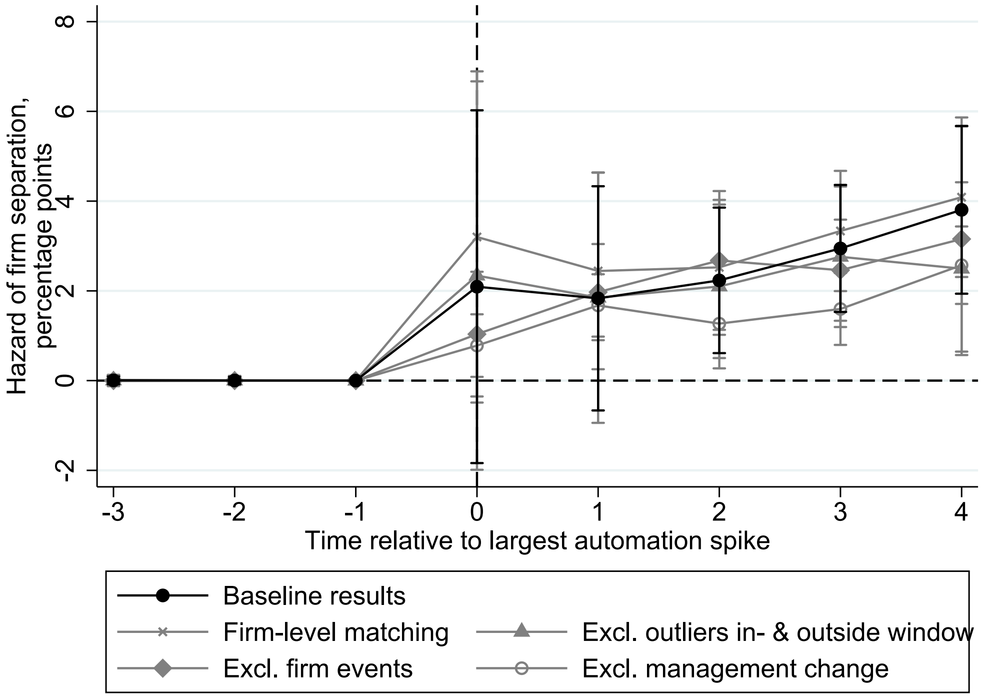

A.3.5 Alternative automation spike definitions

We have considered a range of alternative ways of identifying automation spikes, either by changing the spike definition or the spike threshold.

In particular, rather than using automation cost shares (i.e. automation costs in total costs), we can construct automation events from sharp increases in automation outlays per worker. This is more in the spirit of a literature studying the impact of increasing the number of robots per work. Within this event definition, we then also vary the point(s) in time where we measure employment (i.e. the denominator in the spike variable) – either for the years where we have data on total costs ("AC/worker"); or for the full set of years ("AC/worker, full emp data"); or only for the years pre-dating the candidate automation event ("AC/worker, pre-event emp data"). All variations produce similar results to our baseline estimates, as seen in panel (a) of Figure 27.

Further, we show that results are robust to varying the spike threshold from two to four times the average automation costs (our baseline is thrice the average automation costs). Panel (b) in Figure 27 reveals that estimated effect sizes are somewhat larger the higher the threshold, as expected, but these differences are not statistically significant. This highlights that our results are not driven by the specific spike size cut-off we employ in our baseline estimates.

Figure 27. Robustness to changes in spike definition

(a) Changes in spike definition

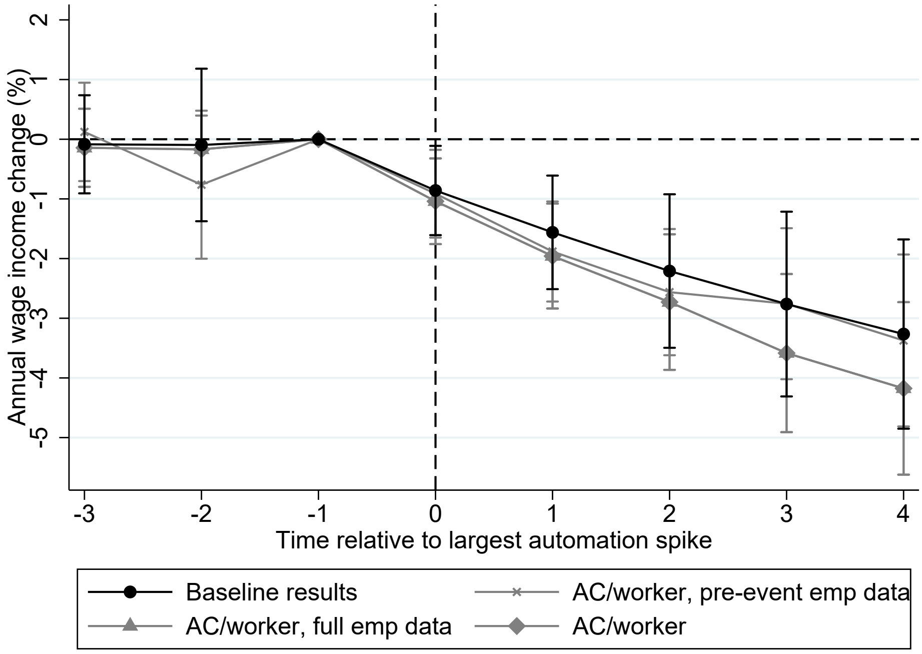

(b) Changes in spike threshold