

What Are the Labor and Product Market Effects of Automation?

New Evidence from France*

Philippe Aghion, *Collège de France*
Céline Antonin, *Sciences-Po*
Simon Bunel, *Paris School of Economics*
Xavier Jaravel, *London School of Economics*

January 2020

Abstract

We use comprehensive micro data in the French manufacturing sector between 1994 and 2015 to document the effects of automation technologies on employment, wages, prices and profits. Causal effects are estimated with event studies and a shift-share IV design leveraging pre-determined supply linkages and productivity shocks across foreign suppliers of industrial equipment. At all levels of analysis — plant, firm, and industry — the estimated impact of automation on employment is positive, even for unskilled industrial workers, which suggests that the productivity effects of automation outweigh its potential displacement effects. We also find that automation leads to higher profits, lower consumer prices, and higher sales. The estimated elasticity of employment to automation is 0.28, compared with elasticities of 0.78 for profits, -0.05 for prices, and 0.37 for sales. Consistent with the importance of business-stealing across countries, the industry-level employment response to automation is positive and significant only in industries that face international competition. These estimates can be accounted for in a simple monopolistic competition model: firms that automate more increase their profits but pass through some of the productivity gains to consumers, inducing higher scale and higher employment. The results indicate that automation can increase labor demand and can generate productivity gains that are broadly shared across workers, consumers and firm owners. In a globalized world, attempts to curb domestic automation in order to protect domestic employment may be self-defeating due to foreign competition.

*For thoughtful comments we thank David Autor, Francesco Caselli, Matt Gentzkow, Larry Katz, Pete Klenow, Alan Manning, Philippe Martin, Guy Michaels, Ben Moll, Enrico Moretti, David Sraer, Michael Webb, Heidi Williams, as well as seminar participants at the Bank of Italy, Berkeley, the College de France, the London School of Economics, the Hebrew University of Jerusalem, Tel Aviv University, and the University of Copenhagen.

I Introduction

What are the effects of automation in the labor and product markets? A host of factors may be at play. Automating the production process may displace certain workers, raising the possibility of technological unemployment (e.g., Keynes (1930), Leontief (1952), Brynjolfsson and McAfee (2014)), but these displacement effects could potentially be offset by a productivity effect (e.g., Bowen and Mangum (1966), Zeira (1998), Autor (2015), Acemoglu and Restrepo (2018)). Automation may induce productivity gains, increase market demand and the scale of production, and in turn increase labor demand. Depending on the extent to which productivity gains are passed through to consumers by producers, consumers could benefit from lower prices or producers could retain higher profits (e.g., Caselli and Manning (2019) and Moll et al. (2019)). Finally, because of business stealing effects from firms that automate and displace their competitors, the industry-level employment, price and profit effects of automation may differ from their firm-level or plant-level impacts.

Because of these multiple and countervailing economic forces, understanding the aggregate and distributional impacts of automation across workers, consumers and producers is fundamentally an empirical question. To design appropriate policy responses, the relative magnitudes of these mechanisms must be estimated in a unified framework.¹ Despite extensive research, the employment effects of automation remain debated, little is known about the impact of automation on consumer prices and profits, and most of the existing evidence is at the industry level rather than at the firm or plant levels, obscuring the channels at play.² Data limitations explain the relative scarcity of evidence on these questions, which can only be answered with comprehensive data on automation and the labor and product markets.

In this paper, we leverage new micro data on the population of firms and plants in the French manufacturing sector to provide a unified analysis of the effects of automation technologies on employment, wages, prices and profits between 1994 and 2015. We use two complementary measures of automation technologies, based on the balance sheet values of industrial equipment as well as plant-level records of the usage of electro-motive force.³

¹A range of policies has been discussed in this respect, including retraining programs, redistribution policies, as well as direct taxation of specific automation technologies like robots. Models of optimal technology regulation are provided by Guerreiro et al. (2017) and Costinot and Werning (2018).

²For example, Chiacchio et al. (2018), Webb (2019) and Acemoglu and Restrepo (2019) find evidence in line with the view that various forms of automation reduce labor demand, while Dauth et al. (2018), Graetz and Michaels (2018) and Mann and Puttmann (2018) document positive employment effects. We discuss the emerging literature on the firm-level effects of robotization at the end of this section.

³The balance sheet measure guarantees broad coverage, but it may include certain machines that are not always

Our analysis proceeds in three steps. we first present descriptive evidence on the population of firms and plants; we then estimate causal effects using a shift-share research design that can be applied to the subset of firms importing industrial equipment from abroad; finally, we study the relationship between automation, employment, prices and profits at the level of industries.

In the first part of the paper, we provide descriptive evidence on the population of firms and plants, using event studies exploiting the timing of adoption of industrial equipment across plants (in the same firm) or across firms (in the same industry). In contrast with the common view that firms that use more automation technologies reduce their labor force, we find that firm-level and plant-level employment increases after automation, including for low-skill industrial workers. The elasticity of employment to automation is +0.2 on impact. The response of employment is then amplified over time, with an elasticity of +0.4 after ten year.

A causal interpretation of these patterns would suggest that the productivity effect may outweigh the displacement effect, resulting in a net increase in firm-level or plant-level labor demand. Moreover, the results also show that the distributional effects of automation in the labor market are subtle. They may occur within each skill group, depending on the set of tasks performed across detailed occupations, but there is not evidence that automation has a broad effect on inequality.

However, potential unobserved shocks may confound the observed relationships. The event studies show no sign of pre trends, which is reassuring and restricts the potential set of confounders that could explain the increase in employment. Counfounding shocks would need to occur simultaneously to the increase in automation. Nonetheless, absent a quasi-experiment potential concerns over omitted factors cannot fully be addressed. For example, demand shocks or competition shocks could be at play. Increased demand or increased competition have a direct impact on employment but may also lead a firm to invest more heavily in automation technologies, exactly when the unobserved shock occurs.

To address these concerns, in the second part of the paper we validate the causal interpretation of the event study results by developing a shift-share IV design. We apply this research design to the subset of firms which import industrial equipment from abroad. Identification stems from

viewed as automation technologies. For example, powered industrial trucks used to lift and move materials over short distances (forklifts) may not be considered an automated technology because they cannot operate independently of a worker and are not tied to a pre-specified set of tasks. Our second measure is based on the observation that automation technologies in the manufacturing sector are typically based on electro-motive force, i.e. the machines used in the production process are set in motion using electric motors. For example, conveyors in the food industry, robotic arms in the automobile industry, or autosamplers in the chemical industry all fall under this definition. The motive power measure only takes into account electric motors that are constantly plugged-in when the production process is ongoing, therefore it excludes machines powered by electric batteries such as an electric forklift or electric car. See Section 2 for a complete discussion.

changes in the productivity of foreign suppliers of industrial equipment over time, which French firms are differentially exposed to through pre-determined importer-supplied relationships. This identification strategy approximate an ideal experiment that would randomly assign the prices of automation technologies across firms. Because changes in machines' quality-adjusted prices are not directly observed, it is convenient to use changes in the market shares of international suppliers over time to infer productivity shocks.

The exclusion restriction underlying this design is that firms linked to increasingly productive suppliers should not be unobservably different from other firms. To test this assumption, we run falsification tests using the lagged outcome variable. Across a range of specifications, we can never reject that there is no relationship.

The results with the shift-share design are closely in line with the event study results. Firms whose international suppliers of machines become more productive increase their usage of automation technologies, and in turn their sales and their labor force. The baseline specification with 2-digit industry by year fixed effects yields an elasticity of firm employment to automation of +0.341 (s.e. 0.121). The point estimates remain comparable in magnitudes with alternative sets of controls. We find that sales increase substantially in response to increased automation, with elasticities ranging from 0.349 to 0.561 across specifications. The relationship is significant at the 5% level in all specifications. In addition, we cannot reject that there is no impact of automation on the labor share. The shift-share instrument is strong in all specifications, as indicated by the first stage F statistic.

These findings are consistent with the role of the productivity effect of automation. Increased automation allows the firm to expand its sales and scale, which requires hiring additional workers for production. However, the firm-level relationships may paint a misleading picture because of business stealing effects across firms may affect the industry-level impacts of automation.

In the third part of the paper, we repeat the analysis at the industry level to account for business stealing and other equilibrium effects. We find that the industry-level relationship between employment and automation is positive on average, but that there is substantial heterogeneity depending on exposure to international trade. While the relationship is positive and significant in sectors that face international competition, there is no significant effect in sectors with low exposure to international competition.

Next, we document that the productivity gains from higher automation benefit both consumers through lower prices as well as firm owners via increased profits. In our baseline specification, at

the industry level a 1% increase in automation leads to a 0.05% (s.e. 0.017) fall in the industry producer price index. The point estimate remain stable with alternative sets of controls. The industry profit elasticity is positive and large, at +0.781 compared with +0.374, compared with elasticities of +0.229 for industry payroll, +0.508 for industry value added and +0.374 for industry sales.

Finally, we show that the estimated industry-level elasticities of sales, employment and prices to automation can be rationalized in a simple monopolistic competition model where consumers reallocate demand toward domestic firms with increased productivity and lower prices. An industry-level consumer demand elasticity of about 8.5 is required to account for the observed responses, which is relatively high but consistent with estimates of substitution elasticities between varieties produced by different countries for the same industry (e.g., Broda and Weinstein (2006)).

In contrast, it would be difficult to rationalize the industry-level results on sales and employment in a closed economy, because industry-level substitution would need to operate between industries (rather than between products produced either domestic firms or by international competitors within the same industry) and would require large price changes which we do not observe in the data. Competition with international suppliers providing close substitutes can explain why the relationship between automation and employment can remain positive even at the industry level, because the response of consumer demand can be large.

This paper builds on and contributes to several strands of literature. A large literature provides estimates of industry-level relationship between employment and various forms of automation, where signs and magnitudes vary across studies (e.g., Chiacchio et al. (2018), Dauth et al. (2018), Graetz and Michaels (2018), Mann and Puttmann (2018), Acemoglu and Restrepo (2019), Webb (2019)). A more recent line of work, parallel to ours, estimates the firm-level employment effects of robotization and documents positive effects (e.g., Acemoglu et al. (2019), Bessen et al. (2019), Chandler and Webb (2019), Dixon et al. (2019), Humlum (2019), Koch et al. (2019)).

We contribute to this literature in three ways. First, we introduce a quasi-experimental shift-share design to provide causal estimates of the effects of automation. In contrast, existing firm-level event study approaches cannot rule out potential unobserved confounding shocks. Second, we extend our analysis to product market outcomes, including consumer prices and firm profits, while the existing literature has focused on labor market impacts. Third, we study industry-level, firm-level and plant-level responses in a unified setting, which helps isolate the relevant mechanisms. Fourth, we examine heterogeneity in the effects depending on the types of automation technologies

(e.g., robots or other forms of automation) or market structure (e.g., exposure to international competition). The large heterogeneity we uncover depending on trade exposure may help reconcile some of the diverging industry-level estimates in prior work.⁴

Furthermore, our estimates can be used by a growing literature that uses quantitative models to assess the macroeconomic impacts of automation on inequality (e.g., Moll et al. (2019)) or to prescribe optimal technological regulations (e.g., Costinot and Werning (2018) and Guerreiro et al. (2017)). Our results provide a set of identified moments at various levels of aggregation (industry, firm and plant) for a large set of automation technologies, which quantitative models can target.

The remainder of the paper is organized as follows. Section 2 describes the data, variables and summary statistics. Section 3 reports the descriptive events from stylized facts and event studies. Section 4 reports the causal estimates from the shift-share design. Finally, Section 5 implements the industry-level analyses. The Online Appendix reports additional results.

II Data, Variable Descriptions and Summary Statistics

In this section, we describe the data sources, define the sample and key variables used in the analysis, and present summary statistics.

II.A Data Sources

To obtain a comprehensive picture of the relationship between automation, employment and firm dynamics, we combine several measures of automation to a matched employer-employee dataset. We then supplement this linked dataset with additional information on trade, prices, and consumption patterns.

Matched employer-employee data set. Detailed information on workers and firms stem from French administrative data, i.e. the DADS and INSEE databases. These databases cover the universe of plants and firms in the manufacturing sector in France from 1994 to 2015. For each firm, we observe total sales, balance sheet records, and detailed industry codes. At the plant level, we observe the composition of the workforce, notably the number of hours worked, total compensation and a detailed occupation code for each worker.⁵

⁴For example, Dauth et al. (2018) find a positive relationship between robotization and employment in Germany, a country which relies heavily on exports. In contrast, Acemoglu and Restrepo (2019) report a negative relationship in the United States, where domestic firms have a larger domestic market and are less exposed to international competition (i.e., business stealing effects operate primarily between domestic firms rather than internationally).

⁵Measures of worker skills are obtained from Charnoz and Orand (2017).

Automation measures. Automation technologies correspond to a subset of capital used in production. We use two complementary measures as proxies for automation, at the firm level and plant level.

Our firm-level proxy for the use of automation technologies leverages detailed balance sheet information available for the universe of French firms. For each firm, we observe the balance sheet value of “industrial equipment and machines” in euros. This subset of capital accounts for a large share (59 %) of total capital in manufacturing, more than the three other categories, namely “land” (2 %), “building” (23 %) and “others” (16 %).

The balance sheet measure of investments in automation has the advantage of being broad and available for all firms, but it may include certain machines that are not always viewed as automation technologies. For example, powered industrial trucks used to lift and move materials over short distances (forklifts) may not be considered an automated technology because they cannot operate independently of a worker and are not tied to a pre-specified set of tasks.

Our second measure of automation technologies is motivated by the Encyclopaedia Britannica (2015), which defines automation as “the class of electro-mechanical devices that are relatively self-operating after they have been set in motion on the basis of predetermined instructions or procedures.” In manufacturing, common automation technologies are typically based on electro-motive force, i.e. the machines used in the production process are set in motion using electric motors. For example, conveyors in the food industry, robotic arms in the automobile industry, or autosamplers in the chemical industry all fall under this definition.

Bearing this motivation in mind, we build another proxy for automation using plant-level records of electricity consumption for motors directly used in the production chain. These records have been assembled by the statistical institute INSEE since 1983. The records distinguish between different uses of electricity: motive power, thermic/thermodynamic, and other uses such as electrolysis. We focus on the motive power measure, which exclude electricity used for heating, for cooling as well as for servers (because servers are not considered to enter directly the production chain). The motive power measure only takes into account electric motors that are constantly plugged-in when the production process is ongoing, therefore it excludes machines powered by electric batteries such as an electric forklift or electric car. The measure is expressed in tonnes of oil equivalent (toe), a common energy metric.

In comparison with the firm-level balance sheet measure, the plant-level motive power measure has the advantage of isolating a more specific set of automation technologies. For example, forklifts

would be included in the firm-level balance sheet measure but excluded from the plant-level motive power, while both measures would include conveyors, robots and autosamplers.

Both measures suffer from the drawback that it is difficult to assess the “efficiency” of an automated technology, i.e. the extent to which it successfully automates the production process. For example, machines may be more expensive or require more motive power in a given industry while still being less efficient than in another industry. To address the potential drawback that balance sheet values or energy consumption may fail to reflect the effective degree of automation across industries, we leverage the panel dimension of the data and conduct our analyses in changes. As discussed in greater detail in Section 3, we use panel data to describe how employment or other outcomes change after a firm or plant increases its investments in machines (in euros) or its reliance on electric power for motive force (in toe), including time and industry fixed effects to control for time-invariant and industry-invariant heterogeneity in automation efficiency.

Because of variation in energy efficiency over time, there could be a non-monotonic relationship between our motive power measure and true automation. By investing in new automated technologies that are more energy efficient, a firm may increase its effective reliance on automation while at the same time decreasing its energy consumption for motive power. Although possible in principle, we find that this case is not relevant in practice: when examining the empirical relationship between the firm-level balance sheet and the motive power measures, we find that firms that increase their investments in industrial equipment also experience an increase in their electric energy use for motive power.

Another potential concern about the proxy based on motive power is that electricity is a variable input. Rather than proxying for investments in automated technologies, changes in motive power could simply correspond to a change in the utilization rate of machines (for example, because of changes in demand that require to adjust variable inputs). To mitigate this concern, instead of relying on the actual electricity consumption for motive power we use a plant’s peak capacity for electric motive power, which is provided by INSEE in the same survey. After major investments in machines, the plant may be required to increase its peak capacity for motive power, while no such change is required when the plant simply varies its factor utilization rate.

Trade. The trade dataset is available from customs records and covers the population of French firms in manufacturing, keeping track of all imports and exports for all firms. We use the trade data to build the shift-share instrument used in Section 4, as well as to isolate the role of robots, focusing on the subset of French firms that import robots. The trade data also provide export

prices (measured as unit values), which we use to measure the productivity effects of automation.

Prices and expenditures. For all detailed industries in our sample we obtain producer price indices from INSEE, which we use to characterize the industry-level impact of automation on productivity. We match these data to consumption spending patterns by income groups, also from INSEE, to describe the distributional effects of automation via the expenditure channel. Using these datasets, we can describe the extent to which the benefits from automation accrue to firm owners via increased profits or to consumers via lower (quality-adjusted) prices.

Robots. To compare whether the patterns differ when using robots as the sole measure of automation, instead of our two broader automation proxies, we use the firm-level trade data described above as well as the industry-level data from the International Federation of Robots (as in Graetz and Michaels (2018) and Acemoglu and Restrepo (2019)), which is available at the level of thirteen sectors.

II.B Summary Statistics

Table 1 and Figure 1 report the main summary statistics.

Table 1 reports the distribution of our main outcome variables, sales and employment, and of our automation proxies, motive power and the value of industrial equipment, at various levels of aggregation - plant level, firm level and industry level. Panel A describes the cross-section while Panel B reports the patterns in changes over the course of our sample. The sample covers 19,448 plants within 16,227 firms operating in 245 manufacturing industries. Both panels show that there is significant heterogeneity across plants, firms and industries in terms of employment, reliance on automation and sales. The following Sections characterize the relationships between these variables using several complementary research designs. Appendix Figure A1 reports similar patterns in a balanced panel of plants and firms, which we use for complementary analysis, although plants and firms in the balanced sample tend to be larger in terms of employment and sales, and rely more on automation.

Figure 1 describes the distribution of automation technologies across industries. Panel A focuses on our motive force proxy and reports the five main industries by usage of electric motive force: chemicals, rubber, paper, glass and ceramics, and finally food and beverages. The distribution of total motive power across these industries is relatively even.

These patterns show that the motive power proxy for automation captures a wide range of relevant machines in multiple industries. Appendix Figure A1 illustrates this finding by reporting

examples of machines using electric motors in the top five industries in Panel A of Figure 1. Pasta machines, conveyors and chemical mixers are all captured by our measures.

A more specific technological focus would miss many of these machines, for example by considering industrial robots only. The International Federation of Robots (IFR) defines industrial robots as “automatically controlled, reprogrammable multipurpose manipulator programmable in three or more axes”. Using the IFR data, Panel B of Figure 1 shows that industrial robots are concentrated in the motor vehicle industry, which accounts for almost 60 % of robots. The other industries in the top five for robots — rubber, machinery, food and beverages, and metal products — account for much smaller shares.

III Descriptive Evidence: Stylized Facts and Event Studies

This section provides descriptive evidence on the relationship between automation, employment, sales and the labor share. We find that firms that use more automation technologies increase their sales, total employment, as well as employment of medium-skill and low-skill workers specifically, while the labor share remains stable. These findings challenge the view that automation technologies lead to a fall in the labor force because workers are replaced by machines, suggesting that the productivity effect may outweigh the displacement effect.

III.A Research Design

When a firm relies more extensively on automation technologies, what happens to sales, employment, demand for worker skills, and the labor share? In this section, we investigate this question in the population of firms and plants. We first report stylized facts on the relationship between these variables in Subsection III.B, then we provide elasticity estimates from event studies in Subsection III.C.

The results in this section should primarily be viewed as descriptive. The event study design can alleviate some of the potential threats to identification (e.g., correlated shocks) thanks to the inclusion of a battery of fixed effects and time-varying controls. Nonetheless, absent a quasi-experiment potential concerns over omitted factors cannot fully be addressed. For example, increased demand or increased competition have a direct impact on employment but may also lead a firm to invest more heavily in automation technologies. It is difficult to sign the potential bias of the estimates of the employment response to automation. In the previous examples, the estimate could be biased upward because of increased demand or biased downward because of increased competition.

After presenting correlational evidence for the population of plants and firms in this section, we validate the causal interpretation of the estimates using a quasi-experimental research design for a subset of firms. Section IV develops a shift-share research design that can be applied to the subset of firms for which exogenous variation in the price of automation technologies is available from trade patterns.

III.B Stylized Facts

We first compare the path of sales, employment and labor share for plants that automate more or less over time. We rank all plants by the change in electric motive power observed in the first three years of the sample, between 1995 and 1998. We then compare the path of outcomes for plants below and above median. All outcomes are normalized to one in the first year of the sample.

Figure 1 present the results. Panel A shows that plants that automate more at the beginning of the sample experience a larger increase in sales over the full sample. By 2013, total (nominal) sales have increased by 100% for plants with automation above median and by only 80% for those below median.

Panels B and C show that plants that automate more expand employment relative to those that automate less. Panel B reports this pattern for high-skill workers. By 2013, the number of high skill workers increases by about 120% for plants above median, compared with 100% for those below median. Panel C shows that the number of low-skill workers decreases in both groups, but more steeply for plants with automation below median. For plants that automate more at the beginning of the sample, low skill employment falls by about 33% by 2013, while the fall is more pronounced and reached about 45% for plants with automation below median.

Finally, Panel D reports the patterns for the labor share, defined as the share of total payroll (inclusive of pensions) in total sales. For both groups of plants the labor share falls over time, but there is not significant different across these groups.

These patterns suggest that automation may not be detrimental to employment or to the labor share. Consistent with the observed increase in sales, the potential productivity effect from automation may more than offset the potential displacement effect on workers. We obtain similar results when repeating this analysis with thresholds other than the median and when using the firm or the industry (rather than the plant) as the level of analysis.

In the remainder of the paper, we refine this analysis to alleviate potential threats to identification. For example, it could be the case that plants that automate more at the beginning of

the sample respond to increased demand, which may on its own explain the observed increase in employment.

III.C Event Studies

To address some of the potential correlated demand or supply shocks that may confound the stylized facts, we introduce an event study design, which can control for time-invariant unobservables as well as industry-year and firm-year fixed effects.

To describe employment dynamics as a firm or plant automates the production process, it would be ideal to use an “extensive margin” event study that isolates investment events in automation technologies. However, in practice we observe that most firms and plants adjust their usage of motive power and their stock of machines on a continuous basis, every year. To leverage the entire variation available in the data, we start by specifying a standard distributed lead-lag model (e.g., Stock and Watson (2015)). In robustness checks, we focus on large investment events, which we infer from large changes in the stock of machines or from a large change in the peak capacity for motive power.

Specification. Indexing plants by i and years by t , our baseline distributed lead-lag model is specified as

$$L_{it} = \sum_{k=-10}^{10} \delta_k \Delta M_{i,t-k} + \mu_i + \lambda_{st} + \epsilon_{it}, \quad (1)$$

with employment L_{it} , the change in electric motive power $\Delta M_{i,t}$, plant fixed effects μ_i , and industry-by-year fixed effects λ_{st} .

This specification allows for delayed responses of employment to changes in automation. The lead-lag coefficient δ_k gives the cumulative dynamic response of the employment outcome (L_{it}) at time $t+k$ to a change in automation at time t , holding fixed the path of the independent variable at all other time horizons. In extensions, we also estimate the impact on outcomes other than employment – sales, wages, and the labor share.

We consider a window of ten years at the plant level for the baseline analysis. As mentioned in Section 2, to ensure that variation in electric motive power does not simply reflect a change in factor utilization intensity, we use changes in the plant’s peak capacity for motive power. To reduce noise we implement the specification by aggregating the data over two-year periods, but the results are similar with yearly changes.

We amend the baseline specification for a variety of robustness checks. First, we repeat the analysis at the firm level, using the peak capacity for motive power and the balance sheet measure

of industrial equipment as alternative proxies for automation. Second, instead of using the full variation in our proxies for automation, we focus on large and discrete investment events by isolating the largest changes in the balance sheet measure and in the peak capacity for motive power. Third, we examine the sensitivity of the estimates to using balanced or unbalanced panels over different time horizons.

Identification. A causal interpretation of the estimates require to satisfy the identification condition:

$$E[\Delta M_{i,t-k} \cdot \epsilon_{it} | \mu_i, \lambda_{st}] = 0 \quad \forall (t, k). \quad (2)$$

The estimated coefficients for the “leads” (i.e., $\hat{\delta}_k$ with $k < 0$) can be used as a standard pre-trend falsification test. If the identification condition (2) holds, we expect the leads to be statistically insignificant and the point estimates to be close to zero.

Although the lack of pre-trends is a necessary condition, it may not be sufficient to guarantee the validity of the identification condition. Correlated demand and supply shocks may occur exactly at the same time as the firm or plant automates the production process. For example, increased demand or increased competition could lead to increased automation with a simultaneous direct impact on employment.

To alleviate this potential concern, we examine the stability of the estimates when including more stringent time-varying controls λ_{st} . We consider in turn 2-digit-industry by year fixed effects, 4-digit-industry by year fixed effects, and firm-year fixed effects. The specification with firm-year fixed effects only exploits variation in automation across plants within the same firm, controlling for all time-varying demand and supply shocks at the firm level. By studying coefficient stability, we can bound the potential role of unobservables as in Oster (2019).

Baseline Results. Figure 3 reports the results of the plant-level event studies. Panel A implements the distributed lead-lag model with 2-digit industry by year fixed effects. We find that employment increase at plants that automate more, using the motive power proxy. The elasticity of plant employment to motive power is +0.2 on impact. The response of employment is then amplified over time, with an elasticity of +0.4 after ten year. The point estimates are precise; the 95% confidence interval rejects an employment elasticity below +0.35 or above +0.42 after ten years.

There is no sign of pre trends: conditional on the controls included in our statistical model, plants that automate more at a given time were on a comparable employment path in prior years and start diverging only afterwards. This finding restricts the potential set of confounder that could

explain the increase in employment — namely, confounding shocks need to occur simultaneously to the increase in automation.

Panels B and C of Figure 3 study the sensitivity of the estimates to changes in the set of industry-year fixed effects. Panel B introduces 4-digit industry by year fixed effects. The estimates are almost identical to Panel A, with slightly larger standard errors. Panel C implements a more stringent specification, introducing firm-by-year fixed effects in the sample of firms with multiple plants. The results remain stable, with point estimates nearly unchanged and no pre-trends. The standard errors increase because we only exploit residual variation within year, but the magnitudes are very similar to Panel A: a 1% increase in automation leads to a 0.2% increase in the labor force on impact. The positive response of the labor force is amplified over time, reaching an increase of 0.4% after ten years.

The stability of the estimates across specification bolsters the plausibility of the statistical model. To explain the patterns we document, potential confounding factors must have precisely the same timing as automation at the plant level and must have stronger explanatory power than firm-year fixed effects.

Heterogeneity. Figure 4 documents heterogeneity across skill groups, using the specification with 4-digit industry by year fixed effects. The three panels show that we find a comparable positive response for high-skill, medium-skill and low-skill workers. As previously, the employment elasticity to automation is about +0.2 on impact and increases to about +0.4 after ten years. The point estimates are estimated less precisely for low skill workers, who account for a smaller share of the total workforce, but they are very similar in magnitude. We also find that relative wages across skill groups remain unchanged after automation.

These results indicate that automation does not have different effects across broad skill groups within the firm. Online Appendix Figure A2 focuses on the subset of unskilled industrial workers, who are more likely to perform routine tasks that may be taken over by automated technologies. We find that the employment elasticity remains positive and comparable in magnitude for industrial unskilled workers.

These findings indicate that the distributional effects of automation in the labor market are subtle. They may occur within each skill group, depending on the set of tasks performed across detailed occupations. In ongoing extensions, we document these distributional effects using data on occupations and tasks.

Robustness checks and additional results. To assess the robustness of our results, we repeat

the analysis at the firm level with the balance sheet measure of the value of industrial equipment. In Figure 5, we isolate investment events by focusing on changes in the balance sheet value of industrial equipment above the 90th percentile of that distribution. We then implement an event study design around these investment events. We find that total firm employment increases after an increase in industrial equipment, with a semi-elasticity of about $+0.2$ after eight years. Online Appendix Figure A3 repeats this analysis but also leverages variation in the amount of investment observed conditional on being above the threshold, rather than a binary indicator. The results are similar, with an elasticity close to $+0.3$ after eight years.

Next, we present an additional robustness check to assess whether our estimates could be confounded by changes in the scale of production. In principle, changes in electric motive force could occur simple because production increases, rather than because of new investments in automation technologies. We already mitigated this potential channel by focusing on peak motive force (rather than actual motive force) as well as by implementing the analysis with the balance sheet measure in Figure 5. To assess more directly the potential relevance of scale effects in our statistical model, we repeat specification (1) but specify the independent variables to be changes in energy usage for production (including fuel and gas) or changes in electricity used for heating, rather than changes in electric motive force.

The results are reported in Figure 6. We find that there is not relationship between employment and changes in energy usage for production or electric heating. Therefore any confounding factor should be correlated specifically with electric motive force and not with other forms of energy used for production or heating.

In a series complementary analyses, we find similar positive employment elasticities when aggregating the motive power measure at the firm level, when using alternative definitions of skill groups, and when focusing on the subset of industries with a large share of IFR robots. Consistent with the stylized facts discussed at the beginning of this section, we also find that sales increase after an increase in automation, while average wages and the labor share remain stable.

Limitations. Despite the robustness of the positive relationship between automation and employment at the firm level, two potential concerns remain. First, because we do not have an explicit quasi-experimental source of variation, it could be that some unobserved factors explain this positive relationship. We address this limitation in Section 4 with a shift-share design. Second, the positive firm-level relationship may be misleading because of business-stealing effects across firms. In Section 5, we conduct a similar analysis at the industry level to incorporate such reallocation

effects.

IV Causal Estimates from Shift-Share IV Research Design

In this section, we introduce a quasi-experimental shift-share design to estimate the causal effects of automation on employments, sales, wages and the labor share. The results validate the findings from Section 3: firm-level employment and sales increase following automation, while the average wage and the labor share remain stable.

IV.A Research Design

To estimate the causal effect of automation on employment, sales and other outcomes, the ideal experiment would randomly assign purchasing prices for automation technologies across firms. We approximate this ideal experiment using a shift-share IV design, which leverages two components: shocks and pre-determined exposure shares. The design leverages variation over time, subdividing the sample into 5-year periods.

The shocks are obtained from variation in the cost of imported machines over time across international trading partners in each industry. Shocks are observed across “trading partners by industry” cells indexed by n (for example, imports of machines from China in the French food industry). The shocks g_{nt} are measured as the aggregate changes in import flows of industrial machines from each trading partner for each 2-digit industry between periods t and $t - 1$.

We define the set of industrial machines based the import categories (HS codes) available from the customs data. The measure includes all HS codes mentioning industrial machines, industrial equipment, spare parts for machines, as well as robots or spare parts for robots.

Using changes in the market shares of international suppliers over time is helpful because changes in machines’ quality-adjusted prices are not directly observed. The customs dataset only provides unit values, which are difficult to adjust for quality. But changes in import shares can be used to infer changes in quality adjusted prices. We can infer that countries with rising market shares become more productive at supplying industrial equipment in specific sectors in specific periods. Standard consumer optimization yields that the quality adjusted price must go down when market shares go up. For example, in the 1990s there is an increase in the share of Italian suppliers in the total imports of machines purchased by French firms in the textile sector. For machines imported by French car manufacturers, the share of German suppliers increases in the 2000s. For food products, Dutch suppliers do particularly well in the 2010s. In the baseline specification, we measures the

shocks across 204 trading partners in 24 2-digit industries

The shift-share design combines this set of shocks with variation in the pre-existing network of international supplier relationships across French firms. The exposure share s_{int} and is computed as the share of trading partner n in firm i 's total imports of machines and robots. Intuitively, because of switching costs, a French firm may be more likely to benefit from a trading partner's productivity shock if it has a pre-existing importing relationship with them. Because contemporaneous shares are liable to reverse causality, we use shares lagged by five years.

The shift-share instrument is built by combining the shocks and exposure shares. Consider log changes in employment ΔL_{it} and log changes in electric motive force ΔM_{it} over five-year periods indexed by t and across firms indexed by i . We estimate by 2SLS

$$\begin{cases} \Delta L_{it} = \beta Z_{it} + \gamma X_{it} + \varepsilon_{it}, \\ \Delta M_{it} = \alpha Z_{it} + \tilde{\gamma} X_{it} + \tilde{\varepsilon}_{it}, \end{cases}$$

where Z_i is the shift-share instrument constructed from shocks g_n and (lagged) exposure shares $s_{in} \geq 0$,

$$Z_{it} = \sum_{n=1}^N s_{int} g_{nt}.$$

We study the sensitivity of the estimates to changes in the set of time-varying controls X_{it} .

Identification. The standard shift-share IV identification assumptions apply (see for example Borusyak et al. (2019)). First, a relevance condition must hold such that the instrument has power, i.e. $E[\Delta M_{it} \cdot Z_{it} | X_{it}] \neq 0$. This can be checked directly in the data by computing the first-stage F statistic. The plausibility of the source of variation can also be assessed more directly by checking that the network of international suppliers is relatively sticky.

Figure 7 reports the length of the relationship between a French firm and its main international supplier, depending on the number of years during which machines are imported. The figure shows that importer-supplier relationships are sticky. For example, firms that import machines for 15 years have the same main supplier for 12.1 years on average.

The exclusion restriction underlying this design is that firms linked to increasingly productive suppliers should not be unobservably different from other firms. To test this hypothesis, one can run a falsification test using the lagged outcome variable. Formally, the exclusion restriction can be expressed equivalently at the firm level or in space of productivity shocks (across foreign suppliers):

$$\left(\frac{1}{I} \sum_i z_i \varepsilon_i \rightarrow^p 0 \right) \iff \left(\frac{1}{N} \sum_n \hat{s}_n g_n \bar{\varepsilon}_n \rightarrow^p 0 \right),$$

with $\bar{\varepsilon}_n = (\sum_i s_{in}\varepsilon_i)/\sum_i s_{in}$ and $\hat{s}_n = \frac{1}{I}\sum_i s_{in}$. As discussed in Borusyak et al. (2019), the expression for the exclusion restriction on the right-hand side is helpful because it highlights that identification “comes from” the shocks (while exposure shares can be endogenous).

Specifications. We report the results of the shift-share IV design for five specifications with alternative sets of controls X_{it} . The first specification only includes 2-digit industry by year fixed effects in X_{it} . The second specification adds a set of pre-determined firm controls including lagged turnover, investment, total asset and employment. The third specification controls for lagged electric motive force, the fourth for the lagged balance sheet value of industrial equipment. Finally, because trade flows play a central role for identification, the final specification adds controls for contemporaneous exports to ensure that export flows do not confound the results.

IV.B Results

The results and falsification tests are reported in Tables 2 through 6, using firm-level electric motive force as our proxy for automation.

We start by reporting the OLS relationship between automation and employment at the firm level. Column (1) includes 2-digit industry by year fixed effects and reports an elasticity of employment to automation of +0.235 (s.e. 0.006), which is slightly smaller than in the event study design from Section 3 over a comparable time horizon. The other columns show that this elasticity remains similar in magnitude as we vary the set of controls. The point estimates hover between 0.199 and 0.215 across specifications.

Table 3 reports the estimates of the impact of automation on employment using the shift-share instrument. The baseline specification with 2-digit industry by year fixed effects yields an elasticity of firm employment to automation of +0.341 (s.e. 0.121). The point estimate is statistically significant at the 1% level and the first stage F statistic of 29.3 indicates that the shift-share instrument is strong. The point estimates remain comparable in magnitudes in columns (2) through (5) as we change the set of controls. The point estimates vary between 0.276 and 0.361, are all significant at the 5% level, and are statistically indistinguishable from one another. The first stage F statistic remains large in all specifications.

These results support the conclusion from Section 3: increases in automation lead to higher employment at the level of the firm. The elasticity of 0.34 are very similar to the point estimates from Section 3 at a five-year horizon in the event study. Relative to these results, the OLS estimates from Table 2 appear to be biased downward, although they remained positive and broadly similar

in magnitude. A potential mechanism explaining this downward bias could be that firms tend to automate in response to increased competition, which could have a direct negative effect of employment.

Table 4 reports the results of pre-trend falsification tests, using the lagged change in employment as the outcome variable in the shift-share IV design. Across all five specifications, we cannot reject that there is no relationship between automation (instrumented with the shift-share instrument) and lagged employment growth. The point estimates switch signs across specifications and are smaller in magnitude than those in Table 2.

Table 5 turns to a different outcome, sales. We find that sales increase substantially in response to increased automation, with elasticities ranging from 0.349 to 0.561 across specifications. The relationship is significant at the 5% level in all specifications. This finding is consistent with the role of the productivity effect of automation. Increased automation allows the firm to expand its sales and scale, which requires hiring additional workers for production.

Finally, Table 6 presents estimates of the impact of automation on the labor share. We define the labor share as the ratio of total labor costs to sales, where labor costs include pensions. In all five specifications, we cannot reject that there is no impact of automation on the labor share. In particular, specifications (2) through (5) yield precisely estimated zero elasticities of the labor share to automation. These findings indicate that the productivity effect may offset the task substitution channel in a way that leaves the labor share unchanged.

IV.C Robustness

We implement several robustness checks. First, we find similar results when measuring automation as the balanced sheet value of industrial equipment, instead of electric motive power. Second, we obtain similar patterns with alternative definitions of the labor share, using value-added in the denominator instead of sales, and using total payroll rather than total labor costs in the numerator. Finally, we find no relationship between the average wage at the firm and automation.

V Industry-Level Analysis

In this section, we study the relationship between automation, employment, prices and profits at the level of industries.

First, we find that the relationship between employment and automation is positive on average across industries, but that there is substantial heterogeneity depending on exposure to international

trade. The positive employment effects are concentrated in industries with higher exposure to international competition, consistent with the importance of business stealing effects across countries. Second, we find that the productivity gains from higher automation benefit both consumers through lower prices as well as firm owners via increased profits. Finally, we show that the estimated elasticities of sales, employment and prices to automation can be rationalized in a simple monopolistic competition model where consumers reallocate demand toward firms with increased productivity and lower prices.

V.A Industry-level Employment Effects and International Business Stealing

The positive plant-level and firm-level relationship between employment and automation could in principle be overturned at the industry level, because firms that automate less may be displaced by firm that automate more. To examine how such business stealing effects may add up, we examine the industry-level relationship between automation and employment.

We conduct the analysis over 5-year periods, as in Section 4. For this analysis, we use changes in the balanced sheet value of industrial equipment as a proxy for automation over time, because this measure is available for all firms and can be aggregated accurately to the industry level. In contrast, while the electric motive power measure is only available in a subsample of plants.

Baseline results. Table 7 reports the results. Column (1) implements our baseline specification, with 2-digit industry by year fixed effects. We find that the elasticity of industry employment to automation remains positive, with a point estimate of 0.406 (s.e. 0.065), which is similar to the firm-level elasticities from Sections 3 and 4. Panel A of Figure 8 reports the binned scatter plot corresponding to this specification and shows that the linear specification provides a good fit to the data.

Column (2) adds controls for other capital investments to ensure that the relationship is not driven by correlated supply shocks. The estimates remain similar to Column (1) at 0.275 (s.e. 0.082). The corresponding binned scatter plot, reported in Panel B of Figure 8, depicts the robustness of this relationship.

It may seem surprising to find that the elasticity of industry-level employment to a change in automation is quantitatively similar to the firm-level employment elasticity. The elasticity of substitution of consumer demand is larger between firms within the same industry than between industries. Therefore in a closed economy we would expect the industry-level employment elasticity to automation to be smaller than at the firm level, because demand reallocation is smaller at the

industry level than at the firm level.⁶

However, in an open economy, the industry-level elasticity of substitution of consumer demand may remain high, because domestic producers compete with foreign suppliers and produce relatively substitutable goods (e.g., Broda and Weinstein (2006)). To assess the role of international trade, Columns (3) to (6) of Table 7 repeat the analysis for subsets of industries with trade exposure above or below median.⁷

Heterogeneity by exposure to international competition. We find that the positive industry-level relationship between automation and employment is driven by industries that face a higher degree of international competition. With higher exposure to international competition, the point estimate in Column (3) is 0.401 (s.e. 0.124) and is similar to firm-level employment elasticities. With additional controls in Column (4), the magnitude of the point estimate remains almost unchanged, with statistical significance at the 1% level.

In contrast, with lower exposure to international competition, the point estimate loses statistical significance and falls in magnitude to 0.111 (s.e. 0.082) with the baseline specification in Column (5), and falls further to 0.016 (s.e. 0.136) with additional controls in Column (6). When exposure to international competition is low, the positive relationship between employment and automation disappears, but it is instructive to note that it does not turn negative.

Figure 9 reports the binned scatter plots corresponding to the specifications in Columns (4) and (6) and illustrates the robustness of these relationships.

Robustness and additional results. The industry-level OLS relationships described above may be confounded by omitted factors. In robustness checks, we repeat the analysis using industry-level event studies (parallel to the firm-level event studies from Section 3), as well as an industry-level shift share design (parallel to the firm-level design from Section 4).

Table 8 reports the industry-level relationships between automation and several other outcomes. The sales elasticity is +0.374, compared with elasticities of +0.229 for payroll, +0.224 for labor cost, +0.508 for value added, and +0.781 for profits. These results show that the increase in profits is particularly pronounced relative to the increases in sales, employment and total payroll. Online Appendix Figure A4 depicts these relationships graphically.

⁶We expect to find larger employment effects when consumers' demand elasticity of substitution is larger because consumers reallocated their spending toward firms or sectors where productivity increases and prices fall. See Section 5.3 for a complete discussion.

⁷We use export shares to measure exposure to international competition. The results are similar with a split based on imports of final products.

V.B Industry-level Price Effects

The positive response of employment to increased automation, at both the firm level and industry level, indicates that the productivity effects of automation may be large, large enough to offset the displacement effect. We now estimate the response of prices to increased automation, using data from INSEE on industry-level producer price indices.

Table 8 reports the results, using changes in the balance sheet value of industrial equipment as a proxy for automation. Using the baseline set of controls in Column (1), we find that at the industry level a 1% increase in automation leads to a 0.05% (s.e. 0.017) fall in the industry producer price index. Columns (2) and (3) indicate that the point estimate remain stable with alternative sets of controls, ranging from -0.062 to 0.058, with statistical significance at the 1% level.

Figure 10 depicts graphically the relationship between industry-level producer price indices and changes in automation. Panel A controls for 2-digit industry by year fixed effects as well as contemporaneous changes in capital investments (as in Column (1) of Table 8). Panel B adds controls for changes in industry value added to ensure that the price response is not driven by scale effects alone (as in Column (2) of Table 8). In both cases, the negative linear fits are robust and provide a good approximation to the underlying data.

V.C The Demand Reallocation Channel

We now assess whether the estimated industry-level increase in employment and sales can be explained by the observed price changes following automation. Intuitively, because we found that prices fall in response to automation (Table 9), consumers should reallocate their expenditures toward industries that automate more. The magnitude of this reallocation effect is governed by consumers' demand elasticity of substitution.

To assess the plausibility of this potential explanatory channel, we present a simple calibration in a CES framework. The goal is to assess whether standard estimates of consumers' demand elasticities can rationalize the positive employment and sales effects from Table 8, together with the negative price effects from Table 9.

Assume consumers have CES preferences over a set of varieties that may be supplied by domestic or foreign industries and are indexed by $k \in \Omega$. Given our focus on industry-level outcomes, we interpret varieties as industry-specific aggregates, which combine all varieties produced in the same industry by a given country (domestic or foreign).

The utility of agent i is given by

$$U_i = \left(\sum_{k \in \Omega} \omega_{k,i} q_{k,i}^{1-\sigma} \right)^{1/(1-\sigma)},$$

where σ is the elasticity of substitution between varieties, $q_{k,i}$ is the quantity index for variety k , and $\omega_{k,i}$ is a taste parameter reflecting the intensity of i 's preference for k . There is also a price index p_k for each variety, which corresponds to the price index studied in Table 8.

Consider a perturbation of the equilibrium: domestic firms adopt automation technologies, which results in changes in prices $\{p_k\}$ and equilibrium quantities $\{q_k\}$. CES preferences yield a convenient relationship between the change in the price index for industry k , p_k , and the change in the total total sales, $p_k \cdot q_k$:

$$\Delta \log(p_k) = -\frac{1}{\sigma - 1} \Delta \log(p_k \cdot q_k). \quad (3)$$

In response to a 1% increase in automation, according to Column (1) of Table 8 we have $\Delta \widehat{\log(p_k \cdot q_k)} = 0.374$, and according to Column (1) of Table 9 we have $\Delta \widehat{\log(p_k)} = -0.050$. To satisfy equation (3), these estimates imply the following demand elasticity of substitution :

$$\hat{\sigma} = 1 - \frac{\Delta \widehat{\log(p_k \cdot q_k)}}{\Delta \widehat{\log(p_k)}} = 8.48.$$

Is the magnitude of $\hat{\sigma}$ in line with existing estimates? A demand elasticity of substitution of 8.48 is relatively high, but it is consistent with estimates of substitution elasticities between varieties produced by different countries for the same industry. For example, Broda and Weinstein (2006) estimates a mean demand elasticity of substitution of 7.5 between internationally traded varieties (within 5-digit SITC industries). This result indicates that the consumer demand substitution channel is plausible in an open economy facing international competition.

In contrast, estimated consumer demand elasticities between domestic industries are much smaller and closer to one (e.g., Costinot and Rodríguez-Clare (2014)). It would be difficult to rationalize the industry-level results on sales and employment in a closed economy, because industry-level substitution would need to operate between industries (rather than between products produced either domestic firms or by international competitors within the same industry) and would require large price changes which we do not observe in the data. Competition with international suppliers providing close substitutes can explain why the relationship between automation and employment can remain positive even at the industry level, because the response of consumer demand can be large.

This observation may also help reconcile some of the diverging industry-level estimates in the literature. For example, Dauth et al. (2018) find a positive relationship between robotization and employment in Germany, a country which relies heavily on exports. In contrast, Acemoglu and Restrepo (2019) report a negative relationship in the United States, where domestic firms have a larger domestic market and are less exposed to international competition (i.e., business stealing effects operate primarily between domestic firms rather than internationally).

VI Conclusion

In this paper, we have leveraged new micro data on the population of firms and plants in the French manufacturing sector to provide a unified analysis of the effects of automation technologies on employment, wages, prices and profits between 1994 and 2015.

At all levels of analysis — plant, firm and industry — the relationship between automation and employment is positive, indicating that in practice the productivity effect tends to outweigh the displacement effects. There is also an increase in sales, a fall in consumer prices, and a substantial increase in firm profits.

At the industry-level, we find that the relationship between employment and automation is positive on average, but that there is substantial heterogeneity depending on exposure to international trade. While the employment response is positive and significant in industries that face international competition, there is no significant effect in sectors with low exposure to international competition.

These patterns can be explained by a simple consumer demand substitution channel. After adopting automation technologies, firm owners increase their profits but pass through some of the productivity gains to consumers, inducing scale effects. Automation can thus lead to higher firm profits, lower consumer prices, increased demand, and in turn to increased firm and industry scale, and to more employment at the expense of foreign competitors. 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.

Taken together, the results suggest that automation can increase labor demand and generate productivity gains that are broadly shared across workers, consumers and firm owners. Because the observed distributional effects of automation are nuanced, training programs targeting specific groups of workers that may be negatively affected by automation (e.g., older workers specializing in routine tasks) may be more appropriate than broader tax instruments (e.g., taxing robots or

capital, or increasing redistribution through the income tax system). Developing and testing such policies is therefore a promising direction for research and policy going forward.

References

- Acemoglu, Daron and Pascual Restrepo**, “The race between man and machine: Implications of technology for growth, factor shares, and employment,” *American Economic Review*, 2018, *108* (6), 1488–1542.
- **and** —, “Robots and jobs: Evidence from US labor markets,” *Journal of Political Economy*, 2019.
- , **Claire Lelarge, and Pascual Restrepo**, “Competing with Robots: Micro Evidence from France,” *Working Paper*, 2019.
- Autor, David**, “Why are there still so many jobs? The history and future of workplace automation,” *Journal of economic perspectives*, 2015, *29* (3), 3–30.
- Bessen, James, Marten Goos, Anna Salomons, and Wiljan van den Berge**, “Automatic Reaction: What Happens to Workers at Firms that Automate?,” *Working Paper*, 2019.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-experimental shift-share research designs,” *Working Paper*, 2019.
- Bowen, Howard R and Garth Leroy Mangum**, *Automation and economic progress*, Prentice-Hall, 1966.
- Broda, Christian and David E Weinstein**, “Globalization and the Gains from Variety,” *The Quarterly journal of economics*, 2006, *121* (2), 541–585.
- Brynjolfsson, Erik and Andrew McAfee**, *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company, 2014.
- Caselli, Francesco and Alan Manning**, “Robot arithmetic: new technology and wages,” *American Economic Review: Insights*, 2019, *1* (1), 1–12.
- Chandler, Daniel and Michael Webb**, “How Does Automation Destroy Jobs? The ‘Mother Machine’ in British Manufacturing, 2000-2015,” *Working Paper*, 2019.

- Charnoz, Pauline and Michael Orand**, “Technical change and automation of routine tasks: Evidence from local labour markets in France, 1999-2011,” *Economie et Statistique*, 2017, 497 (1), 103–122.
- Chiacchio, Francesco, Georgios Petropoulos, David Pichler et al.**, *The impact of industrial robots on EU employment and wages: A local labour market approach*, Bruegel, 2018.
- Costinot, Arnaud and Andrés Rodríguez-Clare**, “Trade theory with numbers: Quantifying the consequences of globalization,” in “Handbook of international economics,” Vol. 4, Elsevier, 2014, pp. 197–261.
- **and Ivan Werning**, “Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation,” *Working Paper*, 2018.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, Nicole Woessner et al.**, “Adjusting to robots: Worker-level evidence,” *Working Paper*, 2018.
- Dixon, Jay, Bryan Hong, and Lynn Wu**, “The Employment Consequences of Robots: Firm-Level Evidence,” *Working Paper*, 2019.
- Graetz, Georg and Guy Michaels**, “Robots at work,” *Review of Economics and Statistics*, 2018, 100 (5), 753–768.
- Guerreiro, Joao, Sergio Rebelo, and Pedro Teles**, “Should robots be taxed?,” *Working Paper*, 2017.
- Humlum, Anders**, “Robot Adoption and Labor Market Dynamics,” *Working Paper*, 2019.
- Keynes, John Maynard**, “Economic possibilities for our grandchildren,” in “Essays in persuasion,” Springer, 1930, pp. 321–332.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka**, “Robots and firms,” *Working Paper*, 2019.
- Leontief, Wassily**, “Machines and Man,” *Scientific American*, 1952, 187 (3), 150–164.
- Mann, Katja and Lukas Puttmann**, “Benign effects of automation: New evidence from patent texts,” *Working Paper*, 2018.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, “Uneven Growth: The Impact of Automation on Income and Wealth Inequality,” *Working Paper*, 2019.

Oster, Emily, “Unobservable selection and coefficient stability: Theory and evidence,” *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.


Stock, James H and Mark W Watson, “Introduction to econometrics: Updated,” 2015.

Webb, Michael, “The Impact of Artificial Intelligence on the Labor Market,” *Working Paper*, 2019.


Zeira, Joseph, “Workers, machines, and economic growth,” *The Quarterly Journal of Economics*, 1998, *113* (4), 1091–1117.

Figure 1: Distribution of Automation Technologies across Industries

A. Top 5 Industries by Usage of Motive Force

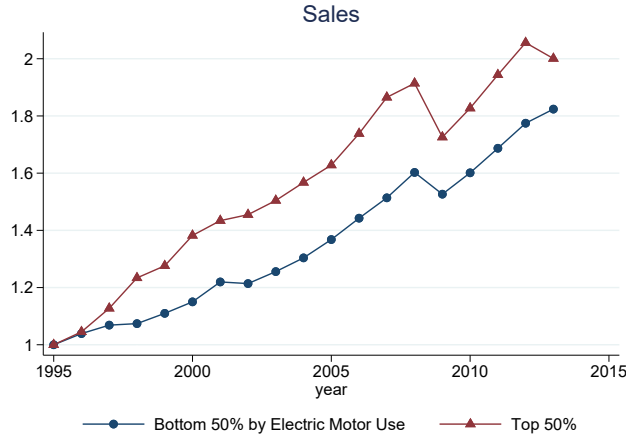


B. Top 5 Industries by Count of Robots




Notes: See Section 2 for a description of the data.


Figure 2: Stylized Facts on the Paths of Sales and Employment after Automation




(a) Sales



(b) High-skill Employment



(c) Low-skill Employment

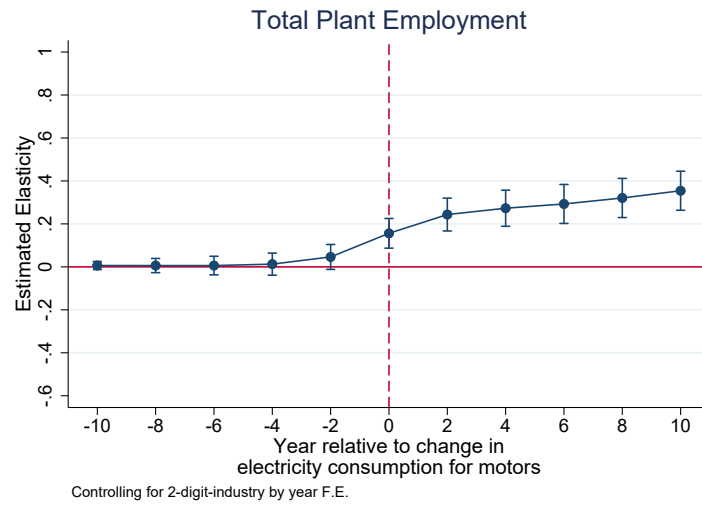


(d) Labor Share


Notes: This figure describes the path of sales, employment and the labor share for plants with different propensities to automate in the first three years of the sample. All outcomes are normalized to one in 1995. See Section 3 for a description of the methodology.

Figure 3: Plant-Level Event Studies

A. With 2-digit industry by year fixed effects



B. With 4-digit industry by year fixed effects



C. With firm by year fixed effects





Figure 4: Heterogeneity across Skill Groups


A. High-Skill Employment

Plant Employment - High Skill




B. Medium-Skill Employment

Plant Employment - Medium skill




C. Low-Skill Employment

Plant Employment - Low Skill




Notes: See Section 3 for the methodology.

Figure 5: Firm-level Event Study with Investment in Industrial Equipment




Notes: See Section 3 for the methodology.

Figure 6: Falsification Test for the Role of Changes in Scale



Notes: The plant-level event study in this figure is similar to Figures 3 and 4, except that the independent variable is the change in energy usage for production other than motive force (including fuel and gas) or the change in electricity used for heating. See Section 3 for the methodology.


Figure 7: Persistence of Importer-Supplier Relationships




Notes: See Section 4 for a description of the methodology.

Figure 8: Industry-level Relationship between Automation and Employment

A. With 2-digit Industry Year F.E.




B. With Controls for Changes in Other Capital Investments




Notes: See Section 5 for a description of the methodology.

Figure 9: Heterogeneity by Exposure to International Competition

A. High International Competition (above median)




B. Low International Competition (below median)




Notes: See Section 5 for a description of the methodology.

Figure 10: Industry-level Relationship between Automation and Prices

A. With Controls for Changes in Other Capital Investments



B. With Controls for Changes in Other Capital Investments and Value Added



Notes: See Section 5 for a description of the methodology.

Table 1: Summary Statistics

A. Average Annual Levels (1995-2013)

	Units N	Units-by-year N	Mean	S.D.	p5	p25	p50	p75	p95
<u>Panel A: Plant level</u>									
Employment			157	270	5	30	86	179	534
Automation - Motive Force (TOE)	19,448	77,489	127	263	2	11	38	119	551
<u>Panel B: Firm level</u>									
Employment			130	220	14	30	68	147	432
Sales (thousands of euros)			28,131	90,957	1,234	3,672	9,891	25,722	96,870
Automation	16,227	67,542							
Motive force (TOE)			80	145	1	8	26	79	361
Industrial machines (thousands of euros)			7,264	17,736	44	449	1,718	6,211	32,206
<u>Panel C: Industry level</u>									
Employment			12,092	16,432	526	2,706	6,866	14,330	40,190
Sales (thousands of euros)	245	4,655	2,999	6,059	126	644	1,660	3,279	9,272
Automation - Industrial machines (thousands of euros)			9,163	19,380	267	1,710	3,931	9,128	29,080

B. Average Annual Changes (1995-2013)

	Units N	Units-by-year N	Mean	S.D.	p5	p25	p50	p75	p95
<u>Panel A: Plant level</u>									
Employment			-4	64	-46	-8	0	4	32
Automation - Motive Force (TOE)	19,448	58,041	0	51	-41	-3	0	5	41
<u>Panel B: Firm level</u>									
Employment			-1	41	-29	-5	0	3	24
Sales (thousands of euros)			484	23,033	-7,398	-677	113	1,406	9,894
Automation	16,227	51,171							
Motive force (TOE)			1	37	-30	-2	0	4	33
Industrial machines (thousands of euros)			325	3,557	-508	0	62	364	2,369
<u>Panel C: Industry level</u>									
Employment			-137	1,244	-1,292	-302	-62	76	861
Sales (thousands of euros)	245	4,410	42	1,415	-368	-33	16	116	570
Automation - Industrial machines (thousands of euros)			79	3,450	-726	-47	39	239	1,245

Notes: See Section 2 for a description of the datasets.

Table 2: Relationship between Automation and Employment, OLS

	Δ_5 Employment				
	(1)	(2)	(3)	(4)	(5)
Δ_5 Motor Cons.	0.235*** (0.00637)	0.207*** (0.00611)	0.215*** (0.00611)	0.199** (0.00608)	0.211*** (0.00630)
Industry-year F.E.	✓	✓	✓	✓	✓
Firm Controls		✓	✓	✓	✓
Lagged Motor Cons.			✓		✓
Lagged Machines				✓	
Exports					✓
N	30,180	30,180	30,180	30,180	30,180

Notes: See Section 4 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Causal Effects of Automation on Employment with Shift-Share IV

	Δ_5 Employment				
	(1)	(2)	(3)	(4)	(5)
Δ_5 Motor Cons.	0.341*** (0.121)	0.361*** (0.1276)	0.410** (0.167)	0.276** (0.138)	0.430*** (0.202)
First-Stage F	29.3	26	16.8	20.6	17.9
Industry-year F.E.	✓	✓	✓	✓	✓
Firm Controls		✓	✓	✓	✓
Lagged Motor Cons.			✓		✓
Lagged Machines				✓	
Exports					✓
N	29,109	29,109	29,109	29,109	29,109

Notes: See Section 4 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Falsification Tests for Shift-Share IV Design

	Lagged Δ_5 Employment				
	(1)	(2)	(3)	(4)	(5)
Δ_5 Motor Cons.	-0.194 (0.185)	-0.0283 (0.177)	-0.156 (0.236)	-0.233 (0.200)	0.120 (0.247)
First-Stage F	25.8	23.8	15.2	20.1	13.1
Industry-year F.E.	✓	✓	✓	✓	✓
Firm Controls		✓	✓	✓	✓
Lagged machines (consumption)			✓		✓
Lagged machines (balanced sheet)				✓	
Exports					✓
N	17,250	16,609	16,609	16,574	15,641

Notes: See Section 4 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Causal Effects of Automation on Sales with Shift-Share IV

	Δ_5 Sales				
	(1)	(2)	(3)	(4)	(5)
Δ_5 Motor Cons.	0.552*** (0.148)	0.422*** (0.148)	0.498** (0.197)	0.349** (0.164)	0.561*** (0.197)
First-Stage F	29.3	26	16.8	20.6	17.9
Industry-year F.E.	✓	✓	✓	✓	✓
Firm Controls		✓	✓	✓	✓
Lagged Motor Cons.			✓		✓
Lagged Machines				✓	
Exports					✓
N	29,109	29,109	29,109	29,109	29,109

Notes: See Section 4 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Causal Effects of Automation on Labor Share with Shift-Share IV

	Δ_5 Labor Cost / Sales				
	(1)	(2)	(3)	(4)	(5)
Δ_5 Motor Cons.	-0.134 (0.0944)	-0.00851 (0.0936)	-0.0298 (0.122)	-0.00607 (0.107)	-0.0691 (0.118)
First-Stage F	29.3	26	16.8	20.6	17.9
Industry-year F.E.	✓	✓	✓	✓	✓
Firm Controls		✓	✓	✓	✓
Lagged Motor Cons.			✓		✓
Lagged Machines				✓	
Exports					✓
N	29,109	29,109	29,109	29,109	29,109

Notes: See Section 4 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Industry-level Relationship between Automation and Employment

	Δ_5 Employment					
	All industries		International Competition			
			Above Median		Below Median	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ_5 Machines	0.406*** (0.065)	0.275*** (0.082)	0.401*** (0.124)	0.382*** (0.150)	0.111 (0.082)	0.016 (0.136)
2-digit industry by year F.E.	✓	✓	✓	✓	✓	✓
Δ_5 Other types of capital		✓		✓		✓
N	972	972	143	143	143	145

Notes: See Section 5 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Industry-level Relationships between Automation and Sales, Payroll, Value Added, and Profits

	Δ_5 Sales	Δ_5 Payroll	Δ_5 Labor Cost	Δ_5 Value Added	Δ_5 Profits
	(1)	(2)	(3)	(4)	(5)
Δ_5 Machines	0.374*** (0.069)	0.229*** (0.088)	0.224*** (0.085)	0.508*** (0.087)	0.781*** (0.174)
2-digit industry by year F.E.	✓	✓	✓	✓	✓
Δ_5 Other types of capital	✓	✓	✓	✓	✓
N	840	840	840	840	840

Notes: See Section 5 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Industry-level Relationship between Automation and Prices

	(1)	(2)	(3)
Δ_5 Machines	-0.050*** (0.017)	-0.062*** (0.018)	-0.058*** (0.020)
2-digit industry by year F.E.	✓	✓	✓
Δ_5 Other types of capital	✓	✓	✓
Δ_5 Value added		✓	✓
Δ_5 Debt and liabilities			✓
N	756	756	756

Notes: See Section 5 for a description of the methodology. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A Online Appendix Figures and Tables

Figure A1: Examples of Automation Technologies



(a) Chemicals



(b) Rubber



(c) Paper



(d) Glass and Ceramics

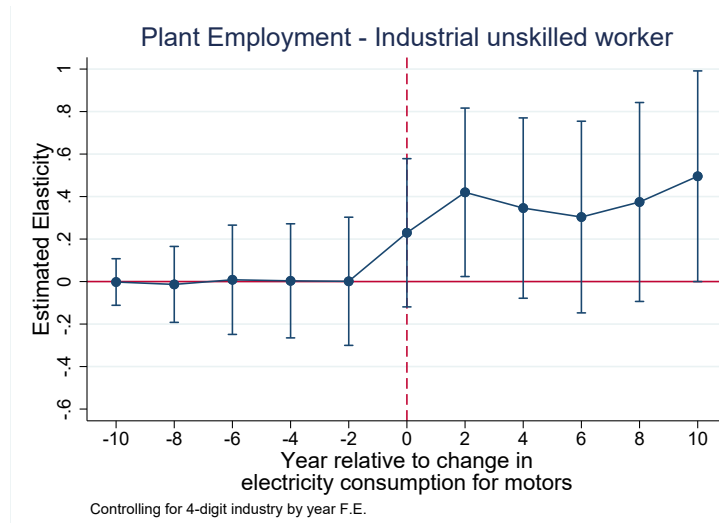


(e) Food

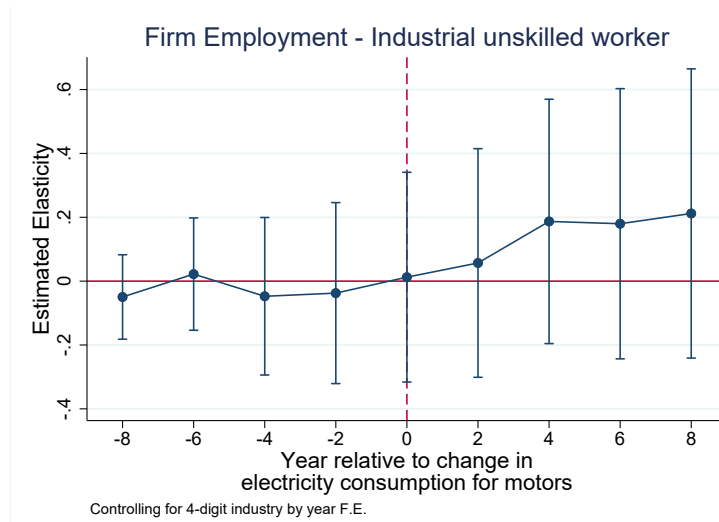
Notes: This figure gives examples of machines for the five industries with the largest usage of motive force, described in Figure 1. See Section 2 for a description of the data.

Figure A2: The Response of Employment of Unskilled Industrial Workers

A. Plant Level




B. Firm Level




Notes: See Section 3 for the methodology.

Figure A3: Firm-level Event Study with Investment in Industrial Equipment, including Intensive Margin



Notes: See Section 3 for the methodology. The figure is similar to Figure 3 in the main text, except that variation in the amount of investment above the threshold is also used, instead of a binary indicator.

Figure A4: Industry-Level Relationships between Automation and Sales, Total Payroll, and Total Labor Cost



Notes: See Section 5 for the methodology.

Table A1: Summary Statistics for Balanced Panel

A. Average Annual Levels (1995-2013, balanced panel)

	Units	N	Units-by-year	N	Mean	S.D.	p5	p25	p50	p75	p95
<u>Panel A: Plant level</u>											
Employment											
Automation - Motive Force (TOE)	553		10,507		332	353	82	131	210	402	963
					247	363	13	47	112	277	975
<u>Panel B: Firm level</u>											
Employment					244	290	52	105	164	283	659
Sales (thousands of euros)					54,171	94,681	7,891	16,373	28,807	56,319	170,662
Automation	500		9,500								
Motive force (TOE)					152	177	11	36	81	207	529
Industrial machines (thousands of euros)					16,884	23,005	1,237	4,153	9,914	19,740	62,646
<u>Panel C: Industry level</u>											
Employment					12,092	16,432	526	2,706	6,866	14,330	40,190
Sales (thousands of euros)	245		4,655		2,999	6,059	126	644	1,660	3,279	9,272
Automation - Industrial machines (thousands of euros)					9,163	19,380	267	1,710	3,931	9,128	29,080

B. Average Annual Changes (1995-2013, balanced panel)

	Units	N	Units-by-year	N	Mean	S.D.	p5	p25	p50	p75	p95
<u>Panel A: Plant level</u>											
Employment					0	42	-45	-9	-1	8	45
Automation - Motive Force (TOE)	553		9,954		1	59	-55	-6	1	9	60
<u>Panel B: Firm level</u>											
Employment					0	39	-32	-7	0	7	34
Sales (thousands of euros)					1,481	20,840	-11,018	-1,232	663	3,334	16,673
Automation	500		9,000								
Motive force (TOE)					2	46	-46	-5	0	8	52
Industrial machines (thousands of euros)					742	3,096	-782	41	306	945	4,000
<u>Panel C: Industry level</u>											
Employment					-137	1,244	-1,292	-302	-62	76	861
Sales (thousands of euros)	245		4,410		42	1,415	-368	-33	16	116	570
Automation - Industrial machines (thousands of euros)					79	3,450	-726	-47	39	239	1,245

Notes: See Section 2 for a description of the datasets.