EMPIRICAL RESEARCH ON AUTOMATION AND "SMART" TECHNOLOGIES[‡]

Competing with Robots: Firm-Level Evidence from France[†]

By DARON ACEMOGLU, CLAIRE LELARGE, AND PASCUAL RESTREPO*

Automation substitutes capital for tasks previously performed by labor, reducing the labor share of value added and increasing value added per worker in the process. While the higher productivity from automation tends to increase labor demand, its displacement effect may outweigh this positive impact and may lead to an overall decline in employment and wages (Acemoglu and Restrepo 2019). Acemoglu and Restrepo (forthcoming) estimates negative effects from the introduction of one of the leading examples of automation technology, industrial robots, across US local labor markets, suggesting that the displacement effects could be significantly larger than the productivity effect.¹ Firm-level evidence is useful as well for understanding how automation is affecting the production process

[‡]*Discussants:* Susan R. Helper, Case Western Reserve University; Robert Seamans, New York University; David Deming, Harvard University; Daniel Rock, Massachusetts Institute of Technology.

*Acemoglu: Massachusetts Institute of Technology (email: daron@mit.edu); Lelarge: University of Paris-Saclay (email: claire.lelarge@universite-paris-saclay.fr); Restrepo: Boston University (email: pascual@bu.edu). We thank our discussant Robert Seamans as well as David Autor and John Van Reenen for their comments. We also thank SYMOP for generously sharing their data with us. Financial support from Accenture, Banque de France, Google, IBM, Microsoft, the National Science Foundation, Schmidt Sciences, the Sloan Foundation, and the Smith Richardson Foundation is gratefully acknowledged. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the Investissements d'Avenir program (reference: ANR-10-EQPX-17—Centre d'accès sécurisé aux données—CASD).

 † Go to https://doi.org/10.1257/pandp.20201003 to visit the article page for additional materials and author disclosure statement(s).

and productivity.² But its interpretation is complicated by the fact that firms adopting automation technologies reduce their costs and may expand at the expense of their competitors.

In this paper, we study firm-level changes associated with robot adoption by using data from France between 2010 and 2015. Consistent with our theoretical expectations (which are developed further in the online Appendix), we find that firm-level adoption of robots coincides with declines in labor shares, increases in value added and productivity, and declines in the share of production workers. In contrast to the market-level effects, however, overall employment increases faster in firms adopting robots.

This positive employment effect may be because firms with greater growth potential are more likely to adopt robots, generating a classic omitted variable bias. Equally important, this positive effect may be a consequence of reallocation of output and labor toward firms that reduce their costs relative to their competitors. We show that such reallocation accounts for the positive firm-level impact of robots. Firms whose competitors adopt robots experience significant declines in value added and employment.³ In fact, the overall impact of robot adoption (combining own and spillover effects) is negative and implies that a 20 percentage point increase in robot adoption (as in our sample) is associated with a 3.2 percent decline in industry employment.

¹Graetz and Michaels (2018) uses variation across industries and countries and finds lower labor share and higher productivity from robots, but negative effects only for unskilled workers. Aghion, Antonin, and Bunel (2019) finds negative regional employment effects in France, while

Dauth et al. (2019) estimates employment declines in manufacturing, but not overall, across German regions.

²For papers using firm-level data on robots, see Dinlersoz and Wolf (2018); Bessen et al. (2019); Dixon, Hong, and Wu (2019); Bonfiglioli et al. (2019); Humlum (2019); and Koch, Manuylov, and Smolka (2019).

³This aligns with Koch, Manuylov, and Smolka's (2019) findings from Spain.

Finally, we use our data to study the decline in the French manufacturing labor share. As in Autor et al. (forthcoming), we find that this decline is explained by a lower covariance between firm-level value added and labor share. However, in our data, this pattern is explained not so much because expanding firms had lower labor shares (or higher markups) but because firms adopting robots are large and expand further as they experience significant relative declines in their labor share.

I. Data on French Robots

Our sample includes 55,390 firms that were active from 2010 to 2015 in the French manufacturing sector. For these firms, we have data on sales, value added, employment (total hours of work), share of production workers, and wages (and can estimate total factor productivity). For firms that export, we also have data on export prices and quantities by detailed product. Further information on the data and the sample are provided in the online Appendix.

We identified 598 manufacturing firms that adopted (purchased) industrial robots during this period by using several sources, including a survey by the Ministry of Industry, information provided by French robot suppliers about their list of clients, customs data on imports of industrial robots by firm, and the French fiscal files, which include information on accelerated depreciation allowances for the purchase of industrial robots. Although only 1 percent of our firms purchased robots in 2010–2015, these firms account for 20 percent of total manufacturing employment. Table A.1 in the online Appendix describes our sample.

Figure 1 presents information on robot adopters. These tend to be the larger firms, as shown by the higher rates of adoption at top percentiles of the size distribution within the 258 four-digit industries in our sample. For example, 13 percent of firms in the top 1 percent of the industry sales distribution adopted robots, while there is almost no robot adoption among firms below the twentieth percentile of the sales distribution. Robot adopters are also likely to be in industries where there are more major advances in robotics technology and more rapid spread of robots in other industrialized economies. In particular, the figure shows that adoption rates are about 50 percent higher in industries with greater *adjusted*



Figure 1. Share of Robot Adopters among Firms in Different Percentiles of the Sales Distribution within Four-Digit Industries; Shown for All Industries and for Industries with High and Low APR

penetration of robots (APR) in other European countries (shown with darker shading).⁴

II. Firm-Level Changes

We first study firm-level changes in value added, productivity, the labor share, employment, and wages associated with robot adoption. Specifically, we estimate the following regression model by ordinary least squares across firms, denoted by f:

(1)
$$\Delta \ln y_f = \beta \cdot Robot_f + \gamma \cdot X_f + \alpha_{i(f)} + \delta_{c(f)} + \varepsilon_f.$$

On the right-hand side we use the change in the log of several firm-level outcomes between 2010 and 2015. The main regressor is $Robot_f$, a dummy for whether the firm adopted robots in 2010–2015. We control for baseline firm characteristics that are likely to be correlated with subsequent changes in our variables of interest (log employment and log value added per

⁴The *APR* measures the common increase in robot use in an industry among advanced economies (excluding France) since 1993 and adjusts for the mechanical effect of industry growth on robot use (see Acemoglu and Restrepo forthcoming). Manufacturing industries with a high *APR* are pharmaceuticals, chemicals, plastics, food and beverages, metal products, primary metals, industrial machinery, and automotive. Industries with a low *APR* are paper and printing, textiles and apparel, electronic appliances, furniture, mineral products, and other transportation vehicles.

	$\Delta \log$ value added	Δ labor share	Δ production employment share	Δ log value added per hour	$\Delta \log$ revenue TFP	$\Delta \log$ employment (in hours)	Δ log mean hourly wage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Unweighted e	estimates						
Robot adopter	0.204 (0.030)	-0.043 (0.009)	-0.016 (0.007)	0.095 (0.018)	0.024 (0.007)	0.109 (0.020)	0.009 (0.004)
R^2	0.083	0.161	0.014	0.222	0.196	0.093	0.024
Panel B. Employment-	weighted estin	nates					
Robot adopter	0.094 (0.025)	-0.027 (0.012)	-0.006 (0.006)	0.040 (0.029)	-0.011 (0.013)	0.054 (0.017)	-0.008 (0.008)
R^2	0.216	0.274	0.080	0.323	0.298	0.188	0.139

TABLE 1-ESTIMATES OF EFFECTS OF ROBOT ADOPTION ON FIRM-LEVEL OUTCOMES

Notes: The sample consists of 55,390 firms, of which 598 are robot adopters. Panel A presents unweighted estimates. Panel B presents estimates weighting each firm by its employment (in hours) in 2010. All specifications control for baseline firm characteristics (log employment and log value added per worker in 2010, as well as dummies for whether the firm is affiliated with a larger corporate group), four-digit industry fixed effects for the main industry in which each firm operates, and fixed effects for the commuting zone that houses each firm's largest establishment. The online Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within four-digit industries are in parentheses.

worker in 2010, as well as dummies for whether the firm is affiliated with a larger corporate group); four-digit industry fixed effects for the main industry in which each firm operates, $\alpha_{i(f)}$; and fixed effects for the commuting zone that houses each firm's largest establishment, $\delta_{c(f)}$. We report standard errors that are robust to heteroskedasticity and cross-firm correlation within four-digit industries.

Table findings 1 reports our using unweighted specifications (in panel A) and employment-weighted specifications (in panel B). The results in panel A show that, consistent with our theoretical expectations, robot adoption is associated with a 20 percent increase in value added from 2010 to 2015 (standard error = 0.030) as well as a 4.3 percentage point decline in the labor share (standard error = 0.009) and a 1.6 percentage point decline in the production worker share of employment (standard error = 0.007). Value added per hour and revenue total factor productivity (TFP) also increase.⁵ Column 5 shows that, in contrast to market-level results in previous works, employment (total hours of work) also increases in firms adopting robots—by 10.9 percent (standard error = 0.020). Hourly wages rise modestly as well (column 6).

The weighted results in panel B are similar, except that there are no longer positive effects on TFP and hourly wages.⁶ The online Appendix documents that these results are robust to controlling for additional covariates in 2010, including sale distribution percentiles, capital intensity, and the share of production workers in employment.

III. Market-Level Spillovers

As noted above, firms adopting robots, by reducing their costs, may gain market share at the expense of their competitors. If so, employment gains in these firms may go hand in hand with employment losses in other firms, and the market-level effects of automation may be very

⁵The value added and TFP results are not driven by price increases but by higher physical productivity. The online Appendix shows that, for the sample of exporting firms where we have detailed price data, robot adoption is associated with price declines.

⁶Even the positive estimate on hourly wages in panel A, which implies a pass-through elasticity from output per worker to wages of about 0.1 percent, is much smaller than estimates in the literature resulting from other sources of productivity increases, such as obtaining a patent (Kline et al. (2019) and references therein), which generate a pass-through elasticity of about 0.35. This is as expected since automation substitutes capital for labor.

	$\begin{array}{c} \Delta \log \\ \text{employment} \\ (\text{in hours}) \\ (1) \end{array}$	$\Delta \log$ value added (2)	Δ labor share (3)	$\begin{array}{c} \Delta \log \\ \text{employment} \\ (\text{in hours}) \\ (4) \end{array}$	$\Delta \log$ value added (5)	Δ labor share (6)	
	Unweighted estimates			Employment-weighted estimates			
Robot adoption by competitors	-0.105 (0.047)	-0.100 (0.051)	0.002 (0.015)	-0.250 (0.107)	-0.209 (0.159)	-0.008 (0.040)	
Robot adopter	0.106 (0.020)	0.201 (0.030)	-0.043 (0.009)	0.035 (0.022)	0.078 (0.029)	-0.027 (0.012)	
R^2	0.093	0.083	0.161	0.190	0.217	0.274	

TABLE 2-ESTIMATES OF EFFECTS OF ROBOT ADOPTION ON COMPETITORS

Notes: The sample consists of 55,388 firms, of which 598 are robot adopters. Columns 1–3 present unweighted estimates. Columns 4–6 present estimates weighting each firm by its employment (in hours) in 2010. All specifications control for baseline firm characteristics (log employment and log value added per worker in 2010, as well as dummies for whether the firm is affiliated with a larger corporate group), four-digit industry fixed effects for the main industry in which each firm operates, and fixed effects for the commuting zone that houses each firm's largest establishment. The online Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within four-digit industries are in parentheses.

different from its firm-level impact. To investigate this issue, we estimate a variant of equation (1) including a measure of a firm's competitors' robot adoption. This measure is defined as

Adoption by competitors_f

$$= \sum_{i} m_{fi} \cdot \sum_{f' \neq f} s_{if'} \cdot Robot_{f'},$$

where the first sum is over all four-digit industries, and m_{fi} is the share of firm f's sales that are in industry *i*, while the second is over all firms other than f, and $s_{if'}$ is the share of industry *i*'s total sales accounted for by firm f'. Thus, the measure of adoption by competitors gives the sales overlap across four-digit industries between a given firm and all robot adopters in the economy. The shares m_{fi} and $s_{if'}$ are constructed by using sales data by firm and four-digit industry from the fiscal files, which cover 85 percent of sales in our sample. We assume that small firms that are not in the fiscal files sell only in their main four-digit industry. Because equation (1) includes four-digit industry fixed effects, the spillovers are identified from the comparison of firms that are in the same main industry but sell different proportions of their products across industries with varying degrees of competition by robot adopters.

Table 2 presents estimates for employment, value added, and the labor share. We report both

unweighted and employment-weighted estimates, but because our main interest is aggregate effects, we now focus on weighted models. Consistent with the notion that automation leads to expansion at the expense of competitors and that the labor share of value added in a firm depends on its own automation decisions, the estimates in columns 4–6 show that a 10 percentage point increase in robot adoption by competitors is associated with a 2.5 percent decline in employment (standard error = 0.0107) and a 2.1 percent decline in value added (standard error = 0.0159) and, consistent with our theory in the online Appendix, competitors' robot adoption has no impact on a firm's labor share.

These results establish that, because of negative spillovers on competitors, firm-level effects do not translate into similar market-level impacts. What is the overall impact of robot adoption on industry employment? Aggregating own and competitors' effects, we find that robot adoption is associated with an overall decline in industry employment: a 20 percentage point increase in robot adoption (which is the average robot adoption in our sample) is associated with a 3.2 percent decline in industry employment.⁷

⁷The online Appendix shows that this effect on employment is $\beta_o \sum_f (\ell_f/\ell) \times Robot_f + \beta_c \sum_f (\ell_f/\ell) \times Robot_f \times \sum_i m_{fi} \cdot (1 - s_{if})$. Here, β_o is the own-firm estimate of robot adoption and β_c the coefficient on competitors, and ℓ_f/ℓ is the baseline employment share in firm *f*. In our data, own-firm

IV. Superstar Effects and the Labor Share

Our estimates in Table 1 suggest that the labor share of a firm that adopts robots declines by 4 to 6.3 percentage points. To explore the contribution of these changes to the aggregate labor share, we follow Autor et al. (forthcoming) and decompose the observed change in an industry's labor share into the change in the unweighted average within firms and the change in the covariance between the share of value added of a firm and the firm's labor share.⁸ Autor et al. documents that the decline in the labor share is driven by a reduction in the covariance term and suggests that these changes may be due to a superstar phenomenon-firms with low labor shares (or high markups) at the baseline expand due to competitive pressures or winner-takes-all dynamics. Our data enable us to investigate whether similar trends are present in French manufacturing and whether industrial automation is responsible for some of these patterns.

Figure 2 presents the decomposition from Autor et al. (forthcoming) for French manufacturing between 2010 and 2015. As in the authors' US results, there is a decline in overall labor share of 0.93 percentage points, which is entirely driven by a declining covariance term. In fact, the average within-firm change in the labor share is positive. To gauge the contribution of automation to these changes, we further decompose these effects between robot adopters and nonadopters. Interestingly, while—analogous to the US results—the labor share increases for firms not adopting robots, it declines for robot adopters. More importantly, about 80 percent of the decline in the covariance term is accounted

⁸Changes in an industry labor share, λ_i^{ℓ} , can be decomposed as $\Delta \lambda_i^{\ell} = \sum_f \Delta \lambda_f^{\ell} + \Delta \sum_f (\lambda_f^{\ell} - \bar{\lambda}_i^{\ell}) \cdot (s_{if}^v - \bar{s}_i^v)$, where λ_f^{ℓ} is the labor share in firm f, s_{if}^v is the share of value added in industry *i* accounted for by firm f, and $\bar{\lambda}_i^{\ell}$ and \bar{s}_i^v are their unweighted averages. The first term is the *unweighted* within-component and the second is the change in the covariance. The decomposition ignores entry and exit since we use a balanced panel of firms.



Figure 2. Changes in the Labor Share of French Manufacturing Industries for 2010–2015 Decomposed as in Autor et al. (2019); the Decomposition Is Extended to Account for Differences between Robot Adopters and Nonadopters

Note: PP is percentage points.

for by the fact that robot adopters are larger from the outset (-2.81 percentage points) and expand (-0.14 percentage points) at the same time as they reduce their relative labor shares. Notably, this is not due to adopters having lower baseline labor shares.⁹ The residual decline in the covariance term, which includes the superstar effect, accounts for 20 percent of the decline in the covariance term. Our results therefore provide a different interpretation of the forces behind the decline in the labor share in manufacturing. As in Autor et al., this decline is not driven by the unweighted within-component but by a decline in the covariance term. However, in French manufacturing, this lower covariance is closely connected to automation: firms adopting robots are large, expand further, and experience significant relative declines in labor share but did not have lower labor shares (or higher markups) at the baseline.

V. Conclusion

How firms change their production structure, employment, labor share, and productivity as they adopt automation technologies can help us understand the wide-ranging effects of automation. Nevertheless, firm-level effects do not

gains account for an increase in employment of 0.7 percent, whereas the second term accounts for a decline in employment of 3.9 percent. Note, however, that these computations do not incorporate any general equilibrium effects (whereby greater productivity in one industry increases employment in other industries). The online Appendix also documents that the cross-industry association between robot adoption and employment is negative.

⁹Though this is conditional on size, robot adopters in an industry have a slightly greater labor share (of about 2 percentage points); unconditionally, they have essentially the same labor share as nonadopters.

correspond to the overall impact of automation because firms that adopt such technologies reduce their costs and expand at the expense of competitors. In this paper, we estimate that French manufacturing firms that adopt robots reduce their labor share and share of production workers and increase their productivity, but also expand their operations and employment. Yet this is more than offset by significant declines in their competitors' employment. Overall, even though firms adopting robots expand their employment, the market-level implications of robot adoption are negative. We also show that robot adoption contributes to the decline in the manufacturing labor share by reducing the covariance between firm-level value added and labor share, and that this is because adopters are large and expand further as they experience sizable relative declines in their labor shares.

REFERENCES

- Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal* of Economic Perspectives 33 (2): 3–30.
- Acemoglu, Daron, and Pascual Restrepo. Forthcoming. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. Forthcoming. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics.*
- Aghion, Philippe, Céline Antonin, and Simon Bunel. 2019. "Artificial Intelligence, Growth and Employment: The Role of Policy." *Economie et Statistique* 510–12: 149–64.
- Bessen, James, Maarten Goos, Anna Salomons, and

Wiljan van den Berge. 2019. "Automatic Reaction—What Happens to Workers at Firms That Automate?" https://www.cpb.nl/sites/default/ files/omnidownload/CPB-Discussion-Paper-390-Automatic-Reaction-What-Happens-to-Workers-at-Firms-that-Automate.pdf.

- Bonfiglioli, A., R. Crinò, H. Fadinger, and G. Gancia. 2019. "Robot Imports and Firm Level Outcomes." https://drive.google.com/file/d/1eVZ2 6RqaaQKntiRfxgxseheDfOvgo92Z/view.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner. 2019. "The Adjustment of Labor Markets to Robots." https://sfndsn.github.io/downloads/ AdjustmentLaborRobots.pdf.
- Dinlersoz, Emin, and Zoltan Wolf. 2018. "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing." https://www2.census.gov/ces/wp/2018/ CES-WP-18-39.pdf.
- Dixon, Jay, Bryan Hong, and Lynn Wu. 2019. "The Employment Consequences of Robots: Firm-Level Evidence." http://content. tcmediasaffaires.com/LAF/lacom2019/robots. pdf.
- Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics* 100 (5): 753–68.
- Humlum, Anders. 2019. "Robot Adoption and Labor Market Dynamics." https:// economics.yale.edu/sites/default/files/ humlumjmp_111419.pdf.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *Quarterly Journal of Economics* 134 (3): 1343–404.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka. 2019. "Robots and Firms." VOX CEPR Policy Portal, July 1, 2019. https://voxeu.org/ article/robots-and-firms.