

Artificial Intelligence, Growth and Employment: The Role of Policy

Philippe Aghion*, Céline Antonin** and Simon Bunel***

Abstract – In this survey paper, we argue that the effects of artificial intelligence (AI) and automation on growth and employment depend to a large extent on institutions and policies. We develop a two-fold analysis. In a first section, we survey the most recent literature to show that AI can spur growth by replacing labor by capital, both in the production of goods and services and in the production of ideas. Yet, we argue that AI may inhibit growth if combined with inappropriate competition policy. In a second section, we discuss the effect of robotization on employment in France over the 1994-2014 period. Based on our empirical analysis on French data, we first show that robotization reduces aggregate employment at the employment zone level, and second that non-educated workers are more negatively affected by robotization than educated workers. This finding suggests that inappropriate labor market and education policies reduce the positive impact that AI and automation could have on employment.

JEL Classification: J24, O3, O4

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

Reminder:

The opinions and analyses in this article are those of the author(s) and do not necessarily reflect their institution's or Insee's views.

* *Collège de France, London School of Economics* (paghion@fas.harvard.edu)

** *Sciences Po OFCE, Collège de France* (celine.antonin@sciencespo.fr)

*** *Insee, Paris School of Economics* (simon.bunel@insee.fr)

The authors thank David Hémons, Adrien Lagouge, Claire Lelarge and Richard Rogerson for helpful comments and suggestions. We thank Simon Margolin for outstanding research assistance.

“The biggest harm that AI is likely to do to individuals in the short term is job displacement, as the amount of work we can automate with AI is vastly bigger than before. As leaders, it is incumbent on all of us to make sure we are building a world in which every individual has an opportunity to thrive.”

Andrew Ng, Stanford University

Artificial intelligence (or AI) is typically defined as the capability of a machine to imitate intelligent human behavior. Hence AI can be seen as the latest form of automation, a fourth automation wave following the steam engine revolution in the 18th century, the combustion engine revolution in the early 20th century and the semiconductor and IT revolution in the 1970s-1980s. In this survey paper, we argue that the effects of AI and automation on growth and employment depend to a large extent on institutions and policies.

Thus, the next section is dedicated to discussing the effects of AI and automation on economic growth: on the one hand, as argued by Zeira (1998), Hémous & Olsen (2014), Acemoglu & Restrepo (2016) and Aghion *et al.* (2017), AI can spur growth by replacing labor which is in finite supply by capital which is in unbounded supply, both in the production of goods and services and in the production of ideas; on the other hand, AI may inhibit growth if combined with inappropriate competition policy.

In a second section, we discuss the effects of AI and automation on aggregate employment. We present and discuss Acemoglu & Restrepo’s results (2017, hereafter AR-2017), and we build on their method to look at the effect of automation on employment in France over the 1994-2014 period. We estimate that the installation of one extra robot reduced aggregate employment by 10 workers at the employment zone level, and find an order of magnitude similar to AR-2017 – who found a loss of 6.2 jobs per extra robot in the US. We also find that non-educated workers are more negatively affected by robotization than educated workers. This in turn suggests that inappropriate labor market and education policies further reduce the positive impact that AI and automation could have on employment.

AI and Economic Growth

The Zeira Model

A benchmark model to think about the relationship between AI, automation, and growth

is Zeira (1998). Here we reproduce the presentation of Zeira (1998) in Aghion *et al.* (2017), henceforth referred to as “AJJ”. Zeira assumes that final output is produced according to:

$$Y = AX_1^{\alpha_1} X_2^{\alpha_2} \dots X_n^{\alpha_n}$$

where $\sum \alpha_i = 1$ and intermediate inputs X_i are produced according to:

$$X_i = \begin{cases} L_i & \text{if not automated} \\ K_i & \text{if automated} \end{cases}$$

While Zeira thought of the X_i as intermediate goods, they can also be viewed as tasks (Acemoglu & Autor, 2011). Hence, tasks that have not yet been automated can be produced one-for-one by labor. Once a task is automated, one unit of capital can be used instead (Aghion *et al.*, 2017).

Automation spurs economic growth as it replaces labor – which is in finite supply – by capital which is in unbounded supply, as a basic production input. Indeed, letting K and L denote aggregate capital stock and labor supply respectively, then final output is ultimately produced (up to a constant) according to:

$$Y = AK^\alpha L^{1-\alpha}$$

where α reflects the overall share of tasks that have been automated.

Hence the rate of growth of per-capita GDP (i.e. of $y = Y/L$) is equal to:

$$g_y = \frac{g_A}{1-\alpha}$$

Automation (e.g. as resulting from the AI revolution) will increase α which in turn will lead to an increase in g_y i.e. to an acceleration of growth. One issue with this model however, is that it predicts a rise in capital share, which in turn contradicts the so-called Kaldor fact that the capital share tends to be stable over time.

The Acemoglu-Restrepo Model

Acemoglu & Restrepo (2016) extend Zeira (1998) by assuming 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 tasks X_i are non automated, produced with labor alone when $i > I$, and are automated *i.e.* capital and labor are perfect substitutes when $i < I$. σ is the constant elasticity of substitution between tasks. With no major loss of insight, we can write:

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

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

In the full-fledged AR-2017 model with endogenous technological change, the dynamics of I and N (*i.e.* the automation of existing tasks and the discovery of new product 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 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.

What makes the capital share remain constant on this BGP is the fact that the automation of existing tasks is exactly offset by the creation of new tasks which require labor, at least initially. One special feature of this model is the assumption that technical progress $\gamma(i) = e^{A_i}$ multiplies labor, but not capital, even after automation takes place. Yet, it seems difficult to conceive of concrete examples where an automated production process would be replaced by a highly productive labor-intensive task.¹ Another feature is that the constancy of the capital share relies entirely on the continuous arrival of new labor intensive tasks. This prediction will be challenged in our section on AI and employment. The model by AJJ (2017), which also extends Zeira (1998), turns out to address these two objections.

Baumol's Cost Disease and the AJJ model

In the following model by Aghion *et al.* (2017), a greater fraction of tasks are being automated over time since there are no new labor-intensive tasks to compensate for the automation of existing tasks. This feature is shared by Zeira's model. Yet the complementarity between

existing automated tasks and existing labor intensive tasks, together with the fact that labor becomes increasingly scarcer than capital over time, allows for the possibility that capital 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):

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

Assuming that a fraction β_t of tasks is automated by date t , we can re-express the above aggregate production function as:

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

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

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

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

Hence an increase in the fraction of automated goods β_t has two offsetting effects on $\frac{\alpha_{K_t}}{\alpha_{L_t}}$: first, a direct positive effect which is captured by the term $\left(\frac{\beta_t}{1-\beta_t} \right)^{1-\rho}$; 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's cost disease: namely, as $\frac{K_t}{L_t}$ increases as a result of 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.

How about long run growth in this model? Let's first consider the case where a constant fraction of not-yet-automated tasks become automated

1. In Hémous & Olsen (2016), new tasks do not feature a higher productivity for labor and add up to existing tasks instead of replacing automated tasks. As a result, their model predicts a decline in the labor share which matches quantitatively the decline observed in the US.

each period, *i.e.* $\beta = \theta(1 - \beta_t)$. In this case, one can show that the growth rate converges to a constant in the long run.

Next, consider the case where all tasks become automated in finite time, *i.e.* where $\beta_t \equiv 1$ for $t > T$. Then, for $t > T$ aggregate final good production becomes $Y_t = A_t K_t$, so that, if capital accumulates over time according to $\dot{K} = sY - \delta K$, we get a long run growth rate equal to $g_Y = g_A + sA - \delta$, which increases unboundedly over time as A grows at the exponential rate g_A .

Automation in the Production of Ideas

AJJ consider the polar case where the production of goods and services uses labor only, whereas automation affects the production of knowledge. Somehow, this gets us closer to what AI is all about, over and beyond automation. Namely, AJJ assume:

$$Y_t = A_t L_t$$

with:

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

where, as before, $\rho < 0$ and

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

Assuming that a fraction β_t of “idea-producing” tasks are automated by date t , then the above knowledge growth equation becomes:

$$\dot{A} = A_t^\phi \left(\beta_t^{1-\rho} K_t^\rho + (1 - \beta_t)^{1-\rho} L^\rho \right)^{\frac{1}{\rho}}$$

Let’s first consider the case where a constant fraction of not-yet-automated tasks become automated at each period, *i.e.* $\beta = \theta(1 - \beta_t)$. In this case, one can show that:

$$g_Y = g_A = -\frac{1-\rho}{\rho} \frac{\theta}{1-\phi}$$

so that even though we assume decreasing returns to knowledge accumulation as in Jones (1995), *i.e.* $\phi > 0$, automation in the production of ideas maintains a positive long-run growth rate of (per capita) GDP.

Next, consider the case where all tasks become automated in finite time, *i.e.* where $\beta_t \equiv 1$ for $t > T$. Then, for $t > T$, the growth of knowledge satisfies the equation:

$$\dot{A} = A_t^\phi K_t$$

where:

$$\dot{K} = sY - \delta K$$

In this case AJJ shows that $A_t = Y_t / L$ becomes infinite in finite time. This extreme form of explosive growth is referred to as a “singularity”.

Explaining the Decline in Growth

Given the predictions of theoretical models, why haven’t we observed a growth outburst in developed countries, particularly in the US over the past decade, even though automation and AI are affecting a growing share of activities? One explanation for the absence of explosive growth may simply be that some essential inputs to production or research cannot be automated: in this case the Baumol’s cost disease effect is back and holds growth down. Yet, this does not account for the fact that productivity growth has actually declined over the past decade.

Aghion *et al.* (2019), henceforth ABBKL, propose the following explanation. Suppose that there are two main sources of heterogeneity across firms in the economy. The first one is “product quality” which improves as a result of innovation on each product line. But on top of product quality, some firms – call them “superstar” firms – may enjoy a persistent “efficiency advantage” over other firms. Natural sources of such an advantage are the organizational capital, the development of networks, or the ability to escape taxation: these help superstar firms to enjoy higher mark-ups than non-superstar firms with the same level of technology. The story developed by ABBKL is that a technological revolution, by reducing the firms’ cost of monitoring each individual activity, will induce all firms to expand their range of activities. However, since superstar firms enjoy higher profits on each product line than non-superstar firms with the same level of technology, the former will end up expanding at the expense of the latter. But this in turn will deter innovation by non-superstar firms, as innovating on a line where the incumbent firm is a superstar firm always yields lower profits than innovating on a line where the incumbent firm is a non-superstar firm. Thus overall, the technological revolution can result in lower aggregate innovation and lower average productivity growth in the long run, following an initial burst of growth associated

with the expansion of superstar firms into more product lines.²

This can explain why productivity growth in the US has declined continuously since 2005, after a burst of growth between 1995 and 2005, in the wake of the AI revolution following the ICT revolution. Moreover, it also accounts for the fact that, over the past decade, the average markup has markedly increased in the US, which was mostly due to a composition effect: namely, the share of higher-markup firms in the economy has gone up, but markups within firms have not shown any significant upward trend.

This explanation illustrates the fact that technological revolutions may have adverse effects on growth if institutions and policies do not adapt. Thus, ICT and AI have helped some superstar firms develop platforms/networks or social capital which in turn have acted as barriers to entry and/or innovation by non-superstar firms. The challenge is then to rethink competition policy so that the ICT and AI revolutions can fully deliver on their growth promises.

The above discussion has stressed the importance of appropriate institutions and policies for the impact of AI on growth. In the next session we analyze the impact of AI on employment, and there again we shall argue that institutions and policies matter.

Automation and Employment

Historical Background

Since AI is only in its infancy, empirical job data with hindsight are not available yet. Hence, empirical studies have focused on automation in a broad sense and on its impact on employment.

Early analyses showed an increase in technological unemployment based on a macroeconomic equilibrium analysis, but without a special focus on automation (Keynes, 1930; Leontief, 1952; Lucas & Prescott, 1974; Davis & Haltiwanger, 1992; Pissarides, 2000).

In the wake of the IT and computer revolution in the 1990s, authors tried to explain the polarization of the labor market. The canonical skill-biased technological change became a major subject of investigation: several studies explained a rising wage gap and a better return

on education by a rising demand for skilled labor versus non-skilled labor (Katz & Murphy, 1992; Krueger, 1993; Autor *et al.*, 1998; Bresnahan *et al.*, 2002; Acemoglu, 2002; Autor & Dorn, 2013).³ This skill-biased technological change hypothesis did not foresee a replacement of labor by capital, but rather supported the idea of complementarity between technology and skilled workers (see Acemoglu & Autor, 2011, for an overview).

Following the critic of Card & DiNardo (2002), and the seminal paper of Autor *et al.* (2003), the theory of skill-biased technology declined, in favor of a “routinization” hypothesis. The academic consensus shifted to a labor-replacing view of automation in routine tasks. The underlying assumption became that “traditional” automation replaces routine jobs, and creates more demand for non-routine jobs that require skills that cannot be performed by machines. Indeed, empirical facts show that automation gave rise to more high-skilled and low-skilled jobs⁴, while crowding out medium-skilled jobs (Goos & Manning, 2007). Several studies highlight the structural change in the labor market and show the disappearance of manufacturing and routine jobs (Autor *et al.*, 2003; Jaimovich & Siu, 2012; Autor & Dorn, 2013; Charnoz & Orand, 2017; Blanas *et al.*, 2019).

Some authors have tried to go beyond the scope of “traditional” automation by questioning the feasibility of automating jobs given current and presumed technological advances. They notably relax the assumption according to which automation could not threaten non-routine jobs. Whereas Autor *et al.* (2003) argued that non-routine tasks such as legal writing, truck driving, medicine, selling, could not be substituted, this view has been questioned by Brynjolfsson & McAfee (2011) who advocate that automation is no longer limited to routine tasks, recalling the example of self-driving cars. Frey & Osborne (2017) have followed this path and estimated the probability of computerization⁵ of 702 jobs. Their main conclusion showed that 47% of

2. On the slowdown of productivity growth and its link with the rise of corporate market power and firm concentration, see also Liu *et al.* (2019).

3. On the same issue, let's also mention Beaudry *et al.* (2013), who highlight the declining demand for non-skilled workers, but through a different mechanism. They argue that the over-qualification of workers induces less demand for qualified workers, who are therefore “forced” to accept underqualified jobs, while non-qualified workers are kicked out of the labor market.

4. Goos & Manning (2007) refer to them as “lovely” and “lousy” jobs.

5. Computerization is defined as job automation by means of computer-controlled equipment.

employment in the US is at risk of automation in the next ten or twenty years, whereas only 33% of jobs have a low risk of automation. They also showed that there was a strong negative relationship between, on the one hand, wages and educational attainment and, on the other hand, the probability of computerization. Frey & Osborne have been under harsh criticism: they ignore the task content of the jobs, and do not factor in the variability of a specific occupation across workplaces. 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. Last, Frey & Osborne's method does not integrate the response of the economy in a general equilibrium model, i.e. the cost of automation, the response of wages, and the creation of new jobs. Despite technological advances, the cost of substitution between machines and labor could prevent firms from automating rapidly, especially if wages adapt. Moreover, other activities could develop and hire the redundant workers.

Being forward-looking without reasoning in a general equilibrium pattern seems unrealistic. Caselli & Manning (2019) criticize the fact that most current papers rely on a partial equilibrium analysis and do not draw on a formal model of the economy as a whole. Instead, they propose a very general framework for thinking about the effects of automation on different types of workers. They notably show that new technology is unlikely to cause wages for all workers to fall and will cause average wages to rise if the prices of investment goods fall relative to consumer goods.

The analyses of automation based on the routine-based technological change share one caveat: since their premise is that automation affects routine jobs, they do not question the measure of automation. Yet, getting an accurate measure of automation is crucial, and this is what recent studies have tried to do. Earlier studies were based on the measure of computers or IT (Krueger, 1993; Autor *et al.*, 1998; Bresnahan *et al.*, 2002), recent papers investigate other measures of automation like automation related patents (Mann & Püttmann, 2017), or the number of robots (Autor & Dorn, 2013; Acemoglu & Restrepo, 2017; Dauth *et al.*, 2017; Graetz & Michaels, 2018; Cheng *et al.*, 2019). It is on this latest strand of literature that we focus in the remaining part of this paper.

The Effect of Robots on Employment in the US

As regards the impact of robots on net employment, evidence is mixed. Chiacchio *et al.* (2018) report negative effects – one more robot per thousand workers reduces the employment rate in six EU countries by 0.16-0.20 percentage points. Yet, Autor *et al.* (2015) and Graetz & Michaels (2018) find no effect of automation on aggregate employment. On German data, Dauth *et al.* (2017) find no evidence that robots cause total job losses, but they show a significant negative effect on employment in the manufacturing industry: each additional robot per thousand workers reduces the aggregate manufacturing employment-to-population ratio by 0.0595 percentage points.

In their paper “Robots and jobs: Evidence from US Labor Markets”, Acemoglu & Restrepo (2017) analyze the effect of the increase in industrial robot usage between 1990 and 2007 on US labor markets. They answer this question using within-country variation in robot adoption. The first part of the paper is dedicated to describing a theoretical model in which robots and humans are substitutes to derive equations and calculate the aggregate impact of robots on employment and wages. They show that, for each labor market, the impact of robots on jobs may be estimated by regressing the change in employment and wages on the exposure to robots and finally 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%.

In detail, AR-2017 focus on the 722 commuting zones covering the US continental territory. For each commuting zone, they gather employment and wage data, and build a measure of the exposure to robots. Then they run regressions on all commuting zones, in order to investigate the impact of this exposure on the change in employment and the change in aggregate wages, i.e. to estimate the following relationships:

$$\begin{cases} d\ln L_c = \beta_L \cdot US\ RobotsExp_c + \varepsilon_c^L \\ d\ln W_c = \beta_W \cdot US\ RobotsExp_c + \varepsilon_c^W \end{cases}$$

The best way to measure local exposure to robots would be to have a direct measure of the stock of robots in each commuting zone. Yet, no such data exist: the main source of data on robotics is provided by the International Federation of Robotics (IFR), which gathers worldwide data from robot producers, on sales,

the destination of sales and their classification by industrial sector. The main advantage of IFR data is to define a robot according to an ISO standard, which provides a homogeneous definition between industries. Indeed, a robot is defined as “an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications”. The main feature of this definition lies in the autonomy of the robot to perform tasks. From these data, they deduce the stock of robots by country and by year from 1993 on⁶, but only on a country – or a group of countries – scale. The IFR provides data on the stock of robots for 19 employment categories, i.e. 2-digit nomenclature data in the non-manufacturing sector and 3-digit data in the manufacturing sector. Graetz & Michaels (2018) were the first to use the IFR data. They estimated that the robotization process between 1993 and 2007 contributed to the annual labor productivity growth by 0.36 percentage points.

AR-2017 build a local index, which is based on the rise in the number of robots per worker in each industry on the one hand, and on the local distribution of labor between different industries on the other hand.

For each commuting zone, the index measuring the exposure to robots between 1990 and 2007 is inspired by the index measuring the exposure to Chinese imports, which has been developed by Autor, Dorn & Hanson (2013). The main idea underpinning this index is to exploit the variation in local industry employment structure before the period of interest, in order to spread a variable (robots, imports, etc.) which is only available at the national level. Autor *et al.* (2013) highlight the fact that the variation of the index stems from two sources: the share of the manufacturing employment, and the specialization in exposed industries within manufacturing. Since we want to capture the second source of variation, it is important to control for the share of manufacturing employment at the beginning of the period, as we will see in detail later. The measure used in the paper to measure the exposure to robots at the commuting zone level is:

$$RobotsExp_{c1993}^{2007} = \sum_{i \in I} l_{ci}^{1970} \left(\frac{R_{i,2007}^{US}}{L_{i,1990}^{US}} - \frac{R_{i,1993}^{US}}{L_{i,1990}^{US}} \right)$$

The sum runs over all the 19 industries i in the IFR data. l_{ci}^{1970} stands for the 1970 share of employment in industry i for a given commuting

zone i . R_i and L_i stand for the stock of robots and the number of people employed in a particular industry i .

The variation of robots exposure between commuting zones is then used in order to explain the observed evolution of employment and wages. Several controls are included in the regressions. An important feature is to take into account changes in trade patterns. Acemoglu & Restrepo therefore use data from Autor *et al.* (2013) on the exposure to Chinese imports, and construct similar measures of the exposure to imports from Mexico. This local labor market exposure to import competition from China (Mexico) is once more calculated in an analogous fashion as in Autor *et al.* (2013), i.e. the national change in import volume from China (Mexico) per worker and per sector, weighted by the sectoral composition of employment in the zone. Another feature is controlling for growth of capital stock (other than robotics) and growth of IT capital. Other controls include the share of employment in routine jobs in 1990, a measure of offshoring of intermediate inputs, baseline differences in demographics in 1990, baseline shares of employment in manufacturing, durable manufacturing and construction, as well as the share of female employment in manufacturing. AR-2017 also construct estimates of the number of robot integrators in each commuting zone.⁷ As they explain in their theoretical development, the empirical estimations are based on two patterns: a quite unrealistic pattern where commuting zones do not trade, and a more realistic pattern where trade between zones is taken into account. The underlying idea is that if an industry in a given zone adopts more robots, then it will become more productive and will export its cheaper product to its neighbors.

As underlined by Acemoglu & Restrepo themselves, a major concern with their empirical strategy is that the adoption of robots in a given US sector could be related to other trends in that sector. Therefore, they use an instrumental method and make two-stage least squares estimates. Their method is similar to the method used by Autor *et al.* (2013) on US data and Bloom *et al.* (2015) on European data in order to estimate the impact of Chinese imports.

6. Yet, for the US, the repartition of robots is not fully detailed by manufacturing industry on the 1993-2004 period. The full detail is given from 2004 on. Outside manufacturing, the number of robots is available for: agriculture, forestry and fishing, mining, utilities, construction, education, research and development and services.

7. Companies that install and program robots for different industrial applications.

In the first stage, they regress the US exposure to robots on the exogenous exposure to robots in the other advanced European countries, calculated using the same approach as on US data, with the industry-level spread of robots in other advanced economies as an instrument for the adoption of robots in US industries. In the second stage, they regress the change in employment (resp. wages) on the instrumented exposure to robots. The main result is that the commuting zones the most exposed to robots have experienced the worst evolutions in terms of employment (resp. wages) between 1990 and 2007.

Under the assumption that there is no trade taking place between commuting zones, AR-2017 estimate that each additional robot per thousand workers reduces aggregate employment to population ratio by 0.37 percentage points and aggregate hourly wages growth by about 0.73%. If they take trade between commuting zones into account, and calibrate a macroeconomic model, they find that the magnitude of the estimates decreases and that one extra robot per thousand workers reduces the aggregate employment to population ratio by 0.34 percentage points and aggregate hourly wages by 0.5%. Adding control variables such as Chinese and Mexican import volumes, the share of routine jobs and offshoring has little impact on the estimates. Among other

robustness checks, AR-2017 run IV regressions where they exclude the commuting zones with the highest exposure to robots (consequence of an important initial automotive employment in these areas). The estimates are quite similar to previous specifications, and they conclude that their results are not solely driven by highly exposed areas.


Robots and Employment in France

In this section, we reproduce the method developed by Acemoglu & Restrepo (2017) on French data over the 1994-2014 period, in order to compare the magnitude of the results obtained in France with those obtained using US data.

Figure I plots how the number of robots evolved in France from 1994 to 2014. Similarly to AR-2017, data on robots are provided by the International Federation of Robotics (IFR). The overall number of robots, pictured by the blue curve, grows steadily between 1994 and 2007, then stagnates from 2007 to 2011, and finally slightly decreases between 2012 and 2014.

In order to ensure the comparability of our results with those of AR-2017 or Dauth *et al.* (2017), we use a very close framework. We then

Figure I
Evolution of the number of robots in France (1994-2014)



Sources: International Federation of Robotics.

define the exposure to robots in a French employment zone⁸ between 1994 and 2014:

$$RobotsExp_{c1994}^{2014} = \sum_{i \in I} \frac{L_{ic,1994}}{L_{c,1994}} \left(\frac{R_{i,2014}}{L_{i,1994}} - \frac{R_{i,1994}}{L_{i,1994}} \right)$$

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 thousands) 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. Data on employment are obtained from the French administrative database DADS.


Our index therefore reflects the exposure to robots per 1,000 workers between 1994 and 2014. Figure II plots the geographical distribution of the exposure to robots. The average exposure in France is 1.16 between 1994 and 2014, well below the average exposure in Germany of 4.64 during the same period. This exposure is also more homogeneous in France, with a standard deviation of 1.42 versus 6.92 in Germany. The order of magnitude of exposure to robots in France is closer to the exposure in the United States between 1993 and 2007.

Figure II shows a fairly marked North/South divide. Indeed, while the North has high exposure rates, most southern employment zones have a rate close to 0. The Northeast, with a strong industrial heritage, but also the West (Normandy and eastern Brittany) are among the highly exposed regions. In the least exposed regions, one finds the entire Atlantic coast and the French Riviera.

In order to measure the impact of exposure to robots on local labor markets, we adopt a strategy similar to the one initiated by Autor *et al.* (2013) to investigate the impact of Chinese imports on local labor markets in the United States. Our variable of interest is the evolution of the employment-to-population ratio between 1990 and 2014. In the first and most naïve specification, we study the impact of exposure to robots on the evolution of employment-to-population ratio. This ratio is constructed from census data. However, it is important to control for other characteristics that may influence the evolution of the employment-to-population ratio. To do

8. According to the official definition provided by 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.

Figure II
Exposure to robots in France (1994-2014)



Sources: IFR; authors' calculations.

so, we construct two other exposure indices. First, an exposure index for information and communication technologies (ICT) $ICTExp$, built in a similar way as the exposure to robots index. The number of robots is replaced by the ICT capital stock in industry i . Data come from the EUKLEMS database. Second, we build an international trade exposure index $TradeExp$ using the COMTRADE database. The number of robots is replaced by net imports from China and Eastern Europe (Bulgaria, Czech Republic, Estonia, Croatia, Hungary, Lithuania, Latvia, Poland, Romania, Slovakia, Slovenia) in industry i . In some regressions, we also add a vector X_c of control for the employment zone c : demographic characteristics in 1990 (population share by level of education and population share between 25 and 64 years old), broad industry shares in 1994 and broad region dummies. Finally, we can write:

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

Even if these control variables partially purge OLS estimations, an instrumental variable approach is necessary to discuss causal impact of robots on employment. In fact, one may imagine a shock, which we do not capture in our controls, but which may impact both the installation of robots at the local level and local labor markets characteristics. Always in a comparability perspective, we adopt the

approach of AR-2017 and Dauth *et al.* (2017), according to which the stocks of robots in industries from n developed countries are used to build n indexes of exposure to robots. These n indexes are built with employment data from 1978, to avoid concerns about reverse causality: those pre-existing levels cannot be impacted by robot installations. We select the following countries: Germany, Denmark, Spain, Italy, Finland, Norway, Sweden and the United Kingdom. Data from North America (US and Canada) are not considered, because we only have information on the total number of robots before 2004, without any industry breakdown.

All the share variables (employment-to-population ratio, population share by level of education, etc.) are considered in percentage points in the following regression. Table 1 displays the results of the OLS regressions.

This table shows a negative correlation between exposure to robots and change in employment-to-population ratio. However, the correlation becomes non-significant in column (6) when we include all the controls and in column (7) when we exclude the commuting zones with the highest exposure to robots. In the first five columns where the correlation is significant, the magnitude of the effect ranges from -1.090 to -0.515.

In the instrumental variable regression shown in Table 2, the coefficients of robots exposure are

Table 1
The effect of robots exposure on employment, 1990-2014, OLS estimates

Dependent variable: Change in employment-to-population ratio 1990-2014 (in percentage points)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$RobotsExp_{1994}^{2014}$	-1.090*** (0.253)	-0.749*** (0.263)	-0.594** (0.239)	-0.515** (0.243)	-0.549* (0.294)	-0.398 (0.244)	-0.430 (0.324)
$ITCExp_{1994}^{2014}$		-3.099* (1.586)	-2.397 (1.594)	-2.495* (1.455)	-0.304 (1.620)	-0.165 (1.576)	-0.154 (1.588)
$TradeExp_{1994}^{2014}$		-0.743*** (0.247)	-0.690*** (0.215)	-0.825*** (0.239)	0.0857 (0.243)	-0.123 (0.278)	-0.124 (0.280)
Demographics			Yes			Yes	Yes
Region dummies				Yes		Yes	Yes
Broad industry shares					Yes	Yes	Yes
Remove Highly exposed areas							Yes
Observations	297	297	297	297	297	297	295
R-squared	0.058	0.090	0.198	0.205	0.249	0.407	0.406

Notes: 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 manufacturing, agriculture, construction, retail 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.

significant whatever the specification chosen, even the one with all the controls. The magnitude of first-stage F-statistics indicates that the weak instrument bias is unlikely to be a problem here. Moreover, we observe that the magnitude of the effects increases in comparison with those obtained by OLS. Column (1) begins with the regression without any control. The negative impact of exposure to robots on employment is massive: one more robot per 1,000 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. If the ICT exposure coefficient is not statistically significant, there is a negative impact of net imports on employment-to-population ratio, as in Autor *et al.* (2013) for the United States. The coefficient for exposure to robots remains of the same order of magnitude. Other controls are successively included in columns (3) to (5): demographic characteristics in column (3), broad region dummies in column (4) and broad industry share before 1994 in column (5). In each specification, the coefficient of exposure to robots remains negative and significant, even if its magnitude decreases slightly. On the contrary, the coefficient of exposure to imports becomes insignificant when we add information about the industry composition of the employment zones. Finally, column (6) combines all the controls and column (7) removes highly exposed areas. The effect of the exposure to robots is

still negative and significant, even though its magnitude has been reduced in comparison with the specification without any control.

In our last specification, we obtain a negative effect of exposure to robots on employment: one more robot per 1,000 workers leads to a drop in the employment-to-population ratio of 0.686 percentage points. A quick calculation allows us to conclude that the installation of one more robot in a commuting zone reduced employment by 10.7 jobs.⁹ The order of magnitude is similar to AR-2017, who found an impact of 6.2 fewer jobs for one more robot. According to the IFR, the number of robots in France increased by around 20,000 between 1994 and 2014. Our result implies a loss of 214,000 jobs ($10.7 \times 20,000$) during this period due to robots.

Results focusing on the 1990-2007 period, like AR-2017, are presented in Table A1 in the Appendix. Using specification with all controls, we conclude that 1 more robot per 1,000 workers led to a drop in the employment-to-population ratio of 0.438 percentage points. This estimation is even closer to that of AR, who estimated a drop of 0.371 percentage points.

9. Our exposure to robots is defined in "robots for 1,000 workers". According to the OECD, the average employment-to-population ratio was 0.64 in 2014. Hence, the installation of one more robot reduced employment by $(0.686/100) \times 1000 / 0.64 = 10.7$ jobs.

Table 2
The effect of robots exposure on employment, 1990-2014, IV estimates

Dependent variable: Change in employment-to-population ratio 1990-2014 (in percentage points)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>RobotsExp</i> ₁₉₉₄ ²⁰¹⁴	-1.317*** (0.325)	-1.010*** (0.322)	-0.974*** (0.271)	-0.737** (0.296)	-0.790*** (0.300)	-0.686*** (0.241)	-0.986*** (0.351)
<i>ITCExp</i> ₁₉₉₄ ²⁰¹⁴		-2.569 (1.618)	-1.699 (1.578)	-2.094 (1.444)	-0.176 (1.590)	-0.0323 (1.518)	0.101 (1.538)
<i>TradeExp</i> ₁₉₉₄ ²⁰¹⁴		-0.670*** (0.242)	-0.589*** (0.211)	-0.773*** (0.230)	0.110 (0.240)	-0.0922 (0.276)	-0.0882 (0.279)
Demographics			Yes			Yes	Yes
Region dummies				Yes		Yes	Yes
Broad industry shares					Yes	Yes	Yes
Remove Highly exposed areas							Yes
Observations	297	297	297	297	297	297	295
First-stage F statistic	53.7	29.4	24.0	25.7	25.1	23.6	46.5
R-squared	0.055	0.087	0.193	0.203	0.248	0.405	0.400

Notes: 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 manufacturing, agriculture, construction, retail 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.

Finally, we investigate the possibility of heterogeneous employment effects of the exposure to robots across education levels. Since we only have this kind of information for individuals between 25 and 54 years old, we restrict our analysis to this population. The results are similar to those presented in Tables 1 and 2. Coefficients estimating the exposure to robots by education level are presented in Figure III (with confidence intervals of 90%). The Certificate of Professional Aptitude (CAP) and the Diploma of Occupational Studies (BEP) are both French professional education degrees. The lower the level of education, the greater the negative impact of exposure to robots. The impact is non-significant for people with high-school diploma. The effect is even positive, but not significant for college graduates. This heterogeneity emphasizes the key role played by education and the need for public policies. In order to limit the negative effects of technical progress on employment, public policies should notably aim at rising the education level and at promoting continuous training.

Discussion

The analyses above raise several potential issues. First, are robots so different from other sources of automation? The IFR definition of robots is quite restrictive and does not include machines like automatic tellers, which replace


human labor as well as robots. Taking a broader measure of technological progress into account would make it possible to use data over a longer period, i.e. to use more evidence from the past.

Another potential concern is that the analysis relies on the hypothesis that the number of robots installed by a given industry, divided by the importance of the industry in the commuting zone, is the same across commuting zones. Yet, robotization by a given industry may be more intense in commuting zone A than in commuting zone B even if the shares of that industry are the same in both regions.

A third potential concern is that variations in the robots exposure index across commuting zones are mostly related to the spatial distribution of automotive activities over the US territory in 1990, since industrial robots are predominant in the automotive industry – automotive robots account for more than one third of total robots. Using this variation to explain employment boils down to asking whether the importance of the automotive industry in a given commuting zone in 1990 can explain the evolution of employment in that zone over the twenty following years.

Indeed, most of the robotization took place in the automobile industry, and in the 1990s and the 2000s, the American automobile market

Figure III
The effects of robots exposure by education level



Notes: Confidence intervals of 90%
Sources: Authors' calculations.

experienced large-scale restructuring. Since the late 1980s, numerous automotive assembly facilities have relocated in the South of the United States, a region which previously had a small automotive presence. This has caused concern among the traditional automotive communities in upper Midwest and southern Ontario (Hill & Brahmst, 2003). Similar outsourcing trends occurred in the French automobile industry. If many closures took place in a commuting zone strongly specialized in automobile in the 1990s, then the negative relationship between exposure to robots and employment might just reflect the relocation of automobile plants. AR-2017 deal with that concern in two different ways. First, they exclude commuting zones with very high exposure to robotization, which presumably are also locations with high initial employment in the automotive industry. Doing so does not affect the basic regression results.¹⁰ Second, they run OLS regressions where exposure to robots in automotive manufacture and exposure to robots in other industries are treated as separate regressors: they find that coefficients on the two regressors are quite similar. Overall, Acemoglu & Restrepo (2017) conclude that there is no concern with the predominance of industrial robots in the automotive industry.

* *
*

In this paper, we have surveyed recent work on artificial intelligence and its effects of economic growth and employment. Our main conclusion is that the effects of AI and automation on growth and employment depend to a large extent on institutions and policies. Yet, despite solid theoretical foundations on how

to model automation and AI, and despite some compelling empirical work, we are still at an early stage in fully understanding all the welfare implications of these technologies.

In the first section, we have argued that while AI can spur growth by replacing labor, which is in finite supply, by capital which is in unbounded supply, it may inhibit growth if combined with inappropriate competition policy.

In the second section, we discussed the effects of AI and automation on aggregate employment: building on Acemoglu & Restrepo (2017), we have looked at the effect of robotization on employment in France over the 1994-2014 period. We find that robotization reduces aggregate employment at the employment zone level. We also find that non-educated workers are more negatively affected by automation than educated workers. This suggests that inappropriate labor market and education policies could reduce the positive impact of AI and automation on employment.

A natural next step would be to bridge the analysis in the two sections. We are currently working to this. Another avenue of research is to investigate how labor market characteristics can affect the nature of innovation, for example, whether the innovation is aimed at automation or the creation of new product lines. The former idea is explored in current work by the present authors, and the latter is explored in recent work by Dechezleprêtre *et al.* (2019). Other extensions, in particular on the effects of AI on consumption and well-being, await future research. □

¹⁰ Similar results are shown for France, cf. Table 2, column (7).

BIBLIOGRAPHY

- Acemoglu, D. (2002).** Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40(1), 7–72.
<https://www.jstor.org/stable/2698593>
- Acemoglu, D. & Autor, D. (2011).** Skills, Tasks and Technologies: Implications for Employment and Earnings. In: Ashenfelter, O. & Card, D. (Eds.), *Handbook of Labor Economics*, Vol. 4, chapter 12, pp. 1043–1171. Amsterdam: Elsevier.
- Acemoglu, D. & Restrepo, P. (2016).** The Race between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment. NBER *Working Paper* N° 22252.
<https://doi.org/10.3386/w22252>
- Acemoglu, D. & Restrepo, P. (2017).** Robots and Jobs: Evidence from US Labor Markets. NBER *Working Paper* N° 23285.
<https://doi.org/10.3386/w23285>
- Aghion, P., Jones, B. & Jones, C. (2017).** Artificial Intelligence and Economic Growth. NBER *Working Paper* N° 23928.
<https://doi.org/10.3386/w23928>
- Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. & Li, H. (2019).** A Theory of Falling Growth and Rising Rents. *Mimeo College de France*.
https://www.college-de-france.fr/media/centre-economie-innovation/UPL1638608134824677946_a_theory_of_falling_growth_and_rising_rents.pdf
- 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.
<https://doi.org/10.1162/003355398555874>
- 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.
<https://doi.org/10.1257/aer.103.5.1553>
- 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.
<https://doi.org/10.1162/003355303322552801>
- 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.
<https://doi.org/10.1111/eoj.12245>
- 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.
<https://doi.org/10.1257/aer.103.6.2121>
- Arntz, M., Gregory, T. & Ziehrhahn, U. (2017).** Revisiting the risk of automation. *Economics Letters*, 159, 157–160.
<https://doi.org/10.1016/j.econlet.2017.07.001>
- Baumol, W. (1967).** Macroeconomics of unbalanced Growth: The Anatomy of Urban Crisis. *American Economic Review*, 57(3), 415–426.
<https://www.jstor.org/stable/i331434>
- Beaudry, P., Green, D. A. & Sand, B. M. (2013).** The great reversal in the demand for skill and cognitive tasks. NBER *Working Paper* N° 18901.
<https://doi.org/10.3386/w18901>
- Blanas, S., Gancia, G. & Lee, S. Y. (2019).** Who is Afraid of Machines? CEPR *Discussion Paper* N° 13802.
<https://cepr.org/content/free-dp-download-20-june-2019-who-afraid-machines>
- Bloom, N., Draca, M. & Van Reenen, J. (2016).** Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity. *The Review of Economic Studies*, 83(1), 87–117.
<https://doi.org/10.1093/restud/rdv039>
- 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.
<https://doi.org/10.1162/003355302753399526>
- 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. Lexington, MA: Digital Frontier Press.
- 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.
- Caselli, F. & Manning, A. (2019).** Robot Arithmetic: New Technology and Wages. *American Economic Review: Insights*, 1(1), 1–12.
<https://doi.org/10.1257/aeri.20170036>

- Charnoz, P. & Orand, M. (2017).** Technical change and automation of routine tasks: Evidence from local labour markets in France, 1999-2011. *Économie et Statistique / Economics and Statistics*, 497-498, 103–122.
<https://www.insee.fr/en/statistiques/3318007?sommaire=3318021>
- Cheng, H., Jia, R., Li, D., & Li, H. (2019).** The Rise of Robots in China. *Journal of Economic Perspectives*, 33(2), 71–88.
<https://doi.org/10.1257/jep.33.2.71>
- 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* N° 2.
<https://bruegel.org/2018/04/the-impact-of-industrial-robots-on-eu-employment-and-wages-a-local-labour-market-approach/>
- Dauth, W., Findeisen, S., Südekum, J. & Wößner, N. (2017).** German robots: The impact of industrial robots on workers. *IAB Discussion Paper* N° 30/2017.
- Davis, S. J. & Haltiwanger, J. (1992).** Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics*, 107(3), 819–863.
<https://doi.org/10.2307/2118365>
- Dechezleprêtre, A., Hémons, D., Olsen, M. & Zanella, C. (2019).** Automating Labor: Evidence from Firm-level Patent Data. Mimeo.
http://morten-olsen.com/DHOZ_Automating_Labor.pdf
- Frey, C. B. & Osborne, M. A. (2017).** The future of employment: How susceptible are jobs to computerization? *Technological Forecasting & Social Change*, 114, 254–280.
<https://doi.org/10.1016/j.techfore.2016.08.019>
- Goos, M. & Manning, A. (2007).** Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *Review of Economics and Statistics*, 89(1), 118–133.
<https://www.jstor.org/stable/40043079>
- Graetz, G. & Michaels, G. (2018).** Robots at work. *Review of Economics and Statistics*, 100(5), 753–767.
https://doi.org/10.1162/rest_a_00754
- Hémons, D. & Olsen, M. (2014).** The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality. *CEPR Discussion Paper* N° 10244.
<https://econpapers.repec.org/RePEc:cpr:ceprdp:10244>
- Hill, K. & Brahmst, E. (2003).** The Auto Industry Moving South: An Examination of Trends. *Center for Automotive Research*.
<https://www.cargroup.org/publication/the-auto-industry-moving-south-an-examination-of-trends/>
- Jaimovich, N. & Siu, H. E. (2012).** Job Polarization and Jobless Recoveries. *NBER Working Paper* N° 18334.
<https://doi.org/10.3386/w18334>
- Jones, C. (1995).** R & D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4), 759–784.
<https://www.jstor.org/stable/2138581>
- Katz, L. & Murphy, K. (1992).** Changes in Relative Wages: Supply and Demand Factors. *Quarterly Journal of Economics*, 107(1), 35–78.
<https://doi.org/10.2307/2118323>
- Keynes, J. M. (1932).** Economic possibilities for our grandchildren. In: Keynes, J. M., *Essays in Persuasion*, pp. 358–373. New York: Harcourt Brace.
- Krueger, A. B. (1993).** How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989. *Quarterly Journal of Economics*, 108(1), 33-60.
<https://doi.org/10.2307/2118494>
- Leontief, W. (1952).** Machines and Man. *Scientific American*, 187(3), 150–64.
<https://www.jstor.org/stable/e24950774>
- Liu, E., Mian, A. & Sufi, A. (2019).** Low Interest Rates, Market Power, and Productivity Growth. *NBER Working Paper* N° 25505.
<https://doi.org/10.3386/w25505>
- Lucas, R. E. & Prescott, E. C. (1974).** Equilibrium search and unemployment. *Journal of Economic Theory*, 7(2), 188–209.
[https://doi.org/10.1016/0022-0531\(74\)90106-9](https://doi.org/10.1016/0022-0531(74)90106-9)
- Mann, K. & Püttmann, L. (2017).** Benign Effects of Automation: New Evidence from Patent Texts. *Unpublished manuscript*.
- Pissarides, C. A. (2000).** *Equilibrium Unemployment Theory*. Cambridge, MA: MIT Press.
- Zeira, J. (1998).** Workers, Machines, and Economic Growth. *Quarterly Journal of Economics*, 113, 1091–1117.
<https://doi.org/10.1162/003355398555847>

Tableau A1
The effect of robots exposure on employment, 1990-2007, IV estimates

Dependent variable: Change in employment-to-population ratio 1990-2007 (in percentage points)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>RobotsExp</i> ₁₉₉₄ ²⁰⁰⁷	-0.382*** (0.119)	-0.344* (0.198)	-0.508*** (0.195)	-0.148 (0.197)	-0.560** (0.217)	-0.438** (0.198)	-0.633** (0.298)
<i>ITCExp</i> ₁₉₉₄ ²⁰⁰⁷		-0.322 (1.613)	0.990 (1.611)	-1.274 (1.571)	2.844 (2.142)	1.845 (2.019)	2.184 (2.056)
<i>TradeExp</i> ₁₉₉₄ ²⁰⁰⁷		-0.217 (0.319)	-0.285 (0.293)	-0.400 (0.324)	0.301 (0.347)	0.107 (0.383)	0.111 (0.391)
Demographics			Yes			Yes	Yes
Region dummies				Yes		Yes	Yes
Broad industry shares					Yes	Yes	Yes
Remove Highly exposed areas							Yes
Observations	297	297	297	297	297	297	295
First-stage F statistic	45.7	24.8	29.3	24.7	22.6	25.7	44.4
R-squared	0.004	0.007	0.075	0.129	0.144	0.293	0.284

Notes: 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 manufacturing, agriculture, construction, retail 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.