Firm-level Automation, Offshoring and Employment

Erling Barth, Marianne Røed, Pål Schøne, Janis Umblijs

Abstract

In this paper we analyse the effects of automation on offshoring and employment, using detailed matched employer-employee data, containing firm-level register information on both automations, offshoring, and employment. Automating the production process may displace certain workers, especially workers conducting tasks overtaken by the automation process. A key question is which workers are potentially displaced when automation occurs. In this paper we look closer at the relationship between domestic workers and workers located abroad. Irrespective of offshore indicator used, the results show a negative effect of automation on offshoring, in the post-period. This result is supportive of a hypothesis that automation substitutes workers from abroad. We also find that automation reduce domestic employment. The reduction is largest among blue-collar workers and low-skilled workers. Our results combined, suggest that automation reduces employment abroad, but does not spur domestic employment.

Index-terms: F14, F15, F16,

Keywords: Automation, Offshoring, employment.

* We thank participants at the annual EALE conference in Prague in September 2023 for valuable comments. We acknowledge funding from the projects: “When robots meet the Nordics: Job creation, job destruction or retraining?” grant # 280307, and “Power, structure and technology.”, grant # 295914. Both financed by the Norwegian Research Council.

1.Introduction

A characteristic feature of globalisation in the last decades is the increase in offshoring (see for example Feenstra and Hanson 2003, Hummels et al. 2018). Dramatic advances in transportation and communication technologies have reduced the costs of offshoring. When components and intermediates can be transported fast and inexpensively across borders, and the output of many tasks can be carried out electronically, firms can take advantage of factor cost differences in different countries. This has motivated firms in high-income countries to outsource part of the production abroad (Grossmann and Rossi-Hansberg 2008), reallocating jobs abroad, often to low-income countries, specialising in manual tasks. Over the last decades we have observed a rapid increase in automation and robotizing in the production processes in high income countries. This new technology can have an impact on firms' decisions on where to locate production, between relocating work abroad or utilizing new technologies to increase productivity domestically.

Automating the production process domestically may displace certain workers, especially workers conducting tasks overtaken by the automation process. A key question is which workers are potentially displaced when automation occurs. In this paper we look closer at the relationship between domestic workers and workers located abroad. Automation may foster a recent phenomenon called reshoring. Reshoring describes the reverse process of offshoring, namely that previously offshored tasks are moved back into the home country. If automation mainly substitutes foreign labour, firm offshoring should fall. One of the reasons why companies may decide to not offshore or even to reshore production is that advances in robotic automation technologies reduce their costs of production, no matter where they produce. This, in turn, increases the attractiveness of domestic production as compared to offshoring. Automation thus has the potential to increase reshoring. Automation may increase productivity and bring back jobs that previously had been located abroad.

In this paper we analyse the relationship between automation and offshoring. This will shed light on the question on which type workers compete most directly with automation. If automation mainly replace tasks undertaken abroad, we expect a negative effect of automation on offshoring activities. A related question is how automation affects domestic employment, and the relationship between automation, offshoring and domestic employment. We also analyse effects of automation on firm-level domestic employment, for all workers and subgroups of workers with different task sets, and for different levels of offshoring in the pre-automation spike periods.

The relationship between automation and offshoring and the concept of reshoring has been used more often in recent years, both in political documents and in the business literature[1], but the concept has not yet received much research attention, firm-level evidence is especially scant.[2]

We analyse the impact of firm-level automation on firm-level offshoring and employment, using detailed Norwegian employer-employee data for the period 2000-2018, matched with firm-level import data on both offshoring, employment, and automation. Information on both offshoring and automation is taken from trade data. The trade data includes information on the value of each firm's imports, broken down by HS-6 digit product codes and source countries.

Methodologically, we employ a difference-in-differences model, limiting treat and control firms to firms that at some point went through a process of automation. Treatment status is defined by a spike approach, identifying the automation spike by the period the

---

[1]See e.g., article in *The Economist* on the German company, Adidas, who traditionally has offshored production to low labour costs locations in China, Vietnam, and Indonesia. Now, new factories are located in Ansbach, Germany and Atlanta, US. They produce shoes using mainly industrial robots, but since robots cannot perform all tasks, Adidas employ about 160 people locally in Ansbach and Atlanta. https://www.economist.com/business/2017/01/14/adidass-high-tech-factory-brings-production-back-to-germany

[2] Aghion et al. (2021) is an overview article covering the effects of automation on labour demand. One avenue for future empirical research they point to is the interaction between automation and outsourcing and international trade.

firm's relative automation costs spike. By limiting treat and control firms to firms that have at some point automated , we use variation in the timing of treatment for identification, thereby increasing the likelihood of firms in the treatment and control groups being on parallel trends prior to the automation spike.

The literature analysing the effects of automation on offshoring are scant, and many of the studies have analysed impact in low-income countries, for example, Artuc et al. (2019), and Faber (2020) both use data on employment and exports in Mexican local labour markets. They find that exposure to increased robot penetration in the US, using data from the International Federation of Robotics (IFR), reduces exports from Mexico to the US, and also reduced employment in Mexico.

One recent relevant contribution analysing impacts in high income countries is Bonfiglio et al (2022), They use US data across industries, occupations and local labour markets to analyse the relationship between automation (measured by robotization) and offshoring. They find that industrial robots lower the incidence of offshoring.

These papers all rely on aggregate-level information on automation and robot use. The literature using firm-level information on automation is very scant. One example of the latter though is Stapleton and Webb (2021). They use data for Spanish manufacturing firms, 1990-2016, to analyse how automation in Spain affects trade and multinational activity in lower-income countries. Their results show, contrary to the assertion that automation in high-income countries will cause reshoring of production, the use of robots in Spanish firms actually had a positive impact on their imports from, and number of affiliates in, lower-income countries. They show that these findings can be explained in a framework that incorporates firm heterogeneity, the choice between automation, offshoring and performing tasks at home and where automation and offshoring both involve upfront fixed costs.

Our paper also relates to the growing literature, that analyses the effect of automation and robot use on different outcomes, using firm-level data, including Acemoglu et al. (2020), Aghion et al. (2020), Humlum (2020) and Koch et al. (2019). One important outcome in these papers is firm level employment, and while Acemoglu et al (2020) find negative effects of robots on employment, both Koch et al. (2019) and Aghion find positive employment effects. Aghion et al. (2019) find the effect of automation on employment is positive, even for unskilled industrial workers. Automation leads to higher sales, and higher profits. They argue that the results can be accounted for in a simple monopolistic competition model: firms that automate more increase their profits but pass through some of the productivity gains to consumers, inducing higher scale and higher labour demand.

We contribute to the literature in several way: first; we add to the very scant literature on in the effect of automation on offshoring, using high quality firm-level data, with firm-level information on both automation and offshoring. Second, we use an empirical approach (stacked difference-in-differences event study) this is well suited to uncover causal relationships. Third, we provide evidence on automation and firm level employment, thereby presenting a broader picture on the relationship between offshoring and domestic employment, shedding light on the reshoring debate.

The paper proceeds as follows: Section 2 presents the empirical approach; section 3 presents the data and variables. Section 4 is the results section, while section 5 concludes.

2.Empirical approach

We analyse the effect of changes in firm-level automation on annual firm-level offshoring and employment. We follow Bessen at al. (2025) and employ a stacked difference-in-differences model, limiting treat and control groups to firms that have automated at some point.[3]

We identify the group of treatment firms by the year *c*, the year they have a spike in their automation investments (we define automation spikes in the next section). These firms are then followed in the event window from three years before, to five years after the spike {-3,…,5}. The control group of firms are those that spike their automation investments in year c+5, or later. These firms are followed in the same period as the treatment group; {-3-, …, 5}. For example, treatment firms that have a spike in automation costs in 2003, are compared to a group of control firms that spike in 2008, or later. Both groups are then evaluated in the same period, i.e., 2000-2011. This exercise is then repeated for each cohort. Finally, we stack all the cohort datasets, limiting all datasets to run in the same time window {-3-, …, 5}. The cohort periods are 2001-2012. This gives us a balanced panel data in event-time, thereby reducing problem related to staggered timing of treatment, extensively discussed in the recent difference-in-differences-literature (see e.g., Sun and Abraham 2021, de Chaisemartin and D'Haultfoeuille, 2020; Goodman-Bacon, 2021).

Concretely, we estimate variants of:

$$(1)\ Y_{jt} = \alpha + Treat_j\ x \sum_{\tau=-3}^{5} \beta_\tau I_\tau + \sum_{\tau=-3}^{5} \beta_\tau I_\tau + \beta_t + \ \beta_{j,k} + \varepsilon_{jt}$$

where Y_{jt} is the firm level outcome for firm *j* at time *t*, and τ is event time (-3, …, 5). *Treat* equals 1 for firms that have a spike in their automation imports at τ=0. I_t denotes event-time indicators, with $\tau = -1$ as the reference category. The key parameters β_τ. These coefficients identify the relative changes in the outcome variable relative to the pre-treatment period τ=0 for automating firms versus later automating firms.

[3] Baker et al. (2022) and Deshpande and Li (2019) are two examples of other paper using similar approaches.

Furthermore, β_t denotes calendar year fixed-effects, which account for general trends in the outcome variable. $\beta_{j,k}$ denotes cohort by firm fixed effects. The cohort by firm fixed effects are included to account for the fact that firms can potentially enter both treatment and control groups between different event cohorts (for example, in event year 2003, firms that have a spike 2008 or later, will serve as control firms in 2003, but will enter as treatment firms in the 2008 cohort). Standard errors are clustered at the firm-level.

The main difference between the difference-in-differences strategy in (1) and an event study design is that the difference-in-differences set up uses a control group. This eliminates event time trends that do not appear in calendar time.

The two important identifying assumptions are parallel trends in the absence of treatments, and no anticipation at firms prior to automation. Regarding parallel trends, we will present whether pre-trends trend differently between treat and control firms. Further, by limiting the sample to all-automation firms, we only use variation in event-timing, rather variation between those who automate and those who do not. This should increase the likelihood of them being on similar rends.

The non-anticipation assumption, at the firm level, implies that firms do not anticipate an automation event. In general, the non-anticipation assumption says that future treatment time does not affect current outcomes (e.g. Abbring and Van Den Berg, 2003). We recognise that this assumption is difficult to fully meet at the firm level, but we shed light on the severeness of the problem by presenting an analysis, using adjusted definitions of the spike measure. Concretely, we use a stricter definition of the spike. In the original definition, the criterion is that the real total operating costs (excluding automation costs) averaged across all years are at least *three times* the average firm-level cost share (definition of spike in the next section). In the adjusted analysis we raise the bar to *four times*. This should make the spike measure sharper, potentially leaving less scope for bias due to anticipation. Still, this exercise is not a “bullet proof” test of the no-anticipation assumption, and therefore the results should therefore be interpretated as partly descriptive.

3. Data and variables

We exploit matched employer-employee register data, collected, and organised by Statistics Norway, consisting of all manufacturing firms observed in the period 2000-2018. We have information on all firms and all employees working in these firms, during this period. All analyses are limited to firms in the manufacturing industry.

The key dependent variable is an annual measure of firm-level offshoring, taken from public trade statistics, linked to each firm. In the trade statistics there is a distinction between intermediate imports, imports of capital goods, and import for consumption. We are primarily interested in the extent to which firms are engaged in offshoring and how this is related to automation, and division of labour domestically and abroad. This raises the question of whether the firm-level imports we observe are final goods or inputs into production, and also whether these inputs are potentially substitutes for labour, domestically or abroad. Our strategy is twofold: i) to use information on the industry location of the imports, and ii) to use information on whether the imported item is classified as an intermediate good. Concretely, we include three measures of offshoring: i) narrow offshoring, ii) offshoring of intermediates, and iii) a combination of the two, i.e., narrow offshoring of intermediates. Regarding narrow offshoring we follow Feenstra and Hanson (1999) and define “narrow offshoring” as purchases of inputs belonging to the same industry as the importing firms. The idea is that the closer the imported products are to the final outputs, the more likely it is that workers employed by the firm, alternatively, could have produced those inputs. Concretely, we limit import within the same two-digit industry as the firm’s output. Regarding import intermediates, the rise of global value chains has made the analytical distinction between trade in intermediates and trade in final goods more important. Intermediate goods, are typically used as inputs in the production of other goods.[4]

The second dependent variable is an annual measure of firm-employment, taken from public administrative registers. We include and construct employment measures for all workers, and for subgroups of workers, defined by their level of education and occupation. To increase robustness, we also present results using different measures of employment, taken from different administrative registers.

Trade data is collected from the Norwegian Trade Statistics Register. For each year, we have the value of imports disaggregated by products. Yearly information from the Trade Statistics Database is merged with all of the firms in the sample. Trade flows are reported according to the 8-digit combined nomenclature. These are aggregated up to the six-digit Harmonised system (HS) to be compatible with the COMTRADE data, which is the United Nations International Trade Statistics Database providing annual trade data for over 170 countries covering our period of interest. The value of imports is reported in Norwegian kroner (NOK). The second dependent variable is annual measure of firm-level employment, taken from Statistics Norway. The main measure is based on aggregate firm-level measure, but to increase robustness, we supplement with three other measures, including a measure based on own calculation from individual register information, and aggregated up at the firm level.

The key explanatory variable in this paper is the firms' automation costs. We identify automatising costs based on the Norwegian Trade Statistics Register, which records all cross-border goods purchased by firms. Following the literature, we define automation machinery using CN-2018 product codes. In particular, we follow the categorization of Acemoglu and

[4] Intermediate goods are defined from United Nations BEC-codes (Broad Economic categories). In general,in these statistics there are three types of goods: Consumption, capital, and intermediates. The following Rev. 4 intermediate codes are included: 1,3,5,6,7,9,11,15. The BEC is an international product classification. Its main purpose is to provide a set of broad product categories for the analysis of trade statistics. The newest BEC codes is Rev.5. This has not been available to us.

Restrepo (2022) and Bessen et al. (2023) and include automatically controlled machines, automation transfer machines, automatic welding machines, numerically controlled machines, and robots. In Appendix, A1., we list all codes included in the measure. Figure A1 in Appendix presents the distribution of automation costs across industries (2-digit industry), for each of the five automation technologies. For each technology we present the five largest industries, measured in total automation costs. The automation technologies are not spread evenly across industries in Norway. The most concentrated technology is *Automatically controlled machines*, with the majority of imports found in just one industry. The other four automation technologies are not as concentrated, but with some industries reappearing in several of the automation technologies. Typically, the five largest industries constitute more than 50 per cent of the automation costs. Two industries; manufacture of machinery equipment, and manufacture of transport equipment, are the two most common industries to introduce new automation technologies.

Ideally, we would like to have a direct firm-level measure of automation, derived from accounting measurements. Unfortunately, such data is not available to us. Nevertheless, we can use the import of automation technologies, including robots, as a reliable measure of automation use (see, for example, Acemoglu and Restrepo 2022). The rationale behind this is that Norway does not produce many industrial robots or advanced automated machines with the majority of this type of technology being imported from abroad.

Even though the majority of the automation technologies we are interested in are imported from abroad we acknowledge that there are some limitations to this approach. Firstly, this measure does not account for those automation technologies that are domestically produced. Secondly, it does not capture automation components purchased from domestic wholesale operators. While our data does not include information on self-production and third-party purchases, a recent representative survey of the Norwegian econo suggests that most companies that use automation technologies, such as robots, typically use those produced abroad. The 2020 survey carried out by the Institute for Social Research shows that 67% of industrial robots, are produced abroad ('DoT2020-survey').

Still, this does not rule out the possibility that these technologies could be purchased from domestic wholesale operators. These purchases would not be captured in our data. Therefore, our measure of automation costs will likely be a conservative estimate of the actual automation costs incurred by firms.

Table A1 in Appendix presents descriptive statistics for the outcome variables, and for the level of automation costs, by event cohort. The level of narrow and intermediate offshoring is generally somewhat higher for the treatment groups than for the control group. Number of employees is also somewhat higher among the treatment group.

Defining automation

We identify automation events using what we call automation spikes. Our definition follows Bessen et al. (2025). Automation spike is defined 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, TCj , are at least three times the average firm-level cost share, excluding year τ. This can be expressed as:

$$(2)\ Spikej\tau = 1\left\{\frac{ACj\tau}{TCj} \geq 3x\,\frac{1}{T-1}\sum_{t\neq\tau}^{T}\left(\frac{ACjt}{TCj}\right)\right\}$$

where 1{. . .} denotes the indicator function. As such, a firm that has automation costs around one percent of all other operating costs for year $t = \tau$ will be classified as having an automation spike in $t \neq \tau$ if its automation costs in τ exceed three percent of average operating costs over years *t*. Some firms have more than one spike. In that case we define an automation event as the year when the firm has its first automation spike.

Figure 1 shows how the automation cost shares develop before and after the year of the spike.

Figure 1. Automation spike

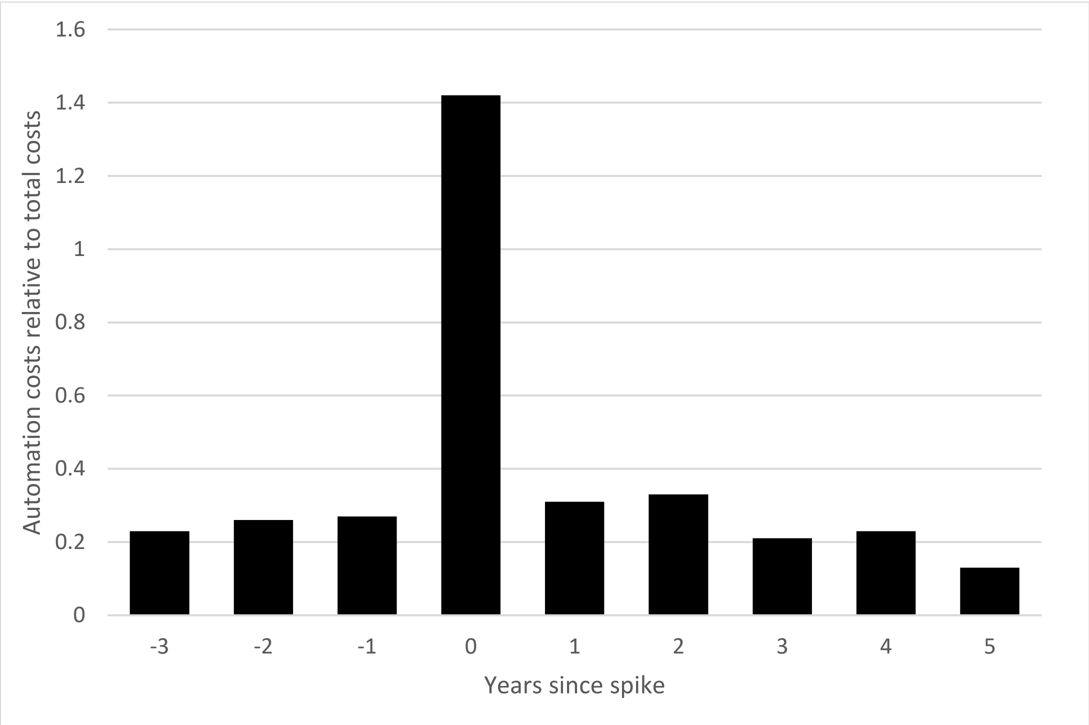

Note: The years since spike is from a data set where all the other years, besides the spike year also is included.

Figure 1 shows a clear one-period increase in the automation cost share when a firm has its automation event. In the robustness check section later, we present results using a moderated version of (2), using median instead of mean, and using the level of automation costs (not the relative shares) as spike measures.

Figure 2 shows how the number of firms is distributed along the spike years.

Figure 2. Number of firms at spike years

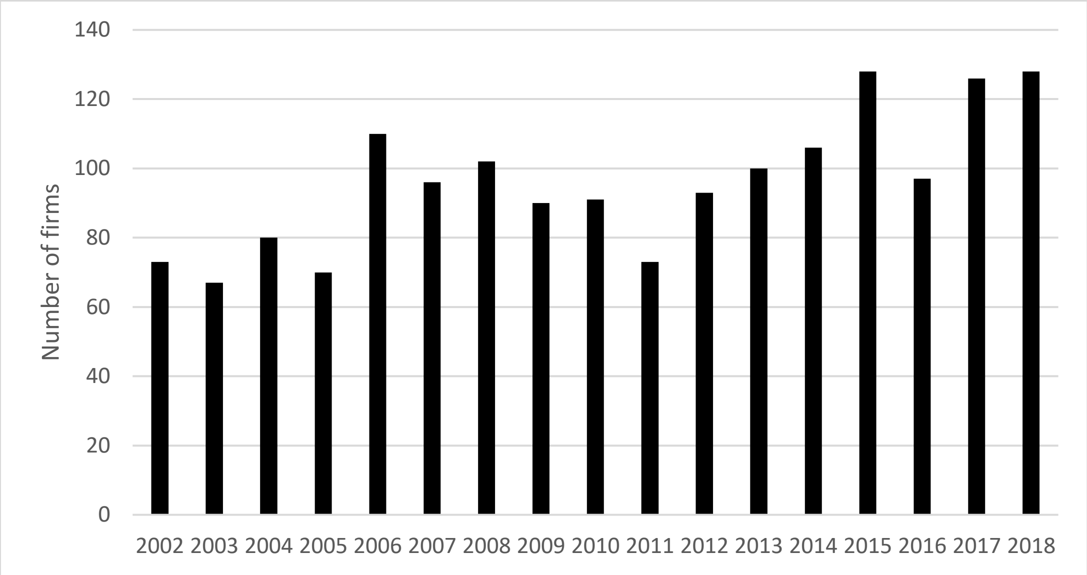

The number of firms at each spike year varies between 67 and 128. There is a small trend towards larger numbers towards the end of the period.

Figure 3a and b show the development in the three offshoring measures in the event window, for treat and control firms.

Figur3a. Log offshoring in event window. Treatment firms.

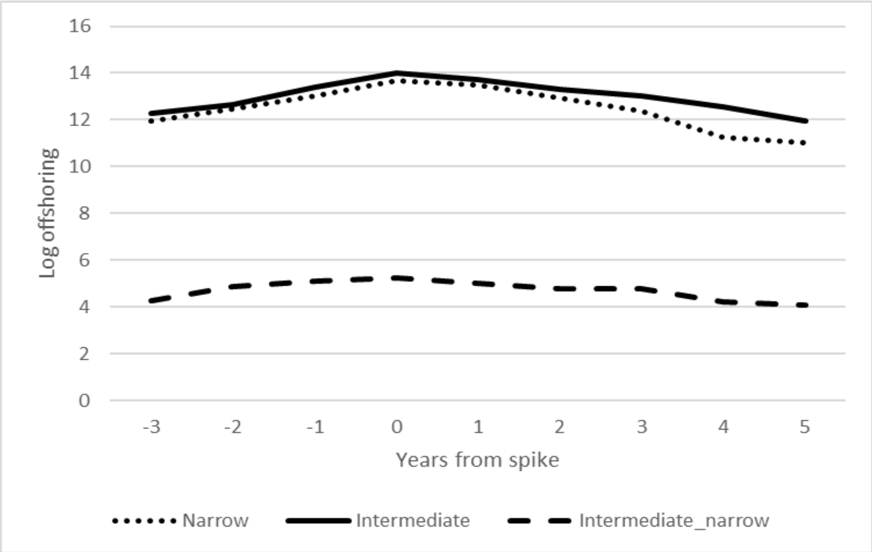

Figur3b. Log offshoring in event window. Control firms.

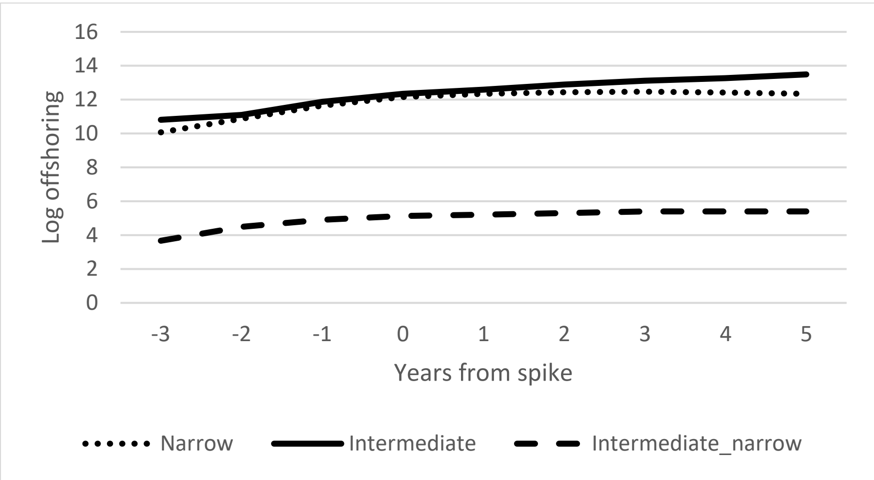

Note: Mean values of log offshoring in the event window (-3,…,5).

Figure 3a show that treatment firms are on an increasing trend prior to the spike, thereafter the mean level of offshoring is decreasing. For control firms, there is also an increasing pre-trend, but no trend shift is present around t=0. There is also evident that the mean level of offshoring is much lower when we consider both narrow and intermediate offshoring.

4. Results

4.1. Main results

Figure 4 presents results from estimating equation (1) for the three different measures of offshoring: narrow, intermediates and narrow and intermediates. All of these are measured in logs.[5]

Figure 4. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring.

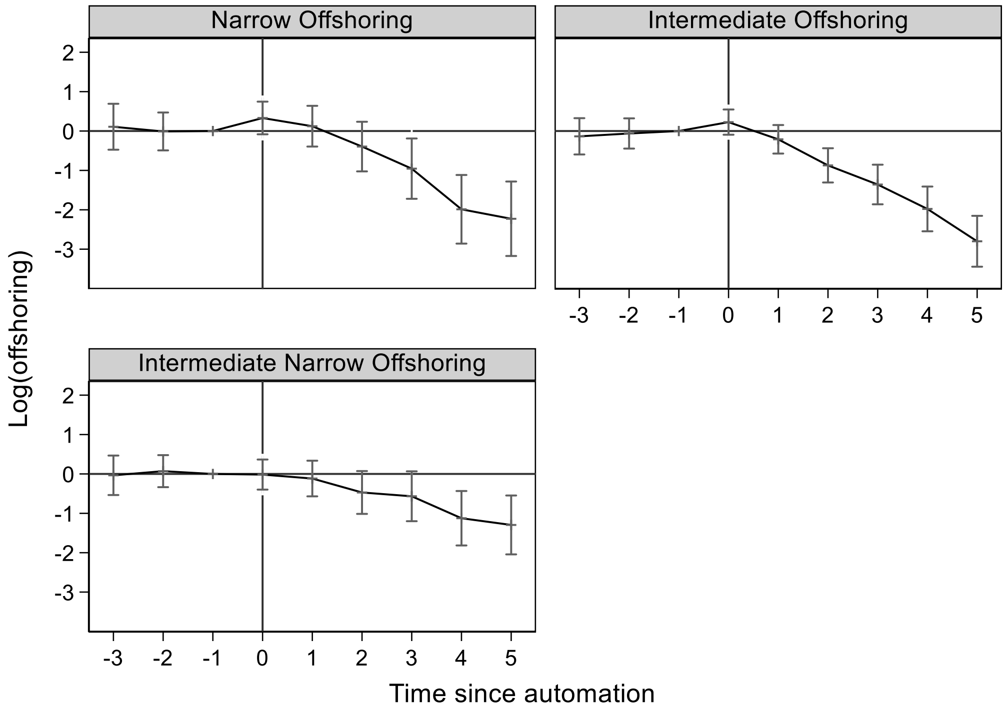

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals.

Irrespective of offshore indicator used, the results show a negative effect of automation on offshoring, in the post-period. For the narrow measure, the DD-coefficients suggest that firms

[5] We include firms with zero offshoring, included as 0+1.

in the treatment group have approximately -3.2 lower log offshoring five years after the spike, compared to the control group. Reassuringly, treat and post groups are not trending differently in the pre-period. These results are consistent with the notion that automation substitutes tasks that previously were undertaken abroad.[6]

Figure 5 presents regression results, using binary measures for offshoring, i.e., whether the firm is conducting offshoring or not. In general, most firms do some offshoring, for example the mean value at time 0 for narrow offshoring, is 0.83.

Figure 5. Automation and offshoring. Dependent variable: p(offshoring). All offshoring, narrow offshoring, and intermediates

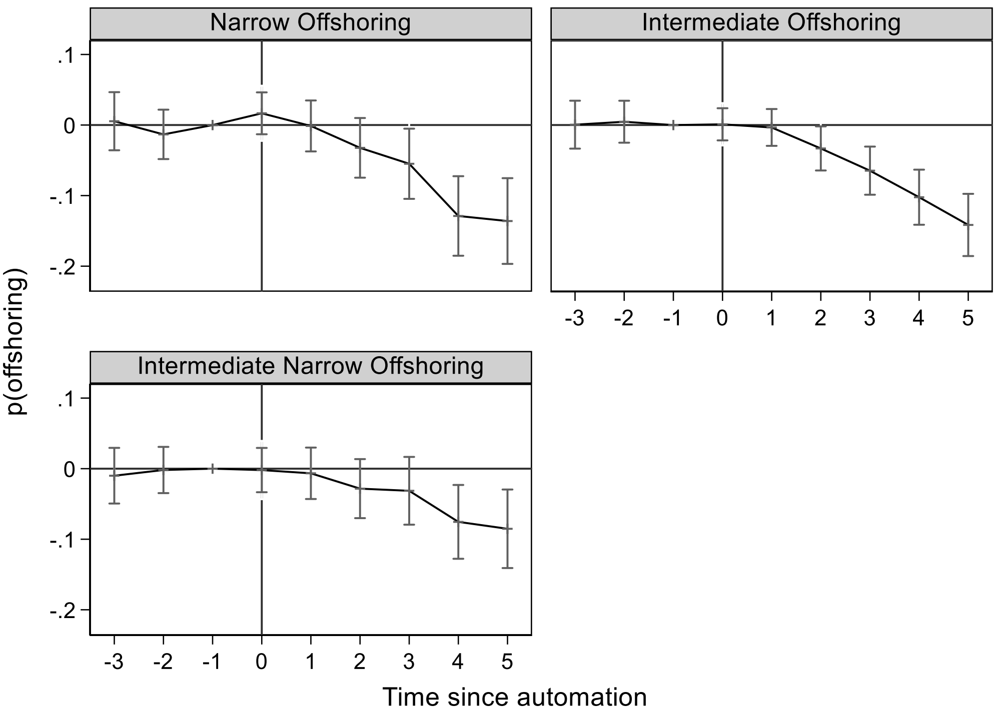

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Results for the binary measures resemble results using the continuous measure, we find a negative development in the likelihood of offshoring in the post-period. For example, using

[6] Full set of DD-coefficients and standard error are shown in Appendix, Table A2 (Figure 4) and A3 (Figure 5). In Appendix, Figure A6, we also present DD-coefficients for comparable to figure 3, using instead the the Inverse Hyperbolic Sine Transformation for the offshore measure (ihst). ihst= ln(x + ((x^2 +1)^0.5)). The results are almost identical. Therefore, we proceed with the simpler measure.

the narrow offshoring measure, we find that treatment firms have approximately 15 percent lower likelihood of conducting offshoring, five years after the event spike, compared to the control group of firms. These results suggest that automation affects offshoring on the extensive margin, i.e., on the decision to offshore or not. Again, reassuringly, the pre-period estimates do not reveal any diverging trends between treat and control observations.

In Appendix, Figure A2, we present DD-coefficient when estimating on the intensive margin, i.e., only for those that have positive values on offshoring. We do find negative effects also on the intensive margin, suggesting that automation also affects the decision on how much to offshore. Focussing on the narrow measure, five years after the spike, automating firms have approximately 40 per cent lower levels of offshoring, compared to firms automating later.

4.2. Heterogeneity analyses

In this section we conduct heterogeneity analyses depending on source of offshoring (OCED vs non-OECD countries), firm size and type of industry. First, we distinguish between offshoring to OECD and non-OECD-countries.[7] If automation processes are substituting task previously undertaken abroad, we suspect the negative development in offshoring should be especially prevalent in non-OECD countries, since these are the countries where many offshoring activities have been placed, due to lower production costs in these countries. First, Table 1 shows that the level of offshoring is considerably higher between OECD countries.

[7] OECD countries include Australia, Austria, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, Great Britain, Greece, Iceland, Ireland, Italy, Japan, South Kora, Luxembourg, Mexico, Netherlands, New Zeeland, Portugal, Sweden, Turkey, USA.

Table 1. Mean values. Offshoring OECD and non-OECD countries. Log of offshoring

	OECD	Non-OECD
Narrow offshoring	11.50	6.02
Intermediate offshoring	12.25	4.41
Intermediate and Narrow offshoring	4.69	1.86

Note: Log offshoring, including zeros as log(x+1).

Figure 6a and 6b present the results, separate for OECD and non-OECD countries, respectively.

Figure 6a. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. OECD countries

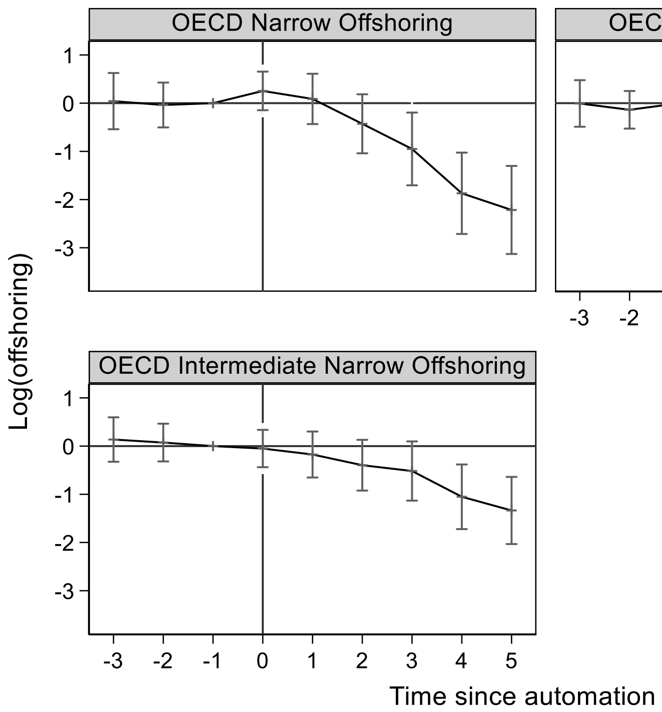

Figure 6b. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Non-OECD countries

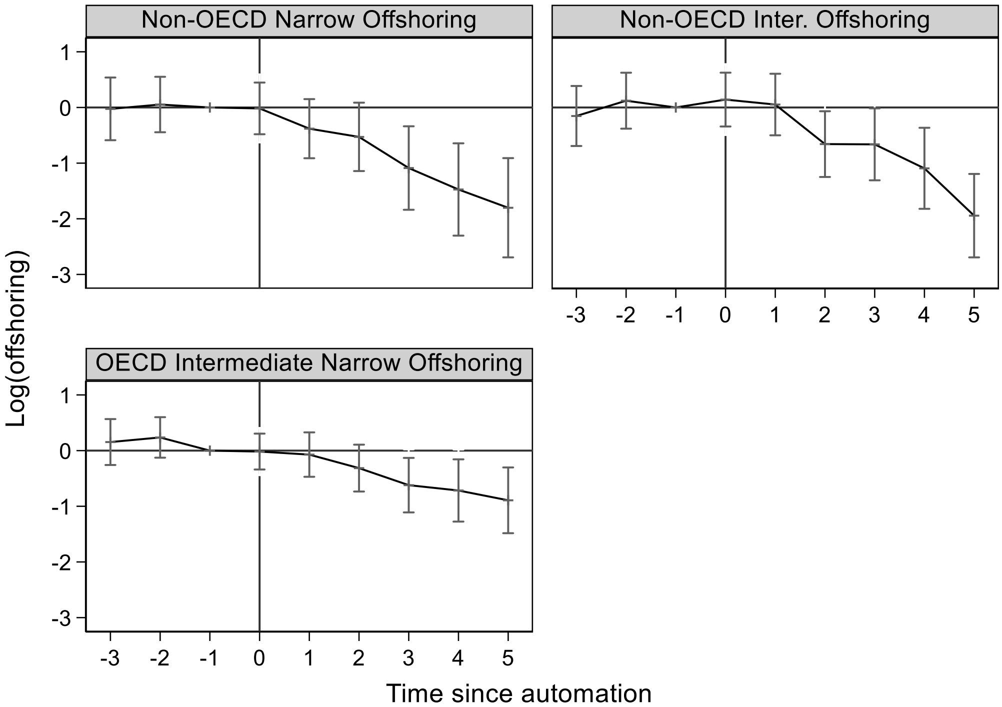

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals.

We find negative developments in offshoring in the post-period for both OECD and non-OECD countries. The difference is not significant. The result that offshoring from low-income countries do go down, suggests negative employment effects in these countries, a result consistent with analyses focusing on employment effects in offshoring countries (see e.g., Faber et al 2020). But this pattern is also visible among high income countries.

Next, we analyse results depending on firm size. Figure 7a and 7b present DD-results from estimating (1) for firms with more than 25, and 50 employees, respectively.

Figure 7a. Automation and offshoring. Dependent variable log(offshoring). More than 25 employees

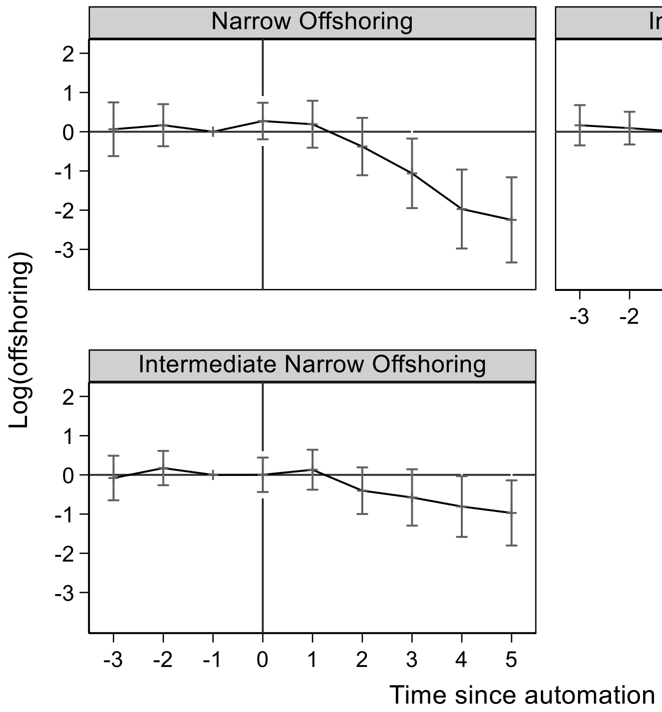

Figure 7b. Automation and offshoring. Dependent variable log(offshoring). More than 50 employees

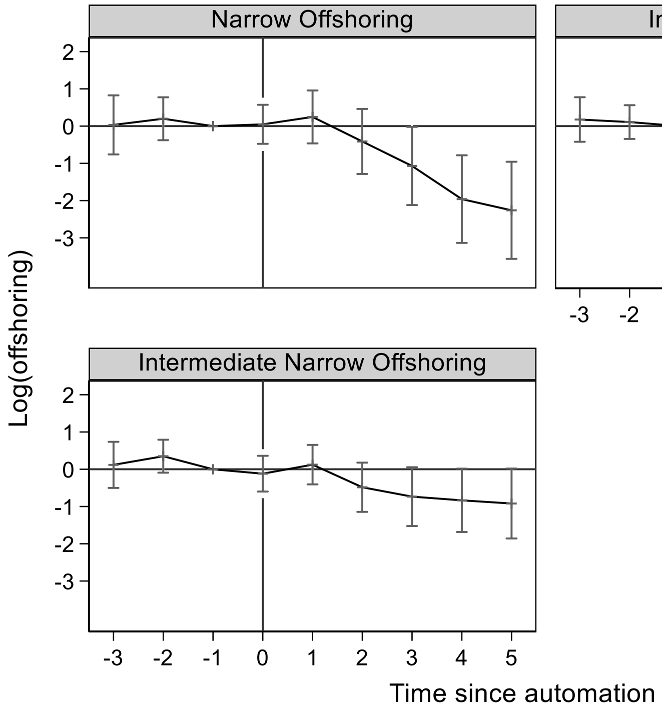

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

The results reveal that the negative effects of automation on offshoring is quite stable across firm size. Next, we conduct analyses separately for the 10 industries within the manufacturing industry that have the largest average automation costs.[8]

[8] The included Nace codes (2-digit Nace-2 codes) are 34, 35, 32, 33, 36, 27, 21, 20, 29, 15.

Figure 8. Automation and offshoring. Dependent variable log(offshoring). 10 largest automation industries within manufacturing

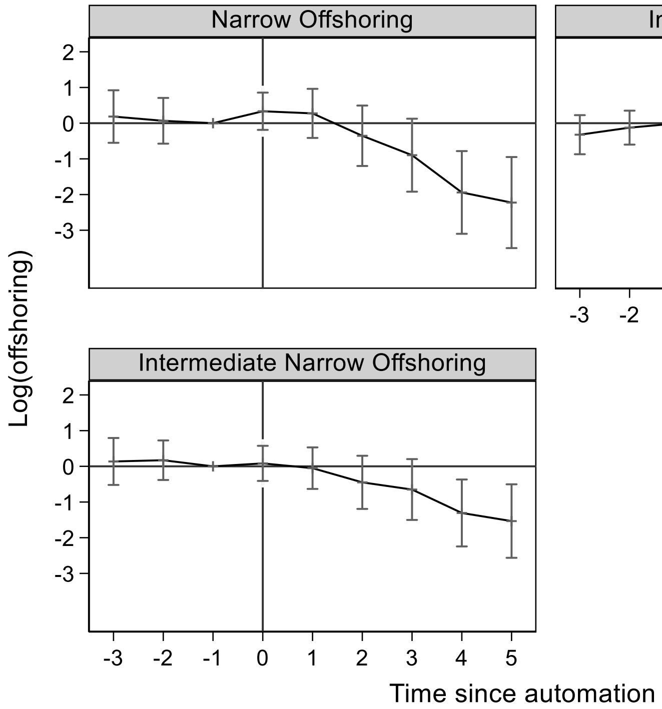

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

The results for these "high-dose" automation industries also reveal negative effects of automation on offshoring in the post-period. The negative effects are comparable in size to the results for all industries.

4.3. Robustness checks

In this section we conduct several robustness checks. First, we analyse if there are differences between early and later automation spike firms. We distinguish between event-periods earlier than 2007, and those from 2007 and later. Figure 9 and 10 presents the results for early and later adopters, respectively.

Figure 9. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Early adopters (earlier than 2007)

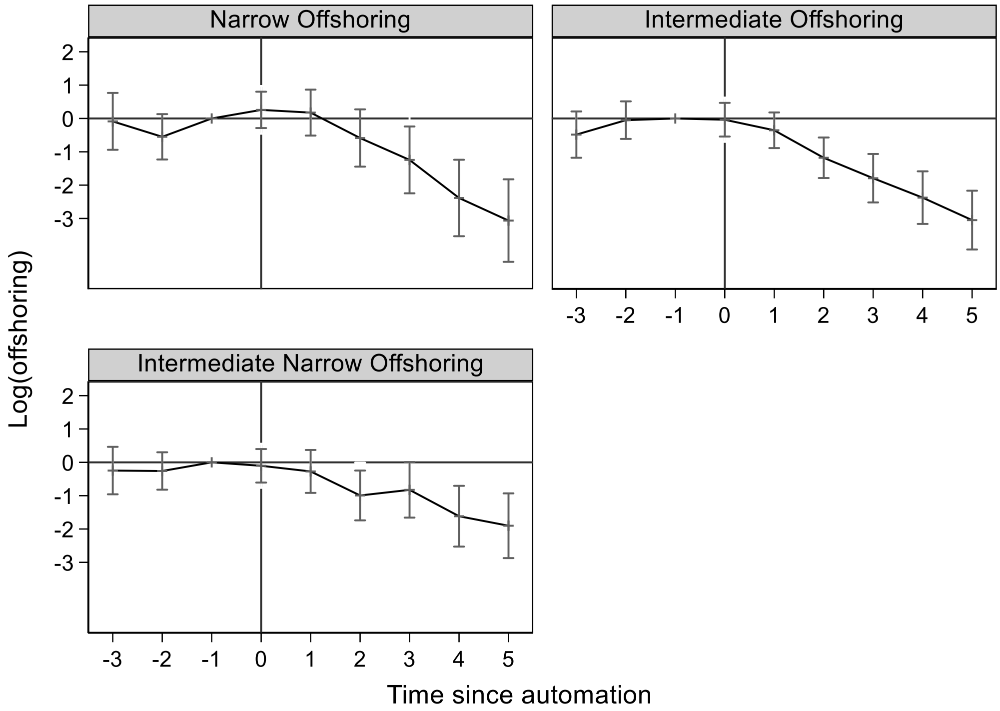

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure 10. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Late adopters (2007 or later)

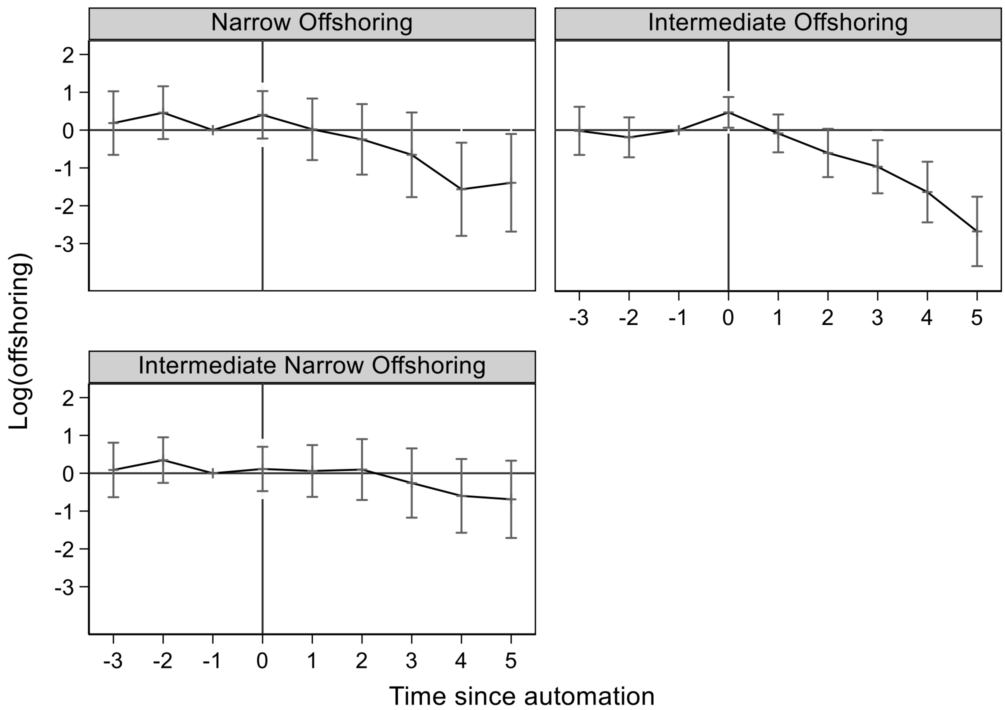

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

The negative effects are visible in both periods, but the negative effects appear sooner and are somewhat stronger in the early periods. The differences are generally not significant, and we conclude that the distinction between early and later adopters are not important.

Second, Figures A3-A5 in the appendix, we run regressions on equation (1) using slightly different spike definitions. Figure A3 uses median instead of the mean. i.e., that the cost share is at least three times the *median* firm-level cost share. Figure A4 uses the level of automation costs only, i.e., we do not use the relative definition, the spike criteria are that the level of automation costs is at least three times the mean level of automation costs. Figure A5 presents results for a more conservative definition of an automation spike, where the level of automation costs is at least four (as opposed to three) times the mean level of automation. In short, the results are very robust to these changes in spike definitions.

Finally, we examine whether the way we define our control group - specifically in terms of the number of years after the treatment group experienced a spike in automation - has any significant impact on our results. In the empirical analyses the control group of firms are those that spike their automation investments in year c+5, or later. This implies that we may compare treatment firms with control group of firms that spike many years later (since the requirements is 5 years or later). This may again mean that we compare automation processes that are quite different, considering the rapid technological development. To check for the potential severity of this problem we re-estimate equation (1), limiting control firms to those that spike maximum 7 years after the treatment firm. Figure 11 presents the results.

Figure 11. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Max 7 years wait for control firms

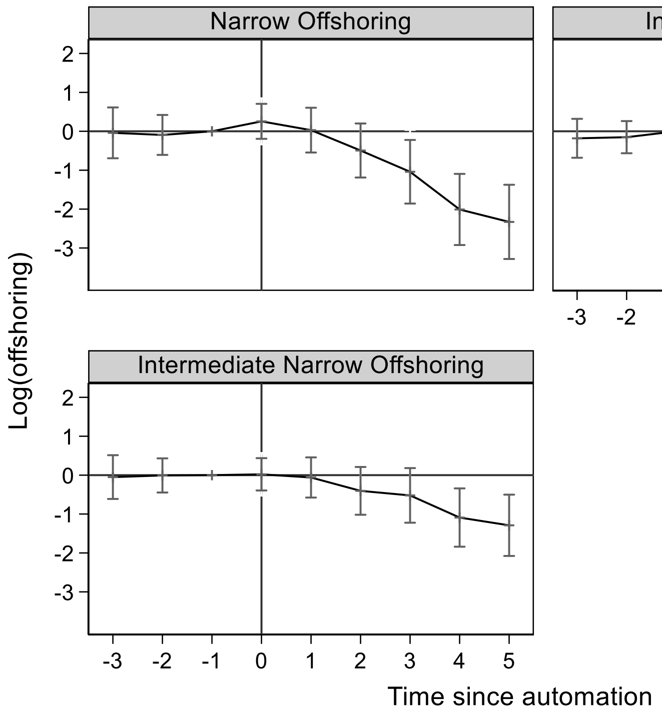

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

The reduction in the number of years for the control firms to be included naturally reduces the number of observations. The sample is reduced with approximately one third. Still, the results are quite stable after this adjustment. We still find a negative development in offshoring in the post-period. As earlier, the reduction is most evident using the narrow and intermediate measure.

4.4. Automation and employment.

We have established a robust negative effect of automation on offshoring, across different specifications. This implies that investments in automation may prompt firms to decrease their dependence on foreign intermediate inputs, which could potentially lead to a reduction in employment for those who were previously engaged in the production of these goods. A closely related question is how automation affects domestic employment, for the same firms.

If reshoring is part of the process, we should expect that some of the previously outsourced jobs is brought back and appear as an increase in domestic employment, especially for workers in occupations with tasks that complements the investments in automation technology.

In this section, we estimate the effect of automation on domestic employment, and hiring and separation patterns, for all workers and subgroups of workers. We also conduct analyses depending on the firm's pre-spike level of offshoring.

First, to get a descriptive picture of the development of the level of employment, Figure 12 presents descriptive mean values of log employment for treat and control firms, before and after an automation spike. Prior to the spike, the trends for both groups of firms appear similar. In the post period, however, their paths diverge. While control firms demonstrate a steady increase, firms that underwent automation appear to hit a plateau around the time of the spike, which is then followed by a slight decline in employment.

Figure 12. Automation and employment. Mean values for treat and control firms

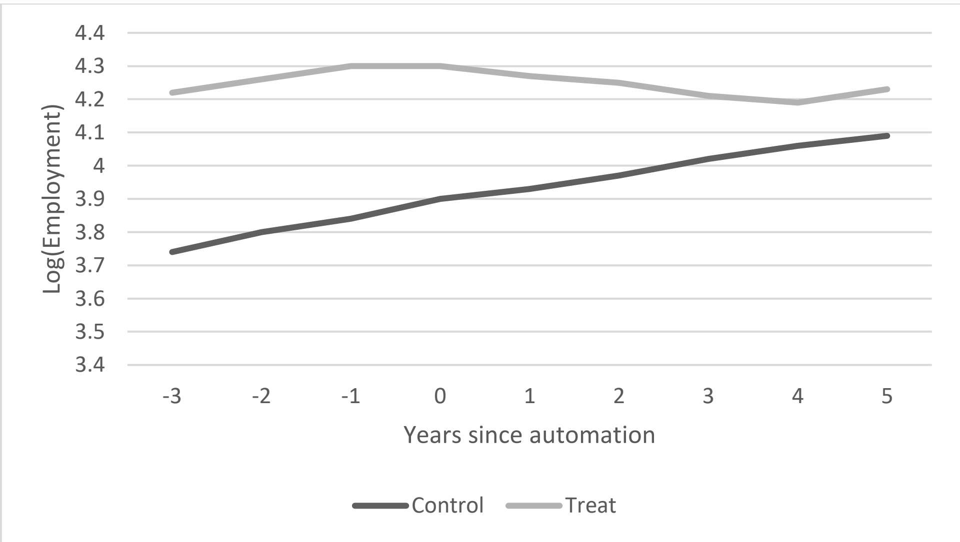

Note: Mean values of log(employment) for treat and control firms (year effects are not controlled for).

Next, Figure 13 presents DD-coefficients, from estimating equation (1), with log employment as dependent variable.

Figure 13. Automation and employment. Dependent variable log(employment)

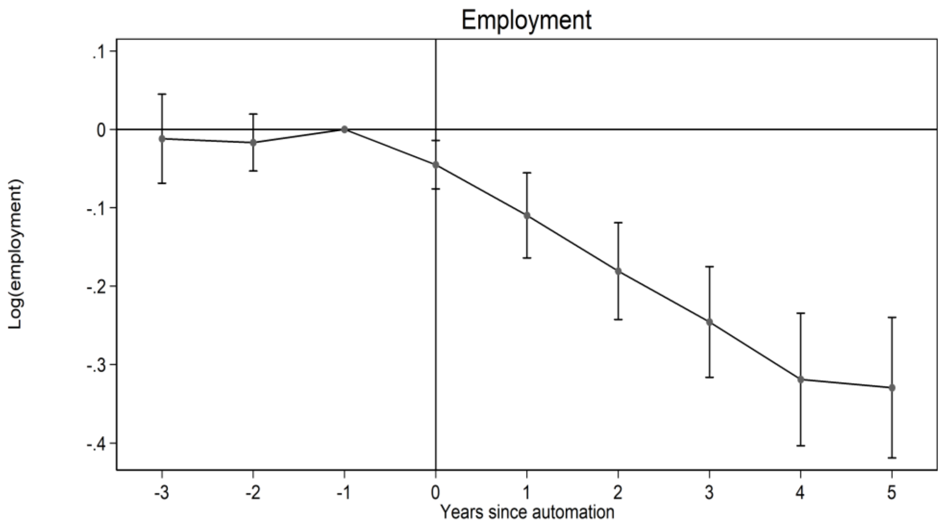

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

We find a clear negative effect of automation on firm-level employment. Five years after the spike, treatment firms have approximately 30 per cent lower employment compared to the control group of firms, that spike minimum five years later. The two groups do not trend differently in the pre-period, which is again reassuring.

Negative employment effects of robotisation are for example found in Acemoglu and Restrepo (2019) but using more aggregate level measures of robotisation. More comparable to our results are firm-level results in Bessen et al (2025). Our negative effects are somewhat larger than results reported in Bessen et al (2025). They find that in firms with fewer than 500 workers, employment is reduced following an automation event by about 20 per cent compared to firms that have an automation event later.[9]

Together with the offshoring results, the employment results suggest that automation replace workers and tasks abroad, but the results do not give any strong support for the reshoring hypothesis, i.e., that automation leads to creation of new domestic jobs. In the following we present evidence for different domestic group of workers with typically different tasks, and we conduct separate analyses depending on the pre-automation spike level of firm-offshoring.

Automating the production process may be especially detrimental to workers conducting tasks that are more easily automatable. Figure 14 presents DD-coefficients for employment for four different groups of workers, constructed to reflect different degrees of exposure to automation.: i) High skilled (university or college degree), ii) Low skilled workers (upper secondary education or lower), iii) STEM workers, and iv) Blue-collar workers. STEM workers and Blue-collar workers are defined from ISCO occupational codes: *STEM*: ISCO-08 code21 (Science and engineering professionals) and ISCO-08 code 31 (Science and engineering associate professionals). *Blue collar* workers: ISCO-08 code 7 and 8 (<83).[10] The employment variables are taken from individual employer-employ register data, merged with individual information on education and occupation, and then aggregated to number of employees for each group and firm, in each year.

[9] Firm employment can be constructed and taken from different administrative registers. In Appendix, Figure A7, we present results using three different measures of firm employment, taken from different registers, and also taken from own calculations, constructed from individual register information and aggregated up at the firm level. Irrespective of source and definition, all measures show the same pattern as in Figure 12.

[10] If the firm does not have any employees in the specific group, employment is set to 0, and Ln Employment is 0+1.

Figure 14. Automation and employment. Different groups. Dependent variable log(employment)

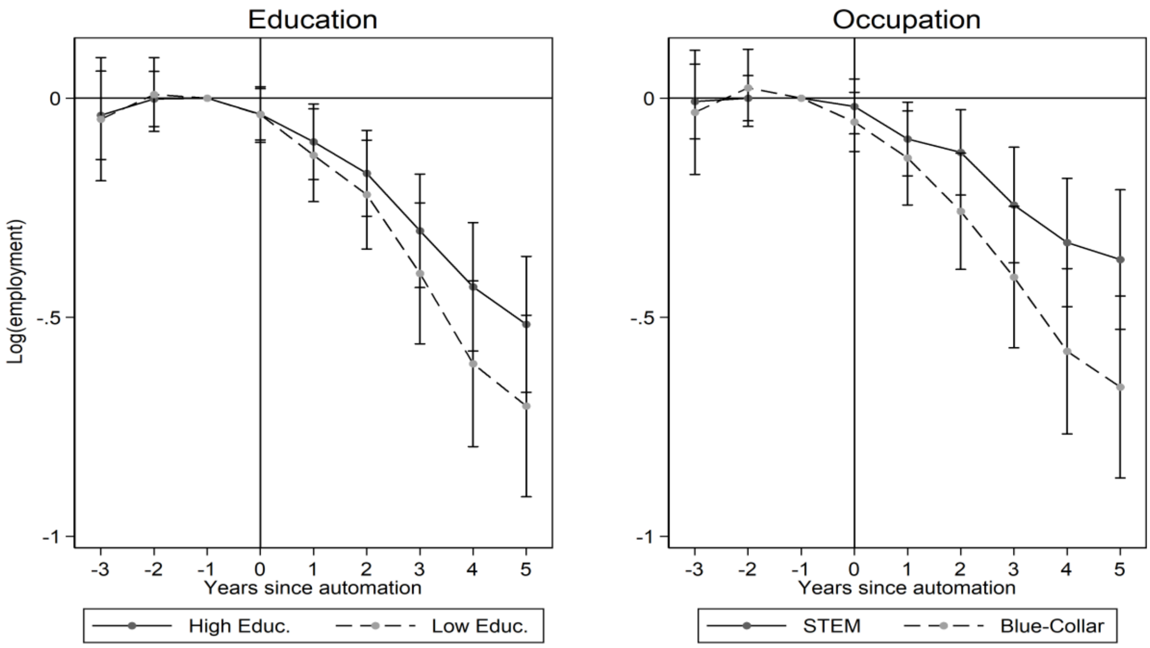

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

First, pre-trends do not trend differently between treat and control groups, which again is reassuring. Second, we find employment reductions for all four groups in the years following the automation spike, but the reduction is larger for low educated workers and blue-collar workers. Blue collar workers and workers with low education typically perform tasks that are more easily automated, and the results suggest that these workers take the hardest hit with respect to their employment development.

Our findings indicate that automation has negative impacts on both offshoring and employment. Our finding of a negative effect on domestic employment does not support the theory of reshoring. For a more comprehensive understanding of this issue, we estimate employment effects for different *pre*-spike measures of offshoring. We split firms in four equally sized quartiles, depending on level of pre-spike level of offshoring, from low to high

(1^{st}q: 0-25; 2^{nd}q: 25-50; 3^{rd}q: 50-75; 4^{th}q: above 75). The first quartile also include firms with zero offshoring in the pre-period. We use two types of pre-measure of offshoring: a level measure and a relative measure (offshoring measured relative to pre-automation spike employment).

Figure 15 and Figure 16 present the results using the two types of pre-spike offshoring measure. If reshoring is a contributing factor, meaning that a portion of the decrease in foreign employment is offset by an increase in domestic employment, we would anticipate a smaller reduction in employment among firms that were heavily engaged in offshoring before the automation spike. That is because for these firms, if reshoring was indeed the chosen strategy, automation would create a greater incentive to replace foreign with domestic labour, compared to firms with less or no prior offshoring.

Still, this is not what we observe. We find negative employment effects in all four quartiles, but the reduction in employment is not smaller among firms in higher quartiles. This adds further weight to the previous interpretations; that automation reduces offshoring but does not come with reshoring.

Figure 15. Automation and employment. Dependent variable log(employment). Level pre-spike measure in offshoring

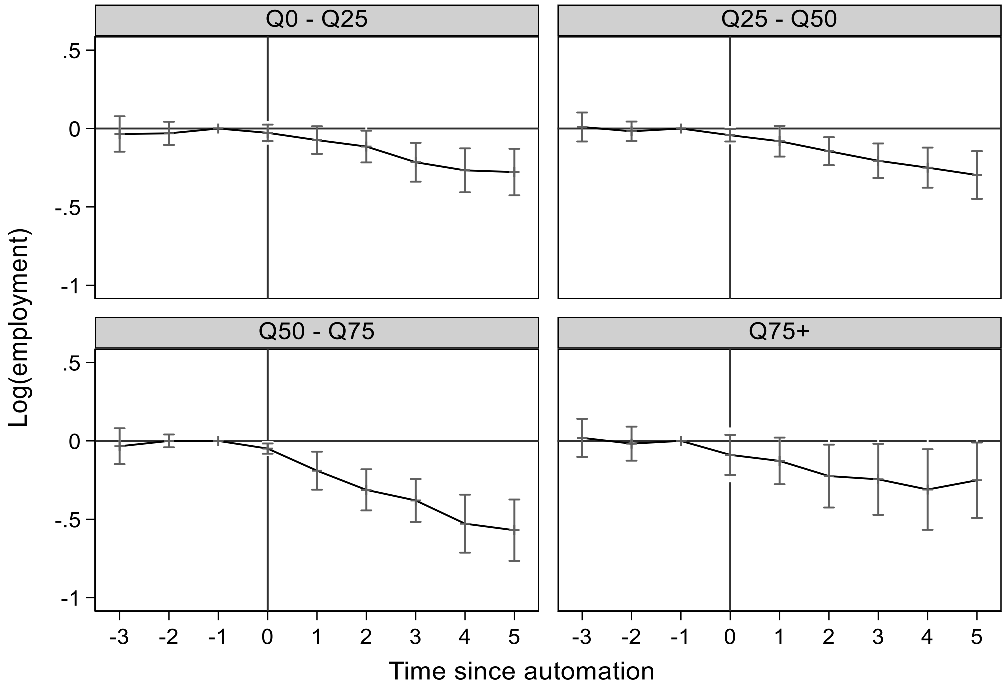

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure 16. Automation and employment. Dependent variable log(employment). Relative pre-spike measure in offshoring

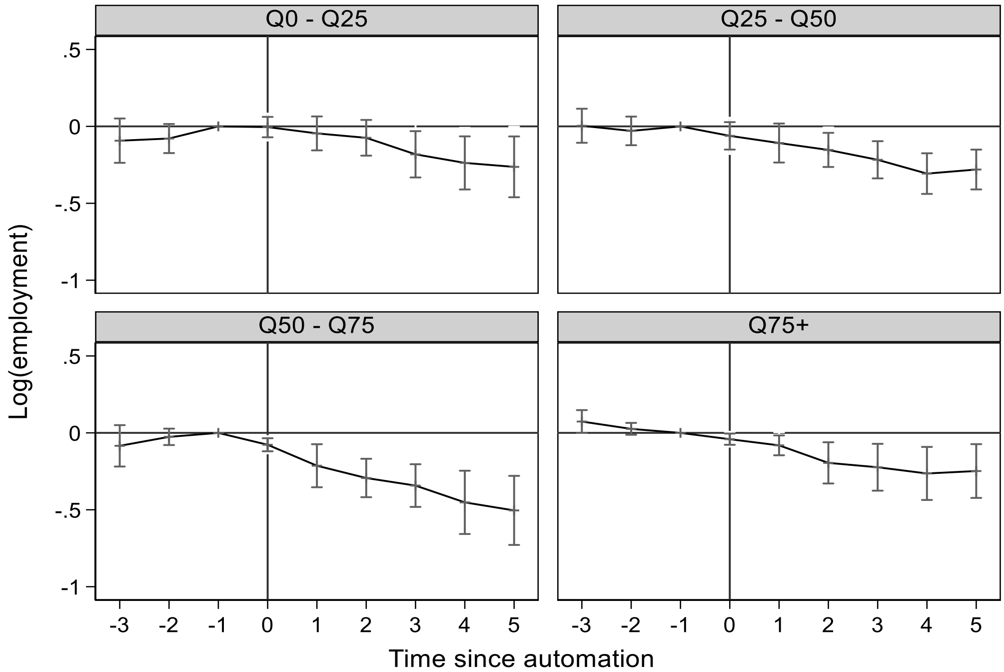

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

A closely related question to employment results, is what happens to firm productivity. Figure 14 presents evidence on this question. We construct a proxy for firm productivity with firm income per employees. As number of employees, we use the measure we use in Figure 13. Firm income is gross income minus cost of goods sold, i.e., subtracted the value of goods and services that are inputs in production.

Figure 14. Automation and firm productivity. Dependent variable: Net Firm income/Number of employees

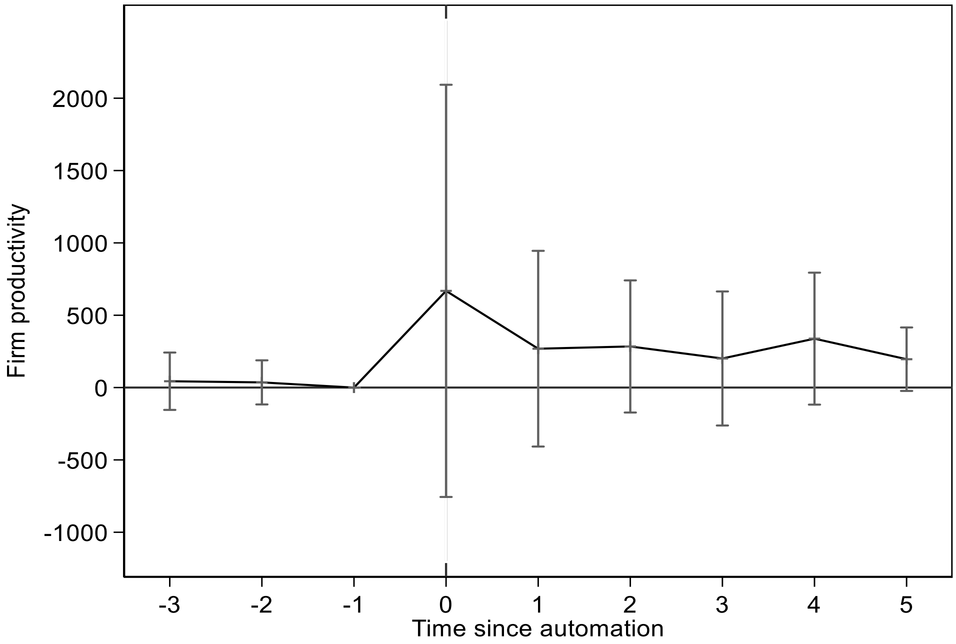

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Again, the pre-trends are small and not significant. The point estimates in the post-period are on the positive side, i.e., above zero, but the effects are not significant. Therefore, the results suggest that automation does not have a significant impact on firm productivity.

Hires and separations.

We end by analysing how the negative employment development can be seen through the pattern of hires and separations, before and after the automation spike. Hires and separations are defined year by year, by comparing the individual-firm identifier each year. For example, a person is defined as hired if registered at firm j in year t+1, but not in year t. Figure 17 presents the DD-results for alle hires and separations, Figure 18 for two groups of workers: STEM workers and blue-collar workers.

Figure 17. Number of hires and separations

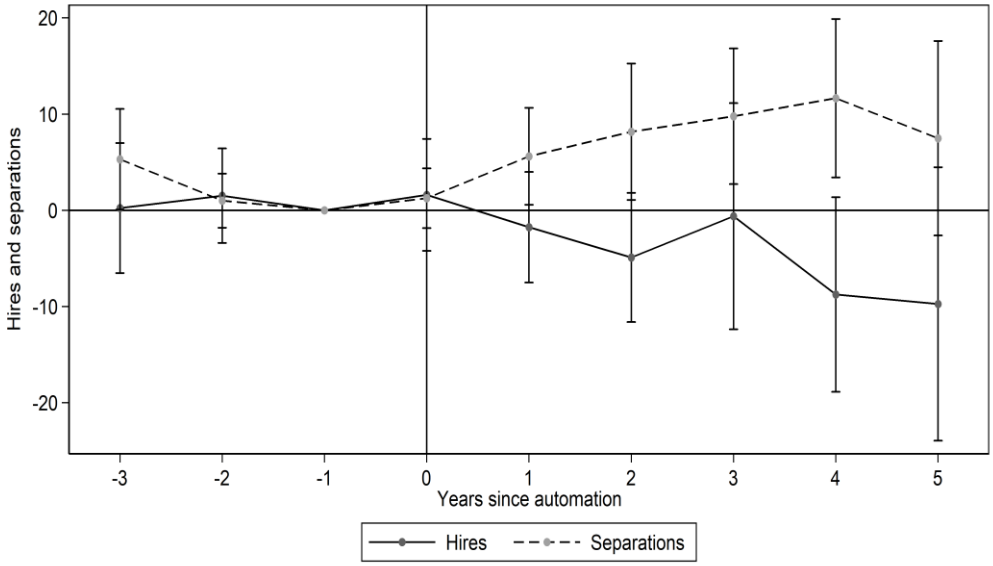

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

We find a statistically significant increase in the number of separations in the post-period, and a smaller decrease in hires. Therefore, this result suggests that increases in separations is an important driver for the previously reported negative employment results.

Figure 18. Number of hires and separations. Blue collar and STEM workers

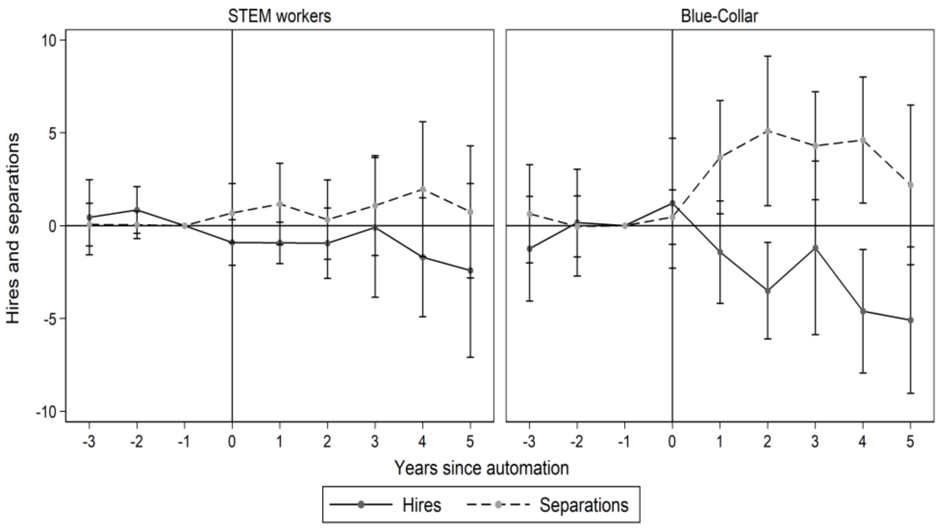

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Both hires and separations are statistically significant for blue-collar workers; number of hires drop in the post-period, and separations increase. For STEM-workers, the pattern of hires and separations are less responsive, and not significant in the post-period.

5. Conclusions

A characteristic feature of globalisation in the last decades is the increase in offshoring (see for example Feenstra and Hanson 2003, Hummels et al. 2018). As a consequence, parts of production in high-income countries are moved offshore, often to low-income countries. However, the rapid increase in automation technologies and their associated increases in production could increase incentives to retain a larger share of production domestically. The choice between producing abroad or domestically may, therefore be swayed by automation, potentially leading to a relocation of production processes back to the home country.

At the same time automating the production process may displace certain workers, especially workers conducting tasks overtaken by the automation process. A key question is, therefore, which workers are potentially displaced when automation occurs. In this paper we look closer at the relationship between domestic workers and workers located abroad.

In this paper we analyse the effects of automation on offshoring, using detailed matched employer-employee data for the period 2000-2018, matched with firm-level import data on both offshoring and automation. The import data enables us to distinguish between different types of offshoring, and between different offshoring countries.

Methodologically, we employ a difference-in-differences model, limiting treat and control firms to all-automation firms. By limiting the sample to firms that have at some point experienced automation, we use variation in event-timing, rather variation between those who automate and those who do not to identify the effect of automation on offshoring and employment. This should increase the likelihood of the treatment and control groups being on similar rends. Identification of treatment is based on a spike measurement approach, defining a treatment event when the relative automation cost is three times the firm average, measured over the whole period.

Irrespective of offshore indicator used, the results show a negative effect of automation on offshoring, in the post-period. These results are consistent with the notion that automation substitutes tasks that previously were undertaken abroad. The results are robust to different definition of spike measures and is stable across different firm sizes. And, also reassuringly, we do not find that treat and control group of firms are trending differently in the pre-period.

If automation typically replace workers abroad that undertake automation competing tasks, we expect the reduction in offshoring to be largest in low-income countries, typically non-OECD countries. Splitting the offshore measure by country, we find that automation reduces offshoring from non-OECD countries, but this pattern is also present among OECD-countries. Therefore, automation seem to reduce offshoring in both low- and high-income countries.

Our results also show that automation reduce domestic employment. Five years after the spike, treatment firms have more than 30 per cent lower employment compared to the control group of firms, that spike five years or later. The negative employment effects are quite sizeable, and larger than for example results for example in Bessen et al (2025). The employment reduction is largest among blue-collar workers and low educated workers, workers that undertake task that potentially are most easily to automate. Offshoring and employment results combined, suggest that automation reduces employment abroad, but does not spur domestic employment.

References

Abbring, J. H. and Van Den Berg, G. J. (2003). The nonparametric identification of treatment effects in duration models. *Econometrica*, 71:1491–1517.

Acemoglu D, and P. Restrepo (2019), "Robots and jobs: Evidence from US labor markets." *Journal of Political Economy*.

Acemoglu, D., C LeLarge, and P Restrepo (2010), "Competing with Robots: Firm-Level Evidence from France," AEA Papers and proceedings, 110:

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

Aghion,P,C. Antonin,S. Bunel, and X. Jaravel (2020)."What Is the Impact of Automation on Employment? New Evidence from France,"CEPR DiscussionPaper14443.

Aghion,P,C. Antonin,S. Bunel, and X. Jaravel (2021),"The Effects of Automation on Labor Demand: A Survey of the Recent Literature," Working Paper, 2021.

Artuc, E, L Christiaensen, and H Jinkler (2019), "Does Automation in Rich Countries Hurt Developing Ones? : Evidence from the U.S. and Mexico," Technical Report 8741, The World Bank February.

Baker, A., D. F. Larcker, C. C. Y. Wang (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144:370-395.

Bessen, J., M. Goos, and A. Salomons (2025), What Happens to Workers at Firms that Automate? *The Review of Economics and Statistics*, 107: 125–141 https://scholarship.law.bu.edu/faculty_scholarship/584

Bonfiglioli, A., R. Crino, G Gancia, I. Papadakis (2022), Robots, offshore and welfare. In: Robots and AI. A new economic area (ed by L Y Ing, and G Grossman)

de Chaisemartin, C. and D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110:2964–96.

Deshpande, M., and Y Li (2019), Who Is Screened Out? Application Costs and the Targeting of Disability Programs, *American Economic Journal: Economic Policy,* 11:213-248.

Faber,M. (2020).“Robots and Reshoring: Evidence from Mexican Labor Markets,”Journal of International Economics,127.

Feenstra, R. C. and G. H., Hanson, 1999. The impact of outsourcing and high-technology capital on wages: Estimates for the United States. The Quarterly Journal of Economics, 114:907–940.

Feenstra, R., and G. Hanson, 2003. Global Production Sharing and Rising Inequality: A Survey of Trade and Wages. In Handbook of International Trade, edited by E. Kwan Choi and James Harrigan, 146–85. Malden, MA: Blackwell Publishing.

Grossman, G. and E. Rossi-Hansberg 2008. Trading Tasks: A Simple Theory of Offshoring. American Economic Review 98:1978-1997.

Hummels, D., R., Jørgensen, J. Munch, and C. Xiang, 2018. The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data. American Economic Review, 104: 1597-1629.

Humlum, A. (2020), “Robot Adoption and Labor Market Dynamics,” Manuscript.

Koch, M, I Manuylov, and M Smolka (2019), “Robots and Firms.” Manuscript.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225:175–199

Appendix

Table A1. Descriptive statistics. Mean values, by event cohort

Cohort	2002		2004		2006		2008		2010		2012	
	Treat	Control	Treat	Control	Treat	Control	Treat	Control	Treat	Control	Treat	Control
Log Offshoring:												
Narrow	14.27	11.88	13.47	12.64	13.63	12.60	12.17	11.39	13.38	11.89	12.70	12.06
Intermediate	13.75	11.82	14.21	12.29	14.22	12.90	14.28	12.69	13.83	12.72	14.33	13.41
Narrow and intermediate	4.75	4.57	4.86	5.17	5.46	5.81	4.08	4.84	3.78	5.56	4.58	5.74
Automation costs	2277444	4183099	1163473	4432347	3901529	5144429	3664423	5553081	1248590	2136083	5731734	2242629
Number of employees	201.4	128.9	163.2	128.5	139.9	140.391	168.8	139.4	196.7	141.5	199.1	250.9

Note: Mean values at t=0, for each event cohort. Automation costs for the control is measured at the time the cohort firm has its spike in automation costs

A.1. Automation codes

We identify automatising costs based on the Norwegian Trade Statistics Register, which records all cross-border goods purchased by firms. For each year, we have the value of imports disaggregated by products. Trade flows are reported according to the 8-digit Harmonised system (HS) combined nomenclature. Concretely, we follow the categorization of Acemoglu and Restrepo (2022) and Bessen et al. (2025) and include automatically controlled machines, automation transfer machines, automatic welding machines, numerically controlled machines, and robots. Examples of descriptions of automatically controlled machines are "Automatic regulating or controlling instruments and apparatus"; examples of automatic transfer machines are "Continuous-action elevators and conveyors, for goods or materials"; examples of automatic welding machines are "Machines and apparatus for arc (including plasma arc) welding of metals"; examples of numerically controlled machines are "Numerically controlled bending, folding, straightening or flattening machines (including presses)"; and robots are described as "Industrial robots, not elsewhere specified or included".

Automation costs are defined from the following codes:

• *Automatically controlled machines*: 90321080, 90321000, 90328100, 90320000, 90321020, 90328900, 90328100, 90329000, 90322000

• *Automatic transfer machines*: 84283100, 84283900, 84573090, 84283300, 84283200, 84283990, 84580000, 84283100, 84283920, 84573000, 84573010

• *Automatic welding machines*: 85153100, 85153100, 85152100, 85152100

• *Numerically controlled machines*: 845811000080, 845811200080, 845811410010, 845811410080, 845811490080, 845811800080, 845891000010, 845891000080, 845891200080, 845891800080, 845921000010, 845921000080, 845931000010, 845931000080, 845941000010, 845941000080, 845951000010, 845951000080, 845961000010, 845961000080, 845961100080, 845961900080, 846012000010, 846012000080, 846022000010, 846022000080, 846023000080, 846024000080, 846031000010, 846031000080, 846040100080, 846221000010, 846221000080, 846221100080, 846221800080, 846231000010, 846231000080, 846241000010, 846241000080, 846241100080, 846241900080

• *Robots*: 84795000

Figure A1. The five largest industries for each automation component

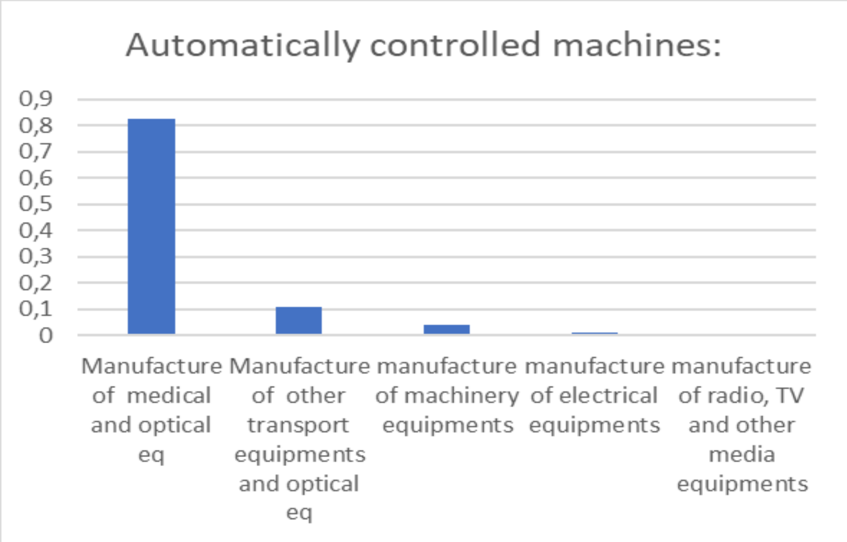

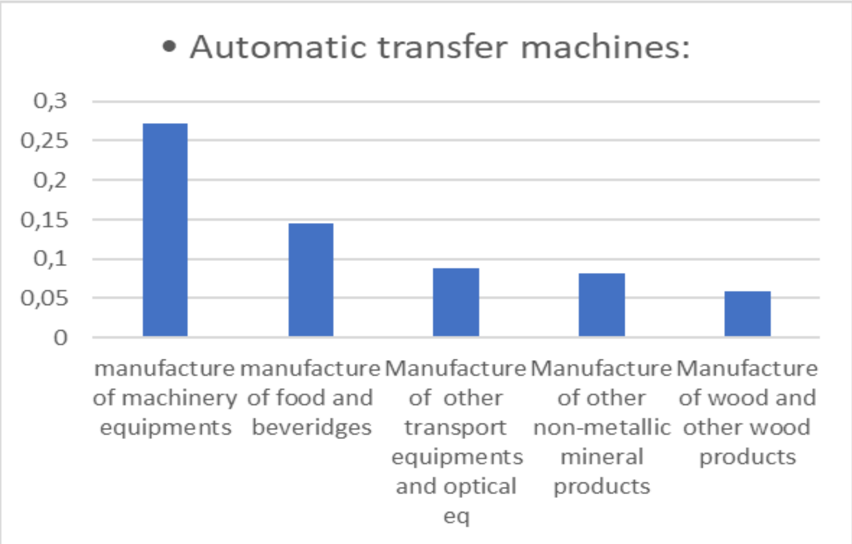

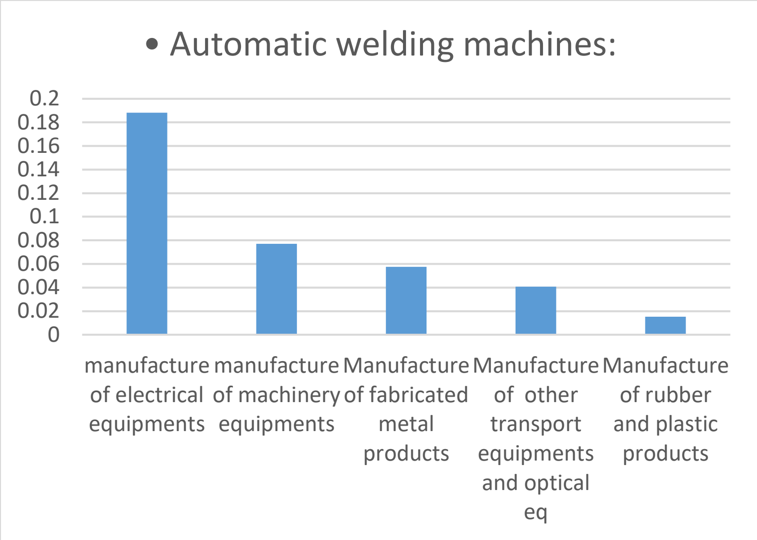
• Automatic welding machines:
0.2
0.18
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
manufacture of electrical equipments
manufacture of machinery equipments
Manufacture of fabricated metal products
Manufacture of other transport equipments and optical eq
Manufacture of rubber and plastic products

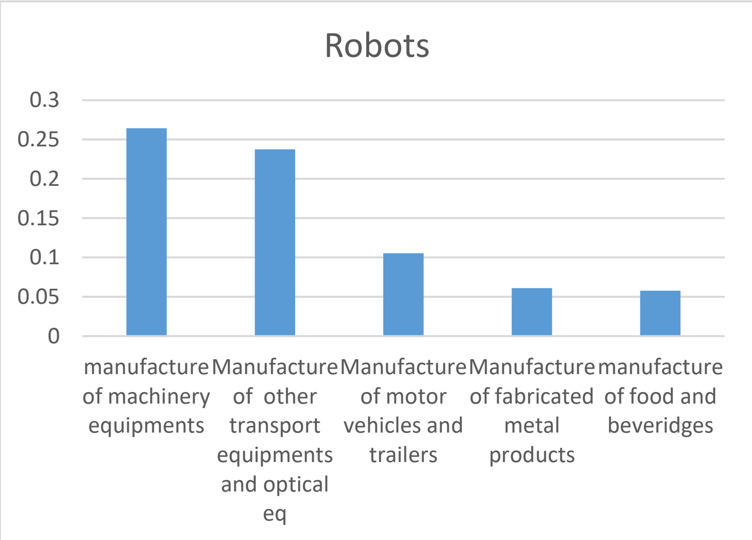
Robots
0.3
0.25
0.2
0.15
0.1
0.05
0
manufacture of machinery equipments
Manufacture of other transport equipments and optical eq
Manufacture of motor vehicles and trailers
Manufacture of fabricated metal products
manufacture of food and beveridges

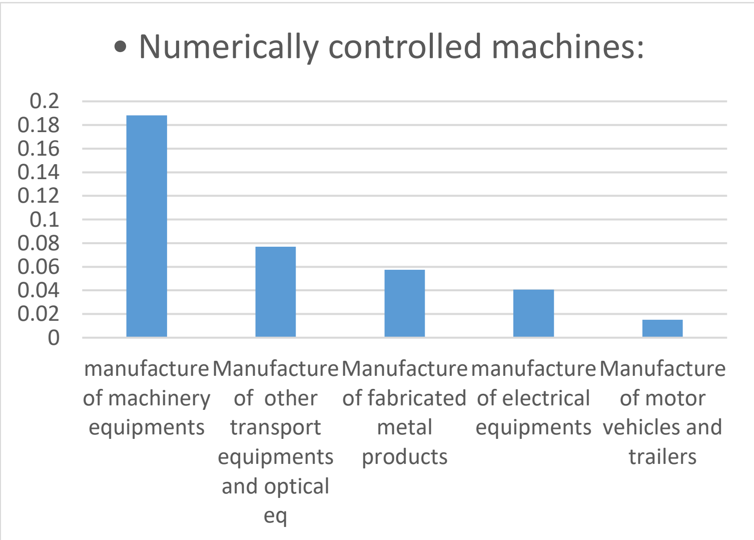
• Numerically controlled machines:
0.2
0.18
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
manufacture of machinery equipments
Manufacture of other transport equipments and optical eq
Manufacture of fabricated metal products
manufacture of electrical equipments
Manufacture of motor vehicles and trailers

Table A2. Automation and offshoring. Dependent variable: Ln(offshoring)

	(1) Narrow offshoring	(2) Intermediate offshoring	(3) Intermediate and narrow
-3	0.109	-0.135	-0.0351
	(0.297)	(0.235)	(0.255)
-2	-0.0106	-0.0622	0.0710
	(0.245)	(0.195)	(0.207)
0	0.331	0.226	-0.0162
	(0.211)	(0.163)	(0.194)
1	0.122	-0.211	-0.117
	(0.263)	(0.184)	(0.231)
2	-0.395	-0.873***	-0.471*
	(0.321)	(0.221)	(0.277)
3	-0.955**	-1.358***	-0.567*
	(0.390)	(0.256)	(0.322)
4	-1.987***	-1.978***	-1.125***
	(0.444)	(0.289)	(0.352)
5	-2.228***	-2.799***	-1.295***
	(0.481)	(0.329)	(0.381)
N	24857	24857	24857
R^2	0.693	0.657	0.761

Note: Standard errors in parentheses add control: firm fixed effects * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3. Automation and offshoring. Dependent variable: p(offshoring)

	(1) Narrow offshoring	(2) Intermediate offshoring	(3) Intermediate and narrow
-3	0.00544	0.000532	-0.0100
	(0.0210)	(0.0173)	(0.0201)
-2	-0.0132	0.00469	-0.00191
	(0.0178)	(0.0151)	(0.0167)
0	0.0166	0.00103	-0.00204
	(0.0151)	(0.0116)	(0.0160)
1	-0.00126	-0.00343	-0.00662
	(0.0184)	(0.0133)	(0.0185)
2	-0.0323	-0.0332**	-0.0283
	(0.0215)	(0.0159)	(0.0213)
3	-0.0549**	-0.0647***	-0.0314
	(0.0253)	(0.0174)	(0.0244)
4	-0.129***	-0.102***	-0.0754***
	(0.0287)	(0.0199)	(0.0266)
5	-0.136***	-0.142***	-0.0852***
	(0.0309)	(0.0224)	(0.0283)
N	24857	24857	24857
R^2	0.603	0.518	0.710

Note: Standard errors in parentheses add control: firm fixed effects * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A2. Automation and offshoring. Dependent variable: Ln(offshoring). Intensive margin; using only firms with positive values on offshoring

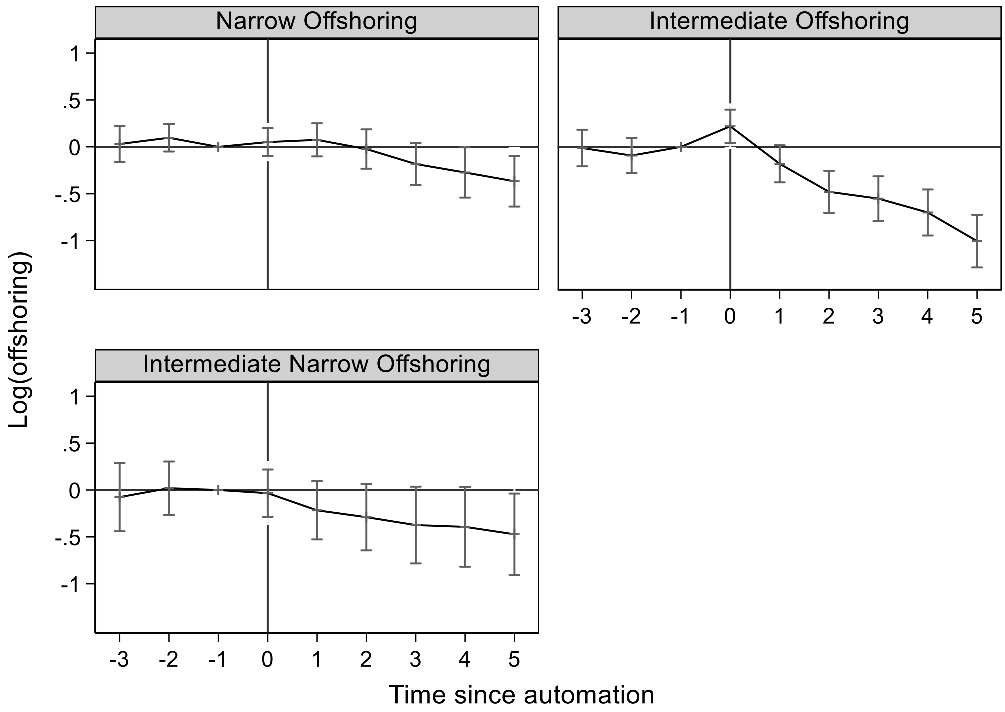

Figure A3. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Median definition of spikes

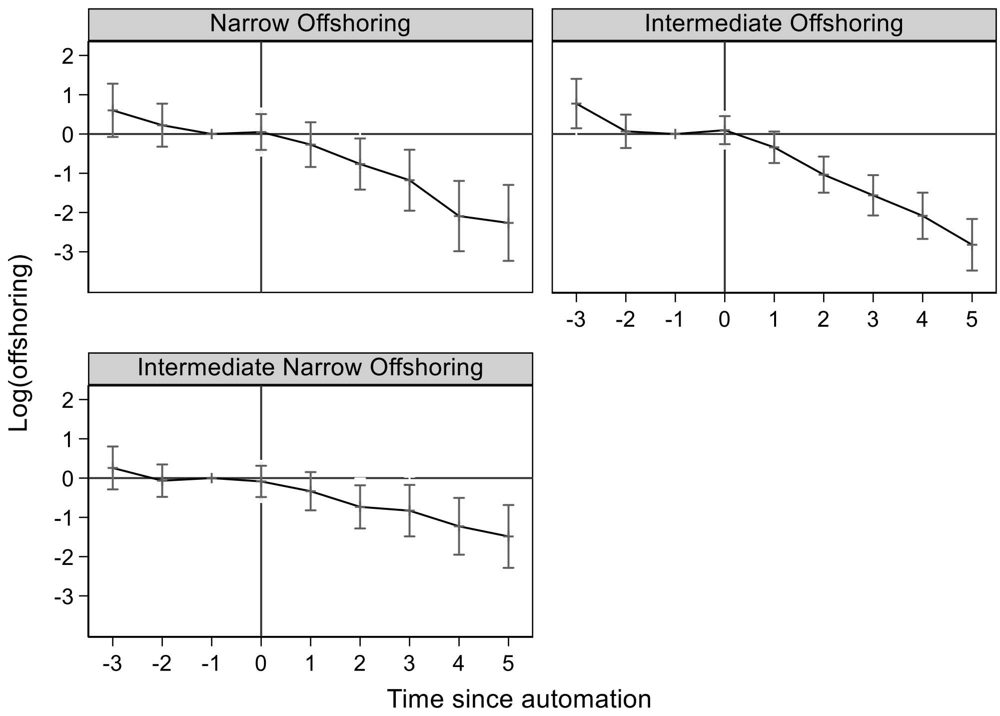

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure A4. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Level of automation definition of spikes

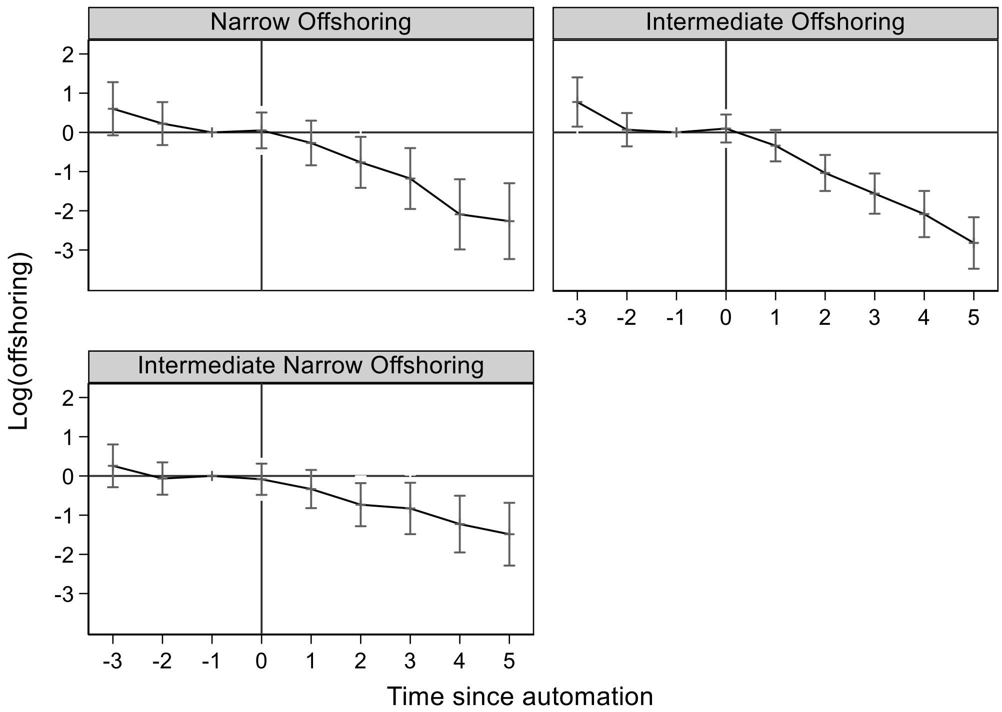

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure A5. Automation and offshoring. Dependent variable log(offshoring). Narrow, intermediate, and narrow and intermediate offshoring. Four times the average spike

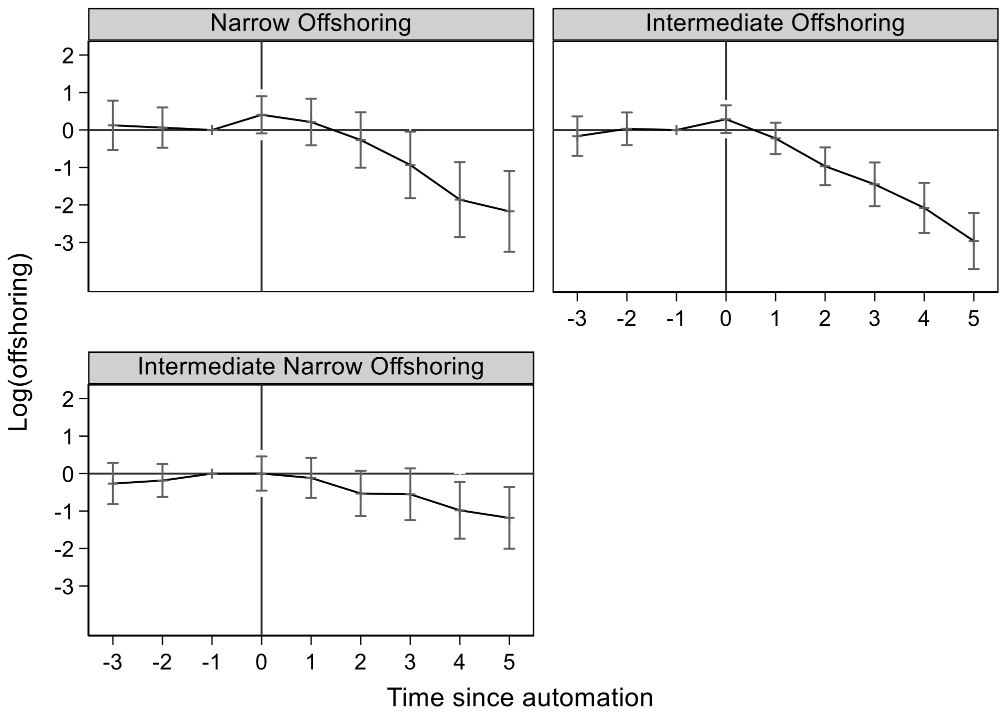

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure A6. Automation and offshoring. Dependent variable: Ln(offshoring) using the Inverse Hyperbolic Sine Transformation

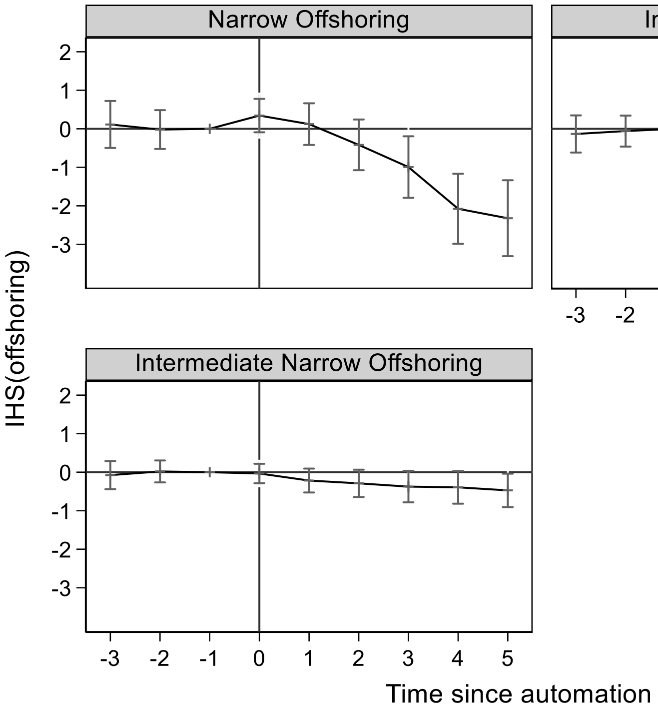

Note: DD-coefficients from estimating equation (1), with 95% confidence intervals

Figure A7. Automation and employment. Different measures of employment

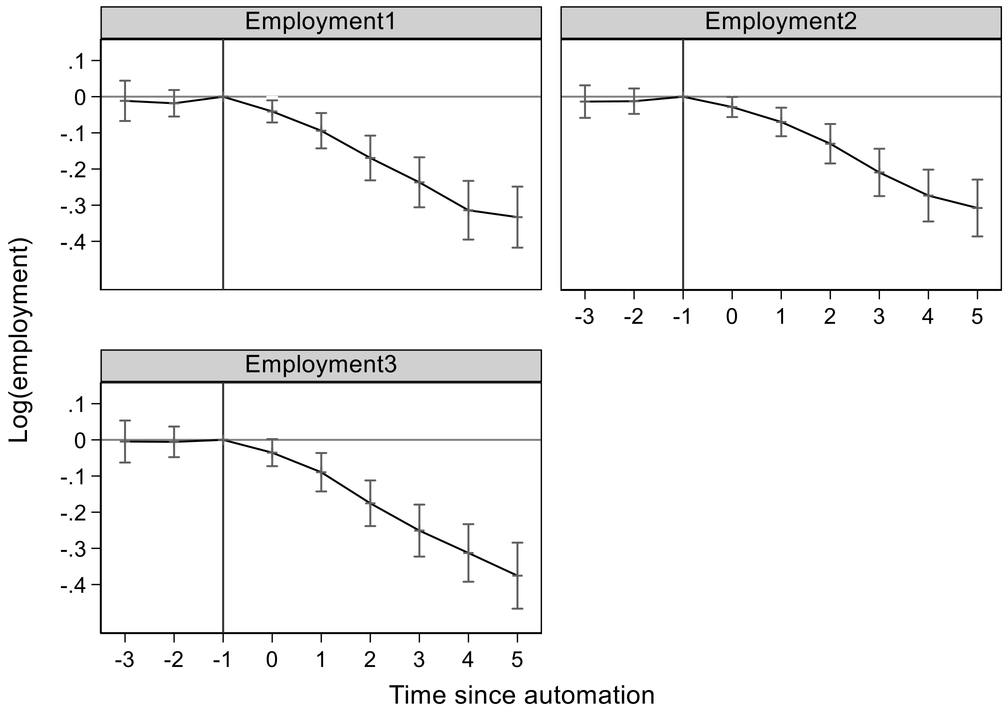

Note: Employment1 is based on information from “Enhetsregisteret”, Employment2 is based on Strukturstatistikken , Enployment3 is based on own calculation, from individual information from Employer-employee register data, and aggregated up to firm level.