What Are the Labor and Product Market Effects of Automation? New Evidence from France*

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

December 2021

Abstract

We use comprehensive micro data in the French manufacturing sector between 1995 and 2017 to document the effects of automation technologies on employment, sales, prices, wages, and the labor share. 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. We find that automation leads to higher sales, higher profits, and lower consumer prices, while it leaves wages, the labor share and within-firm wage inequality unchanged. Consistent with the importance of business-stealing across countries, the estimated industry-level employment response to automation is stronger 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 labor demand. 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 Ajay Agrawal, David Autor, Francesco Caselli, Matt Gentzkow, Anders Humlum, 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, Bar-Ilan University, Berkeley, CIFAR, the College de France, Dartmouth, Georgetown, the Hebrew University of Jerusalem, the London School of Economics, the NBER's Economics of Artificial Intelligence Conference, the NBER Summer Institute, the OFCE, Queen Mary, Tel Aviv University, the University of Copenhagen, the University of Munich, and the WEFI webminar. We thank Eleonore Richard for excellent research assistance.

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: automation may induce productivity gains, increase market demand and the scale of production, and in turn increase labor demand (e.g., Bowen and Mangum (1966), Zeira (1998), Autor (2015), Acemoglu and Restrepo (2018)). 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 policies 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 provided 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 labor and product markets.

In this paper, we leverage micro data on the population of firms and plants in the French manufacturing sector between 1995 and 2017 to provide a unified analysis of the effects of automating technologies and machines[3] on employment, sales, prices, wages, the labor share, and profits at several levels of aggregation — across plants, firms, and industries. We use three complementary

[1] A range of policies has been discussed in this respect, including retraining programs, redistribution policies, as well 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).

[2] For example, Chiacchio et al. (2018), Webb (2019) and Acemoglu and Restrepo (2019) find evidence in line with the view that various forms automation reduce labor demand, while Graetz and Michaels (2018), Mann and Puttmann (2018) and Klenert et al. (2020) document positive employment effects. At the end of this section we discuss the emerging literature on the firm-level effects of robotization, which has grown in recent years and ran parallel to this paper.

[3] We call this set of technologies and machines "automation" in the remainder of the paper.

measures of automation technologies, respectively based on the balance sheet values of industrial equipment, records of the usage of electro-motive force, and imports of automating industrial machines.[4]

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 our measures of automation and 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 firms or across plants in the same 5-digit industry. In contrast with the common view that firms that rely on automation technologies reduce their labor force, we find that firm-level and plant-level employment increases after automation, including for unskilled industrial workers. The elasticity of employment to automation is positive, and it is amplified over time. In line with the hypothesized productivity effect, we find that sales and exports increase when firms automate, while export prices and competitors' employment fall. We find no evidence that automation has an impact on firm-level wage inequality, although it does increase the rate of job creation and destruction within the firm.[5]

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. 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. Confounding shocks would need to occur exactly at the

[4]The balance sheet measure guarantees broad coverage and is available for the population of firms, but it does not provide an explicit list of all technologies and machines. 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. The third measure allows us to draw an exhaustive list of the machines we consider, using detailed product classifications from customs, although this approach can only be applied to the subset of importing firms. See Section 2 for a complete discussion.

[5]Distributional effects may occur within each skill group, depending on the set of tasks performed across detailed occupations, but we do not find that the firm-level wage distribution is affected by the adoption of automation technologies. Other studies studying the effects of technological change have emphasized the importance of within-skill heterogeneity. For example, Hummels et al. (2014) documented that the distributional effects of offshoring occurred primarily within skill groups, rather than across.

same time as the increase in automation. Nonetheless, absent a quasi-experiment potential concerns over omitted factors cannot be fully 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 at the time 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 implement this research design for the subset of firms that import machines from abroad. Identification stems from changes over time in the productivity of foreign suppliers of automation technologies, which French firms are differentially exposed to through pre-determined importer-supplier relationships. This identification strategy approximates 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 across suppliers. For identification, we use productivity shocks across foreign suppliers, which we infer using trade flows across trading partners excluding France. Specifically, we use HS6-level shocks measured in EU countries (except France) and Switzerland as instruments for the adoption of machines in France.

The exclusion restriction underlying this design is that firms linked to increasingly productive suppliers should not have unobservable features affecting our outcomes. 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 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 4-digit product-by-year and 2-digit industry-by-year fixed effects yields an elasticity of firm employment to automation of 0.426 (s.e. 0.0842). 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 0.325 (s.e. 0.131) to 0.346 (s.e. 0.103) across specifications. In addition, we cannot reject that there is no impact of automation on wages, on inequality across workers, or on the labor share.

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 business-stealing effects across firms may affect the industry-level impacts of automation. Indeed, the shift-share IV design shows automation at a firm causes a fall in competitors' employment.

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 first present descriptive patterns, using industry-level event studies, on the relationship between automation and a range of outcomes, including employment, wage inequality and sales. We find that the relationships are the same as in the event studies implemented at the level of firms and plants. Industry-level automation remains associated with higher employment for all skill groups and with higher sales, while average wages, wage inequality and the labor share remain stable.

To address potential confounding factors in the industry-level event studies, we use an industry-level shift-share design to estimate the causal impact of automation on employment at the industry level. The shift-share design leverages the exact same productivity shocks across foreign suppliers of machines as in the firm-level shift-share design, but we now measure outcomes across 5-digit industries rather than at the firm level. We find that the industry-level responses are similar to the firm-level responses described previously. The elasticity of industry-level employment to industry-level automation is positive at 1.011 (s.e. 0.213),[6] compared with 1.063 (s.e. 0.383) for industry sales; these point estimates are statistically indistinguishable from the firm-level estimates. Like in the firm-level analysis, we cannot reject the hypothesis that wages and the share of labor in industry sales remains unchanged, while profits increase.

The finding that the employment response to automation remains positive at the industry level may appear surprising given the potential for business stealing effect. To understand the mechanism, we examine the role of international business stealing effects in two steps. First, we examine the heterogeneity in the industry-level response depending on exposure to international trade. While the relationship between automation and sales or employment is positive and significant in sectors that are exposed to import competition, there is no significant effect in sectors with low exposure to international competition (below median). This finding is consistent with the view that the business-stealing effect induced by automation mainly affects foreign competitors' employment in sectors facing international competition, whereas it mainly affects domestic competitors' employment in less open sectors. At the firm-level, there is no such heterogeneity and the response

[6]When restricting attention to incumbent firms only, instead of accounting for entry and exit, the point estimate is reduced to 0.714 (s.e. 0,314).

of employment and sales remains positive and significant regardless of the degree of exposure to import competition in the firm's industry.

Second, using the elasticities for the response of sales and consumer prices to automation, we assess whether the industry-level patterns can be explained by a demand reallocation channel. We use a simple monopolistic competition model with CES demand, where consumers reallocate demand toward domestic firms with increased productivity and lower prices, at the expense of international competitors. We find that, in an open economy, standard consumer demand elasticities (Broda and Weinstein (2006)) can account jointly for the estimated elasticities of sales and prices. In contrast, it would be difficult to rationalize the industry-level results on sales and employment in a closed economy. Industry-level substitution would need to operate between industries (rather than between products within the same industry, produced either by domestic firms or by international competitors); because demand elasticities of substitution between industries are relatively small (Costinot and Rodríguez-Clare (2014)), explaining the observed sales response would require very large price changes that 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 to a change in productivity and prices 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,[7] where signs and magnitudes vary across studies, potentially due to the empirical challenges raised by causal identification at the industry level (e.g., Autor and Dorn (2013), Chiacchio et al. (2018), Dauth et al. (2018), Graetz and Michaels (2018), Mann and Puttmann (2018), Acemoglu and Restrepo (2019), Aghion et al. (2019), Cheng et al. (2019), Webb (2019), Adachi et al. (2020), Klenert et al. (2020)). A more recent line of work, parallel to ours, uses event studies to estimate the firm-level employment effects of robotization and documents a positive reponse (e.g., Acemoglu et al. (2020), Bessen et al. (2020), Bonfiglioli et al. (2020), Dixon et al. (2019), Domini et al. (2019), Humlum (2019), Koch et al. (2019)).[8] Furthermore, our analysis speaks to a broader literature on the estimation of the micro and macro capital-labor elasticities of substitution (e.g., Oberfield and

[7]Since the 2010s, the International Federation of Robotics (IFR) has provided data on the deployment of robots by country and industry, and machine learning algorithms have made it possible to measure automation using text analysis of patents. While robots have been a focus of much recent work, Benmelech and Zator (2021) show that investment in robots accounts for less than 0.30% of aggregate expenditures on equipment and that recent increases in robotization do not resemble the explosive growth observed for IT technologies in the past.

[8]We provide a review of the literature to date and the divergence between existing estimates in a companion survey paper, Aghion et al. (2021).

Raval (2021), Hubmer (2018), Houthakker (1955)) and capital-skill complementarities (e.g., Goldin and Katz (1998), Doms et al. (1997)).

We contribute to this literature in four ways. First, we introduce a quasi-experimental shift-share design to provide causal estimates 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 sales, 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.[9] The shift-share IV design implemented at the industry-level allows us to quantify the impact of automation on domestic employment accounting for business-stealing across domestic firms. Fourth, we examine the role of exports and we document heterogeneity in the industry-level effects automation, depending on exposure to international competition. The heterogeneity we uncover depending on trade exposure helps reconcile some of the diverging industry-level estimates in prior work.[10]

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 the next generation of 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, at the firm and plant levels. Section 4 reports the causal estimates from the shift-share design at the firm level. 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.

[9] We conjecture that our methodological approach, which uses the same shocks to study outcomes at different levels of aggregation in a unified shift-share IV framework, could be applied to advance research on several other topics beyond automation.

[10] For example, 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). By contrast, exposure to international competition is higher in the sample of European countries studied by Klenert et al. (2020), including France.

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 stems from French administrative data, the DADS and FICUS/FARE databases. These databases cover the universe of plants and firms in the manufacturing sector in France from 1995 to 2017. 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 occupation codes for each worker.[11]

Measuring automation. Automation technologies correspond to a subset of capital used in production. We use three complementary proxies for automation, at the industry, firm and plant levels.

Our first proxy for the use of automation technologies leverages detailed balance sheet information available for the universe of French firms. Following French accounting standards, more specifically the general accounting plan (*Plan Général Comptable*), our balance-sheet measure of automation is the aggregation of (i) industrial equipment and (ii) industrial tools. Industrial equipment includes "all equipment and machines used for extraction, processing, shaping, packaging of materials or supplies or for services", while industrial tools include "instruments which, combined with an industrial equipment, specialize this equipment into a specific task." In short, our balance-sheet measure of automation includes all the machines used during the production process of manufactured products, which are specialized into a specific task by industrial tools. This measure excludes transport equipment, which corresponds to "all vehicles and devices used to transport people and goods, materials and products" as well as office and IT equipment that includes "typewriters, accounting machines, computers, etc.".[12]

Our second measure of automation 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 based on predetermined instructions or procedures." In manufactur-

[11]Measures of worker skills are obtained from Charnoz and Orand (2017).

[12]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 (55 %) of total capital in manufacturing, more than the three other categories, namely "land" (1 %),"building" (12 %) and "others" (32 %).

ing, 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.

With this motivation in mind, we build a 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 French statistical institute INSEE, in the Annual survey on industrial energy consumption, since 1983 for a large representative sample of plants. 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 excludes electricity used for heating and cooling as well as for servers, because servers are not considered to be directly part of the production chain. The motive power measure only takes into account electric motors that are constantly plugged-in when the production process is ongoing; it therefore excludes machines powered by electric batteries such as an electric forklift or electric car. The measure is expressed in tons 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, at the cost of being available only for a sample of plants rather than the full population.

Our third measure of automation focuses on the subset of firms that import machines from abroad. Using customs data for the universe of French firms, we can draw a precise list of imported automating machines. This approach provides a precise understanding of the machines that are included in our automation measure, although the analysis must be restricted to importing firms. Most types of machinery are found in two broad categories, HS84 "Nuclear reactors, boilers, machinery and mechanical appliances; parts thereof" and HS85 "Electrical machinery and equipment and parts thereof; sound recorders and reproducers, television image and sound recorders and reproducers, parts and accessories of such articles". Given our focus on industrial automation, we exclude from these categories a number of more detailed product categories corresponding to household machines (for cooking, washing, cleaning etc.), agricultural machinery, and IT machines.[13] We thus keep 489 categories out of 1338 categories of machinery. For example, industrial robots are included in this list; additional examples are provided in Section II.B.

[13]Examples of machines we want to exclude from the automation measure are: dish washing machines, lifts and escalators, printers and copying machines, calculating machines, microphones, headphones and hearphones, vacuum cleaners, telephone sets.

Our measures of automation have the advantage of covering a wide range of sectors rather than being concentrated on a small number of industries.[14] They may include machines that are not pure substitutes for labor and that instead complement labor,[15] at least at some level of aggregation of the production function (e.g., at the firm level rather than at the task level).[16]

All three proxies for automation 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, import value or peak motive force 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 Sections III, IV and V, we use panel data to describe how employment or other outcomes change after a plant or firm increases its investments in automation (as measured by our proxies), including time-by-industry fixed effects to control for potential time-by-industry changes 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 electric energy used for motive power.

Trade. The trade dataset is available from customs records and covers the population of French

[14] As shown by Benmelech and Zator (2021), investment in robots is small, and highly concentrated in few industries, primarily the automobile industry in France. In the French automobile sector investments in robots account for 0.04% of total investments, and in the other sectors they account for less than 0.015%. Whether we use IFR data or robot imports data, robotization does not capture automation in sectors such as pharmaceuticals, paper, or textiles.

[15] This feature is common to all proxies for automation. In particular, it applies to the IFR measure of robots, which are not necessarily substitutes to labor. For example, so-called "cobots" are collaborative robots that assist workers in some way, either to help them perform a task or as a guide. Unlike autonomous robots which operate alone and without supervision, cobots are programmed and designed to respond to human instructions and actions. "These collaborative robots are not replacing human work, but are increasing the productivity of human workers, whilst simultaneously reducing the risk of workplace injury-for example due to repetitive heavy lifting" (IFR, 2017).

[16] As discussed further in Section V.C, our empirical results show that such complementarity is not the main force at work explaining the increase in employment. If the complementarity between capital and labor was the driving force, (i) we would find an increase in the labor share, while we find that the labor share remains stable; (ii) we would find a positive response of industry-level employment in all industries, whereas we observe a positive response only in sectors that are open to international trade.

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 IV, as well as to isolate the role of specific machines or robots, focusing on the subset of French firms, as discussed previously. 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 in Section V. We match these data to consumption spending patterns by income groups, also from INSEE, to describe the distributional effects of automation via changes in purchasing power. 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.

II.B Summary Statistics

Table 1, Table 2 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 – the balance sheet value of industrial equipment, motive power, and imports of industrial machines – at various levels of aggregation, i.e. 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 analysis is conducted with 2,773 plants and 1,599 firms that operate continuously from 1995 to 2017 in 255 manufacturing industries.[17] 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.

Next, Table 2 provides a series of examples of the types of machines included in the import categories that we selected in the trade data as proxies for automation technologies. These examples highlight the breadth of our measure, which encompasses machines for the production of semiconductors, for metal working, for bending, folding, straightening or flattening, etc.

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, glass and ceramics, food and beverages, and metals. These patterns show that our proxis for automation capture a wide range of relevant machines in multiple industries. Appendix Figure

[17]We focus the analysis on the set of firms available firms but our main results are similar when using an unbalanced panel of firms (unreported).

A1 illustrates this finding by reporting examples of machines using motive power. 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 when 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 50 % of robots. The other industries in the top five for robots — food and beverages, metal products, rubber and "other" — account for much smaller shares.

Summary statistics on expenditure shares across income groups are also instructive to learn about the household groups that are most likely to benefit from automation technologies. Appendix Figure A2 documents that the average (sales-weighted) income of consumers an industry sells to is *lower* in industries that rely more intensively on automation technologies. This pattern indicates that lower-income households are likely to benefit relatively more from potential price declines brought about by automation and increased productivity.[18] We estimate the extent to which productivity gains from automation are passed through to consumers in the following section, using direct price measures or by inferring consumer price changes from changes in sales.

III Descriptive Evidence: Stylized Facts and Event Studies

This section provides descriptive evidence on the relationship between automation, employment, sales, prices, wages, 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 wages and 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

[18]For small shocks, the envelope theorem (Roy's identity) implies that price changes affect each type of consumer in proportion to their spending share across products indexed, regardless of the demand system. The first-order effect comes from the expenses the agent saves by paying lower prices, holding spending shares constant.

these variables in Subsection III.B, then we provide semi-elasticity estimates from event studies in Subsection III.C.

The results in this section should primarily be viewed as descriptive. The event study design does 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 both have a direct impact on employment and may also lead a firm to invest more heavily in automation technologies. It is therefore difficult to sign the potential bias of the estimates of the employment response to automation: 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 Firm-level Stylized Facts

We first compare the path of sales, employment and labor share for firms that automate more or less over time. Considering all firms that operate between 1995 and 2017, we rank them by the change in the balance sheet value of industrial equipment observed at the beginning of the sample, between 1996 and 1999. 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 2 presents the results. Panel (a) shows that firms that automate more at the beginning of the sample experience a larger increase in sales over the full sample. By 2017, total (nominal) sales have increased by 70% for firms with automation above median and by only about 15% for those below median.

Panels (b), (c) and (d) show that plants that automate more expand employment relative to those that automate less. Panel (b) reports this pattern for high-skill workers. By 2017, the number of high-skill workers increases by about 140% at firms above median, compared with 60% for those below median. In panel (c), the number of medium-skill workers increases for firms with automation above median, while it decreases for those below median. Panel (d) shows that the number of low-skill workers decreases in both groups, but more steeply for firms with automation below median. For firms that automate more at the beginning of the sample, low skill employment falls by about 33% by 2017, while the fall is more pronounced and reaches about 45% for firms with automation below median.

Finally, Panel (e) 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 no significant difference across these groups.

Consistent with the observed increase in sales, the potential productivity effect from automation may more than offset the potential displacement effect on workers, resulting in a positive effect on employment. We obtain similar results when repeating this analysis with thresholds other than the median, when using the plant or the industry (rather than the firm) as the level of analysis, and with our other proxies for automation.

In the remainder of the paper, we refine this analysis to provide causal evidence.

III.C Event Studies

III.C.1 Specifications and Identification

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, we use a standard "extensive margin" event study that isolates investment events in automation technologies. We observe that most firms and plants adjust their usage of motive power and their stock of machines or imports of machines every year. We define discrete investment thresholds isolating the largest changes in the use of automation technologies by firms. In the baseline approach, an automation event for a firm corresponds to a change in its level of automation – measured by the balance sheet value of industrial equipment – above a pre-specified threshold in the distribution of possible changes across firms. We consider alternative thresholds and variables, defining the automation event based on different percentiles of the distribution and using our three proxies for automation.

In our baseline specification, we study investment events defined alternatively by p90, p75 or p50 of the change in the balance sheet value of industrial machines. For the sake of robustness, we study the two other proxies for automation – the change in motive power, and the change in the value of imported machines. When analyzing the change in peak capacity for motive power, which is only available for a representative sample of plants surveyed by INSEE, to maximize power we leverage the entire variation available in the data by specifying a standard distributed lead-lag model (e.g., Stock and Watson (2015)).

Specification. Indexing firms by i and years by t, our event study is specified as

$$\Delta log(Y_{it}) = \sum_{k=-10}^{10} \delta_k E_{i,t-k} + \mu_i + \lambda_{st} + \epsilon_{it}, \tag{1}$$

with Y_{it} the outcome of interest, the event $E_{i,t-k}$, firm 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 outcome Y_{it} at time $t+k$ to the automation event at time t. We consider a variety of outcomes at the firm level, including employment, sales, wage, the labor share, and measures of within-firm wage inequality. The specification is similar when analyzing the change in peak capacity for motive power with the distributed lead-lag model.[19]

Identification. A causal interpretation of the estimates requires the identification condition:

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

The estimated coefficients for the "leads" (i.e., $\widehat{\delta_k}$ 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. Standard errors are clustered at the firm level.

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} in robustness checks.

[19] Indexing plants or firms by i and years by t, our baseline distributed lead-lag model is specified as

$$Y_{it} = \sum_{k=-10}^{9} \delta_k \Delta M_{i,t-k} + \delta_{10} M_{i,t-10} + \mu_i + \lambda_{st} + \epsilon_{it},$$

with the outcome Y_{it}, the change in peak capacity for electric motive power $\Delta M_{i,t}$, plant fixed effects μ_i, and industry-by-year fixed effects λ_{st}. δ_k is the cumulative impact of automation on the outcome after k periods (see Stock and Watson (2015), equation (15.7)).

III.C.2 Results

Impacts of automation on firm-level employment. Figure 3 reports the results of the firm-level event studies. We find that employment increases in firms that automate more, using the change in the balance sheet value of industrial equipment as a proxy. Panel A implements the event study with 5-digit industry by year fixed effects, defining the automation event as a change beyond the 90th percentile of the distribution of changes in the balance value of industrial equipment. The semi-elasticity of firm employment to the automation event is +0.2 on impact.[20] The response of employment is then amplified over time, with a semi-elasticity of +0.4 after ten years. The point estimates are precise; the 95% confidence interval rejects an employment elasticity below +0.20 or above +0.50 after ten years. Panels B and C show similar results using alternative thresholds to define the automation event (p75 and p50, respectively), with slightly smaller semi-elasticities.

There is no sign of pre-trends: conditional on the controls included in our statistical model, firms that automate more at a given time were on a comparable employment path in years prior to the event and started diverging only afterwards. This finding restricts the potential set of confounders that could explain the increase in employment — namely, confounding shocks need to occur simultaneously to the increase in automation.

Figure 4 examines the robustness of these findings when using different proxies for automation. In Panel A, we keep the balance sheet value of industrial machines as our proxy for automation, but we restrict attention to firms in the automobile industry, where industrial robots are prevalent. The patterns are very similar. In Panel B, we use the proxy for automation based on imports of industrial automating machines, within the subset of firms that import machines. The semi-elasticity of employment remains positive, at about +0.25 after 10 years. Finally, Panel C reports the elasticity of employment to motive power, which is also positive at +0.30 after 10 years. Thus, we consistently find a positive response of employment to automation events.

In Online Appendix Figure A3, we obtain similar results with an alternative measure of the flow of investment, using investments in industrial equipment as a fraction of the initial balance-sheet value of the stock of machines.[21] Online Appendix Figure A4 shows that the results remain stable in alternative specifications with other sets of interacted industry fixed effects, i.e. 2-digit industry by year fixed effects or 4-digit industry by year fixed effects; the point estimates remain nearly

[20] The average log change in the balance sheet value of industrial machines after the event is close to one, such that the semi-elasticities can also be interpreted as elasticities.

[21] This alternative measure is not sensitive to depreciation.

unchanged, with no pre-trends.

A 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 adjusting variable inputs). To address this concern, instead of relying on the actual electricity consumption for motive power we use the plant's peak capacity for electric motive power, which is provided by INSEE in the same survey. After major investments in machines, the plant should adjust its peak capacity for motive power, while there is no such change when the plant simply varies its factor utilization rate. Online Appendix Figure A5 shows that the results are similar when using peak motive power.

Impact on employment by skill groups. Figure 5 documents heterogeneity across skill groups, using the specification with 5-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 semi-elasticity to automation is about +0.2 on impact and increases to about +0.4 after ten years. The paths of the point estimates for the three skill groups are statistically indistinguishable.

These results indicate that automation does not have different effects across broad skill groups within the firm. Online Appendix Figure A6 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 results remain similar at the firm level and plant level.

Impact on within-firm inequality. Figure 6 reports the impact of the automation event on within-firm wage inequality. Panel (i) show that the average hourly wage remains flat after the automation event, with no sign of pre-trends. Panels (ii) and (iii) show that the average hourly wage remains similarly flat for low-skill workers and medium-skill workers. In Panel (iv), we see that the average wage for high-skill workers stays relatively flat, but tends to increase after four years, with a semi-elasticity of 0.02. Thus, none of the broad skill groups suffers from a fall in wages.

Panels (v) through (vii) complement these findings with several measures of within-firm inequality. In Panel (v), we use the ratio of the 90th to the 10th percentile of wages within the firm as our dependent variable. We find that this ratio remains unresponsive to the automation event, i.e. within-firm inequality remains unchanged. Panel (vi) document a similar pattern by studying two specific occupations, unskilled industrial workers and engineers. We find that the wage ratio of unskilled industrial workers to engineers remains unchanged after the automation event. Finally, Panel (vii) examines the response of the labor share, defined as the ratio of total labor costs over total sales (as in Autor et al. (2020)): the labor share remains flat.

The patterns so far suggest that workers from all skill groups may benefit from automation, on average, as employment increases and wages remain stable, with no impact on within-firm inequality. Figure 7 focuses on the rates of job creation and destruction within the firm. Using the balance sheet value of industrial machines as our proxy for industrial equipment, Panels A(i) and A(ii) show that automation events lead to an increase in both job creation and job destruction. Although there is more job creation than job destruction, these patterns indicate that automation events lead to many instances of job reallocations.[22] In Panel A(iii), we compute a within-firm job dissimilarity index, computed using shares of workers by occupations across consecutive years. We find that there is an increase in the dissimiliarity index exactly at the time of automation, i.e. automation induces a reallocation of occupations within the firm. Panel B reports a placebo test, using investments in real estate: in this case, the event is not related to patterns of job creation and destruction and the job dissimilarity index remains flat.

Overall, the patterns in Figures 6 and 7 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. However, there is no evidence that automation leads to important changes in within-firm inequality.[23]

Market dynamics. Figure 8 documents the impact of the automation event on market dynamics. Panel (i) shows the response of total sales, which increase markedly after the automation event, with a semi-elasticity of 0.2 on impact that increases slightly over time. Panel (ii) shows that export sales also increase. Finally, Panel (iii) uses competitor employment as the outcome, showing a negative impact. The semi-elasticity of competitors' employment (defined as domestic firms in the same 5-digit industry) is negative, at about -0.001.

Together, these patterns highlight the importance of the scale effect brought about by the increase in productivity due to automation, but also the potential for business-stealing effects negatively affecting employment in firms that do not automate. These findings motivate our analysis

[22]This finding speaks to the fact that our measures of automation do capture automation, inducing reallocations of tasks, and not just complementarities between labor and some subsets of machines.

[23]Hummels et al. (2014) obtained similar results when studying automation in Denmark, finding that the wage effects of offshoring vary substantially within skill type.

at the industry level in Section V, which account for these equilibrium effects at the industry level.

Exports prices. Next, we document the relationship between changes in automation and price changes. To do so we use export prices, which are readily available for all exporting firms from customs data. Export prices are measured as the unit value of exported products. To account for potential changes in composition over time, we run the specification at the level of detailed product cells identified by the standard product classification for traded goods, HS6 codes. Our baseline specification is the same event study as in equation (1), but with a different set of fixed effects: we now control for HS6-by-year fixed effects, trading partner by year fixed effects, and firm fixed effects.

The results of the baseline specification are reported in Panel A of Figure 9. We find that export prices fall after an increase in the firm's industrial equipment. In line with the employment patterns, there is no sign of pre-trends. The estimated elasticity of prices reaches -0.10 after four years. The observed fall in prices suggests that firms that automate pass through some of the productivity gains to consumers, leading to higher demand and more employment. We return to this demand reallocation channel in Section V.C.

We conduct additional tests to address the possibility that changes in the composition of products sold by the firm may affect the average unit price of exported goods we observe over time. In Panel B of Figure 9, we implement the same specification using NC8 product category codes rather than HS6 codes. The NC8 product category codes are the most detailed classification used by the French customs. The panel shows that the results are very similar, with a semi-elasticity of -0.10 after 4 years. The finding suggests that composion effects do not afflict our results, which remain stable across product classifications.

In Appendix Figure A7, we document similar results when using the 95th percentile instead of the 90th percentile as the threshold defining the automation event. With this alternative threshold, the semi-elasticity is 0.10 on impact and falls further to -0.20 after 4 years. In unreported robustness checks, we find that the results are similar in a sample restricted to firms that only export products in a single HS6 code.

Limitations. Despite the robustness of the firm-level relationships documented above, in particular between automation and employment, two potential concerns remain. First, because so far 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 IV with a shift-share design. Second, the positive firm-level relationship between automation and employment, along with the negative relationship between automation and prices, may be misleading because of business-stealing effects across firms, which could affect the industry-level relationships. In Section V, we conduct a similar analysis at the industry level to incorporate such reallocation effects.

IV Causal Estimates from Firm-Level Shift-Share IV

In this section, we introduce a quasi-experimental shift-share design to estimate the causal effects of automation on employment, sales, wages and the labor share across firms. The results validate the findings from Section III: 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 shocks are obtained from variation in the cost of imported machines over time across international trading partners across detailed HS6 product categories.[24] We focus on imported automation technologies, our third proxy for automation. Shocks are observed across "trading partners by HS6 product" cells indexed by n (for example, imports of machines from China for the production of semi-conductors). The shocks g_{nt} are measured as the aggregate changes in import flows of industrial machines from each trading partner for each HS6 product category between 5-year periods centered around t.

To obtain an instrument plausibly exogenous to the choices made by French firms, we study trade flows in countries similar to France. Specifically, we use HS6-level trade shocks measured in EU countries (except France) and Switzerland as instruments for the adoption of machines in France. Considering trade flows in these countries, we compute the following symmetric percentage change over time:

$$g_{nt} = \frac{ImportMachines_{n,t+1,t+5} - ImportMachines_{n,t-4,t}}{ImportMachines_{n,t+1,t+5} + ImportMachines_{n,t-4,t}}, \tag{3}$$

where n is a "trading partner by HS6 product" cell. In the baseline specification, we measure the

[24]The Hamonized System (HS) nomenclature, which was developed by the World Customs Organisation (WCO), is an internationally recognized nomenclature of standardized product classification used in over 200 countries. It encompasses over 5,000 commodity groups, every group being labeled with a 6-digit code.

shocks g_{nt} across 149 trading partners in 257 5-digit industries. We conduct the analysis around years $t = 2005$ and $t = 2010$, i.e. we leverage variation across consecutive five-year periods from 2001 to 2015, measuring shocks across HS6 product categories using equation (3).

Using changes in the market shares of international suppliers over time is helpful because changes in the quality-adjusted prices of machines are not directly observed. The customs dataset only provides unit values, which are difficult to adjust for quality. But changes in import flows can be used to infer changes in quality adjusted prices. Indeed, 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, 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.

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_{i0n} 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 initial shares, measured from 1996 to 2000, and we conduct the analysis from 2000 onward, with the trade shocks measured using trade data excluding France.

The shift-share instrument is built by combining the shocks and exposure shares. The outcomes and endogenous variable are log changes across consecutive five year periods, centered around $t = 2005$ and $t = 2010$. We use our first measure of automation, the log change in the balance sheet value of industrial equipment denoted ΔM_{it}[25] across firms indexed by i. Denoting by ΔL_{it} log changes in employment, we estimate by 2SLS

$$\begin{cases} \Delta L_{it} = \beta Z_{it} + \gamma X_{it} + \varepsilon_{it}, \\ \Delta M_{it} = \alpha Z_{it} + \widetilde{\gamma} X_{it} + \widetilde{\varepsilon}_{it}, \end{cases} \tag{4}$$

where Z_{it} is the shift-share instrument constructed from shocks g_{nt} and (initial) exposure shares $s_{i0n} \geq 0$,

$$Z_{it} = \sum_{n=1}^{N} s_{i0n} g_{nt}.$$

[25]In our baseline specification, we do not use the measure of automation based on imports as our endogenous variable because it is measured as a flow rather than a stock.

We study the sensitivity of the estimates to changes in the set of time-varying controls X_{it}. We use a battery of period-specific fixed effects. Specifically, in our baseline specification we use HS4-by-period fixed effects (i.e., all variation arises within HS4 product categories within each 5-year period) as well as 2-digit-industry-by-period fixed effects and trading-partner-by-period fixed effects. In this way, we compare French firms in the same 2-digit industry that source their inputs from different suppliers, within narrowly defined HS4 product categories. Fixed effects for industries and product categories are allowed to be period-specific so that they can flexibly absorb potential demand shocks.

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 10 reports the length of the relationships 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 13.2 years on average.

The exclusion restriction underlying this design is that firms linked to increasingly productive suppliers should not feature unobservable characteristics that affect the outcomes of interest. To test this hypothesis, one can run a falsification test using the lagged outcome variable. Omitting the period subscripts for brevity, the exclusion restriction can be expressed equivalently at the firm level or in space of productivity shocks (across foreign suppliers of different machines):

$$\left(\frac{1}{I}\sum_{i} Z_i \varepsilon_i \xrightarrow{p} 0\right) \Longleftrightarrow \left(\sum_{n} \hat{s}_{0n} g_n \bar{\varepsilon}_n \xrightarrow{p} 0\right),$$

with $\bar{\varepsilon}_n = (\sum_i s_{i0n}\varepsilon_i)/\sum_i s_{i0n}$ and $\hat{s}_{0n} = \frac{1}{I}\sum_i s_{i0n}$. 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, rather than from exposure shares (as in Goldsmith-Pinkham et al. (2020)). The effective number of shocks leveraged by this research design can be gauged by estimating the inverse of the Herfindahl index (HHI) of the weights $\hat{s}_{0n}$. Intuitively, if a few trading partners have most of the market shares, the effective sample size is small and the condition for consistency ($E\left(\sum_{n=1}^{N} (\hat{s}_{0n})^2\right) \rightarrow 0$) may not be met. In practice, we compute that the inverse HHI of the weights $\hat{s}_{0n}$ is 268, indicating that the effective sample size is large.

Inference. In a shift-share IV design, observations cannot be treated as i.i.d. We follow Adao et al. (2019) and Borusyak et al. (2019) to correct standard errors and the first-stage F-statistic appropriately. All results are clustered by trading partner, which allows for correlated shocks within a trading partner over time and across industries. For example, China may experience positive productivity shocks throughout our period of study in a large number of industries.

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 trading partner by period fixed effects, 4-digit product period fixed effects, and 2-digit industry-period fixed effects. The second specification adds a set of pre-determined firm controls including lagged turnover, total asset, employment, and the share of industrial workers in total employment. The third specification controls for the lagged balance sheet value of industrial equipment, and the fourth specification controls for other types of capital (land, buildings, others). Finally, because trade flows play a central role for identification, the final specification adds controls for contemporaneous exports to ensure that potential export shocks do not confound the results. The stability of coefficients across specifications can be viewed as a test of the exclusion restriction, as explained in the discussion of identification in shift-share IV design of Borusyak et al. (2019).

IV.B Results

The results and falsification tests are reported in Tables 3 and 4, using the change in the firm-level balance sheet value of industrial machines as our endogenous variable proxying for automation.[26]

We start by reporting the OLS relationship between automation and employment at the firm level in Table 3. We only keep the set of firms that import machines so that the results can be compared with the shift-share IV design. Panel A focuses on employment growth. Column (1) includes year fixed effects interacted with trading partners, HS4 categories and 2-digit industries. We obtain an elasticity of employment to automation of +0.409 (s.e. 0.0212), which is similar to 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 0.413 and 0.415 across specifications. Panels B of Table 3 shows that the OLS relationships with sales are positive, with elasticities around 0.30. Panel C reports the relationship with hourly wages, with is small statistically significant point estimate around -0.0367 (s.e. 0.00554). Panel D shows the results for the labor share, defined as labor cost over total

[26]The source of variation in the SSIV is trade shocks for imports of industrial automating machines.

sales, which is small and insignificant. When we instead define the labor share as labor cost over value added, we obtain similar results. Panel E shows a positive correlation with profits, with an elasticity close to 0.35 across specifications. Finally, Panel F reports that automation is associated with a decline in competitors' employment within the same 5-digit industry. To assess whether these OLS estimates are biased, we next turn to the shift-share IV design.

Panel A of Table 4 reports the estimates of the impact of automation on employment using the shift-share instrument. The baseline specification in Column (1) yields an elasticity of firm employment to automation of +0.426 (s.e. 0.0842). The point estimate is statistically significant at the 1% level and the first stage F statistic of 17.65 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 0.424 and 0.433, are all significant at the 1% level, and are statistically indistinguishable from one another. The first stage F statistic remains above 20 in all specifications.

These results support the conclusion from Section 3: increases in automation lead to higher employment at the firm level. Relative to these SSIV results, the OLS estimates from Table 3 do not appear to be biased either upward or downward. Offsetting effects may explain why OLS estimates appear to be unbiased: while some firms automate in response to increased competition, which could have a direct negative effect on employment (downward bias), other firms automate in response to increase demand, which could have a direct positive effect on employment (upward bias), such that the net bias in OLS estimates is close to zero.

Next, Panel B of Table 4 takes sales as the outcome. We find that sales increase in response to increased automation, with elasticities ranging from 0.325 to 0.346 across specifications. The relationship is significant at the 1% 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.

Panel C of Table 4 presents estimates of the impact of automation on average hourly wages in the firm. Consistent with the results from Section III, we find no impact on wages. Panel D of Table 4 presents estimates of the impact of automation on the labor share, defining the labor share as the ratio of total labor to sales. The outcome is the difference in the labor share over time (in levels), and we estimate semi-elasticities. In all five specifications, we cannot reject that there is no impact of automation on the labor share. The results are consistent with the OLS semi-elasticity estimates reported in Table 3, which are close to zero, with small standard errors, and where the sign of the semi-elasticity is positive. These findings indicate that the productivity effect may offset the task substitution channel in a way that leaves the labor share unchanged at the firm level. Panel E of Table 4 shows a positive impact of automation on firm profits, with a point estimate of 0.995 (s.e. 0.448) in the baseline specification.

Turning to business-stealing effects, Panel F of Table 4 reports a negative impact of automation on competitors' employment (within the same 2-digit industry).[27] The point estimates are negative and statistically significant in all specifications, varying between -0.00578 and -0.00920.

Finally, Panel G and F of Table 4 reports the results of pre-trend falsification tests, using the exact same shift-share IV specification, but using lagged employment and sales as the outcomes. Conceptually, these lagged outcomes can serve as a proxy for the unobserved error terms ε_{it} in equation (4). Across all five specifications, we cannot reject that there is no relationship between the shocks and lagged employment growth or lagged sales growth. These results lend credibility to a causal interpretation of these estimates.

IV.C Robustness

We implement several robustness checks. First, using the specification from Column (1) of Table 4, Figure 11 shows the first-stage (panel A) and reduced-form relationships underlying the shift-share IV design (panels B(i) and B(ii)), as well as the falsification tests (panels C(i) and C(ii)). Supplier shocks are taken as the unit of observation, using the numerical equivalence in Borusyak et al. (2019). This figure depicts relationships that appear to be robust graphically, with conditional expectation functions close to linear.

Second, we obtain very similar results when using a more stringend set of fixed effects. In Appendix Table A1, we use HS6-by-period fixed effects along with 5-digit-industry-by-period fixed effects and trading partner by period fixed effects. All results are statistically indistinguishable from those of our baseline specification in Table 4. Figure A8 reports the binned scatter plots for the first stage, reduced-form relationships and falsification tests corresponding to Column (1) of Table A1.

Third, in Appendix Table A2, we conduct the analysis with less stringent fixed effects, using HS4-by-period fixed effects and trading partner by period fixed effects, without industry fixed effects. The results remain statistically indistinguishable and are depicted graphically in Figure A9 for the specification of Column (1).

[27]The results are similar when considering competitors' employment within the same 5-digit industry, as reported in Table A1.

Fourth, we obtained similar results when repeating the analysis with an alternative set of shocks, using French customs data to focus on the most detailed item categories available from the French customs data, called NC8 product categories (not reported).

V Industry-Level Analysis

In this section, we study the relationship between automation, employment, prices, and profits at the level of industries. We first present event studies in Section V.A, then the industry-level shift-share IV design in Section V.B., and we assess the role of international business stealing effects and the demand reallocation channel in Section V.C.

V.A Industry-level Event Studies

The positive plant-level and firm-level relationships 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 start by implementing the same event study methodology as in Section III at the 5-digit industry level, with year fixed effects and 5-digit-industry fixed effects. In the baseline specification, we use the 50th percentile of changes in the balance-sheet value of industrial machines as our event threshold. Our measure of employment includes all firms, i.e. we account for entry and exit at the industry level.

The results are reported in Figure 12. Panel (i) shows that industry-level employment increases after the automation event, with a semi-elasticity of about 0.10 on impact, increasing to about 0.20 over 5 years. These patterns indicate that the employment response remains positive at the industry level. Panel (ii) to (v) are also similar to the firm-level estimates showing no impact on inequality: Panel (ii) shows that there is no impact on relative labor demand between high- and low-skill workers at the industry level; Panel (iii) documents that average hourly wages remain unaffected; Panel (iv) reports that the wage ratio between high skill and low skill employees remains constant; and Panel (v) shows that we cannot reject that the labor share remains constant, although it is imprecisely estimated. Finally, Panel (vi) reports the response of sales, with a semi-elasticity of about 0.2 on impact, similar to the firm-level pattern.

Overall, these findings indicate that, despite the potential for business stealing effects, the overall effect of automation on employment remains positive at the industry level through a scale effect, with a large increase in sales. However, like previously, the event study remains liable to correlated demand or supply shocks, which we address next by developing an industry-level SSIV design.

V.B Causal Estimates from Industry-Level Shift-Share IV

To assess whether a causal interpretation of the industry-level event study estimates is warranted, we implement and industry-level shift-share IV design.

Research design. The research design is identical to the shift-share IV presented in Section IV.A with equations (4), except for the fact that i now indexes 5-digit industries rather than firms. We use the same trade shocks, measured across detailed HS6 product categories in the EU (excluding France) and Switzerland. We use the same set of imported inputs as previously, and compute the symmetric percentage change over 5-year periods in equation (3). We can thus examine the response of employment and sales in industries that source their machines from increasingly productive foreign suppliers.

In this design, all outcomes are measured at the level of 5-digit industries, which are narrow and include, for example, "manufacture of plastics plates, sheets, tubes and profiles" or "manufacture of metal structures." Furthermore, the exposure share s_{i0n} is computed as the share of trading partner n in industry i's total imports of machines and robots. As previously, we use initial shares measured from 1996 to 2000, and we conduct the analysis from 2000 onward, with the trade shocks measured using trade data excluding France.

In the baseline specification, we measure the shocks across HS6 categories, with 149 trading partners in the 257 5-digit industries during consecutive 5-years periods centered about $t = 2005$ and $t = 2010$. The inverse HHI of the relevant weights $\hat{s}_{0n}$ is 712, indicating that the effective sample size is large. To address potential correlated demand shocks, we use HS4-by-period fixed effects as well as partner-period fixed effects.

Results. The results and falsification tests are reported in Tables 5 and 6, using the change in the industry-level balance sheet value of industrial machines as our endogenous variable proxying for automation.

Table 5 reports the OLS relationships. The results are similar to the firm level, with an increase in employment whether we consider total industry employment (including entry and exit) in Panel A or only incumbent firms in Panel B. The correlation with sales is positive in Panel C. There is a small positive correlation with wages (Panel D) and a negative correlation with the labor share

(Panel E), while there is a large positive correlation with profits (Panel F).

Panel A of Table 6 reports the estimates of the impact of automation on employment and other outcomes using the industry-level shift-share IV design. The baseline specification yields an elasticity of firm employment to automation of +1.011 (s.e. 0.213). The point estimate is statistically significant at the 1% level and the first stage F statistic of 8.38. The point estimates remain comparable in magnitudes in columns (2) through (5) as we change the set of controls, with similar F statistics. The point estimates vary between 1.003 and 1.019, are all significant at the 1% level, and are statistically indistinguishable from one another. These results support the findings from the previous subsection: increases in automation lead to higher employment at the level of the industry. Given the magnitudes of the standard errors, we cannot reject that the elasticity at the industry level is of the same magnitude as at the firm level. While Panel A accounts for the impact of entry and exit on employment, the point estimates for the industry-level employment elasticities are reduced to about 0.70 in Panel B when focusing on incumbent firms that exit in all periods.

Panel C of Table 6 turns to sales. We find that sales increase in response to increased automation, with elasticities ranging from 0.923 to 1.063 across specifications. The relationship is significant at the 1% level in all specifications. This finding is consistent with the role of the productivity effect of automation. Increased automation allows the industry as a whole to expand its sales and scale, which requires hiring additional workers for production. By assuming a given industry-level demand elasticity of substitution, we can infer the impact on the industry-level price index, which we investigate further in the next subsection. Panels A and B of Figure 13 shows the first-stage and reduced-form relationships underlying the industry-level shift-share IV design for employment and sales, depicting graphically the robustness of the findings.

Next, Panels D and E of Table 6 present estimates of the impact of automation on wages and the labor share. We define the labor share as the ratio of total labor costs to sales at the industry level (including pensions in total labor costs, as previously). In all four specifications, we cannot reject that there is no impact of automation on hourly wages and the labor share. These findings are similar to the firm-level analysis and suggest that the productivity effect may offset the task substitution channel in a way that leaves wages and the labor share unchanged, even at the industry level.

Finally, Panel F documents a positive elasticity of industry profits to automation. The elasticity is large in magnitude at about 2.8 but is imprecisely estimated, with standard errors of about 1.3, such that the results are statistically indistinguishable from the firm-level estimates in Table 4.

Table 4 also reports the results of pre-trend falsification tests, implementing the shift-share IV design taking as outcome the lagged changes in industry employment in Panel G and lagged changes in industry sales in Panel F. As previously, these lagged outcomes can serve as a proxy for the unobserved error terms ε_{it}, now at the industry level. Across all four specifications, we cannot reject that there is no relationship between the shocks and lagged employment growth. Panel C of Figure 13 reports the reduced-form relationships with lagged employment and lagged sales. These results lend credibility to a causal interpretation of the industry-level estimates.

In Online Appendix Table A3, for robustness we implement the industry SSIV with a less stringent set of fixed effects, partner-by-period fixed effects and HS4 fixed effects. The F statistics are now above 15 in all specifications and the point estimates remain similar, hovering between 1.080 and 1.091 across specifications for total employment, between 0.566 and 0.608 for incumbents' employment, and between 1.207 and 1.309 for sales.

V.C International Business Stealing and the Demand Reallocation Channel

Our finding that the elasticity of industry-level employment to a change in automation is quantitatively similar to the firm-level employment elasticity may seem surprising. Indeed, 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.[28]

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, in Table 7 we repeat the analysis for subsets of industries with trade exposure above or below median. We use the share of imports in final consumption, obtained from national accounts, to measure exposure to international competition.

Heterogeneity by exposure to international competition. In Table 7, we document that the positive industry-level relationship between automation and employment or sales is driven by industries that face a higher degree of international competition. To reduce noise in this subsample analysis, we implement OLS specifications with long differences between 1996 and 2017. The results for

[28]We also expect to find larger employment effects when consumers' demand elasticity of substitution is larger because consumers reallocate their spending toward firms or sectors where productivity increases and prices fall.

employment and sales with this specification for all industries, reported in Column (1), are in line with the industry-level estimates from Section V.B.

With higher exposure to international competition, the point estimate for employment in Column (2) is 0.404 (s.e. 0.055) and is similar to firm-level employment elasticities. In contrast, with lower exposure to international competition in Column (3), the point estimate loses statistical significance and falls in magnitude to 0.171 (s.e. 0.133). 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. Likewise, the response of sales in Column (2) is 0.510 (s.e. 0.084) with higher exposure to international competition, while is becomes smaller and statistically insignificant at 0.188 (s.e. 0.121) with lower exposure in Column (3).

The heterogeneity by exposure to international competition is thus consistent with the role of international business stealing. The demand reallocation channel predicts that heterogeneity should be visible only for industry-level outcomes, since at the firm level business stealing will operate regardless of exposure to international trade. With this motivation in mind, in Appendix Table A4 we implement a falsification test by running the firm-level analysis in the same subsamples of exposure to international competition. Consistent with the role of international business stealing, at the firm-level there is no heterogeneity and the employment and sales responses remain positive and similar in magnitudes regardless of the degree of exposure to international competition.

Furthermore, the findings in Table 7 confirm that the positive impact of automation on employment, which we found at both the firm and industry levels, does not primarily stem from the complementarity between labor and some machines included in our various measures of automation: if that were the case, then we would have found a positive correlation between automation and employment at the industry level even in industries that are not exposed to international competition.[29]

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, consumers should reallocate their expenditures toward industries that automate more. The magnitude of this reallocation effect is governed by consumers' demand elasticity of substitution. Appendix Table A5 reports a negative relationship between automation and the industry-level producer price index, with point estimates

[29]Complementarity between labor and machines would also have led to an increase in the labor share, which we can reject in the data.

ranging from -0.113 (s.e. 0.0573) to -0.199 (s.e. 0.0698) across specifications. The magnitudes are very similar to the firm-level price response documented in Figure 9.

To assess the plausibility of the demand reallocation 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 together with the negative price effects.

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 the representative agent is given by

$$U = \left(\sum_{k \in \Omega} \omega_k q_k^{1-\sigma} \right)^{1/(1-\sigma)},$$

where σ is the elasticity of substitution between varieties, q_k is the quantity index for variety k, and ω_k is a taste parameter reflecting the intensity of the representative agent's preference for variety k. p_k denotes the price index for country-industry variety k.

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 log-linear relationship between the change in the price index for industry k, p_k, and the change in total sales, $p_k \cdot q_k$:

$$\Delta \log(p_k) = -\frac{1}{\sigma - 1} \Delta \log (p_k \cdot q_k) + \Omega. \tag{5}$$

In response to a 1% increase in automation, according to Column (4) of Table 6 we have $\Delta \widehat{\log (p_k \cdot q_k)} = 0.923$; according to Column (4) of Table 7 we have $\Delta \widehat{\log(p_k)} = -0.178$. To satisfy equation (5), these estimates imply the following demand elasticity of substitution:[30]

$$\widehat{\sigma} = 1 - \frac{\Delta \widehat{\log (p_k \cdot q_k)}}{\Delta \widehat{\log(p_k)}} = 6.18.$$

Is the magnitude of $\widehat{\sigma}$ in line with existing estimates? A demand elasticity of substitution of 6.18 is consistent with estimates of elasticities of substitution between varieties produced by different countries for the same industry. For example, Broda and Weinstein (2006) estimate a

[30] The implied magnitude for $\widehat{\sigma}$ is similar when using the sales and price estimates from the firm-level analysis. From Column (5) of Table 4 we obtain $\Delta \widehat{\log (p_k \cdot q_k)} = 0.346$ and Figure 9 yields $\Delta \widehat{\log(p_k)} = -0.10$, implying $\widehat{\sigma} = 4.46$.

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 by domestic firms or by international competitors within the same industry) and would require large price changes that 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.[31]

This observation may also help reconcile some of the diverging industry-level estimates in the literature, depending on the degree of import competition in a country. For example, 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). By contrast, Klenert et al. (2020) estimate a positive relationship in a sample of European countries, including France, which are more exposed to international competition.

The analysis presented above is based on the industry-level OLS results on prices reported in Appendix Table A5, because we lacked power to estimate price effects directly in the industry-level shift-share IV design. If we know the demand elasticity of substitution σ, equation (5) can be used to infer price changes from our SSIV estimates for the change in sales. The industry-level SSIV yields an elasticity of sales to automation of 0.923 in Column (4) of Table 6. We can plug this estimate of $\Delta \log (p_k \cdot q_k)$ into equation (5) and use a standard range of empirical estimates for σ. For example, Broda and Weinstein (2006) report a mean of 7.5 between internationally traded varieties (within 5-digit SITC industries), while Simonovska and Waugh (2014) obtain an elasticity of 4.2.

[31] We focus on the impact of automation on domestic labor demand and our results do not speak to the impact on overall labor demand across multiple countries, which we view as an important topic for future research left outside the scope of this paper. Our results are not inconsistent with the idea that automation leads to structural change and labor reallocation across sectors (e.g., Ngai and Pissarides (2007)); rather, they highlight that, perhaps surprisingly, domestic manufacturing employment is better preserved in countries that automate faster (as in Germany, for example) due to a productivity effect and international business-stealing.

Depending on the choice of σ, the implied price elasticity to automation, $\Delta \log(p_k)$, ranges from -0.31 ($= -\frac{1}{4.2-1} \cdot 0.923$) to$-0.142$ ($= -\frac{1}{7.5-1} \cdot 0.923$). This range of implied price elasticities is close to the estimates obtained with the event study at the firm level (with a price elasticity of -0.20 in Panel B of Figure A7) and with the OLS analysis at the industry level (with a price elasticity of -0.113 in Column (1) of Table A5). This confirms that the demand reallocation channel can account for the estimates collected in this paper.[32]

VI Conclusion

In this paper, we have leveraged new micro data on plants, firms and industries in the French manufacturing sector to provide a unified analysis of the effects of automation technologies on employment, wages, prices, sales, and profits between 1995 and 2017.

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 the effect is heterogeneous depending on exposure to international trade, with a stronger employment response in industries that face 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 consumer demand, and in turn to increased firm and industry scale, higher labor demand and higher domestic 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

[32]It may be instructive to note that the demand reallocation channel could affect the optimal design of innovation policies. Domestic policymakers may not internalize the effects of domestic innovations on foreign consumer prices, nor their effects on the disruption of foreign labor markets. These channels create a motive for coordinating innovation policies internationally, which would be fruitful to characterize formally in the next generation of models of optimal technology regulation.

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: Firm-level evidence from france," in "AEA Papers and Proceedings," Vol. 110 2020, pp. 383–88.

Adachi, Daisuke, Daiji Kawaguchi, and Yukiko Saito, "Robots and Employment: Evidence from Japan, 1978-2017," *Discussion papers*, 2020, *20051.*

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Aghion, Philippe, Céline Antonin, and Simon Bunel, "Artificial intelligence, growth and employment: The role of policy," *Economie et Statistique*, 2019, *510* (1), 149–164.

_ **, Celine Antonin, Simon Bunel, and Xavier Jaravel**, "The Effects of Automation on Labor Demand: A Survey of the Recent Literature," *Working Paper*, 2021.

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.

_ **and David Dorn**, "The growth of low-skill service jobs and the polarization of the US labor market," *American economic review*, 2013, *103* (5), 1553–97.

_ **,** _ **, Lawrence F Katz, Christina Patterson, and John Van Reenen**, "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Benmelech, Efraim and Michał Zator, "Robots and Firm Investment," 2021.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, "Firm-level automation: Evidence from the netherlands," in "AEA Papers and Proceedings," Vol. 110 2020, pp. 389–93.

Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia, "Robot imports and firm-level outcomes," 2020.

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.

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.

Cheng, Hong, Ruixue Jia, Dandan Li, and Hongbin Li, "The rise of robots in China," *Journal of Economic Perspectives*, 2019, *33* (2), 71–88.

Chiacchio, Francesco, Georgios Petropoulos, and David Pichler, "The impact of industrial robots on EU employment and wages: A local labour market approach," 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, and Nicole Woessner, "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.

Domini, G, M Grazzi, D Moschella, and T Treibich, "Threats and opportunities in the digital era: automation spikes and employment dynamics (LEM Working Paper No. 2019/22)," *Laboratory of Economics and Management, Sant'Anna School of Advanced Studies. Pisa*, 2019.

Doms, Mark, Timothy Dunne, and Kenneth R Troske, "Workers, wages, and technology," *The Quarterly Journal of Economics*, 1997, *112* (1), 253–290.

Goldin, Claudia and Lawrence F Katz, "The origins of technology-skill complementarity," *The Quarterly journal of economics*, 1998, *113* (3), 693–732.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

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.

Houthakker, Hendrik S, "The Pareto distribution and the Cobb-Douglas production function in activity analysis," *The Review of Economic Studies*, 1955, *23* (1), 27–31.

Hubmer, Joachim, "The race between preferences and technology," *Working Paper*, 2018.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," *Working Paper*, 2019.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The wage effects of offshoring: Evidence from Danish matched worker-firm data," *American Economic Review*, 2014, *104* (6), 1597–1629.

Keynes, John Maynard, "Economic possibilities for our grandchildren," in "Essays in persuasion," Springer, 1930, pp. 321–332.

Klenert, David, Enrique Fernandez-Macias, and José Ignacio Antón Pérez, "Do robots really destroy jobs? Evidence from Europe," *JRC Working Papers Series on Labour, Education and Technology*, 2020.

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.

Ngai, L Rachel and Christopher A Pissarides, "Structural change in a multisector model of growth," *American economic review*, 2007, *97* (1), 429–443.

Oberfield, Ezra and Devesh Raval, "Micro data and macro technology," *Econometrica*, 2021, *89* (2), 703–732.

Simonovska, Ina and Michael E Waugh, "The elasticity of trade: Estimates and evidence," *Journal of international Economics*, 2014, *92* (1), 34–50.

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

Webb, Michael, "The Impact of Artifical 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

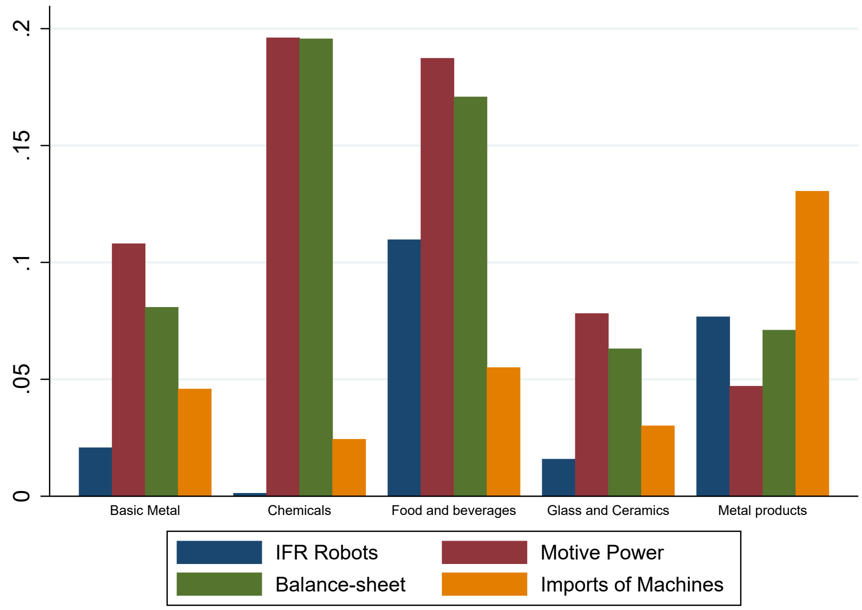

B. Top 5 Industries by Count of Robots

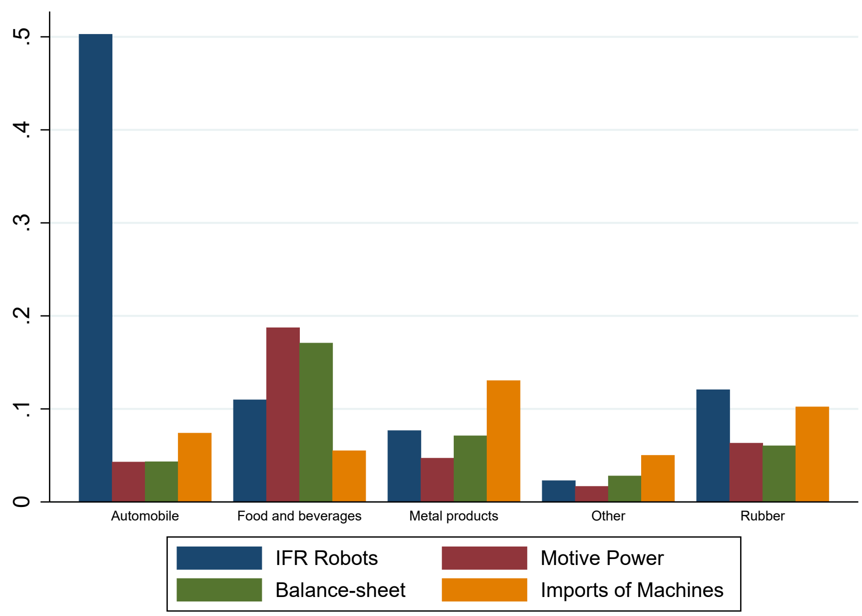

Notes: See Section 2 for a description of the data.

Figure 2: Firm-level Stylized Facts by Use of Industrial Machines

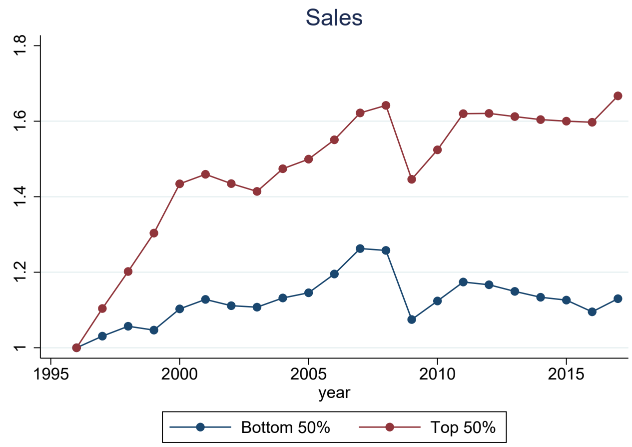

(a) Sales

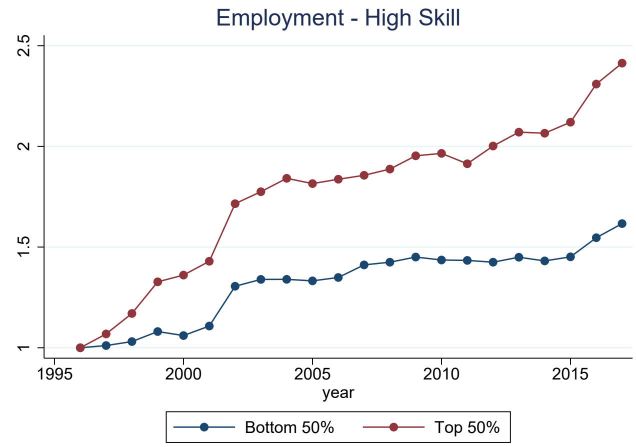

(b) High-skill Employment

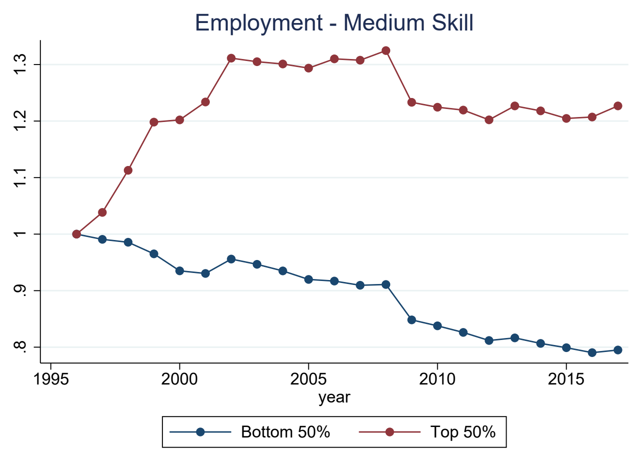

(c) Medium-skill Employment

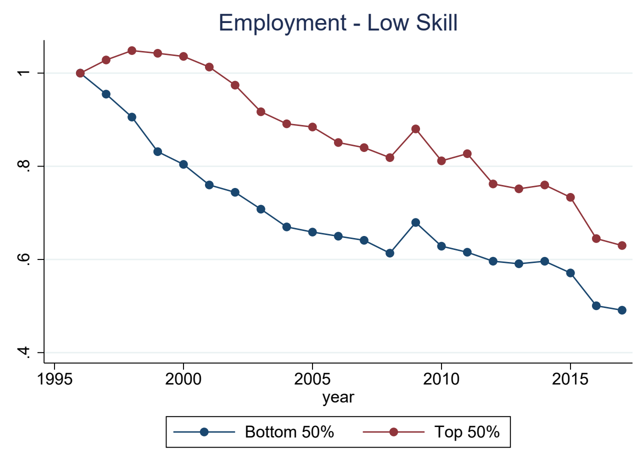

(d) Low-skill Employment

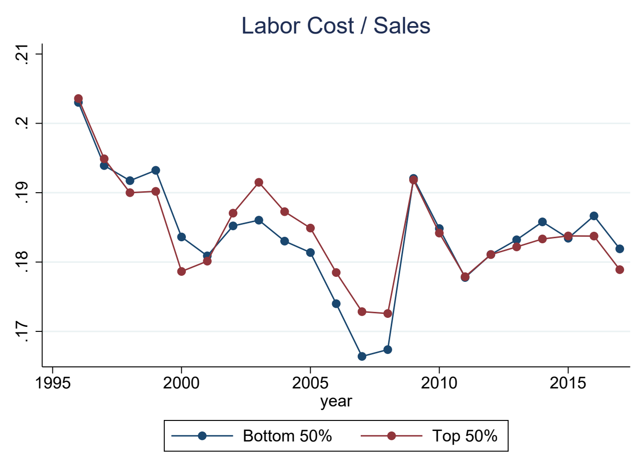

(e) Labor Share

Notes: This figure describes the path of sales, employment and the labor share for firms with a change in industrial machines above median from 1996 to 1999. All outcomes are normalized to one in 1996. See Section 3 for a description of the methodology.

Figure 3: Firm-Level Event Studies for Employment

A. 90th percentile of investment in industrial equipment

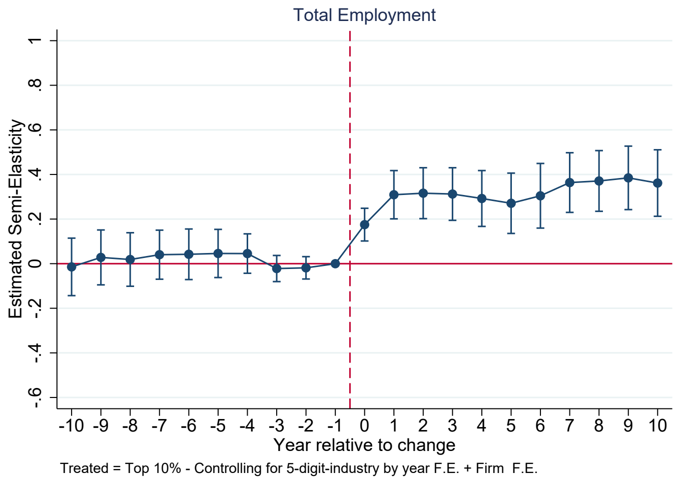

B. 75th percentile of investment in industrial equipment

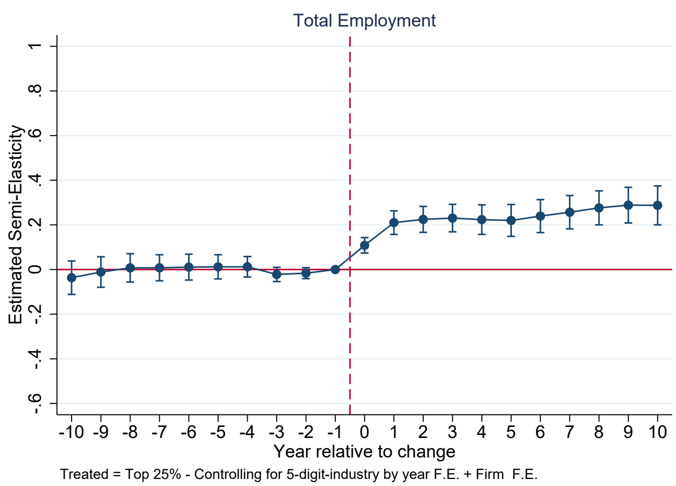

C. 50th percentile of investment in industrial equipment

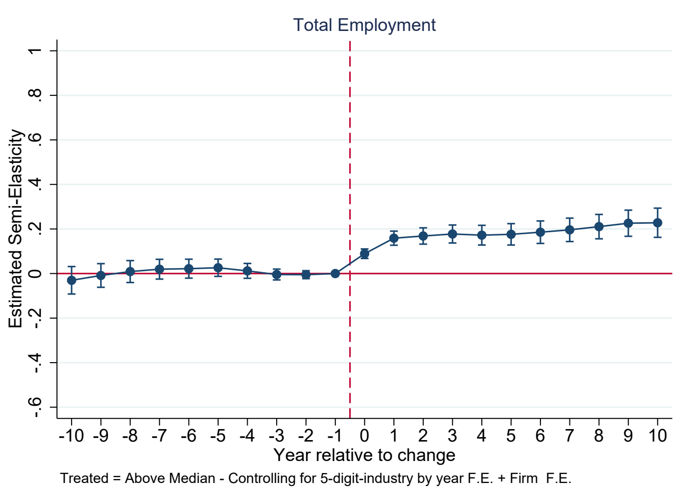

Notes: See Section 3 for the methodology.

Figure 4: Robustness across Measures of Automation

Panel A: Results in the Automobile Industry

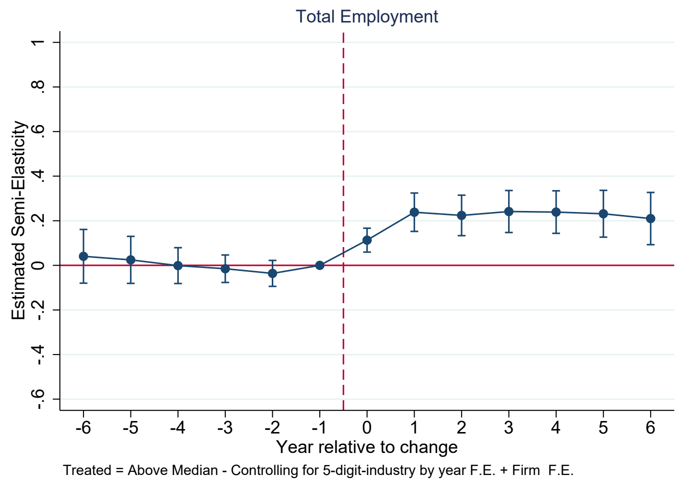

Panel B: Results with Imports of Industrial Automating Machines

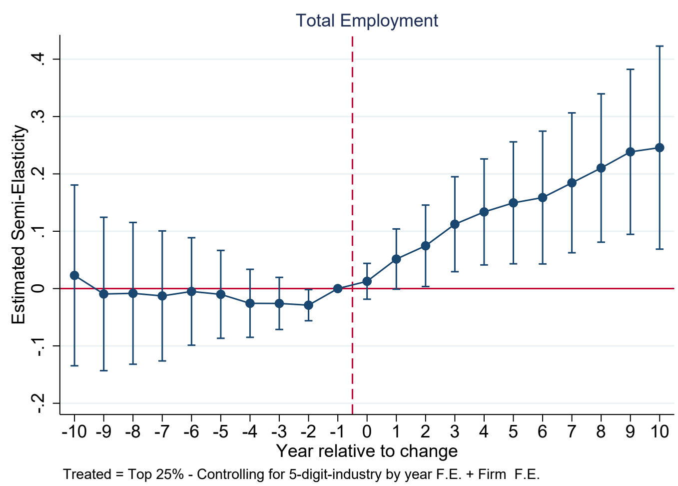

Panel C: Results with Motive Power (Distributed Lag)

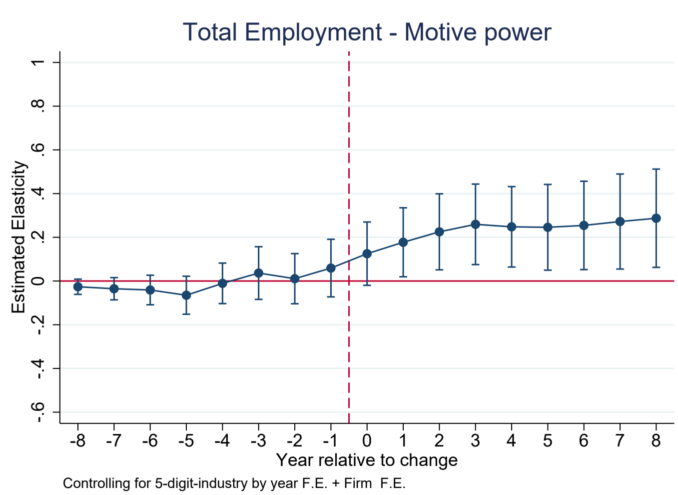

Notes: See Section 3 for the methodology.

Figure 5: Heterogeneity across Skill Groups

A. High-Skill Employment

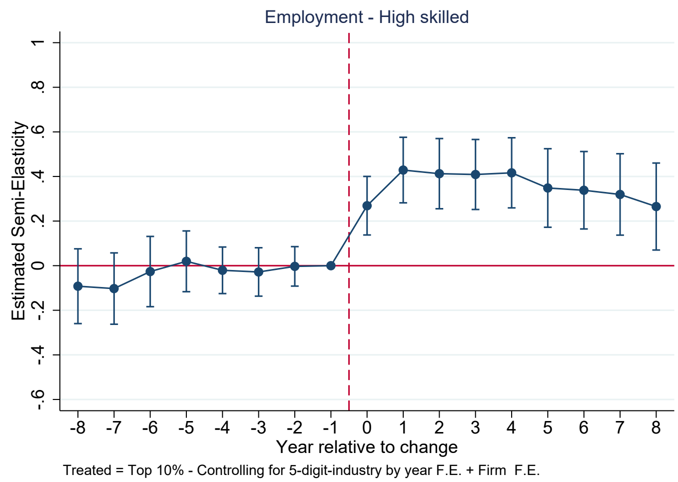

B. Medium-Skill Employment

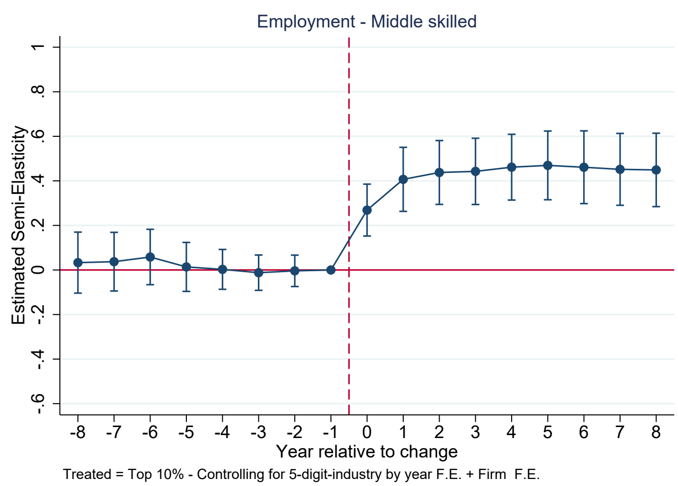

C. Low-Skill Employment

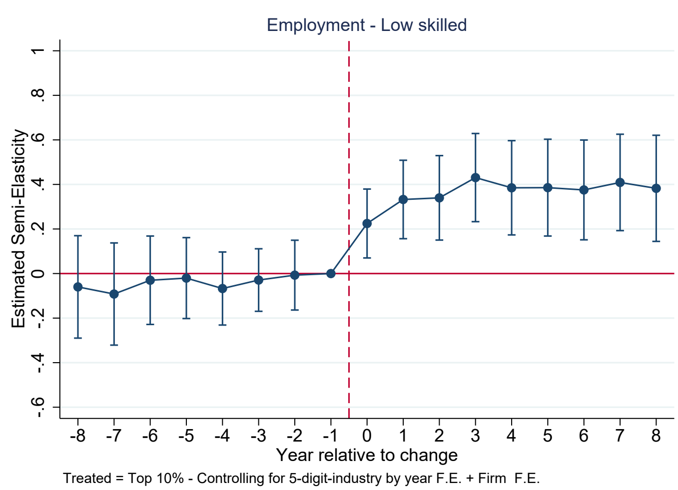

Notes: See Section 3 for the methodology.

Figure 6: Firm-Level Event Studies for Wages and Within-Firm Inequality

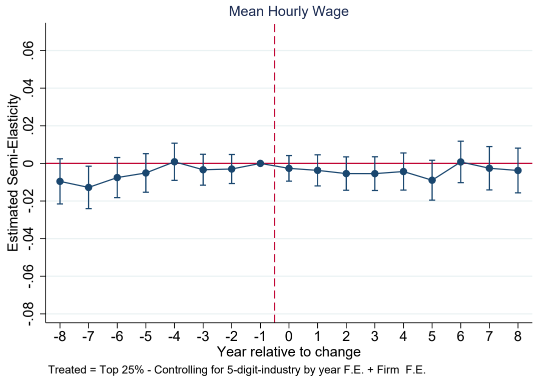

(i) Mean Hourly Wage, All

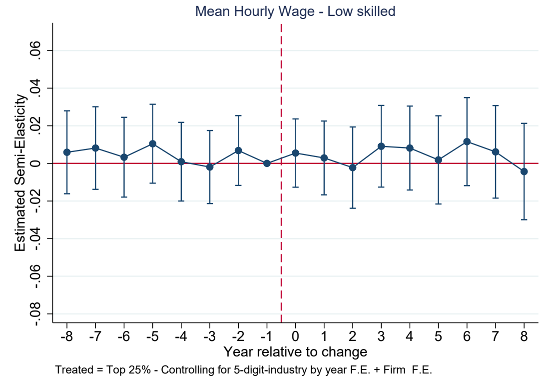

(ii) Mean Hourly Wage, Low Skill

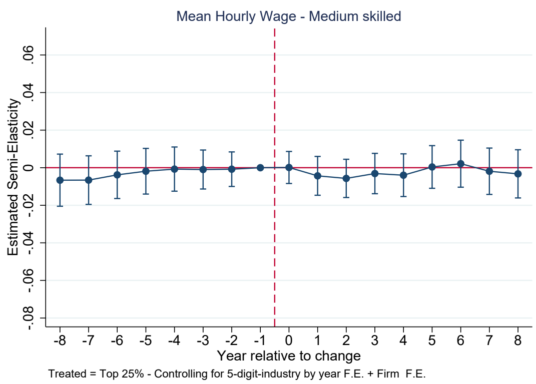

(iii) Mean Hourly Wage, Medium Skill

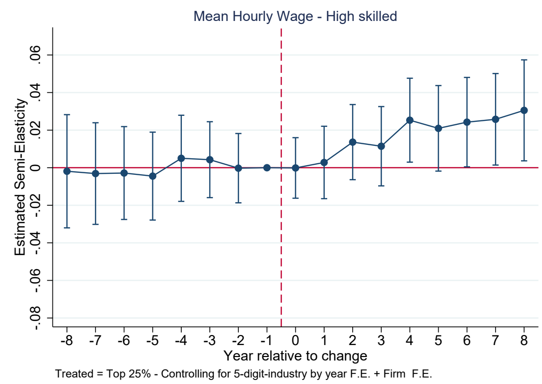

(iv) Mean Hourly Wage, High Skill

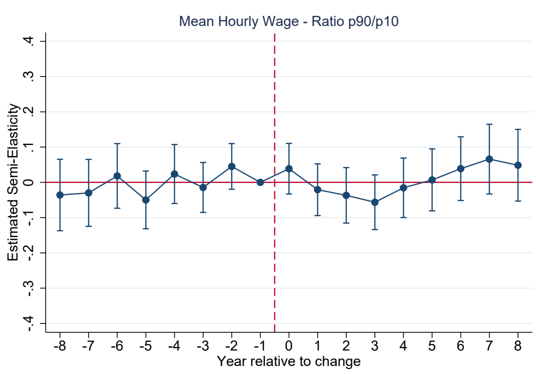

(v) P90/P10 of Wage Distribution

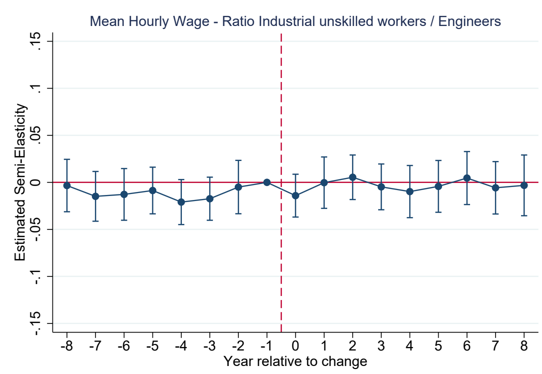

(vi) Wage Ratio Unskilled Workers / Engineers

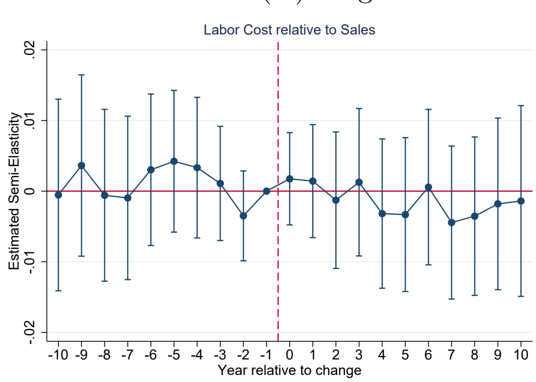

(vii) Labor Cost / Sales

Notes: See Section 3 for a description of the methodology.

Figure 7: Firm-Level Event Studies for Job Creation/Destruction

Panel A: Main Result with Investments in Industrial Equipment

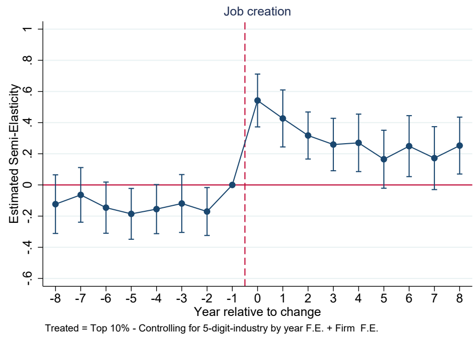

(i) Job Creation

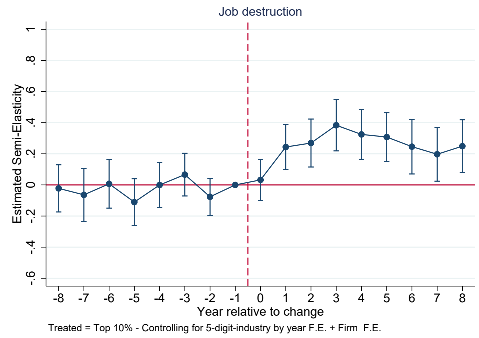

(ii) Job Destruction

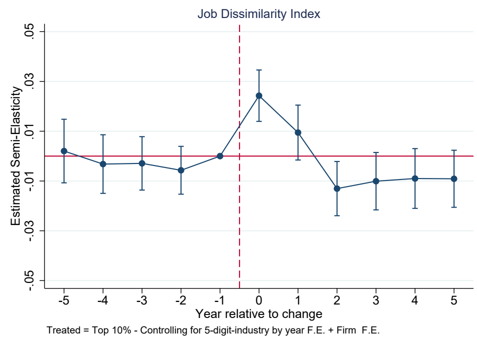

(iii) Job Dissimilarity Index

Panel B: Placebo Test with Investments in Real Estate

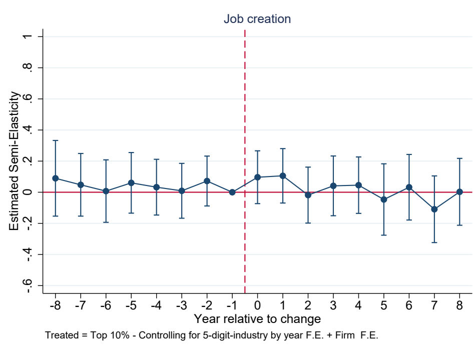

(i) Job Creation

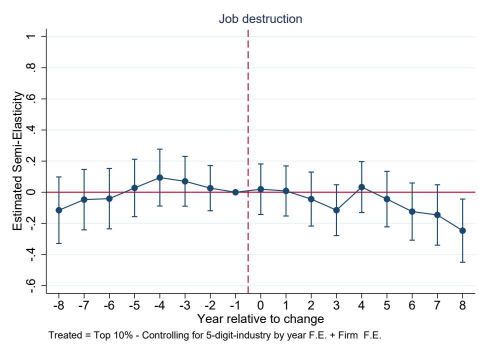

(ii) Job Destruction

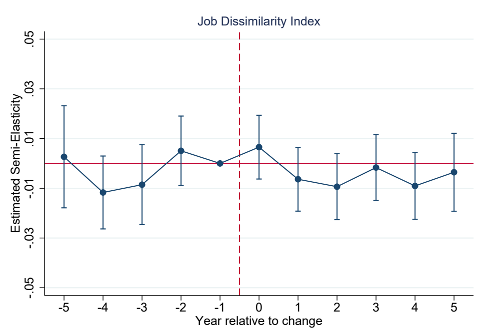

(iii) Job Dissimilarity Index

Notes: See Section 3 for a description of the methodology.

Figure 8: Firm-Level Event Studies for Market Dynamics

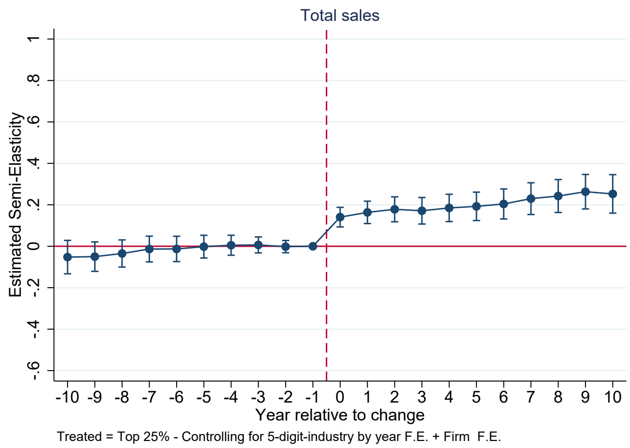

(i) Sales

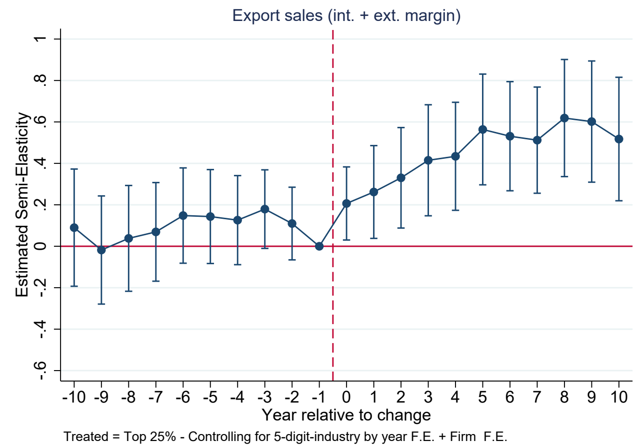

(ii) Log(1+Export Sales)

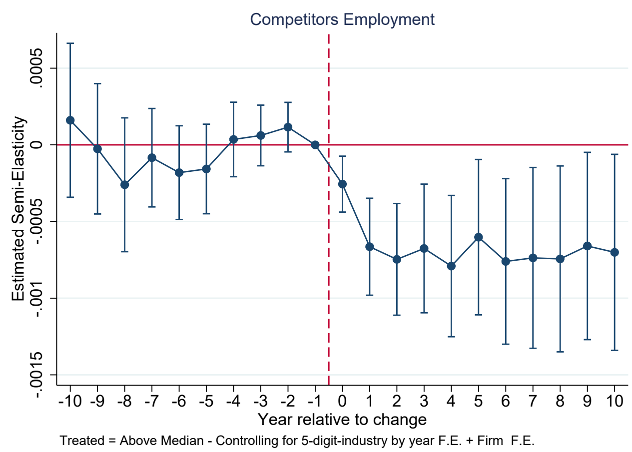

(iii) Business Stealing across Firms

Notes: See Section 3 for a description of the methodology.

Figure 9: Firm-Level Event Studies for Prices

A. 90th percentile of investment for industry equipment, HS6 product level

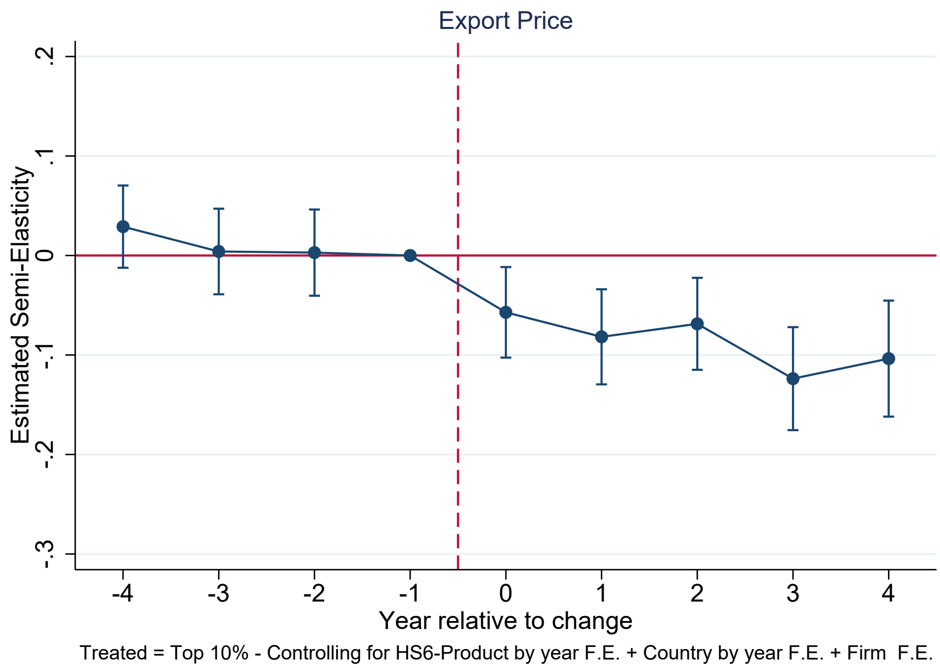

B. 90th percentile of investment for industry equipment, NC8 product level

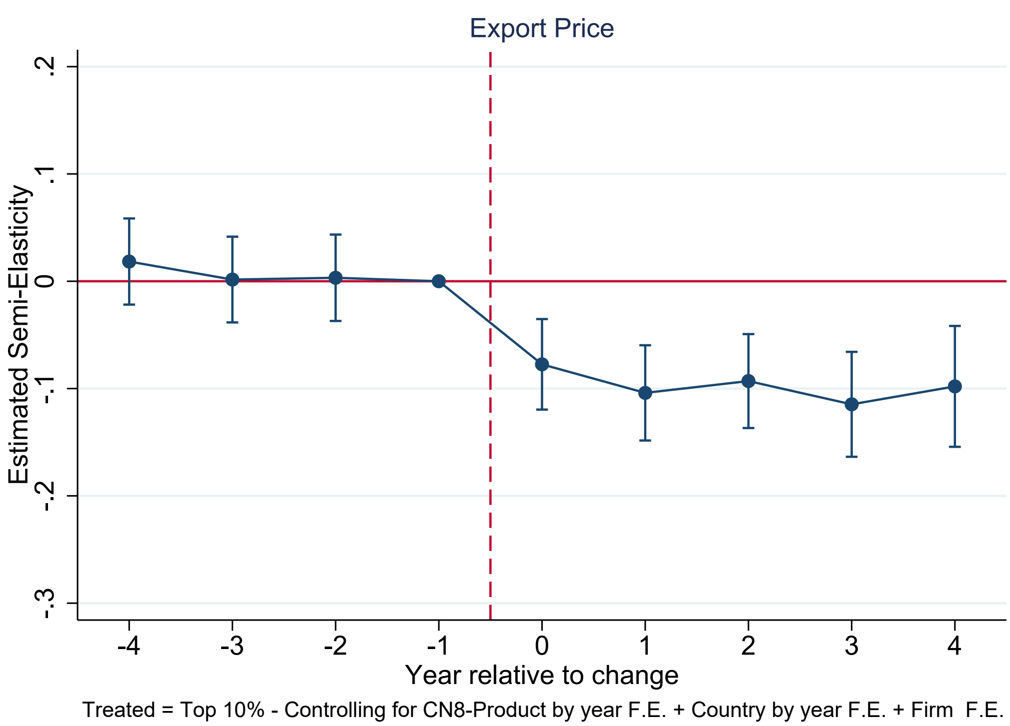

Notes: See Section 3 for a description of the methodology.

Figure 10: Persistence of Importer-Supplier Relationships

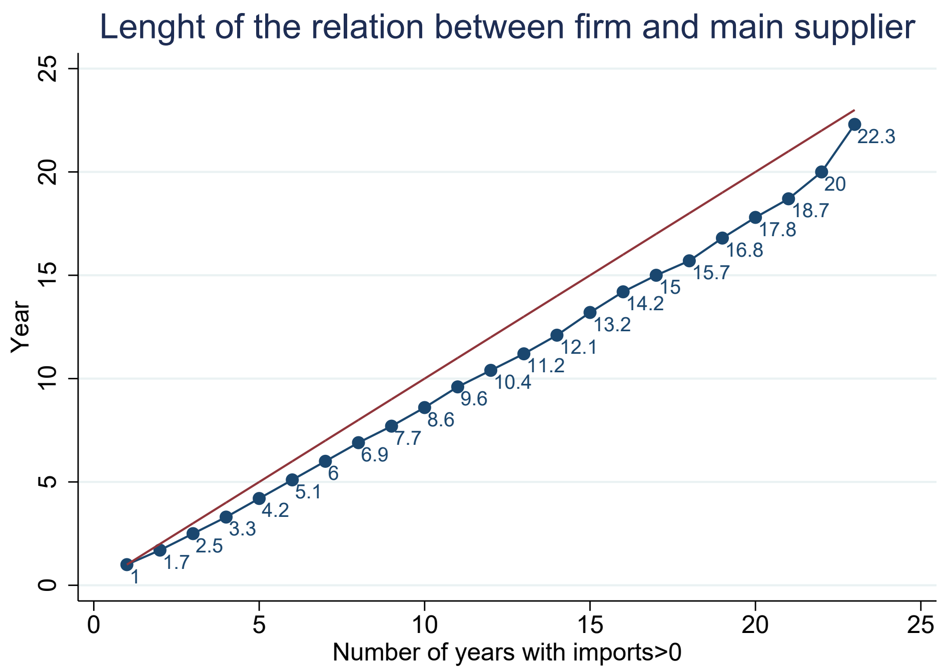

Notes: See Section 4 for a description of the methodology.

Figure 11: Firm-level Shift-Share IV Design

A. First stage

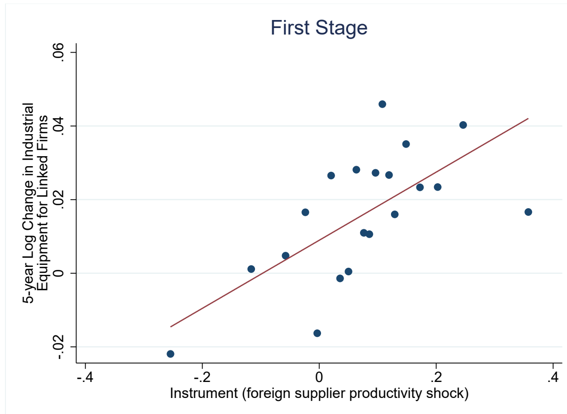

B. Reduced-form

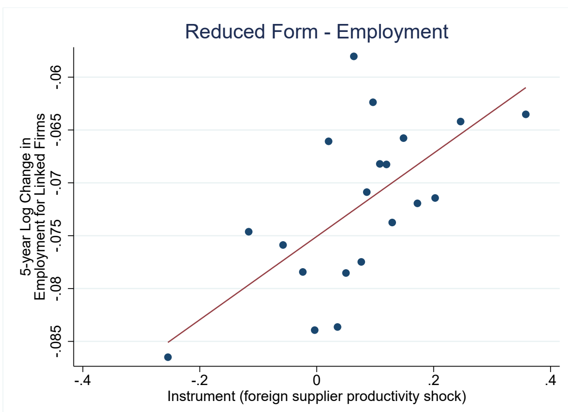

(i) Employment

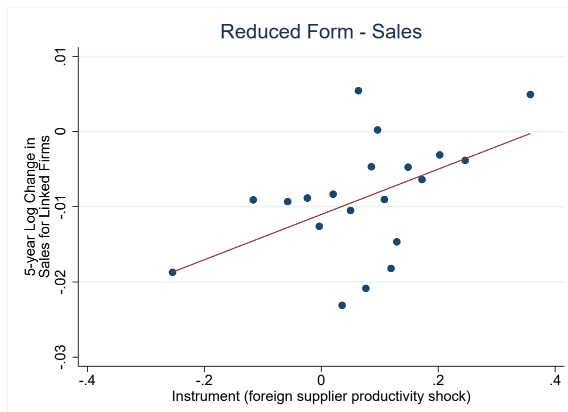

(ii) Sales

C. Falsification test

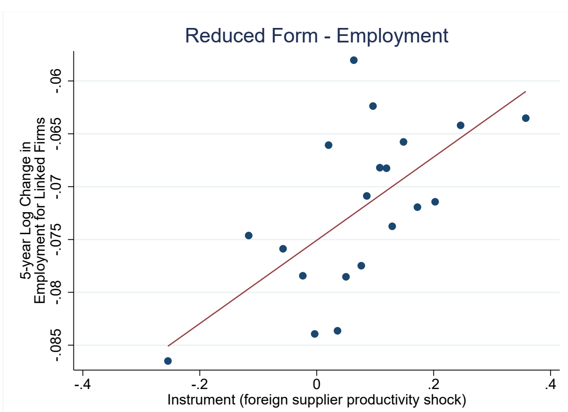

(i) Employment

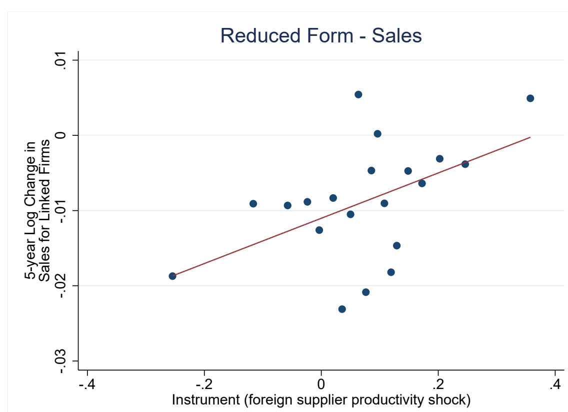

(ii) Sales

Notes: This figure uses HS6-level shocks measured in EU countries (except France) and Switzerland as instruments. Partner-period and 4-digit product by period fixed effects are used. Outcomes are measured at the firm level. Each dot represents 5% of the data. See Section 4 for a description of the shift-share methodology.

Figure 12: Industry-Level Event Studies

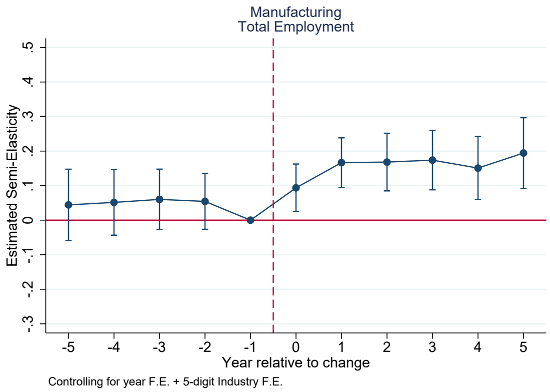

(i) Employment

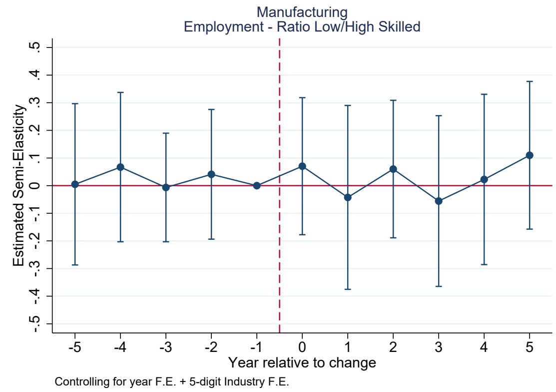

(ii) Employment Ratio, High- vs. Low-Skill

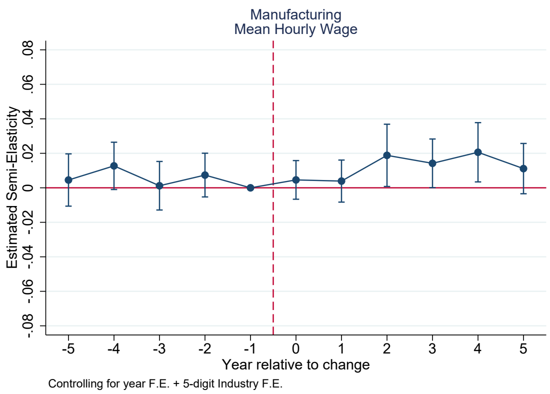

(iii) Mean Hourly Wage, All

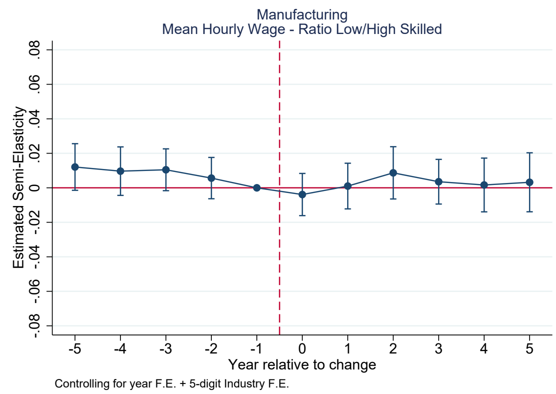

(iv) Hourly Wage Ratio, High- vs. Low-Skill

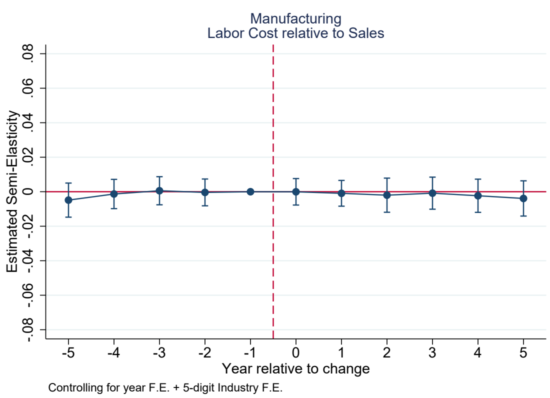

(v) Labor Cost to Sales

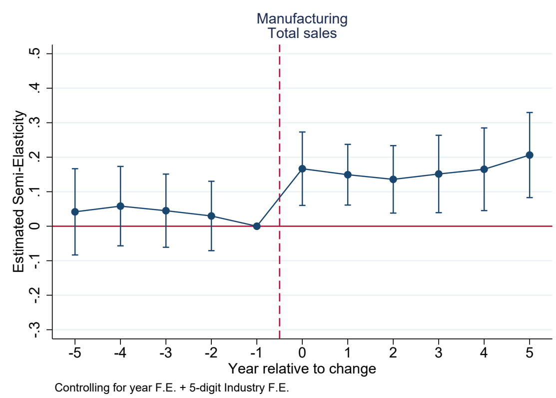

(vi) Sales

Notes: See Section V.A for a description of the methodology. The specifications include 5-digit industry fixed effects and year fixed effects.

Figure 13: Industry-level Shift-Share IV Design

A. First stage

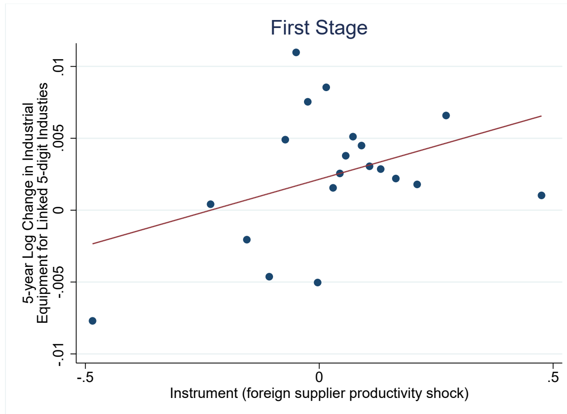

B. Reduced-form

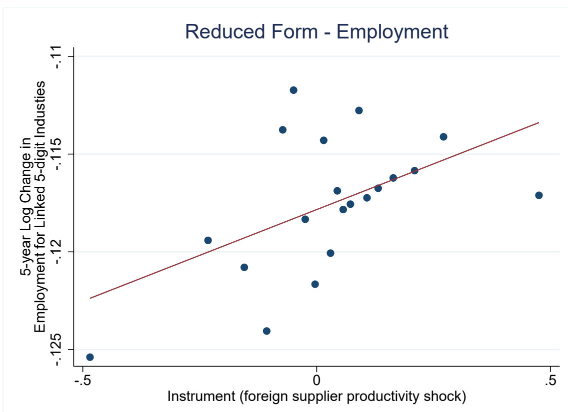

(i) Employment

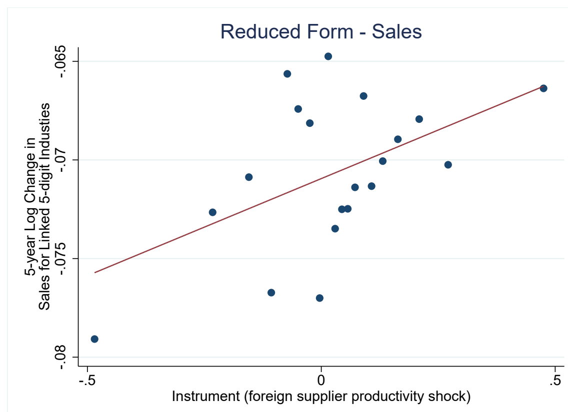

(ii) Sales

C. Falsification Tests

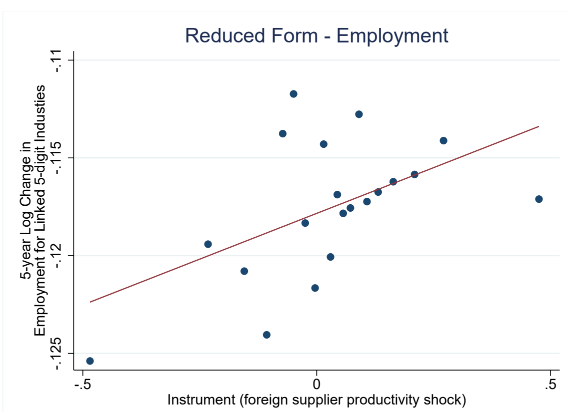

(i) Employment

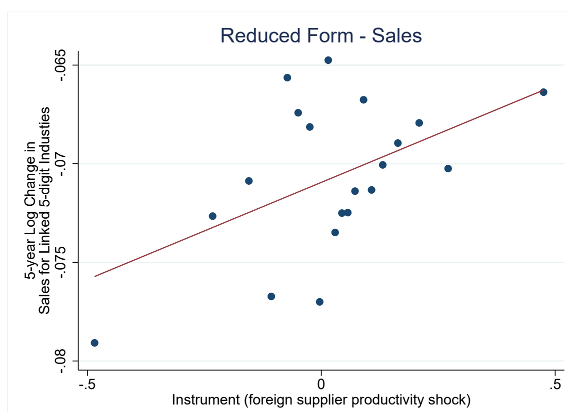

(ii) Sales

Notes: This figure uses HS6-level shocks measured in EU countries (except France) and Switzerland as instruments. Partner-period and 4-digit product by period fixed effects are used. Outcomes are measured at the 5-digit industry level. Each dot represents 5% of the data. See Section 4 for a description of the shift-share methodology.

Table 1: Summary Statistics

Panel A: Annual Levels, 1995–2017

	Units N	Units-by-year N	Mean	S.D.	p5	p25	p50	p75	p95
Plant level									
Employment	2,773	54,755	273	461	24	92	163	309	824
Automation – Motive Force (TOE)	2,773	54,755	1,381	4,943	34	143	353	917	5,489
Firm level									
Employment	1,599	33,579	55	158	3	7	15	40	204
Sales (thousands of euros)	1,599	33,579	18,289	97,344	350	972	2,404	7,801	52,228
Automation									
Industrial Machines (thousands of euros)	1,599	33,579	7,519	114,066	15	92	328	1,266	12,047
Imports of Machines (thousands of euros)	1,599	33,579	77	876	0	0	0	0	143
Automation – Motive Force (TOE)	485	7,910	2,144	8,506	41	147	346	1,041	7,353
Industry level									
Employment	255	5,865	10,868	15,215	491	2,498	6,287	13,192	34,630
Sales	255	5,865	3,463	8,887	117	669	1,764	3,636	9,761
Industrial Machines (millions of euros)	255	5,865	988	2,653	21	142	391	932	2,881

Panel B: Annual Changes, 1995–2017

	Units N	Units-by-year N	Mean	S.D.	p5	p25	p50	p75	p95
Plant level									
Employment	2,773	39,647	-2	48	-46	-9	-1	6	41
Automation – Motive Force (TOE)	2,773	39,647	-1	1,175	-309	-23	1	33	295
Firm level									
Employment	1,599	31,980	0	33	-7	-1	0	1	8
Sales (thousands of euros)	1,599	31,980	228	19,796	-2,041	-142	14	285	3,079
Automation									
Industrial Machines (thousands of euros)	1,599	31,980	142	6,008	-116	-4	1	34	557
Imports of Machines (thousands of euros)	1,599	31,980	3	634	-31	0	0	0	36
Automation – Motive Force (TOE)	485	7,161	2	1,325	-369	-21	2	36	369
Industry level									
Employment	255	5,610	-95	1,124	-1,148	-267	-38	96	876
Sales	255	5,610	22	2,018	-421	-56	4	82	521
Industrial Machines (millions of euros)	255	5,610	16	620	-67	-3	4	25	121

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

Table 2: Examples of Imported Industrial Automating Machines

Panel A: Randomly-Drawn Subset of 10 Machines

Name	Value of Imports, $	Share of Imports
Apparatus for dry-etching patterns on semiconductor materials	430,688	0.0035
Bending, folding, straightening or flattening machines	675,899	0.0056
Letterpress printing machinery, reel fed (excl. flexographic printing machinery)	122,370	0.00101
Machine tools for working any material by removal of material, operated by electro-discharge processes	129,927	0.00107
Machines for butt welding of metals	28,319	0.00023
Machines for preparing textile fibres (excl. carding, combing, drawing or roving machines)	134,543	0.0011
Machines for processing reactive resins	30,259	0.00025
Machining centres for working metal (excl. horizontal machining centres)	2,194,883	0.018
Parts of machinery and apparatus for soldering, brazing, welding or surface tempering	197,589	0.0016
Printing machinery for use in the production of semiconductors	5,056	0.000042

Panel B: Top 10 Machines by Value of Imports

Name	Value of Imports, $	Share of Imports
Machines, apparatus and mechanical appliances	5,640,191	0.046
Parts of machines and mechanical appliances having individual functions (excl. of cast iron or cast steel)	2,860,606	0.023
Parts of machinery for working rubber or plastics	2,811,272	0.023
Machines, apparatus and mechanical appliances	2,739,767	0.022
Parts of machinery for sorting, screening, separating, washing, crushing of modulding mineral susbtances	2,359,772	0.019
Parts of machinery for the industrial preparation or manufacture of food or drinks	2,203,525	0.018
Machining centres for working metal (excl. horizontal machining centres)	2,194,883	0.018
Parts of machines and mechanical appliances having individual functions	2,087,921	0.017
Industrial robots	2,021,562	0.016
Parts and accessories for machine tools for working metal without removing material	2,009,490	0.016

Notes: This table provides examples of our proxy for automation based on the taxonomy for machines in the customs data.

Table 3: Firm-level OLS Relationships with Automation

	(1)	(2)	(3)	(4)	(5)
Panel A: $\Delta_5 Employment$					
$\Delta_5 Machines$	0.409***	0.415***	0.414***	0.414***	0.413***
	(0.0212)	(0.0204)	(0.0206)	(0.0205)	(0.0205)
Panel B: $\Delta_5 Sales$					
$\Delta_5 Machines$	0.309***	0.325***	0.322***	0.320***	0.311***
	(0.0249)	(0.0251)	(0.0252)	(0.0254)	(0.0251)
Panel C: $\Delta_5 Hourly\,Wages$					
$\Delta_5 Machines$	-0.0367***	-0.0365***	-0.0366***	-0.0367***	-0.0373***
	(0.00554)	(0.00476)	(0.00494)	(0.00501)	(0.00497)
Panel D: $\Delta_5 Labor\ Cost\ Over\ Sales$					
$\Delta_5 Machines$	0.00300	0.00301	0.00328	0.00358	0.00502
	(0.00572)	(0.00551)	(0.00538)	(0.00551)	(0.00538)
Panel E: $\Delta_5 Profits$					
$\Delta_5 Machines$	0.349***	0.370***	0.367***	0.361***	0.347***
	(0.0494)	(0.0526)	(0.0529)	(0.0516)	(0.0525)
Panel F: $\Delta_5 Competitors'\ Employment$					
$\Delta_5 Machines$	-0.00513***	-0.00556***	-0.00558***	-0.00557***	-0.00564***
	(0.000851)	(0.000754)	(0.000754)	(0.000747)	(0.000750)
Partner-period F.E.	Yes	Yes	Yes	Yes	Yes
4-digit Product-period F.E.	Yes	Yes	Yes	Yes	Yes
2-digit Industry-period F.E.	Yes	Yes	Yes	Yes	Yes
Lagged Firm Controls		Yes	Yes	Yes	Yes
Lagged Machines			Yes	Yes	Yes
Lagged Other Capital				Yes	Yes
Contemporaneous Exports					Yes
N (Trading partner - Product - Period)	4,460	4,460	4,460	4,460	4,460

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

Table 4: Firm-level Effects of Automation with Shift-Share IV

	(1)	(2)	(3)	(4)	(5)
Panel A: $\Delta_5 Employment$					
$\Delta_5 Machines$	0.426***	0.425***	0.424***	0.433***	0.433***
	(0.0842)	(0.0997)	(0.0995)	(0.0984)	(0.0982)
Panel B: $\Delta_5 Sales$					
$\Delta_5 Machines$	0.325**	0.340***	0.340***	0.345***	0.346***
	(0.131)	(0.123)	(0.121)	(0.114)	(0.103)
Panel C: $\Delta_5 Hourly\,Wages$					
$\Delta_5 Machines$	-0.0718*	-0.0625	-0.0625	-0.0641	-0.0640
	(0.0377)	(0.0390)	(0.0390)	(0.0395)	(0.0391)
Panel D: $\Delta_5 Labor\ Cost\ Over\ Sales$					
$\Delta_5 Machines$	0.00453	0.00604	0.00607	0.00697	0.00686
	(0.0164)	(0.0173)	(0.0172)	(0.0166)	(0.0157)
Panel E: $\Delta_5 Profits$					
$\Delta_5 Machines$	0.995**	0.824*	0.824*	0.827*	0.828**
	(0.448)	(0.432)	(0.432)	(0.424)	(0.412)
Panel F: $\Delta_5 Competitors'\ Employment$					
$\Delta_5 Machines$	-0.00578*	-0.00920***	-0.00920***	-0.00914***	-0.00913***
	(0.00323)	(0.00330)	(0.00327)	(0.00331)	(0.00329)
Panel G: $Lagged\ \Delta_5 Employment$					
$\Delta_5 Machines$	-0.180	-0.198	-0.199	-0.199	-0.200
	(0.219)	(0.220)	(0.223)	(0.220)	(0.218)
Panel H: $Lagged\ \Delta_5 Sales$					
$\Delta_5 Machines$	0.0274	0.166	0.165	0.155	0.155
	(0.202)	(0.209)	(0.218)	(0.214)	(0.211)
First-Stage F	17.65	20.59	21.43	20.88	21.62
Partner-period F.E.	Yes	Yes	Yes	Yes	Yes
4-digit Product-period F.E.	Yes	Yes	Yes	Yes	Yes
2-digit Industry-period F.E.	Yes	Yes	Yes	Yes	Yes
Lagged Firm Controls		Yes	Yes	Yes	Yes
Lagged Machines			Yes	Yes	Yes
Lagged Other Capital				Yes	Yes
Contemporaneous Exports					Yes
N (Trading partner - Product - Period)	4,460	4,460	4,460	4,460	4,460

Notes: Each panel of this table corresponds to separate sets of five specifications with different outcomes. Section 4 describes the shift-share methodology. Standard errors and the first-stage F-statistics are clustered at the partner level. $*p < 0.1$, $**\ p < 0.05$, $***\ p < 0.01$.

Table 5: Industry-level OLS Relationships with Automation

	(1)	(2)	(3)	(4)
Panel A: $\Delta_5 Employment$				
$\Delta_5 Machines$	0.648***	0.647***	0.651***	0.651***
	(0.00933)	(0.00907)	(0.00988)	(0.00985)
Panel B: $\Delta_5 Incumbents' \ Employment$				
$\Delta_5 Machines$	0.324***	0.320***	0.325***	0.325***
	(0.0229)	(0.0211)	(0.0209)	(0.0208)
Panel C: $\Delta_5 Sales$				
$\Delta_5 Machines$	0.741***	0.741***	0.722***	0.721***
	(0.0143)	(0.0144)	(0.0128)	(0.0118)
Panel D: $\Delta_5 Hourly \ Wages$				
$\Delta_5 Machines$	0.0188***	0.0185***	0.0131***	0.0128***
	(0.00297)	(0.00283)	(0.00254)	(0.00239)
Panel E: $\Delta_5 Labor \ Cost \ Over \ Sales$				
$\Delta_5 Machines$	-0.0152***	-0.0153***	-0.0138***	-0.0136***
	(0.00183)	(0.00190)	(0.00180)	(0.00172)
Panel F: $\Delta_5 Profit$				
$\Delta_5 Machines$	0.973***	0.960***	0.979***	0.974***
	(0.0637)	(0.0583)	(0.0575)	(0.0520)
Partner-period F.E.	Yes	Yes	Yes	Yes
4-digit Product-period F.E.	Yes	Yes	Yes	Yes
Lagged Firm Controls	Yes	Yes	Yes	Yes
Lagged Machines		Yes	Yes	Yes
Lagged Other Capital			Yes	Yes
Contemporaneous Exports				Yes
N (Trading partner - Product - Period)				

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

Table 6: Industry-level Effects of Automation with Shift-Share IV

	(1)	(2)	(3)	(4)
Panel A: $\Delta_5 Employment$				
$\Delta_5 Machines$	1.011***	1.003***	1.006***	1.019***
	(0.213)	(0.212)	(0.218)	(0.225)
Panel B: $\Delta_5 Incumbents'\ Employment$				
$\Delta_5 Machines$	0.714**	0.645**	0.668**	0.674**
	(0.314)	(0.276)	(0.297)	(0.299)
Panel C: $\Delta_5 Sales$				
$\Delta_5 Machines$	1.063***	1.062***	0.976**	0.923**
	(0.383)	(0.387)	(0.409)	(0.378)
Panel D: $\Delta_5 Hourly\, Wages$				
$\Delta_5 Machines$	0.0610	0.0555	0.0450	0.0299
	(0.0881)	(0.0900)	(0.0984)	(0.102)
Panel E: Δ_5*Labor Cost Over Sales*				
$\Delta_5 Machines$	-0.0551	-0.0582	-0.0525	-0.0453
	(0.0559)	(0.0586)	(0.0617)	(0.0572)
Panel F: $\Delta_5 Profit$				
$\Delta_5 Machines$	2.806**	2.612**	2.850**	2.721**
	(1.320)	(1.250)	(1.347)	(1.277)
Panel G: *Lagged* $\Delta_5 Employment$				
$\Delta_5 Machines$	-0.00837	0.0159	0.0373	0.0564
	(0.386)	(0.385)	(0.400)	(0.398)
Panel H: *Lagged* $\Delta_5 Sales$				
$\Delta_5 Machines$	0.336	0.295	0.302	0.335
	(0.336)	(0.339)	(0.358)	(0.357)
First-Stage F	8.38	8.58	7.23	7.23
Partner-period F.E.	Yes	Yes	Yes	Yes
4-digit Product-period F.E.	Yes	Yes	Yes	Yes
Lagged Firm Controls	Yes	Yes	Yes	Yes
Lagged Machines		Yes	Yes	Yes
Lagged Other Capital			Yes	Yes
Contemporaneous Exports				Yes
N (Trading partner - Product - Period)	7,687	7,687	7,687	7,687

Notes: Each panel of this table corresponds to separate sets of five specifications with different outcomes. Section 4 describes the shift-share methodology. Standard errors and the first-stage F-statistics are clustered at the partner level. $^*p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: The Role of International Business-Stealing Effects – Industry Level

Panel A: Automation, Employment, and International Competition

	Δ Employment 1996-2017		
		International Competition	
	All industries	Above Median	Below Median
	(1)	(2)	(3)
Δ Machines 1996-2017	**0.345***	**0.404***	**0.171**
	(0.059)	(0.055)	(0.133)
Δ Other types of capital 1996-2017	✓	✓	✓
N	255	121	134

Panel B: Automation, Sales, and International Competition

	Δ Sales 1996-2017		
		International Competition	
	All industries	Above Median	Below Median
	(1)	(2)	(3)
Δ Machines 1996-2007	**0.427***	**0.510***	**0.188**
	(0.066)	(0.084)	(0.121)
Δ Other types of capital 1996-2017	✓	✓	✓
N	255	121	134

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 five industries with high usage of motive force. See Section 2 for a description of the data.

Figure A2: Consumer Income and Use of Automation Technologies

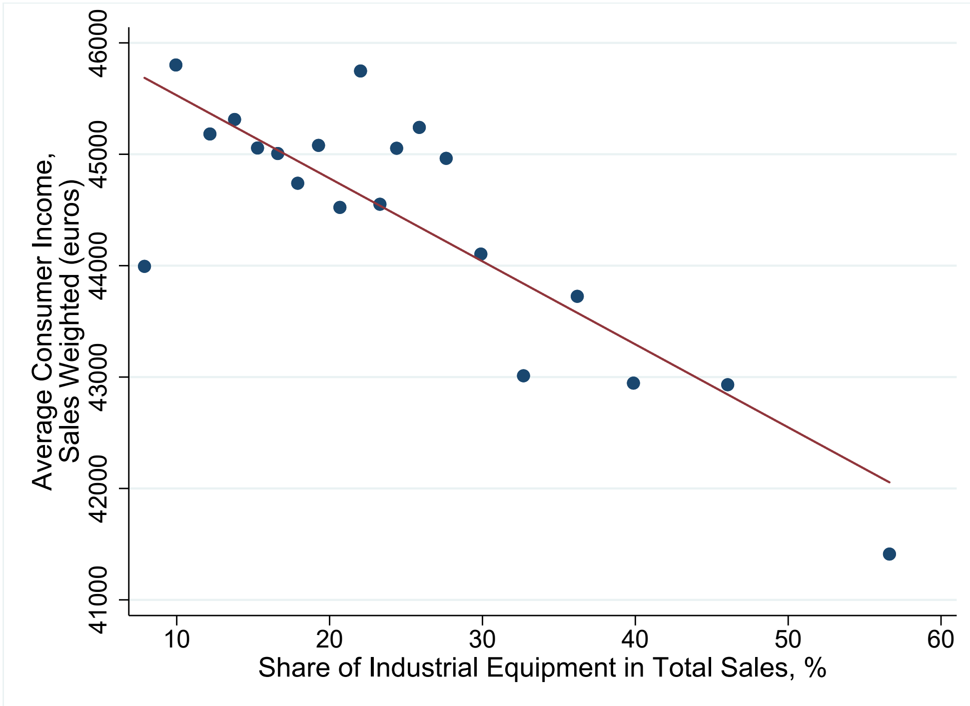

Notes: This figure reports the relationship between the average income of consumers an industry sells to and the value of industrial equipment as a share of total sales for this industry. The average consumer income is computed using sales weights. Similar patterns hold when using average total household expenditures as the outcome, as a proxy for households' permanent incomes.

Figure A3: Robustness of Event Study with Alternative Measure of Investment

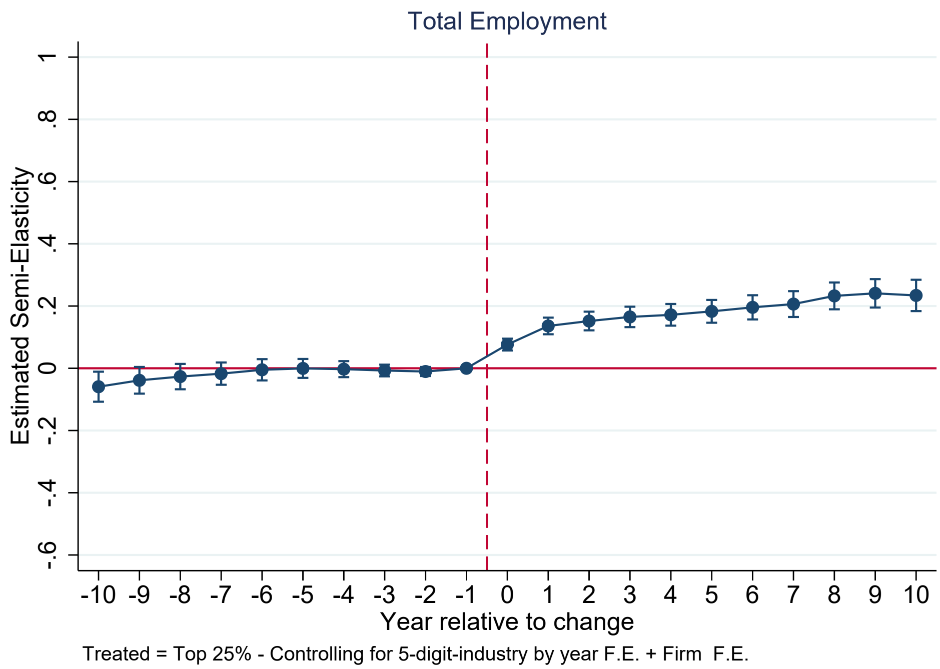

Notes: This figure uses investment in industrial equipment as a fraction of the initial balance-sheet value of the stock of machines. This measure is not sensitive to depreciation.

Figure A4: Robustness of Event Study with Alternative Fixed Effects

Panel A: With 2-digit-by-year Fixed Effects

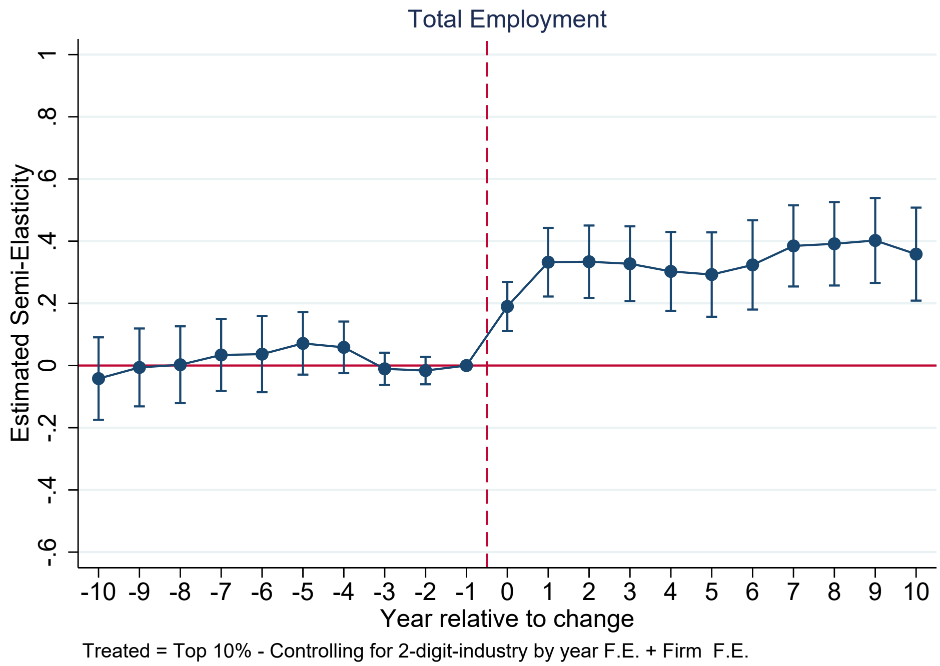

Panel A: With 4-digit-by-year Fixed Effects

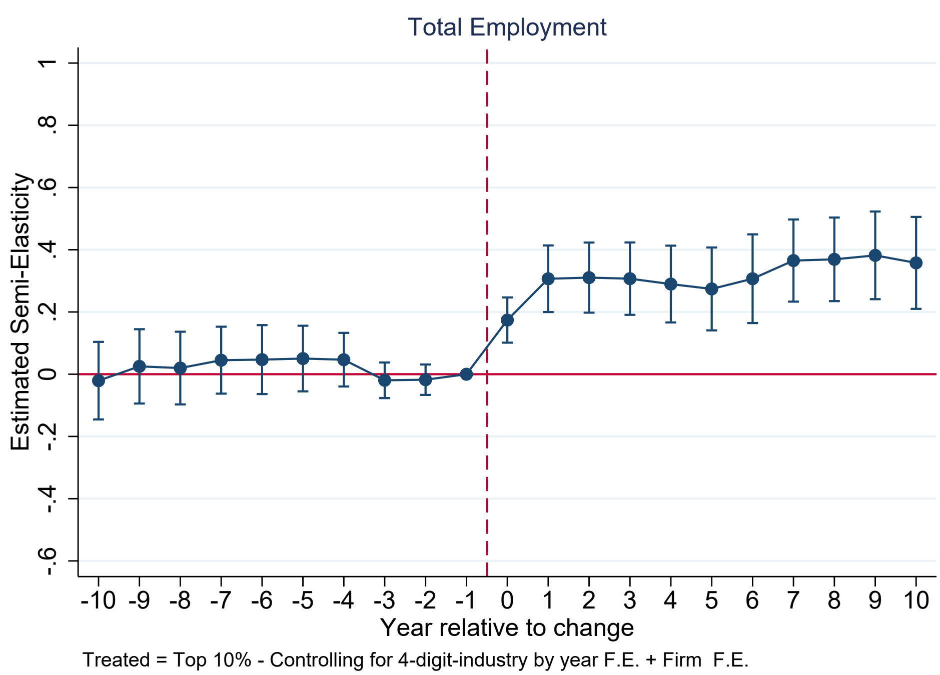

Notes: The specification in this figure is given by equation (1), with 2-digit by year or 4-digit by year fixed effects.

Figure A5: Distributed Lag using Peak Capacity for Motive Power

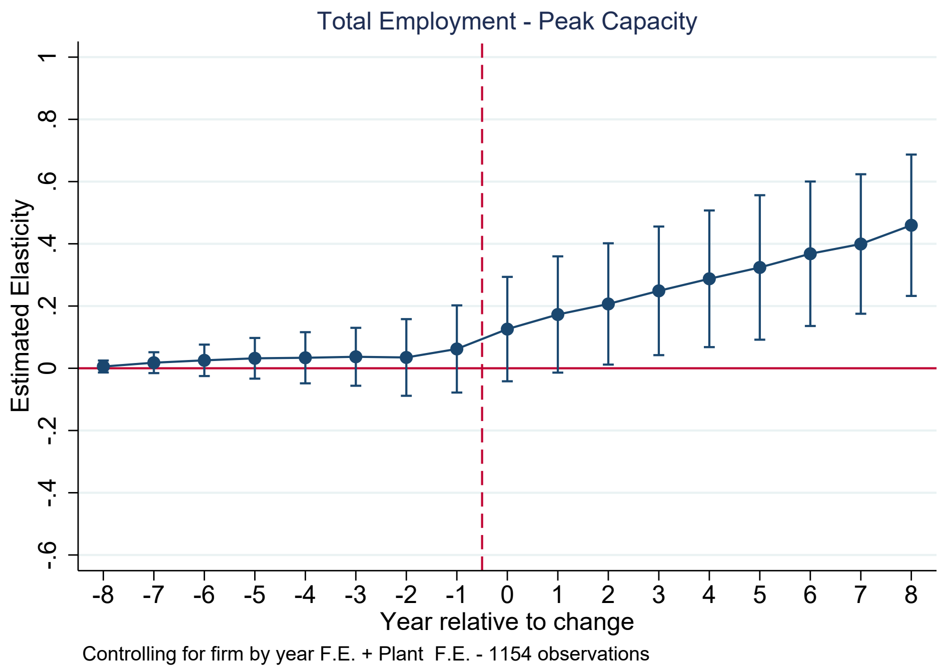

Notes: The specification in this figure uses peak capacity for motive power, instead of actual electricity consumption for motive power as in Panel C of Figure 4. Firm-by-year fixed effects and plant fixed effects are used.

Figure A6: The Employment Response for Unskilled Industrial Workers

Panel A: Plant level (Distributed Lag)

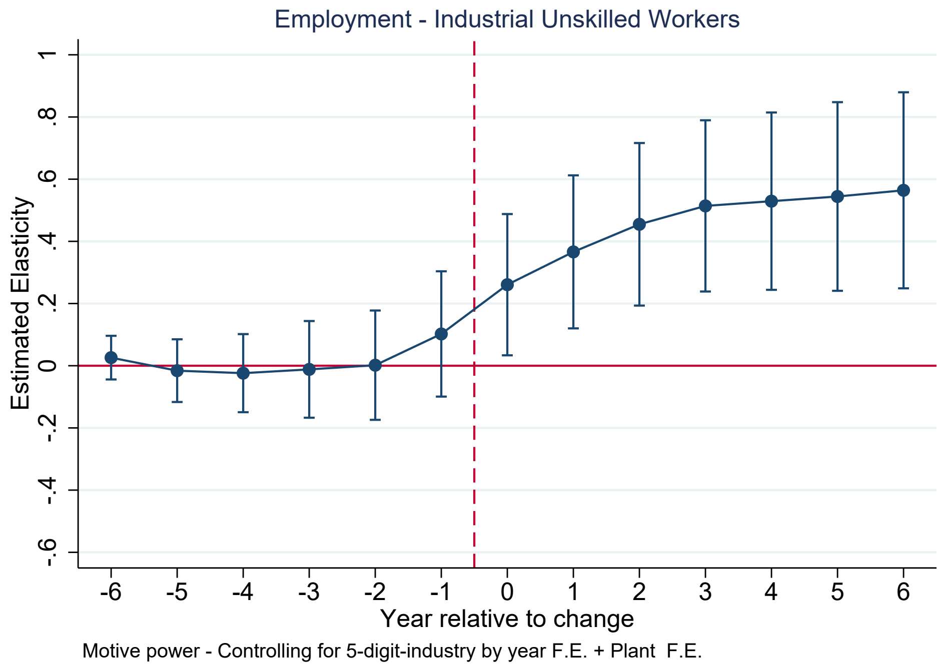

Panel B: Firm level (Event Study)

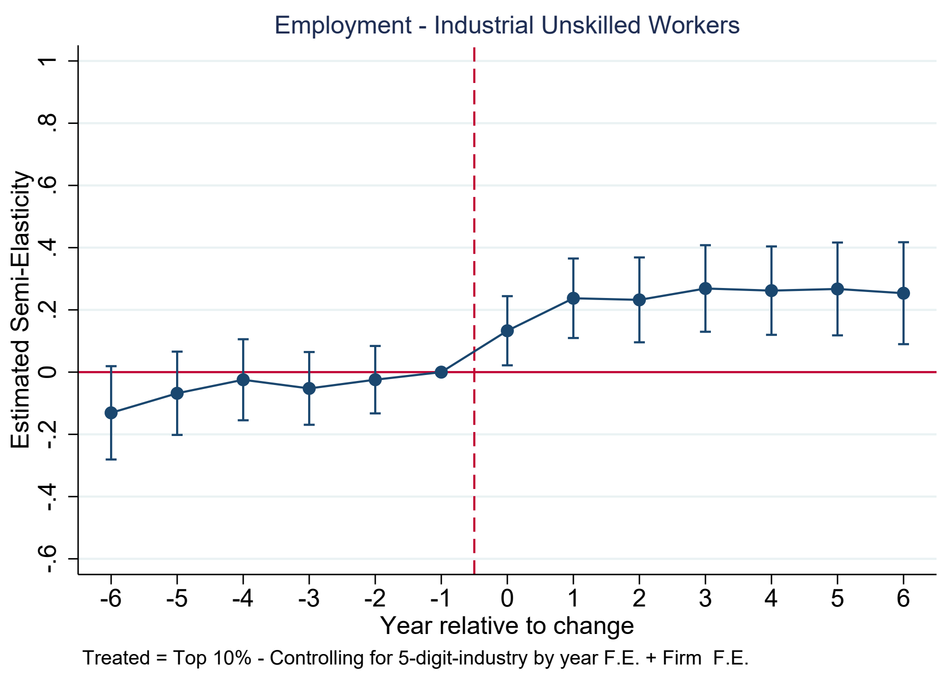

Notes: See Section 3 for the methodology.

Figure A7: Firm-Level Event Studies for Prices, Robustness with p95 Investment Threshold

A. 90th percentile of investment for industry equipment, HS6 product level

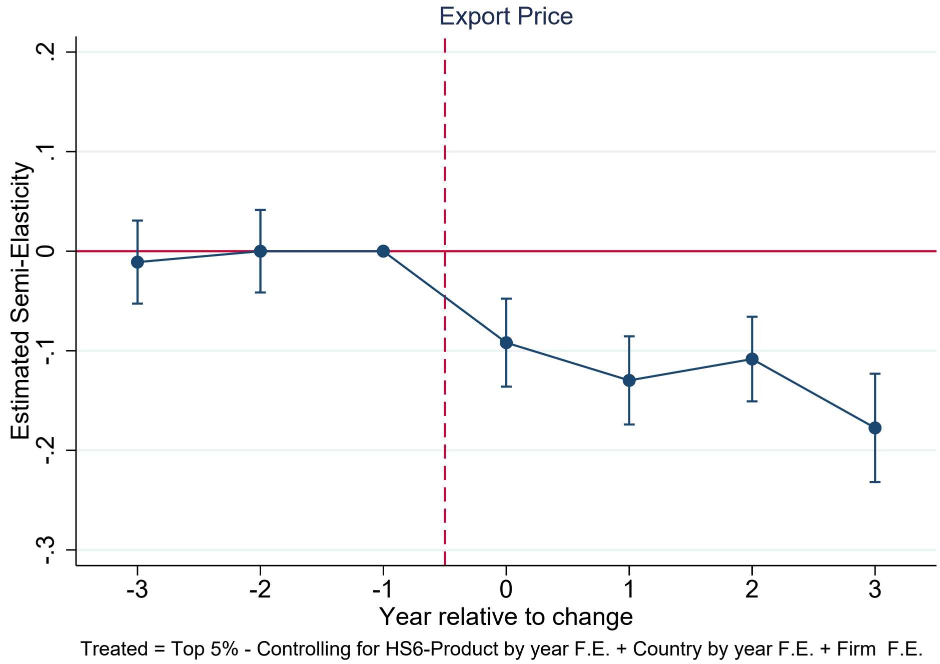

B. 90th percentile of investment for industry equipment, NC8 product level

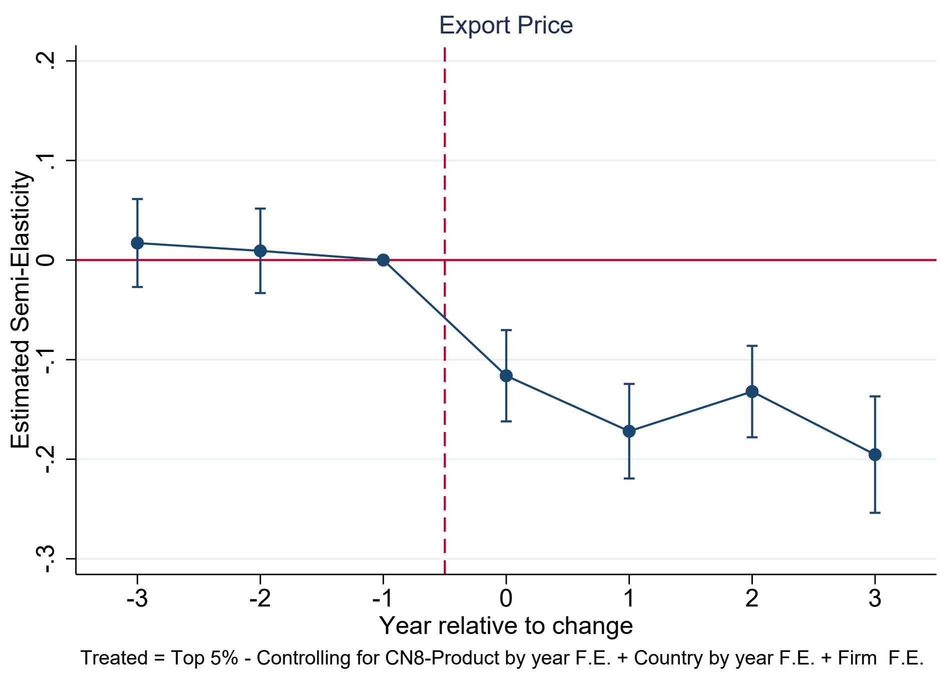

Notes: See Section 3 for a description of the methodology.

Figure A8: Firm-level Shift-Share IV, Robustness with More Stringent Fixed Effects

A. First stage

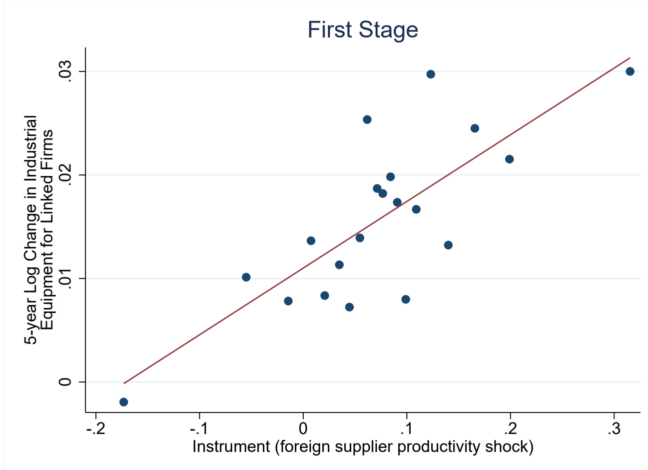

B. Reduced-form

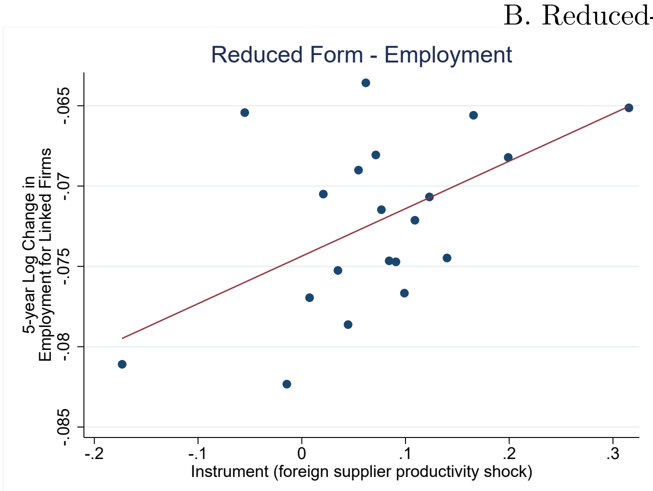

(i) Employment

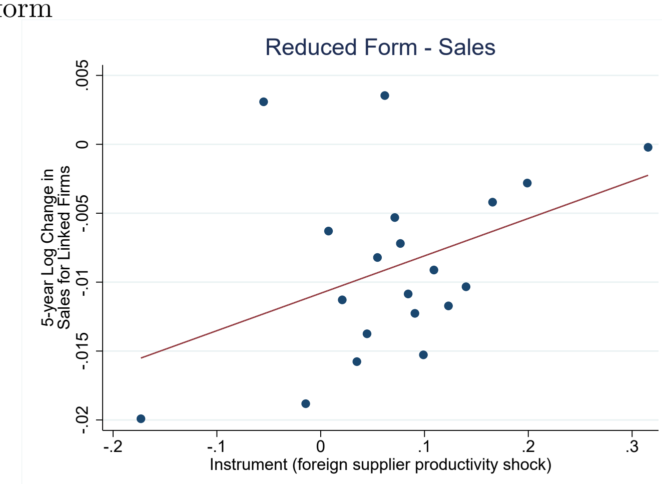

(ii) Sales

C. Falsification test

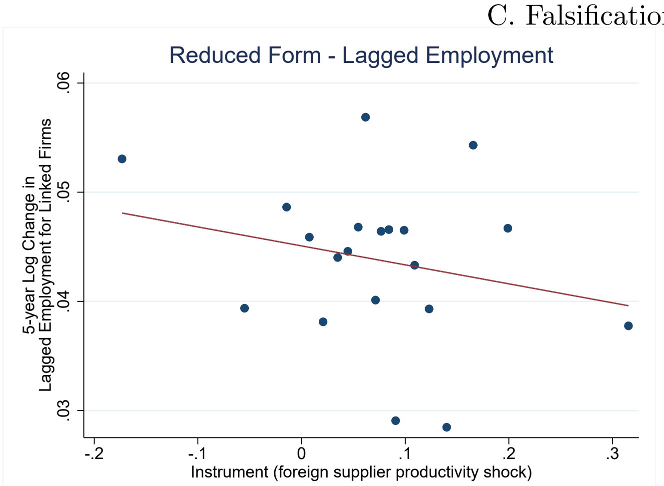

(i) Employment

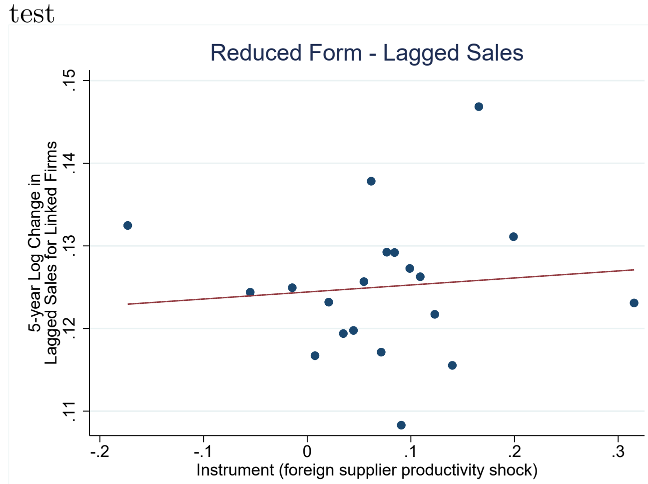

(ii) Sales

Notes: This figure uses HS6-level shocks measured in EU countries (except France) and Switzerland as instruments. Partner-period, 6-digit product by period fixed effects, and 5-digit industry be period fixed effects are used. Each dot represents 5% of the data. See Section 4 for a description of the shift-share methodology.

Figure A9: Firm-level Shift-Share IV, Robustness with Less Stringent Fixed Effects

A. First stage

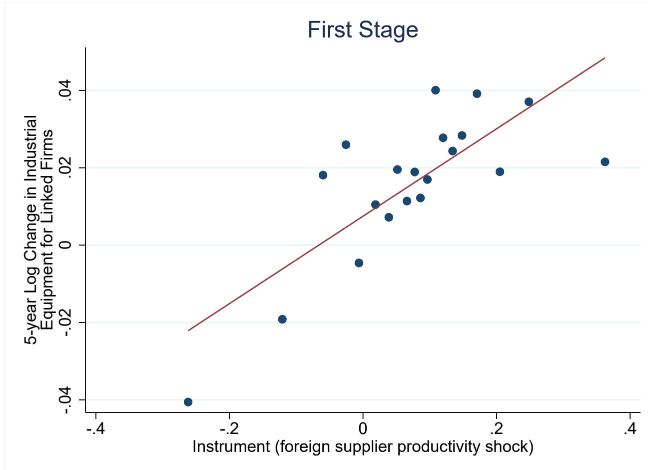

B. Reduced-form

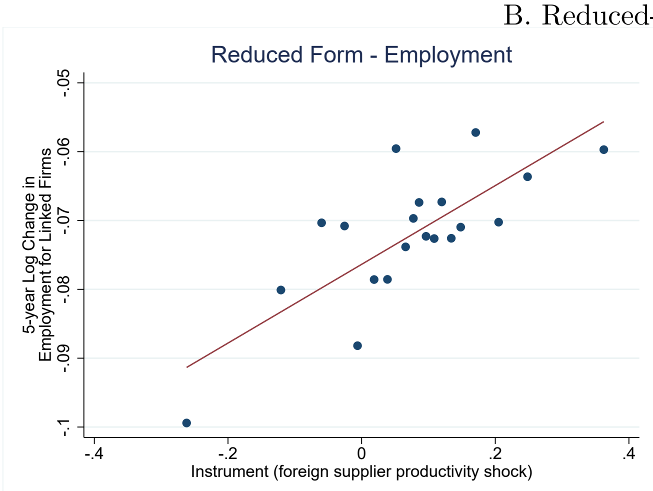

(i) Employment

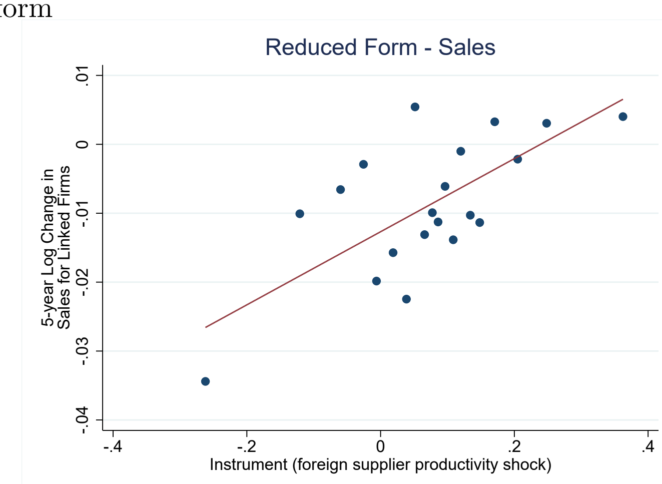

(ii) Sales

C. Falsification test

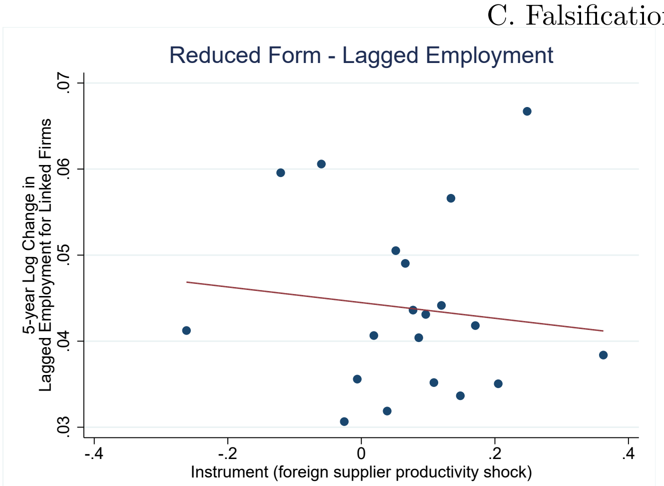

(i) Employment

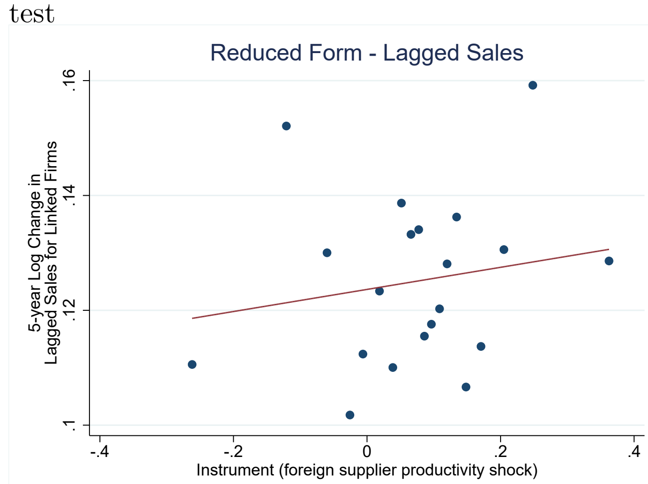

(ii) Sales

Notes: This figure uses HS6-level shocks measured in EU countries (except France) and Switzerland as instruments. Partner-period and 4-digit product-period fixed effects are used. Each dot represents 5% of the data. See Section 4 for a description of the shift-share methodology.

Table A1: Firm-level Shift-Share IV, Robustness with More Stringent Fixed Effects

	(1)	(2)	(3)	(4)	(5)
Panel A: $\Delta_5 Employment$					
$\Delta_5 Machines$	0.459***	0.434***	0.454***	0.456***	0.459***
	(0.0955)	(0.0970)	(0.103)	(0.103)	(0.105)
Panel B: $\Delta_5 Sales$					
$\Delta_5 Machines$	0.421***	0.361***	0.363**	0.359**	0.369***
	(0.156)	(0.140)	(0.148)	(0.146)	(0.141)
Panel C: $\Delta_5 Hourly\, Wages$					
$\Delta_5 Machines$	-0.0739	-0.0754	-0.0873*	-0.0879	-0.0872
	(0.0544)	(0.0489)	(0.0519)	(0.0538)	(0.0537)
Panel D: Δ_5*Labor Cost Over Sales*					
$\Delta_5 Machines$	-0.00487	-0.00745	-0.00553	-0.00426	-0.00562
	(0.0222)	(0.0202)	(0.0219)	(0.0207)	(0.0185)
Panel E: $\Delta_5 Profits$					
$\Delta_5 Machines$	0.559	0.429	0.430	0.413	0.434
	(0.593)	(0.534)	(0.573)	(0.544)	(0.538)
Panel F: $\Delta_5 Competitors'\ Employment$					
$\Delta_5 Machines$	-0.0869***	-0.0610***	-0.0645***	-0.0638***	-0.0641***
	(0.0259)	(0.0165)	(0.0175)	(0.0168)	(0.0166)
Panel G: *Lagged* $\Delta_5 Employment$					
$\Delta_5 Machines$	-0.270	-0.141	-0.190	-0.183	-0.183
	(0.313)	(0.235)	(0.262)	(0.263)	(0.262)
Panel H: *Lagged* $\Delta_5 Sales$					
$\Delta_5 Machines$	0.132	0.262	0.214	0.207	0.210
	(0.299)	(0.261)	(0.278)	(0.278)	(0.276)
First-Stage F	8.01	11.27	11.78	12.18	12.49
Partner-period F.E.	Yes	Yes	Yes	Yes	Yes
6-digit Product-period F.E.	Yes	Yes	Yes	Yes	Yes
5-digit Industry-period F.E.	Yes	Yes	Yes	Yes	Yes
Lagged Firm Controls		Yes	Yes	Yes	Yes
Lagged Machines			Yes	Yes	Yes
Lagged Other Capital				Yes	Yes
Contemporaneous Exports					Yes
N (Trading partner - Product - Period)	4,460	4,460	4,460	4,460	4,460

Notes: Each panel of this table corresponds to separate sets of five specifications with different outcomes. Section 4 describes the shift-share methodology. Standard errors and the first-stage F-statistics are clustered at the partner level. *$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Firm-level Shift-Share IV, Robustness with Less Stringent Fixed Effects

	(1)	(2)	(3)	(4)	(5)
Panel A: $\Delta_5 Employment$					
$\Delta_5 Machines$	0.506***	0.518***	0.517***	0.519***	0.519***
	(0.0647)	(0.0779)	(0.0780)	(0.0805)	(0.0818)
Panel B: $\Delta_5 Sales$					
$\Delta_5 Machines$	0.469***	0.536***	0.533***	0.527***	0.506***
	(0.135)	(0.119)	(0.117)	(0.114)	(0.109)
Panel C: $\Delta_5 Hourly\,Wages$					
$\Delta_5 Machines$	-0.0522	-0.0359	-0.0363	-0.0377	-0.0395
	(0.0392)	(0.0399)	(0.0400)	(0.0407)	(0.0407)
Panel D: $\Delta_5 Labor\ Cost\ Over\ Sales$					
$\Delta_5 Machines$	-0.0154	-0.0181	-0.0179	-0.0153	-0.0117
	(0.0199)	(0.0175)	(0.0174)	(0.0173)	(0.0167)
Panel E: $\Delta_5 Profits$					
$\Delta_5 Machines$	1.273***	1.317***	1.317***	1.287***	1.255***
	(0.411)	(0.394)	(0.397)	(0.403)	(0.403)
Panel F: $Lagged\, \Delta_5 Employment$					
$\Delta_5 Machines$	-0.0805	-0.136	-0.142	-0.140	-0.126
	(0.170)	(0.181)	(0.184)	(0.189)	(0.190)
Panel G: $Lagged\, \Delta_5 Sales$					
$\Delta_5 Machines$	0.170	0.204	0.195	0.191	0.201
	(0.174)	(0.188)	(0.193)	(0.196)	(0.197)
First-Stage F	26.17	32.30	34.58	31.99	32.16
Partner-period F.E.	Yes	Yes	Yes	Yes	Yes
4-digit Product-period F.E.	Yes	Yes	Yes	Yes	Yes
Lagged Firm Controls		Yes	Yes	Yes	Yes
Lagged Machines			Yes	Yes	Yes
Lagged Other Capital				Yes	Yes
Contemporaneous Exports					Yes
N (Trading partner - Product - Period)	4,460	4,460	4,460	4,460	4,460

Notes: Each panel of this table corresponds to separate sets of five specifications with different outcomes. Section 4 describes the shift-share methodology. Standard errors and the first-stage F-statistics are clustered at the partner level. *$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Industry-level Shift-Share IV, Robustness with Less Stringent Fixed Effects

	(1)	(2)	(3)	(4)
Panel A: $\Delta_5 Employment$				
$\Delta_5 Machines$	1.080***	1.076***	1.081***	1.091***
	(0.185)	(0.186)	(0.190)	(0.193)
Panel B: $\Delta_5 Incumbents'\ Employment$				
$\Delta_5 Machines$	0.608***	0.566***	0.583***	0.587***
	(0.208)	(0.184)	(0.197)	(0.197)
Panel C: $\Delta_5 Sales$				
$\Delta_5 Machines$	1.309***	1.312***	1.245***	1.207***
	(0.338)	(0.338)	(0.337)	(0.327)
Panel D: *Lagged* $\Delta_5 Employment$				
$\Delta_5 Machines$	-0.0318	-0.0176	-0.000564	0.0131
	(0.249)	(0.251)	(0.261)	(0.260)
Panel E: *Lagged* $\Delta_5 Sales$				
$\Delta_5 Machines$	0.0811	0.0521	0.0442	0.0655
	(0.264)	(0.257)	(0.271)	(0.268)
First-Stage F	17.98	18.03	15.53	15.53
Partner-period F.E.	Yes	Yes	Yes	Yes
4-digit Product F.E.	Yes	Yes	Yes	Yes
Lagged Industry Controls	Yes	Yes	Yes	Yes
Lagged Machines		Yes	Yes	Yes
Lagged Other Capital			Yes	Yes
Contemporaneous Exports				Yes
N (Trading partner - Product - Period)	7,687	7,687	7,687	7,687

Notes: Each panel of this table corresponds to separate sets of five specifications with different outcomes. Section 4 describes the shift-share methodology. Standard errors and the first-stage F-statistics are clustered at the partner level. *$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: The Role of International Business-Stealing Effects – Falsification Test at the Firm Level

Panel A: Automation, Employment, and International Competition

	Δ Employment 1996-2017		
		International Competition	
	All industries	Above Median	Below Median
	(1)	(2)	(3)
Δ Machines 1996-2017	**0.324*****	**0.316*****	**0.329*****
	(0.010)	(0.017)	(0.012)
5-digit industry F.E.	✓	✓	✓
Δ Other types of capital 1996-2017	✓	✓	✓
N	5,375	1,921	3,454

Panel B: Automation, Sales, and International Competition

	Δ Sales 1996-2017		
		International Competition	
	All industries	Above Median	Below Median
	(1)	(2)	(3)
Δ Machines 1996-2007	**0.388*****	**0.348*****	**0.408*****
	(0.011)	(0.019)	(0.014)
5-digit industry F.E.	✓	✓	✓
Δ Other types of capital 1996-2017	✓	✓	✓
N	5,375	1,921	3,454

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

Table A5: Industry-level Relationship between Automation and Producer Prices, OLS

	(1)	(2)	(3)	(4)
$\Delta_5 Machines$	-0.113**	-0.199***	-0.194***	-0.178**
	(0.0573)	(0.0698)	(0.0699)	(0.0750)
2-digit Industry-period F.E.	Yes	Yes	Yes	Yes
4-digit Industry F.E.	Yes	Yes	Yes	Yes
Lagged Industry Controls	Yes	Yes	Yes	Yes
Lagged Machines		Yes	Yes	Yes
Lagged Other Capital			Yes	Yes
Contemporaneous Exports				Yes
N (Industry by Period)	183	183	183	183

Notes: The level of observation is a 4-digit industry by year. *$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.