Robot Imports and Firm-Level Outcomes*

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

This draft: July 2022

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 various empirical strategies 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 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), 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).

[†]Queen Mary University of London and CEPR, Mile End Road, London E1 4NS, UK. E-mail: a.bonfiglioli@qmul.ac.uk

[‡]University of Bergamo, Dept. of Economics, CEPR and CESifo. Via dei Caniana 2, 24127, Bergamo, ITALY. E-mail: rosario.crino@unibg.it.

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

[¶]Queen Mary University of London and CEPR. Mile End Road, London E1 4NS, UK. E-mail: g.gancia@qmul.ac.uk

1 Introduction

The widespread diffusion of industrial robots has fuelled growing concerns about the future of work. Robots are programmable machines that have 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 worringly 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 to identify 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 effect of automation on employment 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 build a measure of automation *intensity* that is independent of demand shocks and 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, more productive, and have a larger employment share of high-skill professions. Second, looking at the evolution over time of firm-level outcomes, we find that robot import

[1] By 2018, there were an estimated 2.44 million industrial robots performing a variety of tasks that humans used to do. 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).

occurs after periods of expansion in firm size, suggesting that adoption may be driven by demand shocks, but it is followed by a *decline* in employment. Third, we find that an increase in robot intensity, defined as the ratio between cumulated robot imports and the capital stock, is associated with a *fall* in employment. Fourth, using a new measure of exposure to automation based on pre-determined technological characteristics, we find that firms that are more prone to adopt robots experience a stronger reduction in employment 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 labor productivity 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 nascent 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, Dixen, Hong and Wu (2020) for Canada, and Acemoglu, Lelarge and Restrepo (2020) for France. Other papers use dummies from survey data. These include Koch, Manuylov and Smolka (2021) for Spain, Cheng et al. (2019) for China, Dinlersoz and Wolf (2018) for the U.S., and a study by the European Commission (2015) for 7 European countries. None of these papers uses exogenous variation in automation across firms and, as a result, they find positive correlations with employment.

We are aware of two papers that try to identify causal effects using firm-level data. The first is Aghion et al. (2020), who use the same French data as us, but proxy automation with investment in industrial equipment and electricity consumption. 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 that they consider, which is more likely to be complementary to labor. The second paper is Bessen et al. (2019), 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). Yet, 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 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, κ_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 φ_i and automation κ_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/κ_i ($l_i/(1-\kappa_i)$) is capital (workers) per task.

Firms are monopolistically competitive and choose capital, labor and automation so as to maximize profit:

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

where $hf_i(\kappa_i)$ is a fixed cost increasing in automation. This fixed cost is in units of a composite input, which may include managers, scientists and engineers, with price h. The first-order conditions for capital and labor 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, κ_i. Using (3)-(4) into

[2] Proofs are in Appendix A. The model adds firm heterogeneity to theories of automation such as Acemoglu and Autor (2011), Acemoglu and Restrepo (2018), Hemous and Olsen (2022), Aghion, Jones and Jones (2019), but also Acemoglu, Gancia and Zilibotti (2015). See Martinez (2021) for a model of automation embodied in capital goods generating a distribution of technologies.

(2) also shows that output per worker is increasing in κ_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, κ_i, has two opposite effects on the demand for labor. 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 κ_i raises production, which in turn increases the demand for labor. The derivative of l_i with respect to κ_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 σ and w/r are high, i.e., when the productivity effect is strong enough. If σ 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, κ_i. We assume that automating more tasks poses an increasingly difficult challenge. For tractability, we focus on the following functional form:

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

with $\rho_i \in (0,1)$. The parameter $1/\rho_i$ captures the rate at which the marginal cost of automation increases with κ_i.[4] Hence, we interpret ρ_i as an index of replaceability of tasks in the production process and we allow it to vary across firms.[5] The first-order condition for κ_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 φ_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 $1/\rho_i$.[6]

[3] Acemoglu and Restrepo (2018) allows 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] 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 κ_i the lower ρ_i is.

[5] For any given task automation cost, firms with a higher ρ_i have a larger share of tasks below that cost.

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

The model shows that the effect of κ_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 κ_i. Specifically, demand shocks trigger automation but also have a direct positive effect on labor demand. Yet, the model offers possible remedies to this bias. A first possibility is to use a measure of *automation intensity* independent of demand shocks. Using (3) into (6), we can write:

$$\frac{\kappa_i}{rk_i} = \frac{1}{hf_i'(\kappa_i)} \ln\left(\frac{w}{r}\right). \tag{7}$$

This equation shows that normalizing κ_i by capital expenditure isolates the costs and benefits of automation independent of size and demand. The reason is that shocks to demand raise both κ_i and k_i, leaving the ratio unchanged. Yet, automation intensity still depends on variables that could affect employment directly. A second possibility is to focus on exogenous characteristics that exclusively affect the automation choice. In the model, a firm-level parameter that has no effect other than through κ_i is replaceability, ρ_i. Eq. (6) shows that firm-level replaceability interacts with the industry-level characteristics capturing the suitability of the production process to automation as in (7). Based on this insight, we will build a measure of exposure to robots by combining information on which industries are more suitable to automation with firm-level measures of replaceability of employment.

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.[7] We complement the balance sheet data with information on the occupational structure of employment in each firm from DADS Etablissement. For each sample year, this dataset contains employment data disaggregated into five two-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.[8] For the

[7]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.

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

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. We thus have a proxy not only for whether a given firm adopts robots or not but also for the intensity with which it uses robots.

Robot imports are recognized as a good proxy for automation because of the high concentration of this sector.[9] 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 measurement issues. 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.[10] 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. We also restrict the analysis to firms with more than ten employees, for which the reporting threshold is 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 C1 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 Ve-

[9] See, for instance, Acemoglu and Restrepo (2022), Blanas, Gancia and Lee (2019) and Bonfiglioli et al. (2022).

[10] The threshold equals 40,000 Euros before 2002; 100,000 Euros over 2002-2006; 150,000 Euros over 2007-2010; and 460,000 Euros from 2011 onwards. These thresholds are not very high given the average price of a robot.

hicles" industry because our data lack information for the two biggest car manufacturers in France.[11] 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 64,173 manufacturing firms, of which 765 have imported robots at least once over 1994-2013 (henceforth, "robot adopters"). This number is consistent with Acemoglu, Lelarge and Restrepo (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. Between 1994 and 2013, the shares of robot adopters in manufacturing employment and value added have increased steadily to reach 8% and 14%, respectively.[12] 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 accompanied by the adoption of labor-saving technologies.[13]

Appendix Table C1 reports summary statistics separately for robot adopters and non adopters, showing that the former firms are systematically larger, more productive, 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 on value 1 in the first year in which the firm imports robots and in all subsequent periods, and is equal to 0 otherwise. We control for (i) firm fixed effects, α_i, to absorb time-invariant firm characteristics; (ii) 5-digit industry$\times$year fixed effects, α_{jt}, to account for differences in the industry of operation and for industry-specific shocks; and (iii) firm characteristics—namely, log sales and dummies for firms that export or import goods other than robots—measured

[11] 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.

[12] These figures refer to a consistent sample of firms that are active in all years and import robots at least once over 1994-2013.

[13] 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, Robot Adoption and Robot Intensity

	(1)	(2)	(3)	(4)
	Ln Sales	Ln No. of Employees	Ln VA per Worker	Empl. Sh. High Skill
		a) Robot Adoption		
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
		b) Robot Intensity		
Ln RobInt_{it}	-0.129***	-0.144***	0.040***	0.015***
	[-4.150]	[-5.427]	[2.654]	[2.815]
Obs.	5,706	5,711	5,542	5,711
R2	0.97	0.93	0.84	0.89

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. *Adoption* $_{it}$ is a dummy equal to 1 for all years since the firm starts importing robots, and equal to 0 otherwise. *Ln RobInt* $_{it}$ is the log ratio between the stock of robot capital and the total capital stock of the firm. 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.

in the first year in which the firm is observed and interacted with a full set of year dummies, $\mathbf{X}_{it}$. These interactions flexibly control for heterogeneous trends across firms characterized by different initial conditions. We estimate eq. (8) for four major outcomes 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 (scientists, managers, and engineers).[14] The results are reported in panel a) of Table 1; standard errors are corrected for clustering at the firm level and t-statistics are shown in square brackets. All estimates of β 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 analyze how the four outcomes evolve over time in firms that adopt robots relative to the rest. To this purpose, we extend eq. (8) by adding the first five lags

[14] We focus on these outcomes because they can be constructed directly from the data.

and leads of $Adoption_{it}$:

$$Y_{it} = \alpha_i + \alpha_{jt} + \sum_{s=-5}^{5} \beta_s \cdot Adoption_{it-s} + \varepsilon_{it}. \tag{9}$$

The coefficients β_s illustrate how a given outcome evolves over time within robot adopters relative to non adopters, over a ten-year window around the first instance of robot imports ($s = 0$).

The results are displayed in Figure 1. The estimation coefficients corresponding to each plot are reported in Appendix Table C2. The figure shows that robot adoption is antedated by significant differences in the trends of sales and employment between robot adopters and non adopters. In particular, the former group of 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 efficiency 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 efficiency. These results suggest that robot adoption occurs after periods of expansion in firm size, and is followed by employment losses, improvements in efficiency, and labor demand shifts towards high-skill workers, with limited changes in total sales.

The above evidence suggests that the correlations between robot adoption and other firm characteristics may be confounded by demand shocks. To provide additional evidence on this hypothesis, we re-estimate eq. (8) using a continuous measure of robot intensity, $\ln RobInt_{it}$, defined as the log ratio between the stock of robot capital and the total capital stock of the firm. This variable may proxy for the theoretical measure introduced in eq. (7) as long as robot capital is proportional to κ. By scaling robot capital with the total capital stock of the firm, $\ln RobInt_{it}$ neutralizes demand shocks affecting both the numerator and the denominator of the ratio. The log transformation implies that $\ln RobInt_{it}$ is only defined for robot adopters. Because the specification controls for firm and industry$\times$year fixed effects, the coefficients β are identified from changes in robot intensity over time within robot adopters, controlling for common shocks hitting all firms in an industry.

The results are reported in panel b) of Table 1. Compared to panel a), the estimate of β switches sign, from positive to negative, in the regressions for sales and employment, and is highly statistically significant. This pattern is consistent with demand shocks leading

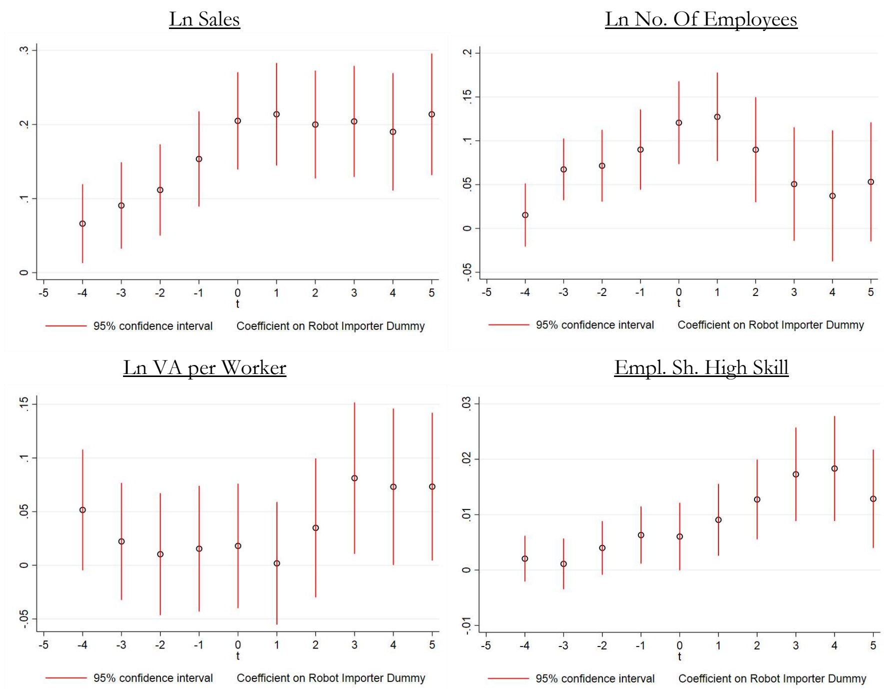

Each graph plots coefficients and confidence intervals on various lags and leads of $Adoption_{it}$ estimated using eq. (9) for a different outcome variable (indicated in the heading of the graph). $Adoption_{it}$ is a dummy that takes on value 1 in the first year in which a firm imports robots and in all subsequent periods, and is equal to 0 otherwise. Lags and leads of $Adoption_{it}$ are indicated on the horizonthal axis of each graph, with *t=0* referring to the first year in which a firm imports robots. The estimated coefficients corresponding to each graph are reported in Appendix Table C2.

Figure 1: Difference-in-Differences Event Studies

firms to both expand and automate, resulting in a spurious positive correlation between robot adoption and firm size. Once demand shocks are neutralized, however, automation may lead to job displacement. The negative effect on sales suggests that $\ln RobInt_{it}$ may be partly driven by increases in wages, which trigger automation but also raise production costs. The estimates of β for the other outcomes remain positive, suggesting that automation is associated with improvements in firm efficiency and labor demand shifts towards high-skill workers.[15]

4 The Effect of Robot Exposure

We now exploit differential cross-firm variation in robot exposure stemming from pre-determined technological characteristics to identify a causal effect of automation on firm-level outcomes.

[15] As shown in Appendix Table C3, these results are very similar if we compute the stock of robot capital with an annual depreciation rate of 15%.

4.1 Variables and Specification

Our model suggests that the predisposition of firms to automate depends on the interplay between cost of machines and replaceability of employment. In particular, a lower cost of machines stimulates robot adoption relatively more at firms whose employment can more easily be replaced by robots. It is well known that differences in the production process make some industries more suitable to automation than others, implying a larger cost-advantage of machines in these industries. At the same time, within any given industry, some firms are more prone to automate production than others, because they perform activities that are relatively easier to assign to robots. Building on these insights, our robot exposure measure, $RobExp$, exploits the interplay between a proxy for automation suitability in a given industry, $RobSuit$, and a proxy for replaceability of worker activities by robots within a given firm, $Repl$. Using this measure, we study how robot exposure affects outcomes and adoption decisions at the firm level.

In a given 5-digit NACE industry j, $RobSuit$ is defined 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, respectively, the initial stock of robots and the initial total capital stock of firm i'. The hyperbolic sine transformation preserves the zeros. Industries in which this ratio is higher are relatively more suitable to automation. Our measure of replaceability, $Repl$, is similar to the indicator proposed by Graetz and Michaels (2018) but is defined at the firm-level rather than at the industry level. We source from Graetz and Michaels (2018) information on whether each of 377 U.S. Census occupations is replaceable or not, where 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. Census occupation into the 29 French occupations for which we have employment data in 1994 and construct firm-level replaceability as follows:

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

where $Repl_o$ is the replaceability of French occupation o and ω_{oi} is the share of occupation

o in firm i's employment in 1994. We finally obtain $RobExp$ as[16]

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

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 ΔY_i is the annualized change in outcome Y for firm i between the first and the last year in which the firm is present in the sample; α_j are 5-digit industry fixed effects; $\mathbf{X}_i$ are start-of-period values of log firm sales and of indicators for exporting and importing firms; and ε_i is an error term. 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. The industry fixed effects absorb differential trends in outcomes across industries, and the covariates remove heterogeneous trends across firms with different initial conditions within the same industry. In particular, these variables account for the fact that larger and more trade-oriented firms may be more exposed to robots and may systematically follow different paths in terms of key outcomes compared to other firms. We correct the standard errors for clustering within 5-digit industries to account for possibly correlated shocks among firms in the same industry.[17]

We believe that neither $Repl_i$ nor $RobSuit_{j-i}$ alone is sufficient to capture a firm's exposure to robots. In particular, 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 the spirit of a difference-in-differences specification. Moreover, 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 the outcomes of interest. Being identified by both firm- and industry-level variation, the interaction coefficients β_1 are less likely to be confounded by omitted firm or industry characteristics than the linear terms.

[16] See Appendix Table C4 for descriptive statistics on $RobExp_i$ and the other variables used in this section.

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

4.2 Baseline Results and Sensitivity Analysis

The baseline estimates of β_1 are shown in panel a) of Table 2, where the change in each outcome is multiplied by 100 to express it in percentages.[18] The coefficient on $RobExp_i$ is negative and precisely estimated in the employment regression, indicating that firms that are more exposed to robots experience a relatively larger and statistically significant reduction in employment over the sample period. As for the other outcomes, β_1 is positive and precisely estimated in the regression for value added per worker, implying that higher robot exposure induces larger efficiency gains within firms. The effect of robot exposure on sales, while positive, is not statistically significant. While reinforcing the view that demand shocks bias the relation between robot adoption and firm size, this result also suggests that productivity gains may not always translate into lower prices. Finally, the results point towards a positive, albeit imprecisely estimated, effect of robot exposure on the skill structure of employment.

Our working hypothesis is that robot exposure affects outcomes by inducing firms to adopt robots. To study this mechanism, in column (5), we estimate eq. (13) with a different dependent variable, namely, a dummy equal to 1 for firms that start importing robots over the sample period, $Adopter_i$. The coefficient on $RobExp_i$ is positive and very precisely estimated, implying that firms that are more exposed to robots do indeed show a greater tendency to adopt robots in subsequent years. The point estimates imply that a change in automation suitability equal to the interquartile range of its distribution (12.12) is associated with a 67% higher probability of adoption and 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 have a 56% higher adoption probability and 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 48% higher adoption probability and a 0.26 p.p. per year employment fall at the 75th percentile of the replaceability distribution relative to the 25th percentile.[19]

[18] In Appendix Table C5, we estimate eq. (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 efficiency, and a relatively faster shift in labor demand towards high-skill workers.

[19] We refrain from interpreting our estimates in a Two-Stage Least Squares (2SLS) framework. We believe that, due to the aforementioned limitations of the import data, $Adopter_i$ provides a qualitative measure of robot adoption but does not accurately capture all instances of adoption and the intensity of automation. This complicates the quantitative interpretation of the estimate in column (5) as a first-stage coefficient.

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	0.069
			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	0.052
			d) Broader Definition of Robot Imports		
$RobExp_i$	0.127	-0.038	0.187*	0.010**	0.261*
	[1.314]	[-0.810]	[1.768]	[2.096]	[1.830]
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 100 x the annualized changes in the firm-level outcomes indicated in columns' headings. In columns (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 employment share of occupations that can be replaced 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 84864. 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. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

The remaining panels of Table 2 contain an extensive sensitivity analysis. In panel b), we weigh the observations by the initial number of employees in each firm. The estimated coefficients are similar to those obtained from unweighted regressions. In panel c), we further exclude firms in the "Manufacture of Motor Vehicles" industry. The qualitative and quantitative pattern of results is unchanged. In panel d), we extend the definition of automation suitability to include not just industrial robots but 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 labor demand towards high-skill professions. These results are consistent with the notion that broader forms of capital equipment are more complementary to labor, as found in Aghion et al. (2019), especially to high-skill workers.

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 (see also Bessen, 2019). We therefore extend eq. (13) by interacting $RobExp_i$, $Repl_i$, and $RobSuit_{j-i}$ with the elasticity of substitution in each sector sourced from Broda and Weinstein (2006). The results in panel e) confirm that robot exposure causes a relatively larger increase in sales in sectors where products are more substitutable. Similarly, robot exposure has a relatively less negative effect on employment in sectors where demand is more elastic, although the interaction coefficient is marginally insignificant. The effect of robot exposure on productivity is also relatively stronger in sectors where products are more substitutable.

Finally, we use an alternative proxy for robot exposure, which is obtained by replacing $RobSuit_{j-i}$ with the initial value of the log stock of installed robots in each firm's sector in the U.S. from the IFR. This variable may be a better proxy for automation suitability in industries where firms predominantly source robots from domestic suppliers. Being based on data for the U.S. rather than France, it also further allays concerns with endogeneity. However, the IFR data are only available for 13 aggregate manufacturing sectors, so variation is much more limited. The qualitative evidence is similar with this alternative proxy for robot exposure. In particular, the coefficient on $RobExp_i$ in the employment regression is not statistically significant, consistent with an upward bias in the correlation between robot adoption and employment. The alternative proxy also confirms the positive effects of robot exposure on robot adoption, efficiency, skill composition of labor demand and now even sales.

4.3 Threats to Identification

Our identification strategy requires that, conditional on the fixed effects and control variables included in eq. (13), $RobExp_i$ is uncorrelated with omitted variables that could influence the outcomes. Because $RobExp_i$ is the interaction between $Repl_i$ and $RobSuit_{j-i}$, this identifying assumption could be violated in two cases: (i) if $Repl_i$ was correlated with other firm characteristics that affect outcomes differentially across industries with varying levels of automation suitability; and (ii) if $RobSuit_{j-i}$ captured other industry characteristics affecting outcomes heterogeneously across firms with different degrees of replaceability. To address these concerns, we extend the specification by adding interactions of $Repl_i$ and $RobSuit_{j-i}$ with some of the most likely confounders, and study how the coefficients on $RobExp_i$ are affected.

In panel a) of Table 3, we add the interaction between $RobSuit_{j-i}$ and the routine intensity of each firm.[20] While routine intensity is known to be correlated with the adoption of skill-intensive technologies such as computers (e.g., Autor, Levy and Murnane, 2003), Cheng et al. (2019) find that robots are more prevalent at firms where employees are commonly doing manual tasks rather than routine tasks. Consistently, we find the new interaction to have no significant effect on robot adoption, and its inclusion to leave the evidence on $RobExp_i$ unaffected. In panel b), we extend the specification by adding interactions between $RobSuit_{j-i}$ and all the control variables included in $\mathbf{X}_i$. While larger and more trade-oriented firms could have different levels of replaceability, the main results are preserved. Similarly, panel c) shows that the results are 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) the average unit values of imports and exports; and (iii) imports of capital and technology goods.[21] Controlling for these interactions leaves the coefficients on $RobExp_i$ close to the baseline estimates. Finally, we add interactions between $Repl_i$ and a full set of 2-digit sector dummies. Contributing to the identification of β_1 is now only the remaining variation in $RobSuit_{j-i}$ across narrow (5-digit) industries within the same 2-digit sector. As

[20] This variable is defined as the share of routine-intensive occupations in the firm's total employment in 1994. Data on routine intensity by occupation are sourced from Autor and Dorn (2013) and matched to the 29 French occupations in our data. The linear term in routine intensity is included in the specification but untabulated.

[21] Similar to $RobSuit_{j-i}$, we construct each of these variables in the initial year by aggregating across firms other than i in each 5-digit industry. Each of these characteristics also enters the specification linearly.

Table 3: Firm-Level Outcomes and Robot Exposure, Threats to Identification

	(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]
$RobSuit_{j-i}$ 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				
$RobExp_i$	0.137	-0.091**	0.283**	0.007	0.167***
	[1.260]	[-2.006]	[2.564]	[1.174]	[2.730]
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				
$RobExp_i$	0.148	-0.095**	0.304**	0.008	0.169***
	[1.245]	[-2.070]	[2.534]	[1.286]	[2.906]
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				
$RobExp_i$	0.173	-0.137***	0.399***	0.012*	0.136**
	[1.591]	[-2.771]	[3.969]	[1.795]	[2.252]
Obs.	36,254	36,537	35,134	36,537	36,903
R2	0.10	0.04	0.07	0.04	0.05
	e) Interactions of Replaceability with Sector Dummies				
$RobExp_i$	0.155	-0.089*	0.306***	0.006	0.175***
	[1.353]	[-1.934]	[2.620]	[0.968]	[2.853]
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 100 x the annualized changes in the firm-level outcomes indicated in columns' headings. In columns (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 employment share of occupations that can be replaced 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 ($RobSuit_{j-i}$). 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 $RobSuit_{j-i}$ 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 $Repl_i$. 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 $RobSuit_{j-i}$, 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. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

shown in panel e), our main results are qualitatively and quantitatively unchanged also in this case.

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. 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.[22] We also view our results as a building block for a comprehensive study of the macroeconomic effects of automation.[23] While we have focused attention to firms that import 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. Since robot adoption is likely to induce a reallocation away from non adopters, it is likely to have further negative effects on employment. Estimating these industry-level effects seems an important avenue for future research.[24] We have also found that, while robot adoption increases productivity, its effect on sales is less strong. 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.

[22] See, for instance, Beraja and Zorzi (2022), Costinot and Werning (2018), Thuemmel (2018) and Guerreiro, Rebelo, and Teles (2022).

[23] See, for instance, Jaimovich et al. (2022) and Moll, Rachel and Restrepo (2022).

[24] See Hubmer and Restrepo (2022), Acemoglu, Lelarge and Restrepo (2020) and Koch, Manuylov and Smolka (2021) for some evidence on this reallocation.

References

[1] Acemoglu, Daron, Gary Anderson, David Beede, Catherine Buffington, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Krof, Pascual Restrepo and Nikolas Zolas (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, Daron and David Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, Volume 4B, Chapter 12, 1043–1171, Elsevier.

[3] Acemoglu, Daron, Gino Gancia and Fabrizio Zilibotti (2015). "Offshoring and Directed Technical Change," *American Economic Journal: Macroeconomics*, 7: 84–122.

[4] Acemoglu, Daron, Claire Lelarge and Pascual Restrepo (2020). "Competing with Robots: Firm-Level Evidence from France," *AEA Papers and Proceedings*, 110: 383–388.

[5] Acemoglu, Daron and Pascual Restrepo (2018). "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment," *American Economic Review*, 108: 1488–1542.

[6] Acemoglu, Daron and Pascual Restrepo (2020). "Robots and Jobs: Evidence from the US Labor Markets," *Journal of Political Economy*, 128: 2188–2244.

[7] Acemoglu, Daron and Pascual Restrepo (2022). "Demographics and Automation," *Review of Economic Studies*, 89: 1–44.

[8] Aghion, Philippe, Celine Antonin, Simon Bunel and Xavier Jaravel (2019). "What Is the Impact of Automation on Employment? New Evidence from France," CEPR DP 14443.

[9] Aghion, Philippe, Benjamin Jones and Charles Jones (2019). "Artificial Intelligence and Economic Growth," *The Economics of Artificial Intelligence: An Agenda* (Ajay Agrawal, Joshua Gans and Avi Goldfarb, eds.), Chicago: University of Chicago Press.

[10] Autor, David and David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103: 1553–1597.

[11] Autor, David, Frank Levy and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118: 1279–1333.

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

[13] Beraja, Martin and Nathan Zorzi (2022). "Inefficient Automation," Mimeo, MIT.

[14] Bessen, James (2019). "Automation and Jobs: When Technology Boosts Employment," *Economic Policy*, 34: 589–626.

[15] Bessen, James, Maarten Goos, Anna Salomons and Wiljan van den Berge (2019). "What Happens to Workers at Firms that Automate?" Boston University School of Law, Law and Economics Research Paper 19-2.

[16] Blanas, Sotiris, Gino Gancia, and Sang Yoon (Tim) Lee (2019). "Who is Afraid of Machines?" *Economic Policy*, 34: 627–690.

[17] Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia and Ioannis Papadakis (2022). "Robots, Offshoring and Welfare," *Robots and AI: a New Economic Era* (Lili Yan Ing and Gene M. Grossman, eds.), Oxon and New York: Routledge, forthcoming.

[18] Broda, Christian and David Weinstein (2006). "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121: 541–585.

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

[20] Cheng, Hong, Ruixue Jia, Dandan Li and Hongbin Li (2019). "The Rise of Robots in China," *Journal of Economic Perspectives*, 33: 71–88.

[21] Costinot, Arnaud and Ivan Werning (2018). "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," NBER WP 25103.

[22] Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum and Nicole Woessner (2021). "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 19: 3104-3153.

[23] Dinlersoz, Emin and Zoltan Wolf (2018). "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing," CES WP 18-39.

[24] Dixen, Jay, Bryan Hong and Lynn Wu (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, Carl and Michael Osborne (2017). "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change*, 114: 254–280.

[27] Graetz, Georg and Guy Michaels (2018). "Robots at Work," *Review of Economics and Statistics*,100: 753–68.

[28] Guerreiro, Joao, Sergio Rebelo and Pedro Teles (2022). "Should Robots Be Taxed?" *Review of Economic Studies*, 89: 279–311.

[29] Hemous, David and Morten Olsen (2022). "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," *American Economic Journal: Macroeconomics*, 14: 179–223.

[30] Hubmer, Joachim and Pascual Restrepo (2022). "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital-Labor Substitution," NBER WP 28579.

[31] Humlum, Anders (2019). "Robot Adoption and Labor Market Dynamics," Mimeo, University of Chicago.

[32] Jaimovich, Nir, Itay Saporta Eckstein, Yaniv Yedid-Levy and Henry Siu (2022). "The Macroeconomics of Automation: Data, Theory, and Policy Analysis," *Journal of Monetary Economics*, forthcoming.

[33] Koch, Michael, Ilya Manuylov and Marcel Smolka (2021). "Robots and Firms," *Economic Journal*, 131: 2553–2584.

[34] Martinez, Joseba (2021). "Putty Clay Automation," CEPR DP 16022.

[35] Moll, Benjamin, Lukasz Rachel and Pascual Restrepo (2022). "Uneven Growth: Automation's Impact on Income and Wealth Inequality," NBER WP 28440.

[36] Thuemmel, Uwe (2018). "Optimal Taxation of Robots," CESIFO WP 7317.

For Online Publication

A Model: Derivations and Proofs

Using equation (5) into (4) yields:

$$l_i = w^{-\sigma}\left(1-\frac{1}{\sigma}\right)^{\sigma} A_i \varphi_i^{\sigma-1}\left(\frac{w}{r}\right)^{\kappa_i(\sigma-1)}(1-\kappa_i)\,. \tag{A1}$$

This equation shows how employment depends on κ_i and other exogenous parameters, and can be used to compute $\frac{d\ln l_i}{d\kappa_i}$.

Substituting (A1) in (5) we can express output as a function of automation and other exogenous parameters:

$$y_i = A_i \varphi_i^{\sigma} w^{-\sigma}\left(1-\frac{1}{\sigma}\right)^{\sigma}\left(\frac{w}{r}\right)^{\kappa_i\sigma}. \tag{A2}$$

This equation confirms that automation raises output as long as capital is cheaper than production workers:

$$\frac{d\ln y_i}{d\kappa_i} = \sigma\ln\left(\frac{w}{r}\right), \tag{A3}$$

and it illustrates that the productivty effect is stronger in industries where demand is more elastic (σ).

Consider now the choice of automation. Firms choose the level of κ_i that maximizes profit:

$$\max_{\kappa_i}\left\{\frac{p_i y_i}{\sigma} - h f_i(\kappa_i)\right\}.$$

The first-order condition for κ_i is:

$$\left(1-\frac{1}{\sigma}\right) p_i y_i \ln\left(\frac{w}{r}\right) = h f_i'(\kappa_i)\,.$$

Using $y_i = A_i p_i^{-\sigma}$, (A2) and $h f_i'(\kappa_i) = h(1-\kappa_i)^{-1/\rho_i}$ yields:

$$\left(1-\frac{1}{\sigma}\right)^{\sigma} A_i\left(\frac{\varphi_i}{w}\right)^{(\sigma-1)}\left(\frac{w}{r}\right)^{\kappa_i(\sigma-1)}\ln\left(\frac{w}{r}\right) = h(1-\kappa_i)^{-1/\rho_i}\,.$$

This expression shows the exogenous determinants of the marginal benefit of automation and can be used to solve implicitly for the equilibrium level of κ_i. The second-order condition is necessarily satisfied if $(\sigma-1)\ln(w/r) < 1/\rho_i$ and the unique solution is interior if:

$$\left(1-\frac{1}{\sigma}\right)^{\sigma} A_i\left(\frac{\varphi_i}{w}\right)^{(\sigma-1)}\ln\left(\frac{w}{r}\right) > h.$$

We assume both conditions to be satisfied. Denote the marginal benefit and the marginal cost of automation as MB_i and MC_i, respectively. Then:

$$\begin{aligned}
\frac{\partial MB_i}{\partial \kappa_i} &= MB_i \times (\sigma - 1) \ln \left(\frac{w}{r} \right) \\
\frac{\partial MC_i}{\partial \kappa_i} &= \frac{MC_i}{\rho_i (1 - \kappa_i)}.
\end{aligned}$$

Under the assumption $(\sigma - 1) \ln \left(\frac{w}{r} \right) < 1/\rho_i$, profits are concave in κ_i so that:

$$\frac{\partial MB_i}{\partial \kappa_i} < \frac{\partial MC_i}{\partial \kappa_i}.$$

We now derive the comparative statics for the optimal level of automation, κ_i^*, with respect to the primitives of the model and prove that:

$$\frac{d\kappa_i^*}{dA_i} > 0; \quad \frac{d\kappa_i^*}{d\varphi_i} > 0; \quad \frac{d\kappa_i^*}{d(w/r)} > 0; \quad \frac{d\kappa_i^*}{d\rho_i} > 0 \quad ; \quad \frac{d\kappa_i^*}{dh} < 0.$$

Differentiating the first-order condition, we obtain the implicit derivative of κ_i^* with respect to any parameter v as

$$\frac{d\kappa_i^*}{dv} = \frac{\frac{\partial MC}{\partial v} - \frac{\partial MB}{\partial v}}{\frac{\partial MB}{\partial \kappa_i} - \frac{\partial MC}{\partial \kappa_i}}.$$

The denominator is always negative. Hence, to find the sign of the derivatives of interest, we just need to compute the numerator of the expression above for A_i, φ_i, (w/r), ρ_i and h as follows:

$$\begin{aligned}
\frac{\partial MC}{\partial A_i} - \frac{\partial MB}{\partial A_i} &= -\frac{MB}{A_i} < 0 \to \frac{d\kappa_i^*}{dA_i} > 0 \\
\frac{\partial MC}{\partial \varphi_i} - \frac{\partial MB}{\partial \varphi_i} &= -(\sigma - 1) \frac{MB}{\varphi_i} < 0 \to \frac{d\kappa_i^*}{d\varphi_i} > 0 \\
\frac{\partial MC}{\partial (w/r)} - \frac{\partial MB}{\partial (w/r)} &= -\frac{MB}{(w/r)} \left[\kappa_i (\sigma - 1) + \frac{1}{\ln (w/r)} \right] < 0 \to \frac{d\kappa_i^*}{d(w/r)} > 0 \\
\frac{\partial MC}{\partial \rho_i} - \frac{\partial MB}{\partial \rho_i} &= \frac{h}{\rho_i^2} \frac{\ln (1 - \kappa_i)}{(1 - \kappa_i)^{1/\rho_i}} < 0 \to \frac{d\kappa_i^*}{d\rho_i} > 0 \\
\frac{\partial MC}{\partial h} - \frac{\partial MB}{\partial h} &= \frac{MC}{h} > 0 \to \frac{d\kappa_i^*}{d\rho_i} < 0.
\end{aligned}$$

B Discrete Choice of Automation

We now consider the case in which firm i can choose whether to keep the current level of automation κ_0 at no additional cost or increase it to $\kappa_1 > \kappa_0$, subject to the cost $\frac{h\kappa_1}{\rho_i}$. The discrete choice problem facing firm i is

$$\max_{\kappa_i \in \{\kappa_0, \kappa_1\}} \left\{ \frac{p_i(\kappa_i)\, y_i(\kappa_i)}{\sigma} - h f_i(\kappa_i) \right\}.$$

The condition for i to choose κ_1 is

$$\frac{p_i(\kappa_1)\, y_i(\kappa_1) - p_i(\kappa_0)\, y_i(\kappa_0)}{\sigma} > \frac{h\kappa_1}{\rho_i},$$

which, after using $y_i = A_i p_i^{-\sigma}$ and (A2), becomes

$$\frac{A_i}{\sigma} \left[\varphi^{\sigma} w^{-\sigma} \left(1 - \frac{1}{\sigma}\right)^{\sigma} \right]^{1-1/\sigma} \left[\left(\frac{w}{r}\right)^{\kappa_1 \sigma} - \left(\frac{w}{r}\right)^{\kappa_0 \sigma} \right]^{1-1/\sigma} > \frac{h\kappa_1}{\rho_i}.$$

The left-hand side captures the benefit of further automation, while the right-hand side corresponds to its cost.

In this case, we can express the comparative statics in terms of the probability that an increase in any parameter induces a switch from κ_0 to κ_1. In particular, we are interested in the effect of an increase in (w/r) and its interaction with A_i, φ_i and ρ_i. It is easy to show that the left-hand side, denoted by B_i, is increasing in (w/r):

$$\frac{\partial B_i}{\partial \left(\frac{w}{r}\right)} = \frac{(\sigma - 1)\, A_i}{\sigma} \left[\varphi_i^{\sigma} w^{-\sigma} \left(1 - \frac{1}{\sigma}\right)^{\sigma} \right]^{1-1/\sigma} \frac{\left[\kappa_1 \left(\frac{w}{r}\right)^{\kappa_1 \sigma - 1} - \kappa_0 \left(\frac{w}{r}\right)^{\kappa_0 \sigma - 1} \right]}{\left[\left(\frac{w}{r}\right)^{\kappa_1 \sigma} - \left(\frac{w}{r}\right)^{\kappa_0 \sigma} \right]^{1/\sigma}} > 0.$$

This means that increasing automation is more likely to be optimal for lower relative cost of capital (r/w).

To characterize the interaction with A_i and φ_i, we compute the cross derivatives of B_i,

$$\begin{aligned} \frac{\partial^2 B_i}{\partial \left(\frac{w}{r}\right) \partial A_i} &= \frac{\partial B_i}{\partial \left(\frac{w}{r}\right)} A_i^{-1} > 0, \\ \frac{\partial^2 B_i}{\partial \left(\frac{w}{r}\right) \partial \varphi_i} &= \frac{\partial B_i}{\partial \left(\frac{w}{r}\right)} \sigma \varphi_i^{-1} > 0, \end{aligned}$$

which imply that the likelihood of further automation increases more with (w/r) for larger and more productive firms.

The derivative of the automation cost with respect to ρ_i,

$$\frac{\partial}{\partial \rho_i} \left(\frac{h\kappa_1}{\rho_i} \right) = -\frac{h\kappa_1}{\rho_i^2} < 0,$$

suggests that an increase in (w/r) increases more the likelihood of further automation for

Table C1: Descriptive Statistics, Whole Sample

	Obs.	No. Firms	Mean	Median	Std. Dev.	Mean Δ (annualized)
Robot Adopters						
Adoption	6,373	765	1	1	1	0
Robot Intensity	6,373	765	0.078	0.005	0.520	0.182
No. of Employees	6,373	765	852	191	3,129	-0.016
Empl. Sh. High Skill	6,373	765	0.153	0.108	0.142	0.006
Sales (€'000)	6,373	765	761,597	46,050	6,812,860	-0.075
VA per Worker (€'000)	6,225	761	178	65	2,715	-0.070
Dummy Importer	6,373	765	0.972	1	0.164	0.001
Dummy Exporter	6,373	765	0.947	1	0.224	0.002
Non Robot Adopters						
Adoption	598,925	63,408	0	0	0	0
Robot Intensity	586,785	63,448	0	0	0	0
No. of Employees	598,925	63,448	78	27	313	-0.030
Empl. Sh. High Skill	598,925	63,448	0.081	0.056	0.106	0.003
Sales (€'000)	598,922	63,448	54,703	7,615	683,130	-0.092
VA per Worker (€'000)	587,342	62,741	190	71	1,973	-0.066
Dummy Importer	598,925	63,448	0.568	1	0.495	0.001
Dummy Exporter	598,925	63,448	0.561	1	0.496	0.004

The whole sample consists of all manufacturing firms with more than 10 employees excluding firms in the "Installation and Repair of Machinery and Equipment" industry (64,173 firms). *Adoption* is a dummy taking on value 1 since the first year in which a firm imports robots. *Robot Intensity* is the ratio between the stock of robot capital and the total capital stock of the firm; the stock of robot capital is constructed as the sum of robot imports over time. *Importer* and *Exporter* are dummies taking on value 1 if the firm imports (resp. exports) goods other than robots in a given year and 0 otherwise. All statistics are computed on firm-level observations for the 1994-2013 period. Changes are computed as annualized log differences, except for *Employment Sh. High Skill*, *Exporter* and *Importer* dummies, for which annualized changes in levels are reported.

firms with higher replaceability ρ_i, since these face a lower cost.

C Additional Empirical Results

Table C1 contains summary statistics on the firm-level variables, separately for firms that import robots at least once over 1994-2013 ("robot adopters") and for firms that do not ("non robot adopters"). The statistics are computed on the whole sample, comprising 64,173 manufacturing firms, of which 765 are robot adopters. Robot intensity, defined as the ratio between the stock of robot capital and the total physical capital stock of the firm, equals 7.8% on average for robot adopters. The average robot adopter is around 11 times larger than the average non adopter in terms of employment and around 14 times larger in terms of sales. The skill composition of employment also differs across robot adopters and non adopters, with the share of employment in high-skill professions roughly twice as high on average in the former group of firms than in the latter. Robot adopters are also more likely to import and export goods other than robots.

Table C1 also reports the average annualized change in each variable over 1994-2013, separately for the two sets of firms. Robot adopters increased robot intensity at an average rate of 0.18 log points per year. While employment decreased in both groups of firms, robot

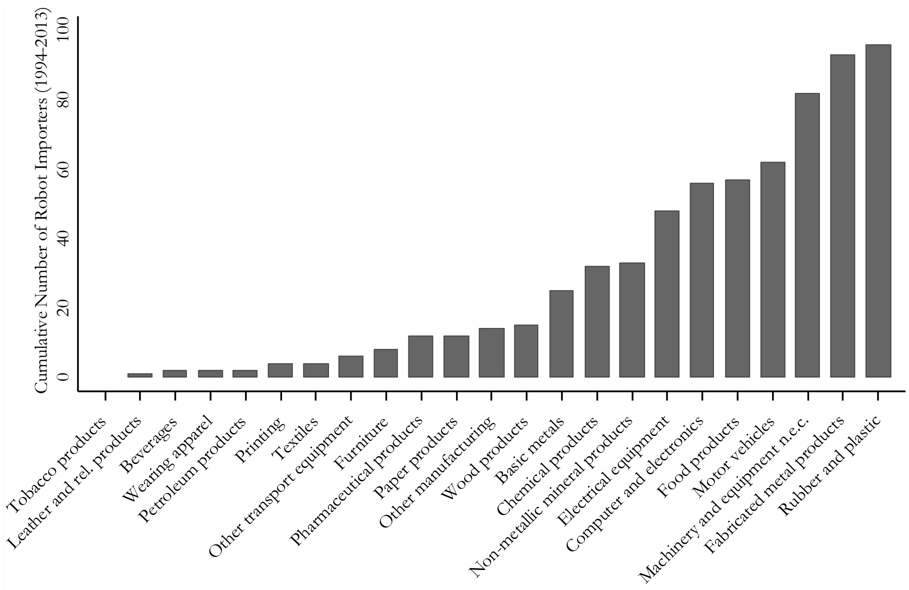

Figure C1: Cumulative Number of French Robot Importers by Sector (1994-2013)

adopters shed workers at a slower rate than non adopters (0.016 vs. 0.03 log points per year, respectively).[25] Robot adopters also experienced a relatively slower reduction in sales and a relatively faster increase in the employment share of high-skill professions.

As for the sectoral distribution of robot adopters, Figure C1 shows that these firms are present in all 2-digit manufacturing sectors but are particularly frequent in the production of motor vehicles, machinery, and electrical equipment.

Table C2 contains the estimation coefficients corresponding to the difference-in-differences event studies shown in Figure 1.

Table C3 replicates the conditional correlations between the outcomes and robot intensity, $\ln RobInt_{it}$, but it does so computing the stock of robot capital with a perpetual inventory method combined with an annual depreciation rate of 15%. The latter falls within the range of depreciation rates normally assumed for robots in manufacturing (see, e.g., Graetz and Michaels, 2018). The results are similar, both qualitatively and quantitatively, to the baseline estimates reported in panel b) of Table 1.

Table C4 contains descriptive statistics for the variables used in the long-differences specifications. Both replaceability and robot exposure are higher on average for robot adopters. Employment, sales and value added per worker have decreased less among robot adopters than among other firms, and robot adopters have experienced a relatively larger increase in the employment share of high-skill professions. As shown in Table C5, these differences persist even after controlling for the industry of operation of the firm and for differences in initial conditions. In particular, the table reports estimates of eq. (13), in which $RobExp_i$, $RobSuit_{j-i}$ and $Repl_i$ are replaced with a dummy, $Adopter_i$, which takes on value 1 if firm i starts importing robots over the sample period, and is equal to 0 both for non-adopters and for firms that were already using robots initially. The control variables are 5-digit industry

[25] Manufacturing employment significantly declined in France during the sample period.

Table C2: Difference-in-Differences Event Studies

	(1)	(2)	(3)	(4)
	Ln Sales	Ln No. of Employees	Ln VA per Worker	Empl. Sh. High Skill
$Adoption_{it-4}$	0.066**	0.015	0.052*	0.002
	[2.434]	[0.833]	[1.798]	[0.974]
$Adoption_{it-3}$	0.091***	0.067***	0.022	0.001
	[3.046]	[3.751]	[0.799]	[0.476]
$Adoption_{it-2}$	0.112***	0.072***	0.010	0.004
	[3.551]	[3.425]	[0.356]	[1.616]
$Adoption_{it-1}$	0.153***	0.090***	0.015	0.006**
	[4.679]	[3.851]	[0.517]	[2.399]
$Adoption_{it}$	0.205***	0.121***	0.018	0.006*
	[6.115]	[5.004]	[0.611]	[1.942]
$Adoption_{it+1}$	0.214***	0.127***	0.002	0.009***
	[6.056]	[4.939]	[0.063]	[2.739]
$Adoption_{it+2}$	0.200***	0.090***	0.035	0.013***
	[5.385]	[2.938]	[1.057]	[3.465]
$Adoption_{it+3}$	0.204***	0.051	0.081**	0.017***
	[5.339]	[1.530]	[2.255]	[4.008]
$Adoption_{it+4}$	0.190***	0.037	0.073**	0.018***
	[4.699]	[0.975]	[1.967]	[3.787]
$Adoption_{it+5}$	0.214***	0.053	0.073**	0.013***
	[5.104]	[1.533]	[2.087]	[2.837]
Obs.	689,846	593,312	581,715	593,312
R2	0.93	0.88	0.82	0.67

The subscripts *i* and *t* denote firms and years, respectively. The dependent variables are annual observations at time *t* of the firm-level outcomes indicated in columns' headings. $Adoption_{it}$ is a dummy that takes on value 1 in the first year in which a firm imports robots and in all subsequent periods, and is equal to 0 otherwise. All specifications include firm fixed effects and 5-digit industry x year fixed effect. Standard errors are robust to heteroskedasticity; t-statistics are reported in square brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

Table C3: Firm-Level Outcomes and Robot Intensity with Depreciation

	(1) Ln Sales	(2) Ln No. of Employees	(3) Ln VA per Worker	(4) Empl. Sh. High Skill
Ln $RobInt_{it}$	-0.085***	-0.098***	0.030**	0.011***
	[-3.458]	[-4.672]	[2.380]	[2.740]
Obs.	5,706	5,711	5,542	5,711
R2	0.97	0.93	0.84	0.89

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. *Ln* $RobInt_{it}$ is the log ratio between the stock of robot capital and the total capital stock of the firm; the stock of robot capital is constructed using a perpetual inventory method with a depreciation rate of 15%. 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.

fixed effects and initial values of log sales and of dummies for importing and exporting firms. The standard errors are corrected for clustering within 5-digit industries to account for possibly correlated shocks among firms in the same industry. The coefficients on $Adopter_i$ reflect cross-sectional differences in the growth of outcomes between robot adopters and other firms. The results show that firms that adopt robots over the sample period experience a relatively larger increase in size, a relatively stronger improvement in efficiency, and a relatively faster shift in labor demand towards high-skill workers.

Table C4: Descriptive Statistics, Sample Used for Specifications in Long Differences

	Robot Adopters			
	Obs.	Mean	Median	Std. Dev.
Δ Ln No. of Employees	497	-0.009	0.003	0.077
Δ Empl. Sh. High Skill	497	0.005	0.003	0.009
Δ Ln Sales	493	-0.093	-0.081	0.093
Δ Ln VA per Worker	470	-0.096	-0.094	0.095
Ln Initial Sales	497	11.778	11.644	1.768
Dummy Initial Importer	497	0.924	1.000	0.266
Dummy Initial Exporter	497	0.889	1.000	0.314
Replaceability	497	0.378	0.416	0.183
Robot Exposure	497	-5.872	-5.330	3.730
	Non Robot Adopters			
Δ Ln No. of Employees	36,087	-0.033	-0.012	0.095
Δ Empl. Sh. High Skill	36,087	0.003	0.001	0.011
Δ Ln Sales	35,808	-0.132	-0.108	0.131
Δ Ln VA per Worker	34,710	-0.104	-0.101	0.141
Ln Initial Sales	36,087	9.882	9.686	1.376
Dummy Initial Importer	36,087	0.550	1.000	0.498
Dummy Initial Exporter	36,087	0.519	1.000	0.500
Replaceability	36,087	0.358	0.360	0.190
Robot Exposure	36,087	-6.681	-5.946	4.300

The sample for the specifications in long differences consists of 36,584 manufacturing firms with more than 10 employees excluding firms in the "Installation and Repair of Machinery and Equipment" industry. Statistics are reported for the annualized changes in the outcomes and for the initial values of the *Importer* and *Exporter* dummies, *Ln Sales*, *Replaceability* and *Robot Exposure*. The latter variable is the product between the initial firm-level employment share of occupations that can be replaced by robots (*Replaceability*) and the initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry (*Robot Suitability*).

Table C5: Firm-Level Outcomes and Robot Adoption, Long Differences

	(1)	(2)	(3)	(4)
	Δ Ln Sales	Δ Ln No. of Employees	Δ Ln VA per Worker	Δ Empl. Sh. High Skill
$Adopter_i$	4.438***	2.434***	1.517***	0.007
	[11.032]	[7.775]	[3.609]	[0.155]
Obs.	36,301	36,584	35,180	36,584
R2	0.10	0.04	0.06	0.04

The subscript i denotes firms. In each regression, the dependent variable is 100 x the annualized change in the firm-level outcome indicated in the corresponding column. $Adopter_i$ is a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non importers. All specifications include 5-digit industry fixed effects as well as initial values of log sales and of dummies for importing and exporting firms. Standard errors are corrected for clustering within 5-digit industries; t-statistics are reported in square brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.