

**The Dynamics of Automation
Adoption: Firm-Level
Heterogeneity and Aggregate
Employment Effects**

Laura Bisio, Angelo Cuzzola, Marco Grazzi, Daniele Moschella

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

The Dynamics of Automation Adoption: Firm-Level Heterogeneity and Aggregate Employment Effects

Abstract

We investigate the impact of investment in automation-related goods on adopting and non-adopting firms in the Italian economy during 2011-2019. We integrate datasets on trade activities, firms', and workers' characteristics for the population of Italian importing firms and estimate the effects on adopters' outcomes within a difference-in-differences design exploiting import lumpiness in product categories linked to automation and AI technologies. We find a positive average adoption effect on the adopters' employment and on the value-added and average wage, whereas sales and productivity increase after an initial drop with a net positive effect five years after adoption. Crucially, the employment effect is heterogeneous across firms: a positive scale effect is predominant among small firms, whereas a negative displacement effect is predominant among medium and large firms. We complete the framework with a 5-digit sector-level analysis showing that adopting automation technologies has an overall negative effect on aggregate employment.

JEL-Codes: D240, J230, L250, O330.

Keywords: automation, employment, firm heterogeneity, imports, technology adoption.

Laura Bisio
Italian National Institute of Statistics
Rome / Italy
bisio@istat.it

Marco Grazzi
Department of Economic Policy, Università
Cattolica del Sacro Cuore, Milan / Italy
marco.grazzi@unicatt.it

Angelo Cuzzola
Institute of Economics & Department
L'EMbeDS, Scuola Superiore Sant'Anna
Pisa / Italy
angelo.cuzzola@santannapisa.it

Daniele Moschella
Institute of Economics & Department
L'EMbeDS, Scuola Superiore Sant'Anna
Pisa / Italy
daniele.moschella@santannapisa.it

1 Introduction

The impact of automation technologies on employment has been a topic of discussion in the economic literature for a long time, from Ricardo to recent debates about robots and the factory of the future.¹ As the world economy is entering a new era of technological revolution, shaped by the pervasive diffusion of robots and Artificial Intelligence and affecting the global organization of jobs, tasks, activities, and industries (Brynjolfsson and McAfee, 2014; Frey and Osborne, 2017; Acemoglu and Restrepo, 2019; Montobbio et al., 2022), the question remains: will it be true that "humans need not apply"? Looking for an answer to this question in recent waves of automation technologies, the empirical literature offers mixed evidence. On the one hand, recent studies based on country and industry-level data tend to show both positive and negative effects (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Grigoli et al., 2020; Dottori, 2021). On the other hand, emerging empirical evidence at the firm level suggests that firms adopting automation technology tend to increase their employment (Koch et al. 2021; Acemoglu et al. 2020; Aghion et al. 2020; Domini et al. 2021), but some studies also find negative firm-level effects (Bessen et al. 2020; Bonfiglioli et al. 2020).

This mixed evidence is not surprising. Economic theory shows that the effects of investments in automation technologies on employment are ambiguous, as both substitution and complementarity forces are at play. In their task-based model, Acemoglu and Restrepo (2022) posit that an initial displacement effect triggered by the adoption of new automation technology (for instance, the introduction of robots) can be counteracted by several forces operating at the economy-wide, sector- or firm-level. At the level of the individual firm, productivity or reinstatement (creation of new tasks) effects might induce firms to expand their production scale and increase labour demand. A similar positive compensation force is the Schumpeterian mechanism of new product introduction, as shown in Dosi et al. (2021). Moreover, as firms exhibit heterogeneous patterns of adoption of new technologies, the sectoral and economy-wide effects of automation technologies are the results of the interaction between different adopting firms and between them and non-adopting firms at each point in time.

In this paper, we provide novel empirical evidence that the effects of investments in automation technologies can be heterogeneous across adopting firms, and that this has important aggregate implications. Specifically, we study the effects of automation technologies at the firm and industry level, using a novel dataset created by integrating multiple official sources from the Italian National Institute of Statistics (ISTAT) on international trade, firm

¹ The body of scholarly research on the employment consequences of automation and technological change is enormous. In-depth discussions are provided in Freeman et al. (1982), Pianta (2005), Autor (2015), among many others.

structural characteristics, balance sheet information, and labour force characteristics.

First, we exploit granular product-level information for traded goods to measure the adoption of automation-related technologies across the population of Italian importing firms, which employ around 40% of the total Italian labour force. Leveraging the fact that firms tend to concentrate their investments in such technologies in single spikes, we set up a difference-in-differences (DiD) framework to investigate the evolution of employment and other firm-level characteristics around the spike. Our identification strategy relies on the variations in adoption timing within adopters and on the comparison between adopting and non-adopting firms. Following the recent advancements related to the DiD econometric literature, we adopt the estimator proposed by [Callaway and Sant’Anna \(2021\)](#) to avoid biases due to the presence of ‘non-clean’ controls, to relax the classical ‘parallel trend assumption’ so that it is possible to assume its validity only conditionally on selected firm-level covariates, and to control for potential anticipating behaviours by the adopters. We consider the dynamics of employment variables, including potential differences in impact across occupational categories and types of contracts, as well as the impact on average wage and wage inequality within adopting firms. Our analysis shows that, on average, adopting firms significantly increase their employment after the adoption event.

Second, we move from studying the average effect of the adoption process on individual adopters to investigating its wider implications on aggregate employment. To do so, we start by examining the impact of adoption on the group of adopters, considering potential heterogeneous effects among different size classes. We find that the positive scale effect is prevalent among small firms, whereas the displacement effect is prevalent among medium and large firms. Subsequently, we analyze the impact on the entire population of importers at the industry level. The results show that adopting automation technologies has an overall negative effect on aggregate employment at the sectoral level, which can be explained by considering the differential effects across various size classes, as observed in the firm-level analysis.

The validity of these findings is corroborated through a series of robustness checks. In particular, we control that our estimates are not driven by firms acting as intermediaries in international trade transactions; we control for adopters’ anticipation; we test different definitions of automation spike to check whether specific technologies drive the results; we exclude the possibility that the presence of synchronous investments in generic machinery drives the automation effects.

Our work contributes to the growing body of literature on the granular effects of automation, exploiting both firm and worker-level data (see, among the latest, [Bessen et al., 2023](#); [Acemoglu et al., 2023](#)). We also aim to contribute to the literature on automation and wage inequality ([Aksoy et al., 2021](#); [Faia et al., 2022](#)) as well as the link between firm-level and industry-level effects of technology adoption ([Acemoglu et al., 2020](#); [Aghion et al., 2020](#)).

The remainder of the paper is structured as follows. Section 2 introduces the data sources, the main variables of interest and their integration into a unique firm-level dataset. Section 3 presents the results from the firm-level difference-in-differences exercise and their aggregate implications. Section 4 presents robustness checks. Section 5 briefly concludes and recaps.

2 Data construction and variables

2.1 Sources

The analysis is based on a novel dataset on Italian firms, which results from the integration of multiple official sources by ISTAT.

The first source is the *Commercio Estero* (COE) register, which contains transaction-level information on the population of Italian exporting and importing firms. Data are available as of 2005 onwards. For each transaction, the register reports the 6-digit Harmonized System product code (HS6), origin/destination country, quantity, and trade value. The information comes from the Single Administrative Document (S.A.D.) for non-EU countries and from the Intrastat models acquired by the Customs Agency (*Agenzia delle Dogane*) for EU countries. We use COE to identify importing firms and their imports of automation-related goods based on the 6-digit HS code.

The second source is the *ASIA-Employment Register* (ASIA-ER), which is a 'Linked Employer-Employee' (LEED) register based on multiple administrative data covering the universe of Italian active firms. The Register, which has been available since 2011, contains information on firms' structural characteristics (5-digit NACE sector, geographical information, size class) and on several employees' attributes, including gender, age, place of birth, contract type (temporary or permanent), occupation (blue-collar, white-collar, middle-managers and managers), working time regime (full-time or part-time), and attained educational degree (classified by ISCED 2011). Additional information on worker wages and tenure has been integrated by the *Labour cost of employees register* (LCER) by ISTAT. Finally, note that in the ASIA-ER, workers are classified as either employee or self-employed² but we exclude the latter from the sample as we are interested in the consequences of automation for the labour demand by firms.

Finally, we gather information on firms' economic performance from the *Structural Business Statistics Register* (Frame-SBS), which is based on both administrative and fiscal sources that are integrated with data from sample surveys on small and medium-sized enterprises (SMEs), for the Italian active firms. Data are released yearly as of 2011.

² According to the European System of Accounts, self-employed are people who work for themselves, i.e. the owners/partners of a company that is not distinct from the natural person(s) who owns it. In 2020, self-employed accounted for 26.5% of total employment by Italian active firms.

Frame-SBS Register includes companies' main income statement items: we will consider total revenues, wage bills, and value added.

2.2 Dataset

The integration of the three sources restricts the analysis to the 2011-2019 period. Since our measure of automation exploits only imported product categories (see Section 2.3), we start by identifying our reference sample as those firms importing at least once in the 2011-2019 period. Then, we perform a series of data-cleaning steps. First, we implement a cleaning strategy to prevent a structural break of the series of importers caused by a change in legislation in 2017. This legislation states that firms must report - for either fiscal or statistical purposes - the value of their transactions from/to an EU country if their amount is above a minimum threshold.³ To address this issue, we filtered the sample from 2018 backwards by excluding firms that imported less than €400,000 annually over the 2011-2017 period. As a consequence of the filtering process, the sample was reduced from 284,086 distinct importers observed in 2011-2019 to 245,191 firms. More details are provided in Appendix A, showing that the cleaning process leaves out marginal importers. Second, we excluded self-employed as discussed in 2.1 and atypical workers without standard employment contracts, inducing a removal of 56,765 (micro) firms.

Then, we removed intermediaries of automation goods to exclude firms importing and reselling such goods on international and/or domestic markets (8,117 firms).⁴ Finally, we are forced to further remove 7,024 observations from our panel dataset due to the lack of information on the NACE division for a subset of firms. As a result of the cleaning procedure, the final sample includes 179,933 active firms over the 2011-2019 period.

Table 1 displays the sample coverage, excluding self-employed individuals and automation intermediaries. During 2011-2019, importing firms represented less than 10% of active firms, but employed around 40% of the labour force in the economy and around 70% in manufacturing. If we consider firms importing automation goods in at least one year, they account for around 45% of manufacturing employment and 21% of total employment. The figures in the table highlight that there are no significant structural changes in the considered window.

³ According to Italian Law n.19 of 27 February 2017, as of January 2018, firms importing less than 200,000 euros per quarter (exporting less than 100,000 euros per quarter) are exempt from declaring the imported (exported) amounts for either fiscal or statistical purposes. These limitations apply only to intra-EU transactions.

⁴ Based on Bernard et al. (2015), we started from wholesalers, identified by Nace rev.2 group codes from 461 to 469. Among them, we identified the relevant 3-digit sectors as defined in Table A3.

Year	TOTAL ECONOMY			MANUFACTURING		NON-MANUF.	
	Importing Firms	Empl % Importers	Empl % Autom imp	Empl % Importers	Empl % Autom imp	Empl % Importers	Empl % Autom imp
2011	8.2	38.9	21.3	66.8	42.6	28.2	13.1
2012	8.3	38.2	21.1	67.4	43.3	27.6	13.0
2013	8.7	39.0	21.5	68.8	44.4	28.3	13.3
2014	8.9	39.6	21.8	69.3	45.1	28.9	13.5
2015	8.9	39.5	21.6	69.1	45.1	29.2	13.4
2016	8.9	39.5	21.6	69.1	45.2	29.4	13.5
2017	8.7	39.5	21.0	68.7	44.8	29.9	13.1
2018	8.5	39.5	20.8	68.6	44.8	29.9	12.9
2019	8.3	39.2	20.7	68.5	44.8	29.5	12.8

Table 1: Share of importing firms; employment share of importing firms; employment share of firms importing automation technology. Percentage shares (%) are reported for the total economy, manufacturing sectors, other non-manufacturing sectors. Time period: 2011-2019. Elaborations on International trade statistics and ASIA-Employment register.

2.3 Variables and descriptive statistics

Automation

Data on the adoption of digital and automation technologies at the the firm level is scant. National statistical offices have started to collect them only recently, and they are not yet included in main innovation surveys such as the Community Innovation Survey. The US Census Bureau introduced a new module in the 2019 Annual Business Survey to study the adoption of advanced technologies (Acemoglu et al., 2022), while Germany collects survey data on automation and robot adoption from 2016 (Deng et al., 2021; Benmelech and Zator, 2022). Notably, the Dutch statistical office (CBS) includes a question on automation costs in their national survey starting from 2000 (Bessen et al., 2023), and in Spain survey data on robot adoption are available from 1990 (Koch et al., 2021). The empirical literature has employed an alternative approach to work around this data limitation by leveraging administrative import data to identify firms that are adopting robots and other automation technologies. Import data have been used for Canada (Dixon et al., 2021), France (Acemoglu et al., 2020; Aghion et al., 2020; Bonfiglioli et al., 2020; Domini et al., 2021), Netherlands (Acemoglu et al., 2023), Denmark (Humlum, 2020), Norway (Barth et al., 2020), and this is also the approach we will follow in our work.

More specifically, we identify imports of goods that embed automation technologies following the taxonomy by Acemoglu and Restrepo (2018). Exploiting the 6-digit Harmonized System (HS) product code, they classify as automation-related imports of industrial robots, dedicated machinery, numerically controlled machines, and a number of other automated capital goods. Additionally, we include 3-D printers (Abeliansky et al., 2020) and other categories of imports expected to embed AI automation, namely, automatic data processing machines and electronic calculating machines. For a full list, see Table A5.


Description	Division	Automation share (%)	Employment share (%)
Manufacture of fabricated metal products, except machinery and equipment	25	7.8	5.0
Manufacture of machinery and equipment n.e.c	28	26.7	12.7
Wholesale and retail trade and repair of motor vehicles and motorcycles	45	5.1	0.8
Wholesale trade, except of motor vehicles and motorcycles	46	12.5	5.0
Computer programming, consultancy and related activities	62	7.6	1.7

Table 2: Top-5 divisions (2-digit NACE rev.2) with an automation share larger than their employment share, year 2019. Elaborations on International Trade and Frame-SBS data.

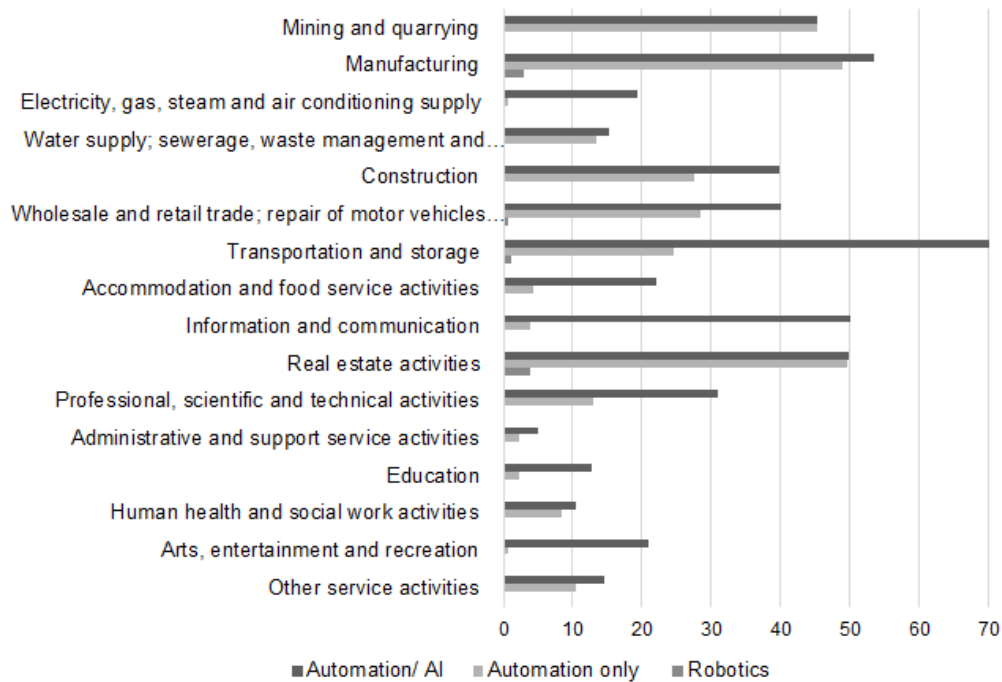
Capital goods are a highly traded category of manufactures (Eaton and Kortum, 2001), so our measure is able to capture large adopters of automation technologies (see Table 1). There are, however, some potential issues which we need to address. First, firms may purchase capital goods embedding automation technologies domestically rather than abroad, leading them to be incorrectly categorized as non-adopters. Limiting the analysis to firms involved in international trade reduces the likelihood that firms in our sample are incorrectly labelled as non-adopters. We further reduce this likelihood when we run our baseline estimations on the restricted sample of importing firms that buy automation goods from abroad for at least one year (adopters). In this setting, the estimates of the adoption effect are obtained by exploiting only the variation in event timing, defining a control group that includes only not-yet-adopters: that is, at any year, firms that have not yet adopted automation technologies but will adopt them later in the time window. Second, firms may use an intermediary rather than import goods directly (Blum et al., 2010; Bernard et al., 2010; Ahn et al., 2011); however, this is less likely for more complex goods (Bernard et al., 2015) that are highly relation-specific, such as those that comprise our measure. Finally, firms that import automation goods may resell them in the domestic or international markets. In our main sample, we removed intermediaries to minimize the risk of firms being incorrectly categorized as adopters (see Section 2.2). Additionally, as manufacturing firms are also known to engage in the (re)export of goods that they do not produce, a practice known as Carry Along Trade, CAT, (Bernard et al., 2019), we also apply some filters to exclude re-exporting firms in our main analysis (see below, Section 3.1).

Automation spikes

Our import-based measure of automation shows that all sectors have automation expenditures, though there is substantial heterogeneity both between and within sectors. Figures 1a and 1b report the share of firms and their employment importing automation goods (over



(a) Share of firms importing automation



(b) Share of employment at firms importing automation

Figure 1: Distribution of imported automated goods across economic sections, year 2019. Percentage values.

the total population of importing firms) in 2019. For each sector, both figures report the corresponding share for our most comprehensive measure of automation, including also goods

expected to embed AI automation ("Automation/AI"), which we will use in the paper; a more restrictive measure, excluding AI-related goods ("Automation only"), and the share related only to the import of robots ("Robotics"). Automation, including AI, is diffused across importing firms in the mining, manufacturing and electricity sectors, where the share of firms importing such technologies is around 20%. IT sector is also important, especially when focussing on the subset of AI-related goods. The figure also reveals that robots are still in their infancy in terms of diffusion, compared to other automation technologies, and that the employment share of adopting firms is heterogeneous across sectors. Table 2 looks at the intensive margin of our measure, reporting the top 5 sectors (at the 2-digits level) in terms of automation imports per employee: there are two manufacturing sectors (fabricated metal products and machinery and equipment), as well as service sectors, in particular, wholesale trade and computer programming.




Figure 2: Fraction of firms with positive imports of automated goods for a given number of years.

Moving to the statistical properties of automation-related imports, it can be seen that they exhibit the typical spiky behaviour of an investment variable (Letterie et al., 2004; Asphjell et al., 2014; Grazzi et al., 2016). This means that, first, among firms importing automation goods, more than 40% does it only once over the period 2011-2019 (see Figure 2), and the frequency drops linearly with greater values, with the exception of a tiny group of firms that import automated goods every year. Second, as shown in Figure 3, the largest yearly event of such imports represents a predominant share of a firm's total across years. When ranking the shares of each year's imports (out of all years) from largest to smallest, it is clear that the top-ranked import event displays a significantly high share (approximately 80% in mean), while lower ranks rapidly decrease in value. This evidence is similar to the import-based measure of Domini et al. (2021, 2022) and to the automation cost measure as in Bessen et al. (2023). Because of the skewed nature of this variable within firms, we identify




Figure 3: Investment shares by rank. Rank 1 is the highest yearly investment share in the firm’s timeline.

Size classes	No spike		Spike in automation	
	N. Firms	Share	N. Firms	Share
1-19	83,568	88.6	10,770	11.4
20-49	14,508	71.6	5,750	28.4
50-249	6,131	49.2	6,331	50.8
250+	760	29.3	1,838	70.7

Table 3: Spike in automation by size, year 2019. Elaborations on International Trade and FrameSBS data.

the greatest event as an automation spike for each firm. More specifically, an automation spike occurs for firm j in year t_i where $i \in [2011, 2019]$ when the cumulative value of imports in automation-related goods in that year is the largest compared to those observed in all the other years t_i of the interval.

Table 3 offers additional information on adopters and non-adopters, showing the size distribution of the different groups of firms. We divide firms into four size categories: 1-19 employees, 20-49, 50-249, and 250+. ⁵ The relative frequency of importing firms with a spike in automation goods is clearly increasing in size: among firms with less than 20 employees, it is relatively rare to observe a spike, whereas among firms with at least 50 employees, the percentage of adopting firms goes from 50% to more than 60%.

⁵ These size classes are the result of an aggregation of more granular size classes defined by the EU Regulation on Structural Business Statistics (see Commission Regulation (EU) 2019/2152 and Commission Implementing Regulation (EU) 2020/1197).

Outcome variables

The impact of investment in automation technologies is analysed with respect to a number of firm-level outcome variables in the years following the automation spike.

In particular, we consider a block of employment-related outcome variables: *i*) the log of total employment emp_j computed as the sum of annual jobs⁶ in the firm j ; *ii*) the share of, respectively, blue-collar, white-collar and middle and top-managers; *iii*) the share of fixed-term/permanent employees; *iv*) the share of full-time employees, *v*) the share of high-educated employees, defined as the share of employees holding at least a tertiary education degree according to ISCED 2011; *vi*) the share of low-educated employees, defined as the share of employees holding at maximum a lower-secondary education degree according to ISCED 2011; *vii*) the share of young/old employees, defined as the share of employees aged under 25 years-old / over 50 years-old, respectively. Then, concerning the labour cost variables, we consider: *viii*) the labour share computed as the firm's total labour cost (the sum of workers' wages and salaries plus social contributions) over the firm's value added; *ix*) the log of the average value of wages and salaries computed as the sum of workers' wages and salaries over firm total employment, and the *x*) the within-firm wage dispersion, computed as the Top10/Bottom10 ratio of wages and salaries at the firm level. Finally, concerning firm performance, we consider in turn: *xi*) the log of sales; *xii*) the log of labour productivity computed as value added-to-employment ratio and *xiii*) the log of value-added.

Table 4 reports the mean values of the different outcome variables, distinguishing between firms with a spike ("adopters") and firms without a spike ("non-adopters") and pooling all the values along the relevant time period. Overall, adopting firms tend to be larger, both in terms of employment and in terms of production value, pay higher wages, have a higher share of managers and white-collar workers, and employ an older, more educated, and stable workforce. The next step of our empirical analysis addresses the extent to which these differences can be attributed to automation spikes.

3 The effects of automation on the employment dynamics

We investigate the evolution of employment and other firm-level outcomes around a spike in the import of automation-related goods with an event-study (see [Berson et al., 2020](#); [Domini et al., 2022](#); [Bessen et al., 2023](#), as recent examples of similar methodology in the context of the adoption of automation technologies). We apply the estimator proposed recently by [Callaway and Sant'Anna \(2021\)](#), which improves upon the standard dynamic TWFE specification by solving the problems associated with differential treatment timing

⁶ From the ASIA-ER, annual jobs are defined as full-time equivalent positions, based on the number of weeks worked in each year.

	No spike	Spike	T-test
Number of employees	18.91	107.81	***
Value added per employee (log)	10.91	11.15	***
Turnover per employee (log)	12.31	12.48	***
Wages and salaries per-employee (log)	9.99	10.30	***
Share of blue-collar employee (%)	56.36	51.61	***
Share of white-collar employee (%)	41.30	43.23	***
Share of managers	1.84	4.69	***
Share of 15/29 years-old employees (%)	18.21	15.33	***
Share of 50 years-old and older employees (%)	23.75	24.18	***
Share of permanent employees (%)	88.69	92.58	***
Share of fixed-term employees (%)	11.31	7.42	***
Share of high-educated workers (%)	12.70	15.24	***
Share of medium-educated workers (%)	48.40	51.081	***
Share of low-educated workers (%)	33.75	30.51	***
Share of full-time employees (%)	72.40	86.84	***
Share of part-timers (%)	27.60	13.16	***
Number of observations	973,069	228,585	
Number of firms	150,514	29,419	

Note: *** significant difference at 1% level.

Table 4: Comparing firms with and without an automation spike, all years (2011-2019). Elaborations on International Trade, ASIA Employment Register and Frame-SBS data.

and heterogeneity of the treatment effect over time (see [Goodman-Bacon, 2021](#)). The key output of the model is the dynamic aggregation of group-time effects, which refer to the time-varying effects of the adoption for a given group, where a group is defined as a cohort of adopters with a spike in importing automation goods in the same year.

In detail, for all the variables, we aim to identify the average adoption effect for adopters (AAA), analogous to the average treatment-effect for the treated (ATT) ([Roth et al., 2023](#)), formally this reads as the quantity:

$$AAA^Y(a, t) = \mathbf{E}[Y_{i,t} - Y_{i,a-1} | A_i = a] - \mathbf{E}[Y_{i,t} - Y_{i,a-1} | A_i \in \mathcal{A}_{comp}] \quad (1)$$

where $AAA^Y(a, t)$ is the average adoption effect on variable Y in year t for an adopter adopting in year a . $Y_{i,t}$ is the dependent variable of interest at time t for firm i , A_i is an indicator variable with value in the adoption year for firm i and \mathcal{A}_{comp} is the most general set of comparisons for an adopter at time a , i.e. all firms that will never adopt or have not yet adopted at time t , in formal terms $\mathcal{A}_{comp} = \{a' | a' > t\}$. The equation generalises the standard difference-in-difference statistics to the context of staggered treatments. The effects are estimated through a two-step procedure obtaining the Doubly-Robust DiD estimators $\widehat{AAA^Y}(a, t)$ for each pair of years a, t in the considered time window.⁷ As a last step, in order to recover an event-study-like interpretation, we compute the so-called ‘dynamic aggregation’

⁷ Details on the characteristics of the DR estimator and the comparison with possible different approaches can be found in Section 4 of [Callaway and Sant’Anna \(2021\)](#).

obtained via a weighted average of the treatment effect s periods after the adoption across different adopting cohorts,

$$\widehat{AAA^Y}_s = \sum_l w_l \widehat{AAA^Y}(l + s, l) \quad (2)$$

where the weights w_l are chosen to account for the relative frequency of each cohort into the adopters' population. We perform our main estimation and plot the elements in (2) vis à vis relative time indicators, with 95% bootstrap confidence intervals (see [Callaway and Sant'Anna, 2021](#)).

This methodology exploits both variations in adoption timing among adopting firms and cross-sectional variation in the adopter status, i.e. comparing adopters vs. non-adopters.⁸ Causal interpretation of our results is granted under the two classic main assumptions of 'parallel trends' and 'no anticipation' (see [Callaway and Sant'Anna, 2021](#); [Borusyak et al., 2022](#)). First, it is necessary that adopters and control units had followed a parallel trend in the absence of adoption. This is an unobservable counterfactual, although insights can be gleaned from trends prior to adoption. For our application, parallel pre-adoption trends are observed conditionally on a set of firm-level covariates when we compare adopters with non-adopters and not-yet-adopters and all the more so if we work with the difference in timing only, using only not-yet-adopters as a control group.⁹ Second, it is required that the outcomes dynamics of the firms do not anticipate the event, or more formally, that adoption at time t does not have a causal impact on the observed variables at any $s < t$. As observed by [Bessen et al. \(2023\)](#), this hypothesis cannot be verified and is less likely to hold, as automation might be part of firms' investment strategies and thus not exogenous to firm decisions. In this respect, we devote a series of robustness checks to detect potential anticipating behaviours by changing the horizons of the estimations of pre-spike and post-spike variables.¹⁰ Moreover, with respect to causal identification, notice that within the proposed framework we can exclude the existence of contemporaneous relevant shocks affecting both the choice to automate and the target variables by testing the dynamic effects of adoption on several firm-level variables. As an example, consider contemporaneous demand shocks driving both employment dynamics and the import of automation goods, to which we refer in the discussion below (this hypothesis is discussed thoroughly in [Bonfiglioli et al., 2020](#)).

⁸ In formal terms, in order to use only one or both of the sources of variation, one has to change the composition of the set \mathcal{A}_{comp} including non-adopters or firms that did not adopt yet (not-yet-adopters), or both.

⁹ The benchmark exercises are obtained conditioning on dummies for the level 1 NACE sections, exporter status, size class, and firm-level growth rates. The log of the initial productivity (first year in the dataset) is used for all outcome variables except productivity itself, where it is replaced by the log sales.

¹⁰ Details on the methodological precautions and the results can be found in section 4.2.

3.1 Results

We discuss the results of the estimations of items defined by Eq. (2) on two different samples: the first is the baseline sample of importers, presented in Section 2; the second is obtained including only firms importing automation technologies, i.e. the adopters. As for the latter, note that the presence of ‘clean controls’ at time t for AAA estimates is guaranteed by those firms that adopt at any time s greater than t . In both cases, we classify as treated those firms that have a spike in automation-related goods within the span of 2011-2019. However, we exclude firms that re-exported a comparable amount in the same product category after the spike (labelled in the following *minimal re-exporter exclusion*). To determine the comparable amount of exports, we consider a cumulative quantity (over a four-year horizon) within the range of 90% to three times the value imported during the initial spike. Suppose, for example, that the import value of automation goods is $q_{\text{imp}}^t = 100$ (in the adoption year t), then a firm is excluded if the sum of the exported values in the following four years $\sum_{s=t}^{t+4} q_{\text{exp}}^s$ is between 90 and 300.¹¹

The statistics presented in Table 4 highlight significant differences between adopters and non-adopters along several dimensions. Thus, comparing the estimates obtained from the two samples serves as a twofold robustness check: (i) on the one hand, firms in the adopters-only group compose a more homogeneous control group. Although we condition using firm covariates for both samples, the restriction of the smaller sample reduces further the likelihood of a significant violation of the parallel trend assumption; (ii) on the other side, the larger sample, including both adopters and non-adopters, ensures the existence of a control group with a sufficient number of observations at any point in time. This, in turn, limits smaller-sample distortions, which become particularly relevant at the end of the time span.¹²

The results are reported in Figures 4 to 8, where the dynamically aggregated average adoption effects are plotted with s ranging from -4 to 5. The error bars represent the bootstrapped confidence intervals at a 5% significance level. Dependent variables are measured either in logs or without any specific transformation when we work with shares capturing the firm occupational structure. This allows for a direct interpretation of the outcomes in terms of semi-elasticities for the logs and in terms of changes in percentage points for the shares.

¹¹ To clean for hidden intermediation behaviours we ran a series of robustness checks using more restrictive filtering criteria. These are described in detail in Section 4.1.

¹² In the adopters-only sample, the control group consists of adopters that have not adopted yet; hence, by definition, its population reduces as we move forward in time (shrinking to zero for the last year in the window).

Firms’ employment and sales. In Figure 4, we investigate the effects of automation on the number of employees and on the sales of adopters. For employment (Fig. 4a), there is a clear increasing trend in the years following a spike: five years after the adoption event, firms are, on average, 10% larger in terms of employment. Such figures are similar for both samples, even if the confidence intervals are larger in the adopters only sample. As for sales (Fig. 4b), a slight decrease is observed in the year immediately after adoption, but it subsequently recovers, settling at around 10% five years after the initial spike.

These results confirm the positive relationship between automation and the average employment at the firm level observed in other countries (Acemoglu et al., 2020; Dixon et al., 2021; Domini et al., 2021; Acemoglu et al., 2023). Furthermore, the sales dynamic after the adoption event leads us to exclude potential contemporaneous demand shocks driving the employment results (Bonfiglioli et al., 2020). If such a mechanism were at play, we expect it to equally affect employment and sales growth during the adoption period and the immediate aftermath. However, our analysis reveals a significant but minimal decline in sales during the two years following the spike, followed by a gradual recovery to levels below those observed for employment.




Figure 4: The effect of automation adoption on firm-level employment, measured as the log of total full-time equivalent jobs in a year, and sales, again in logs. A comparison between the effect when measured in the full sample of importers and in the sample including adopters only (in both cases, with the minimal reexporters exclusion). Note: scatter plot points represent the coefficients of the dynamically aggregated average treatment effect at a given time step, while the error bars delimit the confidence interval at the 5% significance level, with bootstrapped critical values.

Other firms’ variables. In Figure 5, we expand the set of firm covariates potentially affected by automation, looking at value-added, labour productivity (value added over full-time equivalent number of employees) and labour share (total wages over value added).

The value added dynamic closely resembles the employment curve, while productivity is more in line with the trend of sales. Regarding the impact estimate, the value added increases

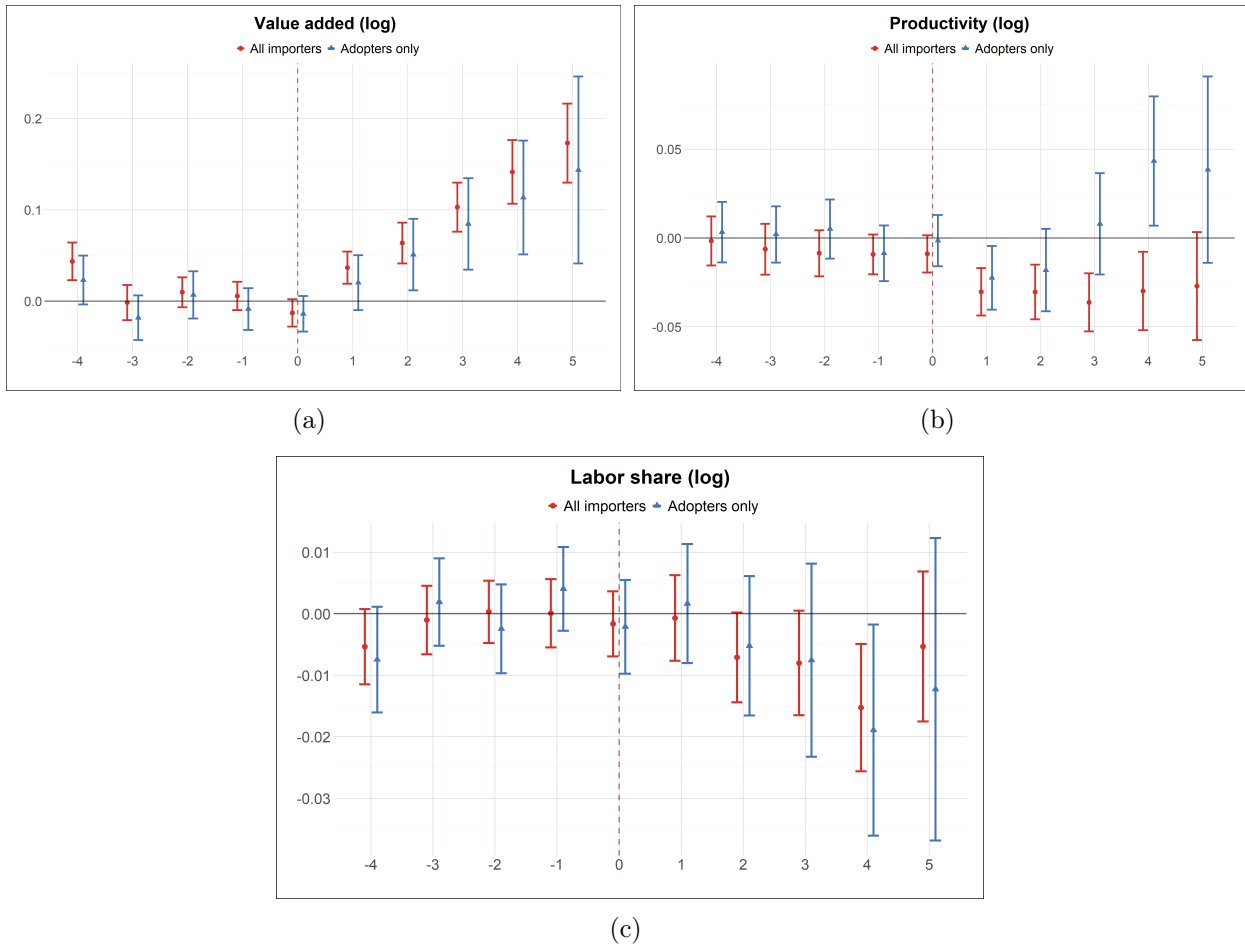


Figure 5: The effect of automation adoption on value added (in logs), labour productivity, measured as the log of the ratio between value added and the total time of full-time equivalent jobs in a year, and labour share (in logs), measured as the ratio between wage bill and value-added. See Figure 4 for the interpretation of the figures' items.

by 15% at five years after the initial spike, with the productivity settling on a decreasing path, though signals of recovery at five years (for a net around 4%) are shown by the estimates obtained from the adopters-only sample. These results suggest that firms' labour force expansion is not synchronized with a proportional growth in value added. In particular, the delay in productivity recovery is consistent with the empirical literature investigating the relation between investment spikes and productivity (Power, 1998; Huggett and Ospina, 2001; Grazzi et al., 2016; Domini et al., 2022) and would be explained by an initial negative effect due to the adoption of the new technology, which requires some adjustment and adaptation of internal production processes through learning-by-doing. Finally, the results on labour share show that the automation event does not necessarily make the production process less labour-intensive. This suggests that the compensation mechanism via new tasks or products may play a role in the average firm-level expansion of employment.

Wages and occupational structure In the following, we first look at the effects of automation on a set of wage-related variables, capturing both the average and the within-firm dispersion of workers’ compensation; then we investigate the effects on the firm’s occupational structure as determined by the share of workers in each occupational category and for different contract types (Figures from 6 to 8). The average wage shows a significant increase after the adoption event, setting at 6% at five years, after a slight decline in the year of the spike. However, we do not find evidence of significant distributional changes within adopters in the direction of increasing wage inequality. Figure-6b shows a non-significant increase of wage polarization indeed - after an initial upswing of the inequality indicator for the whole sample.

The estimates for the average wage are in the ballpark of those obtained by [Barth et al. \(2020\)](#) for Norway and [Humlum \(2020\)](#) for Denmark: they report an increase of, respectively, 4% and 8% in the average wage bill within firms adopting robots, while [Domini et al. \(2022\)](#) find a 1% effect on the average wage for French firms importing automation technologies. Other estimates find negative or non-significant effects,¹³ suggesting that the effect on wages may also depend on the country-level institutional setting. Regarding estimates on wage dispersion, our results are more aligned with [Domini et al. \(2022\)](#), who find that the inequality is substantially unaffected by the adoption of automation technologies in France, than with [Barth et al. \(2020\)](#), who find a positive effect of robot adoption on wage inequality in Norway.

Regarding the occupational structure, we do not find evidence of major changes. Figure 7 shows a slight but not significant increase in the share of managers, compensated for by an opposite trend for the share of blue collars and with white collars largely unaffected. Similarly, small effects are observed for other groups of workers. There is, for example, a small increase in the share of older (aged more than fifty) and less educated workers compared to the other groups. Finally, we also see that the average positive employment growth in firms that automate occurs with the creation of ‘good jobs’, namely regulated by full-time and permanent contracts: the measured effect goes hand in hand with an increase in the share of full-time jobs, around 3%, and of permanent contracts, between 1 and 2% (Figure 8).

3.2 Heterogeneity across size classes and aggregate effects

In this section, we shift our attention from examining the average effect of the adoption process on individual adopters to exploring its broader implications on aggregate employment. To achieve this, we first analyze the impact of adoption on the group of adopters, taking

¹³ [Aghion et al. \(2020\)](#) and [Koch et al. \(2021\)](#) find non-significant effects of robots and automation technologies on wages among Spanish and French firms, whereas [Bessen et al. \(2023\)](#) for the Netherlands find that automating firms have faster employment growth but not faster wage growth than non-automating firms.




Figure 6: The effect of automation adoption on average salary and wage dispersion measured the ratio between the top 10% and the bottom 10% of the distribution. See Figure 4 for the interpretation of the figures' items.





Figure 7: The effect of automation adoption on the firm-level shares for different types of occupations. See Figure 4 for the interpretation of the figures' items.

into account potential heterogeneous effects across different size classes.¹⁴ Then, we analyze

¹⁴Notice that the results obtained on subsamples are not strictly comparable, in terms of measured magnitude, with those obtained using the full sample. A comparison is, instead, always possible when using




Figure 8: The effect of automation adoption on the share of permanent and full-time contracts. See the caption of Figure 4 for a note on the plot interpretation.

the impact on the entire population of importers at the industry level. We find that the adoption of automation technologies has overall negative aggregate effects at the sectoral level, a finding that can be reconciled with the firm-level results in light of the differential effects across different size classes.

3.2.1 Size classes and aggregate effects on adopters

The analysis in Section 3.1 has focused on average effects across firms. However, such average results may not be very informative about aggregate and economy-wide effects if automation does have a heterogeneous impact across firms characterized by different initial levels of employment. This is the hypothesis we test here. We divide firms into three size categories, similarly to Table 3. More specifically, based on the firms' initial (i.e., the first year in the dataset) number of employees, we assigned them to three different size classes, whose employment intervals are thus defined: 1-19, 20-49, 50+. The estimation performed on this partition has two additional advantages in comparison to the one obtained on the full sample. First, we restrict the comparison to firms that are as similar as possible in terms of their propensity to adopt automation technologies, since larger firms are characterised by a higher probability of adoption (see Table 3). Second, we possibly isolate within the first size class (1-19) some effects of the labour market regulation, in particular firing costs, which have a discontinuous change at the size threshold of 15 employees, and which underwent a major change in 2015.¹⁵

the same sample but different estimators as for the case of the comparison between non-weighted and size-weighted estimates of the average adoption effect, which we perform below.

¹⁵ As of March 2015 a major labour market reform took place in Italy (Law 183/2014, known as 'Jobs Act') with the main aim of reducing firing costs for companies with at least 15 employees, besides favouring the use of temporary contracts and the use of hourly-ticket as compensation for work (vouchers). See e.g. Cirillo et al. (2017) and Sestito and Viviano (2018) for a thorough presentation.

Figure 9a displays clear heterogeneous effects of automation spikes across the different size classes. Firms with less than 20 employees appear to be the ones driving the positive average results observed in Figure 4: they increase their employment after adoption and are around 4% larger five years after the spike; firms in the middle size class, between 20 and 49 employees, are not significantly affected by the treatment after five years, and for medium and large firms the effect becomes negative with a value at five years around -4%. These results suggest that among the theoretical mechanisms driving the employment dynamics after the automation event, the displacement effect seems to play a major role for medium and large firms, whereas the scale and productivity effects have a greater impact for micro and small firms.


Given the significant heterogeneity of the effects of automation adoption across different size classes, a natural question arises of whether the positive effect on small firms balances the negative impact on the medium and large firms, or instead, there is a net negative effect in the aggregate. We answer this question by repeating the main exercise by weighting the observations by their initial size. The results reported in Figure 9b point to a negative aggregate effect on employment in the population of adopters settling around -4%. These numbers should be interpreted as a reduction of four percentage points of the total number of employees at the adopting firms in the five years following the adoption.

3.2.2 Sectoral effects of automation


This section investigates the impact of automation on employment by shifting the focus from individual firms to the industry level. In particular, we examine the relationship between automation adoption rates and employment across and within sectors, adding the non-adopters contribution to the aggregate effects measured with the exercise of Figure 9b. Using the spikes specified for the firm-level event study, we define a measure that captures the penetration of automation technologies within 5-digit sectors, SpikeVal_{sy-ey} , as the value of automation goods imported through the spikes in the span between the selected start year sy and the final year ey relative to the total import of the 5-digit sector in the same period.¹⁶

As first descriptive evidence, we look at the distribution of the employment growth rates at the level of industries (2-digits NACE aggregation). In Figure 10, we order the sectors by increasing levels of the relative intensity of automation imports, measured by the employment weighted average of SpikeVal of the 5-digit sectors within each industry in the period 2011-2019, reporting for each sector the aggregate employment growth, including both adopters

¹⁶Notice that the import of automation through spikes excludes only a small fraction of all imports of automation technologies, and its value more likely captures investment events. Nevertheless, we complement the exercise with robustness checks based on different measures of the exposure of sectors to automation: an extensive measure of automation penetration, ShSpike_{sy-ey} , i.e., the share of adopters in the importer population; an intensive measure including the value of all the transactions in automation-related goods relative to the total import in the sector, $\text{ShAutoImport}_{sy-ey}$.



(a)



(b)

Figure 9: The effect of automation adoption on employment at the firm-level. Split samples per size classes (1-19, 20-49, and more than 50 employees) and estimates obtained by weighting the observations by their initial size. See Figure 4 for the interpretation of the figures' items.

and non-adopters. Except for the fact that the 30% of more exposed industries tend to display net positive growth, we do not find any sharp evidence that industries more exposed to the import of automation goods have a significant loss or gain in terms of employment.

This first evidence must be taken with some caution. First, the share of adopting firms might change over time across the different sectors, but any distributional asymmetry in the time span is smoothed out and loses its power for identification when data are pooled together. Second, 2-digits industries may include very heterogeneous economic activities. In light of these considerations, we complement the plots above with a more granular econometric analysis, running a series of linear models with different sets of controls for 5-digit sectors. We first split the time period into two windows; the first, which covers the whole time span (2011-2019), is used to measure the exposure to the adoption of automation in the sector, and the second, restricted to the second half (2015-2019), to assess the variation




Figure 10: Growth rates of employment at the sector level, cumulated from 2011 to 2019. Sectors ordered by increasing intensity of automation imports, measured as the number of spikes in the span 2011-2019 normalized by the firms in the sector.

in the cumulated employment growth of 5-digit sectors.¹⁷ Table 5 shows the results of the estimation of a linear model presented in three specifications: i) without industry fixed effects; ii) with 2-digits industry fixed effects only, and iii) controlling for the variation in sales and the 2-digits industry fixed effects.¹⁸ Results from column 1 align with the main message from Figure 10: there is not a general, significant relationship between sectoral exposure to automation and employment growth. However, within 2-digits sectors (col. 2), there is significant evidence that 5-digits sectors more exposed to automation grow less in employment, suggesting reallocation dynamics. This relationship holds when we additionally control for growth in sales (col. 3).

Overall, this exercise integrates the insights of Figure 9b: there is an aggregate negative effect of automation on employment, both on the population of adopters, and on the industry-level general population of importers (including both adopters and non-adopters).

¹⁷ This choice aims to capture reallocation mechanisms that are likely to be non-contemporaneous and affect sectoral employment only with some lag. This is confirmed by the robustness checks performed varying the time windows (see Table A7).

¹⁸ The measure of automation exposure is standardized to favour the comparison with the robustness checks based on different measures of automation penetration: Table A8 shows variations of the benchmark exercise obtained using an extensive measure of automation (the number of spikes over the number of firms in the sector) and a different intensive measure from all imports of automation goods, not including only goods imported through a spike. All estimations are presented with robust standard errors clustered at NC2 level.

	ΔEmp_{15-19}		
$\Delta\text{Sales}_{15-19}$			0.73*** (0.06)
$\text{SpikeVal}_{11-19}(\text{std.})$	-0.007 (0.005)	-0.012** (0.005)	-0.015** (0.006)
Num.Obs.	750	750	750
R2 Adj.	-0.001	0.058	0.573
R2 Within Adj.		0.001	0.547
Std.Err. by:	NC2	NC2	NC2
FE: NC2		X	X

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: The effect of automation penetration on the cumulated employment growth (2015-2019), measured as log variations for 5-digit sectors. Specifications with 2-digit industry fixed effects and variation in sales are tested. See Tables A7-A8 for robustness checks.

4 Robustness checks

We have highlighted so far a few possible biases affecting the research design through the identification of the adopters and the related automation events. We comment here on a set of robustness checks, whose results are shown in the relative appendices.

4.1 Controlling for hidden intermediation

Our empirical design assumes that imported goods are used and adopted for production by firms. However, it may happen that some firms import automation products with the purpose of selling them domestically or in international markets, engaging in so-called Carry Along Trade (CAT) (Bernard et al., 2019). The direction of the bias depends on the underlying mechanism at work: if the productivity or scale effect is predominant over the displacement effect, i.e. the net impact of automation is positive, the incorrect labelling of the intermediation activities as automation spikes would bias the estimates downward and vice versa if the net effect is negative. With the aim of reducing this bias to the minimum, we leverage the variables available in our dataset to rule out possible potential resellers of automation goods in the international or domestic markets. For the first case, we extend the reexporting exclusion of the main analysis in a more conservative way by removing all the firms that reexported any product in the automation goods categories in any year after the spike. We compare this broader reexport exclusion with the 'minimal' one and with the non-restricted sample. To tackle also the possibility of reselling goods in the domestic market, we exclude: i) persistent importers of automation products (those firms importing

every year), as this might suggest that they are importing in order to resell (domestically or in export markets); ii) oversized importers of automation goods, i.e. those firms that are among the outliers of the distribution of the value of import in automation per employee in any year.¹⁹ Figure A1, shows substantial agreement for the effects measured on the different samples.

4.2 Controlling for adopters' anticipation

As mentioned in Section 3, the causal interpretation of our results is possible under two main assumptions: i) conditional parallel trends between adopters and non-adopters and ii) no anticipation by the adopters. However, since investment decisions might be determined in advance by firms, the second of the two fundamental assumptions might not hold as changes in the outcome variable might occur before the adoption takes place. Typical examples include firms restructuring the labour force, with possible effects on employment and occupation shares, a few years before the automation adoption takes place. Callaway and Sant'Anna (2021) define a generalized version of Eq. (1) allowing the units to anticipate participating in the treatment and adjust their behaviour (or experiencing a change in the observed results) before actually participating in the treatment (adoption in our framework). In essence, the two groups are assumed to follow parallel trends (in the absence of the adoption) except in the period immediately before the adoption, where the horizon defining the anticipation timing is predefined. Under this condition, the formula for the Average Adoption Effect on Adopters reads

$$AAA_{\delta}^Y(a, t) = \mathbf{E} [Y_{i,t} - Y_{i,a-\delta-1} | A_i = a] - \mathbf{E} [Y_{i,t} - Y_{i,a-\delta-1} | A_i \in \mathcal{A}_{comp}^{\delta}]$$

where a parameter δ has been added to model the anticipation horizon. In other words, the difference between trends shaping the measured effect on the outcome variables is measured using pre-adoption and pre-anticipation periods rather than pre-adoption periods only.²⁰ While modelling anticipation allows for a certain degree of flexibility, it comes with a fundamental tradeoff: as we increase the time horizon of the anticipation, we need parallel trends assumption to hold over more periods in advance in order to identify the effects, thus reducing the population of allowed treated units. For this reason, we estimated the main results under the no anticipation assumption, and here we propose a series of robustness checks letting $\delta = 1, 2, 3, 4$, covering anticipation δ from one to four years before the adoption. Figure A2 does not highlight the existence of anticipation effects in the years before the

¹⁹ Excessive mass in the tail of the distribution is identified after the threshold associated with the 993th quantile. We opted for a conservative exclusion, cutting out the last percentile.

²⁰ Notice that the modified equation includes the Eq. 1 as a special case, where $\delta = 0$ and that also the comparison groups change as we modify the anticipation horizon since not-yet-adopters cannot be included in the control group if they are anticipating the treatment.

automation event, with the aftermath’s trends being unaffected.

4.3 Placebo test

We propose two placebo tests, addressing the validity of the identification through variation in adoption timing and through variation in the adoption status. For the former, we randomise adoption timing across adopters in the dataset; for the second, we set to zero the spike for all the adopters and then select firms at random from the population of non-adopters to assign them a fake spike. We proceed then by estimating the identical difference-in-differences model defined by equation (1) and then obtain the dynamic aggregation as defined by (2). Results confirm that the methodology is robust to both placebo tests as no effect is found in the event study plots in Figure A3.

4.4 Dissecting automation spikes.

In our benchmark exercises, we defined spikes based on imports of capital goods related to different automation technologies, including general automation technologies (as automatic machine tools), AI technologies, such as automatic data processing machines, industrial robots or any combination of the three simultaneously.

However, as pointed out in Domini et al. (2021), it is important to consider that these three groups of technologies may have different impacts on employment and other firm-level variables, as they serve different purposes and affect different types of workers (see also Webb, 2019). Moreover, considering separately three spike definitions offers the possibility to compare our results with those of works that focus exclusively on one single technology. Figure A4 shows that trends are similar for the three classes of automation goods. Automating through AI and robots has a larger positive effect with respect to general automating technologies, although estimates for robots are barely significant given the limited diffusion in our sample (see Figure 1).

4.5 Synchronous investments in tangible assets

For the identification of the effect of the adoption of automation technologies, it is important to exclude events correlated in the time dimension that might have an impact on the outcome variables. One of the possible sources of bias is the existence of synchronous investments in capital goods that do not necessarily include automating technologies. This matter is thoroughly discussed in Domini et al. (2021), where they analyze the pattern of synchronization between spikes in capital investment from the balance sheet and spikes of automation from trade sources. We exploit the information available in our sources to show how our measure is distinct from a general measure of the import of machinery. We compare the spikes in

Year	Share of general machinery spikes which are also automation spikes (%)		Share of automation spikes which are also general machinery spikes (%)	
	Sync	Sync and lead/lag	Sync	Sync and lead/lag
2012	10.13	15.39	14.55	22.1
2013	8.7	12.65	12.08	17.57
2014	8.19	12.31	11.74	17.64
2015	9.16	14.33	13.07	20.44
2016	9.74	15.44	14.07	22.31
2017	12.65	19.19	19.37	29.39
2018	9.7	13.68	13.93	19.65

Table 6: A comparison between the patterns of spikes in automation technologies as defined in Table A5 and spikes of import of product codes in the HS2 category 84 including “Nuclear reactors, boilers, machinery and mechanical appliances; part thereof”. Elaborations on International Trade data.

the import of automation goods as defined in the main text and spikes of import in general machinery, namely import transactions registered in the 2-digit HS product category 84, excluding HS6 codes already included in the automation measure and those labelling parts or auxiliary products.²¹ Table 6 shows that the import of automation technology is not happening jointly with the import of general machinery: on average only 16.09% of automation spikes are also general machinery spikes, while only the 10.95% of general machinery spikes are also automation spikes. These percentages increase to 19.87% and 15.93% if we include spikes one year before and one year after. We use the results of this analysis to perform the main exercise on a restricted sample with the aim of identifying the automation event, excluding any source of correlation with the import of other machinery goods. Specifically, we remove adopters from the sample that have a spike in product goods in the category “other machinery” in the same year or in the years before and after. This implies a reduction of the sample of the adopters in between one-fourth and one-fifth. Figure A5 shows the results of the estimation in the event study form: the direction of the impact is confirmed though the reduction of the sample consistently affects the significance of the exercise carried out using only adopters.

4.6 Sectoral patterns

We conclude this robustness section with the results of another split sample exercise defined according to the macro sectors of main activity for each firm. We first assign firms to manufacturing vs. non-manufacturing sectors based on the reported industry of main activity of the firm. Figure A6a shows that the positive effect of automation spikes on firm-level

²¹ The HS2 product code 84 includes “Nuclear reactors, boilers, machinery and mechanical appliances; part thereof”, from which we excluded HS6 product codes included in Table A5 and codes 8409, 8420, 8423, 8425, 8431, 8448, 8466, 8470, 8471, 8472, 8473, 8476, 8481, 8482, 8483, 8484, 8485, 8487.

employment is mostly apparent in non-manufacturing sectors. In manufacturing sectors, it is still positive but close to zero (and not significant) in the first years after adoption, and then becomes positive and weakly significant at the end of the time window, settling at values below 5%. In order to better capture industry-specific technological characteristics, we further classify firms using the revised Pavitt taxonomy (Bogliacino and Pianta, 2010): Figure A6b shows that the increase in employment five years after the spike is not confirmed for scale and information-intensive sectors.

5 Conclusions

We estimate the impacts of adopting automation technologies on firms' outcomes with a particular focus on employment, using firm-level data on the import of automation and AI-related capital goods from the universe of Italian importing firms from 2011 to 2019. Leveraging the most recent advances in the econometrics of difference-in-differences estimators, we set up an event-study-like design to track the evolution of a set of firm-level attributes in the years following the spike in the import of automation goods. We show that, on average, adopters of automation technologies experience consistent employment growth of up to the 10% after five years. Moreover, we find evidence of a transitional drop in sales and productivity and specific hiring patterns favouring full-time and permanent contracts.

We then investigate potential differences in the effects of automation according to firm size cohorts. Our findings show that the positive effect of automation adoption on employment is positive and significant for the sample of small firms, loses statistical significance for the middle category, and becomes negative and positive for larger companies. These results suggest that the scale vs. displacement effects of automation may be heterogeneous across firms. Finally, we assess the aggregate effects of automation with an employment-weighted regression, showing that the overall prevailing effect is negative. Such evidence challenges previous findings at the firm level that reported a positive effect of automation on adopters (Acemoglu et al., 2020; Dixon et al., 2021; Koch et al., 2021).

Finally, we complement the firm-level evidence by studying the sectoral effects of automation, thus considering potential spillover effects for non-adopting firms. We find evidence that 5-digits sectors more exposed to automation grow less in terms of employment, suggesting additional reallocation effects across industries as a consequence of automation adoption.

In conclusion, our findings offer a novel perspective on the effects of automation, which should be considered before any attempt to implement active industrial and labour policies. Previous works have highlighted the differences between adopters and non-adopters, contributing to a substantial reallocation within the industry. In addition, and new to the picture, our results show that even in the relatively more homogeneous group of adopters,

profound differences exist in the effects produced by automation, so additional firm characteristic matters to infer the aggregate effect of adoption.

Acknowledgements

This paper has benefited from the comments of participants at several seminars and conferences: 7th Linked Employer-Employee Data (LEED) Workshop 2022, Porto (June 2022); Workshop on the Impact of Robotisation on Labour Markets and Inequality, University of Bologna (February 2023); ISTAT workshop "Multipurpose data-analysis to detect firm-level capabilities and their reactions to shocks", Rome (April 2023); Eu-SPRI Annual Conference, Brighton, UK (June 2023); CESifo Venice Summer Institute (June 2023). We are also indebted to Tania Treibich for insightful comments. Marco Grazzi gratefully acknowledges financial support from the Italian Ministry of University and Research (MUR) 2022XJHRCJ - REWIND - Resilient Enterprises and Workers: leveraging INtangibles to address Disruptions, CUP: J53D23005020008. Daniele Moschella gratefully acknowledges financial support from the Italian Ministry of Education and Research (MUR) under the PRIN2017 Programme: 201799ZJSN - Technological change, industry evolution and employment dynamics and under the PRIN2022 Programme: 2022ACJ2WS - Automation, Trade, and Global Value Chains: New Empirical Evidence and Theory Development. Declarations of interest: none.

References

- ABELIANSKY, A. L., I. MARTÍNEZ-ZARZOSO, AND K. PRETTNER (2020): “3D printing, international trade, and FDI,” *Economic Modelling*, 85, 288–306.
- ACEMOGLU, D., G. W. ANDERSON, D. N. BEEDE, C. BUFFINGTON, E. E. CHILDRESS, E. DINLERSOZ, L. S. FOSTER, N. GOLDSCHLAG, J. C. HALTIWANGER, Z. KROFF, ET AL. (2022): “Automation and the workforce: A firm-level view from the 2019 Annual Business Survey,” Tech. rep., National Bureau of Economic Research.
- ACEMOGLU, D., H. R. KOSTER, AND C. OZGEN (2023): “Robots and workers: Evidence from the Netherlands,” Tech. rep., National Bureau of Economic Research.
- ACEMOGLU, D., C. LELARGE, AND P. RESTREPO (2020): “Competing with Robots: Firm-Level Evidence from France,” *AEA Papers and Proceedings*, 110, 383–88.
- ACEMOGLU, D. AND P. RESTREPO (2018): “Demographics and Automation,” Working Paper 24421, National Bureau of Economic Research.
- (2019): “Artificial Intelligence, Automation, and Work,” in *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 197–236.
- (2020): “Robots and jobs: Evidence from US labor markets,” *Journal of Political Economy*, 128, 2188–2244.
- (2022): “Demographics and automation,” *The Review of Economic Studies*, 89, 1–44.
- AGHION, P., C. ANTONIN, S. BUNEL, AND X. JARAVEL (2020): “What are the labor and product market effects of automation? New evidence from France,” *CEPR Discussion Paper No. DP14443*.
- AHN, J., A. K. KHANDELWAL, AND S.-J. WEI (2011): “The role of intermediaries in facilitating trade,” *Journal of International Economics*, 84, 73–85.
- AKSOY, C. G., B. ÖZCAN, AND J. PHILIPP (2021): “Robots and the gender pay gap in Europe,” *European Economic Review*, 134, 103693.
- ASPHJELL, M., W. LETTERIE, P. G.A., AND O. NILSEN (2014): “Sequentiality versus simultaneity: Inter-related factor demand,” *The Review of Economics and Statistics*, 95, 986–998.
- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29, 3–30.
- BARTH, E., M. ROED, P. SCHONE, AND J. UMBLIJS (2020): “How Robots Change Within-Firm Wage Inequality,” IZA Discussion Papers 13605, Institute of Labor Economics (IZA).
- BENMELECH, E. AND M. ZATOR (2022): “Robots and firm investment,” NBER Working Papers 29676, National Bureau of Economic Research.
- BERNARD, A. B., E. J. BLANCHARD, I. VAN BEVEREN, AND H. VANDENBUSSCHE (2019): “Carry-along trade,” *The Review of Economic Studies*, 86, 526–563.

- BERNARD, A. B., M. GRAZZI, AND C. TOMASI (2015): “Intermediaries in International Trade: Products and Destinations,” *The Review of Economics and Statistics*, 97, 916–920.
- BERNARD, A. B., J. B. JENSEN, S. J. REDDING, AND P. K. SCHOTT (2010): “Wholesalers and Retailers in US Trade,” *American Economic Review*, 100, 408–13.
- BERSON, C., E. VIVIANO, AND M. DE PHILIPPIS (2020): “Job-to-job flows and wage dynamics in France and Italy,” Working Paper WP756, Banque de France.
- BESSEN, J., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2023): “What happens to workers at firms that automate?” *The Review of Economics and Statistics*, 1–45.
- BESSEN, J. E., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2020): “Firm-Level Automation: Evidence from the Netherlands,” *AEA Papers and Proceedings*, 110, 389–93.
- BLUM, B. S., S. CLARO, AND I. HORSTMANN (2010): “Facts and Figures on Intermediated Trade,” *American Economic Review*, 100, 419–23.
- BOGLIACINO, F. AND M. PIANTA (2010): “Innovation and employment: a reinvestigation using revised Pavitt classes,” *Research Policy*, 39, 799–809.
- BONFIGLIOLI, A., R. CRINO, H. FADINGER, AND G. GANCIA (2020): “Robot imports and firm-level outcomes,” *CEPR Discussion Paper No. DP14593*.
- BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2022): “Revisiting Event Study Designs: Robust and Efficient Estimation,” *Available at SSRN 2826228*.
- BRYNJOLFSSON, E. AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company.
- CALLAWAY, B. AND P. H. SANT’ANNA (2021): “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 225, 200–230.
- CIRILLO, V., M. FANA, AND D. GUARASCIO (2017): “Labour market reforms in Italy: evaluating the effects of the Jobs Act,” *Economia Politica*, 34, 211–232.
- DENG, L., V. PLÜMPE, AND J. STEGMAIER (2021): “Robot adoption at German plants,” IWH Discussion Papers 19/2020, Halle Institute for Economic Research (IWH).
- DIXON, J., B. HONG, AND L. WU (2021): “The robot revolution: Managerial and employment consequences for firms,” *Management Science*, 67, 5586–5605.
- DOMINI, G., M. GRAZZI, D. MOSCHELLA, AND T. TREIBICH (2021): “Threats and opportunities in the digital era: Automation spikes and employment dynamics,” *Research Policy*, 50, 104137.
- (2022): “For whom the bell tolls: the firm-level effects of automation on wage and gender inequality,” *Research Policy*, 51, 104533.
- DOSI, G., M. PIVA, M. E. VIRGILLITO, AND M. VIVARELLI (2021): “Embodied and disembodied technological change: The sectoral patterns of job-creation and job-destruction,” *Research Policy*, 50, 104199.

- DOTTORI, D. (2021): “Robots and employment: evidence from Italy,” *Economia Politica*, 38, 739–795.
- EATON, J. AND S. KORTUM (2001): “Trade in capital goods,” *European Economic Review*, 45, 1195–1235, international Seminar On Macroeconomics.
- FAIA, E., G. OTTAVIANO, AND S. SPINELLIA (2022): “Robot Adoption, Worker-Firm Sorting and Wage Inequality: Evidence from Administrative Panel Data,” *CEPR Discussion Papers No. DP17451*.
- FREEMAN, C., J. CLARK, AND L. SOETE (1982): *Unemployment and technical change: a study of long waves and economic development*, Greenwood Frances Pinter, London.
- FREY, C. B. AND M. A. OSBORNE (2017): “The future of employment: How susceptible are jobs to computerisation?” *Technological Forecasting and Social Change*, 114, 254 – 280.
- GOODMAN-BACON, A. (2021): “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 225, 254–277.
- GRAETZ, G. AND G. MICHAELS (2018): “Robots at work,” *Review of Economics and Statistics*, 100, 753–768.
- GRAZZI, M., N. JACOBY, AND T. TREIBICH (2016): “Dynamics of Investment and Firm Performance: Comparative evidence from manufacturing industries,” *Empirical Economics*, 51, 125–179.
- GRIGOLI, F., Z. KOCZAN, AND P. TOPALOVA (2020): “Automation and labor force participation in advanced economies: Macro and micro evidence,” *European Economic Review*, 126, 103443.
- HUGGETT, M. AND S. OSPINA (2001): “Does productivity growth fall after the adoption of new technology?” *Journal of Monetary Economics*, 48, 173–195.
- HUMLUM, A. (2020): *Essays on automation and labor markets*, Princeton University, chap. Robot adoption and labor market dynamics, 1–101.
- KOCH, M., I. MANUYLOV, AND M. SMOLKA (2021): “Robots and firms,” *The Economic Journal*, 131, 2553–2584.
- LETTERIE, W. A., G. A. PFANN, AND J. POLDER (2004): “Factor adjustment spikes and interrelation: an empirical investigation,” *Economics Letters*, 85, 145 – 150.
- MONTOBBO, F., J. STACCIOLI, M. E. VIRGILLITO, AND M. VIVARELLI (2022): “Robots and the origin of their labour-saving impact,” *Technological Forecasting and Social Change*, 174, 121122.
- PIANTA, M. (2005): “Innovation and employment,” in *The Oxford handbook of innovation*, ed. by J. Fagerberg, D. C. Mowery, and R. R. Nelson, Oxford university press.
- POWER, L. (1998): “The Missing Link: Technology, Investment, And Productivity,” *The Review of Economics and Statistics*, 80, 300–313.
- ROTH, J., P. H. SANT’ANNA, A. BILINSKI, AND J. POE (2023): “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature,” *Journal of Econometrics*.
- SESTITO, P. AND E. VIVIANO (2018): “Firing costs and firm hiring: evidence from an Italian reform,” *Economic Policy*, 33, 101–130.
- WEBB, M. (2019): “The Impact of Artificial Intelligence on the Labor Market,” *Available at SSRN 3482150*.

A Data preparation details

A.1 Structural break in trade reporting

According to Italian Law n.19 of 27 February 2017, as of January 2018, firms importing less than 200,000 euros per quarter (exporting less than 100,000 euros per quarter) are exempt from declaring the imported (exported) amounts for fiscal or statistical purposes. In the original COE data set, we observe not only a drop of the total number of importers in 2018 (Table A2, column 2), but also an upward shift of the whole distribution of total imports (Table A1, panel a) in the same period. To tackle the break in the series, we filtered the sample from 2018 backwards by setting a threshold for annual imports (400,000 euros) below which the firms have been excluded from the sample over the period 2011-2017. This threshold has been empirically determined in such a way as to restore a uniform distribution of total imports over years (Table A1, panel b). Consequently, the filtering process restores a number of importing firms per year without a break in 2018 (Table A2, column 3).²² Due to the filtering process, the sample reduced from 284,086 of distinct importers observed in 2011-2019 to 245,191 firms.

Year	P1	P10	P25	P50	P75	P90	P99
<i>a) Before filtering</i>							
2011	9	129	736	3,786	23,713	123,210	2,033,117
2012	8	116	653	3,487	21,750	112,982	1,859,460
2013	8	108	612	3,345	20,918	109,675	1,830,420
2014	7	103	579	3,220	20,458	108,283	1,803,916
2015	7	99	554	3,155	20,313	108,954	1,833,333
2016	7	94	515	2,986	19,164	103,353	1,727,226
2017	7	90	495	2,899	18,726	102,869	1,781,723
2018	9	166	1,125	5,060	31,383	169,186	2,770,928
2019	8	143	1,048	4,681	29,551	161,120	2,644,618
<i>b) After filtering</i>							
2011	1,042	1,553	3,349	13,097	77,023	580,629	6,160,242
2012	1,040	1,538	3,227	12,406	72,418	559,978	5,771,558
2013	1,038	1,510	3,154	12,006	71,650	569,741	5,831,403
2014	1,039	1,509	3,133	12,100	73,341	578,257	5,864,758
2015	1,038	1,518	3,187	12,533	78,806	599,561	6,009,504
2016	1,039	1,511	3,170	12,277	77,114	589,838	5,843,005
2017	1,038	1,504	3,116	12,135	78,259	613,612	6,184,089
2018	1,039	1,503	3,115	12,040	77,744	621,586	6,511,157
2019	1,037	1,487	3,040	11,809	78,183	629,384	6,610,208

Table A1: Relevant percentiles of the distribution of total imports (euros) by importer for each year in the window 2011-2019. Elaborations on International Trade data.

A.2 Classes of intermediaries

Here, we list the set of 5-digits sectors that are excluded in the analysis to remove potential noise from our key variable coming from intermediation activities. Table A3 reports the list of codes and Table A4 provides the measure of the impact of the cleaning procedure.

²² Dealing with annual data, we could not apply the threshold established by law, which is defined on a quarterly basis.

Year	# importers before filtering	# importers after filtering
2011	122,821	93,179
2012	119,608	88,935
2013	117,993	87,936
2014	118,164	88,331
2015	119,293	89,223
2016	120,428	89,790
2017	120,912	90,519
2018	89,665	89,050
2019	93,201	92,378

Table A2: Number of importers in each year before and after filtering (the same firm may be counted in more than one year). Elaborations on International Trade data.

Classes	Full description
46140	Agents involved in the sale of machinery, industrial equipment, ships and aircraft
46190	Agents involved in the sale of a variety of goods
46500	Wholesale of information and communication equipment
46620	Wholesale of machine tools
46640	Wholesale of machinery for the textile industry and of sewing and knitting machines
46699	Wholesale of other machinery and equipment
46900	Non-specialised wholesale trade

Table A3: 5-digit NACE rev.2 codes referring to classes of retailers/intermediaries of automation\AI goods

Year	Intermediaries share (%)	Import value share intermed. (%)	Export value share intermed. (%)
2011	4.79	4.79	1.37
2012	4.72	4.63	1.42
2013	4.64	4.46	1.29
2014	4.58	4.77	1.30
2015	4.56	5.10	1.24
2016	4.54	5.37	1.25
2017	4.50	5.35	1.21
2018	4.51	5.18	1.19
2019	4.73	5.27	1.24

Table A4: The intermediaries impact on import and export value by year. Elaborations on International Trade data.

Label	HS6 codes
<i>Automation</i>	
Industrial robot	847950
Dedicated machinery	847989
Automatic machine tools	845600-846699, 846820-846899, 851511-851519
Automatic welding machines	851521, 851531, 851580, 851590
Weaving and knitting machines	844600-844699, 844700-844799
Other textile dedicated machinery	844400-844590
Automatic conveyors	842831-842839
Automatic regulating instruments	903200-903299
3-D printers	847780
<i>Artificial Intelligence (AI)</i>	
Automatic data processing machines	847141-847150, 847321, 847330
Electronic calculating machines	847010-847029

Table A5: HS6 product codes referring to automation- and AI-related technologies.

A.3 Product categories

The product categories that we include in our main measure are chosen in line with previous works ([Acemoglu and Restrepo, 2018](#); [Domini et al., 2021, 2022](#)). Table A5 summarises the product codes, while Table A6 offers an overview of the importance of each technology in terms of total imports and employment.

B Event study robustness checks

In this section, we present some robustness checks with reference to the applied methodology for event studies. In particular, we propose various types of analysis to corroborate the validity of our spike identification measure and of our estimator used to measure effects on outcome variables. For the sake of synthesis, we present only the graphs referring to employment, but the others remain available upon request.

	Share of importers according to the imported technology (%)	Share of workers employed in importing firms according to the imported technology (%)
Industrial robot	0.5	3.0
Dedicated machinery	7.1	24.5
Automatic machine tools	10.9	29.2
Automatic welding machines	2.0	10.8
Weaving and knitting machines	0.4	0.6
Other textile dedicated machinery	0.4	0.7
Automatic conveyors	1.5	7.3
Automatic regulating instruments	5.1	20.7
3-D printers	1.0	4.3
Automatic data processing machines	5.8	32.4
Electronic calculating machines	0.4	3.8

Table A6: Share of imports and workers according to the imported technology in year 2019. Elaborations on International Trade and Frame-SBS data.

B.1 Controlling for hidden intermediation - Results

Figure A1 shows the results of the estimation in the event study form for the robustness check described in Section 4.1.




Figure A1: Robustness check for the methodological framework. The impact on employment is shown in the case where the population of adopters is restricted in order to exclude potential hidden intermediaries: over-sized and persistent importers of automation technologies, or firms that export automation technologies after a spike. See Figure 4 for the interpretation of the figures' items.

B.2 Controlling for adopters' anticipation - Results

Figure A2 shows the results of the estimation in the event study form for the robustness check described in Section 4.2.




Figure A2: Robustness check for the methodological framework: controlling for anticipating behaviours by the firm at different time horizons. See Figure 4 for the interpretation of the figures' items.

B.3 Placebo tests - Results

Figure A3 shows the results of the estimation in the event study form for the robustness check described in Section 4.3.

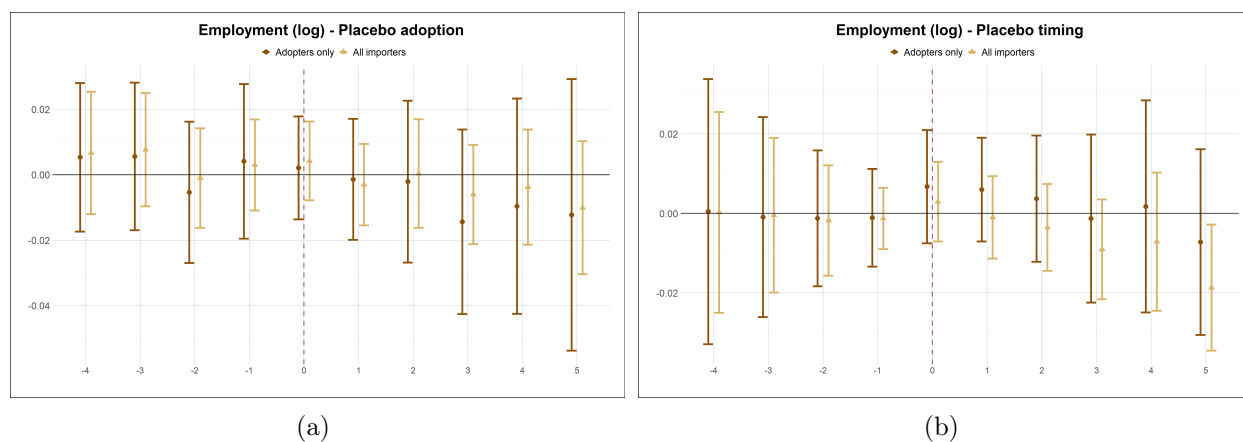


Figure A3: Robustness check for the methodological framework. The results for the employment are shown using i) a dataset where spikes are randomly assigned to firms in equal proportion to real spikes (panel a); ii) a dataset where the timing of the spike for real adopters is randomly assigned (panel b). See Figure 4 for the interpretation of the figures' items.

B.4 Other definitions of automation spike - Results

Figure A4 shows the results of the estimation in the event study form for the robustness check described in Section 4.4.




Figure A4: Robustness check for the methodological framework. The impact on employment is shown in the cases where the spikes are identified isolating products belonging to a particular technological category. See Figure 4 for the interpretation of the figures' items.

B.5 Synchronisation with general investments in machinery - Results

Figure A5 shows the results of the estimation in the event study form for the robustness check described in Section 4.5.






Figure A5: Robustness check for the methodological framework. The impact on employment is shown in the case where the population of adopters is restricted to the firms whose main import event of automation technologies is non-synchronous with import spike in other machinery. See Figure 4 for the interpretation of the figures' items.

B.6 Sectoral patterns

Figures A6 show the results of the estimation in the event study form for the robustness check described in Section 4.6.



(a)



(b)

Figure A6: The effect of automation adoption on employment at the firm-level. Split samples: manufacturing and non-manufacturing and NC2 sector Pavitt taxonomy. See Figure 4 for the interpretation of the figures' items.

C Sectoral effects robustness checks

Tables A7-A8 shows variations of the benchmark exercise of Section 3.2.2 obtained, respectively, changing the reference time windows for the measured employment growth, and modifying the automation penetration measure using instead, $ShSpike_{sy-ey}$, the share of firms having a spike in the import of automation goods in the span between the selected start-year sy and the final year ey , or $ShAutoImport_{sy-ey}$, the value share of automation import over the total import of the sector (i.e. including also goods imported non-necessarily through a spike).

	ΔEmp_{17-19}			ΔEmp_{13-19}			ΔEmp_{11-19}		
$\Delta\text{Sales}_{17-19}$	0.6*** (0.1)								
$\Delta\text{Sales}_{13-19}$				0.73*** (0.06)					
$\Delta\text{Sales}_{11-19}$							0.77*** (0.06)		
SpikeVal $_{11-19}(\text{std.})$	-0.009 (0.005)	-0.012*** (0.004)	-0.009* (0.005)	-0.004 (0.012)	-0.01 (0.01)	-0.007 (0.006)	-0.001 (0.012)	-0.01 (0.01)	-0.005 (0.006)
Num.Obs.	750	750	750	750	750	750	750	750	750
R2 Adj.	0.001	0.056	0.612	-0.001	0.116	0.687	-0.001	0.183	0.755
R2 Within Adj.		0.004	0.591		0.000	0.646		0.000	0.700
Std.Err. by:	NC2	NC2	NC2	NC2	NC2	NC2	NC2	NC2	NC2
NC2 FE		X	X		X	X		X	X

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A7: The effect of automation penetration on the cumulated employment growth. Varying the reference time windows (see benchmark exercise in Table 5).

	ΔEmp_{15-19}			ΔEmp_{11-19}			ΔEmp_{15-19}			ΔEmp_{11-19}		
$\Delta\text{Sales}_{15-19}$	0.73*** (0.06)						0.69*** (0.05)					
$\Delta\text{Sales}_{11-19}$				0.77*** (0.06)						0.72*** (0.06)		
ShSpike $_{11-19}(\text{std.})$							-0.01 (0.01)	-0.007 (0.008)	-0.013 (0.009)	-0.006 (0.012)	0.002 (0.011)	-0.006 (0.006)
ShAutoImp $_{11-19}(\text{std.})$	-0.006* (0.003)	-0.008* (0.004)	-0.010*** (0.003)	0.009 (0.008)	-0.003 (0.006)	-0.006 (0.004)						
Num.Obs.	750	750	750	750	750	750	781	781	781	781	781	781
R2 Adj.	0.001	0.059	0.575	0.001	0.182	0.755	0.000	0.088	0.628	-0.001	0.234	0.784
R2 Within Adj.		0.002	0.549		-0.001	0.700		-0.001	0.592		-0.001	0.718
Std.Err. by:	NC2	NC2	NC2	NC2	NC2	NC2	NC2	NC2	by: NC2	NC2	NC2	NC2
NC2 FE		X	X		X	X		X	X		X	X

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A8: The effect of automation penetration on the cumulated employment growth. Varying the explanatory variables measuring the automation adoption intensity at the sector level: an extensive measure of automation penetration, ShSpike $_{sy-ey}$, i.e., the share of adopters in the importer population; an intensive measure including the value of all the transactions in automation-related goods relative to the total import in the sector, ShAutoImport $_{sy-ey}$ (see benchmark exercise on Table 5).