

Automation, Firm size and Skill Groups*

Julian Tiedtke[†]

July 7, 2025

Abstract

This paper examines the impact of automation investments on employment dynamics and workforce composition using administrative data from Portugal. I exploit the lumpiness of automation imports in a difference-in-differences event study design. My results show that automation creates jobs in small firms but leads to job losses in larger ones. This pattern holds across a wide range of firm types, industries and types of automation technologies. Most importantly, automation favors low-educated, blue-collar workers in routine-intensive jobs over highly skilled workers like STEM professionals. These findings challenge the view of automation as inherently skill-biased.

Keywords: Automation, Employment, Firm heterogeneity, Deskilling.

JEL: J230, L250, O330.

*The author is grateful for helpful comments and suggestions by Michael Koch, Daisuke Adachi, Flavio Calvino, Andrea Mina, Daniele Moschella, Angelo Cuzzola and the participants at seminars and conferences at Global Labor Conference-AASLE-EALE-SOLE (Toronto, 2025), Druid Conference (Toronto, 2025), PSE-CEPR Policy Forum (Paris, 2025), Munich Summer Institute PhD Workshop (Munich, 2025), FIND seminar (Aarhus, 2025), Pisa Lunch seminar (Pisa, 2024), Italian Association of Labour Economics (AIEL, Naples, 2024), Verein für Socialpolitik (Berlin, 2024), Linked Employer-Employee Conference (LEED, Copenhagen, 2024), International Schumpeter Society Conference (Gothenburg, 2024), BeNA 20-Year Jubilee Conference (Berlin, 2024), Astril (Rome, 2024), YSI Pre-conference of Italian Economic Association (Aquila, 2023).

[†]Institute of Economics, Scuola Superiore Sant'Anna di Pisa, Italy. Email: j.tiedtke@santannapisa.it

1 Introduction

Does automation technology steal jobs? And if so, at whose expense? Recent advancements in robotics and artificial intelligence (AI) coupled with the concept of *lights-out manufacturing* — the fully automated production of goods without any human intervention — have raised concerns about technological mass unemployment and growing inequality (Brynjolfsson and McAfee, 2014; Ford, 2015; Arntz et al., 2022).¹ These fears are further fueled by some speculative studies predicting striking job losses in occupations that are susceptible to automation (Frey and Osborne, 2017; Nedelkoska and Quintini, 2018). While automation is designed to perform tasks without human intervention (Nof, 2009), its overall impact on employment, the organisation of work and the skill composition of the workforce remains a subject of intense and ongoing debate. This paper seeks to examine who gains and who loses among workers when firms adopt automation technologies.

From a theoretical standpoint, the effects of adopting automation technology on employment are ambiguous, as both job-creating and job-destroying mechanisms may simultaneously be at play.² Automation can displace workers in specific roles (Keynes, 1930; Leontief, 1952; Brynjolfsson and McAfee, 2014), yet this displacement effect may be mitigated or even reversed through various compensatory effects. For example, automation can generate cost-saving efficiencies prompting firms to expand their production scale or enable the creation of new tasks increasing their demand for labor (Bowen and Mangum, 1966; Zeira, 1998; Autor, 2015; Acemoglu and Restrepo, 2019). Additionally, compensation may occur via a Schumpeterian mechanism, where the introduction of new products stimulates the demand for labor (Dosi et al., 2021; Tuhkuri et al., 2022). Nonetheless, even when displacement effects are counterbalanced, at the industry level there can still be negative spillovers. For example, as adopting firms enhance their productivity and capture greater market shares, non-adopting competitors may suffer declines in both productivity and employment (Nelson and Winter, 1982; Bloom et al., 2013). Importantly, since firms exhibit heterogeneous patterns of technology adoption and the diffusion of these technologies occurs gradually, the sectoral and economy-wide effects of labor-saving technologies are continuously shaped by the dynamic interactions between adopting and non-adopting firms over time (Dosi and Nelson, 2010). Ultimately, the impact of automation on employment and work organization remains an empirical question.

¹According to the Eurobarometer (2017) survey in Portugal—my focus area—over 90% of respondents believe that robots and AI steal jobs. These fears of widespread technological unemployment are far from new. From the Luddite movement in the early 19th century to contemporary debates on AI’s disruptive potential, concerns about automation replacing human labor have persisted for centuries (see e.g., Autor, 2015).

²For a more comprehensive discussion of the theoretical literature, see, e.g., Calvino and Virgillito (2018); Aghion et al. (2022a); Corrocher et al. (2023).

Empirical research on the relationship between automation and employment has grown rapidly, yet remains inconclusive. Most studies focus on industry-level adoption of industrial robots, reporting mixed outcomes—ranging from positive to neutral to negative—for more exposed regions and sectors (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020a; Gentili et al., 2020; Dauth et al., 2021; Dottori, 2021; Adachi et al., 2024; Aghion et al., 2023). At the firm level, a growing body of literature suggests that automation leads to increased employment within adopting firms (Dixon et al., 2021; Domini et al., 2021), though this expansion often comes at the expense of non-adopting competitors experiencing productivity and employment losses (Acemoglu et al., 2020, 2023; Koch et al., 2021). However, some studies also identify a negative relationship between the adoption of automation and employment (Bessen et al., 2020, 2023; Bonfiglioli et al., 2024).

In addition to its impact on employment, it is crucial to understand how automation reshapes the organization of production and the skill composition within firms. The dominant theoretical framework (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019, 2020b) suggests that lower-skilled workers, particularly those engaged in routine-intensive tasks—such as clerical and production roles—are most vulnerable to automation, as these tasks are more easily codifiable (Autor et al., 2003). Conversely, higher-skilled workers, such as STEM professionals, may benefit from automation’s compensatory effects, which create new opportunities for productivity and innovation (Acemoglu and Restrepo, 2019, 2022; Restrepo, 2023). As a result, following an automation event, the organization of work undergoes an upward skill shift. However, empirical evidence on the effects of automation on the skill composition of the workforce remains inconsistent. For instance, using French data, Acemoglu et al. (2020) finds negative employment effects for production workers, while other studies report positive effects for this group and no significant changes in overall workforce composition (Aghion et al., 2022b; Domini et al., 2021). Similarly, Acemoglu et al. (2023) find negative effects for low-educated workers in the Netherlands, while Bessen et al. (2023) highlights adverse effects on middle-educated workers. To take stock, we are far from a consensus about the impact of automation at the industry, firm and worker-level.

Against this backdrop, this paper presents novel firm-level evidence on the impact of automation technology on employment and skill composition within firms in Portugal. I construct a firm-panel dataset spanning the period from 2004 to 2021 by combining linked employer-employee data with firm balance sheets and trade transaction data. The richness of Portuguese administrative data sources provides me with detailed worker-level information on total hours worked and most importantly occupational groups and educational levels, allowing for a more granular analysis than many previous studies, which often suffer from the absence of such detailed information.³

³For instance, Koch et al. (2021) using Spanish data lack access to worker-level information, while

To measure investments in automation, I employ an import-based proxy approach, leveraging fine-grained product-level data on traded goods. I introduce a new classification to capture a broader range of advanced automation technologies, including industrial robots, numerically controlled machines, 3D printers, and technologies related to the Industrial Internet of Things (IIoT), providing a more detailed and accurate picture of automation’s impact across sectors. I identify automation events by leveraging the lumpy nature of automation investments and estimate the effects of these events on employment and workforce composition using a difference-in-difference event study design. To account for well-documented issues with heterogeneous effects in such empirical settings (de Chaisemartin and D’Haultfœuille, 2020a; Goodman-Bacon, 2021; Sun and Abraham, 2021; Borusyak et al., 2022), I employ the approach developed by Callaway and Sant’Anna (2021).

Portugal is particularly well-suited for this empirical strategy, as Portuguese firms heavily rely on imports of automation technologies due to the near absence of domestic manufacturers. In contrast, studies using this approach for countries like France, Italy, and the Netherlands (Acemoglu et al., 2020, 2023; Bisio et al., 2025), where domestic automation manufacturers are more prevalent (Andreoni et al., 2023; Castellani et al., 2022), are more prone to measurement errors, potentially leading to false negatives when identifying automation investments.

I contribute to the literature in two principal ways. First, my analysis suggests that automation leads to a reorganization of production within firms, primarily relying on a lower-skilled workforce. This is particularly evident in the marked increase in total hours worked by low-educated workers, blue-collar workers and those performing routine-manual tasks. In contrast, the results show no significant effects for highly educated workers or skilled professionals such as those in STEM occupations. Importantly, this pattern of workforce reorganization around lower-skilled workers is consistent across all types of automation technologies. These findings challenge the prevailing view that automation inherently favors high-skilled labor.

Second, whereas my overall estimates align with previous studies, that find a positive relationship between automation and employment (Acemoglu et al., 2020; Dixon et al., 2021; Koch et al., 2021), distinguishing by firm size reveals a dual nature in automation with job creation in small firms and job destruction in larger ones. Specifically, automation leads to a 19% increase in employment in small firms, whereas large firms experience a 30% decline in total hours worked within five years of an automation event. Importantly, I show that this pattern holds across a wide range of firm types, industries and types of automation technologies, reinforcing the critical role of firm size in shap-

the Italian (Bisio et al., 2025) and Dutch data sources (Bessen et al., 2023) do not provide occupational codes. Additionally, French data sources do not contain information about workers’ educational levels (Acemoglu et al., 2020; Domini et al., 2021)

ing automation’s labor market consequences. While small firms vastly outnumber large firms, the latter employ the majority of workers. As a result, when weighting by baseline hours worked, I find on average a net decline of 17% in total hours worked, suggesting that automation’s aggregate effect is negative. This is consistent with a recent study by Bisio et al. (2025), which find similar patterns in a different context for Italy. Crucially, this dual nature in automation with job creation in small firms and job destruction in large firms helps reconcile conflicting findings in the literature, as some studies reporting negative effects (Bessen et al., 2020, 2023) rely on weighted estimates that may mask positive effects among smaller firms.

Regarding skill composition in relation to firm size, I find that automation tends to lead to workforce reductions across most skill groups in large firms, whereas small firms experience workforce expansions across the same groups. Despite these divergent trends, it appears that low-skilled workers are favored relatively to skilled workers—such as STEM professionals—across both small and large firms. This is particularly striking given that STEM professionals in large firms, often seen as the primary beneficiaries of automation, experience the strongest negative effects. These patterns—consistent across firm size categories and types of automation technologies—challenge conventional assumptions about who gains from technological change and prompt a reconsideration of the characterization of the nature of technological change.

The rest of this paper proceeds as follows. In Section 2 I present my data sources and provide some novel stylized facts of automation that inform my empirical strategy, which I outline in Section 3. Section 4 discusses my main results, while Section 5 provides a bunch of robustness checks. Section 6 concludes.

2 Data and Descriptive Statistics

This paper’s empirical analysis combines and examines several administrative data sources from the Portuguese National Statistics Institute (INE). Using a unique firm identifier, I combine linked employer-employee records with firm balance sheet data and automation-related imports from trade transaction data. The final panel covers nearly all firms active in the private and non-financial sector in continental Portugal from 2004 to 2021. This section outlines the data sources, panel construction, and key variables, with further details in Appendix A.

2.1 Data sources

Linked employer-employee data. The *Quadros de Pessoal* (QP), administered by the Ministry of Employment of Portugal, draws on a compulsory annual census of all private

firms in Portugal that employ at least one dependent worker and has been run since 1985.⁴ The full sample includes about 350,000 firms and 3 million employees in each year. Each firm is required by law to report information on its characteristics, for each of its plants, and for each of its workers at the end of the census reference month (October) of each year. Entering the database, each firm and each worker are assigned a unique, time-invariant identifying number, making it possible to track firms and workers over time.⁵ The list of variables available in the data set is particularly rich including among other the firm's location, industry, date of creation, employment, sales, ownership and legal basis. At the worker level, the data covers information on worker's demographic characteristic (date of birth, gender, education and so forth) and job characteristics (occupation group, wage, hours worked, type of contract, and so forth).

Balance Sheet data. The next set of data to be used is the *Sistema de Contas Integrado da Empresa* (SCIE) available since 2004. It is a yearly census of firms collected by INE with the main objective to characterize the economic and financial performance of companies.⁶ The data covers all firms (companies, individual entrepreneurs, and self-employed) in the non-financial business sector that carry out an activity producing goods or services during the year.⁷ The large number of variables include, among others, the firm's industry classification (5-digit CAE Rev.3), births and deaths of companies as well as accounting statements, e.g. value added, sales and the wage bill.⁸

Trade transaction-level data. The third data source is the *Comércio Internacional* (CI), which records individual trade transactions (imports and exports) on a monthly basis for firms that are located in Portugal. The information is collected separately for EU partner countries and non-EU partners. Information on extra-EU transactions are collected by customs declarations, which cover the universe of external trade transactions ("Extrastat"). Since 1993, following the creation of the Single European Market, information on intra-EU transactions has come from the Intrastat system. It is a survey that

⁴Note that in the census public administration and non-market services are excluded. QP has been widely used by, amongst others, Blanchard and Portugal (2001), Cabral and Mata (2003) and Card et al. (2015, 2018).

⁵The purpose of this survey is to verify if firms are complying with labor law. Since employers are the one reporting the data, variables such as wages and worker qualifications are less prone to measurement error. The unique worker identifier is based on the worker's social security number. In addition, the Ministry of Employment implements several checks to certify the quality of the data. Overall this implies a high degree of coverage and reliability.

⁶The information gathered in this dataset is the result of a process of integrating information on companies, initially from the Annual Business Survey (IEH) and later from administrative information from the protocol signed with the Tax and Customs Authority

⁷Accordingly, financial and insurance companies are excluded, as are entities that are not market-oriented, namely central and local public administration units and various associative activities.

⁸All the data is subjected to comparisons between year n and year $n-1$, paying special attention to possible deaths and births of companies, as well as changes in activity, region, levels of staff employed, turnover and legal form. Whenever necessary, companies are contacted to provide additional clarifications, which may or may not lead to occasional corrections, in order to guarantee the statistical quality of the information to be made available.

must be answered by firms whose value of intra-EU trade transactions exceeds a legally binding threshold that has changed over time. For example, in 2015 the Portuguese thresholds were set at 250,000 Euro and 350,000 Euro for exports and imports, respectively. Throughout the sample period, the thresholds were set to ensure that the survey overall includes at least 97 percent of intra-EU exports and 93 percent of intra-EU imports each year.⁹ For each transaction the database reports, among other variables, trade value (in euros), trading partner country and an 8-digit Combined Nomenclature (CN) product code, an extension of the 6-digit international Harmonized System (HS) trade classification. I use CI to identify importing firms and their imports of automation-related goods based on the 6-digit HS code.

2.2 Sample Selection

I combine firm-level workforce information from QP with firms' accounting information from SCIE and trade information from CI using the unique, time-invariant firm identifier. I keep only firm observations that appear both in the QP and SCIE data sources. The integration of my three sources restrict the analysis to the period 2004–2021. I assign to each firm a permanent 2-digit sector based on the most frequent occurrence, since firm industry codes may vary across years. I consider in my analysis firms based in continental Portugal and exclude individual entrepreneurs and self-employment enterprises. Regarding workers, I restrict the analysis to dependent workers aged between 18 and 65 years old (working-age population). Observations with unreasonable values (e.g., non-positive turnover or gross value added, non-complete basic remuneration, etc.) were discarded. Since automation adoption is measured through imported product categories, most analyses focus on importing firms. I distinguish between the full sample, the importer sample, which includes firms that import at least once during 2004–2021, and the adopters sample, which includes firms that import automation technologies. I exclude intermediaries of automation goods to prevent misclassification of adopters.¹⁰ All monetary values are being deflated using the Consumer Price Index provided by INE.

⁹Note that this data is the country's official information source on imports and exports used for official trade statistics.

¹⁰Following Bisio et al. (2025), I began with wholesalers, identified by NACE Rev. 2 group codes, and determined the relevant three-digit sectors as defined in Table (4). Further discussion can be found in Section 2.3.

2.3 Automation imports

Firm-level data on automation investments remain scarce.¹¹ To address this issue, I construct a proxy for automation investments using firm-level data on the yearly imports of automation-related capital goods. Using imports as a proxy for technology adoption and diffusion is well-established in the literature (Caselli and Wilson, 2004; Andreoni et al., 2023). Originating from the seminal work by Caselli and Coleman (2001), which investigated the diffusion of computers, this method has been recently adapted to measure automation technology adoption at the firm level across various countries.¹²

The rationale for using an import-based proxy is grounded in the observation that innovation and the production of capital goods, including automation technologies, are concentrated in a few industrialized countries, making international trade a crucial channel for technological diffusion (Eaton and Kortum, 2001; Caselli and Coleman, 2001). For example, Andreoni et al. (2023) show that 14 countries account for over 90% of global digital exports, though their study includes a broader range of technologies beyond automation. However, countries like Denmark, France, Italy, and the Netherlands, frequently analyzed in the literature, are net exporters of automation technologies, which can introduce measurement bias in an import-based approach. This limitation is less of an issue for Portugal, where most automation technologies are imported (Castellani et al., 2022). Due to the absence of domestic production, Portuguese firms rely heavily on foreign equipment, making the import-based proxy an effective measure of their automation investment decisions.

I propose a novel, more comprehensive measure of automation that encompasses a broader range of automation technologies. While previous studies have predominantly focused on industrial robots as a proxy for automation (e.g., Acemoglu et al., 2020; Humlum, 2019; Dixon et al., 2021; Koch et al., 2021), it is important to recognize that investments in industrial robots represent only a small fraction of total automation technologies and are concentrated in a limited number of industries (Montobbio et al., 2022; Domini et al., 2021; Aghion et al., 2022b). Broader measures adopted by other studies (Domini et al., 2021; Bisio et al., 2025) often fail to account for emerging technologies like 3D printers or the Industrial Internet of Things (IIoT), which hold significant potential for advancing automation.

To address this gap, I identify firm-level imports of capital-embodied automation technologies using 6-digit Harmonized System (HS) product codes, categorizing them

¹¹Firm-level data on automation investments is available for only a few countries, such as automation expenditures by Dutch firms (Bessen et al., 2023) and data on robot usage and intensity in Germany and Spain (Deng et al., 2021; Koch et al., 2021). Some cross-sectional evidence exists, e.g., for the United States (Acemoglu et al., 2022), but these datasets are not well-suited for analyzing employment dynamics.

¹²For example, it has been applied for Canada (Dixon et al., 2021), Denmark (Humlum, 2019), France (Acemoglu et al., 2020; Domini et al., 2021, 2022; Aghion et al., 2022b), Italy (Bisio et al., 2025), the Netherlands (Acemoglu et al., 2023), and Norway (Barth et al., 2020).

into four groups: First, industrial robots and conventional automation technologies, such as numerically controlled machines, automatic machine tools, and automatic welding machines, which have been widely studied in previous research (Acemoglu and Restrepo, 2021; Domini et al., 2021; Bisio et al., 2025). Second, additive manufacturing (also known as 3D printing), which refers to technologies that create physical objects through the successive addition of material, in contrast to traditional subtractive manufacturing methods (International Organization for Standardization, 2015). Third, automatic data processing (ADP) machines that automate tasks related to data processing, reducing labor-intensive manual data handling.¹³ Finally, Industrial Internet of Things (IIoT) refers to the use of interconnected smart devices in industrial applications, enabling machine-to-machine communication for the collection, monitoring, and analysis of data from operational processes. The primary objective of IIoT is to leverage sensors and automation to enhance production flexibility, efficiency, and cost-effectiveness (Atzori et al., 2010; Gubbi et al., 2013; Wang et al., 2016). A detailed breakdown of the HS codes and categories is provided in Table (3) in the appendix.

The import-based approach has certain potential limitations, which I address. First, firms may be misclassified as adopters if they resell imported automation goods domestically or internationally. This is by definition true for resellers and intermediaries of automation technologies. To mitigate this issue, I follow Bisio et al. (2025) and exclude firms operating in these sectors from the analysis, as detailed in Table (4). Furthermore, reselling activity is typically associated with regular imports. However, Figure (3) demonstrates that automation imports occur infrequently within firms, making such misclassification unlikely. To further enhance the robustness of the analysis and reduce potential noise in my automation measure, I apply varying restrictions to the treatment definition. These include restricting the analysis to firms that import only once and excluding firms engaged in re-exporting activities, as outlined in section 5. Importantly, these varying restrictions do not alter the results.

Second, firms may be incorrectly labeled as non-adopters if they source automation technologies domestically rather than importing them. However, given the structure of the Portuguese economy, it is unlikely that firms buy automation technologies from domestic producers. As Juchniewicz and Łada (2020) shows, Portugal’s high-tech sector lags significantly behind and is less competitive than those of other European countries. Additionally, patent data from Cséfalvay and Gkotsis (2022) reveal a lack of domestic suppliers of industrial robots in Portugal. As a result, Portugal is a net importer of automation technologies (Castellani et al., 2022). That said, firms may still rely on intermediaries to import goods, as documented by Blum et al. (2010), Bernard et al. (2010), and Ahn et al. (2011). While the use of intermediaries is less common for complex goods


¹³Notably, Domini et al. (2021) argue, based on USPC-to-HS concordances, that this group is linked to artificial intelligence technologies.

requiring high relational specificity—such as automation technologies—firms in smaller, lower-income countries like Portugal may be more likely to use them (Bernard et al., 2015). To explore this further, Table 5 examines the role of intermediaries in automation imports. It shows that, among firms that have imported automation technologies, approximately 8.2% are active in the sectors defined in Table 4, accounting for 8.5% of the total value of automation imports. This suggests that, while intermediaries do play a role in automation diffusion, the vast majority of firms do not rely on them for purchasing goods from abroad. While this may introduce some noise, it is unlikely to significantly affect the accuracy or validity of my measure of automation. Moreover, my empirical strategy, as detailed in Section 3, compares treated firms with those yet to adopt automation. This approach minimizes the risk of misclassifying firms as non-adopters, as it focuses on firms at different adoption stages, effectively eliminating any misclassification bias.

2.4 Stylized Facts about Automation

In this section, I present five stylized facts about automation imports in Portugal. These findings serve to inform the empirical strategy and provide crucial insights into the patterns of automation diffusion.

Figure (1) Automation imports between 2004-2021



Notes: Panel A of Figure (1) depicts the total value of automation-related capital goods imports (red line, in millions of euros) alongside the annual share of importing firms, distinguishing between small (1–49 employees) and large (>49 employees) firms, with their shares adding up to the total. Panel B breaks down the composition of automation imports by technology, including Industrial Robots (IR), Conventional Automation (CA), Additive Manufacturing (AM), Automatic Data Processing Machines (ADP), and the Industrial Internet of Things (IIoT), as a share of total automation imports (see Table (3) for a detailed breakdown of technologies).

Fact 1. Automation Adoption is Slow. Contrary to the prevailing narrative of

rapid and disruptive automation diffusion (Brynjolfsson and McAfee, 2014; Ford, 2015), the adoption of automation technologies in Portugal has been slow. Panel A of Figure (1) shows that only around 2% of firms import automation technologies annually, with smaller firms making up the majority. Automation imports are highly sensitive to economic downturns. The global financial and Eurozone crises triggered a sharp, prolonged decline, stalling adoption for several years. A recovery began only after 2013, though the COVID-19 pandemic caused a temporary setback before imports surged in 2021. By then, about 7% of firms had imported automation technologies at least once, though these firms employ roughly 36% of the workforce, highlighting their significant role in labor market dynamics.

Panel B of Figure (1) shows the composition of automation imports by technology type. Automatic data processing machines remain dominant, though their share has declined from around 50% to 40%. Conventional automation consistently accounts for over 25% of imports, while the share of Industrial Internet of Things (IIoT) technologies has grown from 14% in 2004 to 28% in 2021. In contrast, industrial robots and additive manufacturing remain marginal. These findings challenge the narrative that robots are the most disruptive technological shift. In Portugal, automation adoption has been driven more by traditional technologies than advanced robotics, a trend also observed in Italy (Bisio et al., 2025).


Fact 2. Automation diffusion is skewed. Panel A of Figure (2) illustrates the highly skewed nature of automation adoption across industries, with rates ranging from under 1% to nearly 19%. The highest adoption rates are observed in primary sectors, manufacturing, and high-tech industries like Information and Communication. While only a small proportion of firms adopt automation, Panel B shows that these firms disproportionately account for a significant share of total employment, ranging from 6% to 75% of employment in their respective industries. Notably, small firms, though more numerous, contribute little to total employment. In contrast, large firms, despite being fewer in number, employ a substantial portion of the workforce. This trend holds across all industries, although the balance between small and large adopters varies by sectors.

Panel C of Figure (2) highlights distinct patterns in automation technology imports across sectors. In primary and manufacturing industries, conventional automation, industrial robots, and additive manufacturing are more common, although industrial robots and additive manufacturing remain relatively rare. In contrast, service sectors are dominated by automatic data processing (ADP) technologies, followed by Industrial Internet of Things (IIoT) technologies.


Fact 3. Automation investments are lumpy. Automation imports exhibit the characteristic "lumpy" behavior typical of investment variables, meaning they are (i) rare within firms and (ii) concentrated in a few episodes that account for a large share of total investments (Letterie et al., 2004; Asphjell et al., 2014; Grazzi et al., 2016).

Figure (2) Automation diffusion across industries


Panel A. Share of Adopters



Panel B. Employment Share of Adopters




Panel C. Technology composition



Notes: Panel A of Figure (2) shows the share of firms in 2021 that imported automation technologies. Panel B presents the corresponding share of employment within these firms in 2021, with the red line representing the overall mean. Firms are categorized as small (< 50 employees) and large (≤ 50 employees), with both adding up to the total share. Panel C breaks down automation imports by technology, including Industrial Robots (IR), Conventional Automation (CA), Additive Manufacturing (AM), Automatic Data Processing Machines (ADP), and the Industrial Internet of Things (IIoT) from 2004 to 2021, as a share of total automation imports (see Table (3) for a detailed breakdown).

Panel A of Figure (3) shows that approximately 43% of firms that import automation technology do so only once. Import frequency declines steadily with larger values, with a small group of firms consistently importing automation-related goods every year. Panel

Figure (3) Automation investments occur in spikes



Notes: Panel A shows the distribution of the number of years in which firms import automation technology. Panel B ranks the mean share of each year's automation imports relative to total imports, from largest to smallest. The figure separately distinguishes between small firms (with < 50 employees) and large firms (≥ 50 employees) in their baseline year. Note that the total, small, and large firms represent different samples and are not cumulative. The figures are based on the sample of adopters only.


B illustrates that the largest import episode typically accounts for about 75% of a firm's total automation-related imports, with the median episode representing more than 90%. Subsequent import episodes represent much smaller shares. These findings align with studies on the lumpy nature of automation investments in Denmark (Humlum, 2019), France (Domini et al., 2021), Italy (Bisio et al., 2025), and the Netherlands (Bessen et al., 2023).

Additionally, the lumpiness of automation imports is evident for both small and large firms. However, the degree of lumpiness is more pronounced for small firms. About 47% of small firms import automation technology only once, with an average of 77% of their total imports concentrated in a single year. For large firms, 24% import automation only once, and while the pattern of decreasing frequency remains, lumpiness is less pronounced. Notably, persistent importers—firms that import every year—represent a substantial share, contributing some noise to the measure. Still, 64% of total imports for large firms occur in a single year. In Section (5), I evaluate whether such noise impacts my results; however, the analysis confirms that it does not.

Fact 4. Automation spikes involve more than just robots. While most studies focus on industrial robots, actual automation adoption patterns remain unclear—particularly whether firms implement multiple technologies or standalone solutions. Figure (4) explores this by analyzing the prevalence and combinations of different automation technologies during adoption spikes.

Panel A shows that most firms (73%) adopt only one of five technological groups,

Figure (4) Technology composition of automation spikes



Notes: Panel A of Figure (4) shows the distribution of firms by the number of technologies adopted during an automation spike, distinguishing between small (1–49 employees) and large (≥ 50 employees) firms. Panel B illustrates the technological composition within each adoption group, detailing the shares of Industrial Robots (IR), Conventional Automation (CA), Additive Manufacturing (AM), Automatic Data Processing Machines (ADP), and the Industrial Internet of Things (IIoT). Panel C examines the co-adoption patterns within each technological group, indicating the number of technologies adopted alongside a given technology. The figures are based on the adopters-only sample, focusing on firms in the year of the automation spike.

while 27% combine two or more—though adopting more than two remains rare. This pattern holds across firm sizes, with larger firms being more likely to adopt multiple technologies.

Panel B differentiates between firms adopting a single technology and those adopting multiple. Among single-technology adopters, conventional automation dominates (48%), followed by automation data processing and IIoT, while industrial robots and additive manufacturing are rarely adopted. Among multi-technology adopters, conventional automation, data processing, and IIoT are the most common combinations, whereas robots and additive manufacturing remain marginal. These patterns hold across both small and large firms.

These findings challenge the prevailing focus on industrial robots, raising questions about studies that adopt this narrow perspective (e.g., Acemoglu and Restrepo, 2020a; Acemoglu et al., 2020; Humlum, 2019).

Fact 5. Firms select into automation adoption. Table (1) compares firms that adopt automation technologies with non-adopting importers, revealing significant differences in firm characteristics and workforce composition. Automating firms tend to be larger and offer higher wages than non-adopters. Additionally, their workforce is more

Table (1) Summary Statistics: Selection into Adoption

Variable	Adopters			Non-adopters			P-value
	Mean	SD	<u>Baseline</u>	Mean	SD	<u>Industry</u>	B-I
Employees	55.28	282.20	51.59	14.67	51.52	28.32	***
Total hours worked	9555.21	48282.93	8923.58	2532.78	8890.90	4879.38	***
Labor Productivity	36.21	515.31	34.88	25.34	220.52	25.85	***
Labor Share	0.52	0.18	0.51	0.54	0.19	0.58	***
Hourly Wage	7.30	3.67	7.34	5.94	2.97	6.66	***
Exporter	0.47	0.50	0.46	0.24	0.43	0.32	***
Multinational	0.13	0.34	0.13	0.04	0.20	0.04	***
Share of workers							
Low Education	0.47	0.33	0.45	0.53	0.37	0.59	***
Middle Education	0.31	0.26	0.32	0.29	0.30	0.25	***
High Education	0.21	0.27	0.23	0.17	0.28	0.16	***
Routine-Manual	0.33	0.35	0.31	0.29	0.37	0.35	***
Routine-Cognitive	0.09	0.21	0.09	0.12	0.26	0.11	***
Manual	0.04	0.13	0.04	0.06	0.17	0.05	***
Cognitive	0.53	0.35	0.55	0.52	0.39	0.48	***
Non-professional	0.38	0.33	0.38	0.47	0.38	0.44	***
Bluecollar	0.29	0.33	0.28	0.26	0.35	0.31	***
Professional	0.12	0.21	0.13	0.13	0.25	0.11	***
STEM	0.14	0.25	0.15	0.07	0.20	0.09	***
Manager	0.06	0.14	0.07	0.07	0.18	0.06	***
Number of Firms	208,714		14,258	400,236		14,258	

Notes: Table (1) presents summary statistics comparing adopters and non-adopters. Column (1) and Column (4) report the mean values for the adopters and non-adopters (importers), respectively. Column (3) displays the mean baseline values, measured in the year before adoption. Column (6) presents the average values for a randomly selected group of non-adopters matched to adopters within the same industry-year cell (one-to-one matching). Finally, Column (7) reports p-values for the null hypothesis that the mean values of adopters (Column 3) and matched non-adopters (Column 6) are statistically equivalent. Statistical significance is indicated by *, **, and *** for p-values below 0.1, 0.05, and 0.01, respectively.

skilled, with a higher share of educated and specialized employees.

Crucially, these differences are already evident before adoption. When comparing firms in the year prior to adoption with non-adopters in the same industry-year cell, I observe that adopters were systematically different even before implementing automation. This suggests that firms do not adopt automation randomly; rather, there is a clear selection process into automation adoption. These pre-existing differences underscore the importance of accounting for selection effects when analyzing the impact of automation on firms and labor markets.

3 Empirical Strategy

This section outlines my empirical strategy, which is grounded in three key observations from section 2.4. First, automation investments are lumpy, that means they occur in discrete spikes. This justifies to model automation adoption as a one-off decision and the use of a difference-in-difference (DiD) event study design to estimate its impact on firm employment and workforce composition. Second, firms do not adopt automation randomly; rather, they exhibit systematic differences from non-adopters even before adoption, raising concerns about selection bias. To mitigate this issue, I use firms that have not yet adopted automation as the control group. Third, automation adoption varies across firms, industries, and technologies. This motivates an analysis of how socio-economic factors shape the outcome of automation adoption.

Treatment definition. In Section 2.4, I showed that automation investments tend to occur in spikes, a statistical property that allows for the identification of distinct automation events. Given the highly skewed distribution of this variable, I define a firm’s largest automation investment as an automation event.¹⁴ Figure (11) provides empirical validation of this definition. Panel A shows the share of total automation imports around the event year, revealing a distinct one-period surge in the adoption year. Panel B examines whether this spike is also reflected in balance sheet data by plotting the yearly share of total machinery investment around the automation spike. Here, too, a clear pattern emerges, reinforcing the validity of my definition of treatment. In addition, I provide further robustness exploring various restrictions and alternative definitions of spikes in Section 5. None of these affect my results.

Event study design My empirical setting involves a binary, absorbing treatment where automation adoption occurs at staggered intervals across firms and time. I exploit this variation in treatment timing by employing a difference-in-differences (DiD) approach. Recent literature has highlighted several limitations of the traditional Two-Way Fixed Effects (TWFE) estimator in settings with staggered treatment adoption and treatment effect heterogeneity (Meer and West, 2016; Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020b; Sun and Abraham, 2021; Borusyak et al., 2022). While the TWFE estimator is commonly used in DiD analyses, it can yield biased estimates in such contexts. This bias arises because it computes a weighted average of all possible two-by-two DiD comparisons. These comparisons include both “good comparisons” (with units that have never been treated or are not yet treated) and “bad comparisons” (involving units that have already been treated).¹⁵ To address these methodological issues,

¹⁴See Domini et al. (2021, 2022) and Bisio et al. (2025) for similar definitions.

¹⁵Most previous firm-level studies on automation do not to account for these issues, leading to potentially misleading conclusions (e.g., Acemoglu et al., 2020; Koch et al., 2021; Domini et al., 2021). Some

I use the estimator proposed by Callaway and Sant’Anna (2021), which is more robust in such settings.¹⁶

The Callaway and Sant’Anna (2021) estimator involves a two-step procedure. First, it estimates group-time average treatment effects (ATT) for each cohort of adopters. Then, it aggregates these effects to derive parameters of interest, such as the overall ATT or event study estimates. Formally, I estimate the group-time ATT as follows:

$$\text{ATT}(g, t) = E[Y_{i,t} - Y_{i,g-1} | G_i = g] - E[Y_{i,t} - Y_{i,g-1} | G_i = G_{\text{comp}}] \quad (1)$$

where $\text{ATT}(g, t)$ represents the average treatment effect at time t for the cohort that adopts automation in year g .¹⁷ Here, $Y_{i,t}$ is the dependent variable of interest at time t for firm i , G_i indicates the year of adoption for firm i , and G_{comp} denotes the set of comparison firms for an adopter at time g , which have either never adopted or have not yet adopted by time t . Formally $G_{\text{comp}} = \{g' | g' > t\}$. By construction, this method relies solely on “good comparisons” and does not impose restrictive assumptions on treatment effect heterogeneity.

To compute the event-study parameters, I take a weighted average of the treatment effects l periods after adoption across different cohorts:

$$\widehat{\text{ATT}}_l = \sum_g w_g \text{ATT}(g, g + l) \quad (2)$$

where the weights w_g reflect the relative frequency of each cohort among adopters. I cluster standard errors at the firm level and report 90% and 95% bootstrap confidence intervals. The validity of my DiD approach relies on two key assumptions: (i) parallel trends and (ii) no anticipation.

Parallel trends. The parallel trends assumption requires that, in the absence of treatment, the average outcome for the treatment and control group would have followed parallel trends over time. However, as highlighted in Section 2.4, automating firms differ

exceptions exist: for instance, Bessen et al. (2023) employs a stacked DiD design as proposed by Cengiz et al. (2019). However, despite its appeal, Wing et al. (2024) argue that stacked regressions can be biased due to differential implicit weighting of treatment and control trends across sub-experiments. Another exception is the work by Bisio et al. (2025), which similarly to me employs the approach by Callaway and Sant’Anna (2021).

¹⁶The approach by Callaway and Sant’Anna (2021) also enhances the transparency and objectivity of the analysis, as advocated by Rubin (2007, 2008), by clearly defining the control group. Given the challenges of estimating the effects of automation, such methodological rigor is particularly valuable in my context.

¹⁷For example, $\text{ATT}(2015, 2017)$ would be the average treatment effect in 2017 for firms that adopt automation technology in 2015.

significantly from non-adopters, e.g., in terms of size and productivity. Indeed, automation is a decision by the firm, hence a self-selection into treatment. To account for this issue, I define the comparison group to include only firms that have not yet adopted automation at the time of comparison. Additionally, I condition parallel trends on a set of relevant covariates.¹⁸ Although the parallel trend assumption is not directly testable, its validity can be partially assessed by examining pre-treatment trends. I visually inspect whether the pre-trend coefficients are different from zero and formally test for pre-treatment parallel trends based on group-time average treatment effects (Callaway and Sant’Anna, 2021). Finally, I perform various robustness checks and sensitivity analyses, as detailed in Section 5. My findings remain unchanged across these controls.

No anticipation. The no-anticipation assumption requires that firms do not adjust their behavior in anticipation of automation adoption. This means that adopting automation in time t should not affect the outcome variables before t . Although this assumption is not testable, I conduct robustness checks by varying the horizons of the baseline period to detect any potential anticipatory behaviors (Callaway and Sant’Anna, 2021).

4 Results

This section presents the main findings from my difference-in-differences (DiD) event study design to identify the winners and losers of automation. First, I examine automation’s impact on total hours worked, with a particular focus on how firm size and other socio-economic characteristics shape these outcomes. Second, I explore more granular effects by investigating automation’s impact on the skill composition of the workforce.


Effect on Employment. Figure 5 presents the impact of an automation spike on the log of total hours worked at the firm-level. Panel A estimates indicate a 9% increase in total hours worked five years post-adoption. The visual inspection shows that pre-trends in employment were similar between the treatment and control groups and I cannot reject the null hypothesis of conditional parallel trends. This finding aligns with a growing body of evidence suggesting a positive relationship between automation and firm-level employment, as observed in various studies across different countries.¹⁹ These studies typically argue that automation can enhance productivity, leads to new task or job creation,

¹⁸In my benchmark exercise I control for baseline values of the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status using the doubly-robust (DR) estimation procedure (Sant’Anna and Zhao, 2020; Callaway and Sant’Anna, 2021).

¹⁹For example, for Canada (Dixon et al., 2021), Denmark (Humlum, 2019), Finland (Tuhkuri et al., 2022), France (Acemoglu et al., 2020; Aghion et al., 2022b), Italy (Bisio et al., 2025) and Spain (Koch et al., 2021).

thereby driving overall employment growth within firms. Some studies, however, offer a contrasting narrative by highlighting the job-destroying nature of automation, identifying a negative relationship between automation and employment (Bonfiglioli et al., 2024; Bessen et al., 2020, 2023). These contradictory findings reveal a fundamental puzzle in the current literature.

Figure (5) Event study: Job creating and destroying effects of automation



Notes: The figure provides event study estimates of adopting automation using the approach of Callaway and Sant’Anna (2021) as specified in equation (2). The comparison group comprises firms that have not yet adopted automation, with the analysis conducted under the assumption of conditional parallel trends. I control for the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status at the firm’s baseline ($\tau = -1$). The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004–2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1–49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

Socio-economic context of automation. A potential explanation for the discrepancies in the literature is that the average effect of automation on employment may mask heterogeneous treatment effects across firms. In other words, differences across firms may influence the outcomes of automation (Koch et al., 2021; Bisio et al., 2025). One commonly cited factor, which could influence automation outcomes in opposing ways is firm size. Larger firms may experience greater productivity gains from automation than smaller firms due to economies of scale, which in turn, may result in higher employment growth. (Koch et al., 2021; Humlum, 2019; Acemoglu et al., 2022). Conversely, a substantial body of literature on firm growth determinants suggests an inverse relationship between firm size and employment growth (Birch et al., 1979; Neumark et al., 2011). Moreover, scale effects from automation are frequently linked to export status (Koch et al., 2021) and multinational ownership (Leone, 2023). Some studies also find

that firms with greater technological sophistication benefit more from exposure to international competition (Bloom et al., 2015). Additionally, firm age plays a crucial role in employment dynamics, with multiple studies documenting a negative correlation between firm age and firm-level employment growth (Delmar et al., 2003; Haltiwanger et al., 2013). Beyond firm-level characteristics, sectoral differences may be crucial, as automation could reduce jobs in some industries while fostering employment in others (Bogliacino and Pianta, 2010; Reljic et al., 2023). Similarly, different types of automation technologies may apply across varying contexts and affect distinct groups of workers. While such socioeconomic factors may drive variation in automation’s employment effects, their influence remains unclear. The following analysis investigates their role in shaping automation outcomes.

Table (2) Event study: Job creating and destroying effects of automation

Event time	Unweighted (All)		Weighted (All)		Small firms		Large firms	
	ATT	SE	ATT	SE	ATT	SE	ATT	SE
-4	0.015	(0.009)	-0.010	(0.036)	0.023	(0.012)	-0.031	(0.015)
-3	0.019	(0.008)	0.020	(0.03)	0.026	(0.01)	-0.030	(0.013)
-2	0.005	(0.006)	0.008	(0.028)	0.012	(0.008)	-0.026	(0.01)
0	0.046	(0.005)	-0.016	(0.034)	0.067	(0.006)	-0.046	(0.01)
1	0.084	(0.007)	-0.104	(0.048)	0.127	(0.008)	-0.106	(0.013)
2	0.098	(0.009)	-0.166	(0.049)	0.156	(0.011)	-0.152	(0.016)
3	0.104	(0.011)	-0.203	(0.057)	0.180	(0.014)	-0.204	(0.022)
4	0.107	(0.013)	-0.262	(0.063)	0.191	(0.015)	-0.241	(0.021)
5	0.094	(0.015)	-0.263	(0.083)	0.190	(0.017)	-0.292	(0.028)
Overall ATT	0.089	(0.008)	-0.169	(0.047)	0.152	(0.009)	-0.174	(0.014)
Pre-trend p-value	0.350		0.140		0.530		0.170	
Number of Firms	14720.000		14720.000		11956.000		2764.000	

Notes: This table reports overall ATT and dynamic ATT estimates from the event study in equation (2), following Callaway and Sant’Anna (2021). The outcome variable is the log of total hours worked. The comparison group consists of firms that have not yet adopted automation under the assumption of conditional parallel trends. I control for the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status at the firm’s baseline ($\tau = -1$). Columns (1)-(2) show unweighted results, (3)-(4) present results weighted by baseline total hours worked (net employment effect for adopters), and (5)-(6) and (7)-(8) report results for small and large firms, respectively. The coefficient for the year before adoption is normalized to zero. Bootstrapped standard errors, clustered at the firm level, are in brackets.


Small vs. Large Firms. To assess whether automation’s employment effects vary by firm size, I separately analyze its impact on small (1–49 employees) and large firms (50+ employees) based on their baseline size ($\tau = -1$). Panel B of Figure 5 reveals a stark contrast: automation increased employment in small firms by approximately 15%, while large firms experienced a 30% decline in total hours worked. Since small firms dominate in number among adopters, but large firms account for most employment (see Section

2.4, the aggregate employment effects of adopters likely differ from those in Panel A of Figure (5). To quantify the net impact, I weight the results by total hours worked in the baseline year. As shown in Table (2), the negative effect dominates, leading to an average 17% reduction in total hours worked. This pattern aligns with recent evidence from Italy (Bisio et al., 2025), highlighting the dual nature in automation with job creation in small firms and job destruction in larger firms. Moreover, this divergence between small and large firms may help reconcile conflicting findings in the literature, as some studies reporting negative effects (Bessen et al., 2020, 2023) rely on weighted estimates that may mask positive effects among smaller firms.

Other socio-economic factors. Figure (6) examines the impact of automation across various firm characteristics and industries. Panels A and C present the overall effects, while Panels B and D further break down the results by firm size. Panel A shows positive employment effects for young firms, non-exporters, multinational firms, and domestic firms, while exporters exhibit no effect and old firms experience employment losses. Panel B confirms the dual pattern: small firms see job creation, while large firms face job destruction, with the exception of small old firms, where the effect is positive but not statistically significant. Industry heterogeneity in Panel C reveals predominantly positive employment effects across sectors, though some are not statistically significant. Panel D, which further differentiates by firm size, highlights a consistent pattern across all industry categories—small firms benefit from automation, while large firms experience employment declines. Notably, Table (7) shows that firm types and sectors with statistically insignificant results for the total sample tend to have a higher proportion of large firms, reinforcing the idea that sample composition, particularly firm size, can explain discrepancies in the literature. Finally, I examine heterogeneous effects across different automation technologies. This consistent pattern is also observed across all five automation technologies, with small firms consistently experiencing job creation, while large firms see employment declines or no effect. Taken together, these findings confirm a robust and recurring result: regardless of firm type, industry, or technology, firm size determines the direction of automation’s impact on employment dynamics.

Effect on the workforce composition. Automation influences not only the number of jobs but may also change the composition of the workforce. The prevailing view in the literature suggests that automation is skill-biased, meaning it tends to displace low-skilled workers engaged in simple, repetitive tasks (routine-manual work), while increasing demand for workers in more complex roles that require a broader range of skills, such as managers and STEM professionals (Acemoglu and Restrepo, 2019, 2020b; Restrepo, 2023). However, an alternative perspective argues that automation can lead to “deskilling,” where work is reorganized around lower-skill levels than previously required


Figure (6) Heterogeneity analysis by firm and industry characteristics



Notes: The figure provides overall ATTs based on event study estimates of dynamic effects of adopting automation on employment shares using the approach of Callaway and Sant’Anna (2021) specified in equation (2). The comparison group comprises firms that have not yet adopted automation, with the analysis conducted under the assumption of conditional parallel trends. I control for the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status at the firm’s baseline ($\tau = -1$). The coefficient for the year prior to adoption is normalized to zero. Panel A segments the sample by firm characteristics, distinguishing between young (age ≤ 15) and old firms (age > 15); exporting and non-exporting firms, multinational firms (foreign capital share $\leq 50\%$) and domestic firms ($> 50\%$) based on their baseline values. Panel C differentiates firms by sector, categorizing them into manufacturing, services, and varying levels of digital intensity, following the sectoral taxonomy outlined by Calvino et al. (2018). Panel B and D further disaggregate these groups by firm size. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided at both the 5% and 10% significance levels.

(Braverman, 1974; Kunst, 2020; Downey, 2021; Danieli, 2022). Furthermore, some empirical studies point to skill-neutral effects, where automation does not significantly change

Figure (7) Heterogeneity Analysis by Technology Groups



Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant’Anna (2021), differentiating between various groups of technologies: Industrial Robots (IR), Conventional Automation (CA), Additive Manufacturing (AM), Automatic Data Processing Machines (ADP), and the Industrial Internet of Things (IIoT). For more detailed information refer to table (3). The spike is defined in relation to the largest event of imports for the respective technology type. The comparison group consists of firms that have not yet adopted the respective technology. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004–2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1–49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

the overall skill distribution of the workforce (Domini et al., 2021; Tuhkuri et al., 2022).


Figure (8) illustrates the impact of automation on workforce composition across three key dimensions: (i) education levels, (ii) task groups, and (iii) macro-occupational categories.²⁰ Given the heterogeneous patterns observed earlier, Panel A presents results for the full sample, while Panel B differentiates between small and large firms to uncover potential firm-size-driven disparities. Broadly, an increasing demand for higher-skilled workers aligns with the notion of skill-biased technological change, whereas a shift toward lower-skilled roles supports the deskilling hypothesis.²¹

Looking first at educational groups, Panel A shows a strong positive effect of automation adoption on the total hours worked by low-educated workers. A similar but less pronounced effect is observed for middle-educated workers, while the impact on highly educated workers appears negligible. Panel B, which differentiates by firm size, reveals a downsizing trend across all educational groups in large firms and an expansion in small firms. However, the magnitude differs significantly between lower- and higher educated workers, suggesting a relative shift toward lower educated workers.

²⁰For detailed definitions of these categories, see Table (6).

²¹Note that this analysis focuses on shifts between occupational and educational groups rather than within-group changes.

Figure (8) Impact on Skill Composition



Notes: The figure provides overall ATTs based on event study estimates of dynamic effects of adopting automation on total hours worked by skill groups using the approach of Callaway and Sant’Anna (2021) specified in equation (2). The figure distinguishes between workers according to three key dimensions: (i) educational attainment, (ii) routine-intensity of their tasks and (iii) macro-occupational group. For more detailed information refer to table (6). The comparison group comprises firms that have not yet adopted automation, with the analysis conducted under the assumption of conditional parallel trends. I control for the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status at the firm’s baseline ($\tau = -1$). The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size —distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided at both the 5% and 10% significance levels.

Regarding task groups, Panel A shows that automation significantly increases the working hours of routine-manual workers, mirroring the strong effect observed for low-educated workers. While other task groups also experience positive effects, these are notably smaller in magnitude. Panel B further clarifies the pattern of job creation in small firms and job destruction in large firms across all task categories. Importantly, the demand for routine-manual workers is particularly strong in small firms, reinforcing the notion that automation fosters employment growth in lower-skilled, manual-intensive roles within smaller enterprises.

Turning to macro-occupational groups, Panel A indicates positive effects on total hours worked for non-professional, blue-collar, and professional workers, though the increases are far more pronounced for blue-collar and non-professional workers. In contrast, I find no significant effects for STEM professionals and managers. Once again, Panel B confirms the opposing trends of downsizing in large firms and expansion in small firms. Notably, in large firms, the effect on managers is not statistically significant, and the decline in hours worked is most pronounced for STEM professionals. Conversely, in small firms, the demand for lower-skilled workers is strongest for blue-collar and non-

professional workers.

Overall, these findings challenge the prevailing narrative in the literature. The results in Figure (8) suggest that automation leads to a shift in the skill composition of adopting firms away from highly skilled workers—both in terms of education and occupational groups—toward less-skilled workers. Particularly striking is the strong positive effect for low-educated workers, routine-manual workers, and blue-collar workers in small firms, which contrasts with the widely held view of automation as primarily skill-biased. Notably, estimates for industrial robots and additive manufacturing are relatively noisy and often not statistically significant. Moreover, it is worth highlighting that for automatic data processing and IIoT technologies, the estimates for highly educated workers in small firms turn negative, though they remain statistically insignificant.

5 Robustness

I conduct extensive robustness checks, sensitivity analysis and falsification tests. More specifically, I control for potential issues related to the diff-in-diff approach, i.e. potential violations of the parallel trend and no anticipation assumptions, as well as issues related to my automation proxy. The results are presented in Appendix B. Importantly, none of these exercises result in significant deviations from my key findings.

Sensitivity Analysis of Robustness to Violations of Parallel Trends. I use the approach by Rambachan and Roth (2023) to evaluate the robustness of my findings to potential violations of the parallel trends assumption. Their approach is based on the intuition that pre-trends are informative about violations of parallel trends. Thus, violations of parallel trends in the post-treatment period cannot be much bigger than those in the pre-treatment period. Accordingly, I can impose boundaries such that the post-treatment violation of parallel trends is no more than some constant ($\bar{M}$) larger than the maximum violation of parallel trends in the pre-treatment period. For example, setting the value of $\bar{M}$ to 1 assumes post-treatment violations are, at most, the same size as pre-treatment violations. Figure (12) presents sensitivity analyses for employment effects five years post-adoption. Results for the full sample (Panel A) remain robust even when allowing post-treatment violations up to half the maximum pre-treatment deviation. For small and large firm subsamples (Panels B and C), results hold even when permitting violations as large as the maximum pre-treatment deviation.

Placebo treatment timing. Moreover, one would want to be sure that the dynamic results I report are not the result of unrelated trends. To test this, I re-estimate the effect of an automation spike by considering a ‘placebo’ treatment. More precisely, I randomized the timing of the adoption of automation across adopters. If my results are

driven by spurious trends between adopting cohorts, these trends will now be picked up by my estimation as the treatment effect. Figure (13) shows zero effects of these placebo-treatment.

No anticipation. As pointed out above, a central concern for identification is the validity of the no-anticipation assumption. My identification strategy could be threatened if firms had anticipated their investment in automation in advance. In such cases, results might be biased, as firms could have already begun adjusting their labor force during the pre-treatment period. To address this issue, I follow the approach of Callaway and Sant’Anna (2021) by relaxing the assumption of no treatment anticipation behavior. Essentially, this involves extending the pre-treatment window further back in time and only including periods that are sufficiently distant from when the treatment actually begins. I conduct a sensitivity analysis that accounts for potential anticipation effects, covering periods ranging from one year (baseline) to four years prior to automation adoption. Figure (14) shows no significant evidence of anticipation effects in the years leading up to the automation event.

Single vs. Multiple Establishment Firms. Another potential concern involves the organization of work across different plants within a firm. The displacement effects of automation are expected to be more pronounced in the plants that directly implement automation, whereas non-automating plants within the same firm might experience indirect effects through productivity spillovers. Unfortunately, my analysis is constrained by the nature of my data, which is aggregated at the firm level (from customs and balance sheet data), preventing us from conducting a more granular plant-level investigation. However, I can focus on single-plant firms, where within-plant and within-firm effects are indistinguishable, providing a clearer view of the automation impact. Notably, single-plant firms constitute approximately 80% of my sample. The results presented in Figure (15) align mainly with my main findings. However, it is important to note that for small and multi-plant firms, the observed effects become statistically insignificant.

Controlling for hidden intermediation. Some potential concerns may arise regarding the use of my import-based proxy for automation. Specifically, my empirical design assumes that imported automation-related goods are utilized by firms for production purposes. However, it is possible that some firms may import automation products with the intent to resell them domestically or internationally, a practice known as Carry Along Trade (Bernard et al., 2019). This could potentially lead to a misclassification of adoption. To address the issue of re-exporting, I re-estimated my baseline model using two filters: (i) excluding all firms that re-export automation products within five years of the import spike, and (ii) adopting a more conservative approach by excluding all firms that

export automation products during the entire observation period. To further mitigate the risk of reselling within the domestic market, I (iii) excluded persistent importers — those firms that import every year — as this may indicate a business model focused on reselling, and (iv) restricted my sample to firms that import only once during the observation period, as it is unlikely these firms are engaged in systematic reselling activities. As illustrated in figure (16), my results are robust to these controls. Note, however, that Panel B shows that the re-estimates for small and large firms are (modestly) bigger than my baseline results, suggesting that hidden intermediation may introduce a bias toward zero.

Alternative Specification of Spikes. Additionally, two potential issues may arise regarding the definition of an import "spike." First, accurately identifying the timing of adoption may be challenging for firms that import automation products multiple times. However, my previous analysis, which restricts the sample to firms that import automation only once, suggests that this is not an issue in my case. Second, there is the possibility that technology-adoption events are not sufficiently large, potentially leading to issues of statistical power. To address these concerns, I conducted a sensitivity analysis using alternative specifications of automation spikes. Following a similar approach to Bessen et al. (2020, 2023), I defined a spike based on a relative size threshold rather than the highest rank, specifically defining a spike as occurring when imports are at least three times larger than the firm's average import value in other years. Additionally, I applied more stringent criteria to my baseline spike definition by imposing a minimum import value of 2,500 and 5,000 Euro.²² These alternative definitions not only address potential concerns but also enhance the comparability of my results with the existing literature. Also here, Figure (17) shows that my results are robust.

Event-time balanced panel: Another concern is that the panel of treated firms may be unbalanced in terms of event time. To address this, Figure (18) presents results based on a balanced sample, restricted to a 5-year event window. This ensures that the composition of the treatment and control groups remains consistent across different points in event time. The effects of an automation spike observed in this balanced sample closely mirror those found when using the full dataset. While my main specification utilizes all available data to maximize statistical power, the consistency of results from the balanced panel strengthens my confidence that the observed effects of automation on employment are not driven by shifts in firm composition over event time.

²²Note that Acemoglu et al. (2023) impose a minimum import value of 2,500 Euro in their baseline specification.

Reskilling instead of deskilling. The results presented in Figure (8) suggest that after an automation event, work tends to be reorganized around lower skill levels than before. However, an alternative explanation might be that this apparent shift toward lower skill levels is obscured by companies investing in employee training. In this case, while the educational qualifications or occupational categories of workers may not show a noticeable change, automation could still be associated with higher skill requirements. This hypothesis may be particularly relevant in the context of Portugal, where the workforce’s skill level is below the European average. To investigate this hypothesis, I estimate the impact of automation on training costs per worker. If the reskilling hypothesis holds, I would expect to see an increase in training expenses. However, Figure (19) demonstrates that there is no significant effect of automation on training costs following an automation event.

6 Conclusion

The fear of technological mass unemployment is widespread. Despite a rapidly growing body of empirical research on the adoption of automation technologies, its overall impact on employment and workforce composition remains a topic of intense debate. Recent findings are particularly puzzling: while some studies suggest a positive relationship between firm-level automation and employment (Acemoglu et al., 2020; Koch et al., 2021; Dixon et al., 2021), others report a negative impact (Bessen et al., 2020, 2023; Bonfiglioli et al., 2024). Furthermore, although skill-biased frameworks predict that automation should disproportionately displace low-educated production workers in routine-intensive jobs and favors highly skilled workers like STEM professionals, empirical evidence remains inconclusive.

This paper began by confirming the positive relationship between automation adoption and firm-level employment in Portugal. Distinguishing by firm size reveals a critical dual nature in automation—job creation in small firms and job destruction in large firms. Since small adopters are more numerous but large adopters employ the majority of workers, weighting by total hours worked suggests an overall negative net employment effect. Importantly, I show that this pattern persists across various firm types, industries, and types of automation technology, emphasizing the important role of firm size in shaping automation outcomes. This evidence for Portugal is consistent with the recent findings of Bisio et al. (2025) for Italy. Crucially, this dual nature in automation helps reconcile conflicting findings in the literature. Studies identifying a negative relationship (Bessen et al., 2020, 2023) often rely on weighted estimates, which may obscure positive effects among smaller firms. These findings underscore the urgent need to develop a better understanding of how policies and labor market institutions can mitigate automation’s adverse effects.

Most importantly, focusing on the impact of automation on skill composition, my findings suggest that automation reorganizes work around less-skilled workers. Specifically, I identify a positive effect on total hours worked for low-educated, routine-manual, and blue-collar workers. Differentiating by firm size, the analysis reveals that large firms tend to downsize across most skill groups, while small firms expand employment. Yet, both exhibit a relative shift toward lower-skilled labor over high-skilled workers. Particularly striking are the strong negative estimates for STEM professionals in large firms, challenging the widely accepted view that automation primarily benefits highly skilled workers. Instead, my findings contribute to an emerging debate on characterization of technological change in relation to labor. While traditional perspectives emphasize skill-biased technological change, some recent studies suggest a deskilling trend (Kunst, 2020; Downey, 2021; Danieli, 2022). My results align with this emerging perspective, indicating that automation’s impact on the workforce may be more complex than previously assumed.

My findings have several important implications. First, they challenge the optimistic narrative that automation-driven job creation is widespread across adopting firms. The evidence of job destruction in larger firms, combined with some support for the deskilling hypothesis, suggests that automation may exacerbate inequality. Second, discussions on automation’s diffusion and impact are often shaped by techno-deterministic perspectives. However, my findings highlight the critical role of socio-economic factors—such as firm size, the types of automation technology adopted and institutional or country-specific contexts—in shaping both the spread and consequences of automation. This aligns with qualitative studies documenting significant heterogeneity in automation adoption and outcomes across firms, sectors and countries (Krzywdzinski, 2017, 2021). Future research should further explore the varieties of automation. To determine whether my findings reflect patterns specific to middle-income countries or represent broader trends, rigorous comparative cross-country studies on automation, aligned with recent econometric advancements, are needed. Third, my findings have direct policy implications, particularly for industrial strategies such as Industry 4.0. For example, Portugal’s Industry 4.0 strategy adopts a bottom-up approach targeting small- and medium-sized enterprises (SMEs) (Yang and Gu, 2021). Policymakers must recognize the dual nature of automation’s impact on small and large firms; failure to do so could lead to unintended consequences—either fostering employment growth or exacerbating unemployment, depending on the context. Finally, the net employment effect of automation calls for labor market policies and institutions that mitigate its negative impacts, highlighting the need for research to better understand automation’s interaction with these institutions and inform effective policy responses.

References

- Abeliansky, Ana Lucia, Inmaculada Martínez-Zarzoso, and Klaus Prettnner (2020) “3D printing, international trade, and FDI,” *Economic Modelling*, 85, 288–306, <https://doi.org/10.1016/j.econmod.2019.10.014>.
- Acemoglu, Daron, Gary W Anderson, David N Beede et al. (2022) “Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey,” Working Paper 30659, National Bureau of Economic Research, [10.3386/w30659](https://doi.org/10.3386/w30659).
- Acemoglu, Daron and David Autor (2011) “Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings,” in Card, David and Orley Ashenfelter eds. *Handbook of Labor Economics*, 4, 1043–1171: Elsevier, [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Acemoglu, Daron, Hans R. A Koster, and Ceren Ozgen (2023) “Robots and Workers: Evidence from the Netherlands,” Working Paper 31009, National Bureau of Economic Research, <http://www.nber.org/papers/w31009>.
- Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo (2020) “Competing with Robots: Firm-Level Evidence from France,” *AEA Papers and Proceedings*, 110, 383–88, [10.1257/pandp.20201003](https://doi.org/10.1257/pandp.20201003).
- Acemoglu, Daron and Pascual Restrepo (2019) “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33 (2), 3–30, [10.1257/jep.33.2.3](https://doi.org/10.1257/jep.33.2.3).
- (2020a) “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128 (6), 2188 – 2244, <https://doi.org/10.1086/705716>.
- (2020b) “Unpacking Skill Bias: Automation and New Tasks,” *AEA Papers and Proceedings*, 110, 356–61, [10.1257/pandp.20201063](https://doi.org/10.1257/pandp.20201063).
- (2021) “Demographics and Automation,” *The Review of Economic Studies*, 89 (1), 1–44, [10.1093/restud/rdab031](https://doi.org/10.1093/restud/rdab031).
- (2022) “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 90 (5), 1973–2016, <https://doi.org/10.3982/ECTA19815>.
- Adachi, Daisuke, Daiji Kawaguchi, and Yukiko U. Saito (2024) “Robots and Employment: Evidence from Japan, 1978–2017,” *Journal of Labor Economics*, 42 (2), 591–634, [10.1086/723205](https://doi.org/10.1086/723205).
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel (2022a) “The effects of automation on labor demand,” *Robots and AI: A new economic era*, 15–39.

- Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel (2023) “The Local Labor Market Effects of Modern Manufacturing Capital: Evidence from France,” *AEA Papers and Proceedings*, 113, 219–23, [10.1257/pandp.20231039](https://doi.org/10.1257/pandp.20231039).
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel (2022b) “Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France,” SciencePo Working papers Main hal-03943312, HAL, <https://ideas.repec.org/p/hal/spmain/hal-03943312.html>.
- Ahn, JaeBin, Amit K. Khandelwal, and Shang-Jin Wei (2011) “The role of intermediaries in facilitating trade,” *Journal of International Economics*, 84 (1), 73–85, <https://doi.org/10.1016/j.jinteco.2010.12.003>.
- Andreoni, Antonio, Guendalina Anzolin, Mateus Labrunje, and Danilo Spinola (2023) “Unveiling Structure and Dynamics of Global Digital Production Technology Networks: A new digital technology classification and network analysis based on trade data,” Working Papers 261, Department of Economics, SOAS University of London, UK, <https://EconPapers.repec.org/RePEc:soa:wpaper:261>.
- Arntz, Melanie, Sebastian Blesse, and Philipp Doerrenberg (2022) “The end of work is near, isn’t it? Survey evidence on automation angst,” ZEW Discussion Papers 22-036, ZEW - Leibniz Centre for European Economic Research, <https://dx.doi.org/10.2139/ssrn.4247981>.
- Asphjell, Magne K., Wilko Letterie, Øivind A. Nilsen, and Gerard A. Pfann (2014) “Sequentiality Versus Simultaneity: Interrelated Factor Demand,” *The Review of Economics and Statistics*, 96 (5), 986–998, [10.1162/REST_a_00411](https://doi.org/10.1162/REST_a_00411).
- Atzori, Luigi, Antonio Iera, and Giacomo Morabito (2010) “The internet of things: A survey,” *Computer networks*, 54 (15), 2787–2805.
- Autor, David H. (2015) “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29 (3), 3–30, [10.1257/jep.29.3.3](https://doi.org/10.1257/jep.29.3.3).
- Autor, David H., Frank Levy, and Richard J. Murnane (2003) “The Skill Content of Recent Technological Change: An Empirical Exploration*,” *The Quarterly Journal of Economics*, 118 (4), 1279–1333, [10.1162/003355503322552801](https://doi.org/10.1162/003355503322552801).
- Barth, Erling, Marianne Roed, Pal Schone, and Janis Umblijs (2020) “How Robots Change Within-Firm Wage Inequality,” IZA Discussion Papers 13605, Institute of Labor Economics (IZA), <https://ideas.repec.org/p/iza/izadps/dp13605.html>.

- Bernard, Andrew B., Marco Grazzi, and Chiara Tomasi (2015) “Intermediaries in International Trade: Products and Destinations,” *The Review of Economics and Statistics*, 97 (4), 916–920, [10.1162/REST_a_00495](https://doi.org/10.1162/REST_a_00495).
- Bernard, Andrew, Emily J Blanchard, Ilke Van Beveren, and Hylke Vandenbussche (2019) “Carry-Along Trade,” *The Review of Economic Studies*, 86 (2), 526–563, [10.1093/restud/rdy006](https://doi.org/10.1093/restud/rdy006).
- Bernard, Andrew, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2010) “Wholesalers and Retailers in US Trade,” *American Economic Review*, 100 (2), 408–13, [10.1257/aer.100.2.408](https://doi.org/10.1257/aer.100.2.408).
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge (2020) “Firm-Level Automation: Evidence from the Netherlands,” *AEA Papers and Proceedings*, 110, 389–93, [10.1257/pandp.20201004](https://doi.org/10.1257/pandp.20201004).
- (2023) “What Happens to Workers at Firms that Automate?,” *The Review of Economics and Statistics*, 1–45, [10.1162/rest_a_01284](https://doi.org/10.1162/rest_a_01284).
- Birch, D.L., Massachusetts Institute of Technology. Program on Neighborhood, and Regional Change (1979) *The Job Generation Process*, M.I.T. Program on Neighborhood and Regional Change: M.I.T. Program on Neighborhood and Regional Change, <https://books.google.it/books?id=fsQTAQAAMAAJ>.
- Bisio, Laura, Angelo Cuzzola, Marco Grazzi, and Daniele Moschella (2025) “The dynamics of automation adoption: Firm-level heterogeneity and aggregate employment effects,” *European Economic Review*, 104943, <https://doi.org/10.1016/j.euroecorev.2024.104943>.
- Blanchard, Olivier and Pedro Portugal (2001) “What Hides Behind an Unemployment Rate: Comparing Portuguese and U.S. Labor Markets,” *American Economic Review*, 91 (1), 187–207, [10.1257/aer.91.1.187](https://doi.org/10.1257/aer.91.1.187).
- Bloom, Nicholas, Mirko Draca, and John Van Reenen (2015) “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity,” *The Review of Economic Studies*, 83 (1), 87–117, [10.1093/restud/rdv039](https://doi.org/10.1093/restud/rdv039).
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen (2013) “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 81 (4), 1347–1393, <https://doi.org/10.3982/ECTA9466>.
- Blum, Bernardo S., Sebastian Claro, and Ignatius Horstmann (2010) “Facts and Figures on Intermediated Trade,” *American Economic Review*, 100 (2), 419–23, [10.1257/aer.100.2.419](https://doi.org/10.1257/aer.100.2.419).

- Bogliacino, Francesco and Mario Pianta (2010) “Innovation and Employment: a Reinvestigation using Revised Pavitt classes,” *Research Policy*, 39 (6), 799–809, <https://doi.org/10.1016/j.respol.2010.02.017>.
- Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia (2024) “Robot Imports and Firm-Level Outcomes,” *The Economic Journal*, ueae055, [10.1093/ej/ueae055](https://doi.org/10.1093/ej/ueae055).
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2022) “Revisiting Event Study Designs: Robust and Efficient Estimation,” papers, arXiv.org, <https://arxiv.org/abs/2108.12419>.
- Bowen, Howard Rothmann and Garth L Mangum (1966) *Automation and economic progress*, 147: Prentice-Hall.
- Braverman, Harry (1974) *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*: NYU Press, <http://www.jstor.org/stable/j.ctt9qfrkf>.
- Brynjolfsson, E. and A. McAfee (2014) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*: W. W. Norton, <https://books.google.it/books?id=PMBUAgAAQBAJ>.
- Cabral, Luís M B and José Mata (2003) “On the Evolution of the Firm Size Distribution: Facts and Theory,” *American Economic Review*, 93 (4), 1075–1090, [10.1257/000282803769206205](https://doi.org/10.1257/000282803769206205).
- Caliendo, Lorenzo, Giordano Mion, Luca David Opromolla, and Esteban Rossi-Hansberg (2020) “Productivity and Organization in Portuguese Firms,” *Journal of Political Economy*, 128 (11), 4211–4257, [10.1086/710533](https://doi.org/10.1086/710533).
- Callaway, Brantly and Pedro H.C. Sant’Anna (2021) “Difference-in-Differences with multiple time periods,” *Journal of Econometrics*, 225 (2), 200–230, <https://doi.org/10.1016/j.jeconom.2020.12.001>, Themed Issue: Treatment Effect 1.
- Calvino, Flavio, Chiara Criscuolo, Luca Marcolin, and Mariagrazia Squicciarini (2018) “A taxonomy of digital intensive sectors,” Technical report, OECD, <https://doi.org/https://doi.org/10.1787/f404736a-en>.
- Calvino, Flavio and Maria Enrica Virgillito (2018) “The Innovation-Employment Nexus: A Critical Survey if Theory and Empirics,” *Journal of Economic Surveys*, 32 (1), 83–117, <https://doi.org/10.1111/joes.12190>.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018) “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 36 (S1), S13–S70, [10.1086/694153](https://doi.org/10.1086/694153).

- Card, David, Ana Rute Cardoso, and Patrick Kline (2015) “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women*,” *The Quarterly Journal of Economics*, 131 (2), 633–686, [10.1093/qje/qjv038](https://doi.org/10.1093/qje/qjv038).
- Caselli, Francesco and Wilbur John Coleman (2001) “Cross-Country Technology Diffusion: The Case of Computers,” *The American Economic Review*, 91 (2), 328–335, <http://www.jstor.org/stable/2677783>.
- Caselli, Francesco and Daniel J Wilson (2004) “Importing technology,” *Journal of Monetary Economics*, 51 (1), 1–32, <https://doi.org/10.1016/j.jmoneco.2003.07.004>.
- Castellani, Davide, Fabio Lamperti, and Katiuscia Lavoratori (2022) “Measuring adoption of industry 4.0 technologies via international trade data: insights from European countries,” *Economia e Politica Industriale: Journal of Industrial and Business Economics*, 49 (1), 51–93, [10.1007/s40812-021-00204-](https://doi.org/10.1007/s40812-021-00204-).
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) “The Effect of Minimum Wages on Low-Wage Jobs*,” *The Quarterly Journal of Economics*, 134 (3), 1405–1454, [10.1093/qje/qjz014](https://doi.org/10.1093/qje/qjz014).
- de Chaisemartin, Clément and Xavier D’Haultfoeulle (2020a) “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110 (9), 2964–96, [10.1257/aer.20181169](https://doi.org/10.1257/aer.20181169).
- (2020b) “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110 (9), 2964–96, [10.1257/aer.20181169](https://doi.org/10.1257/aer.20181169).
- Corrocher, Nicoletta, Daniele Moschella, Jacopo Staccioli, and Marco Vivarelli (2023) “Innovation and the labor market: theory, evidence, and challenges,” *Industrial and Corporate Change*, 33 (3), 519–540, [10.1093/icc/dtad066](https://doi.org/10.1093/icc/dtad066).
- Cortes, Guido Matias, Ana Oliveira, and Anna Salomons (2021) “Do technological advances reduce the gender wage gap?,” *Oxford Review of Economic Policy*, 36 (4), 903–924, [10.1093/oxrep/graa051](https://doi.org/10.1093/oxrep/graa051).
- Cséfalvay, Zoltán and Petros Gkotsis (2022) “Robotisation race in Europe: the robotisation chain approach,” *Economics of Innovation and New Technology*, 31 (8), 693–710, [10.1080/10438599.2020.1849968](https://doi.org/10.1080/10438599.2020.1849968).
- Danieli, Oren (2022) “Revisiting U.S. Wage Inequality at the Bottom 50%,” <https://api.semanticscholar.org/CorpusID:259157746>.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner (2021) “The Adjustment of Labor Markets to Robots,” *Journal of the European Economic Association*, 19 (6), 3104–3153, <https://doi.org/10.1093/jeea/jvab012>.

- Delmar, Frédéric, Per Davidsson, and William B Gartner (2003) “Arriving at the high-growth firm,” *Journal of Business Venturing*, 18 (2), 189–216, [https://doi.org/10.1016/S0883-9026\(02\)00080-0](https://doi.org/10.1016/S0883-9026(02)00080-0).
- Deng, Liuchun, Verena Plümpe, and Jens Stegmaier (2021) “Robot adoption at German plants,” IWH Discussion Papers 19/2020, Halle Institute for Economic Research (IWH), <https://ideas.repec.org/p/zbw/iwhdps/192020.html>.
- Dixon, Jay, Bryan Hong, and Lynn Wu (2021) “The Robot Revolution: Managerial and Employment Consequences for Firms,” *Management Science*, 69 (9), <https://doi.org/10.1287/mnsc.2020.3812>.
- Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich (2021) “Threats and opportunities in the digital era: Automation spikes and employment dynamics,” *Research Policy*, 50 (7), 104137, <https://doi.org/10.1016/j.respol.2020.104137>.
- (2022) “For whom the bell tolls: The firm-level effects of automation on wage and gender inequality,” *Research Policy*, 51 (7), 104533, <https://doi.org/10.1016/j.respol.2022.104533>.
- 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 (4), 104199, <https://doi.org/10.1016/j.respol.2021.104199>.
- Dosi, Giovanni and Richard R. Nelson (2010) “Chapter 3 - Technical Change and Industrial Dynamics as Evolutionary Processes,” in Hall, Bronwyn H. and Nathan Rosenberg eds. *Handbook of The Economics of Innovation*, Vol. 1, 1 of Handbook of the Economics of Innovation, 51–127: North-Holland, [https://doi.org/10.1016/S0169-7218\(10\)01003-8](https://doi.org/10.1016/S0169-7218(10)01003-8).
- Dottori, Davide (2021) “Robots and employment: evidence from Italy,” *Economia Politica*, 38, [10.1007/s40888-021-00223-x](https://doi.org/10.1007/s40888-021-00223-x).
- Downey, Mitch (2021) “Partial automation and the technology-enabled deskilling of routine jobs,” *Labour Economics*, 69, 101973, <https://doi.org/10.1016/j.labeco.2021.101973>.
- Eaton, Jonathan and Samuel Kortum (2001) “Trade in capital goods,” *European Economic Review*, 45 (7), 1195–1235, <https://www.sciencedirect.com/science/article/pii/S0014292100001033>.
- Eurobarometer (2017) “Attitudes Towards the Impact of Digitisation and Automation on Daily Life,” Technical report, European Commission, <https://europa.eu/eurobarometer/surveys/detail/2160>.

- Ford, M. (2015) *Rise of the Robots: Technology and the Threat of a Jobless Future*: Basic Books.
- Frey, Carl Benedikt and Michael A. Osborne (2017) “The future of employment: How susceptible are jobs to computerisation?” *Technological Forecasting and Social Change*, 114, 254–280, <https://doi.org/10.1016/j.techfore.2016.08.019>.
- Gentili, Andrea, Fabiano Compagnucci, Mauro Gallegati, and Enzo Valentini (2020) “Are machines stealing our jobs?,” *Cambridge Journal of Regions, Economy and Society*, 13 (1), 153–173, [10.1093/cjres/rsz025](https://doi.org/10.1093/cjres/rsz025).
- Goodman-Bacon, Andrew (2021) “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 225 (2), 254–277, <https://doi.org/10.1016/j.jeconom.2021.03.014>, Themed Issue: Treatment Effect 1.
- Graetz, Georg and Guy Michaels (2018) “Robots at Work,” *The Review of Economics and Statistics*, 100 (5), 753–768, https://doi.org/10.1162/rest_a_00754.
- Grazzi, Marco, Nadia Jacoby, and Tania Treibich (2016) “Dynamics of investment and firm performance: comparative evidence from manufacturing industries,” *Empirical Economics*, 51 (1), 125–179, [10.1007/s00181-015-0991-2](https://doi.org/10.1007/s00181-015-0991-2).
- Gubbi, Jayavardhana, Rajkumar Buyya, Slaven Marusic, and Marimuthu Palaniswami (2013) “Internet of Things (IoT): A vision, architectural elements, and future directions,” *Future Generation Computer Systems*, 29 (7), 1645–1660, <https://doi.org/10.1016/j.future.2013.01.010>, Including Special sections: Cyber-enabled Distributed Computing for Ubiquitous Cloud and Network Services Cloud Computing and Scientific Applications — Big Data, Scalable Analytics, and Beyond.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013) “Who Creates Jobs? Small versus Large versus Young,” *The Review of Economics and Statistics*, 95 (2), 347–361, [10.1162/REST_a_00288](https://doi.org/10.1162/REST_a_00288).
- Humlum, Anders (2019) “Robot Adoption and Labor Market Dynamics,” in *Working paper*.
- International Organization for Standardization (2015) “Additive manufacturing—General principles—Terminology,” <https://www.iso.org/obp/ui/#iso:std:iso-astm:52900:ed-1:v1:en>, Accessed: YYYY-MM-DD.
- Juchniewicz, Małgorzata and Magdalena Łada (2020) “Diversification of the competitive position of the advanced technology sector in EU,” *International Journal of Management and Economics*, 56 (1), 31–41, [doi:10.2478/ijme-2020-0006](https://doi.org/10.2478/ijme-2020-0006).

- Keynes, John Maynard (1930) “17Economic Possibilities for our Grandchildren (1930),” in *Revisiting Keynes: Economic Possibilities for Our Grandchildren*: The MIT Press, [10.7551/mitpress/9780262162494.003.0002](https://doi.org/10.7551/mitpress/9780262162494.003.0002).
- Koch, Michael, Ilya Manuylov, and Marcel Smolka (2021) “Robots and Firms,” *The Economic Journal*, 131 (638), 2553–2584, [10.1093/ej/ueab009](https://doi.org/10.1093/ej/ueab009).
- Krzywdzinski, Martin (2017) “Automation, skill requirements and labour-use strategies: high-wage and low-wage approaches to high-tech manufacturing in the automotive industry,” *New Technology, Work and Employment*, 32 (3), 247–267, <https://doi.org/10.1111/ntwe.12100>.
- (2021) “Automation, digitalization, and changes in occupational structures in the automobile industry in Germany, Japan, and the United States: a brief history from the early 1990s until 2018,” *Industrial and Corporate Change*, 30 (3), 499–535, [10.1093/icc/dtab019](https://doi.org/10.1093/icc/dtab019).
- Kunst, David (2020) “Deskilling among Manufacturing Production Workers,” Tinbergen Institute Discussion Papers 19-050/VI, Tinbergen Institute, <https://EconPapers.repec.org/RePEc:tin:wpaper:20190050>.
- Leone, Fabrizio (2023) “Multinationals, robots and the labor share,” cep discussion papers, Centre for Economic Performance, LSE, <https://EconPapers.repec.org/RePEc:cep:cepdps:dp1900>.
- Leontief, Wassily W. (1952) “Machines and Man,” *Scientific American*, 187, 150–160, <https://api.semanticscholar.org/CorpusID:121768118>.
- Letterie, Wilko A., Gerard A. Pfann, and J.Michael Polder (2004) “Factor adjustment spikes and interrelation: an empirical investigation,” *Economics Letters*, 85 (2), 145–150, <https://doi.org/10.1016/j.econlet.2004.03.017>.
- Meer, Jonathan and Jeremy West (2016) “Effects of the Minimum Wage on Employment Dynamics,” *The Journal of Human Resources*, 51 (2), 500–522, <http://www.jstor.org/stable/24736030>.
- Montobbio, Fabio, Jacopo Staccoli, Maria Virgillito, and Marco Vivarelli (2022) “The Empirics of Technology, Employment and Occupations: Lessons Learned and Challenges Ahead,” *SSRN Electronic Journal*, [10.2139/ssrn.4281286](https://ssrn.com/abstract=4281286).
- Nedelkoska, Ljubica and Glenda Quintini (2018) “Automation, skills use and training,” Technical Report 202, OECD, <https://doi.org/https://doi.org/10.1787/2e2f4eea-en>.

- Nelson, Richard R. and Sidney G. Winter (1982) “The Schumpeterian Tradeoff Revisited,” *American Economic Review*, 72 (1), 114–132.
- Neumark, David, Brandon Wall, and Junfu Zhang (2011) “Do Small Businesses Create More Jobs? New Evidence For The United States From The National Establishment Time Series,” *The Review of Economics and Statistics*, 93 (1), 16–29, <http://www.jstor.org/stable/23015917>.
- Nof, Shimon Y. (2009) “Automation: What It Means to Us Around the World,” in Nof, Shimon Y. ed. *Springer Handbook of Automation*, 13–52, Berlin, Heidelberg: Springer Berlin Heidelberg, [10.1007/978-3-540-78831-7_3](https://doi.org/10.1007/978-3-540-78831-7_3).
- Rambachan, Ashesh and Jonathan Roth (2023) “A More Credible Approach to Parallel Trends,” *The Review of Economic Studies*, 90 (5), 2555–2591, [10.1093/restud/rdad018](https://doi.org/10.1093/restud/rdad018).
- Reljic, Jelena, Valeria Cirillo, and Dario Guarascio (2023) “Regimes of robotization in Europe,” *Economics Letters*, 232, 111320, <https://doi.org/10.1016/j.econlet.2023.111320>.
- Restrepo, Pascual (2023) “Automation: Theory, Evidence, and Outlook,” Working Paper 31910, National Bureau of Economic Research, [10.3386/w31910](https://doi.org/10.3386/w31910).
- Rubin, Donald B. (2007) “The design versus the analysis of observational studies for causal effects: parallels with the design of randomized trials,” *Statistics in Medicine*, 26 (1), 20–36, <https://doi.org/10.1002/sim.2739>.
- (2008) “For objective causal inference, design trumps analysis,” *The Annals of Applied Statistics*, 2 (3), 808 – 840, [10.1214/08-AOAS187](https://doi.org/10.1214/08-AOAS187).
- Sant’Anna, Pedro H.C. and Jun Zhao (2020) “Doubly robust difference-in-differences estimators,” *Journal of Econometrics*, 219 (1), 101–122, <https://doi.org/10.1016/j.jeconom.2020.06.003>.
- Sun, Liyang and Sarah Abraham (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199, <https://doi.org/10.1016/j.jeconom.2020.09.006>, Themed Issue: Treatment Effect 1.
- Tuhkuri, Joonas, Aapo Stenhammar, and Johannes Hirvonen (2022) “New Evidence on the Effect of Technology on Employment and Skill Demand,” *Quarterly Journal of Economics*.
- Wang, Shiyong, Jiafu Wan, Di Li, and Chunhua Zhang (2016) “Implementing Smart Factory of Industrie 4.0: An Outlook,” *International Journal of Distributed Sensor Networks*, 12 (1), 3159805, [10.1155/2016/3159805](https://doi.org/10.1155/2016/3159805).

- Wing, Coady, Seth M. Freedman, and Alex Hollingsworth (2024) “Stacked Difference-in-Differences,” NBER Working Papers 32054, National Bureau of Economic Research, Inc.
- Yang, Fengwei and Sai Gu (2021) “Industry 4.0, a revolution that requires technology and national strategies,” *Complex & Intelligent Systems*, 7 (3), 1311–1325, [10.1007/s40747-020-00267-9](https://doi.org/10.1007/s40747-020-00267-9).
- Zeira, Joseph (1998) “Workers, Machines, and Economic Growth*,” *The Quarterly Journal of Economics*, 113 (4), 1091–1117, [10.1162/0033553985555847](https://doi.org/10.1162/0033553985555847).

A Online-Appendix

A.1 Data Processing and Variable Construction

I start with the matched employer-employee data set (QP), which includes 5,570,419 firm and 50,960,574 worker observations for the period 2004-2021. I first merge firm and worker data for each year and I then apply for each year the following cleaning procedure: (i) I drop a minority of worker observations with an invalid social security number, i.e., if the worker id has less than 6 digits or more than 10 digits. (ii) For workers with multiple jobs within the same year, I keep his or her highest paying job with the highest hours worked (since most likely, this is his or her primary job).²³ I keep (iii) dependent workers (iv) between 18 and 65 (working-age population). I drop worker-year pairs (v) without complete basic remuneration, (vi) whenever the sum of weekly normal and overtime hours is below 25 and above 80 and (vii) their regular earnings are less than 80 percent of the minimum wage.²⁴ (viii) Based on the resulting sample, I trim worker-year pairs whose monthly total wage is outside the top 0.5 percentile. Gross monthly earnings from dependent work are obtained by the sum of the five types of pay available in the data (base wages, tenure-related payments, overtime pay, subsidies). I construct hourly total wages by dividing individuals' total monthly remuneration by total monthly hours worked. The data is then appended to construct my linked employer-employee panel data. After this procedure I have 4,187,954 firm-year observations and 35,889,005 worker-year observations. Finally, as the Portuguese classification of occupations (corresponding to the European ISCO classification) has been revised in 2010, I establish a data-driven one-to-one mapping. To do this, I use the following criterion: if the majority of workers in occupation A in 2010 (CPP/2010) have occupation B in 2009 (CNP/1994), then I map occupation A into occupation B. I then turn to the balance sheet data (SCIE) and recover information on firm's 5-digit industry classification (CAE Rev.3, based on NACE-Rev. 2 Statistical classification of economic activities in the European Community), gross value added at market prices, total sales, wage bill, investment in equipment and machines and expenses for worker training. I discard observations that have unreasonable values, i.e. non-positive turnover or gross value added.

The customs data records individual trade transactions on a monthly basis. I collapse these transactions to the firm, product, year level and retrieve information on total yearly imports and exports of automation-related capital goods as described in section 2 and table 3. I then compute total imports and exports aggregating the data at the firm-year level.

Finally, I merge the the matched employer-employee data QP with the balance sheet

²³Note that others (e.g. Caliendo et al., 2020) drop workers with multiple jobs in the same year.

²⁴Information on minimum wage regulation stems from <https://www.dgert.gov.pt/evolucao-da-remuneracao-minima-mensal-garantida-rmmg>.

data SCIE using the unique, time-invariant identifier based on the firm’s social security number. I keep only firm observations that appear in both data sources. Then I add firm-year trade information from the CI dataset. I assign each firm a permanent 2-digit sector based on the most frequent occurrence, since firms industry code may vary across years, e.g. due to misreporting. I restrict the sample to firms located in mainland Portugal (omitting those in the Azores and Madeira islands). To avoid potential misclassification of firms as adopters, I exclude retailers and intermediaries from my analysis by dropping firms with the respective 5-digit industry codes as detailed in Table (4). Given that my proxy for automation adoption is based on imported product categories, I primarily focus on firms that have imported at least once during the period from 2004 to 2021 (importer sample). This results in a final sample size of 578,764 firm-year observations. All monetary values in the paper are quoted in year 2012 Euros, deflated using the Consumer Price Index from INE.

Identification of automation-related capital goods. Table 3 presents the product codes that I use to identify automation-related imports following the previous work by Acemoglu and Restrepo (2021), Abeliatsky et al. (2020), Domini et al. (2021, 2022), Castellani et al. (2022) and Andreoni et al. (2023).

Table (3) HS product codes referring to automation-related capital goods

Label	HS17 codes (6-digits)		Used in extant literature
01 <i>Industrial robots</i>	847950.		Acemoglu et al. (2020); Dixon et al. (2021); Domini et al. (2021, 2022); Acemoglu and Restrepo (2021); Castellani et al. (2022); Bisio et al. (2025).
02 <i>Conventional Automation technologies</i>			
2.1 Dedicated Machinery	847989		
2.2 Numerically controlled machines (NCM)	845811,	845891,	
	845921,	845931,	
	845941 ¹ ,	845951,	
	845961,	846012 ² ,	
	846022,	846023,	
	846024 ³ ,	846031,	
	846221,	846231,	
	846241.		
2.2 Automatic machine tools (excl. NCM)	845600-	846699,	Acemoglu and Restrepo (2021); Domini et al. (2021, 2022); Bisio et al. (2025).
	846820-846899, 851511-851519		
2.3 Automatic welding machines	851521,	851531,	
	851580, 851590		
2.4 Weaving and knitting machines	844610-844630, 844711-844790.		
2.5 Other textile dedicated machinery	844400-844590		
2.6 Automatic conveyors	842831,	842832,	
	842833, 842839.		
03 <i>Additive Manufacturing (3D-Printing)</i>	847720,	847759,	Abeliansky et al. (2020); Castellani et al. (2022).
	847780.		
04 <i>Automatic Data processing machines</i>			
3.1 Automatic data processing machines	847130,	847141,	Domini et al. (2022); Bisio et al. (2025).
	847149,	847150,	
	847321, 847330.		
3.2 Electronic calculating machines	847010,	847021,	
	847029.		
05 <i>Industrial Internet of Things</i>			
5.1 Network communication equipment	847180, 847190		
5.2 Non-wireles communication equipment	851762 ⁴ .		Castellani et al. (2022)
5.3 Radio navigational aid and remote control apparatus	852691, 852692.		
5.4 Automatic regulating and control instruments	903210,	903220,	Acemoglu and Restrepo (2021); Domini et al. (2021, 2022); Bisio et al. (2025); Castellani et al. (2022); Andreoni et al.
	903281,	903289,	
	903290.		
	42		

Role of Intermediaries in Automation. As discussed in section 2.3, I exclude intermediaries in automation technologies to reduce potential noise in the import-based automation proxy. To identify these sectors, I follow Bisio et al. (2025), who began with wholesalers classified under NACE Rev. 2 group codes 461–469 and then identified the relevant 5-digit sectors, as detailed in Table (4).

Table (4) Industry codes of retailers/intermediaries of automation goods

5-digit codes	Description
46140	Agents involved in the sale of machinery and industrial equipment
46190	Agents involved in the sale of various goods
46500	Wholesale of information and communication equipment
46620	Wholesale of machine tools
46640	Wholesale of machinery for the textile industry
46699	Wholesale of other machinery
46900	Non-specialized wholesale trade

Table (5) Role of Intermediaries in Automation

Year	Share of firms	Share of import value	Net imports
2004	11.8	7.8	86.9
2005	11.8	8.3	86.2
2006	11.7	8.2	76.9
2007	12.5	9.4	77.2
2008	11.9	7.0	78.1
2009	11.0	4.4	78.0
2010	10.4	6.7	77.9
2011	10.5	7.6	83.5
2012	10.6	7.7	79.0
2013	10.6	9.4	83.3
2014	10.3	9.2	88.7
2015	10.1	8.9	90.6
2016	10.2	8.5	84.1
2017	9.7	9.4	85.7
2018	10.0	9.8	87.9
2019	9.9	9.9	87.1
2020	9.6	9.7	88.7
2021	8.2	9.6	91.3
Total	8.2	8.5	79.4

Notes: Table (5) presents the yearly share of resellers among firms that import automation technologies, as well as the corresponding import value share of these technologies. This table also includes the net import share for resellers, which is calculated as the difference between automation imports and exports, divided by total imports. The final row shows the share of firms that act as resellers among those that have imported automation technologies at least once. For more details on the role of intermediaries, refer to Table (4).

Measuring skill groups. To distinguish workers by skill groups I rely on three commonly used measures: First, I distinguish workers by their educational level defined as low, middle and high-educated workers. Second, I classify occupation codes into broad task groups, i.e. non-routine cognitive, routine cognitive, routine manual and non-routine manual, following the work by Cortes et al. (2021). Third, I define five macro-occupation groups: managers, STEM professionals, other professionals, blue-collar workers and other non-professionals, similarly as in Barth et al. (2020). For the description of the codification of the occupation codes see Table (6).


Table (6) Classifications of Occupational codes


Broad Task Group (Cortes et al., 2021)	Occupational classification coding, CNP94 (3-digits)
Non-routine cognitive	112-131, 211-247,311-315,321-515
Routine cognitive	521-523
Routine manual	711-745,811-834,916-933
Non-routine manual	516,911-915
Macro-Occupation Groups	Occupational classification coding, CNP94 (2-digits)
Managers	11-13
STEM professionals	21,31
Other professionals	22-24,32-34
Blue-collar workers	71-83
Other non-professionals	41-42,51-52,61-62,91-93

Notes: i) I map detailed occupation codes to broad task groups following previous work by Cortes et al. (2021). ii) I classify occupation codes into five macro groups similar to Barth et al. (2020).

A.2 Appendix Tables and Figures


Figure (9) Event study - Impact on Skill Groups by Technologies





Notes: The figure provides overall ATTs based on event study estimates of dynamic effects of adopting automation on total hours worked by skill groups using the approach of Callaway and Sant’Anna (2021) specified in equation (2). The figure distinguishes between workers according to three key dimensions: (i) educational attainment, (ii) routine-intensity of their tasks and (iii) macro-occupational group. For more detailed information refer to table (6). I distinguish between various groups of technologies: Conventional Automation (CA) in Panel A and B, Additive Manufacturing (AM) in Panel C and D, Industrial Robots (IR) in Panel E and F, Automatic Data Processing Machines (ADP) in Panel G and H, and the Industrial Internet of Things (IIoT) in Panel I and J. For more detailed information refer to table (3). The spike is defined in relation to the largest event of imports for the respective technology type. The comparison group consists of firms that have not yet adopted the respective technology. I control for the log of labor productivity, log of turnover, log of hourly wages, firm size class and export status at the firm’s baseline ($\tau = -1$). The coefficient for the year prior to adoption is normalized to zero. Panels on the left hand side (A,C,E,G,I) include all firms that eventually adopt automation within the 2004-2021 period, while Panels on the right hand side (B,D,F,H,J) further disaggregates the sample by firm size —distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided at both the 5% and 10% significance levels.

Figure (11) Validation of treatment definition



Notes: Figure 11 illustrates, in Panel A, the average yearly share of automation imports as a percentage of total automation imports over the period 2004–2021. Panel B presents the average yearly share of investments in fixed assets, based on balance sheet data, relative to total investment over the same period, serving as a proxy for investments in machinery. Both panels are centered around the automation spike, defined as the largest automation import episode for each firm.


Table (7) Size Class Composition

Variable	Share of firms	
	Small	Large
Digital-Intensity - High	0.78	0.22
Digital-Intensity - Medium-high	0.86	0.14
Digital-Intensity - Low	0.65	0.35
Digital Intensity - Medium-low	0.69	0.31
Services	0.83	0.17
Production	0.69	0.31
Domestic firms	0.81	0.19
Multinational firms	0.56	0.44
Non-exporter	0.88	0.12
Exporter	0.68	0.32
Young firms	0.86	0.14
Old firms	0.65	0.35

Notes: Table (7) presents the composition of firm size classes in the baseline year, relative to the subsamples used across different firm types and industry categories in Figure (6).

B Robustness Checks - Figures


Figure (12) Sensitivity Analysis of Robustness to Violations of Parallel Trends




Notes: The figure presents a sensitivity analysis using the approach from Rambachan and Roth (2023) for the event study that uses the approach from Callaway and Sant'Anna (2021) depicted in Figure (5) when $e = 5$. The red line represents the 95% confidence interval based on the original estimates. The black lines show the 95% confidence sets, allowing for violations of the parallel trends assumption up to $\bar{M}$ times as large as the maximum observed in pre-treatment periods. Panels A, B, and C illustrate the results of this sensitivity analysis for all firms, small firms, and large firms, respectively.

Figure (13) Placebo-Test

Panel A. Total hours worked (ln)




Panel B. Total hours worked (ln), by firm size




Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), conducting a placebo test. In this test, I randomize the timing of automation adoption among the adopters. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

Figure (14) Controlling for adopters' anticipation



Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), this time relaxing the assumption of no treatment anticipation behavior. To account for potential anticipation effects, I extend the pre-treatment window further back in time, ranging from one year (baseline) to four years prior to the adoption, and include only those periods that are sufficiently distant from the actual onset of treatment. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.


Figure (15) Single vs. Multi-establishments




Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), distinguishing between single and multi-establishment firms w.r.t. their baseline year. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

Figure (16) Controlling for hidden intermediation

Panel A. Total hours worked (ln)




Panel B. Total hours worked (ln), by firm size




Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), with adjustments to account for hidden intermediation. The following filters are applied: (i) exclusion of all firms that re-export within five years of the import spike, (ii) exclusion of all firms that eventually export, (iii) exclusion of persistent importers, and (iv) inclusion of only those firms that import automation products a single time. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

Figure (17) Controlling for alternative spike definitions



Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), incorporating alternative specifications of automation spikes. Specifically, I define a spike as occurring when (i) imports reach a minimum value of at least 2,500 Euros, (ii) imports exceed 5,000 Euros or (iii) imports are at least three times larger than the firm's average import value in other years. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.


Figure (18) Event-time balanced panel




Notes: The figure presents re-estimated event study results from Figure (5) using the approach of Callaway and Sant'Anna (2021), applied to a balanced sample restricted to a 5-year event window. The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.

Figure (19) Effect on training costs per worker

Panel A. Training costs per worker (ihs)



Panel B. Training costs per worker (ihs), by firm size



Notes: The figure provides event study estimates of dynamic effects of adopting automation on training costs using the approach of Callaway and Sant'Anna (2021) specified in equation (2). I restrict the panel to the period 2010-2021 - excluding firms that adopt automation before 2011, because information on training costs from the SCIE data is not available before. The comparison group comprises firms that have not yet adopted automation, with the analysis conducted under the assumption of conditional parallel trends. I control for the log of labor productivity and the log of total hourly wage at the firm's baseline ($t = -1$). The coefficient for the year prior to adoption is normalized to zero. Panel A includes all firms that eventually adopt automation within the 2004-2021 period, while Panel B further disaggregates the sample by firm size—distinguishing between small firms (1-49 employees) and large firms (50 or more employees) based on their baseline values. Bootstrapped standard errors are clustered at the firm level, with confidence intervals provided for both the 5% and 10% significance levels.