

**The Direct and Indirect Effects of Automation on Employment:  
A Survey of the Recent Literature**

Philippe Aghion (Collège de France, London School of Economics, INSEAD)

Céline Antonin (Sciences Po OFCE, Collège de France)

Simon Bunel (Banque de France, Paris School of Economics)

Xavier Jaravel (London School of Economics)

**Summary**

In this article, we survey the recent literature and discuss two contrasting views on the direct and indirect impacts of automation on employment. A first view predicts that firms that automation reduce employment (a negative “direct effect”), even if this may ultimately result in new job creations taking advantage of the lower equilibrium wage induced by job destructions (a positive “indirect effect”). A second approach emphasizes the market size and business stealing effects of automation. Automating firms become more productive, which enables them to lower their quality-adjusted prices, and therefore to increase the demand for their products. The resulting increase in scale translates into higher employment by automating firms (a positive “direct effect”), potentially at the expense of their competitors (a negative “indirect effect” through business stealing). Drawing from our empirical work on French firm-level data and a growing literature covering multiple countries, we provide empirical support for this second view: automation has a positive direct effect on employment at the firm level. We discuss the implications of these results for the taxation of automation technologies such as robots.

JEL codes J24, O3, O4

Keywords: artificial intelligence, growth, policy, automation, robots, employment.

## 1. Introduction

Should we fear or wish automation? On the one hand, we may fear it because it replaces human workers by machines to perform their tasks, thereby increasing unemployment and reducing the labor share. On the other hand, we may welcome it as spurs growth and prosperity, as illustrated by the big technological revolutions – steam engine in the early 1800s, electricity in the 1920s – none of which generated the mass unemployment anticipated by some.

The fear that machines will destroy human jobs began long ago. Already in 1589, when William Lee invented a machine to knit stockings, the working class was so fearful of the consequences that he was rejected everywhere and even threatened. Then came the first industrial revolution, the “steam engine revolution”, and with it the so-called Luddite movement. Despite a 1769 law protecting machines from being destroyed, destruction intensified as the weaving loom became widespread, culminating with the Luddite rebellion in 1811–1812.

The second industrial revolution, the “electricity revolution”, occurred first in the US in the 1920s. Starting in the 1930s, economists began to express concern about the unemployment that this revolution would generate. In 1930, Keynes wrote, “We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment.”<sup>1</sup> Once again, the prediction of a large-scale increase in unemployment did not materialize.

More recently, the information technologies (IT) and artificial intelligence (AI) revolutions, have raised similar fears: by creating further opportunities to automate tasks and jobs, IT and AI may increase unemployment and reduce wages. Hence the idea suggested by some, that one should tax robots to avert that danger.

In this paper, we discuss the effects of automation on employment, appealing to both the existing literature on AI and automation and our recent empirical work using French firm-level data (Aghion et al., 2019, 2020). We first spell out the two contrasting views on the subject. A first view sees automation as primarily destroying jobs, even if this may ultimately result in new job creations taking advantage of the lower equilibrium wage induced by the job destruction. The prediction is that automation should reduce employment and the aggregate labor share. An alternative approach emphasizes the market size effect of automation: namely, automating firms become more productive, which enables them to lower their quality-adjusted prices and therefore to increase the demand for their products; the resulting

---

<sup>1</sup> Keynes, “Economic Possibilities for Our Grandchildren.”

increase in market size translates into higher employment by these firms. We provide empirical support for the second approach, drawing from our empirical work on French firm-level data and a growing literature covering multiple countries.

The paper is organized as follows. Section 2 presents the debate. Section 3 describes the emerging empirical consensus towards the more positive view of automation, with positive direct effects on employment at the firm level. Drawing on our recent empirical work, Section 4 describes the main methodological approaches and main findings from the literature using data on French plants, firms, and labor markets in recent years. Section 5 concludes.

## 2. The debate: what are the direct and indirect effects of automation on employment?

In this section we briefly present the two contrasting views on automation and employment.

### a. The “old view”: negative direct effects and positive indirect effects

The “negative” view is that the most direct effect of automation is to destroy employment and push wages downward. This direct effect may then be counteracted. In Acemoglu and Restrepo (2016) it is counteracted by the fact that automation depresses the equilibrium wage, which in turn encourages the creation of activities which initially employ labor (before being themselves subsequently automated); this in turn increases the demand for labor and therefore limits the wage decline. In Aghion, Jones and Jones (2017), the direct effect is counteracted by a “Baumol Cost Disease” effect whereby labor becomes increasingly scarcer than capital over time, which pushes wage upward (due to the complementarity between labor and capital at the aggregate level).

More formally, Acemoglu and Restrepo (2016) assume that final output is produced by combining the services of a unit measure of tasks  $X \in [N - 1, N]$ , according to:

$$Y = \left( \int_{N-1}^N X_i^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}$$

where: (i) tasks  $X_i$  with  $i > I$  are non automated, produced with labor alone; (ii) tasks  $X_i$  with  $i < I$  are automated, i.e. capital and labor are perfect substitutes within tasks, with  $\sigma > 1$  denoting the constant elasticity of substitution between tasks; (iii)

$$X_i = \alpha(i)K_i + \gamma(i)L_i$$

where: (a)  $\alpha(i)$  is an index function with  $\alpha(i) = 0$  if  $i > I$  and  $\alpha(i) = 1$  if  $i < I$ ; (b)  $\gamma(i) = e^{Ai}$ .

In the full-fledged Acemoglu-Restrepo model with endogenous technological change, the dynamics of  $I$  and  $N$  (i.e., the automation of existing tasks and the discovery of new lines) results from endogenous directed technical change. Under reasonable parameter values guaranteeing that innovation is directed towards using the cheaper factor, there exists a unique and (locally) stable Balanced Growth Path (BGP) equilibrium.

Stability of this BGP follows from the fact that an exogenous shock to  $I$  or  $N$  will trigger forces which bring the economy back to its previous BGP with the same labor share. The basic intuition for this result is that, if a shock leads to too much automation, then the decline in labor costs will encourage innovation aimed at creating new (more complex) tasks which exploit cheap labor, i.e. it will lead to an increase in  $N$ . In other words, the direct negative effect of automation of employment is mitigated by an indirect – general equilibrium – effect, whereby the depressing effect of automation on wages encourages entry of new activities which initially take advantage of labor becoming cheaper.

Aghion, Jones and Jones (2017) point to another counteracting force, namely the “Baumol Cost Disease” effect, which prevents automation from depressing wages too much. There it is the complementarity between existing automated tasks and existing labor-intensive tasks, together with the fact that labor becomes increasingly scarcer than capital over time, which allows for the possibility that the labor share remains constant over time.

More formally, final output is produced according to:

$$Y_t = A_t \left( \int_0^1 X_{it}^\rho di \right)^{\frac{1}{\rho}}$$

where  $\rho < 0$  (i.e., tasks are complementary),  $A$  is knowledge and grows at constant rate  $g$  and, as in Zeira (1998):

$$X_{it} = \begin{cases} L_{it} & \text{if not automated} \\ K_{it} & \text{if automated} \end{cases}$$

Under the assumption that a fraction  $\beta_t$  of tasks is automated by date  $t$ , we can re-express the above aggregate production function as:

$$Y_t = A_t (\beta_t^{1-\rho} K_t^\rho + (1 - \beta_t)^{1-\rho} L^\rho)^{1/\rho}$$

where  $K_t$  denotes the aggregate capital stock and  $L_t \equiv L$  denotes the aggregate labor supply.

In equilibrium, the ratio of capital share to labor share is equal to:

$$\frac{\alpha_{K_t}}{\alpha_L} = \left( \frac{\beta_t}{1 - \beta_t} \right)^{1-\rho} \left( \frac{K_t}{L_t} \right)^\rho$$

Hence an increase in the fraction of automated goods  $\beta_t$  has two offsetting effects on  $\frac{\alpha_{K_t}}{\alpha_L}$ : (i) first, a direct positive effect which is captured by the term  $\left( \frac{\beta_t}{1 - \beta_t} \right)^{1-\rho}$ ; (ii) second, a negative indirect effect captured by the term  $\left( \frac{K_t}{L_t} \right)^\rho$ , as we recall that  $\rho < 0$ . This latter effect relates to the well-known Baumol Cost Disease: namely, as  $\frac{K_t}{L_t}$  increases due to automation, labor becomes scarcer than capital which, together with the fact that labor-intensive tasks are complementary to automated tasks (indeed we assumed  $\rho < 0$ ), implies that labor will command a sustained share of total income.

While the above two models emphasize different counteracting forces which limit the wage decline induced by automation, both have in common that the direct effect of automation is to destroy employment. In particular, this direct effect would be observed within firms that automate.

### **b. The “new view”: positive direct effects and negative indirect effects**

Recent work suggests a more “positive” view of automation: the direct effect of automation may be to increase employment at the firm level, not to reduce it.<sup>2</sup> The reason is that firms and plants that automate become more productive. This allows them to offer a better quality-adjusted price than their competitors, and therefore to “steal business” away from their competitors, and more generally to expand the size of their markets (domestic and foreign). This in turn increases their demand for labor.

Note that this channel does not exclude the possibility that aggregate employment, at the national, industry, regional or commuting zone level may not respond so positively to automation and may even

---

<sup>2</sup> See Acemoglu et al. (2020), and Aghion et al. (2020).

react negatively to it. There may be an overall negative effect if automating firms induce a sufficiently large decline in employment for non-automating firms, and their exit. But a main difference with the “negative view”, is that, here, the direct dominant effect of automation is the positive productivity effect, which may then be counteracted by an indirect “creative destruction” or “eviction” effect. Furthermore, the negative indirect effect is partly borne by international competitors, which has implications for the desirability of taxing robots.

### **c. Implications for the taxation of robots**

Adjudicating between the two views is important for the debate of taxation of robots. As discussed in Aghion et al. (2020), the second view implies that unilateral taxation of robots by a given country could be counterproductive for employment in that country, because of business stealing effects across countries. According to that view, the direct positive effect of automation will benefit countries that keep automating, while the indirect negative effect will be shared across countries, given that competition operates in world markets. Therefore, as explained by Aghion et al. (2020), unilateral taxes on robots or other automation technologies may be detrimental to domestic employment: *“without international coordination, in a globalized world attempts to curb domestic automation in an effort to protect domestic employment may be self-defeating because of foreign competition.”*

In the next section, we confront the two view to recent evidence from the literature, covering many countries and time periods.

### **3. A survey of the empirical evidence from the recent literature**

Early analyses would point toward the former story based on macroeconomic equilibrium analyses (Keynes, 1930; Leontief, 1952; Lucas & Prescott, 1974), however these lacked empirical support. A next generation of studies were able to confront theoretical models to the data. Their analyses have been primarily run at the national or industry level and have mostly conveyed the idea of automation having a negative impact on aggregate employment and aggregate wages: automation is primarily labor saving. Yet these analyses fall short of describing the process that goes on within firms. It is only over the past few years, thanks to the increasing availability of new empirical datasets, that analyses of the effects of automation on employment could be performed at the firm level.

In this section, we provide an overview of the recent empirical literature on automation and employment. As our literature survey illustrates, the profession has evolved from the more “negative” view of automation as primarily destroying jobs (through its direct effect), towards the more “positive” view of automation as enhancing productivity, market size, and therefore labor demand and employment.

### **a. Mixed evidence from research designs using variation across industries and labor markets**

How should automation be measured? Until recently, the number of reliable sources on which empirical analyses of automation could be built was limited<sup>3</sup>. But since the 2010s, the International Federation of Robotics (IFR) has provided data on the deployment of robots by country and industry, and machine learning algorithms have made it possible to measure automation using text analysis of patents. Therefore, recent papers investigate new measures of automation, and notably the number of robots (Autor & Dorn, 2013; Acemoglu & Restrepo, 2020; Cheng et al., 2019; Dauth et al., 2021; Graetz & Michaels, 2018), or automation related patents (Mann and Püttmann, 2017; Webb, 2020).

Based on IFR aggregate data, the empirical findings in Acemoglu and Restrepo (2020) suggest that the job destruction effect of automation dominates. More precisely, the authors analyze the effect of the increase in industrial robot usage between 1990 and 2007 on US labor markets. Using within-country variation in robot adoption they estimate the local labor market effects of robots by regressing the change in employment and wages on the exposure to robots.<sup>4</sup> The authors find that one more robot per thousand workers reduces the employment to population ratio by about 0.2 percentage points and wage growth by 0.42 %, while productivity increases, and labor share decreases. According to their estimates, each robot installed in the US replaces six workers.

The Acemoglu-Restrepo methodology has been applied in several other countries. Chiacchio et al. (2018) find a displacement effect between three and four workers per robot in six European countries, but do not point to robust and significant results for wage growth. Aghion et al. (2019) find a displacement effect of ten workers per robot on French administrative data. However, using German data, Dauth et al. (2021) report a null effect of exposure to robots on aggregate employment. For low- and mid-skilled workers, they report lower wages, while at the aggregate level the use of industrial robots contributes to the fall in the labor share.

Other measures of automation also yield mixed results. For instance, Webb (2020)'s measure of automation relies on the analysis of patent texts, applied to two historical case studies, software and industrial robots. He highlights the displacement effect: jobs which were highly exposed to previous automation technologies saw declines in employment and wages over the relevant periods. However, the results of Mann et Püttmann (2017), who also measure automation using patent data, paint a different

---

<sup>3</sup> Earlier studies used the measure of computers or IT as a proxy (Krueger, 1993; Autor et al., 1998; Bresnahan et al., 2002).

<sup>4</sup> The local exposure to robots is an indirect measure of robot penetration at the local level, which is based on the rise in the number of robots per worker in each national industry on the one hand, and on the local distribution of labor between different industries on the other hand.

picture. Linking automation patents to industries and labor markets, they find a positive effect of automation on employment.

Several studies provide results in line with the first view whereby, as the equilibrium wage declines, industries may increase their demand for labor. Indeed, these studies report that following robot adoption, firms expand their demand for complementary tasks. The decline in manufacturing employment is thus offset by positive employment spillovers on other local industries in the service sector (Dauth et al., 2021; Mann et Püttmann, 2017; Gregory et al., 2016).

Distinctive predictions of the “job destruction” story include a fall in employment and/or in the equilibrium wage and a decline of the labor share. However, Graetz & Michaels (2018), use the robot count from IFR data on a panel of seventeen developed countries, find no effect of automation on aggregate employment, despite a reduction of the low-skilled workers’ employment share. On the contrary, they show that robot densification is associated with increases in both total factor productivity and wages, and with decreasing output prices. Using the same measure on a panel of fourteen European countries, Klenert et al. (2020) find that robot use is correlated with an increase in total employment.

Thus, overall, aggregate studies on automation/robotization and employment provide mixed evidence in favor of the more “negative” story.

#### **b. Firm-level research designs provide causal evidence supporting the “new view”**

A large number of recent studies using firm-level data supports the prediction a direct positive effect of automation on employment in automating firms: in France (Acemoglu et al., 2020, Aghion et al., 2020), in the United States (Bessen et al., 2019), in the United Kingdom (Chandler and Webb, 2019), in Canada (Dixon et al., 2019), in Denmark (Humlum, 2019), and in Spain (Koch et al., 2021). This positive effect may reflect either a net creation of jobs by automating firms or lower separation rates by these firms (Domini et al. 2019). Several of these studies provide quasi-experimental evidence to establish that automation *causes* an increase in employment at the firm level. In the next section, we describe the methodology in detail, focusing on our own empirical work on automation and employment at the plant and firm levels.

Thus, the “negative” story faces difficulties when confronted to firm-level analyses. And at odds with the predictions of the “pessimistic” story, most of the above-mentioned studies do not find evidence of a falling equilibrium wage nor of a decreasing labor share (e.g. Bessen et al., 2019; Dixon et al., 2019; Humlum, 2019; Koch et al., 2021; Aghion et al., 2020).



Overall, these studies support the view that automation inside a firm fosters labor productivity. It drives quality-adjusted prices down for consumers,<sup>5</sup> increases product demand and market share of the firm, which can result in net job growth. Provided that demand is elastic enough to prices, then growth in demand will offset job losses. The increase in the market share will only last until markets become saturated (Bessen, 2019).

Yet, the productivity effect may contribute to the crowding-out of non-automating firms by automating firms. Since the productivity effect inside the automating firm causes an increase in product demand, the market share of the firm goes up at the expense of its non-automating competitors. Firms whose competitors adopt robots experience significant declines in value added and employment (Acemoglu, 2020; Aghion et al. (2020), Koch et al., 2021). For example, Koch et al. find that robot-adopting firms create new jobs, expand the scale of their operations, while non-adopters incur negative output and lose employment because of tougher competition with high technology firms. Using industry-level analysis, Aghion et al. (2020) find that, for domestic employment the negative indirect effect is not large enough to offset the direct effect; indeed, business stealing occurs in part at the expense of other countries.

Babina et al. (2020) bring out a similar result with firm-level investment in AI technology. Firms that invest more in AI experience faster growth in sales and employment both at the firm- and industry-levels. AI allows the expansion of the most productive firms *ex ante*: they grow larger, gain sales, employment and market share. The authors report a null effect on productivity in the short run, perhaps because of the novelty of AI technologies, which have not been fully appropriated by workers.

Thus, drawing on different measures of automation, different countries, and various time periods, recent micro studies consistently point to the importance of the productivity effect, with positive employment effects within automating firms and potential displacement effects across firms. As Autor (2015) states it, *“journalists and even expert commentators tend to overstate the extent of machine substitution for human labor and ignore the strong complementarities between automation and labor that increase productivity, raise earnings, and augment demand for labor”*. At the firm level, automation has a positive impact on employment, which highlights the market size effect of automation. Automating firms become more productive, which enables them to lower their quality-adjusted prices; the resulting increase in market size translates into higher employment at the firm level.

### **c. Which workers benefit or lose from automation?**

---

<sup>5</sup> Aghion et al. (2020) provide direct empirical evidence on the response of consumer prices. Bonfiglioli et al. (2020) suggest that productivity gains from automation may not be entirely passed on to consumers in the form of lower prices.

Separate from the debate about the nature of direct and indirect effects of automation on employment, there is a debate about the types of jobs that are created or destroyed and the distribution effects of automation. The economic literature has long considered technological change to be labor augmenting and favorable to skilled workers. In the wake of the IT and computer revolution in the 1990s, the emphasis was given to the skill-biased technological change hypothesis. This hypothesis indeed supported the idea of complementarity between technology and skilled workers (see Acemoglu & Autor, 2011, for an overview). Technological change would result in the polarization of the job market, i.e., the slower increase in mid-wage occupations compared to both high-wage and low-wage occupations.

In the 2000s, following the critic of Card & DiNardo (2002), and the seminal paper of Autor et al. (2003), the academic consensus shifted to a labor-replacing view of automation in routine tasks. According to this idea, “traditional” automation replaces routine jobs, and creates more demand for non-routine jobs that cannot be performed by machines. Several studies have documented the disappearance of manufacturing and routine jobs (Autor et al., 2003; Jaimovich & Siu, 2012; Autor & Dorn, 2013; Charnoz & Orand, 2017; Blanas et al., 2019).

Coming back to firm-level studies, some of them highlight a reallocation of workers between occupations (Bessen, 2019; Bonfiglioli et al. 2020; Humlum, 2019; Acemoglu et al., 2020). Humlum (2019) notably reports a shift from low-skilled to high-skilled workers in Denmark: labor demand shifts from production workers toward tech workers, such as skilled technicians, engineers, or researchers. In the same vein, Bonfiglioli et al. (2020) show that robot imports by French firms increase productivity and the employment share of high-skill professions. Similarly, Bessen (2019) shows that computer automation causes growth in well-paid jobs and decreases in low-paid jobs. Using Canadian data, Dixon et al. (2019) document a polarization effect: investments in robotics are associated with shrinking employment for mid-skilled workers, but with increasing employment for low-skilled and high-skilled workers, notably managerial activities. This shift from low-skilled to high-skilled workers may also contribute to boosting productivity (Humlum, 2019; Acemoglu et al., 2020).

Yet, some studies do not find any reallocation effect between different types of workers and occupational categories (Domini et al. 2019; Aghion et al., 2020). This could be explained by a reallocation effect within jobs, since automation technologies generally do not replace entire jobs but only a certain number of tasks (Acemoglu and Autor 2011). Perhaps paradoxically, some human skills may become more valuable than ever in the presence of machines (Brynjolfsson & McAfee, 2011). Automation may thus lead to a restructuring of the task content of different jobs “within worker” (Aghion et al., 2020), enhancing labor productivity and, potentially, employment, but without any change in the skill structure of employment.

This is precisely the issue that Arntz et al. (2017) raise when they question Frey and Osborne (2017)'s analysis on the future of AI. Frey and Osborne (2017) tried to forecast the probability of computerization of 702 jobs and concluded that 47 % of employment in the US was at risk of automation in the next ten or twenty years, while only 33 % of jobs had a low risk of automation. But their analysis disregards the task content of jobs. Arntz et al. (2017) show that, when factoring in the heterogeneity of tasks within occupations, only 9 % of all workers in the US face a high risk of automation.

#### 4. Recent empirical evidence from France

We illustrate the main points from the preceding literature review using French data, drawing from our recent work (Aghion et al. 2019, 2020). We first show that labor market level analysis using IFR data provides mixed support in favor of the first view. Second, we show that firm level and plant level analyses using alternative measures of automation provides quasi-experimental evidence supporting the second view. We present the methodology and main results from our existing work, as well as novel complementary specifications.

##### a. Labor market level analysis using IFR data

Aghion et al. (2019) reproduce the method developed by Acemoglu and Restrepo (2017, hereafter AR) on French data over the 1994-2014 period, analyzing the impact on increased robotization on employment at the aggregate employment zone level.<sup>6</sup>

To measure exposure to robots at the labor market – defined as commuting zone – level, AR build a local exposure index, which combines two elements: (i) the number of robots per worker in each of industry on the one hand and (ii) the pre-existing share of employment in industry  $i$  for a given commuting zone  $c$ . Thus, this local exposure index exploits the initial heterogeneity in industry employment structures across commuting zone to distribute cross-industry variation in the stocks of robots in the various industries, observed nationwide during the sample period. More formally, the increases in robot exposure at the commuting zone level is defined as:

$$US \text{ robot exposure } 1993\_2007_c = \sum_{i \in I} l_{ci}^{1990} \left( \frac{R_{i,2007}^{US}}{L_{i,1990}^{US}} - \frac{R_{i,1993}^{US}}{L_{i,1990}^{US}} \right)$$

---

<sup>6</sup> AR analyze the effect of the increase in industrial robot usage between 1990 and 2007 on US local labor markets. They find that one more robot per thousand workers reduces the employment to population ratio by about 0.37 percentage points and wage growth by 0.73 percent.

where the sum is over all the 19 industries  $i$  in the IFR data;  $l_{ci}^{1990}$  stands for the 1990 share of employment in industry  $i$  for a given commuting zone  $c$ ;  $R_i$  and  $L_i$  stand for the stock of robots and the number of people employed in a particular industry  $i$ .

In line with AR, in Aghion et al. (2019) measure the increase in robot exposure in a French employment zone<sup>7</sup> between 1994 and 2014 as:

$$Robot\ exposure\ 1994\_2014_c = \sum_{i \in I} \frac{L_{ic,1994}}{L_{c,1994}} \left( \frac{R_{i,2014}}{L_{i,1994}} - \frac{R_{i,1994}}{L_{i,1994}} \right)$$

where  $L_{ic,1994}$  refers to employment in the employment zone  $c$  in industry  $i$  in 1994,  $L_{c,1994}$  refers to employment in employment zone  $c$  in 1994 and  $L_{i,1994}$  refers to employment in industry  $i$  in 1994.  $R_{i,1994}$  and  $R_{i,2014}$  respectively stand for the total number of robots in industry  $i$  in 1994 and 2014. This index reflects the exposure to robots per worker between 1994 and 2014. Our outcome variable of interest is the evolution of the employment-to-population ratio between 1990 and 2014.

In the baseline OLS specification, we study the impact of exposure to robots on the evolution of employment-to-population ratio. Then, we add controls such as an exposure index for information and communication technologies (ICT)  $TICExp_r$ , built in a similar way as the exposure to robot index and an international trade exposure index  $TradeExp_r$  from China and Eastern Europe. In some regressions, we also add a vector  $X_c$  of control for the employment zone  $c$ : demographic characteristics, manufacturing shares, broad industry shares, broad region dummies and specific industry shares within manufacturing. We can write:

$$\Delta \frac{L_{c,1994}}{Pop_{c,1994}} = \alpha + \beta_1 RobotsExp_c + \beta_2 TradeExp_c + \beta_3 TICExp_c + \gamma X_c + \epsilon_c$$

To measure the impact of exposure to robots on local labor markets, the strategy adopted is similar to the one initiated by Autor et al. (2013): the observed change in robot exposure in U.S. industries is instrumented with changes in robot exposure in the same industries in other developed economies. This approach helps address U.S.-specific threats to identification affecting the OLS approach: one may imagine a shock, which we do not captured in our controls, but which may impact both the installation of robots at and local labor markets dynamics. Following AR, the stocks of robots in industries from other developed countries (Germany, Denmark, Spain, Italy, Finland, Norway, Sweden, and the United

---

<sup>7</sup> According to the official definition provided by the INSEE, an employment zone is a geographical area within which most of the labor force lives and works. It provides a breakdown of the territory adapted to local studies on employment.

Kingdom) are used to build other indexes of exposure to robots. These new indexes are then used to instrument the exposure index built on French stock of robots.

In this shift-share IV research design, identification arises from the heterogeneity in robotization shocks across industries, which is projected to the regional level. Indeed, as described in Borusyak et al. (2021), the employment shares  $l_{ci}^{1990}$  are not tailored to exposure to robotization: they are “generic”, in that they could conceivably measure an observation’s exposure to multiple shocks, both observed and unobserved. Identification stems from the robotization shocks  $\frac{R_{i,2007}^{US}}{L_{i,1990}^{US}} - \frac{R_{i,1993}^{US}}{L_{i,1990}^{US}}$  and  $\frac{R_{i,2014}}{L_{i,1994}} - \frac{R_{i,1994}}{L_{i,1994}}$ .

Accordingly, it is important to control for industry-level characteristics that may contaminate the industry-level identifying variation, such as whether an industry belongs to manufacturing. Absent such controls, we would conflate the potential effects of robotization with broad sectoral trends.

**Table 1: Effect of robot exposure on employment-to-population ratio, 1990-2014, OLS estimates**

TABLE 1 – OLS estimates - Robots exposure on employment

	Dependant variable : Change in employment-to-population ratio 1990-2014 (in % points)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>RobotsExposure</i> <sub>1994–2014</sub>	-1.090*** (0.253)	-0.749*** (0.263)	-0.594** (0.239)	-0.515** (0.243)	-0.169 (0.239)	-0.549* (0.294)	-0.398 (0.244)	-0.430 (0.324)	-1.074 (0.768)	-1.035 (0.783)
<i>TICExposure</i> <sub>1994–2014</sub>		-3.099* (1.586)	-2.397 (1.594)	-2.495* (1.455)		-0.304 (1.620)	-0.165 (1.576)	-0.154 (1.588)	1.519 (1.641)	1.493 (1.648)
<i>TradeExposure</i> <sub>1994–2014</sub>		-0.743*** (0.247)	-0.690*** (0.215)	-0.825*** (0.239)		0.0857 (0.243)	-0.123 (0.278)	-0.124 (0.280)	0.200 (0.335)	0.201 (0.337)
Demographics			Yes				Yes	Yes	Yes	Yes
Region dummies				Yes			Yes	Yes	Yes	Yes
Manufacturing industry share					Yes	Yes	Yes	Yes	Yes	Yes
Other broad industry shares						Yes	Yes	Yes		
Specific manufacturing industry shares									Yes	Yes
Removes highly exposed areas							Yes			Yes
Observations	297	297	297	297	297	297	297	295	297	295
R-squared	0.058	0.090	0.198	0.205	0.174	0.249	0.407	0.406	0.409	0.408

Demographics control variables are population share by level of education and population share between 25 and 64 years old. Broad industry shares cover the share of workers in agriculture, construction, retail and the share of women in manufacturing in 1994. Specific manufacturing industry shares cover the share of workers in automotive, rubber, food and the share of women in manufacturing in 1994. Broad region dummies refers to the 13 metropolitan regions of France. Highly exposed areas are Poissy and Belfort-Montbéliard-Héricourt. Robust standard errors in parentheses. Levels of significance : \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Sources : IFR, COMTRADE, EUKLEMS, DADS, Census data.

Source: Data from Aghion et al. (2019).

Table 1 displays the results of the OLS estimates. This table shows a negative correlation between exposure to robots and change in employment-to-population ratio. However, we observe that the level of significance decreases as more controls are added. Significance is lost in column 5 once a control for the local manufacturing industry share is included and the point estimate fall substantially, indicating that broad sector trends play an important role. The correlation is marginally significant in column 6 and non-significant in columns 7 through 10, where we add several types of controls simultaneous or exclude the commuting zones with the highest exposure to robots.

In the instrumental variable regression shown in Table 2, the coefficients of robot exposure are significant while we consider broad controls from column (1) to (4). Column (1) begins with the

regression without any control and finds a negative effect: one more robot per 1000 workers leads to a drop in the employment-to-population ratio of 1.317 percentage points. Column (2) adds controls on ICT and imports exposures and the magnitude remains the same. Then, columns (3) and (4) successively test the impact of demographic characteristics and broad region dummies, leaving the results almost unaffected. In column (5), adding a control for the manufacturing share alone is sufficient to lose significance and substantially reduce the point estimate. The result highlights again the importance of controlling for broad industry trends, as emphasized by Borusyak et al. (2021).

Combining different sets of controls, the specifications in columns (6) through (8) deliver negative and statistically significant IV estimates. However, in columns (9) and (10), we replace broad industry shares controls by controls for specific industry shares with manufacturing at commuting zone level. Specifically, we control for the three 2-digit industries that have the highest number of robots at the end of the period, which account for 74% of the total number of robots in 2014: automotive, rubber and food industries. Thus, these are key industries relative to the construction of the index. The coefficients remain negative but become non-significant. These last two columns highlight that the results are sensitive to the inclusion of a few highly robotized industries. Table 3 in appendix reproduces the same estimations over the period 1990-2007.

**Table 2: Effect of robot exposure on employment-to-population ratio, 1990-2014, IV estimates**

TABLE 2 – IV estimates - Robots exposure on employment

	Dependant variable : Change in employment-to-population ratio 1990-2014 (in % points)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>RobotsExposure</i> <sub>1994–2014</sub>	-1.317*** (0.325)	-1.010*** (0.322)	-0.974*** (0.271)	-0.737** (0.296)	-0.389 (0.248)	-0.790*** (0.300)	-0.686*** (0.241)	-0.986*** (0.351)	-1.305 (0.799)	-1.221 (0.812)
<i>TICExposure</i> <sub>1994–2014</sub>		-2.569 (1.618)	-1.699 (1.578)	-2.094 (1.444)		-0.176 (1.590)	-0.0323 (1.518)	0.101 (1.538)	1.590 (1.601)	1.547 (1.609)
<i>TradeExposure</i> <sub>1994–2014</sub>		-0.670*** (0.242)	-0.589*** (0.211)	-0.773*** (0.230)		0.110 (0.240)	-0.0922 (0.276)	-0.0882 (0.279)	0.198 (0.322)	0.199 (0.323)
Demographics			Yes				Yes	Yes	Yes	Yes
Region dummies				Yes			Yes	Yes	Yes	Yes
Manufacturing industry share					Yes	Yes	Yes	Yes	Yes	Yes
Other broad industry shares						Yes	Yes	Yes		
Specific manufacturing industry shares									Yes	Yes
Removes highly exposed areas								Yes		Yes
Observations	297	297	297	297	297	297	297	295	297	295
First-stage F statistic	57.2	42.6	45.8	46.0	32.6	28.7	35.1	18.9	16.5	16.3
R-squared	0.055	0.087	0.193	0.203	0.172	0.248	0.405	0.400	0.409	0.408

Demographics control variables are population share by level of education and population share between 25 and 64 years old. Broad industry shares cover the share of workers in agriculture, construction, retail and the share of women in manufacturing in 1994. Specific manufacturing industry shares cover the share of workers in automotive, rubber, food and the share of women in manufacturing in 1994. Broad region dummies refers to the 13 metropolitan regions of France. Highly exposed areas are Poissy and Belfort-Montbéliard-Héricourt. Robust standard errors in parentheses. Levels of significance : \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Sources : IFR, COMTRADE, EUKLEMS, DADS, Census data.

Source: Data from Aghion et al. (2019).

Thus, the IFR data at the industry level provides mixed evidence on the effects of employment, due to the small number of industries that are used as the source of identifying variation. Furthermore, the finding of a negative or non-significant effect of robotization on employment at the aggregate employment zone level could be consistent with either the “new view” or “old view” on automation and employment. Indeed, this result could reflect either the fact that robotizing firms destroy employment,

and that this direct effect is not fully offset by the counteracting general equilibrium effect working through wage reduction and the resulting entry of new activities; or the fact that the positive market size effect of automation at the firm level is more than offset by the job destructions in the non-automating firms that are partly or fully driven out of the market by the automating firms. To alleviate the limitations of the research design and find out more about which of these two stories applies, we need to move to a more microeconomic analysis of the effect of automation on employment.

#### **b. Firm-level and plant-level analyses**

In Aghion, Antonin, Bunel and Jaravel (2020), henceforth AABJ, we use two complementary measures as proxies for automation, at the firm level and plant level. At the firm-level, we use the balance sheet value of *industrial equipment and machines* in euros, which is available for all French firm between 1995 and 2017. This type of capital is defined as “the equipment and machines used for the extraction, processing, shaping and packaging of materials and supplies or for carrying out a service” (industrial machines) and “instruments or tools that are added to an existing machine in order to specialize it in a specific task” (industrial equipment). Within manufacturing industry, this type of capital accounts for 59% of total capital. Our second measure of automation follows the Encyclopaedia Britannica (2015), which defines automation technology as a “class of electromechanical equipment that is relatively autonomous once it is set in motion on the basis of predetermined instructions or procedures”.<sup>8</sup> In the manufacturing industry, the French statistical office (Insee) records electricity consumption for motors directly used in the production chain (motive power) since 1983. It distinguishes motive power from other potential uses of electricity such as thermic/thermodynamic, or electrolysis. Thus, we are able to proxy automation by motive power, which exclude heating, cooling or servers uses.

We perform two types of event studies: (i) “extensive margin” event studies at firm level, exploiting the timing of the large investment in industrial equipment and machines for each firm as an automation event, and (ii) distributed lead-lag analysis at firm and plant level that allows for delayed responses changes in automation and takes into account continuous changes in the stock of machines.

Our main finding from the event studies is that the impact of automation on employment is positive, and in fact increases over time: namely, a 1 percent increase in automation in a plant today increases employment by 0.2 percent immediately and by 0.4 percent after ten years (see Figure 1 below, reproduced from AABJ with permission). Results are similar at firm level. In other words, conditionally on surviving, automation leads to a net increase in employment by automating firms and plants.

---

<sup>8</sup> Definition from Encyclopaedia Britannica (2015), “Automation.”




Figure 1: Plant-Level Event Study of Automation on Employment  
 Source: reproduced from AABJ with permission.

Next, Figure 2 (reproduced from AABJ with permission) shows that automation also translates into an increase in a firm’s total sales in the years after the firm automates. The effect remains stable from year of investment in automation to eight years after.




Figure 2: Firm-Level Event Study of Automation on Sales  
 Source: reproduced from AABJ with permission.

All these findings speak to a “productivity” effect of automation, in line with the “positive view” spelled out in the previous section: namely, firms that automate more become more productive. This enables them to obtain larger market shares, because their products offer consumers better value for money than



their competitors. The resulting gain in market share prompts those firms that automate to produce at a larger scale, and therefore to hire more employees.

However, the event study research design does not fully address potential correlated demand and supply shocks that could occur exactly at the same time as the increase in automation. Thus, to estimate the causal effects of automation on employment, sales, wages, and the labor share across firms, AABJ use a quasi-experimental shift-share design.

In fact, the ideal design would randomly assign purchasing prices for machines across firms. In AABJ, our idea is to approximate this hypothetical experiment using a shift-share instrument, which leverage two components: (i) the time variation in the implicit cost of imported machines over time across international trading partners (the “shift” component); (ii) the heterogeneity in pre-existing supplier relationships across French firms (the “exposure shares” component). The ideal “shock” variable would be the expected quality-adjusted price of imported machines by French manufacturing firms. However, we cannot directly observe these prices and that is why, instead, we infer changes in quality-adjusted price from changes in export flows of these foreign machines.

The intuition behind the shift-share instrument is that firms will be differentially exposed to these changes in quality-adjusted price of machines from different trading partners due to their sticky pre-existing relationships. For instance, if two French firms A and B import respectively 80% and 20% of their machines from Italy, and machines produced Italy suddenly have a better quality-adjusted price, firm A will have more incentives to automate than firm B due to its strong established relationship with Italian suppliers of machines.

The estimates of the impact of automation on employment using the shift-share instrument are in line with the previous findings from the event studies. The elasticity of firm employment to automation that we find ranges between 0.397 and 0.444 (Table 3A of AABJ), significant at the 5% or 1% level depending on the set of controls, and the first stage F statistic remains close to 10 in all specifications.

Next, we conduct the same exercise with sales and the labor share at firm level. We find that sales increase in response to increased automation, with elasticities ranging from 0.395 to 0.512 (Table 3B of AABJ) across specifications. Using the same specifications, we cannot reject that there is no impact of automation on the labor share, which in turn suggests that the productivity effect may offset the task substitution channel in a way that leaves the labor share unchanged at the firm level.

What happens when we move from firm or plant level to industry level? AABJ find a positive effect of automation on employment also at industry level, particularly in industries that more exposed to

international trade (in terms of their export ratios). This again speaks to the importance of the productivity effect: in industries that are more exposed to international trade, French firms that automate expand their export market at the expense of foreign firms. This in turn explains why, particularly in these industries, the productivity effect is the dominant effect, as it is mostly foreign firms in foreign markets which suffer from the resulting business stealing. In a closed economy, domestic non-automating firms would suffer from the business-stealing by the automating firms; the increase in employment by the automating domestic firms would be more likely to be counteracted by the job destruction by non-automating domestic firms.

Figure 3, using data from AABJ, illustrates this business-stealing – or eviction - effect: firms that invest significantly in new industrial equipment substantially lower their likelihood of going out of business over the following ten years compared to firms that do not make such an investment.




Figure 3: Effect of a substantial investment in industrial equipment on probability of firm exit. Source: Data from AABJ (2020).

**5. Conclusion**

In this paper, we relied on both the existing literature and our own empirical work to discuss the effects of automation on employment. We pointed to two contrasting views on the subject. A first view sees automation as primarily destroying jobs, even if this may ultimately result in new job creations taking advantage of the lower equilibrium wage induced by the job destruction. A second view emphasizes the productivity effect of automation as the main direct effect: namely, automating firms become more productive, which enables them to lower their quality-adjusted prices and therefore to increase the demand for their products; the resulting increase in market size translates into higher employment by

these firms. We provided direct empirical evidence supporting the second view, and we showed that the empirical literature on automation and employment was also leaning in that direction.

Overall, automation is thus not in itself an enemy of employment. By modernizing the production process, automation makes firms more competitive, which enables them to win new markets and therefore to hire more employees in a globalized world.

We can think of several avenues for further empirical research on automation and the labor market. One would be to explore how automation interacts with outsourcing and international trade. Another avenue would be to distinguish between different types of sectors and industries. A third avenue would be to introduce the distinction between routine and non-routine jobs. These and other extensions of the analyses surveyed in this paper are promising directions for future research.

## **Bibliography**

**Acemoglu, D. & Autor, D. (2011).** Skills, Tasks and Technologies: Implications for Employment and Earnings. In: O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Vol. 4, Elsevier, June 2011, chapter 12, 1043–1171.

**Acemoglu, D. & Restrepo, P. (2016).** The Race between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment. *NBER Working Paper*, 22252, National Bureau of Economic Research.

**Acemoglu, D., Lelarge, C. & Restrepo, P. (2020).** Competing with Robots: Firm-Level Evidence from France. *AEA Papers and Proceedings*, 110, 383-88.

**Acemoglu, D. & Restrepo, P. (2020).** Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6), 2188-2244.

**Aghion, P., Jones, B. & Jones, C. (2017).** Artificial Intelligence and Economic Growth. *NBER Working Paper*, 23928, National Bureau of Economic Research.

**Aghion, P., Antonin, C. & Bunel, S. (2019).** Artificial Intelligence, Growth and Employment: The Role of Policy. *Economie et Statistique / Economics and Statistics*, 510-511-512, 149–164.

**Aghion, P., Antonin, C., Bunel, S. & Jaravel, X. (2020).** What are the Labor and Product Market Effects of Automation? New Evidence from France. *CEPR Discussion Paper*, 14443.

**Autor, D., Katz, L.F. & Krueger, A. B. (1998).** Computing inequality: have computers changed the labor market? *Quarterly Journal of Economics*, 113(4), 1169-1213.

**Autor, D. & Dorn, D. (2013).** The growth of low skill service jobs and the polarization of the US labor market, *American Economic Review*, 103(5), 1553-1597.

**Autor, D., Levy, F. & Murnane, R. J., (2003).** The skill content of recent technological change: an empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279-1333.

**Autor, D. H., Dorn, D. & Hanson, G. H. (2015).** Untangling Trade and Technology: Evidence from Local Labor Markets. *The Economic Journal*, 125(584), 621-646.

**Autor, D. (2015).** Why are there still so many jobs? The history and future of workplace automation. *Journal of economic perspectives*, 2015, 29 (3), 3–30.

**Autor, D. H., Dorn, D. & Hanson, G. H. (2013).** The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6), 2121–68.

**Arntz, M., Gregory, T. & Ziehrhahn, U. (2017).** Revisiting the risk of automation. *Economics Letters*, 159, 157-160.

**Babina, T., Fedyk, A., He, A., & Hodson, J. (2020).** Artificial Intelligence, Firm Growth, and Industry Concentration, *Working Paper*.

**Bessen, J. (2019).** Automation and jobs: when technology boosts employment, *Economic Policy*, 34 (100), October 2019, 589–626.

**Bessen, J., Goos, M., Salomons, A. & van den Berge, W. (2019).** Automatic Reaction: What Happens to Workers at Firms that Automate?, *Working Paper*.

**Blanas, S., Gancia, G. & Lee, S. Y. (2019).** Who is Afraid of Machines? *CEPR Discussion Paper*, 13802.

**Bonfiglioli, A., Crinò, R., Fadinger, H. & Gancia, G. (2020).** Robot Imports and Firm-Level Outcomes, *CEPR Discussion Paper*, 14593.

**Bresnahan, T.F., Brynjolfsson, E. & Hitt, L.M. (2002).** Information technology, workplace organization, and the demand for skilled labor: firm-level evidence. *Quarterly Journal of Economics*, 117(1), 339-376.

**Brynjolfsson, E. & McAfee, A. (2011).** Race against the machine: how the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy. *Digital Frontier Press*, Lexington, MA.

**Card, D. & DiNardo, J. E. (2002).** Skill biased technical change and rising wage inequality: some problems and puzzles. *Journal of Labor Economics*, 20, 733-783.

**Chandler, D. & Webb, M. (2019).** How Does Automation Destroy Jobs? The 'Mother Machine' in British Manufacturing, 2000-2015, *Working Paper*.

**Charnoz, P. & Orand, M. (2017).** Technical change and automation of routine tasks: Evidence from local labour markets in France, 1999-2011. *Économie et Statistique*, 497, 103-122.

**Cheng, H., Jia, R., Li, D., & Li, H. (2019).** The Rise of Robots in China. *Journal of Economic Perspectives*, 33(2), 71-88.

**Chiacchio, F., Petropoulos, G. & Pichler, D. (2018).** The impact of industrial robots on EU employment and wages: A local labor market approach. *Bruegel Working paper*, 2.

**Dauth, W., Findeisen, S., Südekum, J. & Wößner, N. (2021).** The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, forthcoming.

**Dixon, J., Hong, B. & Wu, L. (2019).** The Employment Consequences of Robots: Firm Level Evidence, Working Paper.

**Domini, G., Grazzi, M., Moschella, D., Treibich, T. (2019).** Threats and opportunities in the digital era: Automation spikes and employment dynamics, *LEM Working Paper Series*, 2019/22.

**Frey, C. B. & Osborne, M. A. (2017).** The future of employment: How susceptible are jobs to computerization? *Technological Forecasting & Social Change*, 114, 254-280.

**Graetz, G. & Michaels, G. (2018).** Robots at work. *Review of Economics and Statistics*, 100(5), 753-767.

**Gregory, T., Salomons, A. & Zierahn, U. (2016).** Racing with or Against the Machine? Evidence from Europe. *ZEW Discussion Paper*, 16-053.

**Humlum, A. (2019).** Robot Adoption and Labor Market Dynamics. *Working Paper*, 2019.

**Jaimovich, N. & Siu, H.E. (2012).** The trend is the cycle: job polarization and jobless recoveries. *NBER Working Paper*, 18334, National Bureau of Economic Research.

**Keynes, J. M. (1932).** Economic possibilities for our grandchildren (1930). In: *Essays in Persuasion* (New York: Harcourt Brace, 1932), 358-373.

- Klenert, D., Fernández-Macías, E. & Antón-Pérez, J. I. (2020).** Do robots really destroy jobs? Evidence from Europe. *JRC Working Papers on Labour, Education and Technology*, 2020/01.
- Koch, M., Manuylov, I., Smolka, M. (2021).** Robots and Firms, *The Economic Journal*, forthcoming.
- Krueger, A.B. (1993).** How computers have changed the wage structure: evidence from microdata, 1984–1989. *Quarterly Journal of Economics*, 108(1), 33-60.
- Leontief, W. (1952).** Machines and man. *Scientific American*, 187(3), 150-64.
- Lucas, R.E. & Prescott, E.C. (1974).** Equilibrium search and unemployment. *Journal of Economic Theory*, 7(2), 188-209.
- Mann, K. & Püttmann, L. (2017).** Benign Effects of Automation: New Evidence from Patent Texts. *Unpublished manuscript*.
- Webb, M. (2020).** The Impact of Artificial Intelligence on the Labor Market. *Working Paper*.
- Zeira, J. (1998).** Workers, Machines, and Economic Growth. *Quarterly Journal of Economics*, 113, 1091-1117.

## APPENDIX

**Table 3: The effect of robot exposure on employment, 1990-2007, IV estimates**

TABLE 3 – IV estimates - Robots exposure on employment

	Dependant variable : Change in employment-to-population ratio 1990-2007 (in % points)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>RobotsExposure</i> <sub>1994–2007</sub>	-0.382*** (0.119)	-0.344* (0.198)	-0.508*** (0.195)	-0.148 (0.197)	-0.123 (0.149)	-0.560** (0.217)	-0.438** (0.198)	-0.633** (0.298)	-0.939 (1.043)	-0.703 (1.197)
<i>TICExposure</i> <sub>1994–2007</sub>		-0.322 (1.613)	0.990 (1.611)	-1.274 (1.571)		2.844 (2.142)	1.845 (2.019)	2.184 (2.056)	2.143 (2.090)	1.925 (2.101)
<i>TradeExposure</i> <sub>1994–2007</sub>		-0.217 (0.319)	-0.285 (0.293)	-0.400 (0.324)		0.301 (0.347)	0.107 (0.383)	0.111 (0.391)	0.477 (0.462)	0.467 (0.458)
Demographics			Yes				Yes	Yes	Yes	Yes
Region dummies				Yes			Yes	Yes	Yes	Yes
Manufacturing industry share					Yes	Yes	Yes	Yes	Yes	Yes
Other broad industry shares						Yes	Yes	Yes		
Specific manufacturing industry shares									Yes	Yes
Removes highly exposed areas								Yes		Yes
Observations	297	297	297	297	297	297	297	295	297	295
First-stage F statistic	53.1	43.0	44.5	47.8	39.1	38.3	46.4	17.8	7.3	5.5
R-squared	0.004	0.007	0.075	0.129	0.039	0.144	0.293	0.284	0.261	0.264

Demographics control variables are population share by level of education and population share between 25 and 64 years old. Broad industry shares cover the share of workers in agriculture, construction, retail and the share of women in manufacturing in 1994. Specific manufacturing industry shares cover the share of workers in automotive, rubber, food and the share of women in manufacturing in 1994. Broad region dummies refers to the 13 metropolitan regions of France. Highly exposed areas are Poissy and Belfort-Montbéliard-Héricourt. Robust standard errors in parentheses. Levels of significance : \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Sources : IFR, COMTRADE, EUKLEMS, DADS, Census data.

Source: Data from Aghion et al. (2019).