Robot Imports and Firm-Level Outcomes*

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

This draft: October 2020

Abstract

We use French data over the period 1994-2013 to study how imports of industrial robots affect firm-level outcomes. Compared to other firms operating in the same 5-digit sector, 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 changes in automation lead to job losses. We also find that robot imports increase productivity and the employment share of high-skill professions, but have 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 Paula Bustos, David Hémous, Joseba Martinez, Pascual Restrepo and seminar participants at the Barcelona GSE Summer Forum (2019), the ECB conference "Challenges in the digital era" and the University of Mannheim for useful comments.

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

‡Università Cattolica del Sacro Cuore, Dept. of Economics and Finance, CEPR and CESifo. Via Necchi 5, 20123, Milan, ITALY. E-mail: rosario.crino@unicatt.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 period 1994-2013 and to use a novel identification strategy to identify causality. Recently, researchers have turned to import data as a source of information on the usage of robots.

[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).

Although 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. Compared to other proxies used in the literature, such as dummies collected from survey data, the key advantage of robot imports is that they provide a precise measure of automation intensity that is available for the near universe of firms. With this rich data, we develop various empirical strategies to identify the causal effects of robot adoption on sales, productivity and employment within French firms.

To guide the analysis, we build a simple model in which heterogeneous firms invest in automation, whose effect is to replace workers with capital in a set of tasks. Automation saves on production workers, but it also requires non-production workers such as engineers and managers. As the cost of capital 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. Importantly, 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 also allows 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 their profits.

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 shows that a simple measure of automation *intensity,* namely expenditure on automation over the cost of capital, is independent of demand shocks and hence is more likely to capture the negative effect on employment. Besides the impact on the demand for labor, automation increases productivity, the relative demand for non-production workers and possibly markups.

We then take these predictions to the data. We start by documenting some descriptive patterns, focusing on the manufacturing sector, where the use of industrial robots is more prevalent. First, we show that robot adopters differ significantly from non-adopters. In particular, they are larger, more productive, and have a larger employment share of high-skill professions. But do robot adopters differ from other firms before importing robots, or do they start diverging afterwards? To shed light on this question, we use a difference-in-differences event study approach to analyze how 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 the productivity gains do not translate entirely into a fall in prices.

To identify the causal effects of robots, we next use two complementary strategies. First, we exploit yearly variation within firms, and regress various outcomes on a measure of robot intensity, which is defined as the ratio between the stock of robot imports and the total capital stock of the firm. According to our model, this measure purges away demand shocks. Second, we 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, which should determine the predisposition to automate.

In particular, we construct a novel instrument 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. Consistent with our model, this instrument 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 that are more prone to automate, such as firms whose production is more intensive in tasks that can be performed by robots. 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 robot adoption across firms that operate in industries with varying suitability for automation and exhibit a heterogeneous prevalence of automatable tasks in production.

The results for employment are particularly interesting. We find that while robot adoption and employment growth are correlated, an increase in robot intensity is followed by job losses. Similarly, 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 consistently find that importing robots leads to an increase in the employment share of high-skill professions and various measures of productivity, while the effects on total sales are much weaker.

The first-stage results also confirm the predictions of the model, thereby lending more credibility to our identification strategy: we find that firms performing more replaceable tasks in industries with a higher robot suitability, as well as larger firms, are more likely to start adopting robots in subsequent years. Finally, we show that the IV results are robust to controlling for other phenomena, such as trade and offshoring, that could have affected firms differentially depending on the replaceability of employment.

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. This implies that part of the gains for consumers must be muted by an increase in markups. To our knowledge, 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, 2019, Korinek and Ng, 2018, Caselli and Manning, 2019).

To our knowledge, this is the first paper that uses an IV strategy to identify the 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 International Federation of Robotics (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 period 1990–2005 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 robot exposure across local labor markets in Germany led to job losses in manufacturing that were however offset by gains in the service sector.[2]

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 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. Firm-level robot imports have been used by Humlum (2019) for Denmark, Dixen, Hong and Wu (2019) for Canada, and Acemoglu, Leclerc and Restrepo (2020) for France. Importantly, none of these papers uses a firm-level instrument 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 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 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 IV design to identify causality. The key difference is that they consider a much broader measure of capital inputs, which is likely to be complementary to labor. More 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 use various identification strategies to estimate the

[2] 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).

effect of robot imports on sales, employment, labor productivity, and the employment share of high-skill workers. 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.[3]

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

[3] The model adds firm heterogeneity to earlier contribution 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).

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_i the rental rate of capital and with w the wage of production workers. Since in reality different production processes use very different types of machines, we let the cost of capital equipment, r_i, be firm specific. We also assume $r_i < w$, 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 non-production workers (managers and engineers) with wage h. For now, we take f 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 to maximize profit,

$$\max_{k_i, l_i} \{p_i y_i - r_i k_i - w l_i - hf\},$$

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:

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

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

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

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

Substituting k_i back into the production function yields:

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

which shows that output per production worker is increasing in κ_i if $w > r_i$, 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:

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

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:

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

This derivative is positive for values of κ_i lower than $1 - \left[(\sigma-1)\ln(w/r_i)\right]^{-1}$. This condition is more likely to be satisfied when σ and w/r_i 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_i is high, the cost saving of automation is stronger. If instead $(\sigma - 1)\ln(w/r_i) < 1$, then the displacement effect always dominates.[4]

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

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

This equation confirms that automation raises output as long as capital is cheaper than production workers, $w/r_i > 1$. Moreover, substituting (4) and (5) into the profit function yields

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

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

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 to take the following form:

$$hf(\kappa_i, \rho_i) = \frac{h\kappa_i^{\delta}}{\rho_i \delta}, \qquad \delta > 1,$$

where δ denotes the convexity of the automation cost. The parameter ρ_i captures heterogeneity across firms in the (inverse) cost of automation. In particular, it can interpreted as an index of replaceability of task in the production process of firm i.

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

$$\max_{\kappa_i} \left\{ \frac{p_i y_i}{\sigma} - \frac{h\kappa_i^{\delta}}{\rho_i \delta} \right\}.$$

Automation poses a trade-off between its fixed cost and the reduction in the variable cost it generates. The first-order condition for κ_i is:

$$\left(1 - \frac{1}{\sigma}\right) p_i y_i \ln\left(\frac{w}{r_i}\right) = \frac{h\kappa_i^{\delta - 1}}{\rho_i}. \tag{10}$$

[4] 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.

The left-hand side of (10) is the marginal benefit of automation. It shows that the benefit of automation is increasing in the demand elasticity (σ), revenues ($p_i y_i$) and in the cost saving entailed by machines (w/r_i). The right-hand side is instead the marginal cost.

Substituting y_i from (9), the first-order condition for automation (10) becomes:

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

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

$$0 < A_i \left(1-\frac{1}{\sigma}\right)^{\sigma} \left(\frac{\varphi_i}{r_i}\right)^{\sigma-1} \ln\left(\frac{w}{r_i}\right) < \frac{h}{\rho_i}. \tag{11}$$

Clearly, if $w/r_i < 1$ there is no benefit of automation, hence the optimal κ_i is zero. If instead the cost of automation is too low, the firm will choose full automation, i.e., $\kappa_i = 1$. For $w > r_i$ and a sufficiently high cost of automation, as in (11), instead, there is an interior optimal level of κ_i. As we show in the Appendix, the comparative statics of the equilibrium choice of automation, κ_i^*, to changes in the exogenous parameters are:

$$\frac{d\kappa_i^*}{dA_i} > 0; \quad \frac{d\kappa_i^*}{d\varphi_i} > 0; \quad \frac{d\kappa_i^*}{d(w/r_i)} > 0; \quad \frac{d\kappa_i^*}{d(\rho_i/h)} > 0. \tag{12}$$

These results are intuitive. 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_i) and decreasing in its own cost h/ρ_i.

2.4 Extension: Automation and Market Power

We now extend the model to incorporate the notion that automation may increase market power (e.g., Martinez, 2019, and 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 $\mu'(\kappa_i) > 0$.[5]

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

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

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:

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

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:

$$\max_{\kappa_i} \left\{ \mu(\kappa_i) p_i y_i - \frac{h\kappa_i^{\delta}}{\rho_i \delta} \right\}.$$

The first-order condition for automation becomes:

$$(\sigma - 1) p_i y_i \left[\mu(\kappa_i) \ln\left(\frac{w}{r_i}\right) + \left(\frac{1}{\sigma-1} - \frac{\mu(\kappa_i)}{1-\mu(\kappa_i)}\right) \mu'(\kappa_i) \right] = \frac{h\kappa_i^{\delta-1}}{\rho_i}. \tag{13}$$

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

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

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

2.5 Empirical Implications

The model has clear predictions for the determinants of automation. These are summarized by the comparative statics results in (12). In sum, automation is increasing in demand (A_i and φ_i), in replaceability (ρ_i) and decreasing in the relative cost of capital (r_i/w) and the cost

[5] 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.

of non-production workers (h). These results are intuitive and consistent with the existing literature.[6] As to what may have caused the generalized increase in automation observed in aggregate data, the model highlights the pervasive decline in the relative cost of capital (r_i/w) as a natural candidate. However, it also suggests that the effect should be heterogeneous. Rearranging (10),

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

it can be seen that a decline in r_i/w has a stronger effect on automation in firms where tasks are more replaceable, as captured by the parameter ρ_i.

The implications of the model regarding the relationship between automation and employment are more nuanced. First, (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. Second, the model also illustrates the key challenge that the econometrician faces in answering such a question, which hinges on the endogeneity of κ_i: demand shocks, captured by A_i and φ_i, have a direct positive effect on employment, but they also trigger automation. Hence, demand shocks may generate a positive correlation between automation and employment, even if, conditional on demand, and increase in κ_i would lead to job losses. Firm and sector-year fixed effects are not sufficient to solve the problem because demand shocks are likely to vary both across firms and over time.[7]

Fortunately, the model also offers possible remedies to this bias. Exogenous shocks to the costs and benefits of automation (ρ_i and r_i) can be used to isolate variation in κ_i that is orthogonal to demand shocks. To identify firm-specific shocks to r_i, the model suggests to use *automation intensity* defined as the cost of automation, $h\kappa_i^\delta/(\rho_i\delta)$, over capital expenditure, $r_i k_i$. Using the first-order condition for k_i (5), into the first-order condition for automation (10), we can write:

$$\frac{h\kappa_i^\delta}{\rho_i \delta r_i k_i} = \frac{1}{\delta} \ln\left(\frac{w}{r_i}\right). \tag{14}$$

This equation shows immediately that automation intensity can be used to identify variation in automation that is independent of demand shocks. Controlling for firm and sector-year fixed effects should also purge this measure from any variation that is not driven by firm-

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

[7]See, for instance, Hottman, Redding and Weinstein (2016) and Bonfiglioli, Crinò and Gancia (2019) for the importance of firm-level demand shocks.

specific changes in the cost of capital, r_i. With this new proxy for r_i at hand, one can test the prediction of the model for the effect of changes in r_i on l_i:

$$\frac{d\ln l_i}{d\ln r_i^{-1}} = \kappa_i\,(\sigma-1) + \left[(\sigma-1)\ln r_i^{-1} - \frac{1}{1-\kappa_i}\right]\frac{d\kappa_i}{d\ln r_i^{-1}}. \tag{15}$$

This equation illustrates once more the tension between the productivity effect and the displacement effect. Compared to (8), equation (15) factors in the positive effect of a fall in r_i on capital, and hence captures a somewhat broader effect of automation.[8]

Alternatively, exogenous differences in the replaceability of tasks across firms, ρ_i, can also be used to identify variation in automation that is independent of demand shocks. The literature has shown how to build such proxies, which typically do not exhibit time variation. However, the model suggests that the decline in r_i should have a stronger effect on automation in firms with a higher ρ_i. Based on this insight, in the next sections, we build an instrument for robot adoption by combining information on which industries are more suitable for automation and firm-level measures of replaceability of employment.

What are the implications for other firm-level outcomes? Automation should clearly have a positive correlation with measures of productivity, although causality may run in both directions. It should also increase the demand for non-production workers. The relationship between automation and markups, instead, is possibly ambiguous. The benefit of automation is higher in more competitive markets. Hence, if markups are exogenous, they may exhibit a negative correlation with automation. However, the extension with endogenous market power has shown that, other things equal, automation may increase markups. Once again, the latter effect can be tested exploiting exogenous variation in automation.

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 the Appendix, 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.

[8] If the sign of (15) is negative, the sign of (8) must be negative *a fortiori*.

3 Data and Aggregate Facts

Our empirical analysis uses firm-level data for France over the 1994-2013 period and combines several firm-level data sets 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) in Euros and total employment.[9] We use this information to compute firm-level value added[10] 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 NACE level. We use the Wooldridge (2009) estimator for estimating the production-function coefficients[11]

The balance sheet data are complemented with information on the occupational structure 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. We use this information to construct a firm-level proxy for the extent to which employment is replaceable by robots, as explained below. For the descriptive analysis,

[9] 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.

[10] 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.

[11] 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 variable.

we use the full set of years (1994-2013), while for the Instrumental Variables (IV) regressions we focus on the 1996-2013 period and use 1994 as a pre-sample year. Finally, for each firm and year, we also have customs data on exports and imports from the French customs authority (DOUANE). In particular, 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 proxy for the use of robots at the firm level. The CN classification classifies trade of industrial robots into a specific product code, CN 84795000 (CN 84798950 before 1996). Accordingly, we identify firms that import robots in a given year as firms with positive imports for this product code in that year. In the empirical analysis, we use this information to build a proxy for the adoption of foreign robots by each firm. We also measure the stock of robot capital employed by a firm at a given point in time as the cumulative sum of robot imports by the firm up to that point. Using this information, we build a second proxy for automation measuring the intensity with which the firm uses robots.[12]

Figure 1 plots the time series of total robot imports into France obtained by summing robot imports across all firms in our sample (hollow circles). For comparison, the figure also plots the time series of total French robot imports obtained from the Comext database (full circles). Our 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 appear to be quite volatile, consistent with the lumpy nature of this investment.

Figure 2 shows the cumulative number of robot importers and the stock of imported robot capital over time. Both variables have rapidly trended upward. In particular, the number of firms with at least one year of robot imports rose from 121 in 1994 to more than 1700 in 2013. Similarly, the stock of imported robot capital increased from 63 million Euros in 1994 to around 1,8 billion Euros in 2013. Overall, these numbers suggest that automation has become increasingly widespread across French firms over the sample period.

Finally, Figure 3 reports the cumulative number of robot importers by two-digit sector. The figure shows that, while robot importers are observed in many different sectors, they are particularly frequent in manufacturing, especially in industries such as production of

[12] All our findings are qualitatively similar when computing the stock of robot capital using a perpetual inventory method with an annual depreciation rate of 15%, which falls within the range of depreciation rates normally assumed for robots in manufacturing (see, e.g., Graetz and Michaels, 2018).

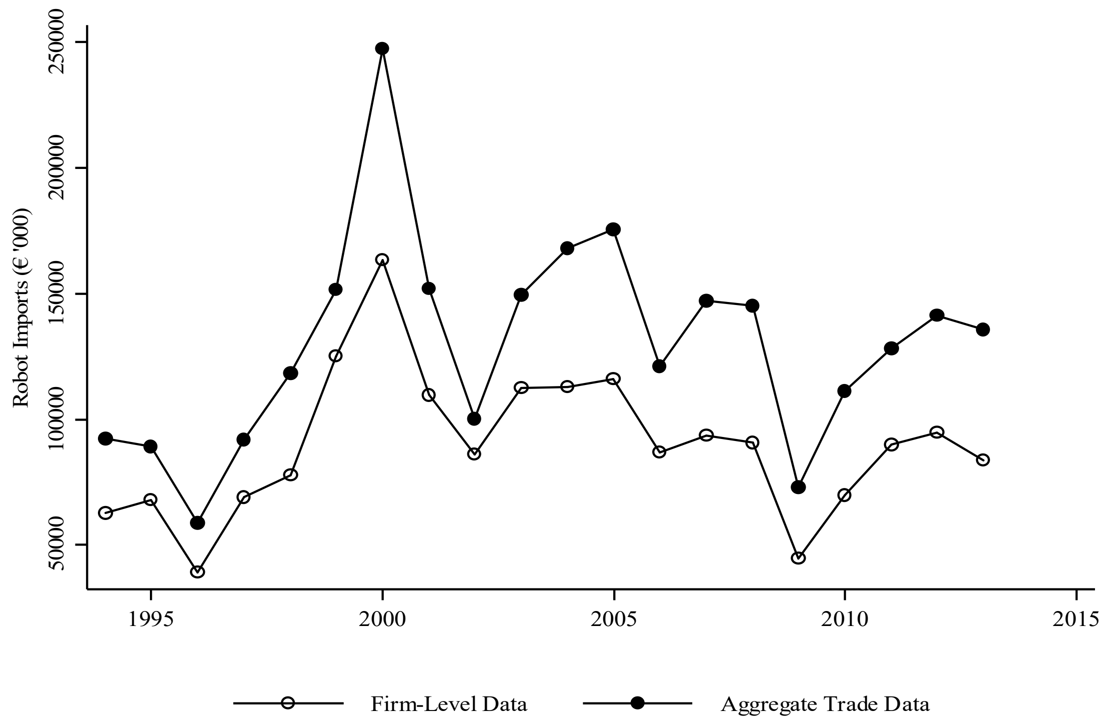

Figure 1: Robot Import, 1994-2013

motor vehicles, machinery, and electrical equipment. Consistent with this evidence, in our econometric analysis we use a baseline sample consisting of manufacturing firms only. Overall, there are more than 800 different manufacturing firms importing robots at least once over the period of analysis. 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 estimation sample.

4 Preliminary Evidence

We start by providing descriptive evidence on how firms that adopt robots compare to firms that do not in terms of various characteristics. Table **??** reports summary statistics on a number of firm-level variables, separately for firms that have imported robots at least once over 1994-2013 ("robot adopters") and for firms that have never imported robots over this period ("non robot adopters"). Our sample consists of 64,760 manufacturing firms. Of these, 746 are robot adopters, corresponding to 1.15 percent of all firms and 0.96 percent of all firm-

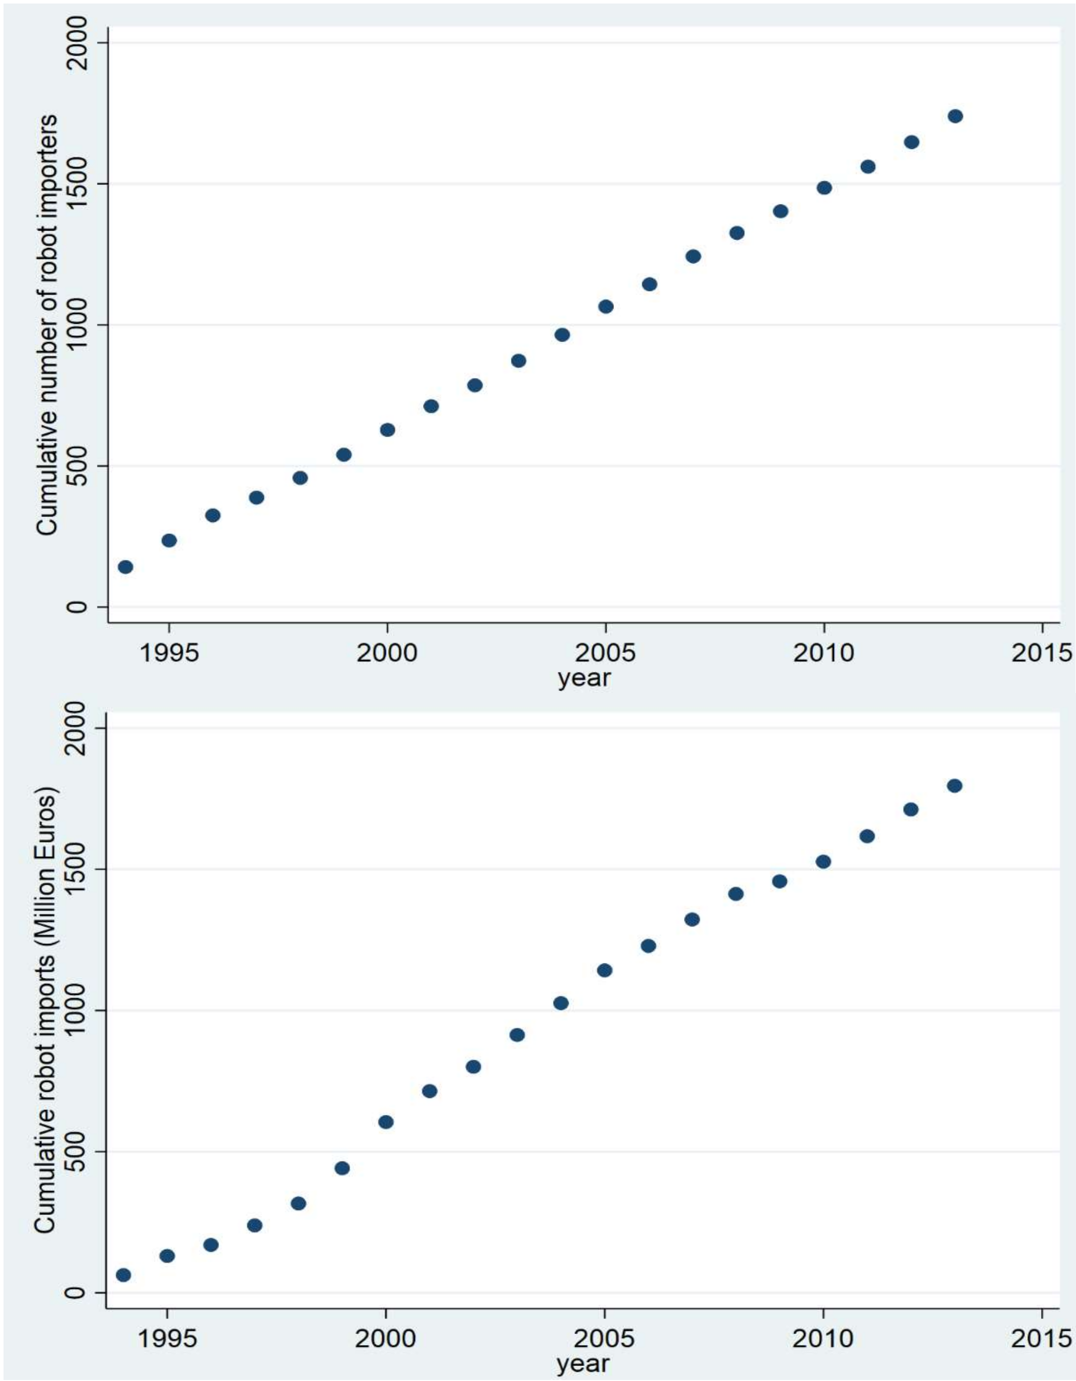

Figure 2: Cumulated Robot Importers and Import

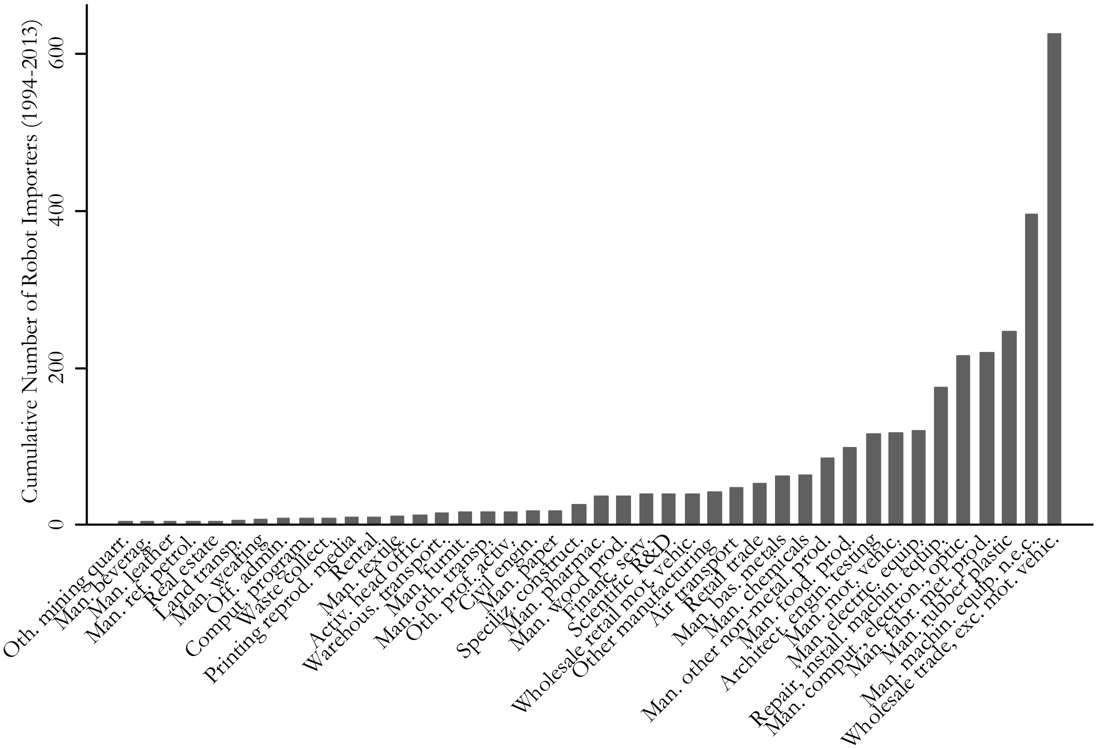

Figure 3: Cumulative number of robot importers by two-digit industry

year observations in our data set.[13] Robot intensity, defined as the ratio between the stock of robot capital and the total physical capital stock of the firm, equals 11 percent on average for robot adopters. The average robot adopter is around 11 times larger than 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 Total Factor Productivity, on average. The skill composition of employment also differs across robot adopters and non robot adopters, the share of employment in high-skill professions being twice as high on average in the former group of firms than in the latter. Robot adopters are also more likely to import and export goods other than robots, and are characterized by a larger share of employment performing tasks that can be replaced

[13] The relatively low number of firms adopting industrial robots is consistent with other existing studies. For instance, Acemoglu, Lelarge and Restrepo (2020), who collected information on robot adoption in France from multiple sources, find that only 1% of the firms in their sample purchased robots over the 2010-2015 period. While robot adopters are a minority, they nevertheless account for a large fraction of employment and sales.

by robots.[14]

Table 1 also reports the average annualized change in each variable over 1994-2013, separately for the two sets of firms. Robot adopters have increased robot intensity at an average rate of 0.2 log points per year. While employment has decreased in both groups of firms, robot adopters have shed workers at a slower rate than non robot adopters (0.02 vs. 0.03 log points per year, respectively).[15] Robot adopters have also experienced a relatively slower reduction in sales, sales per worker and TFP, and a relatively faster increase in the employment share of high-skill professions.

To further shed light on the differences between the two groups of firms, we now estimate conditional correlations between robot adoption and firm-level characteristics, by running OLS regressions of the following form:

$$Y_{ijt} = \alpha_i + \alpha_{jt} + \beta \cdot Rob\_Adoption_{ijt} + \mathbf{X}'_{ijt} \cdot \boldsymbol{\gamma} + \varepsilon_{ijt}, \tag{16}$$

where i denotes a firm, j indicates the 5-digit NACE industry in which the firm operates, and t stands for time. Y_{ijt} is an outcome and $Rob\_Adoption_{ijt}$ 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. (16). In the first version, we control for firm fixed effects, α_i, and for 5-digit industry$\times$year fixed effects, α_{jt}. The "robot adopter premia", β, are then identified by comparing outcomes, in deviations from within-firm means, across firms belonging to the same 5-digit industry and year. This approach ensures that the coefficients β are not contaminated either by time-invariant firm characteristics that could be correlated with adoption and outcomes or by differences in the distribution of adopters and non adopters across industries. In the second version of eq. (16), we add controls for observable firm characteristics (log sales and dummies for exporter and importer status). We measure each characteristic at baseline, that is, in the first year in which the firm is observed in the sample, and interact its first-year value with a full set of year dummies. The resulting interactions, contained in the vector $\mathbf{X}_{ijt}$, 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.

The results are reported in Table 2. Odd-numbered columns show the estimates of the

[14] The replaceability of tasks by robots is constructed following Graetz and Michaels (2018) and is explained in details below. This variable refers to the year 1994 and is computed for the 36,972 firms used in the IV regressions.

[15] 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					
Robot adopter	6,003	746	1	1	1	0
Robot intensity	6,003	746	0.108	0.005	0.635	0.190
No. of employees	6,003	746	838	184	3,107	-0.017
Empl. sh. high skill	6,003	746	0.157	0.111	0.142	0.006
Sales (€'000)	6,003	746	758,388	42,911	6,965,072	-0.073
Sales per worker (€'000)	6,003	746	2,002	221	108,120	-0.058
VA per worker (€'000)	5,855	742	183	164	2,802	-0.069
TFP	5,848	741	422	164	2,702	-0.066
Importer	6,003	746	0.973	1	0.163	0.001
Exporter	6,003	746	0.950	1	0.218	0.002
Replaceability	513	513	0.372	0.403	0.185	-
	Non Robot Adopters					
Robot adopter	616,798	64,014	0	0	0	0
Robot intensity	604,409	64,014	0	0	0	0
No. of employees	616,798	64,014	77	27	309.54	-0.029
Empl. sh. high skill	616,798	64,014	0.082	0.056	0.107	0.003
Sales (€'000)	616,794	64,014	53,465	7,385	673,610	-0.091
Sales per worker (€'000)	616,794	64,014	653	223	11,554	-0.063
VA per worker (€'000)	604,960	63,307	187	69	1,945	-0.066
TFP	593,795	62,571	287	128	1,343	-0.071
Importer	616,798	64,014	0.560	1	0.4963	0.001
Exporter	616,798	64,014	0.554	1	0.4971	0.004
Replaceability	36,459	36,459	0.356	0.358	0.190	-

The whole sample consists of all manufacturing firms with more than 10 employees (64,760 firms). Robot adopter 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 cumulative sum of robot imports. Importer and Exporter are dummies taking on value 1 if the firm imports (resp. exports) in a given year and 0 otherwise. Replaceability is the share of firm employment in occupations that can be replaced by robots. All statistics are computed on firm-level observations for the 1994-2013 period, except for Replaceability, which is observed in 1994 and is computed for 36,972 firms used in the instrumental variables regressions. Changes are computed as annualized log differences, except for Employment sh. high skill, Exporter and Importer, for which we report annualized changes in levels.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	ln Sales		ln Employment		ln Sales per Worker	
Rob_Adoption	0.130***	0.198***	0.093***	0.114***	0.039***	0.087***
	[6.113]	[9.546]	[4.622]	[5.664]	[2.614]	[5.797]
Obs.	615,785	614,427	617,229	615,595	615,785	614,427
R2	0.949	0.95	0.878	0.878	0.89	0.891
	ln VA per Worker		ln TFP		Empl. Sh. High Skill	
Rob_Adoption	0.011	0.051***	0.030**	0.067***	0.011***	0.003
	[0.707]	[3.155]	[2.042]	[4.492]	[4.195]	[0.973]
Obs.	605,217	603,926	593,996	592,746	617,229	615,595
R2	0.815	0.815	0.857	0.858	0.677	0.679
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry ×year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes

The dependent variables are annual observations of the firm-level outcomes indicated in columns' headings. Rob_Adoption is a dummy equal to 1 for all years since the firm starts importing robots, and equal to 0 otherwise. Industry refers to 5-digit industries. The control variables included in columns (2) and (4) 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 clustered at the firm level, t-statistics are reported in square brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

specification including only firm and industry×year fixed effects. Even-numbered columns report the results obtained by adding interactions between year dummies and the initial-period value of three firm characteristics: log sales and 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 industry of operation.

The differences between robot adopters and non robot adopters just documented may have two interpretations: either robot adopters differ from other firms before adopting robots, or they start diverging afterward. To shed light on this question, we now 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. (16) by adding the first five lags and leads of $Rob\_Adoption_{ijt}$:

$$Y_{ijt} = \alpha_i + \alpha_{jt} + \sum_{s=-5}^{5} \beta_s \cdot Rob\_Adoption_{ijt-s} + \varepsilon_{ijt}. \tag{17}$$

The coefficients β_s estimated from eq. (17) 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. As before, we correct the standard errors for clustering at the firm level.

The results are reported in Figure 4, where each graph refers to a different outcome. Note 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 trends in employment is reversed: while robot adopters still grow faster than non robot adopters, the differential gradually vanishes. Robot adopters also experience a more marked shift in the skill composition of the workforce toward 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, which suggests that the efficiency gains from robot adoption do not translate into lower prices. 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, labor demand shifts toward high-skill workers and, possibly, increases in firm markups.

5 Identifying the Effects of Robots on Firm-Level Outcomes

Both our theoretical model and the preliminary evidence suggest that the correlations between robot adoption and other firm-level characteristics may be confounded by demand shocks, which are likely to influence both the outcomes of a firm and its choice to automate. In this section, we use two complementary strategies to purge away demand shocks and identify the effects of robots. First, we exploit yearly within-firm variation and regress outcomes on robot intensity, a variable that should not be influenced by demand shocks according to our model (Section 5.1). Second, we focus on long-run changes in outcomes within firms, and exploit variation in the decision to adopt robots as driven by pre-existing differences in technological characteristics determining the predisposition to automate (Section 5.2).

5.1 Robot Intensity

In our first approach, we re-estimate eq. (16) replacing the dummy $Rob\_Adoption_{ijt}$ with our proxy for the intensity with which a firm uses robots, $\ln Rob\_Intensity_{ijt}$, defined as the log ratio between the stock of robot capital and the total capital stock of the firm. This variable is the empirical counterpart of the theoretical measure introduced in eq. (14). By scaling robot capital with the total capital stock of the firm, $\ln Rob\_Intensity_{ijt}$ neutralizes demand shocks, as the latter should affect both the numerator and the denominator of the ratio. The log transformation implies that $\ln Rob\_Intensity_{ijt}$ is defined only for firms that import robots. Because the specification controls for firm and 5-digit industry$\times$year fixed effects, the coefficients β are identified from changes in robot intensity over time within robot adopters, controlling for common shocks hitting all firms in a narrow sector.

The results are reported in Table 3. Odd-numbered columns refer to the specification that only controls for firm and industry-year fixed effects, while even-numbered columns refer to the specification that also includes the interactions of the year dummies with initial firm size and indicators for the initial import and export status of the firm. Compared to the results reported in Table 2, the estimate of β switches from positive to negative in the regressions for sales and employment, and is highly statistically significant. This pattern suggests that demand shocks lead firms to both expand and automate, resulting in a spurious positive

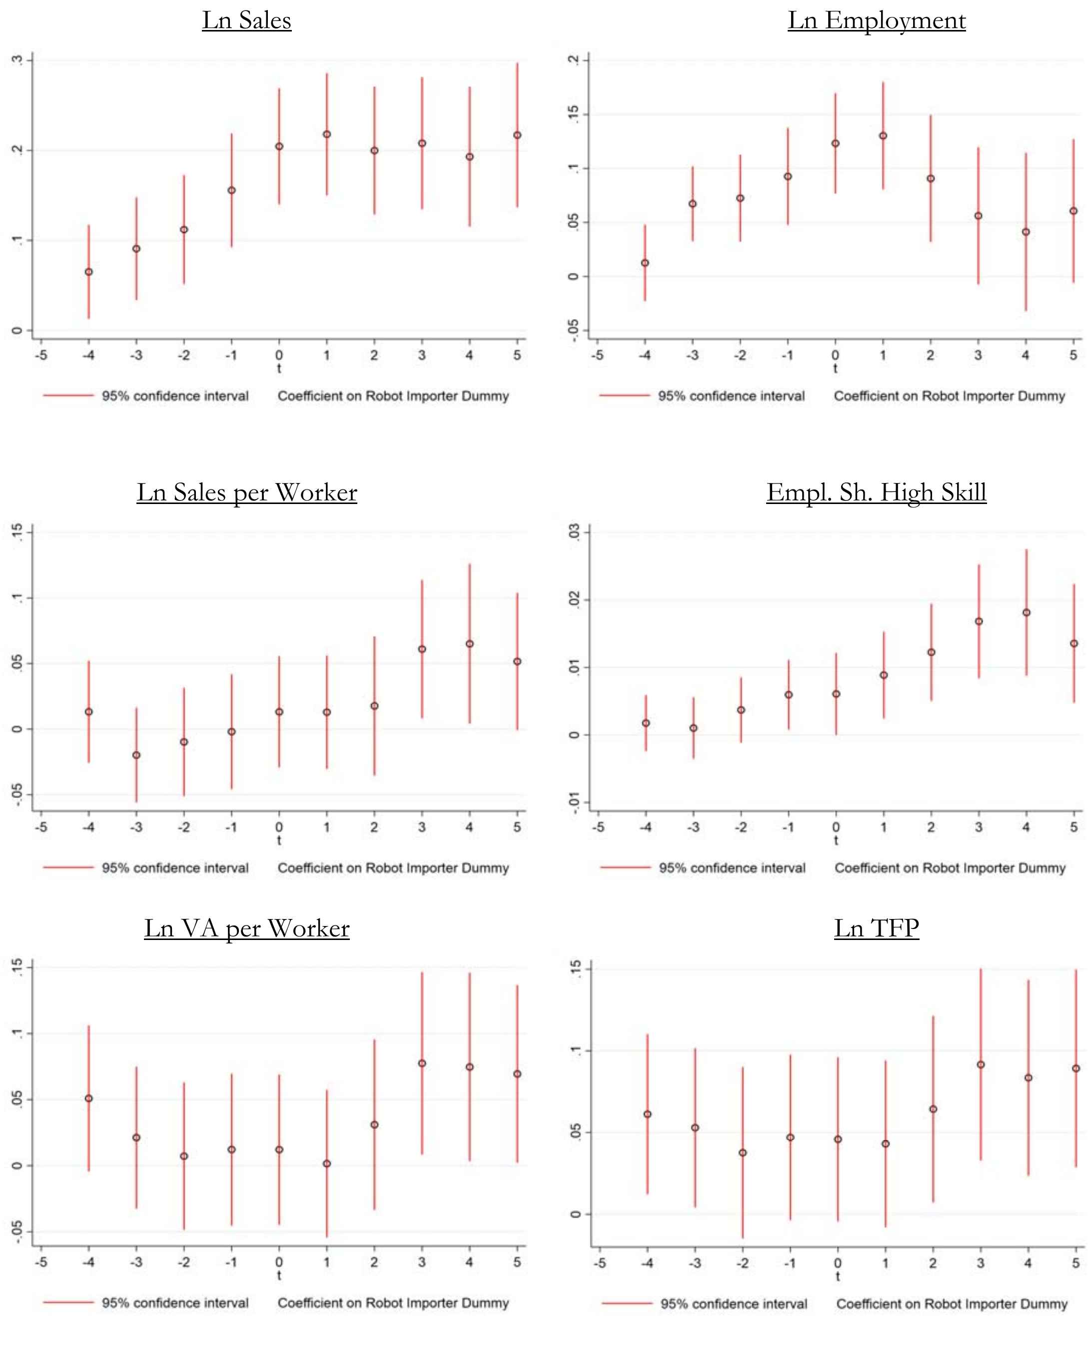

Figure 4: Evolution of Outcomes across Robot Importers and Non-Importers

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

	(1)	(2)	(3)	(4)	(5)	(6)
	ln Sales		ln Employment		ln Sales per Worker	
Ln Rob_Intensity	-0.141***	-0.138***	-0.191***	-0.186***	0.033*	0.029
	[-4.396]	[-4.253]	[-5.882]	[-5.661]	[1.861]	[1.589]
Obs.	5,998	5,948	6,003	5,953	5,998	5,948
R2	0.982	0.982	0.955	0.956	0.885	0.886
	ln VA per Worker		Ln TFP		Empl. Sh. High Skill	
Ln Rob_Intensity	0.052***	0.056***	0.026*	0.032**	0.018***	0.018***
	[3.074]	[3.265]	[1.668]	[2.119]	[2.936]	[2.708]
Obs.	5,823	5,773	5,817	5,767	6,003	5,953
R2	0.795	0.798	0.883	0.885	0.876	0.877
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry ×year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes

The dependent variables are annual observations of the firm-level outcomes indicated in columns' headings. Ln Rob_Intensity is the log ratio between the cumulative stock of robot capital and the total capital stock of the firm. Industry refers to 5-digit industries. The control variables included in columns (2) and (4) 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 clustered at the firm level, t-statistics are reported in squared brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

correlation between robot adoption and firm size. That the relation between robot adoption and employment turns negative once demand shocks are neutralized is consistent with the idea that automation leads to job displacement. In terms of magnitude, multiplying the coefficient on ln $Rob\_Intensity_{ijt}$ in the employment regression by the average annual change in robot intensity reported in Table 1 (0.2 log points) implies that the observed increase in robot intensity explains an average fall in employment equal to 3.5 percent per year among robot adopters. Regarding the other outcomes, Table 3 continues to show positive estimates of β across the board. While some of these coefficients are only marginally significant, the qualitative pattern of results suggests that robots tend to improve firm efficiency and to shift labor demand in favor of high-skill workers.

5.2 Instrumental Variables

Our second approach to identify the effects of automation on firm-level outcomes consists of using IV to isolate the variation in robot adoption that is not contaminated by demand shocks. To operationalize this approach, we estimate long-difference specifications of the following form:

$$\Delta Y_{ij} = \alpha_s + \beta \cdot \Delta Rob\_Adoption_{ij} + \mathbf{X}'_{ij} \cdot \boldsymbol{\gamma} + \Delta \varepsilon_{ij}, \tag{18}$$

where i indexes firms and j denotes 5-digit industries; ΔY_{ij} 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; $\Delta Rob\_Adoption_{ij}$ takes on value 1 if firm i has adopted robots over the sample period, and is equal to 0 both for non-adopters and for firms that were already using robots initially; $\mathbf{X}_{ij}$ are start-of-period values of control variables (described below); and α_s are fixed effects for 3-digit sectors.[16] By eliminating year-on-year variation, the use of long differences implies that the coefficient β is identified from cross-sectional differences in the growth of outcomes between robot adopters and other firms. The sector fixed effects, α_s, absorb differential trends in adoption and outcomes across sectors, while the covariates, $\mathbf{X}_{ij}$, remove heterogeneous trends across firms characterized by different initial conditions within the same sector.

Demand shocks could bias the OLS estimate of β from eq. (18) if they both induced firms to adopt robots and influenced the outcomes. Hence, we instrument $\Delta Rob\_Adoption_{ij}$ using a variable that is meant to eliminate the effect of demand shocks by isolating the variation in adoption occurring for technological reasons. Because most of the variation in robot adoption is across firms, finding a strong instrument for $\Delta Rob\_Adoption_{ij}$ in the context of eq. (18) is easier than explaining the exogenous within-firm variation in $Rob\_Adoption_{ijt}$ in the context of eq. (16).

To construct the instrument, we follow the insights of our theoretical model. The latter shows that a reduction in the cost of machines should affect robot adoption relatively more in firms that are more prone to automate. To capture this idea, we exploit the fact that the different nature of the production process across industries makes production easier to automatize in some industries than in others, implying that the cost of automation should fall relatively more over time in the former industries. At the same time, within a given industry, some firms are more prone than others to automatize production, because they perform activities that are relatively easier to assign to robots. Accordingly, our instrument, labeled $Rob\_Exposure_{ij}$, is obtained by interacting a proxy for how suitable production is for automation in a given industry, $Rob\_Suitability_{ij}$, with a proxy for the ease with which robots can replace worker activities within a given firm, $Replaceability_{ij}$.

For each firm i, $Rob\_Suitability_{ij}$ is defined as the average robot intensity of all firms

[16] Given the significant reduction in degrees of freedom due to the use of long differences, in this specification we define the industry fixed effects at the slightly more aggregated 3-digit level.

$i' \neq i$ in the same 5-digit industry j in the initial year, and is constructed as follows:

$$Rob\_Suitability_{ij} = \ln \frac{1 + \sum_{i' \neq i \in j} Rob\_Stock_{i'j}}{\sum_{i' \neq i \in j} Cap\_Stock_{i'j}},$$

where $Rob\_Stock_{i'j}$ and $Cap\_Stock_{i'j}$ denote, respectively, the initial stock of robots and the initial total capital stock of firm $i' \neq i \in j$. Industries in which this ratio is higher should be relatively more suitable for automation and should thus experience a relatively larger fall in the cost of robots in subsequent years.

As for $Replaceability_{ij}$, we follow Graetz and Michaels (2018) and exploit differences across firms 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 it, 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 International Federation of Robotics, such as welding, painting and assembling.[17] Then, we manually map each US Census occupation into the 29 French occupations for which we have employment data in 1994. With this information in hand, we construct the firm-level replaceability measure as follows:

$$Replaceability_{ij} = \sum_{o=1}^{29} \omega_{oij} \times Replaceability_o,$$

where $Replaceability_o$ is the replaceability of French occupation o and ω_{oij} is the share of occupation o in firm i's employment in 1994.

The instrument, $Rob\_Exposure_{ij}$, is finally obtained as

$$Rob\_Exposure_{ij} = Rob\_Suitability_{ij} \times Replaceability_{ij}.$$

Accordingly, our instrument 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.

To be a valid instrument for $\Delta Rob\_Adoption_{ij}$, $Rob\_Exposure_{ij}$ must be uncorrelated

[17]Previous studies have investigated the effect of new technologies on occupations involving routine tasks (e.g., Autor, Levy, and Murnane, 2003). However, Cheng et al. (2019) find that robots are more prevalent at firms where employees are commonly doing manual tasks, but not those that require routine tasks.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ ln Sales		Δ ln Employment		Δ ln Sales per Worker	
Δ Rob_Adoption	0.025***	0.023***	0.021***	0.044***	0.043***	0.021***
	[6.450]	[7.145]	[6.490]	[10.565]	[10.570]	[4.572]
Obs.	36,666	36,950	36,972	36,666	36,666	36,666
R2	0.057	0.032	0.030	0.075	0.075	0.055
	Δ ln VA per Worker		Δ ln TFP		Δ Empl. Sh. High Skill	
Δ Rob_Adoption	0.001	0.017***	0.005	0.019***	0.001***	0.000
	[0.190]	[4.091]	[1.350]	[5.277]	[3.250]	[0.303]
Obs.	35,534	35,534	33,964	33,964	36,972	36,950
R2	0.029	0.043	0.036	0.050	0.020	0.032
Controls	Industry FE	All Controls	Industry FE	All Controls	Industry FE	All Controls

The dependent variables are the annualized changes in the firm-level outcomes indicated in columns' headings. Δ Rob_Adoption is a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non-importers. Industry fixed effects are dummies for 3-digit industries. The control variables included in columns (2) and (4) are the employment share of occupations that can be replaced by robots in 1994 (Replaceability), the log ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry in 1994 (Rob_Suitability), and the initial values of log sales and of dummies for importing and exporting firms. Standard errors clustered at the 5-digit industry level, t-statistics are reported in squared brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

with $\Delta\varepsilon_{ij}$ in eq. (18) conditional on the covariates. In this respect, the sector fixed effects α_s remove sector characteristics that could be correlated with $Rob\_Exposure_{ij}$ and affect the evolution of outcomes across all firms in a sector. Moreover, we control for the start-of-period values of log firm sales and of indicators for exporting and importing firms, included in the vector $\mathbf{X}_{ij}$ along with $Rob\_Suitability_{ij}$ and $Replaceability_{ij}$. These predetermined firm-level controls account for the fact that, within a given sector, larger and more trade-oriented firms may be more prone to adopt robots and may systematically follow different paths in terms of key outcomes. Then, our identification strategy exploits differential exposure to robot adoption across firms and 5-digit industries within 3-digit sectors: firms that are most exposed to robots in a sector are those with high levels of $Replaceability_{ij}$ operating in industries with high levels of $Rob\_Suitability_{ij}$.

Table 4 reports the OLS estimates of eq. (18), together with standard errors corrected for clustering within 5-digit industries. As indicated in the columns' headings, the dependent variables are the 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.[18] For each

[18] We winsorize the change in each outcome at the top and bottom 5 percent of the distribution to prevent

Table 5: Firm-Level Outcomes and Robot Adoption, Long Differences (IV)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Δ Rob_Adoption	Δ ln Sales	Δ ln Employment	Δ ln Sales per Worker	Δ ln VA per Worker	Δ ln TFP	Δ Empl. Sh. High Skill
Δ Rob_Adoption		0.192	-0.557**	1.019**	1.188*	0.816	0.047**
		[0.422]	[-2.006]	[2.100]	[1.864]	[1.521]	[2.024]
Rob_Exposure	0.002***						
	[2.898]						
Replaceability	0.033***	-0.021***	-0.021***	0.002	-0.005	-0.013**	-0.003***
	[2.657]	[-3.423]	[-6.324]	[0.291]	[-0.666]	[-2.102]	[-7.431]
Rob_Suitability	0.344**	0.196	-0.085	0.289	0.244	0.113	0.085**
	[2.334]	[0.504]	[-0.308]	[0.743]	[0.516]	[0.253]	[2.407]
ln Initial Sales	0.013***	-0.017***	0.007**	-0.027***	-0.029***	-0.023***	0.000
	[7.308]	[-2.665]	[2.111]	[-4.458]	[-3.764]	[-3.260]	[0.272]
Dummy Initial Importer	0.000	0.015***	0.001	0.014***	0.015***	0.014***	0.001**
	[0.127]	[6.561]	[0.787]	[5.704]	[5.365]	[6.172]	[2.380]
Dummy Initial Exporter	0.001	0.006***	-0.004***	0.011***	0.009***	0.008***	0.001***
	[0.724]	[3.124]	[-2.762]	[4.557]	[3.605]	[4.001]	[3.610]
Obs.	36,950	36,666	36,950	36,666	35,534	33,964	36,950
KP F-Statistic		8.745	8.399	8.745	7.370	7.081	8.399

The dependent variables are indicated in columns' headings and are: Δ Rob_Adoption, a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non-importers (column 1); the annualized changes in log sales (column 2), log employment (column 3), log sales per worker (column 4), log value added per worker (column 5), log TFP (column 6) and the employment share of high-skill professions (column 7). Rob_Exposure is the product between the firm-level employment share of occupations that can be replaced by robots in 1994 (Replaceability) and the log ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry in 1994 (Rob_Suitability). All regressions also include 3-digit industry fixed effects. Standard errors are clustered at 5-digit industry level, t-statistics are reported in squared brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

outcome, the table presents results from a specification including only the industry fixed effects (odd-numbered columns) and from the complete specification including also the control variables (even-numbered columns). Consistent with the findings presented in the previous section, Table 4 shows that firms that adopt robots over the sample period experience a relatively larger increase in size, a relatively stronger improvement in efficiency and a relatively faster shift in labor demand toward high-skill workers.

The IV estimates of eq. (18) are reported in Table 5. Column (1) shows the first-stage results. The coefficient on the instrument $Rob\_Exposure_{ij}$ is positive and statistically significant: firms that are more exposed to robots due to pre-existing technological characteristics do indeed show a greater tendency to adopt robots in subsequent years. As for the other regressors, the positive and precisely estimated coefficients on $Replaceability_{ij}$ and $Rob\_Suitability_{ij}$ imply that robot adoption is relatively higher in firms that perform more automatable tasks in the pre-sample period and in industries in which production is more suitable for automation. $\Delta Rob\_Adoption_{ij}$ is also positively correlated with initial firm sales.

results from being driven by extreme observations.

As predicted by the model, this result implies that initially larger firms tend to adopt more robots in the future, and highlights the importance of controlling for size in eq. (18) to eliminate a potential source of correlation between $Rob\_Exposure_{ij}$ and $\Delta\varepsilon_{ij}$.

The second-stage estimates of β are reported in columns (2)-(5) of Table 5. Each column refers to a different outcome, as indicated in the column's heading. The corresponding reduced-form coefficients, obtained by regressing each outcome on $Rob\_Exposure_{ij}$ and the full set of fixed effects and controls, are reported in panel a) of Table 6. In the regression for the log change in employment, the reduced-form coefficient on $Rob\_Exposure_{ij}$ is equal to -0.001 and is very precisely estimated. This implies that firms that are more exposed to robots, owing to the interplay between their industry's initial suitability for automation and their pre-sample specialization in automatable tasks, experience a relatively larger reduction in employment over the sample period. Taking the ratio between the reduced-form and the first-stage coefficient on $Rob\_Exposure_{ij}$ yields the second-stage coefficient on $\Delta Rob\_Adoption_{ij}$. The latter is equal to -0.557 and is precisely estimated, implying that exogenous robot adoption leads firms to shed workers.

By comparing the OLS and second-stage coefficients on $\Delta Rob\_Adoption_{ij}$, we can have a sense of 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 $\Delta Rob\_Adoption_{ij}$, 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 $\Delta Rob\_Adoption_{ij}$, (σ_{IV}^2/σ^2) is the fraction of the overall variance of $\Delta Rob\_Adoption_{ij}$ 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 for most firms (95.7 percent) robot adoption results from demand shocks. Using these numbers along with the estimates of β_{OLS} and β_{IV} reported in Tables 4 and 5, respectively, yields $\beta_{RES} = 0.049$. Accordingly, exogenous adoption explains an average annual fall in employment equal to 2.4 percentage points 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 4.7 percentage points in the former group of firms relative to the latter (i.e., $\beta_{RES} \times (\sigma_{RES}^2/\sigma^2)$).

Turning to the other outcomes, Table 5 exhibits positive and statistically significant coefficients on $\Delta Rob\_Adoption_{ij}$ in the regressions for log sales per worker and log value added per

Table 6: Firm-Level Outcomes and Robot Adoption, Long Differences (IV and RF)

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ ln Sales	Δ ln Employment	Δ ln Sales per Worker	Δ ln VA per Worker	Δ ln TFP	Δ Empl. Sh. High Skill
	a) Reduced Form (RF)					
Rob_Exposure	0.0003	-0.001***	0.001**	0.002**	0.001*	0.0001**
	[0.411]	[-2.733]	[2.561]	[2.268]	[1.687]	[2.256]
Obs.	36,666	36,950	36,950	35,534	33,964	36,950
R2	0.074	0.032	0.033	0.043	0.050	0.033
	b) Additional Interactions of Replaceability, IV					
Δ Rob_Adoption	0.380	-0.824**	1.528**	1.719**	1.250*	0.080**
	[0.641]	[-2.293]	[2.289]	[2.221]	[1.958]	[2.264]
Obs.	36,903	36,903	36,903	36,903	36,903	36,903
KP F-Statistic	7,333	7,333	7,333	7,333	7,333	7,333
	c) Additional Interactions of Replaceability, RF					
Rob_Exposure	0.001	-0.001***	0.002**	0.002**	0.002**	0.0001***
	[-0.619]	[-3.226]	[1.788]	[2.696]	[2.153]	[3.320]
Obs.	36,903	36,903	36,903	36,903	36,903	36,903
R2	0.076	0.033	0.076	0.038	0.050	0.034

The dependent variables are the annualized changes in the firm-level outcomes indicated in columns' headings. Δ Rob_Adoption is a dummy equal to 1 for firms that start importing robots over the sample period and equal to 0 for non-importers. Rob_Exposure is the product between the firm-level employment share of occupations that can be replaced by robots in 1994 (Replaceability) and the log ratio between the overall stock of robots and the total capital stock of all other firms in each 5-digit industry in 1994 (Rob_Suitability). Coefficients in panels a) and c) are estimated with OLS. Coefficients in panel b) are estimated with IV as in Table 5. Panels b) and c) also control for the initial values of: sectoral exports; sectoral imports; sectoral export unit value; and sectoral import unit value. These controls enter both lineraly and interacted with Replaceability. All specifications include 3-digit industry fixed effects. Standard errors are clustered by industry at 5-digit level, t-statistics are reported in squared brackets. ***, **, *: denote significance at the 1, 5 and 10% level, respectively.

worker, and a positive and marginally insignificant coefficient in the regression for log TFP. The table also shows a precisely estimated and positive coefficient on $\Delta Rob\_Adoption_{ij}$ in the regression for the employment share of high-skill professions. Hence, the IV results confirm our previous evidence, according to which robot adoption induces firms to raise efficiency and shifts labor demand in favor of high-skill workers. The effect of robot adoption on total sales, while positive, is not statistically significant, suggesting again that the productivity gains from automation may not always translate into lower prices.

A possible concern with our identification strategy is that the suitability of an industry for automation, $Rob\_Suitability_{ij}$, could be correlated with other industry-level factors that influence outcomes differentially across firms with heterogeneous levels of $Replaceability_{ij}$. To raise confidence in our IV results, we therefore augment the specification by adding other industry-level characteristics, both linearly and interacted with $Replaceability_{ij}$. In particular, we consider: (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; and (iii) the average unit value of exports, to account for cross-industry differences in product characteristics such as quality. Similar to $Rob\_Suitability_{ij}$, we construct each of these variables in the initial year by aggregating across firms other than i. The results are reported in panels b) and c) of Table 6, which contain second-stage and reduced-form coefficients, respectively. Reassuringly, the estimated coefficients are similar to our baseline estimates, suggesting that our evidence is unlikely to be confounded by other industry-level characteristics that could interact with replaceability.

6 Conclusions

In this paper, we have documented how the adoption of industrial robots affects a series 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 saves on production workers, but they also requires 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. We also allow for the possibility that automation, by fostering the technological advantage of top firms, increases market power.

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 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.[19]

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 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.[20]

Second, investigating more the dynamic effects of automation seems equally important. For instance, while we have found evidence consistent with the hypothesis that automation may lead to higher markups, the effect on market power might be transitory. For instance, potential competitors may learn from robot adopters, thereby eroding the technological gap, or it could simply be that firms adjust prices slowly to changes in productivity. Third, studying more the labor-market adjustments to automation seems crucial 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.

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

[20] See Acemoglu, Leclerc and Restrepo (2020), and Koch, Manuylov and Smolka (2019), for some evidence on this reallocation.

References

[1] 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.

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

[3] 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.

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

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

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

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

[8] 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.

[9] 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

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

[11] 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.*

[12] 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.

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

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

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

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

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

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

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

[20] 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.

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

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

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

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

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

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

[27] 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.

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

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

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

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

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

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

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

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

[36] Martinez, Joseba (2019). "Automation, Growth and Factor Shares" working paper.

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

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

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

Appendix A Choice of Automation: Comparative Statics

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

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

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 $(\sigma-1)\ln\left(\frac{w}{r_i}\right) < \delta - 1$, 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:

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

Differentiating the first-order condition (10), 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 (11) 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_i) and (ρ_i/h) as follows:

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

Appendix B Automation and the Labor Share

We now study the effect of automation on the labor share. Recall that automation affects both the demand for production and non-production workers. The labor share, denoted by λ_i, is then:

$$\lambda_i \equiv \frac{wl_i + hf(\kappa_i, \rho_i)}{p_i y_i} = \left(1 - \frac{1}{\sigma}\right)(1 - \kappa_i) + \frac{h\kappa_i^\delta}{\rho_i \delta p_i y_i}.$$

After using equations (5) and (14) we obtain:

$$\lambda_i = \left(1 - \frac{1}{\sigma}\right)\left[1 + \kappa_i\left(\frac{1}{\delta}\ln\left(\frac{w}{r_i}\right) - 1\right)\right]. \tag{19}$$

This equation shows that the labor share falls with automation when $\ln(w/r_i) < \delta$.

Appendix 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

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

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

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

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 $\frac{w}{r_i}$ 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 $\frac{w}{r}$:

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

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

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

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

which imply that the likelihood of further automation increases more with (w/r_i) 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_i) increases more the likelihood of further automation for firms with higher replaceability ρ_i, since these face a lower cost.