# Robot Imports and Firm-Level Outcomes*

Alessandra Bonfiglioli† Rosario Crinò‡ Harald Fadinger§
Gino Gancia¶

This draft: March 28, 2024

**Abstract**

We use French data over the 1994-2013 period to study how imports of industrial robots affect firm-level outcomes. Guided by a simple model, we develop a novel empirical strategy to identify the causal effects of robot adoption. Our results suggest that, while demand shocks generate a positive correlation between robot imports and employment at the firm level, exogenous exposure to automation leads to job losses. We also find that robot exposure increases labour productivity and some evidence that it may raise the relative demand for high-skill professions.

**JEL Classification:** J23, J24, O33, D22

**Keywords**: Automation, Displacement, Firms, Robots

*We thank Philippe Aghion, Kirill Borusyak, Paula Bustos, David Hémous, Joseba Martinez, Pascual Restrepo and seminar participants at the Barcelona GSE Summer Forum (2019), ERWIT (Milan, 2020), the SED Meeting (Minneapolis, 2021), the CRC TR224 Conference, the CRC-ZEW Workshop "Labour Market Adjustments to New Technologies and Globalisation", the ECB Conference "Challenges in the Digital Era", the Third Advances in Economics Winter Symposium (Barzio, 2022), the AISSEC Conference (Pescara, 2022), the CEFES Workshop on "Europe in Unchartered Territories: Challenges and Policy Responses" (Milan, 2023), the Workshop on "Automation, Growth and Labor Markets" (Zurich, 2023), the EALE Conference (Padua, 2022), IHS Vienna, the OECD, University of Mannheim, University of Pavia, University of Southampton and WIIW (Vienna) for useful comments. Harald Fadinger gratefully acknowledges financial support from the DFG via CRC TR 224 (Project B06). Rosario Crinò gratefully acknowledges financial support from the Italian Ministry for University and Research under the PRIN 2022 program (project title: AUTOmation, PROductivity and Wage INequality (AUTOPROWIN): Firms and Workers in Times of Economic Turmoil; project number: 2022FLBY7J), D.D. MUR 104, February 2, 2022.

†University of Bergamo, Dept. of Economics, Queen Mary University of London and CEPR. E-mail: alessandra.bonfiglioli@unibg.it

‡University of Bergamo, Dept. of Economics, CEPR and CESifo. E-mail: rosario.crino@unibg.it.

§University of Mannheim and CEPR. E-mail: harald.fadinger@uni-mannheim.de.

¶University of Milan-Bicocca, Queen Mary University of London and CEPR. E-mail: ginoalessandro.gancia@unimib.it

# 1 Introduction

The widespread diffusion of industrial robots has fuelled growing concerns about the future of work. Robots are programmable machines with the capability to move on at least three axes. As such, unlike other pieces of equipment, they are designed to replicate human actions. The first prototype, the Unimate, was introduced in 1961 at General Motors to perform basic welding and carrying tasks. Other machines of this type were developed to assist human workers with a wide array of tasks, including heavy lifting, as well as hazardous or repetitive work, and their diffusion has grown at a staggering rate.[1] Industrial robots are technologies adopted by firms. To understand their effects, one must know how they affect the firms using them in the first place. Do robots substitute or complement humans in firms that automate? Are the effects heterogeneous across firms and workers? Do robots increase the productivity of firms using them? From a theoretical perspective, the answer to these questions is ambiguous. From an empirical perspective, the available evidence is scarce and often limited to correlations.

This paper is one of the first attempts to fill this gap. Our main innovations are to measure automation using detailed imports of industrial robots by French manufacturing firms and to use a novel empirical strategy for identifying causality. To guide the analysis, we build a simple model in which heterogeneous firms invest in automation, whose effect is to replace workers with capital in a set of tasks. Consistent with the conventional view, the employment effect of automation is potentially ambiguous: while robots displace some workers, they also increase productivity, which raises the demand for all factors. More importantly, the model shows that demand shocks are likely to increase employment and automation simultaneously, thereby generating a positive correlation between these variables. To overcome this bias, the model also illustrates how to isolate exogenous variation in firm-level *exposure to automation* that can be used to identify causal effects.

Our empirical results are consistent with the predictions of the theory. Focusing on the manufacturing sector, where automation is more prevalent, we first find that robot adopters are larger and have a larger employment share of high-skill professions. Second, looking at the evolution of firm-level outcomes over time, we find that robot import occurs after periods of expansion in firm size, suggesting that adoption may be driven by demand shocks, but is followed by a *decline* in employment. Third, using a new measure of exposure

[1] In 2021, there was an estimated stock of 3.5 million industrial robots. The future scale of the phenomenon is difficult to predict. Frey and Osborne (2017) argue that almost half of U.S. employment is at risk of being automated over the next two decades. See also Brynjolfsson and McAfee (2014) and Baldwin (2019).

to automation based on pre-determined technological characteristics, we find that firms that are more prone to adopt robots experience a stronger employment reduction than other firms. We also confirm that our proxy for exposure to automation is a significant predictor of robot imports. Throughout all specifications, we find that robots increase value added per worker and some evidence that they may raise the relative demand for high-skill professions.

These results suggest that demand shocks lead firms to both expand and automate, resulting in a positive correlation between robot adoption and employment. However, exogenous changes in automation lead to job displacement. Hence, they warn that caution should be exercised in interpreting the positive correlation between robot adoption and employment often found in the literature. In particular, there is a growing body of work that studies automation at the firm level. Some papers, like ours, measure automation with robot imports. These include Humlum (2019) for Denmark, Acemoglu *et al.* (2020) for France and Dixen *et al.* (2020) for Canada. Other papers use dummies from survey data. These include Dinlersoz and Wolf (2018) for the U.S., Cheng *et al.* (2019) for China, Koch *et al.* (2021) for Spain and a study by the European Commission (2015) for 7 European countries. None of these papers uses exogenous variation in automation across firms.

We are aware of two papers that try to identify causal effects using firm-level data. The first is Aghion *et al.* (2021), who use the same French data as us, but proxy automation with investment in industrial equipment. Employing a shift-share IV design, they find positive employment effects. As shown in our sensitivity analysis, we believe this result to be driven by the broader measure of capital inputs they consider, which is more likely to be complementary to labour. The second paper is Bessen *et al.* (2023), who use matched employer-employee data for the Netherlands. In line with our findings, they show that spikes in expenditure on "third-party automation services" increase job separations. Finally, our findings are consistent with Acemoglu and Restrepo (2020b) and Dauth *et al.* (2021), who identify the causal effects of automation across commuting zones using data from the International Federation of Robotics (IFR). By comparing firms within industries, our results reveal a new dimension of heterogeneity that cannot be observed in more aggregated data.

## 2 The Model

To guide the empirical analysis, we build a partial equilibrium model of endogenous automation across heterogeneous firms.[2] Consider a firm $i$ facing a demand function with a

[2]Proofs are in Appendix A.1. The model adds firm heterogeneity to task-based theories of endogenous automation such as Acemoglu and Restrepo (2018), Dechezleprêtre *et al.* (2021) and Hemous and Olsen

constant price-elasticity, $y_i = A_i p_i^{-\sigma}$. Production requires a unit measure of symmetric tasks. Tasks $z \in [0, \kappa_i]$ are automated, and thus can be performed by capital. The remaining tasks, $z \in (\kappa_i, 1]$, can only be performed by workers. Hence, $\kappa_i$ represents the extent of automation. Let $(k_i, l_i)$ denote the quantity of capital and workers, respectively, used by firm $i$. Denote with $r$ the rental rate of capital and with $w$ the wage of workers. We assume $r < w$, which implies that automated tasks are performed by capital only. Production of task $z$ is:

$$x_i(z) = \begin{cases} k_i(z) & \text{for } z \in [0, \kappa_i] \\ l_i(z) & \text{for } z \in (\kappa_i, 1] \end{cases} . \tag{1}$$

The production function of a firm with productivity $\varphi_i$ and automation $\kappa_i$ is:

$$y_i = \varphi_i \exp\left(\int_0^1 \ln x_i(z) \mathrm{d}z\right) = \varphi_i \left(\frac{k_i}{\kappa_i}\right)^{\kappa_i} \left(\frac{l_i}{1-\kappa_i}\right)^{1-\kappa_i}, \tag{2}$$

where $k_i/\kappa_i$ ($l_i/(1-\kappa_i)$) is capital (workers) per task.

Firms are monopolistically competitive and choose capital, labour and automation so as to maximize profits:

$$\max_{k_i, l_i, \kappa_i} \left\{p_i y_i - r k_i - w l_i - h f_i(\kappa_i)\right\},$$

where $h f_i(\kappa_i)$ is the cost of automation, including managers and engineers, with price $h$. The first-order conditions for capital and labour are:

$$r k_i = \left(1 - \frac{1}{\sigma}\right) \kappa_i p_i y_i \tag{3}$$

$$w l_i = \left(1 - \frac{1}{\sigma}\right) (1-\kappa_i) p_i y_i. \tag{4}$$

Eq. (3) shows that the demand for capital is increasing in automation, $\kappa_i$. Using (3)-(4) into (2) also shows that output per worker is increasing in $\kappa_i$ if $w > r$, as assumed:

$$y_i = \varphi_i \frac{l_i}{1-\kappa_i} \left(\frac{w}{r}\right)^{\kappa_i}. \tag{5}$$

Eq. (4) shows that automation, $\kappa_i$, has two opposite effects on the demand for labour. First, there is a direct negative displacement effect, given by the fact that more tasks are performed by capital. Second, as (5) shows, there is a positive productivity effect: an increase in $\kappa_i$

(2022), and it shares similarities with Acemoglu *et al.* (2015). See, instead, Martinez (2021) for a model of automation embodied in capital goods generating a distribution of technologies.

raises production, which in turn increases the demand for labour. The derivative of $l_i$ with respect to $\kappa_i$ is:

$$\frac{d\ln l_i}{d\kappa_i} = (\sigma - 1)\ln\left(\frac{w}{r}\right) - \frac{1}{1-\kappa_i},$$

which is positive for $\kappa_i < 1 - [(\sigma - 1)\ln(w/r)]^{-1}$. This condition is more likely to be satisfied when $\sigma$ and $w/r$ are high, i.e., when the productivity effect is strong enough. If $\sigma$ is high, production can be scaled up without a large countervailing fall in prices; and if $w/r$ is high, the cost saving of automation is stronger. If $(\sigma - 1)\ln(w/r) < 1$, instead, the displacement effect always dominates.[3]

Finally, consider the choice of automation, $\kappa_i$. We assume that automating more tasks poses an increasingly difficult challenge. For tractability, we focus on the functional form:

$$hf_i(\kappa_i) = h\frac{\rho_i}{1-\rho_i}\left[(1-\kappa_i)^{-\frac{1-\rho_i}{\rho_i}} - 1\right],$$

with $\rho_i \in (0,1)$. The parameter $1/\rho_i$ captures the rate at which the marginal cost of automation increases with $\kappa_i$. To see this, note that the marginal cost of automation, $hf_i'(\kappa_i) = h(1-\kappa_i)^{-1/\rho_i}$, increases at a faster rate with $\kappa_i$ the lower $\rho_i$ is. Hence, we interpret $\rho_i$ as an index of replaceability of tasks in the production process and we allow it to vary across firms. The first-order condition for $\kappa_i$ is:

$$\frac{1}{1-\kappa_i} = \left[\left(1-\frac{1}{\sigma}\right)\frac{p_i y_i}{h}\ln\left(\frac{w}{r}\right)\right]^{\rho_i}. \tag{6}$$

Larger firms (higher $A_i$ and $\varphi_i$) have a stronger incentive to pay the fixed automation cost to save on the variable production cost; automation is also increasing in the cost-saving it entails $(w/r)$ and decreasing in its own cost $h$ and in $1/\rho_i$.[4]

The model shows that the effect of $\kappa_i$ on $l_i$ is potentially ambiguous, so that whether or not automation raises employment is ultimately an empirical question. It also illustrates that the key empirical challenge is the endogeneity of $\kappa_i$. Specifically, demand shocks trigger automation but also have a direct positive effect on labour demand. A way to overcome this problem is to focus on exogenous determinants of automation that have no effects on firm outcomes other than through $\kappa_i$. These are the parameters capturing automation costs, namely, the industry-level cost shifter, $h$, and firm-level replaceability, $\rho_i$. However, eq. (6)

[3] Acemoglu and Restrepo (2018) allow new tasks to be created when others are automated. We abstract from this additional mechanism, which would reinforce the positive productivity effect on employment.

[4] We show in Appendix A.3 that a variant of the model where automation is a discrete choice yields qualitatively similar results.

shows that $h$ and $\rho_i$ do not operate separately: a lower automation cost has no effects on firms without replaceable tasks ($\rho_i = 0$). Likewise, replaceability is immaterial in an industry where automation costs are prohibitive. Rather, $h$ and $\rho_i$ interact with each other (proof in Appendix A.1):

$$\frac{\partial^2 \kappa_i}{\partial \ln h \partial \rho_i} < 0. \tag{7}$$

Based on this insight, the model suggests to identify firms as more exposed to automation when they employ a larger share of replaceable workers (high $\rho_i$) while operating in industries that are better suited for automation (low $h$).

## 3 Data and Preliminary Evidence

Our empirical analysis uses firm-level data for France over the 1994-2013 period and combines several datasets administered by the French statistical agency (INSEE), covering the universe of French firms (legal entities) that report a complete balance sheet. For each firm, we have data on sales, material purchases, capital stock (value of physical assets) and total employment from the BRN and FARE datasets; using this information, we also compute firm-level value added. We complement the balance sheet data with information on the occupational structure of employment in each firm from DADS Etablissement. For each year, this dataset contains employment data disaggregated into five 2-digit occupations. For the year 1994, it also contains a finer employment disaggregation into 29 occupations, which we exploit when constructing our proxy for robot exposure.[5] For the descriptive analysis, we use the full set of years (1994-2013) while for identification we focus on the 1996-2013 period and take 1994 as a pre-sample year.

For each firm and year, we also have data on values and quantities of exports and imports for all 8-digit products of the Combined Nomenclature (CN) classification from the French customs authority (DOUANE). The CN classification records trade in industrial robots into a specific product code, CN 84795000 (CN 84798950 before 1996). Accordingly, we identify firms that import robots in a given year as firms with positive imports for this product code. We also measure the stock of robot capital employed by a firm at a given point in time as the sum of robot imports by the firm up to that point.

Robot imports are recognized as a good proxy for automation because of the high concentration of this sector; see, e.g., Blanas *et al.* (2019), Acemoglu and Restrepo (2022) and

[5] Yearly employment data for the 29 occupations are only available starting from the 2010s, and are thus missing for most of our sample period.

Bonfiglioli *et al.* (2023). For instance, Japan and Germany alone account for 50% of the total volume of global exports, while France's share is about 5% only. Yet, the use of import data is subject to some limitations. On the one hand, they include imports by robot integrators or resellers, which do not represent actual instances of adoption. On the other, they do not include purchases of robots from domestic suppliers. Moreover, in the case of intra-EU transactions, firms are not required to report the list of imported products as long as their overall intra-EU imports are below a given threshold. To mitigate these issues, we restrict the sample to the manufacturing sector, where robot users are more prevalent, and drop the "Installation and Repair of Machinery and Equipment" industry. The sector of operation and the characteristics of robot importers, such as sales and size, in our final sample make it unlikely that these are just robot integrators. Although the reporting thresholds are not very high given the average price of a robot, we further restrict the analysis to firms with more than ten employees, for which the thresholds are less likely to be binding. More importantly, our identification strategy will circumvent all the limitations of import data by exploiting variation in proxies for robot exposure based on technological characteristics that are observed for all firms and not just importers.

Consistent with other studies, Appendix Figure B1 shows that robot importers are particularly frequent in the production of motor vehicles, machinery, and electrical equipment. However, robot importers are likely to be undercounted in the "Manufacture of Motor Vehicles" industry because our data lack information for the two biggest car manufacturers in France.[6] After removing this industry, the correlation between the number of robot importers and the stock of installed robots from the IFR is 0.79.

Our baseline sample is an unbalanced panel of 57,110 manufacturing firms, of which 859 have imported robots at least once over 1994-2013 (henceforth, "robot adopters"). This number is consistent with Acemoglu *et al.* (2020), who collected information on robot adoption in France from multiple sources for the 2010-2015 period. While robot adopters are a small minority, they account for a large and growing fraction of manufacturing activity. Considering firms that are active in all years between 1994 and 2013, the shares of robot adopters in manufacturing employment and value added have increased steadily to reach 8% and 14%, respectively. This indicates that robot adopters are faring better than other manufacturing firms. Moreover, the value added share has grown significantly more than the employment share, suggesting that the expansion of robot adopters may have been

[6]For large multinational firms (e.g., Peugeot Société Anonym and Renault), INSEE reports only consolidated balance sheets of the entire group. Since the identity and composition of these groups is not constant across periods, they cannot be included for comparisons over time.

accompanied by the adoption of labour-saving technologies.[7]

Appendix Table B1 reports summary statistics separately for robot adopters and non adopters, showing that the former firms are larger and more skill-intensive than the latter, on average. To gain further insight into the differences between the two groups of firms, we estimate conditional correlations between robot adoption and firm-level characteristics by running OLS regressions of the following form:

$$Y_{it} = \alpha_i + \alpha_{jt} + \beta \cdot Adoption_{it} + \mathbf{X}'_{it} \cdot \boldsymbol{\gamma} + \varepsilon_{it}, \tag{8}$$

where $i$ denotes a firm; $j$ indicates the 5-digit NACE industry in which the firm operates; and $t$ stands for time. $Y_{it}$ is an outcome and $Adoption_{it}$ is a dummy that takes value 1 from the first year in which the firm imports robots onwards, and is equal to 0 otherwise. We control for (i) firm fixed effects, $\alpha_i$, to absorb time-invariant firm characteristics; (ii) 5-digit industry×year fixed effects, $\alpha_{jt}$, to account for differences in the industry of operation and for industry-specific shocks; and (iii) firm characteristics—log sales and dummies for firms that export or import goods other than robots—measured in the first year in which the firm is observed and interacted with a full set of year dummies, $\mathbf{X}_{it}$. We estimate (8) for four major outcomes that can be directly constructed from the data and on which we focus throughout the paper: (i) log sales, (ii) log employment, (iii) log value added per worker and (iv) the employment share of high-skill professions. The results are in Table 1; standard errors are corrected for clustering at the firm level and $t$-statistics are shown in square brackets. All estimates of $\beta$ are positive and, with the exception of the regression for the employment share of high-skill workers, they are also highly statistically significant.

Do robot adopters differ from other firms already before adopting robots, or do they start diverging afterwards? To shed light on this question, we use a difference-in-differences event study approach to analyse how the four outcomes evolve over time in robot adopters relative to other firms. We estimate the following specification:

$$Y_{it} = \alpha_i + \alpha_{jt} + \sum_{s=-5}^{-2} \beta_s \cdot D_{it}^s + \sum_{s=0}^{5} \beta_s \cdot D_{it}^s + \varepsilon_{it}, \tag{9}$$

[7]Preliminary evidence from a 2019 survey run by the U.S. Census shows similar patterns. In particular, Acemoglu *et al.* (2021) report that about 2% of firms use robotics for automation and these firms account for about 15% of employment.

Table 1: Firm-Level Outcomes and Robot Adoption

| | (1) Ln Sales | (2) Ln No. of Employees | (3) Ln VA per Worker | (4) Empl. Sh. High Skill |
|---|---|---|---|---|
| $\text{Adoption}_{it}$ | 0.230*** | 0.106*** | 0.057*** | 0.003 |
| | [10.458] | [5.763] | [3.630] | [1.030] |
| Obs. | 596,166 | 597,282 | 585,886 | 597,282 |
| R2 | 0.95 | 0.87 | 0.85 | 0.70 |

The subscripts $i$ and $t$ denote firms and years, respectively. The dependent variables are annual observations of the firm-level outcomes indicated in columns' headings. Value added is computed as sales minus changes in inventories minus purchases of final goods minus purchases of materials plus changes in material inventories minus other purchases. High-skill professions are scientists, managers and engineers. $\textit{Adoption}_{it}$ is a dummy equal to 1 for all years since the firm starts importing robots, and equal to 0 otherwise. All specifications include firm fixed effects and 5-digit industry x year fixed effects. They also control for log sales and dummies for whether the firm is an importer or an exporter; each control variable is observed in the first year in which the firm appears in the sample and is interacted with a full set of year dummies. Standard errors are corrected for clustering within firms; t-statistics are reported in square brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

where $D_{it}^{s}$ are event-study dummies, which are equal to 1 if firm $i$ is $s$ periods away from the first instance of robot imports at time $t$, and 0 otherwise. The coefficients $\beta_s$ illustrate how a given outcome evolves over time within robot adopters relative to non adopters over a ten-year window around initial adoption ($s = 0$).

The results are in Figure 1 and the estimation coefficients in Appendix Table B2. Robot adoption is antedated by significant differences in the trends of sales and employment between robot adopters and non adopters. Specifically, the former firms grow faster than the latter in terms of both variables over the five-year period preceding adoption. Conversely, no clear differential pre-trend is detected in terms of value added per worker and the skill composition of the workforce. After adoption, the diverging trend in employment is reversed: while robot adopters still grow faster than non adopters, the differential gradually vanishes. Robot adopters also experience a relatively stronger shift in the skill composition of the workforce towards high-skill professions, and a faster increase in value added per worker. These results show that robot adoption occurs after periods of expansion in firm size, and is followed by employment losses, improvements in efficiency, and labour demand shifts towards high-skill workers, with limited changes in total sales.[8]

[8] In Appendix Figure B2, we repeat the analysis using the estimator developed by Borusyak *et al.* (2022) to exploit the staggered nature of adoption across firms. While this estimator fails to converge when we control for 5-digit industry×year fixed effects, a more parsimonious specification without these fixed effects delivers results qualitatively similar to those in Figure 1.

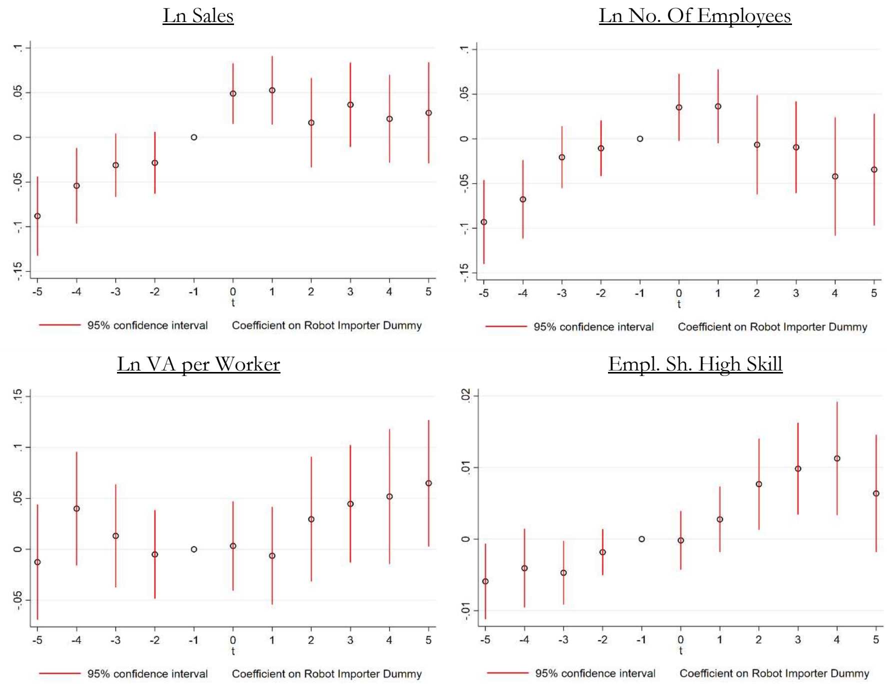


Each graph plots coefficients and confidence intervals on event-study dummies, estimated using eq. (9) for a different outcome variable (indicated in the heading of the graph). Each event-study dummy $D^s_{it}$ is equal to 1 if firm $i$ is $s$ periods away from the first instance of robot imports at time $t$ and is equal to 0 otherwise. The estimated coefficients corresponding to each graph are reported in Appendix Table B2.

Figure 1: Difference-in-Differences Event Studies

# 4 The Effect of Robot Exposure

The results in Figure 1 suggest that the positive correlation between automation and employment may be spurious. We now exploit differential cross-firm variation in robot exposure stemming from pre-determined technological characteristics to identify causal effects.

## 4.1 Variables and Specification

Our model suggests that exposure to automation depends on the interaction between industry-level suitability to automation and firm-level replaceability of employment. On the one hand, some industries have lower automation costs than others (low $h$ in the model) due to the nature of their production process. This is the case of industries that use assembly lines compared to those that rely more on batch production or make customized items. On the other hand, within any given industry, firms differ in the extent to which their activities can be assigned to robots ($\rho_i$ in the model). For example, some firms outsource the assembly stage or are specialized in complex tasks that cannot be automated. Building on these insights, our robot exposure measure, $RobExp$, exploits the interplay between a proxy for automation suitability in each industry, $RobSuit$, and a proxy for employment replaceability within each firm, $Repl$. With this variable, we study how robot exposure affects firm-level outcomes and adoption decisions.

In a given 5-digit industry $j$, we measure $RobSuit$ as the average robot intensity of all firms $i' \neq i \in j$ in the initial year, namely,

$$RobSuit_{j-i} = \sinh^{-1}\left(\frac{\sum_{i' \neq i \in j} RobStock_{i'}}{\sum_{i' \neq i \in j} CapStock_{i'}}\right), \tag{10}$$

where $RobStock_{i'}$ and $CapStock_{i'}$ denote firm $i'$'s initial stocks of imported robots and total capital, respectively. Sectors with a lower cost of automation will naturally have a higher stock of robots. Normalizing it by total capital ensures that $RobSuit_{j-i}$ is not affected by differences in scale or capital intensity and the hyperbolic sine transformation preserves the zeros.

Our replaceability measure, $Repl$, is similar to the indicator of Graetz and Michaels (2018) but is defined at the firm-level. We source from Graetz and Michaels (2018) information on whether each of 377 U.S. Census occupations is replaceable or not. An occupation is defined as replaceable if its title corresponds to at least one of the robot application categories identified by the IFR (e.g., welding, painting, assembling). We map each U.S. occupation into the 29 French occupations for which we have employment data in 1994, and construct firm-level replaceability as

$$Repl_i = \sum_{o=1}^{29} \omega_{oi} \cdot Repl_o, \tag{11}$$

where $Repl_o$ is the replaceability of French occupation $o$ and $\omega_{oi}$ its share in firm $i$'s employment in 1994. The values of $Repl_o$ are in Appendix B (Table B3), where we argue that mapping from U.S. Census occupations induces only a minimal loss of variation in $Repl_i$ across firms. Finally, we obtain $RobExp$ as

$$RobExp_i = RobSuit_{j-i} \cdot Repl_i. \tag{12}$$

Appendix Table B4 provides descriptive statistics on $RobExp_i$ and the other variables used in this section.

We focus on long-run changes and estimate specifications of the following form:

$$\Delta Y_i = \alpha_j + \beta_1 \cdot RobExp_i + \beta_2 \cdot RobSuit_{j-i} + \beta_3 \cdot Repl_i + \mathbf{X}_i' \cdot \boldsymbol{\gamma} + \varepsilon_i, \tag{13}$$

where $\Delta Y_i$ is the annualized change in outcome $Y$ for firm $i$ between the first and last year in which the firm is present in the sample. The use of long differences implies that identification exploits cross-sectional (across firms) variation in the pre-determined level of robot exposure and in the long-run growth of outcomes. We control for 5-digit industry fixed effects, $\alpha_j$, to absorb differential trends in outcomes across industries. Moreover, we control for start-of-period values of log firm sales and of indicators for exporting and importing firms, included in $\mathbf{X}_i$. While firms with high and low replaceability are comparable in terms of these characteristics in the first year (Appendix Table B5), it is possible that initially larger and more trade-oriented firms will follow different paths in terms of adoption and outcomes in subsequent periods; controlling for $\mathbf{X}_i$ absorbs possible heterogeneous trends across firms within the same industry. We correct the standard errors for clustering within 5-digit industries to account for possibly correlated shocks among firms within industries.[9]

As suggested by the model (eq. 7), we believe that neither $Repl_i$ nor $RobSuit_{j-i}$ alone fully captures a firm's exposure to robots. For instance, replaceability of employment cannot trigger automation in industries where robots are not available. Recognizing this, our empirical approach goes beyond the level effect of these variables and focuses instead on their interaction in a difference-in-differences specification. While $Repl_i$ and $RobSuit_{j-i}$ are pre-determined and thus do not respond to subsequent changes in firm-level outcomes, they could still be correlated with other firm or industry variables affecting outcomes. In Section 4.3, we discuss the main threats to identification and perform an extensive sensitivity analysis.

### 4.2 Baseline Results and Robustness Checks

The baseline estimates of $\beta_1$ are in Table 2, panel a). Coefficients are multiplied by 100 to express them in percentages.[10] The full list of coefficients is in Appendix Table B7. The coefficient on $RobExp_i$ is negative and precisely estimated in the employment regression,

[9] We winsorise the change in each outcome at the top and bottom 5% of the distribution to prevent results from being driven by extreme observations.

[10] In Appendix Table B6, we estimate (13) replacing $RobExp_i$, $RobSuit_{j-i}$ and $Repl_i$ with a dummy for firms that adopt robots over the sample period. Consistent with our preliminary evidence, these firms experience a relatively larger increase in size and value added per worker, and a somewhat faster shift in labour demand towards high-skill workers.

indicating that more robot exposed firms experience a relatively larger and statistically significant reduction in employment over the sample period. As for the other outcomes, $\beta_1$ is positive and precisely estimated in the regression for value added per worker, suggesting that higher robot exposure may induce efficiency gains within firms. The effect of robot exposure on sales, albeit positive, is not statistically significant. This suggests that productivity gains need not always translate into lower prices relative to competitors. Finally, the results point towards a positive, albeit imprecisely estimated, effect of robot exposure on the skill structure of employment. As for the other regressors, Appendix Table B7 shows that the coefficient on $Repl_i$ generally has the same sign as that on our preferred measure of robot exposure, $RobExp_i$.[11]

The estimates in Table 2 imply that a change in automation suitability equal to the interquartile range of its distribution (12.12) is associated with an employment fall of 0.36 percentage points (p.p.) per year in firms at the 75th percentile of the replaceability distribution (0.52) relative to firms at the 25th percentile (0.20). As an example, the firm with average replaceability in the "Manufacture of Parts and Accessories for Motor Vehicles" industry (high suitability) would experience a 0.30 p.p. per year employment fall relative to the firm with average replaceability in the "Manufacture of Wine from Grape" industry (low suitability). Moreover, the average increase in $RobSuit_{j-i}$ over 1994-2013 (8.55) would induce a 0.26 p.p. per year employment fall at the 75th percentile of the replaceability distribution relative to the 25th percentile.

Our working hypothesis is that robot exposure affects outcomes by inducing firms to adopt robots. In column (5), we study this mechanism by estimating (13) with a different dependent variable, $Adopter_i$, a dummy equal to 1 for firms that start importing robots over the sample period. The coefficient on $RobExp_i$ is positive and precisely estimated, implying that firms that are more exposed to robots do indeed show a greater tendency to adopt robots in subsequent years. Appendix Table B8 shows that $RobExp_i$ is also associated with an increase in robot intensity (stock of robot capital over total capital stock) among adopters, with a coefficient $\beta_1$ equal to 0.097 ($t$-statistics 2.197). Hence, firms with higher robot exposure also exhibit a more intensive use of robots in production. Because $Adopter_i$ only captures the extensive margin of automation and is driven by a small number of firms, we

[11] Appendix Table B8 also reports small and insignificant effects of $RobExp_i$ on the labour share of value added. As we show in Appendix A.2, this result is consistent with the model if the automation cost is in terms of high-skill workers.

Table 2: Firm-Level Outcomes and Robot Exposure, Main Results and Robustness

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | Δ Ln Sales | Δ Ln No. of Employees | Δ Ln VA per Worker | Δ Empl. Sh. High Skill | Adopter |
| | | | **a) Baseline Regressions** | | |
| $RobExp_i$ | 0.148 | -0.094** | 0.302*** | 0.006 | 0.174*** |
| | [1.343] | [-2.095] | [2.702] | [1.106] | [2.893] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | | | **b) Weighted Regressions** | | |
| $RobExp_i$ | 0.142 | -0.108** | 0.310*** | 0.008 | 0.224*** |
| | [1.192] | [-2.230] | [2.629] | [1.396] | [2.666] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.06 | 0.05 | 0.07 |
| | | | **c) Excluding Manufacturing of Motor Vehicles** | | |
| $RobExp_i$ | 0.148 | -0.095** | 0.303*** | 0.005 | 0.171*** |
| | [1.329] | [-2.101] | [2.695] | [0.837] | [2.847] |
| Obs. | 35,759 | 36,040 | 34,647 | 36,040 | 36,040 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | | | **d) Broader Definition of Robot Imports** | | |
| $RobExp_i$ | 0.120 | -0.044 | 0.191* | 0.011** | 0.208 |
| | [1.234] | [-0.946] | [1.788] | [2.282] | [1.454] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.11 |
| | | | **e) Interactions with Demand Elasticity** | | |
| $RobExp_i$ | -0.160 | -0.203** | -0.061 | 0.001 | 0.065 |
| | [-0.737] | [-2.020] | [-0.270] | [0.111] | [0.414] |
| $RobExp_i$ x $Elast_h$ | 0.069* | 0.023 | 0.076* | 0.002 | 0.023 |
| | [1.963] | [1.405] | [1.955] | [0.774] | [0.838] |
| Obs. | 32,427 | 32,679 | 31,365 | 32,679 | 32,679 |
| R2 | 0.11 | 0.04 | 0.07 | 0.04 | 0.05 |
| | | | **f) Alternative Proxy for Robot Exposure (IFR)** | | |
| $RobExp_i$ | 3.331*** | 0.248 | 3.537*** | 0.070** | 0.625*** |
| | [9.669] | [1.043] | [11.543] | [2.537] | [3.469] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |

The subscript $i$ denotes firms. In columns (1)-(4), the dependent variables are annualized changes in the firm-level outcomes indicated in columns' headings. In column (5), the dependent variable is $Adopter_i$, a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non importers. With the exception of panel f), $RobExp_i$ is the product between the initial firm-level replaceability of employment by robots ($Repl_i$) and the initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry $j$ ($RobSuit_{j\text{-}i}$). In panel f), $RobExp_i$ is constructed by replacing $RobSuit_{j\text{-}i}$ with the log stock of installed robots in each firm's sector in the U.S., based on data from the International Federation of Robotics (IFR) for 13 manufacturing sectors. The regressions in panel b) are weighted by the initial number of employees in each firm. The sample in panel c) excludes firms in the "Manufacture of Motor Vehicles" industry. In panel d), robot imports include CN codes 842489, 842890, 851580, 847950, 851531, 851521 and 848640. In panel e), $Elast_h$ is the elasticity of demand, defined at the 3-digit sector level, $h$; the specification also includes interactions of $Elast_h$ with $Repl_i$ and $RobSuit_{j\text{-}i}$ (coefficients unreported). All regressions also include the linear terms in $Repl_i$ and $RobSuit_{j\text{-}i}$, initial values of log sales and of dummies for importing and exporting firms, and 5-digit industry fixed effects. Standard errors are corrected for clustering within 5-digit industries; t-statistics are reported in square brackets. All coefficients are multiplied by 100 to express them in percentages. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

refrain from interpreting the coefficient in column (5) as the first stage in a 2SLS framework. Nevertheless, in Appendix C, we perform a quantification exercise to have a sense of the magnitude of the effect of exogenous robot adoption. According to our estimates, the latter would explain an average annual fall in employment equal to 2.81% in robot adopters relative to non adopters.

The remaining panels of Table 2 contain an extensive sensitivity analysis. In panel b), we weigh the observations by each firm's initial number of employees. In panel c), we further exclude firms in the "Manufacture of Motor Vehicles" industry. The results are largely unchanged. In panel d), we extend the definition of automation suitability to include all types of machinery designed for lifting, handling, loading, unloading and welding. Robot exposure no longer has a statistically significant effect on employment but induces a stronger shift in labour demand towards high-skill professions. These results are consistent with the finding in Aghion *et al.* (2021) that broader forms of capital equipment are more complementary to labour.

The model predicts that in industries where demand is more elastic, the productivity effect of automation should be stronger because firms can scale up production without large reductions in prices (Bessen, 2019). We therefore extend (13) by interacting $RobExp_i$, $Repl_i$, and $RobSuit_{j-i}$ with the elasticity of substitution in each sector (Broda and Weinstein, 2006). The results confirm that robot exposure causes a relatively larger increase in sales in sectors where products are more substitutable (panel e)). Similarly, it has a relatively less negative employment effect in sectors where demand is more elastic, although the interaction coefficient is marginally insignificant. The effect of robot exposure on value added per worker is also relatively stronger in sectors where products are more substitutable.

Finally, we use an alternative proxy for robot exposure obtained by replacing $RobSuit_{j-i}$ with the log initial stock of installed robots in each firm's sector in the U.S., sourced from the IFR (International Federation of Robotics, 2022). This variable includes domestically sourced robots and, being based on U.S. rather than French data, may be more exogenous. However, the IFR data are only available for 13 aggregate manufacturing sectors, so variation is much more limited. The coefficient on $RobExp_i$ in the employment regression is not statistically significant, though the coefficient on $Repl_i$ alone remains strongly negative (not reported). The effect of $RobExp_i$ on all other outcomes, including now sales, are strongly positive. These results suggest that the IFR sectors may be too coarse to dispel the concern of correlations with positive shocks.

### 4.3 Threats to Identification

We now discuss the main threats to identification. A first concern is that the fixed effects and controls in (13) might not fully account for heterogeneous trends across firms. In Table 3, we thus augment (13) with additional covariates and fixed effects, and study how the coefficient on $RobExp_i$ is affected. In panel a), we add each firm's initial employment share of high-skill workers and capital-labour ratio. Initial differences in skill and capital intensities may reflect various factors leading to heterogeneous trends, such as differences in production organization, product quality and firms' stage within the life cycle. Yet, these controls do not significantly alter our main findings.

Despite our highly disaggregated industry fixed effects, some differences might persist in the types of goods produced by firms within an industry. For the sub-sample of exporting firms, the customs data provide us with the full list of exported products according to the highly detailed CN classification. In panel b), we thus focus on firms that export goods at least once, and include a fixed effect for each firm's main export product (by total export value) at the CN 6-digit level, along with a dummy for multi-product exporters. Accordingly, we compare firms that export similar goods within an industry and control for product-specific trends. The main results are preserved in this very demanding specification.

Differences in initial characteristics may give rise to heterogeneous trends also by introducing variation in the quality of imported robots or by influencing firms' decisions to source robots from domestic integrators rather than foreign suppliers. In a first exercise, we leverage again our rich customs data augmenting the specification with the average unit value of imported robots as a proxy for their quality (panel c)). In a second exercise, we re-estimate (13) on firms that never import robots (panel d)). This is possible because our robot exposure variable, $RobExp_i$, is defined for all firms, not just for importers. Our evidence is confirmed in both cases. In panel e), we interact $RobExp_i$, $Repl_i$, and $RobSuit_{j-i}$ with $Adopter_i$; consistent with panel d), all coefficients on the $RobExp_i \times Adopter_i$ interaction are statistically insignificant.

A second identification concern is that, even after accounting for fixed effects and covariates, $RobExp_i$ might be correlated with omitted variables that impact outcomes. In particular, identification would be threatened if: (i) $Repl_i$ was correlated with other firm characteristics that differentially affect outcomes across industries with varying levels of

Table 3: Threats to Identification: Additional Controls and Alternative Samples

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | Δ Ln Sales | Δ Ln No. of Employees | Δ Ln VA per Worker | Δ Empl. Sh. High Skill | Adopter |
| | **a) Controls for Initial Skill and Capital Intensities** | | | | |
| $RobExp_i$ | 0.124 | -0.092** | 0.263** | 0.004 | 0.184*** |
| | [1.186] | [-2.015] | [2.505] | [0.494] | [2.939] |
| Obs. | 35,747 | 36,023 | 34,663 | 36,023 | 36,023 |
| R2 | 0.10 | 0.04 | 0.07 | 0.07 | 0.05 |
| | **b) Exported Products Fixed Effects, Sub-Sample of Exporting Firms** | | | | |
| $RobExp_i$ | 0.033 | -0.142*** | 0.252** | 0.008 | 0.148** |
| | [0.376] | [-2.610] | [2.529] | [1.235] | [2.128] |
| Obs. | 28,539 | 28,761 | 27,534 | 28,761 | 28,761 |
| R2 | 0.24 | 0.19 | 0.20 | 0.17 | 0.17 |
| | **c) Control for Average Robot Import Price** | | | | |
| $RobExp_i$ | 0.148 | -0.095** | 0.303*** | 0.006 | 0.164*** |
| | [1.349] | [-2.104] | [2.707] | [1.106] | [3.575] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.79 |
| | **d) Sub-Sample of Non Robot Adopters** | | | | |
| $RobExp_i$ | 0.153 | -0.094** | 0.310*** | 0.008 | - |
| | [1.366] | [-2.034] | [2.752] | [1.246] | - |
| Obs. | 35,729 | 36,008 | 34,633 | 36,008 | - |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | - |
| | **e) Robot Adopters vs. Non Robot Adopters** | | | | |
| $RobExp_i$ | 0.146 | -0.095** | 0.306*** | 0.007 | - |
| | [1.319] | [-2.101] | [2.734] | [1.162] | - |
| $RobExp_i$ x $Adopter_i$ | -0.184 | -0.131 | -0.243 | -0.015 | - |
| | [-1.176] | [-0.893] | [-1.472] | [-0.883] | - |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | - |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | - |

The subscript $i$ denotes firms. In columns (1)-(4), the dependent variables are annualized changes in the firm-level outcomes indicated in columns' headings. In column (5), the dependent variable is $Adopter_i$, a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non importers. $RobExp_i$ is the product between the initial firm-level replaceability of employment by robots ($Repl_i$) and the initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry $j$ ($RobSuit_{j\text{-}i}$). The specification in panel a) includes the initial skill intensity (employment share of high-skill workers) and the initial capital intensity (capital-to-labor ratio) of each firm. The specification in panel b) is estimated on the sub-sample of firms that export goods at least once over the sample period and includes both a fixed effect for each firm's main export product (by total export value) at the CN 6-digit level and an indicator variable for multi-product exporters. The specification in panel c) includes the average unit value of imported robots (hyperbolic sine transformation). The sample used in panel d) consists of firms that do not import robots over the sample period. The specification in panel e) also includes the interactions of $Repl_i$ and $RobSuit_{j\text{-}i}$ with $Adopter_i$ (coefficients unreported). Column (5) is missing in panels d) and e) because the dependent variable, $Adopter_i$, is always zero (one) among non robot adopters (adopters). All regressions also include the linear terms in $Repl_i$ and $RobSuit_{j\text{-}i}$, initial values of log sales and of dummies for importing and exporting firms, and 5-digit industry fixed effects. Standard errors are corrected for clustering within 5-digit industries; t-statistics are reported in square brackets. All coefficients are multiplied by 100 to express them in percentages. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

automation suitability; or (ii) $RobSuit_{j-i}$ captured other industry characteristics that heterogeneously affect outcomes across firms with different degrees of replaceability. We now extend the specification by adding interactions of $Repl_i$ and $RobSuit_{j-i}$ with some of the most plausible confounders, and study how this influences the coefficient on $RobExp_i$.

In Table 4, panel a), we interact $RobSuit_{j-i}$ with each firm's routine intensity—the share of routine-intensive occupations in total employment in 1994. We construct this variable by mapping Autor and Dorn's (2013) data on occupational routine intensity to the 29 French occupations; the linear term in routine intensity is included but untabulated. While routine intensity is correlated with the adoption of skill-intensive technologies such as computers (e.g., Autor *et al.* 2003), Cheng *et al.* (2019) find that robots are more prevalent at firms where employees commonly perform manual rather than routine tasks. Consistently, the new interaction has no significant effect on robot adoption, and its inclusion leaves the evidence on $RobExp_i$ unaffected. In panel b), we add interactions between $RobSuit_{j-i}$ and all the control variables in $\mathbf{X}_i$. Consistent with size and trade orientation being comparable across firms with high and low replaceability (Appendix Table B5), the main results are preserved. Similarly, panel c) shows that our evidence is largely unchanged when controlling for the interaction between each variable in $\mathbf{X}_i$ and $Repl_i$.

Next, we consider the possibility that $Repl_i$ interacts with industry characteristics other than $RobSuit_{j-i}$. In panel d), we add interactions between $Repl_i$ and: (i) total imports and exports; (ii) average unit values of imports and exports; and (iii) imports of capital and technology goods. We construct these variables in the initial year by aggregating across firms other than $i$ in each 5-digit industry. These characteristics also enter the specification linearly. The coefficients on $RobExp_i$ remain close to the baseline estimates. Finally, we add interactions between $Repl_i$ and 2-digit sector dummies (panel e)). Contributing to the identification of $\beta_1$ is now only the remaining variation in $RobSuit_{j-i}$ across narrow (5-digit) industries within the same 2-digit sector. The main patterns are preserved.

## 5 Conclusions

We have studied the effects of industrial robots using data for French firms between 1994 and 2013. Our results suggest that, while robot adopters are growing in employment relative to other firms, exogenous exposure to automation leads to significant job losses. There is also some evidence that automation may increase the relative demand for high-skill professions.

Table 4: Threats to Identification: Additional Interactions

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | Δ Ln Sales | Δ Ln No. of Employees | Δ Ln VA per Worker | Δ Empl. Sh. High Skill | Adopter |
| | **a) Interaction of Robot Suitability with Routine Intensity** | | | | |
| $RobExp_i$ | 0.151 | -0.090** | 0.297*** | 0.006 | 0.181*** |
| | [1.385] | [-1.994] | [2.676] | [1.005] | [3.055] |
| $RobSuit_{j\text{-}i}$ x $Routine_i$ | -2.934 | 4.864 | 9.589 | 1.193*** | 2.545 |
| | [-0.129] | [0.829] | [0.433] | [2.781] | [0.355] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | **b) Interactions of Robot Suitability with Firm Characteristics** | | | | |
| $RobExp_i$ | 0.147 | -0.094** | 0.301*** | 0.006 | 0.172*** |
| | [1.335] | [-2.093] | [2.695] | [1.098] | [2.876] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | **c) Interactions of Replaceability with Firm Characteristics** | | | | |
| $RobExp_i$ | 0.190* | -0.096** | 0.348*** | 0.007 | 0.186*** |
| | [1.697] | [-2.117] | [3.080] | [1.140] | [2.966] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | **d) Interactions of Replaceability with Industry Characteristics** | | | | |
| $RobExp_i$ | 0.176 | -0.136*** | 0.402*** | 0.012* | 0.155*** |
| | [1.608] | [-2.752] | [4.023] | [1.808] | [2.616] |
| Obs. | 36,254 | 36,537 | 35,134 | 36,537 | 36,537 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |
| | **e) Interactions of Replaceability with Sector Dummies** | | | | |
| $RobExp_i$ | 0.398*** | -0.118** | 0.633*** | 0.017** | 0.235*** |
| | [3.025] | [-2.233] | [5.811] | [2.338] | [3.465] |
| Obs. | 36,301 | 36,584 | 35,180 | 36,584 | 36,584 |
| R2 | 0.10 | 0.04 | 0.07 | 0.04 | 0.05 |

The subscripts $i$ and $j$ denote firms and 5-digit industries, respectively. In columns (1)-(4), the dependent variables are annualized changes in the firm-level outcomes indicated in columns' headings. In column (5), the dependent variable is $Adopter_i$, a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non importers. $RobExp_i$ is the product between the initial firm-level replaceability of employment by robots ($Repl_i$) and the initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry $j$ ($RobSuit_{j\text{-}i}$). In panel a), $Routine_i$ is the initial firm-level employment share of routine-intensive occupations; the specification also includes the linear term in $Routine_i$ (coefficient unreported). The specifications in panels b) and c) include interactions of $RobSuit_{j\text{-}i}$ and $Repl_i$, respectively, with the initial values of log sales and of dummies for importing and exporting firms. The specification in panel d) includes the initial values of sectoral exports and imports, export and import unit values, capital goods and technology goods imports, as well as the interactions of these variables with *Repl* The specification in panel e) includes interactions of $Repl_i$ with a full set of 2-digit sector dummies. All regressions also include the linear terms in $Repl_i$ and $RobSuit_{j\text{-}i}$, initial values of log sales and of dummies for importing and exporting firms, and 5-digit industry fixed effects. Standard errors are corrected for clustering within 5-digit industries; t-statistics are reported in square brackets. All coefficients are multiplied by 100 to express them in percentages. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

These results are important because the normative literature has shown that automation may call for corrective measures if it displaces workers and/or increases inequality (e.g., Beraja and Zorzi, 2022, Guerreiro *et al.* 2022, Costinot and Werning, 2023, and Thuemmel, 2023). We also view them as a building block in studying the macroeconomic effects of automation (e.g., Moll *et al.* 2022). While we have focused attention to firms that use robots, so as to shed light on the micro adjustment, it would be interesting to study what happens to other firms in the same industry. Robot adoption is likely to induce a reallocation away from non adopters, with further negative effects on employment. Estimating these industry-level effects seems an important avenue for future research.[12] We have also found that, while robot adoption increases labour productivity, its effect on sales is weaker. This suggests that the efficiency gains may be partly offset by an increase in markups. Since automation is prevalent among top firms, investigating its relationship with market power seems another important avenue for future research.

[12] See Acemoglu *et al.* (2020), Koch *et al.* (2021) and Hubmer and Restrepo (2022) for some evidence on this reallocation.

## References


[1] Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Krof, Z., Restrepo, P. and Zolas, N. (2021). "Technology and the size and composition of workforce in U.S. firms: First evidence from the 2019 Annual Business Survey," ASSA Meetings – January 5, 2021.

[2] Acemoglu, D., Gancia, G. and Zilibotti, F. (2015). "Offshoring and directed technical change," *American Economic Journal: Macroeconomics*, vol. 7(3), pp. 84–122.

[3] Acemoglu, D., Lelarge C. and Restrepo, P. (2020). "Competing with robots: Firm-level evidence from France," *AEA Papers and Proceedings*, vol. 110(May), pp. 383–88.

[4] Acemoglu, D. and Restrepo, P. (2018). "The race between man and machine: Implications of technology for growth, factor shares and employment," *American Economic Review*, vol. 108(6), pp. 1488–1542.

[5] Acemoglu, D. and Restrepo, P. (2020). "Robots and jobs: Evidence from the US labor markets," *Journal of Political Economy*, vol. 128(6), pp. 2188–2244.

[6] Acemoglu, D. and Restrepo, P. (2022). "Demographics and automation," *Review of Economic Studies*, vol. 89(1), pp. 1–44.

[7] Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. (2021). "What are the labor and product market effects of automation? New evidence from France," CEPR DP 14443.

[8] Autor, D. and Dorn, D. (2013). "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, vol. 103(5), pp. 1553–97.

[9] Autor, D., Levy, F. and Murnane, R.J. (2003). "The skill content of recent technological change: An empirical exploration," *Quarterly Journal of Economics*, vol. 118(4), pp. 1279–1333.

[10] Baldwin, R. (2019). *The Globotics Upheaval: Globalisation, Robotics and the Future of Work*, Oxford: Oxford University Press.

[11] Beraja, M. and Zorzi, N. (2022). "Inefficient automation," Working Paper, MIT.

[12] Bessen, J. (2019). "Automation and jobs: When technology boosts employment," *Economic Policy*, vol. 34(100), pp. 589–626.

[13] Bessen, J., Goos, M., Salomons, A. and van den Berge, W. (2023). "What happens to workers at firms that automate?" *Review of Economics and Statistics*, forthcoming.

[14] Blanas, S., Gancia, G, and Lee, S.Y. (2019). "Who is afraid of machines?" *Economic Policy*, vol. 34(100), pp. 627–90.

[15] Bonfiglioli, A., Crinò, R., Gancia, G. and Papadakis, I. (2023). "Robots, offshoring and welfare," in (Lili Yan Ing and Gene M. Grossman, eds.) *Robots and AI: a New Economic Era*, Oxon and New York: Routledge.

[16] Borusyak, K., Jaravel, X. and Spiess, J. (2022). "Revisiting event study designs: Robust and efficient estimation," Working Paper, University of California Berkeley, arXiv:2108.12419.

[17] Broda, C. and Weinstein, D. (2006). "Globalization and the gains from variety," *Quarterly Journal of Economics*, vol. 121(2), pp. 541–85.

[18] Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, New York: W.W. Norton.

[19] Cheng, H., Jia, R., Li, D. and Li, H. (2019). "The rise of robots in China," *Journal of Economic Perspectives*, vol. 33(2), pp. 71–88.

[20] Costinot, A. and Werning, I. (2023). "Robots, trade, and luddism: A sufficient statistic approach to optimal technology regulation," *Review of Economic Studies*, vol. 90(5), pp. 2261–91.

[21] Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. (2021). "The adjustment of labor markets to robots," *Journal of the European Economic Association*, vol. 19(6), pp. 3104-53.

[22] Dechezleprêtre, A., Hemous, D., Olsen, M. and Zanella, C. (2021). "Induced automation: Evidence from firm-level patent data," University of Zurich, Department of Economics, Working Paper No. 384.

[23] Dinlersoz, E. and Wolf, Z. (2018). "Automation, labor share, and productivity: Plant-level evidence from U.S. manufacturing," CES WP 18-39.

[24] Dixen, J., Hong, B. and Wu, L. (2020). "The employment consequences of robots: Firm-level evidence," Statistics, Canada, Analytical Studies Branch Research Paper Series 454.

[25] European Commission (2015). "Analysis of the impact of robotic systems on employment in the European Union," Technical Report.

[26] Frey, C. and Osborne, M. (2017). "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, vol. 114(January), pp. 254–80.

[27] Graetz, G. and Michaels, G. (2018). "Robots at work," *Review of Economics and Statistics*, vol. 100(5), pp. 753–68.

[28] Guerreiro, J., Rebelo, S. and Teles, P. (2022). "Should robots be taxed?" *Review of Economic Studies*, vol. 89(1), pp. 279–311.

[29] Hemous, D. and Olsen, M. (2022). "The rise of the machines: Automation, horizontal innovation and income inequality," *American Economic Journal: Macroeconomics*, vol. 14(1), pp. 179–223.

[30] Hubmer, J. and Restrepo, P. (2022). “Not a typical firm: The joint dynamics of firms, labor shares, and capital-labor substitution,” NBER WP 28579.

[31] Humlum, A. (2019). “Robot adoption and labor market dynamics,” Working Paper, University of Chicago.

[32] International Federation of Robotics (2022). “World robotics reports,” https://ifr.org/worldrobotics/ (last accessed: 15 April 2021)

[33] Koch, M., Manuylov, I. and Smolka, M. (2021). “Robots and firms,” *Economic Journal*, vol. 131(638), pp. 2553–84.

[34] Martinez, J. (2021). “Putty clay automation,” CEPR DP 16022.

[35] Moll, B., Rachel, L. and Restrepo, P. (2022). “Uneven growth: Automation’s impact on income and wealth inequality,” *Econometrica*, vol. 90(6), 2645–83.

[36] Thuemmel, U. (2023). “Optimal taxation of robots,” *Journal of the European Economic Association*, vol. 21(3): 1154–90.