DISCUSSION PAPER SERIES

DP14593
(v. 3)

Robot Imports and Firm-Level Outcomes

Alessandra Bonfiglioli, Rosario Crinò, Harald Fadinger and Gino Gancia

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

MACROECONOMICS AND GROWTH

CEPR

ISSN 0265-8003

Robot Imports and Firm-Level Outcomes

Alessandra Bonfiglioli, Rosario Crinò, Harald Fadinger and Gino Gancia

Discussion Paper DP14593
First Published 09 April 2020
This Revision 16 November 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Trade and Regional Economics
- Macroeconomics and Growth

Robot Imports and Firm-Level Outcomes

Abstract

We use French data over the 1994-2013 period to study how imports of industrial robots affect firm-level outcomes. Compared to other firms operating in the same 5-digit industry, robot importers are larger, more productive, and employ a higher share of managers and engineers. Over time, robot import occurs after periods of expansion in firm size, and is followed by improvements in efficiency and a fall in demand for labor. 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, exogenous exposure to automation leads to job losses. We also find that robot exposure increases productivity and the employment share of high-skill professions, but has a weak effect on total sales. The latter result suggests that productivity gains from automation may not be entirely passed on to consumers in the form of lower prices.

JEL Classification: J23, J24, O33, D22

Keywords: automation, Displacement, firms, robots

Alessandra Bonfiglioli - a.bonfiglioli@qmul.ac.uk
Queen Mary University of London and CEPR

Rosario Crinò - rosario.crino@unibg.it
Università di Bergamo and CEPR

Harald Fadinger - harald.fadinger@uni-mannheim.de
University of Mannheim and CEPR

Gino Gancia - g.gancia@qmul.ac.uk
Queen Mary University of London and CEPR

Robot Imports and Firm-Level Outcomes*

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

This draft: November 2021

Abstract

We use French data over the 1994-2013 period to study how imports of industrial robots affect firm-level outcomes. Compared to other firms operating in the same 5-digit industry, robot importers are larger, more productive, and employ a higher share of managers and engineers. Over time, robot import occurs after periods of expansion in firm size, and is followed by improvements in efficiency and a fall in demand for labor. 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, exogenous exposure to automation leads to job losses. We also find that robot exposure increases productivity and the employment share of high-skill professions, but has a weak effect on total sales. The latter result suggests that productivity gains from automation may not be entirely passed on to consumers in the form of lower prices.

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), the CRC TR224 Conference, the CRC-ZEW Workshop "Labour Market Adjustments to New Technologies and Globalisation", the ECB conference "Challenges in the digital era", ERWIT (2020), IHS Vienna, the OECD, University of Mannheim, University of Pavia 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

Humans have always been afraid of competing against machines. Back in the 19th century, the Luddites protested violently against automated textile equipment fearing it would destroy their jobs. In the 1930s, John Maynard Keynes warned of the risk of "technological unemployment". Today, amid growing concerns, economists and politicians alike are discussing the opportunity of introducing a robot tax. While changes in the production process did not lead to mass unemployment, at least yet, stagnation in wages and productivity growth, and soaring inequality, are fuelling the view that new technologies failed to deliver the promised prosperity.

In this debate, the rise of industrial robots has gained special attention. Robots are programmable machines that have the capability to move on at least three axes. As such, robots, unlike other pieces of equipment, 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. Yet, thanks to several recent technological advancements, today's robots have a much higher degree of autonomy. As a result, the adoption of these technologies has grown at a staggering rate.[1]

Industrial robots are technologies adopted by firms. To understand their effect on the economy, one must know how they affect the firms adopting 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? And if so, are these productivity gains passed on to consumers or rather used to consolidate market power? From a theoretical perspective, the answer to all these questions is ambiguous. From an empirical perspective, unfortunately, the available evidence is worryingly limited due to the lack of firm-level data on the use of robotics (Raj and Seamans, 2018).

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 firms over the 1994-2013 period and a novel identification strategy to identify causality. Recently, researchers have turned to import data as a source of information on the usage of robots. Although

[1] By 2018, there were an estimated 2.44 million industrial robots performing a variety of tasks that humans used to do. This number is expected to reach 4 million by 2022 and 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).

they do not include domestic purchases, robot imports are widely recognized as a good proxy for automation because of the high concentration of this very specialized sector. For instance, in 2017, the top six leading companies, ABB (Switzerland), Omron (US), Fanuc (Japan), Kawasaki Robotics (Japan), KUKA (Germany) and Yaskawa (Japan) accounted for 44 percent of global revenue. Global exports are also dominated by few suppliers, with Japan and Germany alone accounting for 50 percent of the total volume, while France's share is about 5 percent only. Compared to other proxies used in the literature, such as dummies collected from surveys, the advantage of robot imports is that they provide a precise and comparable measure of automation intensity that is available for the near universe of firms. With this rich micro data, we develop various empirical strategies to identify the causal effects of robot adoption on sales, productivity and employment within French firms and we carry out an extensive set of robustness checks.[2]

To guide the empirical 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. Automation saves on production workers, but it also requires non-production workers such as engineers and managers. If the cost of robots declines, firms choose to invest more in automation, with ambiguous effects on employment. On the one hand, machines displace workers; on the other hand, the increase in productivity raises the demand for all factors. These effects vary across firms: since automation saves on the variable cost, firms facing a higher demand invest more aggressively in automation and are more likely to shed workers.

The model yields a number of testable predictions. First, it shows that positive demand shocks are likely to increase employment and automation simultaneously, thereby generating a spurious positive correlation between these variables in the data. Negative shocks to the cost of machines, instead, trigger automation and are more likely to reduce employment, especially in firms that are more prone to automate. The model also yields a measure of automation *intensity* that is independent of demand shocks and hence is more likely to capture the negative effect on employment. Besides the impact on labor demand, automation increases productivity and the relative demand for non-production workers.

We then take these predictions to the data. We start by documenting some descriptive patterns. We focus on the manufacturing sector, where automation is more prevalent, and exclude industries in which robot importers are more likely to be resellers rather than final users. These data show that robot adopters differ significantly from non-adopters. In particular, consistently with the model, they are larger, more productive, and have a larger

[2] We also validate that our firm-level proxies are consistent with commonly-used industry-level measures.

employment share of high-skill professions. We also find that over time the value-added share of robot adopters has grown significantly more than their employment share.

We next carry out two preliminary exercises aimed at gauging the role of demand shocks in driving these patterns. First, we use a difference-in-differences event study approach to analyze how firm-level outcomes evolve over time for firms that start to adopt robots relative to firms that do not. The results show that robot adoption occurs after periods of expansion in firm size, and is followed by improvements in productivity and labor demand shifts towards high-skill professions. However, the upward trend in employment reverses and sales stop diverging after adoption, suggesting that workers start to be displaced and that productivity gains do not translate entirely into a fall in prices. Second, we regress the firm-level outcomes on a measure of robot intensity, the ratio between cumulated robot imports and the capital stock of the firm, which should purge away demand shocks. The results indicate that an increase in robot intensity is associated with a fall in employment, and an increase in labor productivity and in the employment share of high-skill professions. These findings suggest that demand shocks may be responsible for the positive correlation between employment and robot adoption in the data.

To identify the causal effects of robots, we next focus on long-run changes in outcomes within firms and exploit variation in the decision to adopt robots driven by pre-existing differences in technological characteristics that determine the predisposition to automate. More precisely, we construct a novel identification strategy by interacting a proxy for how suitable production is for automation in a given industry with a proxy for the ease with which robots can replace worker activities within each firm. This variable captures the idea that a reduction in the cost of machines, which should be relatively larger in industries whose production is more suitable for automation, should affect robot adoption relatively more in firms whose production is more intensive in tasks that can be performed by robots. Compared to shift-share designs, this difference-in-differences approach is based on entirely pre-determined firm and industry characteristics that cannot respond to demand shocks.

Our proxy for an industry's suitability for automation is the initial average robot intensity of all other firms in the same 5-digit industry. Our firm-level proxy for replaceability is instead the pre-sample share of employment that can be replaced by robots in each firm, and is constructed by combining the classification of tasks performed by robots in Graetz and Michaels (2018) with detailed firm-level occupational data. Accordingly, our identification strategy exploits differential exposure to robots across firms that operate in industries with varying suitability for automation and exhibit a heterogeneous prevalence of automatable tasks in production. An important advantage of this strategy is that we can build an exposure variable even for firms that do not import robots directly.

We find that while robot adoption and employment growth are correlated, firms with initially more replaceable tasks operating in industries more suitable for automation experience a stronger reduction in employment than other firms. Regarding other outcomes, we find that robot exposure leads to an increase in the employment share of high-skill professions and in various measures of productivity, while the effects on total sales are positive, but not statistically significant. We then perform an extensive sensitivity analysis, including building industry's suitability from the US stock of installed robots from the International Federation of Robotics (IFR), which is however available for a coarser industry classification only. We also consider possible threats to identification by controlling for a large set of other firm- or industry-level characteristics that may have differential effects depending on the industry's suitability for automation or the firm's replaceability of employment.

To lend further credibility to our approach and shed more light on the mechanism at play, we show that robot exposure is a significant predictor of robot adoption and use these results to quantify the effect that the variation in robot adoption induced by robot exposure, as opposed to demand shocks, exerts on firm-level outcomes. We find that exogenous adoption explains an average annual fall in employment equal to 2.13 percent in robot adopters relative to the remaining firms.

These patterns suggest that demand shocks lead firms to both expand and automate, resulting in a positive spurious correlation between robot adoption and employment. Once demand shocks are neutralized, however, the relationship turns negative, confirming the hypothesis that exogenous changes in automation lead to job displacement. Hence, our results warn that caution should be exercised in interpreting the positive correlation between robot adoption and employment often found in the literature. The weaker results on sales also suggests that, while robot adoption increases productivity, the higher efficiency does not necessarily lead to a fall in prices. Consistently with this interpretation, we also show some evidence that robot exposure leads to an increase in reported profits, but has weak effects of export prices. This suggests that part of the gains for consumers may be muted by an increase in markups. To date, this is the first evidence lending support to the hypothesis that investment in robots may give firms market power. It also raises the concern that firms may have had an incentive to choose an "excessive" level of automation (see, for instance, Acemoglu and Restrepo, 2018a, Martinez, 2021, Korinek and Ng, 2018, Caselli and Manning, 2019).[3]

To our knowledge, this is the first paper that identifies the causal effect of industrial robots at the firm level. In doing so, it contributes to the growing literature on the labor market impact of automation. Several influential papers use data from the IFR, which provides information on purchases of industrial robots for a set of countries and industries. The results are mixed. Acemoglu and Restrepo (2019) find that US commuting zones that were more exposed to robots during the 1990–2005 period experienced negative effects on employment and wages. However, in a panel of 17 countries, Graetz and Michaels (2018) find that, while robots reduced the employment share of low-skill workers, they only had a small effect on total employment and positive effects on productivity. Dauth, Findeisen, Suedekum and Woessner (2018) find that higher robot exposure across local labor markets in Germany led to job losses in manufacturing that were however offset by gains in the service sector.[4]

To overcome the limitations of the IFR data, some recent papers have started to focus on imports of industrial robots. Acemoglu and Restrepo (2018b) and Blanas, Gancia and Lee (2019) use robot imports at the country level. The former paper shows that robot imports behave similarly to other proxies for investment in automation and uses them to study the demand for robots; the latter paper finds that sectors more prone to automation in countries importing more from leading suppliers of robots experienced a fall in demand for low-skill, young and female workers. Bonfiglioli et al. (2021) map US robot imports to commuting zones and find that automation lowered manufacturing employment but also offshoring. Firm-level robot imports have been used by Humlum (2019) for Denmark, Dixen, Hong and Wu (2019) for Canada, and Acemoglu, Lelarge and Restrepo (2020) for France. Importantly, none of these papers uses exogenous variation across firms to isolate the causal effect of robot adoption and, as a result, they tend to find positive correlations with employment.

Finally, there is a growing number of papers using alternative proxies for automation at the firm level. Some use dummies from survey data. These include Koch, Manuylov and Smolka (2019) for Spain, Cheng et al. (2019) for China, Dinlersoz and Wolf (2018) for the US, and a study by the European Commission (2016) for 7 European countries. They

[3] In an extension of the model, we allow for the possibility that automation, by fostering technological lead, increases market power. In this case, the cost savings are partly offset by an increase in markups and, besides efficiency considerations, firms have an incentive to invest in automation just to increase market power.

[4] Other papers showing that alternative measures of automation leads to employment losses in some sectors that are offset by employment gains in others include Mann and Puttman (2017) and Autor and Salomons (2017).

find that robots are generally more likely to be used in larger and more productive firms, and are associated with positive or non-negative changes in employment. Once more, these papers document mostly conditional correlations. Positive employment effects are also found by Aghion et al. (2020), who proxy automation with investment in industrial equipment and electricity consumption of French firms, and use a shift-share Instrumental Variables design to identify causality. As we show in our sensitivity analysis, a key difference is that they consider a broader measure of capital inputs, which is likely to be complementary to labor. In line with our findings, instead, Bessen et al. (2019) use matched employer-employee data from the Netherlands to show that spikes in expenditure on "third-party automation services" increase job separations.

The remainder of the paper is organized as follows. In Section 2, we build a partial equilibrium model in which heterogeneous firms invest in automation, and we use it to derive empirical implications. In Section 3, we discuss the French firm-level data and the main aggregate facts regarding robot imports. In Section 4, we provide descriptive evidence on how robot adopters differ from other firms and we study what happens after a firm in the sample starts importing robots. In Section 5, we develop a novel identification strategy to estimate the effect of robots on firm-level outcomes. Section 6 concludes.

2 The Model

To guide the empirical analysis, we build a model of monopolistic competition in which heterogeneous firms combine production workers, non-production workers and capital to produce differentiated goods. Firms can also invest in automation, which allows capital to perform tasks that used to be performed by labor. The model illustrates the causes and consequences of automation, and the main challenges when testing its empirical predictions. It also suggests some possible identification strategies. The analysis is in partial equilibrium and is deliberately kept as simple as possible.[5]

[5] The model adds firm heterogeneity to earlier contributions combing the task-based approach and endogenous automation. See, for instance, Zeira (1998), Acemoglu and Autor (2011), Acemoglu and Restrepo (2018a), Hemous and Olsen (2018), 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.

2.1 The Basic Set-Up

Consider a sector producing differentiated varieties ω with preferences over these varieties exhibiting constant elasticity of substitution:

$$C = \left[\int_{\omega \in \Omega} c\left(\omega\right)^{\frac{\sigma-1}{\sigma}} \mathrm{d}\omega \right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1.$$

Firm i producing a single variety faces a demand function with a constant price elasticity σ:

$$y_i = A_i p_i^{-\sigma}, \tag{1}$$

where p_i is the price charged and A_i is a parameter capturing demand conditions.

To produce y_i, a firm with productivity φ_i must employ capital and production workers in a unit measure of tasks z:

$$y_i = \varphi_i \exp\left(\int_0^1 \ln x_i(z) \mathrm{d}z \right). \tag{2}$$

Tasks $z \in [0, \kappa_i]$ are automated, and hence can be performed by capital. The remaining tasks, $z \in (\kappa_i, 1]$, can only be performed by production workers. Hence, κ_i represents the extent of automation. Let (k_i, l_i) denote the quantity of capital and labor, respectively, used for the production of y_i. Denote with r the rental rate of capital and with w the wage of production workers. We assume $r < w$, which will guarantee that automation raises productivity. Since machines are cheaper than workers, there is complete specialization, in the sense that tasks $z \in [0, \kappa_i]$ are performed by machines. Hence, given symmetry we obtain:

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

Substituting these into (2) yields:

$$y_i = \varphi_i \left(\frac{k_i}{\kappa_i} \right)^{\kappa_i} \left(\frac{l_i}{1-\kappa_i} \right)^{1-\kappa_i}. \tag{3}$$

To produce, the firm must also hire f_i non-production workers (managers and engineers) with wage h. For now, we take f_i as given, later we will assume it a function of automation, κ_i.

2.2 Exogenous Automation

We now solve the problem of the firm for a given level of κ_i. Firms are monopolistically competitive and choose labor and capital so as to maximize profit,

$$\max_{k_i, l_i} \{p_i y_i - r k_i - w l_i - h f_i\},$$

subject to the demand schedule (1), given the production function (3) and taking automation, κ_i, as given. The first-order condition for labor is:

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

Equation (4) shows automation, κ_i, to have two opposite effects on the demand for labor. First, there is a negative displacement effect, captured by $(1 - \kappa_i)$ and given by the fact that more tasks can be performed by machines (capital). Second, as we will see shortly, there is a positive productivity effect, since an increase in κ_i raises production, which in turn increases the demand for labor.

The first-order condition for capital is:

$$r k_i = \left(1 - \frac{1}{\sigma}\right) \kappa_i p_i y_i. \tag{5}$$

Intuitively, the demand for capital is increasing in the set of tasks it can perform. Taking the ratio of (4) and (5), we obtain:

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

which shows that the capital to labor ratio is also increasing in automation, κ_i.

Substituting k_i back into the production function yields:

$$y_i = \varphi_i \frac{l_i}{1 - \kappa_i} \left(\frac{w}{r}\right)^{\kappa_i}, \tag{6}$$

which shows that output per production worker is increasing in κ_i if $w > r$, as assumed. Intuitively, if labor is more expensive than capital, replacing workers with machines through automation reduces the marginal cost and increases productivity. Finally, using equation

(6) into the demand for labor (4) yields:

$$l_i = w^{-\sigma}\left(1-\frac{1}{\sigma}\right)^{\sigma} A_i\varphi_i^{\sigma-1}\left(\frac{w}{r}\right)^{\kappa_i(\sigma-1)}(1-\kappa_i)\,. \tag{7}$$

This equation shows how employment depends on κ_i and other exogenous parameters. It can be used to study how the productivity effect and the displacement effect depend on the level of κ_i. In the limit case of full automation ($\kappa_i \to 1$), it is immediate to see that $l_i \to 0$. This is intuitive, since in this case workers become useless for the firm, because capital can perform all tasks at a lower cost. Hence, the displacement effect must eventually dominate for high levels of automation. However, at low levels of automation, the productivity effect may dominate the displacement effect. To see this, take the derivative of (7) with respect to κ_i:

$$\frac{d\ln l_i}{d\kappa_i} = (\sigma-1)\ln\left(\frac{w}{r}\right) - \frac{1}{1-\kappa_i}. \tag{8}$$

This derivative is positive for values of κ_i lower than $1-[(\sigma-1)\ln(w/r)]^{-1}$. This condition is more likely to be satisfied when σ and w/r are high, i.e., when the productivity effect is strong enough. In particular, 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 instead $(\sigma-1)\ln(w/r) < 1$, then the displacement effect always dominates.[6]

Finally, substituting (7) in (6) we can express output as a function of automation and other exogenous parameters:

$$y_i = A_i\varphi_i^{\sigma} w^{-\sigma}\left(1-\frac{1}{\sigma}\right)^{\sigma}\left(\frac{w}{r}\right)^{\kappa_i\sigma}. \tag{9}$$

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

$$\frac{d\ln y_i}{d\kappa_i} = \sigma\ln\left(\frac{w}{r}\right), \tag{10}$$

and it illustrates that the productivty effect is stronger in industries where demand is more elastic (σ). Moreover, substituting (4) and (5) into the profit function yields

$$\pi_i = \frac{p_i y_i}{\sigma} - hf_i,$$

[6] Acemoglu and Restrepo (2018a) emphasize another possible effect, namely, that new tasks are created when others are automated. We abstract from this additional mechanism which would reinforce the positive productivity effect on employment.

which shows the familiar result that operating profit is a constant share $1/\sigma$ of revenue.

We summarize the impact of exogenous automation on firm-level employment and production in the following proposition (proof in the text).

Proposition 1 *Suppose $w > r$. Other things equal, an increase in automation, parameterized by a rise in κ_i:*
(i) increases production of firm i, y_i, with a stronger effect the higher is σ;
(ii) decreases employment of firm i, l_i, if $\kappa_i > 1 - [(\sigma - 1)\ln(w/r)]^{-1}$, and increases it otherwise.

2.3 Endogenous Automation

We now allow firms to choose the level of automation, κ_i. Substituting workers with machines requires a costly change in technology, and automating more and more tasks poses an increasingly difficult challenge. Hence, we assume that automation entails a cost in terms of non-production workers (i.e., managers and engineers), which is increasing and convex in κ_i. For convenience, we assume the cost hf_i to take the following form:

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

where the parameter $\rho_i \in (0, 1)$ captures heterogeneity across firms in the ease with which each task can be replaced by machines. To see this, note that the marginal cost of automation is

$$hf_i'(\kappa_i) = h(1-\kappa_i)^{-1/\rho_i}. \tag{11}$$

Equation (11) shows that automating the marginal task costs h at $\kappa_i = 0$ and this cost tends to infinity for $\kappa_i = 1$. Moroeover, automation costs increase at a faster rate with κ_i the lower ρ_i is. Hence, ρ_i can interpreted as an index of replaceability of tasks in the production process of firm i.[7]

In this set-up, firms choose the level of κ_i that maximizes profit given the choice of factors derived in the previous section:

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

Automation poses a trade-off between its fixed cost and the reduction in the variable cost it

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

generates. The first-order condition for κ_i is:

$$\left(1-\frac{1}{\sigma}\right)p_i y_i \ln\left(\frac{w}{r}\right)=hf_i'\left(\kappa_i\right). \tag{12}$$

The left-hand side of (12) is the marginal benefit of automation. It shows that the benefit of automation is increasing in the demand elasticity (σ), in revenues ($p_i y_i$) and in the cost saving entailed by machines (w/r). The right-hand side is instead the marginal cost.

Substituting y_i from (9) and (11), the first-order condition for automation (12) becomes:

$$\left(1-\frac{1}{\sigma}\right)^{\sigma}A_i\left(\frac{\varphi_i}{w}\right)^{(\sigma-1)}\left(\frac{w}{r}\right)^{\kappa_i(\sigma-1)}\ln\left(\frac{w}{r}\right)=h\left(1-\kappa_i\right)^{-1/\rho_i}. \tag{13}$$

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. We can show that the second-order condition is necessarily satisfied if $(\sigma-1)\ln(w/r)<1/\rho_i$ and the unique solution is interior if:

$$\left(1-\frac{1}{\sigma}\right)^{\sigma}A_i\left(\frac{\varphi_i}{w}\right)^{(\sigma-1)}\ln\left(\frac{w}{r}\right)>h. \tag{14}$$

Clearly, if $w/r<1$ there is no benefit of automation, hence the optimal κ_i is zero. We assume both conditions to be satisfied.

We summarize in the next proposition the determinants of automation (proof in Appendix A).

Proposition 2 *Suppose that condition (14) is satisfied. Then, there exists a unique equilibrium choice of automation $\kappa_i^* \in (0,1)$ implicitly defined by equation (13). The comparative statics of κ_i^* to changes in the exogenous parameters are:*

$$\frac{d\kappa_i^*}{dA_i}>0;\quad \frac{d\kappa_i^*}{d\varphi_i}>0;\quad \frac{d\kappa_i^*}{d(w/r)}>0;\quad \frac{d\kappa_i^*}{d\rho_i}>0;\quad \frac{d\kappa_i^*}{dh}<0. \tag{15}$$

These results are intuitive and consistent with the existing literature.[8] Larger firms (high 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 $1/\rho_i$.

[8]See, for instance, Dechezlepretre et al. (2019), Cheng et al. (2019), Hemous and Olsen (2018), Koch, Manuylov and Smolka (2019).

2.4 Empirical Implications

Effect of Automation on Firm-Level Outcomes. The model has sharp predictions for the effect of automation on some firm-level outcomes. In particular, automation should clearly have a positive effect on measures of productivity and increase the demand for non-production workers. The implications of the model regarding the relationship between automation and employment are instead more nuanced. In particular, (8) shows that the effect of κ_i on l_i is potentially ambiguous, and possibly heterogeneous across firms and sectors. Hence, whether or not automation raises employment may ultimately be an empirical question.

Challenges to Identification. The model also illustrates the key challenge that the econometrician faces in identifying the causal effect of automation. The main difficulty hinges on the endogeneity of κ_i. As shown in Proposition 2, automation depends on parameters that capture the exogenous costs and benefits of machines, namely, replaceability (ρ_i), the cost of capital (r) and the cost of non-production workers (h), but it also depends on shocks that have a direct effect on firm-level outcomes. In particular, both demand shocks, captured by A_i, and productivity shocks, captured by φ_i, trigger automation, but they also have a direct positive effect on production and employment. Hence, these shocks may bias upwards the estimates of the effect of automation on sales and productivity; even worse, they may generate a positive correlation between automation and employment, even if, conditional on them, an increase in κ_i would lead to job losses. Firm and sector-year fixed effects are not sufficient to solve the problem because these shocks are likely to vary both across firms and over time. For instance, a recent literature has highlighted the quantitative importance of firm-specific demand shocks for explaining sales.[9] Fortunately, the model also offers possible remedies to this bias. Specifically, exogenous shocks to the costs and benefits of automation can be used to isolate variation in κ_i that is orthogonal to demand shocks.

Netting Out Demand Shocks: Automation Intensity. To identify firm-specific shocks to the cost of automation, the model suggests to use *automation intensity* defined as the level of automation, κ_i, over capital expenditure, rk_i. Using the first-order condition for k_i

[9] See, for instance, Hottman, Redding and Weinstein (2016) and Bonfiglioli, Crinò and Gancia (2019).

(5), into the first-order condition for automation (12), we can write:

$$\frac{\kappa_i}{rk_i} = \frac{1}{hf_i'(\kappa_i)} \ln\left(\frac{w}{r}\right). \tag{16}$$

This equation shows that automation intensity captures variation in the marginal cost and benefits from automation that is independent of demand. The reason is that shocks to demand raise both κ_i and k_i leaving the ratio unchanged. Controlling for firm and sector-year fixed effects should also purge this measure from any variation that is not driven by firm-specific changes in the cost of automation. Nevertheless, a bias will still remain if fixed effects do not fully absorb the impact of r and w, because factor prices have a direct effect on the demand for labor, and not just through automation. In particular, an increase in wages will lower the demand for labor for any given level of κ_i:

$$\left.\frac{d\ln l_i}{d\ln w}\right|_{\kappa_i} = -\sigma(1-\kappa_i) - \kappa_i < 0.$$

On the other hand, a decline in r, will increase the capital stock and the demand for labor for any given level of κ_i:

$$\left.\frac{d\ln l_i}{d\ln r^{-1}}\right|_{\kappa_i} = \kappa_i\,(\sigma-1) > 0.$$

Identifying Causal Effects: Robot Exposure. To mitigate the concern that proxies for automation may still be correlated with other characteristics that may affect the outcomes of interest, the model suggests to focus on the *interaction* between industry- and firm-level proxies of automation opportunities. It is well-known that the use of automation technologies vary dramatically across industries. This suggests that the parameters capturing the cost and benefits of automation, $h^{-1}\ln(w/r)$ in the model, are likely to have an important industry-level component. At the same time, even within industries, firms differ significantly in the structure of employment and hence in the replaceability of the tasks they perform, as captured by the parameter ρ_i. The model then predicts variation in the costs and benefits of automation, $h^{-1}\ln(w/r)$, to have a stronger effect on robot adoption, κ_i in firms performing more replaceable tasks, ρ_i, as can be seen by rearranging (12):

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

Based on this insight, in the next sections, we build a measure for exposure to robots by combining information on which industries are more suitable for automation with firm-level measures of replaceability of employment. The interaction between these variables allows us to isolate the impact of the automation intensity of a sector on firms whose workers are more substitutable with machines, while simultaneously netting out the direct effect that each variable may have in isolation.

Finally, all these predictions have been derived in a model where the choice of automation is continuous. In the data, however, the decision to automate is often measured by binary variables. Nevertheless, as we show in Appendix C, a variant of the model where automation is a discrete choice yields qualitatively similar predictions: a decline in the cost of capital increases the probability that firms adopt a higher automation intensity, and the increase in this probability is higher if tasks are easier to replace with machines.

3 Data and Aggregate Facts

Our empirical analysis uses firm-level data for France over the 1994-2013 period and combines several firm-level datasets administered by the French statistical agency (INSEE). We observe the universe of French firms (defined as legal entities) that report a complete balance sheet in the manufacturing, services and primary sectors (roughly 500,000 firms per year), excluding the government sector. Each firm is uniquely defined by a firm-level identifier (SIREN number) common across all data sets. For each firm that reports a complete balance sheet, we have data on sales, material purchases, capital stock (value of physical assets) and accounting profits in Euros, as well as on total employment.[10] We use this information to compute firm-level value added[11] and revenue TFP. We compute revenue TFP from a Cobb-Douglas value-added production function with labor and physical capital as inputs and output elasticities of inputs that vary at the 2-digit level of the NACE industrial classification. We use the Wooldridge (2009) estimator for estimating the production-function coefficients.[12]

The balance sheet data are complemented with information on the occupational structure

[10] For the years 1994 to 2009 the source of this information is BRN. For 2011-2013 the data source is FARE, which substitutes BRN and is more comprehensive in terms of coverage. This dataset is prepared by INSEE and combines administrative data with survey information and also uses imputation. Compared to BRN, it additionally includes firms that do not report a full balance sheet. We use the subset of FARE that is consistent in terms of sample with BRN.

[11] 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.

[12] The Wooldrige estimator is based on the Levinsohn-Petrin (2003) methodology but uses a one-step GMM estimator instead of a two-step approach. This estimator solves the problem that the labor coefficient may be unidentified in the first stage if labor is freely adjustable (see Ackerberg, Caves and Frazer, 2015). We consider labor endogenous and use lagged labor as an instrument.

of employment from DADS Etablissement. For each sample year, DADS Etablissement contains plant-level employment data disaggregated in five two-digit occupations: (1) firm owners receiving a wage; (2) high-skill professions (i.e., scientists, managers, and engineers); (3) intermediate-skill professions (e.g., teachers, administrative assistants, and technicians); (4) low-skill white-collar workers; and (5) blue-collar workers. We aggregate the occupational employment data from DADS across all plants belonging to the same firm using the SIREN identifier, thereby obtaining the occupational structure of employment for each firm in a given year. For the year 1994, DADS contains more disaggregated information on employment for 29 occupations. As explained in Section 5, we exploit this information for measuring the extent to which employment is replaceable by robots in each firm; we use this variable for the construction of a proxy for robot exposure, which we employ to identify the effects of automation on firm-level outcomes. For the descriptive analysis, we use the full set of years (1994-2013) while for identification we focus on the 1996-2013 period and use 1994 as a pre-sample period.

For each firm and year, we also have customs data on exports and imports from the French customs authority (DOUANE). We observe quantities and values of imports and exports for all 8-digit products of the Combined Nomenclature (CN) classification by origin and destination country. We leverage the detailed information on firm-level imports by product to build proxies for the use of robots within firms. The CN classification records trade in industrial robots into a specific product code, CN 84795000 (CN 84798950 before 1996). We identify firms that import robots in a given year as firms with positive imports for this product code in that year. 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. For each firm, we thus have a proxy not only for whether it adopts robots or not but also for the intensity with which it uses robots in production.

The wide coverage of the import data, coupled with the possibility of constructing continuous measures of robot intensity, represents the key strengths of our automation variables compared to the typical binary indicators of robot adoption based on firm-level surveys. At the same time, the use of import data requires two *caveats*. First, these data do not include purchases of robots from domestic suppliers. As a consequence, the results obtained from this data may not generalize to firms that source robots mostly domestically. We refer to this concern as the problem of "missing robots". While these instances of adoption are relatively rare in a country like France, our identification strategy will exploit variation in proxies for robot exposure based on technological characteristics that are observed for all firms, in-

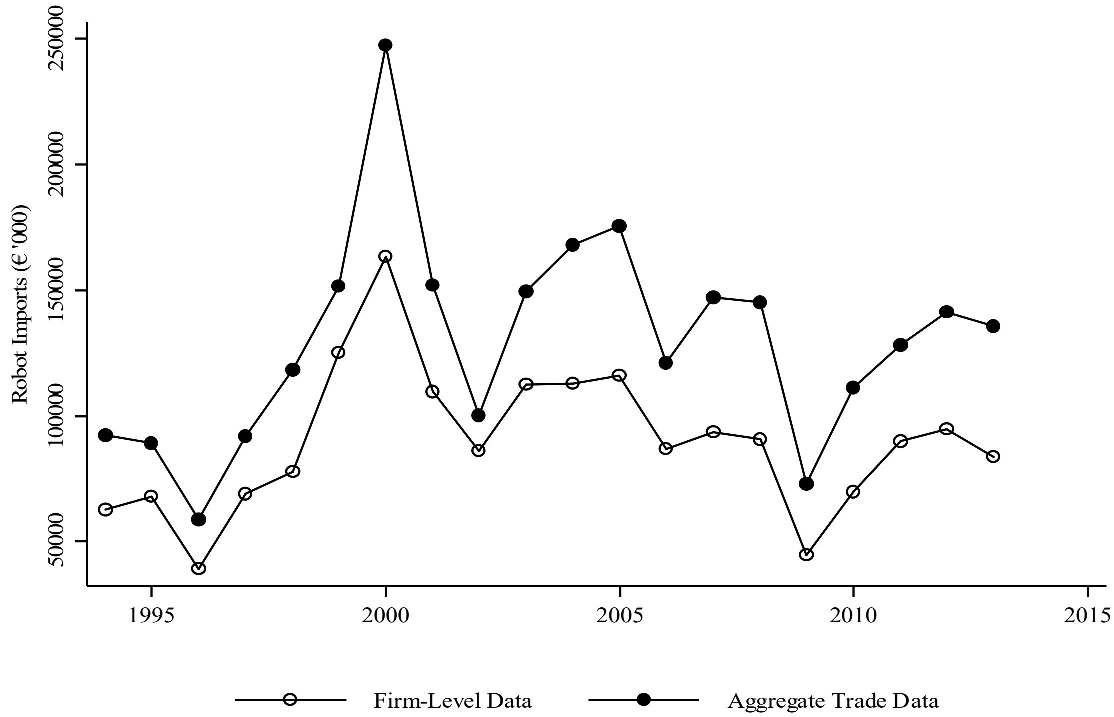

Figure 1: Robot Imports, 1994-2013

cluding those that do not purchase robots from abroad. Second, our data may also include imports of robots by robot integrators or resellers, which do not represent actual instances of adoption. We refer to firms that import but do not use robots as "false positives". To allay the concern that these firms may bias the results, we exclude robot resellers from the analysis by restricting the sample to manufacturing firms and we 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. Consistently, we find that our firm-level proxies align well with the commonly-used industry-level measure of the stock of installed robots from the IFR.

Figure 1 plots the time series of total robot imports into France obtained by summing up robot imports across all French firms (hollow circles). For comparison, the figure also plots the time series of total French robot imports obtained from the Comext database (full circles). The firm-level data follow quite closely the evolution of aggregate French robot imports implied by official statistics, and account for the majority of these imports in any given year. Interestingly, robot imports are quite volatile, consistent with the lumpy nature of this investment. Yet, due to this steady investment, the stock of imported robot capital has markedly increased over time, from 63 million Euros in 1994 to around 1.8 billion Euros

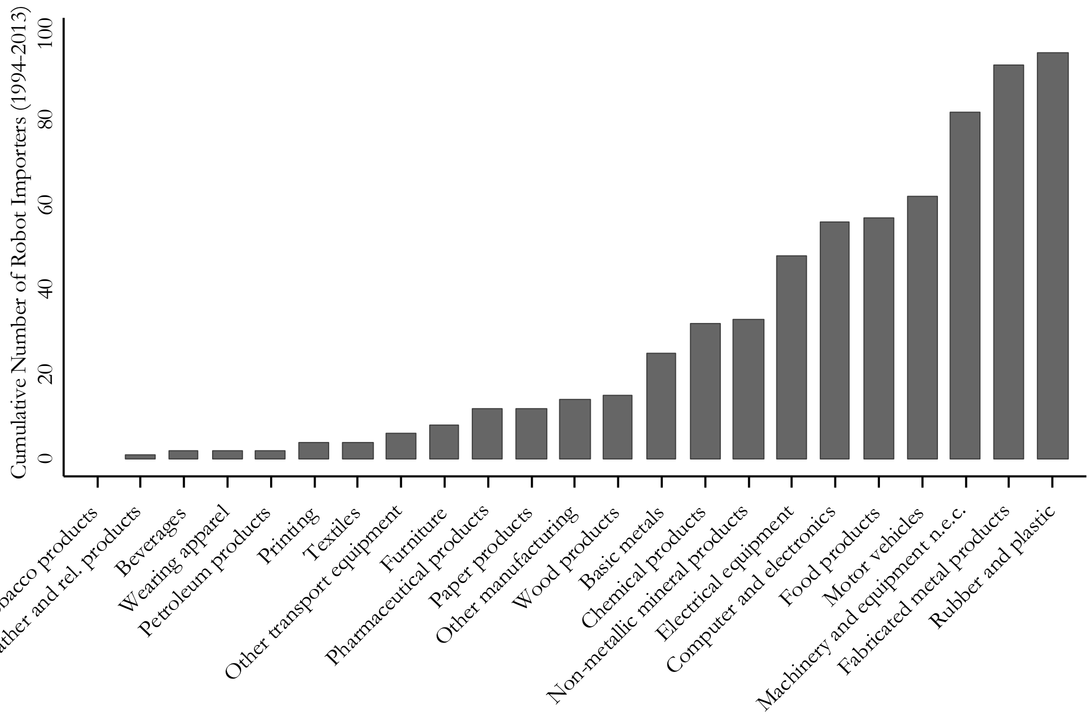

Figure 2: Cumulative Number of Robot Importers by Two-Digit Sector

in 2013. Overall, these numbers suggest that automation has become increasingly widespread in France over the sample period.

Figure 2 reports the cumulative number of robot importers by 2-digit manufacturing sector. While robot importers are observed in many different sectors, they are particularly frequent in the production of motor vehicles, machinery, and electrical equipment. Since our data lack information for the two biggest car manufacturers in France, robot importers are undercounted in the "Manufacturing of Motor Vehicles" industry.[13] Yet, we will show that our main results are unchanged if this industry is excluded. The cross-sectional distribution of robot importers align well with other data sources. In particular, after removing "Manufacturing of Motor Vehicles", the correlation between the number of robot importers and the stock of installed robots from the IFR is 0.79.

Overall, our baseline sample of manufacturing firms includes roughly 800 different enterprises importing robots at least once over the period of analysis. This number is consistent with other existing studies. For instance, Acemoglu, Lelarge and Restrepo (2020), who

[13] For large multinational firms (e.g., Peugeot Sociétè 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.

collected information on robot adoption in France from multiple sources, find that only 1 percent of the firms in their sample purchased robots over the 2010-2015 period. We further focus on firms with more than ten employees given that robots are typically used at relatively large firms and adoption decisions by small firms tend to be more noisy and lumpy. However, the qualitative pattern of our results is largely insensitive to the choice of the sample.

While robot adopters, i.e., firms importing robots at least once over the sample period, are a small minority of firms, they account for a large and increasing fraction of manufacturing activity. This trend clearly emerges from Figure 3, which shows the share of employment and value added accounted for by robot adopters in any given year. To avoid the increasing trends being driven by the growth in the number of adopters over time, all shares are computed for a consistent sample of firms that are active in all years and import robots at least once over 1994-2013.[14] The share of robot adopters in manufacturing employment and value added has rapidly increased over time to reach 8 and 14 percent, respectively, in 2013. This indicates that robot adopters are faring better than other manufacturing firms. Moreover, the fact that the value added share has grown significantly more than the employment share suggests that the expansion of robot adopters may have been accompanied by the adoption of labor-saving technologies.[15]

4 Preliminary Evidence: Robot Adoption and Firm-Level Outcomes

In this section, we provide descriptive evidence on the relationship between robot adoption and firm-level outcomes. We start by studying how firms that adopt robots compare to firms that do not in terms of various characteristics. Table 1 reports summary statistics on a number of firm-level variables, separately for firms that import robots at least once over 1994-2013 ("robot adopters") and for firms that do not import robots over this period ("non robot adopters"). Our sample consists of 64,173 manufacturing firms. Of these, 765 are robot adopters, corresponding to approximately 1 percent of the total number of firms and firm-year observations in the dataset. Robot intensity, defined as the ratio between the stock of robot capital and the total physical capital stock of the firm, equals 7.8 percent on average for robot adopters. The average robot adopter is around 11 times larger than

[14] This however implies that we undercount the number of robot importers in earlier years, as some of these firms may exit.

[15] Preliminary evidence from a 2019 survey run by the US 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.

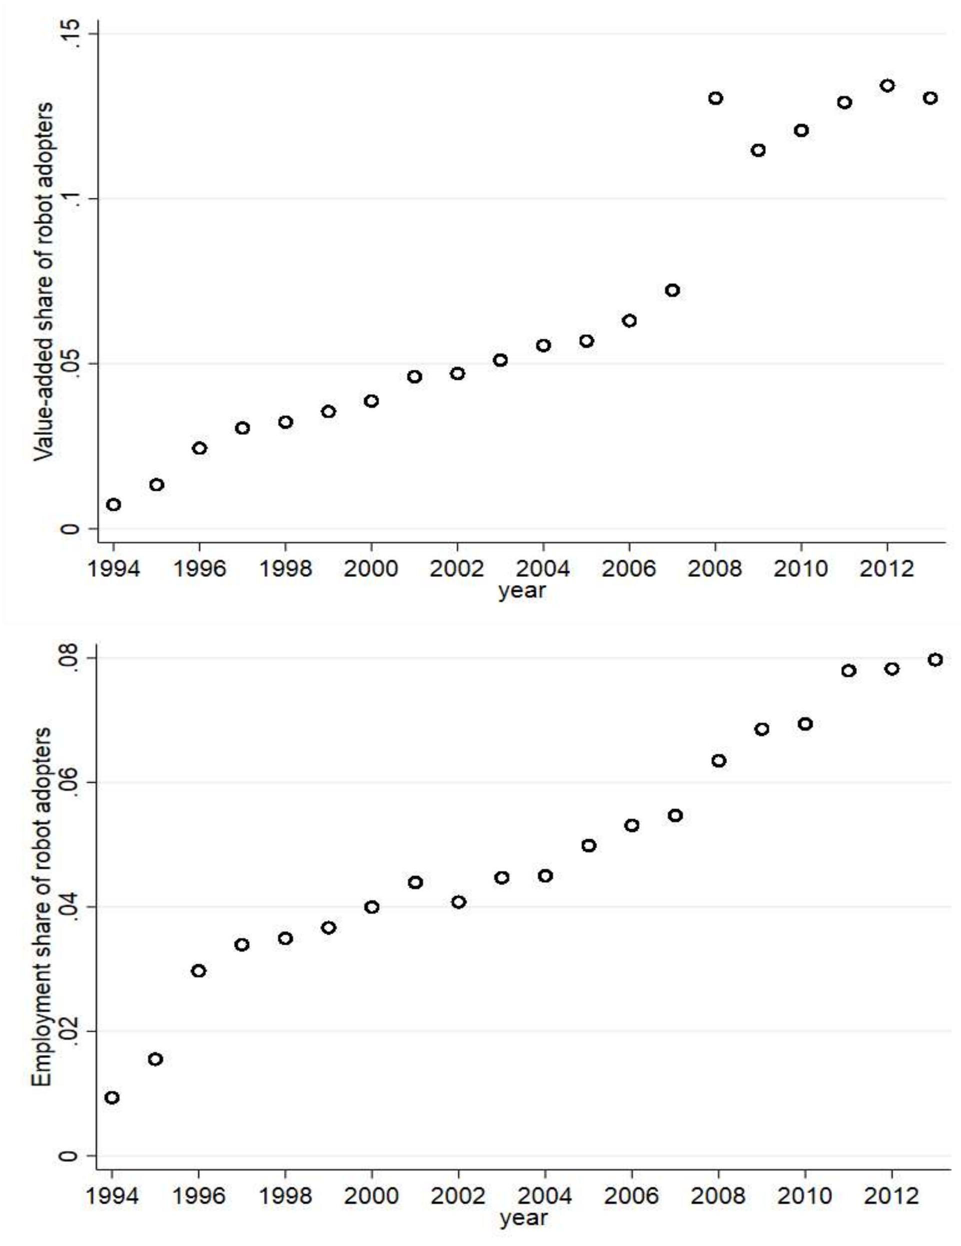

The figure plots the shares of French manufacturing value added and employment accounted for by a consistent sample of firms that import robots at least once over 1994-2013 and are active in all years.

Figure 3: Value Added and Employment Shares of Robot Adopters

the average non robot adopter in terms of employment and around 14 times larger in terms of sales. Robot adopters also exhibit around 3 times higher levels of sales per worker and around 1.5 times higher levels of TFP, on average. The skill composition of employment also differs across robot adopters and non robot 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. Finally, robot adopters are more likely to import and export goods other than robots, and make 16 times higher profits than non robot adopters, on average.

Table 1 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 adopters shed workers at a slower rate than non robot adopters (0.016 vs. 0.03 log points per year, respectively).[16] Robot adopters also experienced a relatively slower reduction in sales, sales per worker, TFP, and profits, and a relatively faster increase in the employment share of high-skill professions.

To gain further insight on 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:

$$Y_{it} = \alpha_i + \alpha_{ht} + \beta \cdot Adoption_{it} + \mathbf{X}'_{it} \cdot \boldsymbol{\gamma} + \varepsilon_{it}, \tag{17}$$

where i denotes a firm; h indicates the 3-digit NACE sector 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 estimate two versions of eq. (17). In the first version, we control for firm fixed effects, α_i, and for 3-digit sector×year fixed effects, α_{ht}. The "robot adoption premia", β, are then identified by comparing outcomes, in deviations from within-firm means, across firms belonging to the same 3-digit sector and year. This approach ensures that the coefficients β are not contaminated by time-invariant firm characteristics that could be correlated with adoption and outcomes, by differences in the distribution of adopters and non adopters across sectors and by sector-specific shocks. In the second version of eq. (17), we add controls for observable firm characteristics, namely, log sales and dummies for firms that export or import goods other than robots. We measure each of these three characteristics at baseline, that is, in the first year in which the firm is observed in the sample, and interact its first-year

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

Table 1: Descriptive Statistics

	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
Sales per Worker (€'000)	6,373	765	1,912	226	104,935	-0.061
VA per Worker (€'000)	6,225	761	178	65	2,715	-0.070
TFP	6,218	760	426	170	2,625	-0.066
Profits (€'000)	6,373	765	19,855	529	223,342	-0.052
Dummy Importer	6,373	765	0.972	1	0.164	0.001
Dummy Exporter	6,373	765	0.947	1	0.224	0.002
Export Price Index	6,039	750	242	22	2,045	0.014
	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
Sales per Worker (€'000)	598,922	63,448	666	231	11,725	-0.063
VA per Worker (€'000)	587,342	62,741	190	71	1,973	-0.066
TFP	576,404	62,005	292	132	1,362	-0.071
Profits (€'000)	598,925	57,293	1,256	98	36,795	-0.065
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
Export Price Index	335,886	42,346	280	14	16,844	0.012

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. *Export Price Index* is the export-value-weighted average of export unit values across 8-digit products of the Combined Nomenclature classification. 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*, for which annualized changes in levels are reported.

Table 2: Firm-Level Outcomes and Robot Adoption, Panel (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Ln Sales		Ln No. of Employees		Ln Sales per Worker	
$\text{Adoption}_{\text{it}}$	0.225***	0.227***	0.084***	0.100***	0.044***	0.082***
	[10.034]	[10.210]	[4.622]	[5.487]	[3.074]	[5.665]
Obs.	597,497	596,174	598,887	597,290	597,497	596,174
R2	0.95	0.95	0.87	0.87	0.91	0.91
	Ln VA per Worker		Ln TFP		Empl. Sh. High Skill	
$\text{Adoption}_{\text{it}}$	0.022	0.055***	0.062***	0.074***	0.012***	0.003
	[1.445]	[3.525]	[4.068]	[4.804]	[4.478]	[1.048]
Obs.	587,151	585,896	576,168	574,954	598,887	597,290
R2	0.85	0.85	0.87	0.87	0.69	0.70
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes

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. Sector refers to 3-digit sectors. The control variables included in columns (2), (4) and (6) are log sales and dummies for whether the firm is an importer or an exporter, observed in the first year in which the firm appears in the sample and 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.

value with a full set of year dummies. The resulting interactions, contained in the vector $\mathbf{X}_{it}$, flexibly control for heterogeneous trends across firms characterized by different initial conditions. We correct the standard errors for clustering at the firm level to account for serially correlated shocks within firms, and we report t-statistics in square brackets.

The results are reported in Table 2. Odd-numbered columns show estimates from the specification including only firm and sector×year fixed effects. Even-numbered columns report results from the specification containing also the interactions between year dummies and the initial-period values of log sales and of the indicators for importing and exporting firms. Both specifications are estimated for six major outcomes on which we focus throughout the paper: (i) log sales, (ii) log employment, (iii) log sales per worker, (iv) log value added per worker, (v) log TFP, and (vi) the employment share of high-skill professions. All estimates of β are positive and, except for two cases, they are also highly statistically significant. These results confirm that robot adopters are larger, more productive, and more skill-intensive than non robot adopters, even when accounting for time-invariant firm characteristics, firm-specific trends and the sector of operation.

The differences between robot adopters and non robot adopters may have two interpretations: either robot adopters differ from other firms already before adopting robots, or they start diverging afterwards. To shed light on this question, we use a difference-in-differences event study approach to analyze how the six outcomes evolve over time in firms that adopt robots relative to firms that do not. To this purpose, we extend eq. (17) by adding the first five lags and leads of $Adoption_{it}$:

$$Y_{it} = \alpha_i + \alpha_{ht} + \sum_{s=-5}^{5} \beta_s \cdot Adoption_{it-s} + \varepsilon_{it}. \tag{18}$$

The coefficients β_s estimated from eq. (18) illustrate how a given outcome evolves over time within robot adopters relative to non robot adopters, over a ten-year window around the first instance of robot imports ($s = 0$).

The results are reported in Figure 4, where each graph refers to a different outcome. We report in Appendix D the estimation coefficients corresponding to each plot. The figure shows that robot adoption is antedated by significant differences in the trends of sales and employment between robot adopters and non robot 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 robot 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, which takes approximately two years to unfold. No differential trend is instead observed in terms of sales after adoption, as sales stop diverging in robot adopters relative to non robot adopters after $s = 0$. This suggests that the efficiency gains from robot adoption do not translate into higher revenues; we provide an interpretation for this fact in Section 5.5. Overall, these results suggest that robot adoption occurs after periods of expansion in firm size, and is followed by employment losses, improvements in firm efficiency, and labor demand shifts towards high-skill workers, but limited changes in total sales.

Both the model and the preliminary evidence presented so far suggest that the correlations between robot adoption and other firm characteristics may be confounded by demand shocks, which are likely to influence both the outcomes of a firm and its choice to automate. The

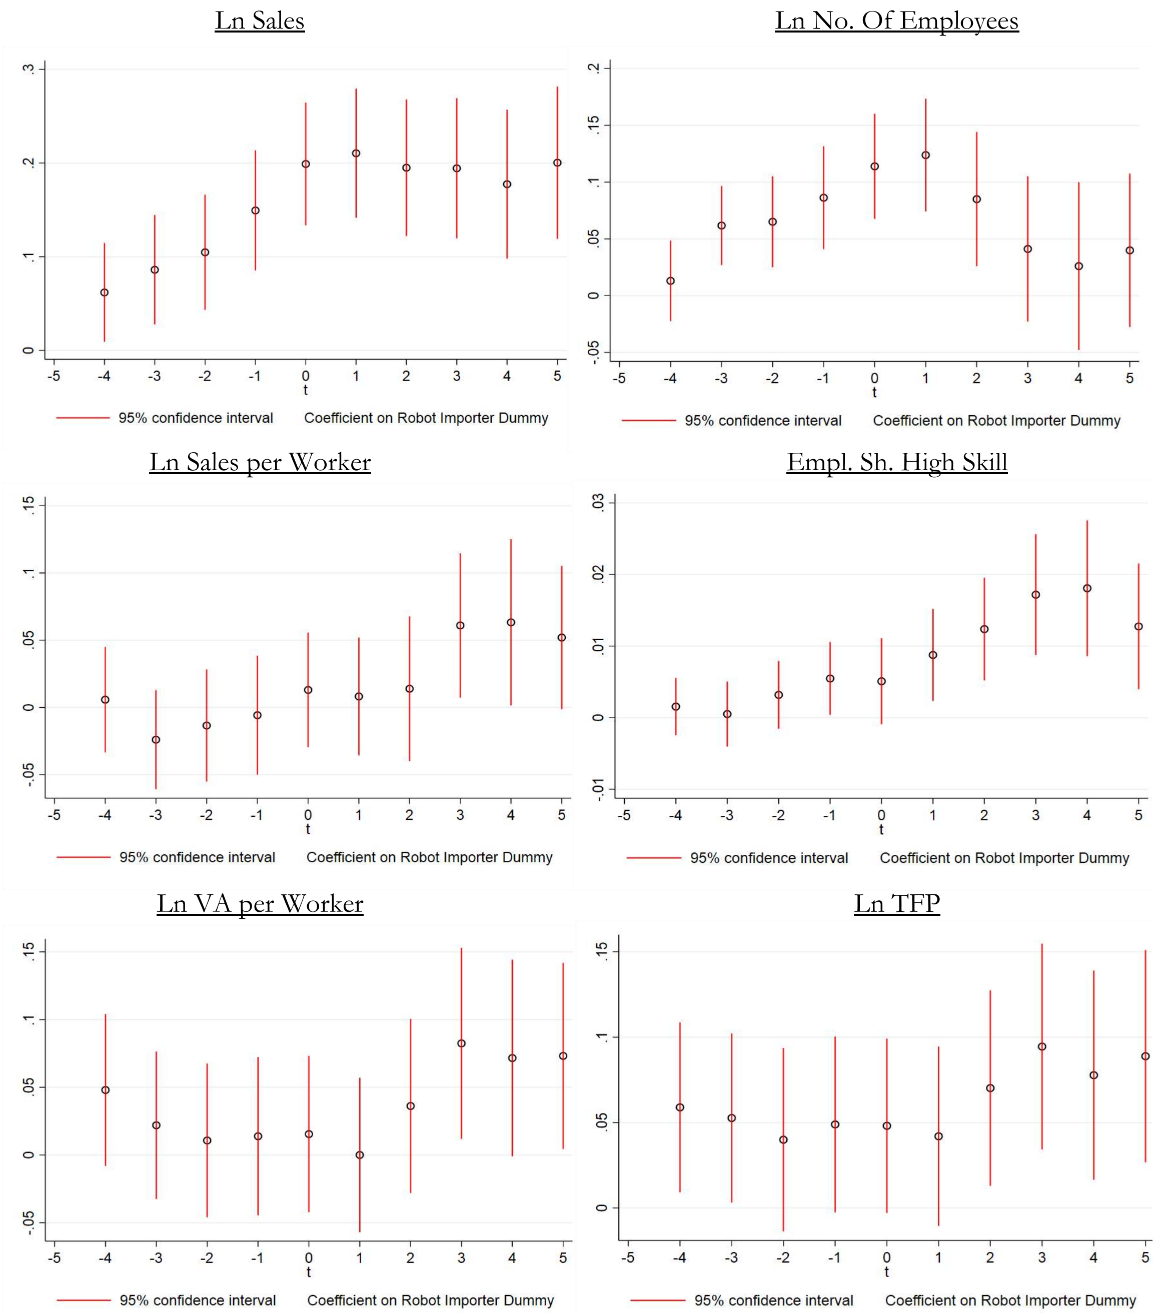

Each graph plots coefficients and confidence intervals on various lags and leads of $Adoption_{it}$ estimated using eq. (18) for a different outcome variable (indicated in the heading of the graph). $Adoption_{it}$ is a dummy that takes on value 1 in the first year in which a firm imports robots and 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 Table A1.

Figure 4: Difference-in-Differences Event Studies

difference-in-differences event studies confirm the existence of pre-trends for employment and sales, but also suggest that these trends may stop or even reverse after adoption. We close this section by providing additional evidence on the possible role played by demand shocks. To this purpose, we follow the model and regress outcomes on a continuous measure of robot intensity, a variable that should be less contaminated by demand shocks compared to the dicothomous indicator of adoption used until now. Then, we compare the results with the patterns documented above.

Specifically, we estimate eq. (17) replacing the dummy $Adoption_{it}$ with a continuous measure of the intensity with which a firm uses robots, namely, the log ratio between the stock of robot capital and the total capital stock of the firm. This variable, labeled $\ln RobInt_{it}$, is a proxy for the theoretical measure introduced in eq. (16). By scaling robot capital with the total capital stock of the firm, $\ln RobInt_{it}$ neutralizes demand shocks, as long as the latter proportionately affect 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 sector$\times$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 a sector.

The results are reported in Table 3. Odd-numbered columns refer to the specification that only controls for firm and sector-year fixed effects, while even-numbered columns refer to the specification with control variables (interactions of year dummies with initial firm size and with indicators for the initial import and export status of the firm). Compared to the results reported in Table 2, 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 firms to both expand and automate, resulting in a spurious positive correlation between robot adoption and firm size. However, once demand shocks are neutralized, automation may lead to job displacement. In terms of magnitude, multiplying the coefficient on $\ln RobInt_{it}$ in the employment regression by the average annual change in robot intensity reported in Table 1 (0.18 log points) implies that the observed increase in robot intensity is associated with an average fall in employment equal to 2.6 percent per year among robot adopters. 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. Regarding the other outcomes, Table 3 continues to show positive estimates of β across the board. While some coefficients are imprecisely estimated, the qualitative pattern of the results suggests that automation is associated with improvements

Table 3: Firm-Level Outcomes and Ln Robot Intensity, Panel (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Ln Sales		Ln No. of Employees		Ln Sales per Worker	
Ln $RobInt_{it}$	-0.118***	-0.136***	-0.144***	-0.144***	0.015	0.014
	[-3.936]	[-4.675]	[-5.802]	[-5.780]	[1.093]	[0.944]
Obs.	6,368	6,324	6,373	6,329	6,368	6,324
R2	0.96	0.97	0.93	0.93	0.90	0.90
	Ln VA per Worker		Ln TFP		Empl. Sh. High Skill	
Ln $RobInt_{it}$	0.031**	0.036***	0.014	0.014	0.014***	0.013***
	[2.427]	[2.732]	[1.083]	[1.082]	[3.165]	[2.923]
Obs.	6,200	6,155	6,195	6,150	6,373	6,329
R2	0.81	0.82	0.87	0.87	0.88	0.88
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes

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. Sector refers to 3-digit sectors. The control variables included in columns (2), (4) and (6) are log sales and dummies for whether the firm is an importer or an exporter, observed in the first year in which the firm appears in the sample and interacted with a full set of year dummies. Standard errors are corrected for clustering within firms; t-statistics are reported in squared brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

in firm efficiency and shifts in labor demand towards high-skill workers.[17]

While $\ln RobInt_{it}$ is less likely to be influenced by demand shocks than the binary indicator $Adoption_{it}$, it may not fully neutralize these shocks, for instance, if the latter do not affect robot capital and total capital proportionately. In addition, $\ln RobInt_{it}$ could be correlated with other firm characteristics that affect outcomes directly. Accordingly, the estimates in Table 3 do not have a causal interpretation. Neveretheless, they suggest that demand shocks seem to play an important role, significantly contaminating the estimated relationship between automation and firm-level outcomes. Motivated by this evidence, in the next section, we exploit differential cross-firm variation in robot exposure, exclusively stemming from pre-determined technological characteristics, to purge away demand shocks and identify a causal effect of automation on firm-level outcomes.

[17] As shown in Appendix D, these results are very similar if we compute the stock of robot capital using a perpetual inventory method with an annual depreciation rate of 15 percent, which falls within the range of depreciation rates normally assumed for robots in manufacturing (see, e.g., Graetz and Michaels, 2018).

5 The Effect of Robot Exposure on Firm-Level Outcomes

Our model suggests that the predisposition of firms to automate depends on the interaction of two technological characteristics: cost of machines and replaceability of employment. In particular, a higher cost-advantage of machines should stimulate robot adoption, and through this mechanism affect outcomes, relatively more in firms whose employment can more easily be replaced by robots. In this section, we draw on this insight to build a proxy for robot exposure whose variation exclusively depends on this interplay and not on demand shocks. We first study how robot exposure affects outcomes and discuss the main threats to identification. Then, we show that robot exposure is a significant predictor of robot adoption by firms. Using these results, we quantify the effect that the variation in robot adoption induced by robot exposure, as opposed to demand shocks, exerts on firm-level outcomes. Consistent with the lumpy nature of the investment in robots, most of the variation in robot adoption in our data is across firms rather than within firms over time. Accordingly, from now on, we focus on long-run changes in adoption and outcomes, and exploit cross-firm variation.

5.1 Variables and Specification

To construct the proxy for robot exposure, we follow two insights. First, the different nature of the production process across industries makes production easier to automate in some industries than in others. Second, within a given industry, some firms are more prone to automate production than other firms, because they perform activities that are relatively easier to assign to robots. Accordingly, we obtain the variable $RobExp$ as the interaction between a proxy for how suitable production is for automation in an industry, $RobSuit$, and a proxy for the ease with which robots can replace worker activities within a given firm, $Repl$. We now explain these variables in detail.

For each firm i operating in a given 5-digit NACE industry j, $RobSuit$ is defined as the average robot intensity of all firms $i' \neq i \in j$, namely,

$$RobSuit_{j-i} = \sinh^{-1}\left(\frac{\sum_{i' \neq i \in j} RobStock_{i'}}{\sum_{i' \neq i \in j} CapStock_{i'}}\right), \tag{19}$$

where $RobStock_{i'}$ and $CapStock_{i'}$ denote, respectively, the initial stock of robots and the initial total capital stock of firm $i' \neq i \in j$. The hyperbolic sine transformation preserves the zeros. Industries in which this ratio is higher are relatively more suitable for automation.

As for *Repl*, we follow Graetz and Michaels (2018) and exploit cross-firm differences in the prevalence of tasks that can be assigned to robots. Our measure is similar to the Graetz and Michaels (2018) indicator but is defined at the firm-level rather than at the industry level. To build this measure, we start by sourcing from Graetz and Michaels (2018) information on whether each of 377 US Census occupations is replaceable or not. The authors define an occupation as replaceable if its title corresponds to at least one of the robot application categories identified by the IFR, such as welding, painting, and assembling. Then, we manually map each US Census occupation into the 29 French occupations for which we have employment data in 1994. Using this information, we construct the firm-level replaceability measure as follows:

$$Repl_i = \sum_{o=1}^{29} \omega_{oi} \cdot Repl_o, \tag{20}$$

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.

Finally, *RobExp* is obtained as

$$RobExp_i = RobSuit_{j-i} \cdot Repl_i. \tag{21}$$

This variable captures variation in robot exposure across firms that operate in industries with different suitability for automation and exhibit a different prevalence of automatable tasks in production. Because $RobExp_i$ is solely based on pre-determined technological characteristics, its variation is not contaminated by demand shocks concurrent to the growth in outcomes.[18]

With these variables in hand, our empirical approach consists of estimating specifications of the following form:

$$\Delta Y_i = \alpha_h + \beta_1 \cdot RobExp_i + \beta_2 \cdot RobSuit_{j-i} + \beta_3 \cdot Repl_i + \mathbf{X}_i' \cdot \boldsymbol{\gamma} + \varepsilon_i, \tag{22}$$

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_h are 3-digit sector 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 eliminates year-to-year variation, so

[18] See Appendix D (Table D3) for descriptive statistics on $RobExp_{j-i}$ and the other variables used in this section.

identification exploits cross-sectional (across firms) variation in the pre-determined level of robot exposure and in the long-run growth of outcomes. Specifically, because the sector fixed effects, α_h, absorb differential trends in outcomes across sectors, the coefficients β_1 are identified from differential robot exposure across firms and 5-digit industries within the same 3-digit sector: firms that are more exposed to robots in a sector are those with higher levels of $Repl_i$ operating in industries with higher levels of $RobSuit_{j-i}$. The covariates, $\mathbf{X}_i$, further remove heterogeneous trends across firms characterized by different initial conditions within the same sector. 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.

Being based on pre-determined technological characteristics, $Repl_i$ and $RobSuit_{j-i}$ do not respond to subsequent changes in firm-level outcomes. However, these variables could still be correlated with other firm or industry variables affecting the outcomes of interest. Moreover, neither $Repl_i$ nor $RobSuit_{j-i}$ alone fully captures a firm's exposure to robots. For instance, the fact of operating in an industry in which production is highly suitable for automation may not matter much for firms whose tasks cannot be easily assigned to robots. For these reasons, our empirical approach goes beyond the simple level effects of these variables and identifies instead their interaction in the spirit of a difference-in-differences specification. This better reflects the idea that robot exposure depends on the interplay between replaceability and automation suitability, rather than on each of these characteristics in isolation. Moreover, 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. After presenting the baseline results, we go back to our empirical strategy and show that the estimates of β_1 are unlikely to be driven by the main remaining threats to identification.

5.2 Baseline Results and Sensitivity Analysis

To help comparing the long-differences results with our preliminary evidence, we first estimate eq. (22) replacing $RobExp_i$ with a binary indicator for adoption, $Adopter_i$. This is a dummy that 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 coefficients on $Adopter_i$ reflect cross-sectional differences in the growth of outcomes between robot adopters and other firms. The results are reported in Table 4, where standard errors are corrected for clustering within 5-digit industries to account for possibly correlated shocks among firms in the same industry. The dependent variables, indicated in the columns' head-

Table 4: Firm-Level Outcomes and Robot Adoption, Long Differences (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Ln Sales		Δ Ln No. of Employees		Δ Ln Sales per Worker	
$Adopter_i$	2.590***	4.549***	2.185***	2.366***	0.457	2.128***
	[6.482]	[10.898]	[6.837]	[7.505]	[1.128]	[4.626]
Obs.	36,301	36,301	36,606	36,584	36,301	36,301
R2	0.06	0.08	0.03	0.03	0.04	0.06
	Δ Ln VA per Worker		Δ Ln TFP		Δ Empl. Sh. High Skill	
$Adopter_i$	0.031	1.719***	0.493	1.976***	0.140***	0.004
	[0.075]	[3.986]	[1.308]	[5.328]	[3.009]	[0.100]
Obs.	35,180	35,180	33,623	33,623	36,606	36,584
R2	0.03	0.04	0.04	0.05	0.02	0.03
Controls	Sector FE	All Controls	Sector FE	All Controls	Sector FE	All Controls

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. Sector fixed effects are dummies for 3-digit sectors. The control variables included in columns (2), (4) and (6) are the initial firm-level employment share of occupations that can be replaced by robots ($Repl_i$), the initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry j ($RobSuit_{j\text{-}i}$), and the initial values of log sales and 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.

ings, are log changes in sales, employment, sales per worker, value added per worker and TFP, and the change in the employment share of high-skill professions; all changes are multiplied by 100 to express them in percentages.[19] For each outcome, the table presents results both from a specification including only the sector fixed effects α_h (odd-numbered columns) and from the complete specification including also the control variables $\mathbf{X}_i$ (even-numbered columns). Consistent with our preliminary evidence, 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.

These results may be biased by demand shocks if the latter influence both adoption and outcomes. Hence, we now turn to estimating eq. (22). The results are reported in Table 5, where each column refers to a different outcome. In the regression for log employment, the coefficient on $Repl_i$ is negative and statistically significant in line with the model. More importantly, the coefficient on $RobExp_i$ is also negative and precisely estimated, implying that firms that are more exposed to robots, owing to the interplay between their pre-sample

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

Table 5: Firm-Level Outcomes and Robot Exposure, Long Differences

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Ln Sales	Δ Ln No. of Employees	Δ Ln Sales per Worker	Δ Ln VA per Worker	Δ Ln TFP	Δ Empl. Sh. High Skill
$RobExp_i$	0.043	-0.094**	0.185**	0.197**	0.138*	0.009***
	[0.517]	[-2.589]	[2.381]	[2.276]	[1.726]	[2.688]
$Repl_i$	-1.203	-3.873***	3.685**	3.171*	1.338	-0.116*
	[-0.697]	[-5.572]	[2.371]	[1.868]	[0.839]	[-1.673]
$RobSuit_{j-i}$	45.832	-47.570***	99.543	68.418	39.260	13.211***
	[0.754]	[-4.584]	[1.521]	[0.994]	[0.649]	[9.548]
Ln Initial $Sales_i$	-1.444***	-0.009	-1.374***	-1.392***	-1.216***	0.071***
	[-11.245]	[-0.167]	[-10.165]	[-10.616]	[-10.345]	[10.250]
Dummy Initial $Importer_i$	1.501***	0.127	1.409***	1.458***	1.410***	0.048**
	[6.370]	[0.927]	[6.766]	[6.887]	[7.049]	[2.369]
Dummy Initial $Exporter_i$	0.637***	-0.480***	1.145***	1.039***	0.914***	0.067***
	[3.057]	[-3.523]	[5.655]	[4.852]	[4.910]	[3.699]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.07	0.03	0.06	0.04	0.05	0.03
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes

The subscripts *i* and *j* denote firms and 5-digit industries, respectively. In each regression, the dependent variable is 100 x the annualized change in the firm-level outcome indicated in the corresponding column. $RobExp_i$ 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* ($RobSuit_{j-i}$). Sector fixed effects are dummies for 3-digit sectors. 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.

specialization in automatable tasks and their industry's initial suitability for automation, experience a relatively larger and statistically significant reduction in employment over the sample period. Comparing this result with the regressions reported in Table 4 indicates that this effect is masked when using direct measures of adoption, due to the confounding role of demand shocks.

Turning to the other outcomes, Table 5 exhibits positive and statistically significant coefficients on $RobExp_i$ in the regressions for the log changes in sales per worker, value added per worker, and TFP. The table also shows a positive and precisely estimated coefficient on $RobExp_i$ in the regression for the change in the employment share of high-skill professions. Hence, firms that are more exposed to robots experience relatively larger increases in efficiency and relatively stronger shifts in labor demand towards high-skill workers. Interestingly, the effect of robot exposure on total sales, albeit positive, is not statistically significant. While reinforcing the view that demand shocks bias the relation between robot adoption and firm size, this result further suggests that the productivity gains from automation may not translate into higher sales.

Table 6 contains an extensive sensitivity analysis on the previous results. In panel a), we weigh the observations by the initial number of employees in each firm. The estimated coefficients are very similar to those obtained from unweighted regressions, suggesting that our evidence does not change if large firms are given greater weight. In panel b), we further exclude firms in the "Manufacturing of Motor Vehicles" industry, as the latter is one of the largest users of robots and our data lack information for the two biggest French car manufacturers. The qualitative and quantitative pattern of results is however unchanged. In panel c), we extend the notion of robots in the definition of automation suitability to include not just industrial robots (CN 84795000) but all types of machinery designed for lifting, handling, loading, unloading and welding (CN 842489, 842890, 851580, 851531, 851521 and 84864). In this case, the estimated coefficients on robot exposure lose precision and the coefficient in the employment regression turns positive, although it is not statistically significant. Robot exposure still leads to stronger shifts in labor demand towards high-skill professions. These results are consistent with the notion that broader forms of capital equipment are in general complementary to labor, as found in Aghion et al. (2019), and especially to high-skill workers.

Next, we study how the effect of robot exposure varies with the elasticity of demand. 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. As a result, employment should be less likely to fall at firms facing a more elastic demand. While a similar point has been made in Bessen (2019), who also provides some historical evidence from US industries, our rich firm-level data seem especially well suited to address this question. Hence, we extend eq. (22) by interacting $RobExp_i$, $Repl_i$, and $RobSuit_{j-i}$ with the elasticity of substitution in each sector.[20] Consistently with the model, the results reported in panel d) show 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. The point estimate implies that the effect of robot exposure switches sign, from negative to positive, when the elasticity of demand is higher than 10, i.e., in roughly 10 percent of the sectors in our sample. Finally, the effect of robot exposure on productivity is also relatively stronger in sectors where products are more substitutable.

In the last robustness check, we use an alternative proxy for robot exposure, which is obtained by replacing $RobSuit_{j-i}$ in eq. (21) with the initial value of the log stock of installed

[20] We source estimates of the elasticity of substitution from Broda and Weinstein (2006). We use estimates based on US import data over the 1990-2001 period for 5-digit SITC Rev. 3 codes and convert them to 3-digit NACE sectors. The linear term in the elasticity of substitution is subsumed by the sector fixed effects α_h.

Table 6: Firm-Level Outcomes and Robot Exposure, Long Differences (Robustness Checks)

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Ln Sales	Δ Ln No. of Employees	Δ Ln Sales per Worker	Δ Ln VA per Worker	Δ Ln TFP	Δ Empl. Sh. High Skill
	a) Weighted Regressions					
$RobExp_i$	0.027	-0.117***	0.191**	0.209**	0.144*	0.009***
	[0.299]	[-3.263]	[2.269]	[2.259]	[1.691]	[2.931]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.07	0.04	0.05	0.04	0.04	0.04
	b) Excluding Manufacturing of Motor Vehicles					
$RobExp_i$	0.042	-0.095***	0.183**	0.196**	0.136*	0.008**
	[0.498]	[-2.602]	[2.356]	[2.263]	[1.708]	[2.460]
Obs.	35,759	36,040	35,759	34,647	33,109	36,040
R2	0.07	0.03	0.06	0.04	0.05	0.03
	c) Broader Definition of Robot Imports					
$RobExp_i$	0.105	0.007	0.128	0.127	0.081	0.010***
	[1.287]	[0.163]	[1.596]	[1.530]	[1.080]	[3.143]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.07	0.03	0.06	0.04	0.05	0.03
	d) Interactions with Demand Elasticity					
$RobExp_i$	-0.472***	-0.207**	-0.241	-0.389**	-0.427**	-0.000
	[-2.693]	[-2.309]	[-1.389]	[-2.003]	[-2.241]	[-0.052]
$RobExp_i$ x $Elast_h$	0.123***	0.025*	0.104***	0.138***	0.132***	0.002
	[3.546]	[1.698]	[2.962]	[3.382]	[3.242]	[1.189]
Obs.	32,427	32,679	32,427	31,365	29,956	32,679
R2	0.08	0.03	0.06	0.05	0.06	0.04
	e) Alternative Proxy for Robot Exposure (IFR)					
$RobExp_i$	3.545***	0.320	3.424***	3.557***	3.294***	0.066**
	[12.243]	[1.491]	[13.569]	[12.659]	[11.887]	[2.506]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.08	0.03	0.06	0.05	0.05	0.03

The subscripts i and h denote firms and 3-digit sectors, respectively. In each regression, the dependent variable is 100 x the annualized change in the firm-level outcome indicated in the corresponding column. With the exception of panel e), $RobExp_i$ 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 e), $RobExp_i$ is constructed by replacing $RobSuit_{j\text{-}i}$ with the log stock of installed robots in each firm's sector; data on the stock of installed robots are sourced from the International Federation of Robotics (IFR) and available for 13 manufacturing sectors. The regressions in panel a) are weighted by the initial number of employees in each firm. The sample in panel b) excludes firms in the "Manufacturing of Motor Vehicles" industry. In panel c), robot imports include CN codes 842489, 842890, 851580, 847950, 851531, 851521 and 84864. In panel d), $Elast_h$ is the elasticity of demand, defined at the 3-digit sector level; 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 dummies for importing and exporting firms, and 3-digit sector 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.

robots in each firm's sector in the US. We source this variable from the IFR. Being based on overall robot installations, rather than robot imports, this variable is likely to provide a better proxy for automation suitability in industries where firms predominantly source robots from domestic suppliers. In particular, it is immune to the issues of "missing robots" or "false positives". Morever, being based on data for the US rather than France, this variable further allays concerns related to endogeneity and omitted industry characteristics. At the same time, a key drawback of this variable is that the IFR data are only available for 13 aggregate manufacturing sectors, so variation is much more limited than for the baseline measure of automation suitability.

Despite these differences, the qualitative evidence is preserved also with the alternative proxy for robot exposure. In particular, compared to the OLS estimates for robot adoption shown in Table 4, the coefficient on $RobExp_i$ in the regression for the log change in employment is not statistically significant, consistent with demand shocks inducing an upward bias in the relation between robot adoption and employment. At the same time, the alternative proxy confirms the positive effects of robot exposure on efficiency, the skill composition of labor demand and now even on sales.

Overall, this sensitivity analysis reassures that the main results do not hinge on specification details, sample composition and definitions of variables, and that they square reasonably well with additional predictions of the model. In the next section, we deal with the main remaining threats to identification and discuss how these threats could influence the results.

5.3 Threats to Identification

Our identification strategy requires that, conditional on the fixed effects and control variables included in eq. (22), $RobExp_i$ 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 discuss the implications of these threats, we now extend the baseline 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 7, we add the interaction between $RobSuit_{j-i}$ and the routine in-

Table 7: Firm-Level Outcomes and Robot Exposure, Long Differences (Threats to Identification)

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Ln Sales	Δ Ln No. of Employees	Δ Ln Sales per Worker	Δ Ln VA per Worker	Δ Ln TFP	Δ Empl. Sh. High Skill
	a) Interaction of Robot Suitability with Routine Intensity					
$RobExp_i$	0.044	-0.092**	0.183**	0.193**	0.134*	0.008**
	[0.529]	[-2.547]	[2.372]	[2.251]	[1.686]	[2.571]
$RobSuit_{j-i}$ x $Routine_i$	11.123	8.213	3.297	24.093	22.625	0.527
	[0.436]	[1.559]	[0.124]	[0.908]	[0.971]	[1.166]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.07	0.03	0.06	0.04	0.05	0.03
	b) Interactions of Robot Suitability with Firm Characteristics					
$RobExp_i$	0.023	-0.094**	0.165**	0.171**	0.110	0.010***
	[0.284]	[-2.578]	[2.134]	[1.986]	[1.386]	[2.890]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.07	0.03	0.06	0.04	0.05	0.03
	c) Interactions of Replaceability with Firm Characteristics					
$RobExp_i$	0.082	-0.096**	0.230***	0.244***	0.182**	0.010***
	[0.903]	[-2.594]	[2.624]	[2.639]	[2.130]	[2.886]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.08	0.03	0.06	0.05	0.05	0.03
	d) Interactions of Replaceability with Industry Characteristics					
$RobExp_i$	0.062	-0.114***	0.226**	0.252***	0.186**	0.012***
	[0.685]	[-3.092]	[2.593]	[2.761]	[2.224]	[3.327]
Obs.	36,254	36,537	36,254	35,134	33,579	36,537
R2	0.08	0.03	0.06	0.05	0.05	0.03
	e) Interactions of Replaceability with Sector Dummies					
$RobExp_i$	0.026	-0.094**	0.166**	0.172**	0.118	0.010***
	[0.306]	[-2.506]	[2.092]	[1.975]	[1.471]	[2.875]
Obs.	36,301	36,584	36,301	35,180	33,623	36,584
R2	0.08	0.03	0.06	0.04	0.05	0.03

The subscripts i and j denote firms and 5-digit industries, respectively. In each regression, the dependent variable is 100 x the annualized change in the firm-level outcome indicated in the corresponding column. $RobExp_i$ 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 panel b) and c) include interactions of $RobSuit_{j-i}$ and $Repl_i$, respectively, with the initial values of log sales and dummies for importing and exporting firms. The specification in panel d) includes the initial values of sectoral exports, imports, export unit value, import unit value, capital goods imports 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 dummies for importing and exporting firms, and 3-digit sector 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.

tensity of each firm.[21] While routine intensity is known to be correlated with the adoption of new 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. Accordingly, we expect the new interaction to have no significant effect on outcomes, and its inclusion to leave the evidence on $RobExp_i$ unaffected. The results do indeed show that the coefficients on the new interaction are not precisely estimated for any outcome, and that the main findings remain unchanged after controlling for this variable.

The previous exercise suggests that our main results are not capturing the possible correlation of $Repl_i$ with other firm characteristics associated to the introduction of technologies other than robots. In a similar spirit, 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, suggesting that the coefficients on $RobExp_i$ is not driven by the potential interplay of $RobSuit_{j-i}$ with firm size and trade orientation. 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 $RobSuit_{j-i}$, and that such an interplay confounds the effects of $RobExp_i$. In a first exercise, presented in panel d), we add interactions between $Repl_i$ and various industry characteristics, namely: (i) total imports and exports, to account for differences in import competition and export opportunities across industries; (ii) the average unit value of imports, to accommodate cross-industry differences in the cost of sourcing inputs from abroad; (iii) the average unit value of exports, to account for cross-industry differences in product characteristics such as quality; (iv) capital goods imports and (v) technology goods imports, to control for other forms of foreign technologies. 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.[22] Controlling for these interactions leaves the coefficients on $RobExp_i$ close to the baseline estimates, suggesting that our evidence is not confounded by other industry characteristics possibly interacting with replaceability.

[21] 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.

[22] Each of these characteristics also enter the specification linearly (coefficients unreported).

To further strengthen this view, in a final exercise, we add interactions between $Repl_i$ and a full set of 2-digit sector dummies. Contributing to the identification of the coefficients β_1 is now only the remaining variation in $RobSuit_{j-i}$ across narrow (5-digit) industries within the same 2-digit sector. As shown in panel e), our main results are qualitatively and quantitatively unchanged also in this case. Overall, the results presented this section largely allay the concern that our findings could be driven by omitted firm or industry characteristics interacting with the two components of robot exposure, and thereby give more credibility to the key identifying assumption underlying our empirical strategy.

5.4 Mechanism and Economic Magnitude

So far, the analysis has unveiled a robust effect of robot exposure on firm-level outcomes but has remained silent on the mechanisms through which this effect could take place. According to the model, the key channel through which robot exposure works is by inducing firms to adopt robots. In this section, we show that this mechanism is at work in our data.

To this purpose, we estimate eq. (22) using $Adopter_i$, the dummy for firms that adopt robots over the sample period, as the dependent variable. The baseline results are presented in column (1) of Table 8. The coefficients on $Repl_i$ and $RobSuit_{j-i}$ are both positive and precisely estimated, implying that robot adoption is relatively higher at firms performing more automatable tasks in the pre-sample period as well as in industries in which production is more suitable for automation. Robot adoption is also positively affected by initial firm sales, showing that initially larger firms adopt more robots in subsequent years, as predicted by the model. Finally, and crucially, the coefficient on $RobExp_i$ is positive and very precisely estimated, implying that firms that are more exposed to robots due to the interplay between replaceability and automation suitability do indeed show a greater tendency to adopt robots in subsequent years.

In the remainder of the table, we submit these results to a series of robustness checks. In column (1), we report estimates from a weighted regression using the initial number of employees in each firm as weights; in column (2), we exclude firms in the motor vehicle industry; and, in column (3), we use the alternative proxy for robot exposure based on the stock of installed robots in the US from the IFR.[23] In all cases, the coefficient on $RobExp_i$ is positive, very precisely estimated, and essentially unchanged compared to the baseline

[23] In column (3), the linear term in $RobSuit_{j-i}$ is absorbed by the sector fixed effects, as the stock of installed robots from the IFR is only available at a higher level of industry aggregation (13 aggregate sectors).

Table 8: Robot Exposure and Robot Adoption, Long Differences

	(1)	(2)	(3)	(4)
$RobExp_i$	0.002***	0.003***	0.002***	
	[3.128]	[3.132]	[3.110]	
$RobExp_i$ (IFR)				0.007***
				[3.704]
$Repl_i$	0.037***	0.051***	0.037***	0.001
	[2.960]	[3.007]	[2.967]	[0.225]
$RobSuit_{j-i}$	0.327**	0.483***	0.331**	
	[2.253]	[2.831]	[2.257]	
Ln Initial $Sales_i$	0.013***	0.017***	0.013***	0.013***
	[7.347]	[7.424]	[7.308]	[7.315]
Dummy Initial $Importer_i$	-0.001	-0.003*	-0.001	-0.001
	[-0.427]	[-1.685]	[-0.474]	[-0.474]
Dummy Initial $Exporter_i$	0.001	-0.001	0.001	0.001
	[0.444]	[-0.376]	[0.544]	[0.613]
Obs.	36,584	36,584	36,040	36,584
R2	0.04	0.06	0.04	0.04
Sector FE	Yes	Yes	Yes	Yes
Specification	Baseline	Weighted	No Motor Vehicles	Robot Exposure (IFR)

The subscripts i and j denote firms and 5-digit industries, respectively. 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 initial ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry j ($RobSuit_{j-i}$). $RobExp_i$ *(IFR)* is constructed analogously, replacing $RobSuit_{j-i}$ with the log stock of installed robots in each firm's sector from the International Federation of Robotics. Sector fixed effects are dummies for 3-digit sectors. 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.

estimate.

The above results lend support to the predictions of the model, by showing that robot exposure is a significant and robust predictor of robot adoption by firms. This does not rule out the possibility that robot exposure affects outcomes through other channels. However, if one is willing to *assume* that this is the only mechanism at work (exclusion restriction), then our results can be interpreted in a standard Two-Stage Least Squares (2SLS) framework, in which $Adopter_i$ is the endogenous regressor in eq. (22) and $RobExp_i$ is the excluded instrument. Recast in this way, the estimates in Table 5 would correspond to the reduced-form coefficients and those in Table 8 to the first-stage coefficients. The Kleibergen-Paap F-statistic for excluded instruments would be around 10, the conventional rule-of-thumb threshold for instrument relevance. Then, the ratio between the coefficients on $RobExp_i$ in Table 8 and the corresponding coefficient in Table 5 would give the 2SLS estimate of the effect of robot adoption on firm-level outcomes.

With the usual *caveat* about the exclusion restriction and taking into account the moderate strength of the instrument, we now use this framework to have a sense of the magnitude of the effect of "exogenous" adoption relative to that driven by demand shocks. Considering the specification for the log change in employment, the 2SLS coefficient on $Adopter_i$ would be equal to -49.481, with a t-statistic of -2.025. By comparing this coefficient with its OLS counterpart in Table 4 (2.366), we can measure how much of the correlation between robot adoption and employment changes is due to exogenous automation and how much reflects instead demand shocks. Following Autor, Dorn and Hanson (2013), the OLS coefficient on $Adopter_i$, β_{OLS}, can be decomposed as follows:

$$\beta_{OLS} = \beta_{IV} \times \frac{\sigma_{IV}^2}{\sigma^2} + \beta_{RES} \times \frac{\sigma_{RES}^2}{\sigma^2},$$

where β_{IV} is the second-stage coefficient on $Adopter_i$, (σ_{IV}^2/σ^2) is the fraction of the overall variance of $Adopter_i$ explained by the fitted values of the first-stage regression (exogenous adoption), and $(\sigma_{RES}^2/\sigma^2)$ is the residual fraction explained by demand shocks (endogenous adoption).

We estimate (σ_{IV}^2/σ^2) to be equal to 4.3 percent in our data, implying that robot adoption is largely driven by demand shocks. Using these numbers along with the estimates of β_{OLS} and β_{IV} yields $\beta_{RES} = 4.696$. Accordingly, exogenous adoption explains an average annual fall in employment equal to 2.13 percent in robot adopters relative to non robot adopters (i.e., $\beta_{IV} \times (\sigma_{IV}^2/\sigma^2)$). Residual adoption due to demand shocks is instead associated with an average annual increase in employment equal to 4.49 percent in the former group of firms relative to the latter (i.e., $\beta_{RES} \times (\sigma^2_{RES}/\sigma^2)$).

5.5 Additional Results

An interesting pattern that emerged from the preliminary analysis and was confirmed by the results presented in the last section is that, in spite of a positive and robust effect on firm efficiency, robot exposure has a modest and generally imprecisely estimated effect on total sales. A possible explanation for this result is that the productivity gains from automation do not entirely translate into lower prices (hence higher sales) because automation induces an increase in the firm's market power. In Appendix B, we formalize this explanation using an extension of the model in which markups are endogenous and respond to automation. In this section, we provide suggestive empirical evidence, focusing on a smaller sample of firms for which we have information on additional outcomes.

We start by studying how firm profits respond to automation. If the latter increased market power, it should also lead to an increase in profits. In column (1) of Table 9, we estimate eq. (22) using the log change in profits as the dependent variable and the dummy for adopters, $Adopter_i$, in place of $RobExp_i$. The coeffient on $Adopter_i$ is positive and statistically significant, implying that firms that adopt robots over the sample period experience a faster increase in profits compared to other firms. In columns (2) and (3), we replace $Adopter_i$ with $RobExp_i$. The coefficient on this variable is also positive and precisely estimated, consistent with a positive effect of robot exposure on profits.

These results are consistent with the view that automation increases market power. The rise in market power should compensate the reduction in marginal costs, resulting into a moderate response of prices to automation. In the remaining columns of Table 9, we thus study the relation between automation and prices. Since our data, like the typical firm-level dataset, does not contain information on prices for domestic sales, we exploit data on export unit values and compute an average export price for each firm.[24] We then estimate the same specifications as in columns (1)-(3), using the log change in export prices as the dependent variable; the sample is now restricted to exporting firms. The coefficients on $Adopter_i$ and $RobExp_i$ are generally imprecisely estimated, suggesting that automation has little impact on firm prices. While admittedly suggestive, the evidence in this section is consistent with

[24] This variable is constructed as the export-value-weighted average of the unit values of all 8-digit CN products exported by the firm.

Table 9: Robots, Profits and Export Prices, Long Differences

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Ln Profits			Δ Ln Export Price Index		
$Adopter_i$	2.844***			1.051*		
	[3.207]			[1.914]		
$RobExp_i$		0.199**			0.098	
		[1.993]			[1.579]	
$RobExp_i$ (IFR)			1.850***			-0.112
			[3.488]			[-0.176]
$Repl_i$	-1.672*	2.032	-1.762**	-0.114	1.678	-0.164
	[-1.887]	[0.910]	[-2.008]	[-0.183]	[1.233]	[-0.263]
$RobSuit_{j-i}$	-167.876	-175.135		125.806***	124.635***	
	[-1.265]	[-1.349]		[4.703]	[4.784]	
Ln Initial $Sales_i$	-1.339***	-1.263***	-1.197***	-0.077	-0.048	-0.070
	[-8.217]	[-8.071]	[-7.769]	[-0.936]	[-0.587]	[-0.867]
Dummy Initial $Importer_i$	1.845***	1.842***	1.799***	-0.129	-0.129	-0.129
	[4.240]	[4.218]	[4.160]	[-0.386]	[-0.388]	[-0.386]
Dummy Initial $Exporter_i$	0.601	0.566	0.554			
	[1.245]	[1.178]	[1.147]			
Obs.	18,572	18,572	18,572	16,018	16,018	16,018
R2	0.030	0.03	0.03	0.01	0.01	0.01
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes

The subscripts i and j denote firms and 5-digit industries, respectively. 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. $RobExp_i$ 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 ($RobSuit_{j-i}$). $RobExp_i$ *(IFR)* is constructed analogously, replacing $RobSuit_{j-i}$ with the log stock of installed robots in each firm's sector from the International Federation of Robotics. Sector fixed effects are dummies for 3-digit sectors. 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.

a view in which the productivity gains from automation are at least partly offset by an increase in market power, and thereby result into a moderate change in prices and sales and a concomitant increase in profits.

6 Conclusions

In this paper, we have documented how the adoption of industrial robots affects a number of firm-level outcomes using data from the universe of French firms observed between 1994 and 2013. To better inform our empirical strategy, we have built a model in which heterogeneous firms invest in automation. Robots save on production workers, but they also require non-production workers such as engineers and managers. A decline in the cost of capital induces firms to invest more in automation, with ambiguous effects on employment. On the one hand, machines displace workers; on the other hand, the increase in productivity raises the demand for all factors. Importantly, these effects vary across firms: since automation saves on the variable cost, firms facing a higher demand invest more in automation and are more likely to shed workers.

The model illustrates one challenge in testing the effect of automation on employment: demand shocks tend to generate a positive correlation between automation and employment even when exogenous changes in automation would lead to job losses. A second key challenge that researchers have faced so far is the measurement of automation at the firm level. The main contribution of this paper is to propose a solution to these difficulties. We have shown how data on firm imports of industrial robots can be used to build proxies for automation that are independent of demand shocks. Our rich data set allows us to document a number of empirical patterns.

First, we have shown that robot adopters differ significantly from other firms: they are larger, more productive and employ a higher share of high-skill workers. Over time, robot adoption occurs after periods of expansion in firm size, and is followed by improvements in firm efficiency and an increase in demand for low-skill workers. Guided by our theoretical model, we have then developed various empirical strategies to identify the causal effects of robot adoption. Our results suggest that, while demand shocks generate a positive correlation between robot adoption and employment, exogenous changes in automation lead to job losses, especially for low-skill workers.

We have also found that, while robot adoption increases significantly sales per worker, its effect on total sales is much less strong, suggesting that the efficiency gains do not always translate into an equivalent fall in prices. These results raise concerns on some possible negative effects of automation: besides the costly displacement of workers emphasized in the literature, our findings suggest that the productivity gains from automation may be partly offset by an increase in markups and that the widespread diffusion of automation, especially among already large firms, may have contributed to the rise of market power.[25]

While this paper is a first attempt at identifying the firm-level effect of the adoption of industrial robots, much remains to be done. First, in this paper we have focused attention to firms that import robots. However, it would also be interesting to study what happens to other firms in the same industry. In particular, robot adoption is likely to induce a reallocation of market shares away from non adopters. Given that these firms differ markedly in many dimensions, such a reallocation is likely to have significant effects on the demand for labor and welfare. Estimating and quantifying these industry-level adjustments seems an important step to fully understand the aggregate impact of automation.[26] Second, investigating more the effects of automation on market power seems also important. While we have found evidence consistent with the hypothesis that automation may grant market power, more direct evidence is needed. Third, this paper focuses on the micro-level adjustment, but we see it as a building block for a comprehensive study of the macroeconomic effects of automation (e.g., Jaimovich et al. 2021). In particular, more evidence on labor-market effects is likely to be useful for designing policies that could guarantee the benefits from new technologies to be fully realized and broadly shared. Given the speed of technological progress and its potentially disruptive effects, this is likely to become one of the most pressing challenges for advanced economies in the near future.

[25] On the recent rise of market power, see for instance De Loecker and Eeckhout (2017) and Autor et al. (2017).

[26] See Acemoglu, Lelarge and Restrepo (2020) and Koch, Manuylov and Smolka (2019), 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." volume 4B, *Handbook of Labor Economics*, chapter 12, 1043-1171. Elsevier.

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

[4] Acemoglu, Daron and Pascual Restrepo (2018a). "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment." *American Economic Review*, 108(6):1488–1542.

[5] Acemoglu, Daron, Claire Lelarge and Pascual Restrepo (2020). "Competing with Robots: Firm-Level Evidence from France," *American Economic Review, Papers and Proceedings*, forthcoming.

[6] Acemoglu, Daron and Pascual Restrepo (2018b). "Demographics and Automation," NBER Working Paper 24421.

[7] Acemoglu, Daron and Pascual Restrepo (2019). "Robots and Jobs: Evidence from the US Labor Markets," *Journal of Political Economy*, forthcoming.

[8] Ackerberg, Daniel, Kevin Caves and Garth Frazer (2015). "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83(6), 2411-2451.

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

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

[11] Autor, David and Anna Salomons (2017). "Robocalypse Now - Does Productivity Growth Threaten Employment?" Working paper, MIT.

[12] Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen (2017). "The Fall of the Labor Share and the Rise of Superstar Firms," *NBER Working Paper No. 23396.*

[13] Autor, David, Frank Levy and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4):1279-1333.

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

[15] Bessen, James (2019). "Automation and jobs: when technology boosts employment," *Economic Policy*, 34(100), 589–626.

[16] Bessen, James, Maarten Goos, Anna Salomons and Wiljan van den Berge (2019). "What happens to workers at firms that automate," working paper.

[17] Blanas, Sotiris, Gino Gancia, and Sang Yoon (Tim) Lee (2019). "Who is Afraid of Machines?" *Economic Policy*, 34(100), 627–690.

[18] Bonfiglioli, Alessandra, Rosario Crinò and Gino Gancia (2019). "Firms and Economic Performance: A View from Trade," CEPR Discussion Paper 12829.

[19] Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia and Ioannis Papadakis (2021). "Robots, Offshoring and Welfare" in *Robots and AI: a New Economic Era* (Lili Yan Ing and Gene M. Grossman eds.), Oxon and New York: Routledge, forthcoming.

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

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

[22] Caselli, Francesco and Alan Manning (2019). "Robot arithmetic: new technology and wages". *American Economic Review: Insights*, 1 (1). pp. 1-12.

[23] Cheng, Hong, Ruixue Jia, Dandan Li, and Hongbin Li (2019). "The Rise of Robots in China," *Journal of Economic Perspectives*, 33(2): 71–88.

[24] Chernozhukov, Victor and Christian Hansen (2008). "The Reduced Form: A Simple Approach to Inference with Weak Instruments," *Economics Letters*, 100(1): 68-71.

[25] Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum and Nicole Woessner (2018). "Adjusting to Robots: Worker-Level Evidence," Technical report, Opportunity & Inclusive Growth Institute, Federal Reserve Bank of Minneapolis.

[26] De Loecker, Jan and Jan Eeckhout (2017). "The Rise of Market Power and the Macroeconomic Implications," *NBER Working Paper No 23687.*

[27] De Loecker, Jan and Frederic Warzynski (2012). "Markups and Firm-Level Export Status," *American Economic Review*, 102(6), 2437-2471.

[28] Dechezlepretre, Antoine, David Hemous, Morten Olsen and Carlo Zanella (2019). "Automating Labor: Evidence from Firm-level Patent Data" working paper.

[29] Dinlersoz, Emin and Zoltan Wolf (2018) "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing," Census.

[30] Dixen, Jay, Bryan Hong and Lynn Wu (2019) "The Employment Consequences of Robots: Firm-level Evidence." Statistics Canada.

[31] European Commission (2016). "Analysis of the Impact of Robotic Systems on Employment in the European Union," Technical report.

[32] Frey, Carl B. and Michael A. Osborne (2017). "The Future of Employment: How Susceptible are Jobs to Computerisation?" Technological Forecasting and Social Change, 114(C):254-280.

[33] Graetz, Georg, and Guy Michaels (2018). "Robots at Work," *Review of Economics and Statistics* 100(5): 753–68.

[34] Hemous, David and Morten Olsen (2018). "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," working paper.

[35] Hottman, Colin, Stephen Redding and DavidWeinstein (2016). "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131, 1291–1364.

[36] Humlum, Anders (2019). "Robot Adoption and Labor Market Dynamics," working paper.

[37] Jaimovich, Nir, Itay Saporta Eckstein, Yaniv Yedid-Levy and Henry Siu (2021). "The Macroeconomics of Automation: Data, Theory, and Policy Analysis," *Journal of Monetary Economics,* forthcoming.

[38] Koch, Michael, Ilya Manuylov and Marcel Smolka (2019). "Robots and Firms," working paper.

[39] Korinek, Anton and Ding Xuan Ng (2018). "Digitization and the Macro-Economics of Superstars".

[40] Levinsohn, James and Amil Petrin (2003). "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 70(2), 317–341.

[41] Mann, Katja and Lukas Puttman (2017). "Benign Effects of Automation: New Evidence from Patent Texts," Unpublished manuscript.

[42] Martinez, Joseba (2021). "Putty Clay Automation," CEPR DP16022.

[43] Raj, Manav and Robert Seamans (2018). "AI, Labor, Productivity and the Need for Firm-Level Data" NYU Stern Working Paper.

[44] Wooldridge, Jeffrey M. (2009). "On estimating firm-level production functions using proxy variables to control for unobservables," *Economics Letters*, 104(3), 112-114.

[45] Zeira, Joseph (1998). "Workers, Machines, And Economic Growth," *Quarterly Journal of Economics*, 113(4), 1091–1117.

A Choice of Automation: Comparative Statics

Denote the marginal benefit and the marginal cost of automation as MB_i and MC_i, respectively. Then:

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

Profits are globally concave in κ_i when:

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

Under the assumption $(\sigma - 1) \ln \left(\frac{w}{r} \right) < 1/\rho_i$, this condition is alwasy satisfied at κ_i^*.

We derive here 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 (12), we obtain the implicit derivative of κ_i^* with respect to any parameter v as

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

As noted above, condition (14) implies that 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 Automation and Market Power

We now extend the model to incorporate the notion that automation may increase market power (e.g., Korinek and Ng, 2018). To keep the analysis as simple as possible, we consider a case in which firms set their price so as to keep potential competitors out of the market (limit pricing). Potential competitors can copy existing varieties, but they are less productive than the original producer. To make the equilibrium markup a function of κ_i, we assume that the production process of firms that use automation more intensively is harder to imitate. As a result, the wedge between the limit price and the marginal cost increase in κ_i. To capture the implications of this setup, we denote with $\mu(\kappa_i) \in (0, 1/\sigma)$ the profit share of revenue and assume $\mu'(\kappa_i) > 0$.[27]

Then, the labor demand in equation (7) becomes:

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

This expression shows that automation affects labor demand not only via the productivity and the displacement effects, but also through the increase in the markup, as it is made clear by the derivative:

$$\frac{dl_i/l_i}{d\kappa_i} = -\frac{\sigma \mu'(\kappa_i)}{1 - \mu(\kappa_i)} + (\sigma - 1) \ln \frac{w}{r} - \frac{1}{1 - \kappa_i}.$$

The endogenous reaction of markups dampens the productivity effect because the cost saving generated by automation is only partially transferred to prices.

The impact of κ_i on markups also affects the incentives to automate. In particular, κ_i is chosen to solve:

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

The first-order condition for automation becomes:

$$(\sigma - 1) p_i y_i \left[\mu(\kappa_i) \ln \left(\frac{w}{r}\right) + \left(\frac{1}{\sigma - 1} - \frac{\mu(\kappa_i)}{1 - \mu(\kappa_i)} \right) \mu'(\kappa_i) \right] = h (1 - \kappa_i)^{-1/\rho_i}. \tag{B1}$$

This equation shows that, as long as the markup is below the one that would be chosen without limit pricing ($\mu(\kappa_i) < 1/\sigma$), and $\mu'(\kappa_i) > 0$, then firms have an incentive to automate to increase their market power. This case introduces the possibility of "excessive" automation. For instance, if

$$\frac{\mu'(\kappa_i)}{1 - \mu(\kappa_i)} = \ln \left(\frac{w}{r}\right),$$

automation would be chosen only to increase profits, with no effect on prices and sales, and hence no gains to consumers.

[27] The main results would be qualitatively similar if we considered other models of imperfect competition in which the perceived demand elasiticy is a function of market shares.

C 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 (1) and (9), 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 firms with higher replaceability ρ_i, since these face a lower cost.

D Additional Results

Table D1 reports the estimation coefficients corresponding to the plots in Figure 4.

Table D1: Difference-in-Differences Event Studies

	(1)	(2)	(3)	(4)	(5)	(6)
	Ln Sales	Ln No. of Employees	Ln Sales per Worker	Ln VA per Worker	Ln TFP	Empl. Sh. High Skill
$Adoption_{it-4}$	0.062**	0.013	0.006	0.048*	0.059**	0.002
	[2.299]	[0.722]	[0.289]	[1.679]	[2.321]	[0.759]
$Adoption_{it-3}$	0.086***	0.062***	-0.024	0.022	0.053**	0.001
	[2.889]	[3.477]	[-1.277]	[0.792]	[2.085]	[0.217]
$Adoption_{it-2}$	0.105***	0.065***	-0.013	0.011	0.040	0.003
	[3.343]	[3.187]	[-0.630]	[0.372]	[1.458]	[1.319]
$Adoption_{it-1}$	0.149***	0.086***	-0.006	0.014	0.049*	0.005**
	[4.566]	[3.730]	[-0.255]	[0.466]	[1.862]	[2.108]
$Adoption_{it}$	0.199***	0.114***	0.013	0.015	0.048*	0.005*
	[5.960]	[4.821]	[0.599]	[0.526]	[1.847]	[1.660]
$Adoption_{it+1}$	0.210***	0.124***	0.008	0.000	0.042	0.009***
	[5.984]	[4.890]	[0.367]	[0.002]	[1.567]	[2.673]
$Adoption_{it+2}$	0.195***	0.085***	0.014	0.036	0.070**	0.012***
	[5.241]	[2.817]	[0.507]	[1.104]	[2.402]	[3.380]
$Adoption_{it+3}$	0.194***	0.041	0.061**	0.082**	0.094***	0.017***
	[5.093]	[1.261]	[2.227]	[2.293]	[3.075]	[3.991]
$Adoption_{it+4}$	0.177***	0.026	0.063**	0.072*	0.078**	0.018***
	[4.380]	[0.690]	[2.007]	[1.934]	[2.488]	[3.734]
$Adoption_{it+5}$	0.200***	0.040	0.052*	0.073**	0.089***	0.013***
	[4.827]	[1.158]	[1.914]	[2.088]	[2.803]	[2.848]
Obs.	689,855	593,320	591,939	581,726	570,812	593,320
R2	0.93	0.88	0.89	0.81	0.86	0.67
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

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. Sector refers to 3-digit sectors. Standard errors are robust to heteroskedasticity; t-statistics are reported in square brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

Table D2 replicates the estimates in Table 2 using $\ln RobInt_{it}$, with the stock of robot capital now constructed assuming a depreciation rate of 15 percent.

Table D2: Firm-Level Outcomes and Ln Robot Intensity with Depreciation, Panel (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Ln Sales		Ln No. of Employees		Ln Sales per Worker	
Ln $\text{RobInt}_{\text{it}}$	-0.074***	-0.090***	-0.097***	-0.097***	0.011	0.009
	[-3.111]	[-3.979]	[-5.096]	[-5.076]	[0.981]	[0.793]
Obs.	6,368	6,324	6,373	6,329	6,368	6,324
R2	0.96	0.97	0.93	0.93	0.90	0.90
	Ln VA per Worker		Ln TFP		Empl. Sh. High Skill	
Ln $\text{RobInt}_{\text{it}}$	0.022**	0.026**	0.013	0.013	0.010***	0.010***
	[2.097]	[2.409]	[1.180]	[1.168]	[3.051]	[2.761]
Obs.	6,200	6,155	6,195	6,150	6,373	6,329
R2	0.81	0.82	0.87	0.87	0.88	0.88
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes

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 depreciation rate of 15%. Sector refers to 3-digit sectors. The control variables included in columns (2), (4) and (6) are log sales and dummies for whether the firm is an importer or an exporter, observed in the first year in which the firm appears in the sample and 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.

Table D3 shows descriptive statistics for the main variables included in the long-differences specifications presented in Section 5.

Table D3: Descriptive Statistics - Long Differences Sample

	Obs.	Mean	Median	Std. Dev.
	Robot Adopters			
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 Sales per Worker	493	-0.083	-0.081	0.090
Δ Ln VA per Worker	470	-0.096	-0.094	0.095
Δ Ln TFP	460	-0.092	-0.086	0.083
Δ Ln Profits	300	-0.053	-0.061	0.189
Δ Ln Export Price Index	424	0.017	0.018	0.148
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 Sales per Worker	35,808	-0.097	-0.093	0.131
Δ Ln VA per Worker	34,710	-0.104	-0.101	0.141
Δ Ln TFP	33,163	-0.107	-0.098	0.123
Δ Ln Profits	18,272	-0.084	-0.077	0.377
Δ Ln Export Price Index	15,594	0.005	0.011	0.285
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 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*).