

DISCUSSION PAPER SERIES

DP14971

COMPETING WITH ROBOTS: FIRM- LEVEL EVIDENCE FROM FRANCE

Daron Acemoglu, Claire Lelarge and Pascual
Restrepo

INDUSTRIAL ORGANIZATION

LABOUR ECONOMICS



COMPETING WITH ROBOTS: FIRM-LEVEL EVIDENCE FROM FRANCE

Daron Acemoglu, Claire Lelarge and Pascual Restrepo

Discussion Paper DP14971

Published 30 June 2020

Submitted 29 June 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization
- Labour Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Daron Acemoglu, Claire Lelarge and Pascual Restrepo

COMPETING WITH ROBOTS: FIRM-LEVEL EVIDENCE FROM FRANCE

Abstract

Using several sources, we construct a data set of robot purchases by French manufacturing firms and study the firm-level implications of robot adoption. Out of 55,390 firms in our sample, 598 have adopted robots between 2010 and 2015, but these firms account for 20% of manufacturing employment and value added. Consistent with theory, robot adopters experience significant declines in labor share and the share of production workers in employment, and increases in value added and productivity. They expand their overall employment as well. However, this expansion comes at the expense of their competitors (as automation reduces their relative costs). We show that the overall impact of robot adoption on industry employment is negative. We further document that the impact of robots on overall labor share is greater than their firm-level effects because robot adopters are larger and grow faster than their competitors.

JEL Classification: J23, J24, L11

Keywords: automation, Competition, Labor Share, Manufacturing, productivity, reallocation, robots, Tasks

Daron Acemoglu - daron@mit.edu
MIT and CEPR

Claire Lelarge - claire.lelarge@gmail.com
Paris-Saclay University and CEPR

Pascual Restrepo - pascual@bu.edu
Boston University

Acknowledgements

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 NSF, 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).

Competing with Robots: Firm-Level Evidence from France*

Daron Acemoglu

MIT and NBER

Claire Lelarge

Paris-Saclay, CREST and CEPR

Pascual Restrepo

Boston University

January 21, 2020

Abstract

Using several sources, we construct a data set of robot purchases by French manufacturing firms and study the firm-level implications of robot adoption. Out of 55,390 firms in our sample, 598 have adopted robots between 2010 and 2015, but these firms account for 20% of manufacturing employment and value added. Consistent with theory, robot adopters experience significant declines in labor share and the share of production workers in employment, and increases in value added and productivity. They expand their overall employment as well. However, this expansion comes at the expense of their competitors (as automation reduces their relative costs). We show that the overall impact of robot adoption on industry employment is negative. We further document that the impact of robots on overall labor share is greater than their firm-level effects because robot adopters are larger and grow faster than their competitors.

Keywords: automation, competition, labor share, manufacturing, productivity, reallocation, robots, tasks.

JEL codes: J23, J24, L11.

*Acemoglu: MIT, daron@mit.edu. Lelarge: Paris-Saclay, claire.lelarge@universite-paris-saclay.fr. Restrepo: Boston University, 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 NSF, 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)

Introduction

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, 2019a). Acemoglu and Restrepo (2019b) estimate 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 impacting the production process and productivity.² But its interpretation is made 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 using data from France between 2010 and 2015. Consistent with our theoretical expectations (which are developed further in the 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 their 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 towards 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% decline in industry employment.

Finally, we use our data to study the decline in the French manufacturing labor share. As in Autor et al. (2019), 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),

¹Graetz and Michaels (2018) use variation across industries and countries and find lower labor share and higher productivity from robots, but negative effects only for unskilled workers. Aghion et al. (2019) find negative regional employment effects in France, while Dauth et al. (2019) estimate 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), Dixen et al. (2019), Bonfiglioli et al. (2019), Humlum (2019), and Koch et al. (2019).

³This aligns with Koch et al.'s (2019) findings from Spain.




FIGURE 1: Share of robot adopters among firms in different percentiles of the sales distribution within 4-digit industries. Shown for all industries, and industries with high *APR* and low *APR*.

but because firms adopting robots are large and expand further as they experience significant relative declines in their labor share.

1 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 adapted (purchased) industrial robots during this period 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% of our firms purchased robots in 2010–2015, these firms account for 20% of total manufacturing employment. Table A.1 in the 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 4-digit industries in our sample. For example, 13% of firms in the top 1% of the industry sales distribution adopted robots, while there is almost no robot adoption among firms below the 20th 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% higher in industries with greater *adjusted penetration of robots (APR)* in other European countries (shown in darker color).⁴

2 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 OLS across firms, denoted by f :

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

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 worker in 2010, as well as dummies for whether the firm is affiliated to a larger corporate group), 4-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 4-digit industries.

Table 1 reports our findings using unweighted (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% increase in value added from 2010 to 2015 (s.e.=0.030) as well as a 4.3 percentage point decline in the labor share (s.e.=0.009) and a 1.6 percentage point decline in the production worker share of employment (s.e.= 0.007). Value added per hour and revenue TFP also increase.⁵ Column 5 shows that, in contrast to market-level results in previous works, employment (total hours

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, 2019b). 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.

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

TABLE 1: Estimates of robot adoption on firm-level outcomes.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Δ log value added	Δ labor share	Δ production employment share	Δ log value added per hour	Δ log revenue TFP	Δ log employment (in hours)	Δ log mean hourly wage
<i>Panel A—Unweighted estimates</i>							
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
<i>Panel B—Employment-weighted estimates</i>							
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

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 to a larger corporate group), 4-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 Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within 4-digit industries are in parentheses.

of work) also increases in firms adopting robots—by 10.9% (s.e.= 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 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.

3 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 different than 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

$$\text{Adoption by competitors}_f = \sum_i m_{fi} \cdot \sum_{f' \neq f} s_{if'} \cdot \text{Robot}_{f'},$$

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%, 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 the references therein), which generate a pass-through elasticity of about 0.35. This is as expected since automation substitutes capital for labor.

TABLE 2: Estimates of robot adoption on competitors

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ log em- ployment (in hours)	Δ log value added	Δ labor share	Δ log em- ployment (in hours)	Δ log value added	Δ labor share
	<i>Unweighted estimates</i>			<i>Employment-weighted estimates</i>		
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

Notes— The sample consists of 55,388 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 to a larger corporate group), 4-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 Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within 4-digit industries are in parentheses.

where the first sum is over all 4-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 4-digit industries between a given firm and all robot adopters in the economy. The shares m_{fi} and $s_{if'}$ are constructed using sales data by firm and 4-digit industry from the fiscal files, which cover 85% of sales in our sample. We assume that small firms that are not in the fiscal files only sell in their main 4-digit industry. Because equation (1) includes 4-digit industry fixed effects, the spillovers are identified from the comparison of firms in the same main industry, but selling 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 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% decline in employment (s.e.=0.0107) and a 2.1% decline in value added (s.e.=0.0159) and, consistent with our theory in the 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 the own and the competitors' effects, we find that robots 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% decline in industry employment.⁷

4 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. (2019) 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. document that the decline in the labor share is driven by a reduction in the covariance term, and suggest 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. for French manufacturing between 2010 and 2015. As in their 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 non-adopters. Interestingly, while, analogously to the US, the labor share increases for firms not adopting robots, it declines for robot adopters. More importantly, about 80% of the decline in the covariance term is accounted for by the fact that robot adopters are larger from the outset (-2.81 pp) and expand (-0.14 pp) 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,

⁷The Appendix shows that this effect on employment is: $\beta_o \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f + \beta_c \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \sum_i m_{fi} \cdot (1 - s_{if})$. 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 gains account for an increase in employment of 0.7%, whereas the second term accounts for a decline in employment of 3.9%. Note, however, that these computations do not incorporate any general equilibrium effects (whereby greater productivity in one industry increases employment in other industries). The Appendix also documents that the cross-industry association between robot adoption and employment is negative.

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 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.

⁹Though conditional on size, robot adopters in an industry have slightly greater labor share (of about

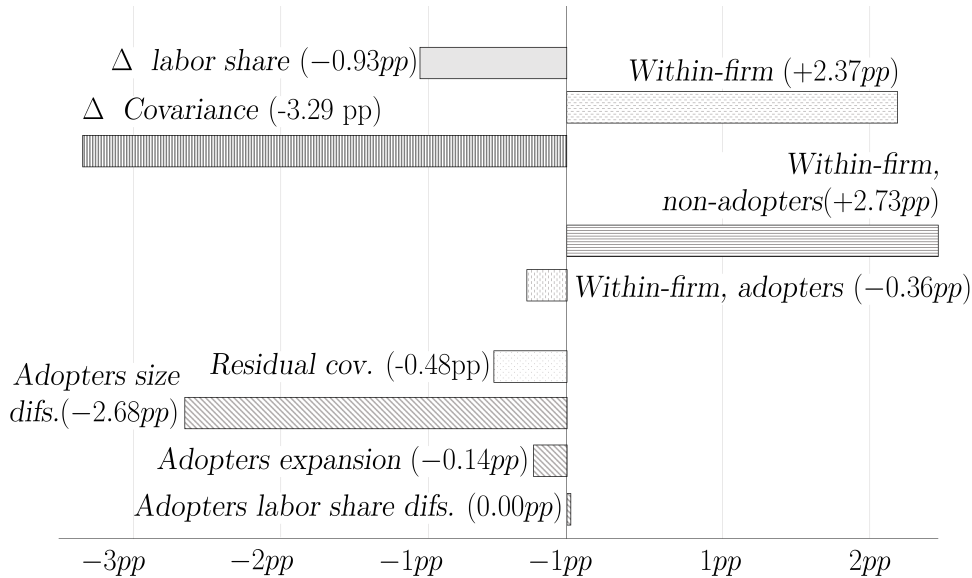


FIGURE 2: Changes in the labor share of French manufacturing industries for 2010–2015 decomposed as in Autor et al. (2019). The decomposition is extended to account for differences between robot adopters and non-adopters.

accounts for 20% 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.

5 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 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 (2 p.p.), unconditionally they have essentially the same labor share as non-adopters.

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, D., & Restrepo, P. (2019a) “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*. 33(2): 3–30.

Acemoglu, D., & Restrepo, P. (2019b) “Robots and Jobs: Evidence from US Labor Markets” in press, *Journal of Political Economy*.

Autor, D., Dorn, D., Katz, L.F., Patterson, C., & Reenen, J.V. (2019) “The Fall of the Labor Share and the Rise of Superstar Firms,” in press *Quarterly Journal of Economics*.

Aghion, P., Antonin, C., & Bunel, S. (2019) “Artificial Intelligence, Growth and Employment: the Role of Policy.” Paris.

Bessen, J. E., Goos, M., Salomons, A., & Van den Berge, W. (2019) “Automatic Reaction-What Happens to Workers at Firms that Automate?” Mimeo. Boston University.

Bonfiglioli, A. Crinò, R., Fadinger, H., & Gancia, G. (2019) “Robot Imports and Firm Level Outcomes.” QMUL.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2019) “The Adjustment of Labor Markets to Robots.” University of Würzburg.

Dinlersoz, E., & Wolf, Z. (2018) “Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing,” Census.

Dixen, J., Hong, B., & Wu, L (2019) “The Employment Consequences of Robots: Firm-level Evidence.” Statistics Canada.

Graetz, G., & Michaels, G. (2018) “Robots at Work,” *Review of Economics and Statistics*, 100(5): 753–768.

Humlum, A. (2019) “Robot Adoption and Labor Market Dynamics.” Princeton University.

Kline, P., Petkova, N., Williams, H., Zidar, O. (2019) “Who Profits from Patents? Rent-Sharing at Innovative Firms.” *The Quarterly Journal of Economics*, 134(3): 1343–1404.

Koch, M., Manuylov, I., & Smolka, M. (2019) “Robots and Firms.” Aarhus University.

Online Appendix for “Competing with Robots”

A. Data Description

Data on robot adopters: Our data sources and sample structure are summarized in Table [A.1](#).

TABLE A.1: Sample description

Size bins (emp. 2010)	All firms	Robot adopters			Sources of robot purchases data		
		Total number	Share of adopters in bin	Share hours among adopters	DGE survey	Customs data	SYMOP and fiscal files
> 5,000 workers	21	12	57.1%	78.0%	< 5	9	8
250 to 5,000 workers	1,114	169	15.2%	21.3%	8	95	82
10 to 250 workers	19,975	380	1.9%	4.2%	100	158	180
≤ 10 workers	34,280	37	0.1%	0.2%	11	13	20
Total	55,390	598	1.1%	19.8%	.	275	290

Notes—The table reports the composition of our sample for firms of different sizes. The Appendix describes the sources used.

Data on purchases of robots for 2010–2015 are assembled from the following sources:

- SYMOP—the French Association of Producers and Importers of Industrial Machinery—we obtained an extract of a subset of the firms who purchased domestically produced or imported industrial robots from SYMOP.
- A survey collected by the French Ministry of Industry (Direction Générale des Entreprises, or DGE), which includes information on robot purchases among small and medium enterprises. This survey sampled firms recognized as clients of SYMOP members.
- From French customs data, we obtained firm imports of industrial robots, which are coded under the NCS product code 84798950. All imports of industrial robots sourced from outside of the European Union are reported. Imports of robots from other countries in the European Union are not recorded at the transaction level. Instead, firms that imported at least 460,000 Euros worth of intermediate inputs and capital goods (including industrial robots) during that year from all sources in the European Union must report their purchases. The 460,000 Euros threshold is the cost of approximately four or five industrial robots. Thus, the customs? data miss firms that imported three or fewer robots from other European Union countries and small amounts of intermediate inputs (so as not to exceed the combined 460,000 Euros threshold), as well as firms that buy imported robots through subsidiaries of foreign robot producers.

- From French fiscal files, we identified firms that used an accelerated amortization scheme dedicated to industrial robots. Eligibility was restricted to small and medium enterprises and to transactions occurring between October 2013 and December 2015. We also incorporated public information on 40 small and medium enterprises which benefited from a subsidy program entitled “Robot Start PME” that was in effect between 2013 and 2016.

Firm accounting information: We obtained detailed accounting information for the firms in our sample from French fiscal files. In particular, we made use of two different files: the BRN (Bénéfices Réels Normaux) and the RSI (Régime Simplifié d’Imposition). The BRN contains the balance sheet of all firms in manufacturing with sales above 730,000 Euros. The RSI is the counterpart of the BRN for firms with sales below 730,000 Euros. Their union covers nearly the entire universe of French manufacturing firms.

Corporate groups: In our regressions, we control for a dummy for firms that belong to larger corporate groups. We obtained data on the ownership structure of firms from the LIFI files (Liaisons Financières Entre Sociétés) supplied by INSEE. This survey is complemented with information on ownership structure available from the DIANE (BvDEP) files, which are constructed using the annual mandatory reports to commercial courts and the register of French firms. Using these data, we constructed dummies for firms that are affiliates of larger corporate groups. In regressions we also control for a dummy that indicates when observations in the fiscal files are an aggregate of several affiliates of a corporate group.

Detailed sales information: The data on sales by firm across 4-digit industries used in the construction of the measure of adoption of robots by competitors come from these French fiscal files as well. In particular, we use the FARE files (Fichier Approché des Résultats d’Esane), which contain sales by firm and industry for over 85% of the sales in our sample. The FARE does not break down sales by industry for small firms, and so we assume that small firms only sell in their assigned 4-digit industry. The FARE also contain data on total sales by industry, which we use to compute the weights s_{if} used in our formula for adoption among competitors.

Data on firm exports and prices: We have detailed data on firm exports by totals and unit values for each NC8 product category. The data come from the French Customs and cover every transaction between a French firm and a foreign importer located in the European Union.

Worker-level information: We incorporate information from the French matched employer-employee administrative dataset (Déclarations Annuelles des Données Sociales, DADS) to retrieve worker-level information on occupation, wages, and hours worked.

Variable definitions: We constructed value added at the firm level as sales minus expenditure on intermediates. For employment, we have data on the count of employees, total hours of work, and full-time equivalent workers. In the main text we focus on total hours of work as our main measure of employment, value added per hour worked as our main measure of labor productivity, and mean hourly wage for the average wage rate at the firm. To measure wages, we use the wage bill of the firm, which accounts for all wage payments to workers. We obtained very similar results using total compensation, which also includes payroll taxes and other benefits.

We define production workers using the DADS data as those employed in unskilled industrial jobs (categories 67 and 68 in the INSEE classification of professions).

We measure changes in (revenue) TFP for the 2010–2015 period as

$$\Delta \ln \text{TFP}_f = \Delta \ln y_f - \lambda_f^\ell \cdot \Delta \ln \ell_f - \lambda_f^m \cdot \Delta \ln m_f - (1 - \lambda_f^\ell - \lambda_f^m) \cdot \Delta \ln k_f.$$

Here, λ_f^ℓ and λ_f^m denote the shares of wages and intermediates in revenue, respectively. These shares are measured for each firm in 2010. Alternative measures using detailed industry shares instead of firm-level ones yield very similar results. In addition, $\Delta \ln y_f$ is the percent change in sales, $\Delta \ln \ell_f$ is the percent change in hours, $\Delta \ln m_f$ is the percent change in materials, and $\Delta \ln k_f$ denotes the percent change in the capital stock during 2010–2015. Since we do not have data on material prices, we assume that these are common across firms.

For exporting firms, we also have information on prices, which enables us to investigate whether productivity changes are related to price changes or changes in physical productivity. In particular, we construct a price index for an exporting firm as follows:

$$\Delta \ln p_f = \sum_{\omega} e_{f\omega} \cdot \Delta \ln p_{\omega f},$$

where the sum is taken over all NC8 product categories ω , $e_{f\omega}$ denotes the export share of ω in firm f , and $\Delta \ln p_{\omega f}$ is the observed change in unit values of the exports of firm f in product category ω .

B. Robustness Checks

This section provides additional own-firm estimates of robot adoption and robustness checks for the estimates in the main text. Table A.2 presents estimates for additional outcomes, including log sales and the share of wages in sales. These results show that the results on Table 1 in the main text hold when we focus on sales rather than value added. Columns 3–5 present results for additional measures of labor productivity, including sales per hour, sales per worker, and value added per worker (as opposed to the per hour measure in the main text). Finally, columns 6 and 7 present results for the percent change in the number of employees (not hours) and the number of production workers (as opposed to their share).

TABLE A.2: Estimates of robot adoption on additional firm-level outcomes.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Δ log sales	Δ labor share in sales	Δ log sales per hour	Δ log sales per worker	Δ log value added per worker	Δ log employment (in total employees)	Δ log employment production workers
<i>Panel A—Unweighted estimates</i>							
Robot adopter	0.142 (0.021)	-0.007 (0.002)	0.032 (0.012)	0.062 (0.018)	0.123 (0.025)	0.078 (0.012)	0.046 (0.032)
R^2	0.064	0.092	0.142	0.079	0.130	0.058	0.024
<i>Panel B—Employment-weighted estimates</i>							
Robot adopter	0.121 (0.019)	-0.012 (0.003)	0.066 (0.021)	0.077 (0.021)	0.050 (0.028)	0.044 (0.014)	-0.084 (0.090)
R^2	0.196	0.164	0.237	0.202	0.277	0.174	0.144

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, and dummies for whether the firm belongs to a larger corporate group), 4-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 Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within 4-digit industries are in parentheses.

As mentioned in the main text, the increase in labor productivity and TFP (in the unweighted specification) are not driven by price increases among firms adopting robots, but reflect changes in quantities (physical productivity). Table A.3 provides evidence in support of this claim. The table uses the sample of exporters to estimate the association between robot adoption and changes in export prices. We provide estimates using different weighting schemes (unweighted, weighted by employment hours as in the main text, or weighting by firm exports) and controlling for 2-digit or 4-digit industry dummies. The sample now is much smaller, and the estimates are less precise. But overall, we find uniformly negative point estimates, which suggest that firms that adopt robots reduce prices from 1% to 5.7% (using the estimates with 4-digit industry fixed effects in columns 2 and 4).

TABLE A.3: Robot adoption and firm-level export prices. Estimates for the subset of exporters.

	Dependent variable: Δ log export price index					
	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Unweighted estimates</i>		<i>Employment-weighted</i>		<i>Export-weighted</i>	
Robot adopter	-0.009 (0.021)	-0.009 (0.021)	-0.066 (0.028)	-0.057 (0.028)	-0.064 (0.048)	-0.051 (0.052)
R^2	0.058	0.092	0.178	0.229	0.242	0.301

Notes— The sample consists of 6,614 firms for which we have data on export prices, of which 372 are robot adopters. Panel A presents unweighted estimates. Panel B presents estimates weighting each firm by its employment (in hours) in 2010. Panel C presents estimates weighting each firm by its exports in 2010. All specifications control for baseline firm characteristics (log employment and log value added per worker in 2010, dummies for whether the firm belongs to a larger corporate group, the sales percentile of the firm in its main 4-digit industry, the share of production workers, and the log of capital per worker), and fixed effects for the commuting zone that houses each firm’s largest establishment. Also, columns 1, 3, 5 control for 2-digit industry-fixed effects; whereas columns 2, 4, 6 control for 4-digit industry-fixed effects. The Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within 4-digit industries are in parentheses.

Finally, Table A.4 shows that the findings in Table 1 in the text are robust to the inclusion of additional covariates. Specifically, we control for dummies for firms in the top 0.1%, top 1%, top 5%, top 10%, top 20% and top 40% of sales in each 4-digit industry as well as log capital stock per worker and the share of production workers in 2010.

TABLE A.4: Robustness checks for estimates of robot adoption on firm-level outcomes. Includes additional covariates.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Δ log value added	Δ labor share	Δ production employment share	Δ log value added per hour	Δ log revenue TFP	Δ log employment (in hours)	Δ log mean hourly wage
	<i>Panel A—Unweighted estimates</i>						
Robot adopter	0.168 (0.024)	-0.035 (0.008)	-0.014 (0.006)	0.079 (0.017)	0.012 (0.006)	0.089 (0.017)	0.008 (0.004)
R^2	0.094	0.166	0.236	0.224	0.207	0.101	0.031
	<i>Panel B—Employment-weighted estimates</i>						
Robot adopter	0.086 (0.026)	-0.023 (0.011)	-0.004 (0.005)	0.029 (0.027)	-0.016 (0.014)	0.057 (0.017)	-0.010 (0.007)
R^2	0.226	0.285	0.231	0.333	0.307	0.192	0.141

Notes— The sample consists of 55,359 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, dummies for whether the firm belongs to a larger corporate group, the sales percentile of the firm in its main 4-digit industry, the share of production workers, and the log of capital per worker), 4-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 Appendix describes the construction of all variables used as outcomes. Standard errors robust to heteroskedasticity and correlation within 4-digit industries are in parentheses.

C. Market-Level Effects

In this section of the Appendix, we aggregate the estimates from Table 2 to obtain market-level effects. Recall that the estimating equation for the models in Table 2 is

$$\Delta \ln \ell_f = \beta_o \cdot \text{Robot}_f + \beta_c \cdot \text{Adoption by competitors}_f + \gamma \cdot X_f + \alpha_{i(f)} + \delta_{c(f)} + \varepsilon_f.$$

We now show that the contribution of robot adoption to overall employment can be approximated (to the first order) as

$$(A.1) \quad \Delta \ln \ell \approx \beta_o \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f + \beta_c \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot \sum_{f' \neq f} s_{if'} \cdot \frac{y_f}{\ell_f} / \frac{y_{f'}}{\ell_{f'}} \right),$$

where the sum is taken over all 4-digit industries and m_{fi} is the share of firm f 's sales that are in industry i , while $s_{if'}$ is the share of industry i 's total sales accounted for by firm f' . Under the additional assumption that firms have similar baseline labor shares, this expression can be further simplified to

$$(A.2) \quad \Delta \ln \ell \approx \beta_o \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f + \beta_c \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot (1 - s_{if}) \right).$$

The numbers given in the main text are obtained from this equation.

We now provide more details on how these numbers are obtained. With a first-order approximation, the change in manufacturing employment can be expressed as

$$\Delta \ln \ell \approx \sum_f \frac{\ell_f}{\ell} \Delta \ln \ell_f.$$

The contribution of robot adoption to aggregate employment is therefore:¹⁰

$$\begin{aligned} \sum_f \frac{\ell_f}{\ell} \Delta \ln \ell_f &= \sum_f \frac{\ell_f}{\ell} (\beta_o \cdot \text{Robot}_f + \beta_c \cdot \text{Adoption by competitors}_f) \\ &= \beta_o \sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f + \beta_c \sum_f \frac{\ell_f}{\ell} \cdot \sum_i m_{fi} \cdot \sum_{f' \neq f} s_{if'} \cdot \text{Robot}_{f'}. \end{aligned}$$

Changing the order of summation, the term multiplying β_c can be expressed as

$$\sum_f \text{Robot}_f \cdot \sum_i s_{if} \cdot \sum_{f' \neq f} m_{f'i} \cdot \frac{\ell_{f'}}{\ell}.$$

¹⁰Note, however, that this computation ignores any general equilibrium effects from robot adoption that can lead to an expansion or contraction of overall manufacturing employment. Such general equilibrium effects cannot be identified with our methodology (or with other approaches based on cross-industry comparisons). Consequently, the estimate of -1% below should be compared with industry-level estimates of the impact of robot adoption on employment.

Multiplying and dividing by ℓ_f , and then rearranging terms, we obtain

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \sum_i s_{if} \cdot \sum_{f' \neq f} m_{f'i} \cdot \frac{\ell_{f'}}{\ell_f}.$$

Denoting the sales of firm f by y_f and using the definitions of s_{if} and $m_{f'i}$, we can write the previous expression as

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \sum_i \frac{y_{fi}}{y_i} \cdot \sum_{f' \neq f} \frac{y_{f'i}}{y_{f'}} \cdot \frac{\ell_{f'}}{\ell_f}.$$

Dividing and multiplying by y_f , this is equivalent to

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \sum_i \frac{y_{fi}}{y_f} \cdot \sum_{f' \neq f} \frac{y_{f'i}}{y_i} \cdot \frac{y_f}{y_{f'}} \cdot \frac{\ell_{f'}}{\ell_f}.$$

Using the definition of $s_{if'}$ and m_{fi} one more time and regrouping terms, we obtain

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot \sum_{f' \neq f} s_{f'i} \cdot \frac{y_f}{\ell_f} / \frac{y_{f'}}{\ell_{f'}} \right).$$

In the special case where firms have similar baseline labor productivity (or equivalently, similar levels of baseline labor shares if wages are common across firms), we would also have $\frac{y_f}{\ell_f} = \frac{y_{f'}}{\ell_{f'}}$, and this can be further simplified to the simpler expression used in the main text

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot (1 - s_{if}) \right).$$

In our data, we have

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f = 0.20,$$

and

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot (1 - s_{if}) \right) = 0.156.$$

Moreover,

$$\sum_f \frac{\ell_f}{\ell} \cdot \text{Robot}_f \cdot \left(\sum_i m_{fi} \cdot \sum_{f' \neq f} s_{f'i} \cdot \frac{y_f}{\ell_f} / \frac{y_{f'}}{\ell_{f'}} \right) = 0.193.$$

Using the estimates from the weighted specification for β_o and β_c (reproduced in column 1 of Table A.6), we estimate aggregate declines in employment associated with robot adoption in the range of 3.2%–4.1%. Alternatively, if we use the specification including 2-digit industry dummies instead of 4-digit industry dummies, the estimates in column 2 of Table A.6 imply somewhat smaller aggregate employment effects, of about -1.2%.

D. Industry-Level Estimates

An alternative strategy to assess the aggregate implications of robot adoption is to exploit only industry-level variation in robot adoption (and is thus different from the approach used in the main text). In particular, we start by estimating an industry-level variant of equation (1) in the main text:

$$(A.3) \quad \Delta \ln \ell_i = \beta_m \cdot \text{Robot adoption}_i + \varepsilon_i,$$

where Robot adoption_i is the employment-weighted share for firms adopting robots in industry i . We focus on industry employment (total hours among the firms in our sample whose main industry is i) as the left-hand side variable, and as in the text, on estimates weighted by industry employment, which are more informative about aggregate effects.

Table A.5 shows that robot adoption is associated with a robust decline in employment across industries. Columns 1 and 2 present estimates of equation (A.3) for 240 4-digit industries. Column 1 shows the unconditional relationship (without any covariates). The estimate in this column suggests that a 20 percentage point increase in robot adoption in an industry is associated with a 2.56% decline in industry employment. Column 2 controls for 2-digit industry fixed effects and leads to somewhat smaller estimates. Now the same 20 percentage point increase in robot adoption is associated with a decline in industry employment of 1.44%. Finally, columns 3 and 4 reproduce the same estimates but for 95 3-digit industries, and show that a 20 percentage point increase in robot adoption among firms in an industry is associated with a decline in employment of 1.96%.

TABLE A.5: Industry-level estimates of robot-adoption on employment.

	Dependent variable: $\Delta \log$ employment (hours)			
	4-digit industries		3-digit industries	
	(1)	(2)	(3)	(4)
Robot adoption	-0.128 (0.081)	-0.072 (0.042)	-0.144 (0.088)	-0.098 (0.058)
R ²	0.121	0.559	0.178	0.713
<i>Covariates:</i>				
2-digit industry fixed effects		✓		✓

Notes—The sample consists of $N = 240$ 4-digit industries (columns 1–2) and $N = 88$ 3-digit industries (columns 3–4). All models weight industries by their employment (in hours) in 2010. Columns 2 and 4 control for 2-digit industry fixed effects. Standard errors robust to heteroskedasticity are in parentheses.

We can further decompose these negative industry-level estimates into own-firm and

TABLE A.6: Additional estimates of spillovers on employment of other firms.

	Dependent variable: Δ log employment (hours)							
	Adoption among competitors defined as in the main text		Adoption among competitors defined as employment-weighted average among firms in the same 4-digit industry			Adoption among competitors defined as employment-weighted average among firms in the same 3-digit industry		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot adoption by competitors	-0.250 (0.107)	-0.110 (0.046)	-0.117 (0.081)	-0.075 (0.042)	-0.061 (0.034)	-0.128 (0.088)	-0.102 (0.052)	-0.083 (0.046)
Robot adopter	0.035 (0.022)	0.046 (0.017)	0.035 (0.020)	0.035 (0.020)	0.054 (0.016)	0.033 (0.021)	0.034 (0.019)	0.052 (0.016)
R ²	0.190	0.155	0.005	0.007	0.154	0.006	0.008	0.152
<i>Covariates:</i>								
Firm covariates	✓	✓			✓			✓
4-digit industry fixed effects	✓							
2-digit industry fixed effects		✓		✓	✓		✓	✓

Notes—The sample consists of $N = 55,388$ firms, of which 598 are robot adopters. All models weight firms by their employment (in hours) in 2010. Columns 3–5 present estimates for the adoption of robots by firms in the same 4-digit industry. Columns 6–8 present estimates for the adoption of robots by firms in all the 3-digit industries in which a firm sells some of its products (weighted by share sales). The set of industry-fixed effects used in each specification is indicated at the bottom rows. Additional covariates in column 1–2, 5 and 8 include: baseline firm characteristics (log employment and log value added per worker in 2010, as well as dummies for whether the firm is affiliated to a larger corporate group), and fixed effects for the commuting zone that houses each firm’s largest establishment. Standard errors robust to heteroskedasticity and correlation within 4-digit (and 3-digit industries in columns 6–8) industries are in parentheses.

spillover effects by estimating the following variant of equation (1) at the firm level:

$$(A.4) \quad \Delta \ln \ell_{if} = \beta_o \cdot \text{Robot adopter}_f + \beta_c \cdot \text{Robot adoption}_i + \varepsilon_{if}.$$

Here, β_c captures spillovers of robot adoption on other firms in the same industry. For this particular specification of spillovers, the estimate of $\beta_c + \beta_o$ corresponds to the industry-level estimate of robots on employment, at least to a first-order approximation.

Columns 3–8 in Table A.6 presents estimates of equation (A.4). Column 3 presents estimates of (A.4) focusing on spillovers among firms in the same 4-digit industry and without any additional covariate. The estimate in column 3 indicates that a 10 percentage point increase in adoption is associated with a 1.17% decline in employment for firms that do not adopt robots and a 0.35% increase in employment at firms that do. The net result is a reduction in employment of 0.82% (s.e.=0.0078), which is similar to the industry-level estimate in Table A.5. Column 4 adds a full set of 2-digit industry dummies as covariates and column 5 further includes the firm-level covariates used in the main text, which lead to more precise estimates of the spillover effect. Finally, columns 6–8 reproduce the same estimates but focusing on spillovers among firms in the same 3-digit industries.

E. Decomposing Changes in the Labor Share

This section provides the details for the decomposition used in Figure 2. Following Autor et al. (2019), we decompose changes in the labor share of industry i as

$$(A.5) \quad \Delta\lambda_i^\ell = \Delta\bar{\lambda}_i^\ell + \Delta \sum_f (\lambda_f^\ell - \bar{\lambda}_i^\ell) \cdot (s_{if}^v - \bar{s}_i^v),$$

where λ_i^ℓ is the labor share in industry i , 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 correspond to unweighted averages of these terms among firms in the industry. The first term in the above decomposition is what Autor et al. (2019) term the *within component* (which is the *unweighted mean* change). The second term is a *covariance term* which accounts for reallocation to firms with lower labor shares, reallocation to firms with declining labor shares, and larger reductions of the labor share among larger firms. We use a balanced panel of firms and ignore entry and exit.

We can explore the contribution to changes in the aggregate labor share arising from robot adoption as follows. Let \mathcal{R}_i be the set of robot adopters in an industry and \mathcal{N}_i be the remaining set of firms. Also, denote the number of adopters by R_i , the number of non-adopters by N_i , and the total number of firms in the industry by F_i . Finally, for a set of firms, \mathcal{X} , define the following unweighted averages

$$\bar{\lambda}_{\mathcal{X}}^\ell = \frac{1}{|\mathcal{X}|} \sum_{f \in \mathcal{X}} \lambda_f^\ell \quad \bar{s}_{\mathcal{X}}^v = \frac{1}{|\mathcal{X}|} \sum_{f \in \mathcal{X}} s_{if}^v.$$

We can decompose the within-firm change component in equation (A.5) as:

$$\Delta\bar{\lambda}_i^\ell = \frac{R_i}{F_i} \Delta\bar{\lambda}_{\mathcal{R}_i}^\ell + \frac{N_i}{F_i} \Delta\bar{\lambda}_{\mathcal{N}_i}^\ell.$$

The first term accounts for the within-firm change in the labor share among adopters. The second term accounts for the within-firm change in the labor share among non-adopters. (Both of those are still unweighted following Autor et al., 2019).

We next decompose the superstar effect in (A.5) as:

$$\begin{aligned} \Delta \sum_f (\lambda_f^\ell - \bar{\lambda}_i^\ell) \cdot (s_{if}^v - \bar{s}_i^v) &= R_i \cdot \Delta(\bar{\lambda}_{\mathcal{R}_i}^\ell - \bar{\lambda}_i^\ell) \cdot (\bar{s}_{\mathcal{R}_i}^v - \bar{s}_i^v) + N_i \cdot \Delta(\bar{\lambda}_{\mathcal{N}_i}^\ell - \bar{\lambda}_i^\ell) \cdot (\bar{s}_{\mathcal{N}_i}^v - \bar{s}_i^v) \\ &\quad + \Delta \sum_{f \in \mathcal{R}_i} (\lambda_f^\ell - \bar{\lambda}_{\mathcal{R}_i}^\ell) \cdot (s_{if}^v - \bar{s}_{\mathcal{R}_i}^v) + \Delta \sum_{f \in \mathcal{N}_i} (\lambda_f^\ell - \bar{\lambda}_{\mathcal{N}_i}^\ell) \cdot (s_{if}^v - \bar{s}_{\mathcal{N}_i}^v). \end{aligned}$$

The first line in the above equation captures how differences between adopters and non-adopters contribute to changes in the covariance term. The second line captures the residual changes in the covariance term that are unrelated to automation (for example, due

to the changes in the allocation of economic activity within robot adopters and separately within non-robot adopters).

Finally, we can further decompose the contribution of robot adoption to the change in the covariance term in three terms:

$$\begin{aligned}
R_i \cdot \Delta(\bar{\lambda}_{\mathcal{R}_i}^\ell - \bar{\lambda}_i^\ell) \cdot (\bar{s}_{\mathcal{R}_i}^v - \bar{s}_i^v) + N_i \cdot \Delta(\bar{\lambda}_{\mathcal{N}_i}^\ell - \bar{\lambda}_i^\ell) \cdot (\bar{s}_{\mathcal{N}_i}^v - \bar{s}_i^v) &= \left(s_{\mathcal{R}_i} - \frac{R_i}{F_i} \right) \times \Delta(\bar{\lambda}_{\mathcal{R}_i}^\ell - \bar{\lambda}_{\mathcal{N}_i}^\ell) \\
&+ (\bar{\lambda}_{\mathcal{R}_i}^\ell - \bar{\lambda}_{\mathcal{N}_i}^\ell) \times \Delta s_{\mathcal{R}_i} \\
&+ \Delta(\bar{\lambda}_{\mathcal{R}_i}^\ell - \bar{\lambda}_{\mathcal{N}_i}^\ell) \times \Delta s_{\mathcal{R}_i},
\end{aligned}$$

where $s_{\mathcal{R}_i}$ denotes the share of value added accounted for by adopters. These terms capture three potential mechanisms via which industrial automation can lower the covariance between value added and labor shares across firms in an industry. The first term accounts for the fact that robot adopters are larger to begin with. Because Autor et al.'s (2019) within component is unweighted, the covariance between value added and the labor share also includes the size difference between adopters and non-adopters. In particular, this covariance declines as adopters automate and reduce their labor share relative to non-adopters. The second term captures the possibility that adopters had a different labor share to begin with. The third term captures the fact that adopters increase their share of value added in their industry as they simultaneously experiencing a reduction in their labor share.

Figure 2 in the main text implements this decomposition using data from French manufacturing firms for 2010–2015. We first obtain the components for each 4-digit industry, and we then aggregate across industries using their average share of value added during this period as weights.

F. A Model of Automation and Reallocation across Firms

This section presents a model that builds and extends on Acemoglu and Restrepo (2019b). Our aim is to clarify the conditions under which robot adoption will be associated with increases in own-firm employment but declines in aggregate employment.

Consider an economy with a single industry consisting of multiple firms with imperfectly substitutable products. In particular, industry output is

$$y = \left(\sum_f \alpha_f^{\frac{1}{\sigma}} y_f^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where y_f is the output produced by firm f and $\sigma > 1$ is the elasticity of substitution across firms.

Firm production is given by

$$y_f = A_f \left(\frac{k_f}{\theta_f} \right)^{\theta_f} \left(\frac{\ell_f}{1-\theta_f} \right)^{1-\theta_f},$$

where θ_f denotes the extent of automation at firm f . We think of improvements in industrial automation technologies as generating an increase in θ_f for the firms that adopt it.

Capital and labor are perfectly mobile across firms. Capital is produced using the final good at a cost $\Gamma_k \cdot k^{1+1/\varepsilon_k}/(1+1/\varepsilon_k)$. Labor is supplied by households, who have quasi-linear preferences and face a disutility from working given by $\Gamma_\ell \cdot \ell^{1+1/\varepsilon_\ell}/(1+1/\varepsilon_\ell)$. These assumptions ensure that a competitive equilibrium maximizes

$$\begin{aligned} & \max_{k, \ell, \{k_f, \ell_f\}_f} \left(\sum_f \alpha_f^{\frac{1}{\sigma}} y_f^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - \frac{\Gamma_k}{1+1/\varepsilon_k} k^{1+1/\varepsilon_k} - \frac{\Gamma_\ell}{1+1/\varepsilon_\ell} \ell^{1+1/\varepsilon_\ell} \\ \text{subject to: } & y_f = A_f \left(\frac{k_f}{\theta_f} \right)^{\theta_f} \left(\frac{\ell_f}{1-\theta_f} \right)^{1-\theta_f} \\ & \sum_f k_f = k \text{ and } \sum_f \ell_f = \ell. \end{aligned}$$

Therefore, an equilibrium is given by factor prices $\{w, r\}$, an allocation $\{k_f, \ell_f\}_f$, and aggregates $\{y, k, \ell\}$ such that:

- the ideal-price index condition holds

$$(A.6) \quad 1 = \sum_f \alpha_f \cdot \left(\frac{r^{\theta_f} w^{1-\theta_f}}{A_f} \right)^{1-\sigma};$$

- aggregate labor demand satisfies

$$(A.7) \quad w\ell = \sum_f (1 - \theta_f) \cdot y \cdot \alpha_f \cdot \left(\frac{r^{\theta_f} w^{1-\theta_f}}{A_f} \right)^{1-\sigma};$$

- aggregate capital demand satisfies

$$(A.8) \quad rk = \sum_f \theta_f \cdot y \cdot \alpha_f \cdot \left(\frac{r^{\theta_f} w^{1-\theta_f}}{A_f} \right)^{1-\sigma};$$

- aggregate labor supply satisfies

$$(A.9) \quad \ell = (w/\Gamma_\ell)^{\varepsilon_\ell};$$

- aggregate capital supply satisfies

$$(A.10) \quad k = (r/\Gamma_k)^{\varepsilon_k};$$

Let w be the equilibrium wage and r the rate at which capital is rented to firms. To ensure that automation technologies are adopted, we assume that for all firms we have

$$\pi \equiv \ln\left(\frac{w}{r}\right) > 0.$$

This equation implies that producing automated tasks with industrial automation technologies is cheaper than producing it with labor. Hence, whenever it can, a firm will adopt automation technologies and this would reduce its costs.

PROPOSITION A1: Suppose that $\theta_f = \theta$ and technological improvement increase θ_f for firm f by $d\theta_f > 0$.

- Own-firm employment changes by

$$(A.11) \quad d\ln \ell_f = \left(-\frac{1}{1-\theta} + (\sigma-1) \cdot \pi \right) d\theta_f + m,$$

where m is common to all firms in the industry.

- Aggregate employment changes by

$$(A.12) \quad d\ln \ell = \frac{\varepsilon_\ell}{\theta\varepsilon_\ell + (1-\theta)\varepsilon_k + 1} \left(-\frac{1}{1-\theta} + (1+\varepsilon_k) \cdot \pi \right) \sum_f s_f^\ell \cdot d\theta_f,$$

where s_f^ℓ denotes the share of employment accounted for by firm f .

- The labor share of firm f declines by $d\theta_f$ and the labor share of other firms remains constant.
- A necessary and sufficient condition for relative employment in firm f to increase and for industry employment to decline is

$$\frac{1}{(1 + \varepsilon_k) \cdot (1 - \theta)} > \pi > \frac{1}{(\sigma - 1) \cdot (1 - \theta)}.$$

PROOF. First, note that labor demand in firm f satisfies

$$w\ell_f = (1 - \theta_f) \cdot y_f \cdot p_f = (1 - \theta_f) \cdot y \cdot \alpha_f \cdot \left(\frac{r^{\theta_f} w^{1-\theta_f}}{A_f} \right)^{1-\sigma}.$$

Taking a log derivative of this equation around an equilibrium with $\theta_f = \theta$ yields

$$d \ln \ell_f = \frac{1}{1 - \theta} (-1 + (\sigma - 1) \cdot (1 - \theta) \cdot \pi) d\theta_f \\ \underbrace{-d \ln w + d \ln y + (1 - \sigma)\theta d \ln r + (1 - \sigma)(1 - \theta)d \ln w}_{\equiv m},$$

which coincides with the formula in equation (A.11).

For aggregates, we can take a log derivative of (A.6), (A.7), (A.8), (A.9) and (A.10) to obtain a system of equations in $\{d \ln \ell, d \ln k, d \ln w, d \ln r, d \ln y\}$. When $\theta_f = \theta$ the system simplifies to

$$(1 - \theta)d \ln w + \theta d \ln r = \pi \sum_f s_f^\ell d\theta_f \\ d \ln w + d \ln \ell = d \ln y - \frac{1}{1 - \theta} \sum_f s_f^\ell d\theta_f \\ d \ln r + d \ln k = d \ln y + \frac{1}{\theta} \sum_f s_f^\ell d\theta_f \\ d \ln \ell = \varepsilon_\ell d \ln w \\ d \ln k = \varepsilon_k d \ln r$$

Solving this system of equations yields the formula in equation (A.12) and establishes the second part of the proposition.

The third part follows from the fact that labor share in firm f' is simply $\theta_{f'}$, and thus the labor share in firm f declines and there is no impact on the labor shares in other firms.

The fourth part follows readily from equations (A.11) and (A.12). \square \blacksquare

The assumption that initially $\theta_f = \theta$ is imposed for simplicity. If, for example, θ_f

and $d\theta_f$ were positively correlated, then economic activity would be reallocated to firms that start with the lower labor share and there would be a larger decline in aggregate employment. As noted above, in French manufacturing, there is little baseline difference in the labor shares of adopters and non-adopters, and thus the equations presented here, where θ_f and $d\theta_f$ are initially uncorrelated, appear to be a good approximation in this context.