

Robots and firms

Michael Koch, Ilya Manuylov, Marcel Smolka

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Robots and firms

Abstract

We study the implications of robot adoption at the level of individual firms using a rich panel data-set of Spanish manufacturing firms over a 27-year period (1990-2016). We focus on three central questions: (1) Which firms adopt robots? (2) What are the labor market effects of robot adoption at the firm level? (3) How does firm heterogeneity in robot adoption affect the industry equilibrium? To address these questions, we look at our data through the lens of recent attempts in the literature to formalize the implications of robot technology. As for the first question, we establish robust evidence that ex-ante larger and more productive firms are more likely to adopt robots, while ex-ante more skill-intensive firms are less likely to do so. As for the second question, we find that robot adoption generates substantial output gains in the vicinity of 20-25% within four years, reduces the labor cost share by 5-7%-points, and leads to net job creation at a rate of 10%. These results are robust to controlling for non-random selection into robot adoption through a difference-in-differences approach combined with a propensity score reweighting estimator. Finally, we reveal substantial job losses in firms that do not adopt robots, and a productivity-enhancing reallocation of labor across firms, away from non-adopters, and toward adopters.

JEL-Codes: D220, F140, J240, O140.

Keywords: automation, robots, firms, productivity, technology.

Michael Koch
University of Bayreuth
Department of Law and Economics
Universitätsstr. 30
Germany – 95447 Bayreuth
michael.koch@uni-bayreuth.de

Ilya Manuylov
Aarhus University, Department of Economics
and Business Economics
Fuglesangs Allé 4, Building 2632
Denmark – 8210 Aarhus V
imanuylov@econ.au.dk

Marcel Smolka
Aarhus University, Department of Economics and Business Economics
Fuglesangs Allé 4, Building 2632
Denmark – 8210 Aarhus V
msmolka@econ.au.dk

April 7, 2019

Financial support by the Tuborg Foundation and the Carlsberg Foundation is gratefully acknowledged.

1 Introduction

The rise of robot technology has sparked an intense debate about the labor market effects of robot adoption.¹ A key concern in this debate is that robots “steal” jobs from humans. A recent literature fuels this concern, finding large negative effects of robots on employment and wages across U.S. commuting zones (Acemoglu and Restrepo, 2017). However, a considerable challenge in this literature is the lack of *micro-level* information on actual robot use. The few existing studies all resort to *macro-level* information by industry to construct measures of local robot exposure. While this approach is an important first step in gauging the local labor market effects of robots, it makes the crucial assumption that all firms in a given industry have the same ability and willingness to adopt robots. It does not take seriously the possibility that the same set of tasks is performed by human labor in some firms, while it is performed by robots in others. The literature thus provides little insight into the following central questions: Which firms adopt robots? What are the labor market effects of robot adoption at the firm level? And how do differences in robot adoption across firms affect the industry equilibrium? Addressing these questions and drawing policy lessons clearly requires an empirical analysis based on micro-level data (Raj and Seamans, 2018).

Our paper is the first attempt in the literature to investigate differences in robot adoption across firms, and analyze the implications of these differences for the labor market effects of robots. To do so, we draw upon a unique panel data-set of Spanish manufacturing firms from the Encuesta Sobre Estrategias Empresariales (ESEE) over a 27-year period (1990-2016).² In contrast to existing studies, our paper uses explicit information on robot use in the production process of individual firms. Figure 1 constructed from the ESEE data-set provides a first indication that firm heterogeneity in the adoption of robots matters greatly for the labor market effects of robot technology. It demonstrates that those firms that adopted robots between 1990 and 1998 (“robot adopters”) *increased* the number of jobs by more than 50% between 1998 and 2016, while those firms that did not adopt robots (“non-adopters”) *reduced* the number of jobs by more than 20% over the same period.^{3,4} From macro-level information on robot use, as employed in the existing literature, it is

¹Industrial robots differ from other technologies or capital equipment in that robots are automatically controlled and capable of doing different tasks (see UNCTAD, 2017, Ch.III p.38). In a broad sense, industrial robots are defined as “automatically controlled, reprogrammable, multipurpose manipulators, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications” (ISO 8373, for details see <https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en> accessed on Feb 23, 2019.)


²The Spanish economy is an interesting case to look into, since it is one of the industrialized nations with the highest robot density within Europe (see UNCTAD, 2017; Jäger et al., 2015).

³To construct the figure, we balance the sample across the entire sample period from 1990 to 2016 and thus abstract from entry into, and exit from, the sample. Moreover, we only keep those firms in the sample that did not use robots in 1990, and had either started to use robots by 1998, or never used robots throughout the sample period. The thus constructed sample consists of almost 100 firms with 675 and 1701 firm-year observations for the group of robot adopters and the group of non-adopters, respectively. In an alternative approach, we keep all firms in the sample except for those that already use robots in the first year they appear in our sample. This gives us 644 robot adopters defined as firms that, at some point in time, start using robots, and 3802 non-adopters defined as firms that never use robots during our sample period. Using this approach, we find that, consistent with the evidence from Figure 1, robot adopters exhibit dramatically higher annual employment growth, on average, than non-adopters.

⁴At the same time, our data reveal that robot adopters were able to reduce their labor cost shares relative to non-adopters; see Appendix A.1.

impossible to identify and investigate this striking pattern in the data.

Figure 1: Evolution of firm-level employment (1990-2016)



Notes: The figure depicts the evolution of average firm employment (measured by the number of workers) in a balanced sample of firms from 1990-2016, separately for robot adopters (solid black line) and non-adopters (dashed grey line). Robot adopters are defined as firms that entered the sample in 1990 and had adopted robots by 1998. Non-adopters are firms that never use robots over the whole sample period.

Source: Authors’ computations based on ESEE data.

To provide a suitable lens through which to interpret our data, we begin our analysis by developing a theoretical framework of firm-level robot adoption. Following Acemoglu and Restrepo (2018a), we combine a monopolistic competition framework with a task-based approach in which robots and labor are perfect substitutes for one another in a specific set of low complexity tasks (“automatable tasks”).⁵ To study across-firm differences in the incentives to adopt robots, we augment the model to allow for firm heterogeneity in terms of productivity, as in Melitz (2003). In its most basic form, our model generates two connected insights that are, in our view, instrumental for understanding the labor market effects of robots. First, robot adoption is characterized by positive selection. This means that firms with higher productivities are more likely to adopt robots. Secondly, since robots are productivity-enhancing, they raise firm-level output and market shares of robot adopting firms, and magnify productivity differences between adopters and non-adopters.

⁵There is a striking similarity between modeling automation and offshoring. In the offshoring literature, foreign labor is assumed to be a perfect substitute for domestic labor in offshorable tasks (e.g. Grossman and Rossi-Hansberg, 2008; Egger et al., 2015). Offshoring thus “parallels [the] analysis of machines replacing tasks” (see Acemoglu and Autor, 2011, p.69).

While this opens up the possibility for net job creation in high-productivity robot adopting firms, it also implies that the least productive non-adopters are forced to exit the market, and that surviving non-adopters lose market shares and reduce employment. These insights suggest the existence of two sources of aggregate productivity gains due to robot technology: (1) direct efficiency gains in those firms that adopt robots; and (2) indirect gains through labor reallocation that benefits those workers employed in robot adopting firms, while hurting those in non-adopting firms.⁶

In our empirical analysis, we demonstrate the importance of this mechanism and quantify the associated effects. We first focus on the adoption decision, and reveal strong evidence for positive selection, i.e., we show that those firms that adopt robots in their production process are larger and more productive than non-adopters already before adopting robots. We also establish robust evidence that, conditional on productivity, more skill-intensive firms are *less* likely to adopt robots. This finding is consistent with a version of our model featuring two skill types of labor as well as firm heterogeneity in the complexity of the production process. Intuitively, a more complex production process requires a larger share of high-skilled workers; since these workers are more difficult to replace, there is a negative relationship between the skill intensity of the firm and the gains from automation (see also Autor et al., 2003). Finally, our data show that exporters are *more* likely to adopt robots than non-exporters, and we provide some evidence that this might reflect internal scale economies that can be harvested by serving foreign markets in addition to the domestic market (Bustos, 2011).

We then proceed by investigating the *effects* of robot adoption at the firm level. Since the adoption decision is not random, but instead governed by, among other things, the firm's productivity and skill intensity, this analysis faces a fundamental endogeneity problem. To tackle this problem and credibly control for non-random selection into robot adoption, we closely follow the methodology proposed by Guadalupe et al. (2012) and combine a difference-in-differences approach with a suitable propensity score reweighting estimator. This allows us to establish the following robust results. First, we find positive and significant output effects of robot adoption. Our estimates imply that the adoption of robots in the production process raises output by almost 25% within four years. Secondly, we find that robots raise firm-level employment by around 10 percent. While positive employment effects turn out to be especially pronounced for high-skilled workers, they also apply to other types of workers, namely low-skilled workers as well as workers employed in the firm's manufacturing establishments. Finally, we estimate a significant decline in the labor cost share by almost 7 percentage points following robot adoption. These results are consistent with our theoretical framework, where robot adopters reduce their labor cost shares, while the impact on employment is ambiguous and depends on the relative strength of the displacement effect and the productivity effect of robot adoption.⁷

Our data are rich enough to also investigate the impact of alternative technologies that firms

⁶There is a resemblance of this mechanism to the model in Melitz (2003) which studies the impact of trade liberalization on aggregate productivity in the presence of firm heterogeneity.

⁷The offshoring literature similarly emphasizes the importance of a productivity effect and a displacement effect on labor demand; see Grossman and Rossi-Hansberg (2008) and Antràs et al. (2006); Egger et al. (2015) in settings with heterogeneous firms.

adopt in their production process, and contrast it with the impact of robot technology. We focus on computer-digital machine tools, computer-assisted design, as well as a combination of some of the technologies through a central computer. While we find that these alternative technologies raise firm output and employment, it turns out that the effects are smaller in magnitude than what we estimate for robots. Moreover, when including all technologies simultaneously in the estimation, we find that only the effects of robot adoption are fully robust, while the other effects, by and large, disappear. Another striking difference between robot technology and other technologies we identify is that only robots lead to a significant reduction in the firm's labor cost share.

In a final step of our empirical analysis, we investigate how non-adopting firms, i.e., firms that do not start using robots, are affected by the rise of robot technology. We also shed light on the contribution of robots to gains in aggregate productivity. We reveal significant job losses in non-adopting firms. Our estimates imply that 10% of jobs in non-adopting firms are destroyed when the share of sales attributable to robot-using firms in their industries increases from zero to one-half. The same logic applies to changes in output, but the implied magnitude is even more pronounced. Looking at survival probabilities, we document significantly higher exit rates among non-adopters due to an increase in the industry's robot density, which is consistent with the predicted increase in the survival cut-off productivity in our theoretical framework. Importantly, our results are robust to using different measures of robot density, including the industry-specific stock of robots from the International Federation of Robotics (IFR). To shed light on the contribution of robots to gains in aggregate productivity, we distinguish between (1) direct efficiency gains within robot adopters and (2) indirect gains through reallocation of labor from non-adopters to adopters. We find that without the availability and adoption of robot technology, aggregate productivity would have doubled over the period from 1990 to 2016 rather than tripled. Moreover, we see that direct technical efficiency gains explain about two thirds of the total productivity gains attributable to robots, while the gains due to labor reallocation explain the remaining one third.

Our paper contributes to a recent literature that investigates the labor market implications of robot technology. The influential paper by Frey and Osborne (2017) was one of the first to examine how susceptible jobs are to computerization. They argue that almost 47% of total U.S. employment can be automated in the nearest future. In their paper, computerization is defined as a job automation by means of computer-controlled equipment. Three recent contributions focus specifically on robot adoption by using variation across countries and industries employing data from the IFR. Focusing on the period from 1993 to 2007 and covering 17 different countries, Graetz and Michaels (2018) find that the growing intensity of robot use accounted for 15% of aggregate economy-wide productivity growth, contributed to significant growth in wages, and had virtually no aggregate employment effects. Acemoglu and Restrepo (2017) and Dauth et al. (2018) use a local labor market approach to estimate the effects of robots on employment, wages, and the composition of jobs. Focusing on the U.S. between 1990 and 2007, Acemoglu and Restrepo (2017) find that one more robot per thousand workers reduces the employment to population ratio by about 0.2 percentage points and wages by 0.37 percent within commuting zones. Looking at Germany

between 1994 and 2014, Dauth et al. (2018) find no effects on total employment, but identify a substantial shift in the composition of jobs, away from manufacturing jobs and towards business service jobs. Moreover, they show how the use of robots increases local labor productivity, but depresses the labor share in total income.

While these studies provide important and novel evidence on robot adoption, using statistics at the industry level precludes an in-depth analysis within and between firms. In our study, we document selection based on observable firm characteristics (size, productivity, skill intensity, and exporting) and reveal positive employment and output effects in those firms that start to use robots, while negative employment (and output) effects arise from lower market shares for non-adopting firms. Furthermore, we demonstrate that the productivity gains documented in Graetz and Michaels (2018) or Dauth et al. (2018) might be partly explained by a reallocation of workers from low-productivity non-adopting firms to high-productivity robot adopters. In other words, with selection of more productive firms into robot adoption, increased exposure to robots reduces market shares of non-adopters and forces the least productive firms to exit. This across-firm reallocation affects aggregate industry productivity and speaks to “enormous and persistent measured productivity differences across producers, even within narrowly defined industries” (Syverson, 2011, p.326). Taking stock, by using detailed firm-level panel data from Spain for an extensive period of time, our paper allows to fill an important gap in recent attempts to investigate how automation affects productivity and labor markets.⁸

The remainder of our paper is organized as follows. In Section 2 we describe the ESEE data-set and provide first descriptive evidence on the use of robots across firms, industries, and time. In Section 3, we provide a theoretical perspective on firm-level robot adoption that guides us in our subsequent empirical analysis. In Section 4, we analyze the robot adoption decision of firms, and in Section 5 we investigate the firm-level effects of robot adoption. In Section 6 we bring our results together by shedding light on the aggregate implications of robot adoption. Section 7 concludes.

2 Data


Our empirical analysis is based on data collected by the Encuesta Sobre Estrategias Empresariales (ESEE) and supplied by the SEPI foundation in Madrid. The ESEE is an annual survey covering around 1,900 Spanish manufacturing firms each year with rich and very detailed information about firms’ manufacturing processes, costs and prices, technological activities, employment, and so forth. For the purposes of our research, the key aspect that sets the ESEE data-set apart from other

⁸By documenting a positive link between robot adoption and globalization, our paper also speaks to a large literature on technology upgrading in the global economy. Bustos (2011) provides evidence that exporters intensify their investments in technology after a trade liberalization process. Lileeva and Trefer (2010) similarly document how improved foreign market access prompted plants in Canada to adopt more advanced technologies. Our study complements these findings by looking at a specific type of technology upgrading, viz. the adoption of robot technology in the production process. Even though we do not claim to establish a causal effect of exporting on robot adoption, we document that exporting is an important determinant of a firm’s decision to adopt robots. This finding speaks to the view that serving more consumers allows firms to scale up their production generating incentives to install cost-saving robots.

data-sets is that it contains firm-level information on the use of robots in production. Hence, it provides a unique opportunity for studying the incentives for, as well as the consequences of, robot adoption at the firm level. In the following, we provide details on the specific data we exploit in our analysis and we document novel facts, drawn from our data, about robot diffusion and robot adoption in Spanish manufacturing.

Our study exploits data across 27 years spanning the years from 1990 to 2016. This is the complete sample period currently available from the ESEE. It provides a unique opportunity to investigate the drivers and consequences of profound changes in robot diffusion over roughly the last three decades. The initial sampling of the data in 1990 had a two-tier structure, combining exhaustive sampling of firms with more than 200 employees and stratified sampling of firms with 10-200 employees. In the years after 1990, special efforts have been devoted to minimizing the incidences of panel exit as well as to including new firms through refreshment samples aimed at preserving a high degree of representativeness for the manufacturing sector at large.⁹ In total, our data-set represents an unbalanced sample of some 5,500 different firms. In the data, we can distinguish between 20 different industries at the 2-digit level of the NACE Rev. 2 classification and six different size groups defined by the average number of workers employed during the year (10-20; 21-50; 51-100; 101-200; 201-500; >500); combinations of industries and size groups serve as stratas in the stratification. We express all value variables in constant 2006 prices using firm-level price indices derived from the survey data or, where necessary, industry-level price indices derived from the Spanish Instituto Nacional de Estadística (INE).

Figure 2: Evolution of robot diffusion in Spain (1990-2014)



Notes: The left panel depicts the share of firms using robots in their production process. The right panel depicts the share of total employment in firms using robots. The solid black lines consider all firms in the sample, while the dashed grey lines consider, respectively, large firms (those with more than 200 employees) and small firms (those with up to 200 employees).

Most importantly for our analysis, we exploit information on whether a firm uses robots in the production process. The survey asks firms: “State whether the production process uses any of the

⁹For details see <https://www.fundacionsepi.es/investigacion/esee/en/spresentacion.asp> (accessed on Feb 19, 2019).

following systems: 1. Computer-digital machine tools; 2. Robotics; 3. Computer-assisted design; 4. Combination of some of the above systems through a central computer (CAM, flexible manufacturing systems, etc.); 5. Local Area Network (LAN) in manufacturing activity”.¹⁰ Based on this question, we construct a 0/1 robot indicator variable equal to one if the firm uses robots and zero otherwise. We also use information on the other systems and generate indicators for CAM, CAD, and FLEX, respectively (more on this below).¹¹ The robot information is available every four years, starting in 1990. In addition, firms report the use of robots in the year 1991, as well as in the first year they enter the sample.¹² Before describing other variables we use in our empirical analysis, we document some patterns of robot use across time and industries.

Figure 2 depicts the evolution of robot diffusion in the Spanish manufacturing sector over the period 1990-2014. The left panel shows that, among all firms, just about 8% were using robots in 1990. This share has grown considerably over time, to more than 20% in 2014. The figure also reveals very significant differences in robot use between small firms (those with up to 200 employees) and large firms (those with more than 200 employees). For example, in 1990 already one third of large firms had adopted robots, while the same number for small firms was just 6%. The difference between these shares has grown over time, such that in 2014 about 60% among large firms use robots vs. 20% among small firms. The right panel of the figure shows the evolution of employment shares corresponding to robot firms. In 2014, almost 50% of all workers were employed in firms using robots, while the same number was more than 70% (37%) when only considering employment in large (small) firms. Taking stock, robot firms represent a highly significant part of modern Spanish manufacturing, especially among large businesses.

Our data also reveal a high degree of heterogeneity in robot diffusion and robot adoption rates across industries. Figure 3 depicts the share of firms using robots for 20 different industries, separately for the years 1990 and 2014. In 1990, the top-3 robot-using industries were Ferrous & Non-Ferrous Metals (18%), Machinery & Electrical Equipment (18%), and Motorized Vehicles (16%). By 2014, this ranking had changed and the top-3 industries were then Motorized Vehicles (57%), Furniture (31%), and Plastic & Rubber Products (30%). Thus, we see huge cross-industry differences in the share of firms using robots at a given point in time, as well as in the adoption rates between 1990 and 2014. Robot adoption at the industry level occurs with varying pace and magnitude.


We now continue by describing in more detail our data-set and the variables we employ in our

¹⁰The original questionnaire is distributed in Spanish. The question in Spanish is: “Indique si el proceso productivo utiliza cada uno de los siguientes sistemas: 1. Máquinas herramientas de control numérico por ordenador; 2. Robótica; 3. Diseño asistido por ordenador (CAD); 4. Combinación de algunos de los sistemas anteriores mediante ordenador central (CAM; sistemas flexibles de fabricación, etc.); 5. Red de Área Local (LAN) en actividad de fabricación”. In 1990, the possible answers were slightly different: “1. CAD/CAM; 2. Robótica; 3. Sistemas flexibles de fabricación; 4. Máquinas herramientas de control numérico”.

¹¹CAM, CAD, and FLEX are 0/1 indicator variables equal to one if the firm uses, respectively, computer-digital machine tools (CAM), computer-assisted design (CAD), and a combination of systems through a central computer (FLEX). We do not use information on Local Area Network adoption since it is only available from 2002 onwards.

¹²This means that we have robot information available in 1990, 1991, 1994, 1998, 2002, 2006, 2010, and 2014 for all firms included in the sample in the respective years. Moreover, we have robot information available in the remaining years (i.e., 1992, 1993, 1995,...) for those firms that appear in the sample for the first time in the respective years.

Figure 3: Share of robot firms by industry (1990 vs. 2014)



Notes: The figure shows the share of robot firms by industry. Black bars show data for 1990 and gray bars for 2014.

empirical analysis. Since we are interested in the effects of robot adoption (i.e., firms switching from non-robot use to first-time robot use), we restrict our sample to firms that do not use robots in the first year they appear in our data.¹³ Moreover, we drop sample observations after a firm undergoes a major restructuring due to changes in corporate structure (e.g. following a merger with another firm). This allows us to eliminate from the analysis situations connected with huge employment or output changes that are unrelated to robot adoption. In total, we have 4,446 different firms in the thus restricted sample. 646 (15%) of these firms adopt robots at some point in time within our sample period (“robot adopters”) and 3,800 (85%) never adopt robots (“non-adopters”). Furthermore, 62% among robot adopters keep on using robots throughout, while 30% report the use of robots for a certain period of time and abandon them afterwards. Finally, less than 10% of robot adopters switch back and forth several times. For our purposes, it is unclear how to interpret these multiple switches and we therefore drop this last group of firms from our analysis.

In our empirical analysis, we employ a rich array of firm-specific variables. These include output, labor productivity, employment, average wage, labor cost share, capital intensity, R&D intensity, skill intensity, export status, import status, and firm’s ownership structure (foreign vs. domestic). Output is given by the market value of the firm’s total annual production.¹⁴ Labor productivity is defined as value added per worker. Employment measures are total employment, given by the average number of workers during the year, and manufacturing employment, given by the number of workers in the firm’s manufacturing as opposed to non-manufacturing establishments. The average wage is constructed as total labor costs (gross salaries and wages, compensations, social security contributions paid by the company) divided by total employment. The labor cost share is calculated

¹³Recall that firms always report whether or not they use robots in the year of sample entry.

¹⁴This variable as well as all other value variables are expressed in constant 2006 prices using firm-level price indices. Thus, changes in our output measure over time within a firm reflect changes in physical output rather than changes in prices (see Ornaghi, 2006).

as labor costs divided by output. Moreover, we use direct firm-level information on the workforce composition by education to compute measures of the firm’s skill intensity, defined as the share of workers with a five-year university degree. Capital intensity is the value of the firm’s capital stock divided by effective work-hours. R&D intensity is the ratio of total expenses in R&D over total sales. Exporter and importer status dummies are equal to one if the firm reports positive export or import values, respectively. Foreign ownership indicates whether the firm is in foreign or domestic ownership (applying a threshold for foreign-owned capital of 50%). All variables are available on a yearly basis, except for the information on workers’ education levels, which are available every four years.

Table 1 presents descriptive statistics on the various variables we employ in our empirical analysis. We pool the data across all years and then sort observations into groups of firms that adopt robots at some point in time and those that never use robots within our sample period. The table reveals some suggestive differences between the two types of firms. Robot adopters turn out to be superior firms in many dimensions. They produce more output, they are more productive, and they employ more workers, even when focusing on just workers in manufacturing jobs or just low-skilled workers. Moreover, while robot adopters pay a higher average wage, they have a lower average labor cost share than non-adopters. In addition, robot adopters are more “globalized”, in the sense that they are more likely to export, import, and be in foreign rather than domestic ownership. Of course, these differences may be caused by factors unrelated to the adoption of robots. In the empirical analysis that follows later on, we will try to sort out which of the differences between robot adopters and non-adopters already existed before firms started to adopt robots, and which are causally associated with robot adoption.

3 A theoretical perspective on firm-level robot adoption

This section provides a theoretical framework for our empirical analysis. It draws from recent attempts in the literature to formalize the implications of robot technology, and serves to reveal the main economic trade-offs that we can expect to be at play at the firm level. We use our theoretical framework to derive hypotheses about the decision of firms to adopt robots, as well as about the implications of robot adoption for output, labor costs, labor demand, and aggregate productivity.

3.1 Basic set-up

Consider an industry in which a large number of monopolistically competitive firms produce horizontally differentiated goods. A firm ω is selling its unique variety at price $p(\omega)$ to consumers, facing an iso-elastic demand $q(\omega)$ of the form

$$q(\omega) = Ap(\omega)^{-\frac{1}{1-\beta}}, \tag{1}$$

Table 1: Descriptive statistics

	Robot adopters (1)	Non-adopters (2)	Observations (1)/(2)
Output (in logs)	15.973 (2.482)	14.407 (3.049)	7,547/25,169
Labor productivity (in logs)	10.552 (0.650)	10.316 (0.673)	7,382/24,010
Total employment (in logs)	4.475 (1.368)	3.510 (1.187)	7,501/24,578
Manufacturing employment (in logs)	4.421 (1.344)	3.461 (1.163)	7,361/24,138
Share of manufacturing employment	0.961 (0.129)	0.964 (0.118)	7,367/24,151
# low-skilled workers (in logs)	4.410 (1.316)	3.508 (1.121)	2,371/9,266
# high-skilled workers (in logs)	1.570 (1.345)	0.857 (1.101)	2,371/9,266
Average wage (in logs)	10.136 (0.447)	9.967 (0.488)	7,420/24,187
Labor cost share	0.285 (0.214)	0.342 (0.476)	7,447/24,127
Capital intensity (in logs)	3.432 (0.987)	2.852 (1.147)	7,081/23,176
Skill intensity (in logs)	0.051 (0.069)	0.043 (0.069)	2,392/9,371
R&D intensity (in logs)	0.343 (0.618)	0.189 (0.495)	7,444/24,564
Exporter status	0.704 (0.456)	0.484 (0.500)	7,487/24,614
Importer status	0.692 (0.462)	0.473 (0.499)	7,470/24,358
Foreign owned	0.155 (0.362)	0.081 (0.272)	7,523/24,697

Notes: The table reports means and standard deviations (in parentheses) of firm-specific variables for robot adopters (i.e. firms that start using robots at some point in time; column (1)) vs. non-adopters (i.e. firms that never use robots; column (2)). The numbers of observations reported in the final column correspond to the firm-year observations in columns (1) and (2). The sample spans the years 1990-2016 and is restricted to firms that do not use robots in the first year they enter the sample. Output is a firm's total production value. Labor productivity is value added per worker. Total employment is the average number of workers during the year. Manufacturing employment is the workforce employed at manufacturing as opposed to non-manufacturing establishments. Share of manufacturing employment is the number of workers employed at manufacturing establishments divided by the total number of workers employed by the firm. High-skilled workers are defined as workers with a five-year university degree, while low-skilled workers are all other workers. Average wage is computed as labor costs divided by the total number of workers. Labor cost share is labor costs divided by the total production value. Capital intensity is the value of the firm's capital stock divided by effective work-hours. Skill intensity is the share of high-skilled workers. R&D intensity is the ratio of total expenses in R&D over total sales volume. We add one to all factor intensity variables as well as the number of high- and low-skilled workers before taking logs in order to keep zero observations. Exporter (importer) status is a dummy variable equal to one if the firm reports positive exports (imports). Foreign ownership indicates whether a firm is foreign owned by more than 50%.

where β controls the (constant) elasticity of substitution $1/(1 - \beta) > 1$ between any two varieties, and A is a demand shifter.¹⁵ As for the production side, we follow Acemoglu and Restrepo (2018b)

¹⁵As is well known, the demand function in Eq. (1) with $A \equiv EP^{1-\beta}$ and $P = \left[\int_{\omega \in \Omega} p(\omega)^{\frac{-\beta}{1-\beta}} d\omega \right]^{-\frac{1-\beta}{\beta}}$ is implied by a standard utility maximization problem where consumers have a CES utility function $U = \left[\int_{\omega \in \Omega} q(\omega)^\beta d\omega \right]^{\frac{1}{\beta}}$ and

in writing output as a composite of different tasks combined in a constant elasticity of substitution (CES) aggregate. However, we depart from Acemoglu and Restrepo (2018*b*) by introducing two types of firm heterogeneity into their framework. The first type is the standard Melitz (2003) heterogeneity meaning that firms differ in their exogenous (baseline) productivity denoted by $\phi(\omega)$. We index tasks by i and assume that they can be ordered according to their complexity where a higher index i reflects higher complexity. Specifically, output of firm ω is given by

$$x(\omega) = \phi(\omega) \left(\int_{N(\omega)-1}^{N(\omega)} x(\omega, i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \quad (2)$$

where σ denotes the elasticity of substitution between any two tasks and $x(\omega, i)$ is the output of task i in firm ω . The parameter $N(\omega)$ generates a second type of firm heterogeneity in the model. It is given exogenously and governs the set of tasks the firm has to perform, with the task range normalized to one and with the limits of integration given by $N(\omega) - 1$ and $N(\omega)$.¹⁶ An increase in $N(\omega)$ reflects quality upgrading, in the sense that new and more complex tasks appear and replace old tasks in the production process (the least complex ones). Crucially, we assume that the simpler tasks with index numbers $i \leq I$ can be performed by robots or human labor, while the more complex tasks with index numbers $i > I$ are bound to be performed by human labor. The parameter $I \in [N(\omega) - 1, N(\omega)]$ thus reflects the ability level of robots in performing complex tasks. This parameter is likely to vary across industries and through time as technology advances. Output at the task level is given by

$$x(\omega, i) = \mathbb{1}[i \leq I] \eta(i)k(\omega, i) + \gamma(i)l(\omega, i), \quad (3)$$

where $\mathbb{1}[i \leq I]$ is a 0/1 indicator equal to one if $i \leq I$ and zero otherwise, and $\gamma(i)$ and $\eta(i)$ denote, respectively, the productivity of labor l and robot capital k in task i . Crucially, robot capital and labor are perfect substitutes for one another in all tasks $i \leq I$. This view highlights an important aspect of automation, namely that machines are used to substitute for human labor (Acemoglu and Restrepo, 2018*a*, p.2).¹⁷

As in Acemoglu and Restrepo (2018*a*), we assume that the ratio of $\eta(i)/\gamma(i)$ is strictly decreasing in i , which formalizes a comparative advantage of labor in more complex tasks. Moreover, we assume that the effective robot capital costs (at rental rate r) are strictly below the effective labor costs (at wage rate w) for all automatable tasks. Formally, we have $r/\eta(I) < w/\gamma(I)$. These assumptions reflect the view that human labor is more valuable in performing complex tasks than robot capital.

face a budget constraint $E = \int_{\omega \in \Omega} p(\omega)q(\omega)d\omega$ with E being the total expenditure on the set of available varieties Ω .

¹⁶Capuano et al. (2017) provide evidence for substantial heterogeneity in the type of tasks performed by German plants even if they are operating in the same industry. Most of the results we derive in our theoretical analysis do not depend on differences in $N(\omega)$. However, allowing for this heterogeneity in a simple extension of our model will generate differences in the skill intensity of firms that are consistent with our data, as will become evident below.

¹⁷Notice the striking similarity to the offshoring literature, where foreign labor is often assumed to be a perfect substitute for domestic labor in offshorable tasks (e.g. Grossman and Rossi-Hansberg, 2008; Egger et al., 2015). This is also true for Groizard et al. (2014), who consider, as we do, the case of a CES production technology.

Accordingly, we can write the unit production costs of a firm using robots to perform all tasks $i \leq I$ as

$$c^a(\phi(\omega), N(\omega), I) = \frac{1}{\phi(\omega)} [\eta(N(\omega), I)r^{1-\sigma} + \gamma(N(\omega), I)w^{1-\sigma}]^{\frac{1}{1-\sigma}}, \quad (4)$$

where $\eta(N(\omega), I) \equiv \left(\int_{N(\omega)-1}^I \eta(i)^{\sigma-1} di\right)^{\frac{1}{\sigma}}$ and $\gamma(N(\omega), I) \equiv \left(\int_I^{N(\omega)} \gamma(i)^{\sigma-1} di\right)^{\frac{1}{\sigma}}$ summarize the productivity over all tasks performed by robots and labor, respectively.¹⁸ The superscript a indicates that the production process has been automated. However, this decision is endogenous and requires incurring a fixed cost, denoted by $F^a > 0$. Not paying the fixed cost means that the firm has to perform all tasks using human labor with corresponding unit cost of $c(\phi(\omega), N(\omega), N(\omega) - 1) = \frac{1}{\phi(\omega)} [\gamma(N(\omega), N(\omega) - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}}$.

Using constant mark-up pricing, we can write a firm's profit gain from robot adoption, defined as $\Delta\pi(\omega) \equiv \pi^a(\omega) - \pi(\omega)$, as follows:¹⁹

$$\Delta\pi(\omega) = (1 - \beta)A \left\{ \frac{1}{\beta} \frac{1}{\phi(\omega)} [\gamma(N(\omega), N(\omega) - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} [\kappa(N(\omega), I) - 1] - F^a, \quad (5)$$

where $\kappa(\cdot)$ is defined as

$$\kappa(N(\omega), I) \equiv \left(\frac{\left(\int_{N(\omega)-1}^I \eta(i)^{\sigma-1} di\right)^{1/\sigma} r^{1-\sigma} + \left(\int_I^{N(\omega)} \gamma(i)^{\sigma-1} di\right)^{1/\sigma} w^{1-\sigma}}{\left(\int_{N(\omega)-1}^I \gamma(i)^{\sigma-1} di\right)^{1/\sigma} w^{1-\sigma} + \left(\int_I^{N(\omega)} \gamma(i)^{\sigma-1} di\right)^{1/\sigma} w^{1-\sigma}} \right)^{\frac{1}{\sigma-1} \frac{\beta}{1-\beta}}. \quad (6)$$

This expression is (weakly) larger than one and reflects the marginal cost savings from robot adoption. Given that labor has a comparative advantage in performing more complex tasks and the fact that $r/\eta(I) < w/\gamma(I)$, we find that $\kappa(\cdot)$ is, ceteris paribus, increasing in the level of robot technology I and decreasing in the complexity of tasks $N(\omega)$. If firms face a highly complex production process such that $I = N(\omega) - 1$, then all tasks must be performed by labor and there are consequently no cost savings from robot adoption, $\kappa(N(\omega), N(\omega) - 1) = 1$.

3.2 The robot adoption decision

We are now ready to investigate explicitly the decision of firms to adopt robots, and derive hypotheses on the selection of firms into robot adoption that we can confront with our Spanish firm-level

¹⁸As shown in Acemoglu and Restrepo (2018b), we can write output as $x(\omega) = [\eta(N(\omega), I)K(\omega)^\rho + \gamma(N(\omega), I)L(\omega)^\rho]^{\frac{1}{\rho}}$, with $\rho = (\sigma - 1)/\sigma$, where $\int_{N(\omega)-1}^I k(\omega, i) di = K(\omega)$ and $\int_I^{N(\omega)} l(\omega, i) di = L(\omega)$. The corresponding unit-cost function is thus given by (4).

¹⁹Note that profits of firm ω using robots can be written as $\pi^a(\omega) = (1 - \beta)A \left[\frac{1}{\beta} c^a(\phi(\omega), N(\omega), I) \right]^{-\frac{\beta}{1-\beta}} - F^a - F$ while profits for the same firm using human labor instead of robots are $\pi(\omega) = (1 - \beta)A \left[\frac{1}{\beta} c(\phi(\omega), N(\omega), N(\omega) - 1) \right]^{-\frac{\beta}{1-\beta}} - F$, where F denotes overall fixed costs of production. Computing the difference between the two gives Eq. (5).

data. In the interest of space and readability, we confine ourselves to an intuitive discussion in the main text. Where necessary, we support this discussion with detailed analytical derivations corresponding to this section as well as the following section in Appendix A.2.

3.2.1 Productivity

First note that for $\kappa(\cdot) > 1$ the profit gain from robot adoption in Eq. (5) is increasing in a firm's baseline productivity $\phi(\omega)$. As is standard in the literature, we can define a cut-off productivity level $\phi^r(\omega)$ at which a firm is exactly indifferent between adopting and not adopting robots (for a given level of complexity $N(\omega)$). This cut-off productivity level is implicitly defined through $\Delta\pi(\omega)|_{\phi(\omega)=\phi^r(\omega)} = 0$. Hence, comparing two firms with an equally complex production process, ex-ante more productive firms are more likely to adopt robots in production.

3.2.2 Exporting

Suppose firms can choose to serve consumers not only in the domestic but also in the foreign economy. While the foreign economy is fully symmetric to the domestic economy, exporting requires the payment of a fixed export cost and per-unit iceberg type transport costs, denoted by F^x and τ , respectively. As is well-known, the introduction of a fixed export cost generates (sharp) selection of ex-ante more productive firms into exporting. Due to symmetry of the two countries, operating profits of exporting firms are now scaled by a constant factor $1 + \tau^{-\beta/(1-\beta)}$. This is similar to Bustos (2011), and we can conclude that exporters have stronger incentives to adopt robots as the gains from doing so—the reduction in variable production costs—can be scaled up to a larger customer base in home *and* foreign.

3.2.3 Skill intensity

Suppose there are two types of human labor, low-skilled and high-skilled workers, referenced by subscripts l and h , respectively. Following Acemoglu and Autor (2011), we assume that high-skilled workers have a comparative advantage over their low-skilled coworkers in the performance of more complex tasks. Specifically, we assume that the relative efficiency of high- to low-skilled labor, $\gamma_h(i)/\gamma_l(i)$, is strictly increasing in i . In such an environment, firms will not only compare the production costs of robots and human labor across tasks, but also consider the skill-specific effective labor costs in each task, i.e., the firm will benchmark $w_l/\gamma_l(i)$ against $w_h/\gamma_h(i)$. Given that high-skilled workers have a relative advantage in performing more complex tasks, this results in a cut-off task at which firms are exactly indifferent between hiring high-skilled and low-skilled workers for the performance of that task. Comparing two otherwise identical firms that differ only in the complexity of their production process, we find that the firm with higher $N(\omega)$ employs a higher share of high-skilled workers. Since firms with higher $N(\omega)$ are less likely to adopt robots, as discussed above, we have established a negative link between the skill intensity of firms and their

propensity to adopt robots.²⁰

3.3 The effects of robot adoption

Having discussed the decision of firms to adopt robots, we now focus attention to the *effects* of robot adoption, both at the firm and the industry level.

3.3.1 Firm-level effects

First of all, since robots have a comparative advantage in the production of automatable tasks, it is straightforward that robot adoption raises firm output. Moreover, due to our assumptions on the task production function in Eq. (3), it follows immediately that robot adoption reduces the labor cost share, as robots substitute human labor in automated tasks. The overall impact of automation on labor demand within firms is, however, ambiguous. It depends on two opposing effects: on the one hand, the *displacement effect* reduces demand for labor since part of the workforce is substituted by robots. On the other hand, the *productivity effect* entails that robots raise the efficiency in production, and thus output and employment. Similar to the offshoring literature (see Grossman and Rossi-Hansberg, 2008), the productivity gains may be strong enough to outweigh the losses, so that total firm-level employment increases. Clearly, the strength of the displacement effect depends on the share of automatable tasks, and thus the parameter I along with $N(\omega)$, while the magnitude of the productivity effect depends on the variable cost savings from robot adoption, determined by the efficiency parameters for robots and workers, $\eta(i)$ and $\gamma(i)$, respectively, as well as factor prices. A final question is which skills (and thus workers) are specifically affected by automation. Using the model with two skill types of labor from above, it is clear that low-skilled workers are more likely to be affected by automation, since they perform the less complex tasks which are the ones being automated. However, as long as the low-skilled workers are not fully replaced by robots, the productivity effect is also working in their favor. This is the case as long as the level of robot technology, I , is below the cut-off task at which firms are indifferent between employing high- and low-skilled labor.

3.3.2 Industry-level effects

We can use the model to study the industry-level effects of changes in the fixed cost of adopting robots, F^a , or changes in the level of robot technology, I . A decrease in F^a or an increase in I , as we can expect to occur over time, both decrease the cut-off productivity that separates robot adopters from non-adopters (for given N) and thus raise industry-level robot exposure. Similar to Melitz (2003), this has important implications for the industry equilibrium. As ex-ante more productive firms gain market shares by reducing marginal costs due to robot adoption, it raises the cut-off productivity at which firms are able to survive in the market. Put differently, increasing

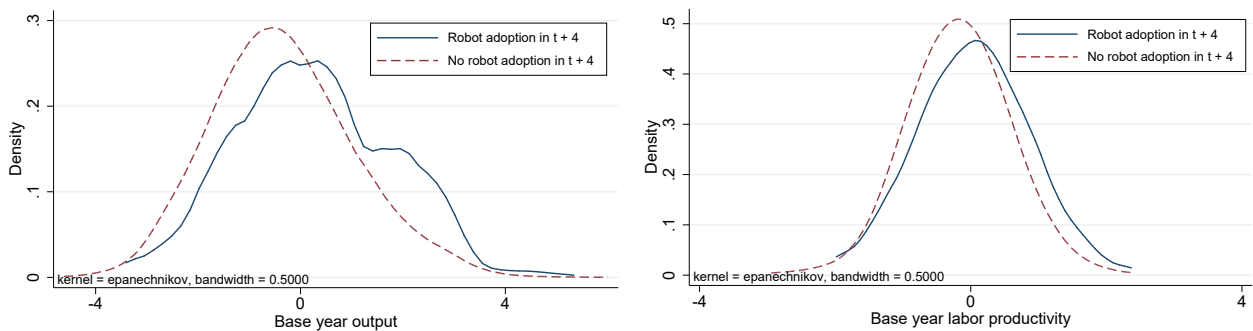
²⁰The motivation for this extension is given by the fact that we do not observe tasks, and thus task complexity in the ESEE data-set. However, following the extension, we can proxy task-complexity by the skill composition of firms in the subsequent empirical analysis.

robot exposure at the industry level prompts the least productive firms to exit, and the surviving non-robot firms to reduce their output and employment. This mechanism, along with the direct firm-level efficiency gains due to the use of robots, raises the industry’s aggregate productivity.

4 Which firms adopt robots?

We now turn to our empirical analysis and begin by investigating which firm-specific characteristics influence the decision to adopt robots. Our theoretical framework generates several predictions that we now bring to our Spanish firm-level data. The most important prediction concerns the relationship between the likelihood of robot adoption and the productivity and size of the firm. The prediction is consistent with arguments in the literature that more efficient firms benefit more from the adoption of higher levels of technology so that we should expect to find more productive firms to be more likely to adopt robots in their production (positive selection). Identifying whether positive selection is indeed at work in the data can help in understanding the large and persistent productivity differences across firms within industries (Syverson, 2011). In fact, if we find evidence for negative selection in the data, then this would point towards an alternative scenario with a potential catching-up of low-productivity firms through the use of robot technology.²¹

Figure 4: Distribution of base year output and productivity for robot adopters vs. non-adopters



Notes: In the left panel, the dashed red line shows the empirical probability density function of base year output of firms that do not use robots when they first appear in the sample at time t and will not have adopted robots four years later, i.e. at time $t + 4$. The solid blue line shows the same function of base year output of firms that do not use robots when they first appear in the sample at time t but will have adopted robots four years later (i.e. at time $t + 4$). The base year output is given in logs, deflated, and demeaned by industry. The right panel shows the same as the left panel but for labor productivity instead of output. The base year labor productivity is given by the log of (deflated) value added per worker demeaned by industry.

Before analyzing robot adoption more formally, we use our data to provide graphical evidence on the relationship between firm size/productivity and robot adoption. The left panel of Figure 4 plots the distribution of base year output (deflated and in logs) for robot adopters vs. non-adopters, i.e., for firms that have adopted robots four years after they first appear in the sample vs. firms that have not adopted robots. The figure reveals that the distribution of robot adopters (solid blue

²¹One argument implying negative selection is that more efficient firms are larger and thus require a more complex degree of bureaucracy that can hamper decision making about new technology and skills.

line) clearly dominates the distribution of non-adopters (dashed red line). Since we compute our measure of output relative to the industry mean, differences in firm size across industries are not driving this observation. Moreover, firms using robots already in the base year are not included in the figure, so the differences that we see are not explained by the *effects* of adopting robots. Importantly, we get a similar picture when using base year labor productivity instead of output, i.e., the productivity distribution of robot adopters clearly dominates the one of non-adopters; see the right panel of Figure 4.

We now proceed by investigating the adoption decision through the use of regression analysis. Specifically, we adopt the following basic empirical framework to describe the decision of firms to adopt robots:

$$\text{Robots}_{it} = \beta\phi_{it-1} + \gamma\mathbf{F}_{it-1} + \delta\mathbf{G}_{it-1} + \mu_{st} + \varepsilon_{it}, \quad (7)$$

where the dependent variable is a 0/1 indicator variable for robot use in the production process of firm i at time t , and where we focus on different sets of explanatory variables: (1) a firm-specific productivity variable ϕ_{it-1} ; (2) a vector of factor intensity variables \mathbf{F}_{it-1} ; and (3) a vector of globalization variables \mathbf{G}_{it-1} (with corresponding parameters to be estimated collected in β , γ , and δ , respectively). We also include industry-year fixed effects given by μ_{st} . These account for the increase in the supply and quality of robots, as well as the evolution of wages and adoption costs that can change the incentives to adopt robots over time. Finally, ε_{it} is the error term. The firm's productivity is measured as the log of labor productivity given by the firm's value added per worker (deflated). The factor intensity variables we use are the firm's capital intensity, skill intensity, and R&D intensity (all in logs). The globalization variables we use are 0/1 indicator variables for whether the firm is an exporter, an importer, and a foreign-owned firm, respectively.

In addition to applying this empirical framework to our panel data-set, we consider a simplified version of Eq. (7) and collapse the data into a single cross-section measuring all explanatory variables in the base year (i.e., in the first year the firm appears in the sample). The dependent variable is equal to one if the firm adopts robots at some point in time during our sample period, and zero otherwise. Note that firms using robots already in the year of sample entry are excluded from the analysis. In Panel A of Table 2 we present OLS estimates of the simplified (cross-sectional) version of Eq. (7).²² Standard errors are robust to arbitrary forms of heteroskedasticity. In column (1a) we use the most parsimonious specification including productivity as the only explanatory variable alongside industry-year fixed effects. Our estimates provide evidence that the more productive firms are significantly more likely to adopt robots. This is in line with our previous observation that the output and productivity distributions of robot adopters dominate those of non-adopters already before first-time adoption. The estimated coefficient is equal to +0.038 in the cross-section and implies that an increase by one standard deviation in the firm's base year labor productivity raises its probability of subsequently adopting robots by 2.5 percentage points.

²²In the interest of space we restrict our attention to OLS estimates in the text. In the appendix, we report estimates obtained with the non-linear Probit model. The results we obtain with this alternative estimator are very similar to the OLS estimates; see Table A.1 in Appendix A.3.

Table 2: Robot adoption I: Productivity-based selection

	Robot adoption (0/1 indicator)			
<i>PANEL A: Cross-sectional specification</i>	(1a)	(2a)	(3a)	(4a)
Base year labor productivity	0.0384*** (0.00907)	0.0147 (0.0107)	0.0150 (0.00950)	0.00285 (0.0110)
Base year skill intensity		-0.330*** (0.0973)		-0.398*** (0.100)
Base year exporter status			0.0716*** (0.0144)	0.0580*** (0.0161)
Observations	4053	3488	3986	3443
R-squared	0.110	0.151	0.130	0.161
		Robot adoption (0/1 indicator)		
<i>PANEL B: Panel specification</i>	(1b)	(2b)	(3b)	(4b)
Lagged labor productivity	0.0392*** (0.00546)	0.0173*** (0.00576)	0.0209*** (0.00550)	0.00856 (0.00578)
Lagged skill intensity		-0.103* (0.0617)		-0.148** (0.0631)
Lagged exporter status			0.0345*** (0.00795)	0.0253*** (0.00819)
Observations	7368	6934	7300	6879
R-squared	0.039	0.053	0.052	0.059
Industry(-base)-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	No	Yes	Yes

Notes: In Panel A the dependent variable in all columns is a 0/1 indicator variable equal to one if the firm adopts robots during our sample period and zero otherwise, while in Panel B it is equal to one if the firm uses robots in a specific year and zero otherwise. Labor productivity is the firm's deflated value added per worker (in logs). Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. All estimates in Panel A (B) include industry-base-year (industry-year) fixed effects. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. In Panel A all explanatory variables are measured in the base year defined as the first year the firm appears in the sample. In Panel B all explanatory variables are lagged by one year. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Moreover, in Panel B we condition on the firm not using robots in the previous year (or the most recent year for which robot data are available for that firm). Robust standard errors are given in parentheses and clustered by firm in Panel B. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

In columns (2a) and (3a) we augment the specification to include factor intensity and globalization variables, respectively. In order to focus on the central predictors suggested by our theoretical framework, we report coefficient estimates for the firm's skill intensity and its export status (along with productivity), while including capital intensity and R&D intensity in column (2a) and import status and a foreign ownership dummy in column (3a) as control variables. Doing so renders the effect of labor productivity insignificant, which reflects significant pairwise correlations among the different variables. Skill intensity enters negatively and significantly. This finding is consistent with the idea that higher skill requirements in the production process reduce the scope for economic benefits through robotization. The coefficient of the firm's export status is positive and significant. These results are upheld in column (4a) where we include all variables simultaneously. The estimated coefficients in this specification imply that exporting makes firms 6 percentage points more likely to adopt robots later on (controlling for productivity, factor intensities, and other globalization variables). These results provide compelling evidence for a fundamental complementarity

between exporting and robot adoption at the firm level. Those firms active on international markets through exporting are considerably more likely to adopt advanced technology in the form of robots.

By extending the analysis to estimate the panel version of Eq. (7), we find a picture that largely resembles our cross-sectional estimates; see Panel B of Table 2. As before, the sample we use includes only firms that do not use robots in the first year they appear in our data-set. We restrict attention to observations of those firms that never adopt robots as well as those that adopt robots for the first time during the sample period. Since we are interested in first time robot adoption, once a firm decides to use robots in the production process, we exclude subsequent observations from the estimation sample. Robust standard errors are clustered by firm. As in the cross-sectional specification, we find that more productive firms are more likely to become first time robot adopters. A firm's skill intensity contributes negatively, while the firm's export status enters the equation positively and significantly.

Table 3: Robot adoption II: Output-based selection

	Robot adoption (0/1 indicator)			
<i>PANEL A: Cross-sectional specification</i>	(1a)	(2a)	(3a)	(4a)
Base year output	0.0483*** (0.00405)	0.0408*** (0.00511)	0.0405*** (0.00504)	0.0346*** (0.00599)
Base year skill intensity		-0.408*** (0.0996)		-0.419*** (0.101)
Base year exporter status			0.0374*** (0.0143)	0.0349** (0.0161)
Observations	4221	3599	4149	3551
R-squared	0.139	0.165	0.142	0.167
	Robot adoption (0/1 indicator)			
<i>PANEL B: Panel specification</i>	(1b)	(2b)	(3b)	(4b)
Lagged output	0.0362*** (0.00229)	0.0335*** (0.00277)	0.0360*** (0.00296)	0.0328*** (0.00331)
Lagged skill intensity		-0.195*** (0.0635)		-0.190*** (0.0638)
Lagged exporter status			0.00363 (0.00800)	0.00150 (0.00833)
Observations	7535	7057	7461	6997
R-squared	0.071	0.074	0.071	0.073
Industry(-base)-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	No	Yes	Yes

Notes: In Panel A the dependent variable in all columns is a 0/1 indicator variable equal to one if the firm adopts robots during our sample period and zero otherwise, while in Panel B it is equal to one if the firm uses robots in a specific year and zero otherwise. Output is the firm's deflated output value (in logs). Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. All estimates in Panel A (B) include industry-base-year (industry-year) fixed effects. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. In Panel A all explanatory variables are measured in the base year defined as the first year the firm appears in the sample. In Panel B all explanatory variables are lagged by one year. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Moreover, in Panel B we condition on the firm not using robots in the previous year (or the most recent year for which robot data are available for that firm). Robust standard errors are given in parentheses and clustered by firm in Panel B. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

To further understand what motivates firms to install robots in the production process, we also use firm output instead of labor productivity in estimates of Eq. (7). The results we report in Table 3 provide evidence that output is a strong and significant predictor of subsequent robot adoption. The estimated coefficient of +0.048 in the cross-sectional specification without controls (see column 1a) implies that an increase by one standard deviation in the firm’s base year output raises its probability of adopting robots later on by as much as 8 percentage points. The effect remains economically meaningful, standing at 6 percentage points, after controlling for the firm’s factor intensities and globalization variables. What is interesting is that the globalization variables are rendered insignificant (at least in the panel specification in columns 3b and 4b) when using firm output instead of labor productivity. Since exporting firms serve a larger market than non-exporting firms, this is evidence that the scale of operations is a critical channel through which globalization supports robot adoption.²³

Finally, in another set of estimates we allow for non-linearity and non-monotonicity in the effects of productivity and output on robot adoption. We do this by replacing the productivity/output variable with dummy variables for each productivity/output quartile; see Tables A.3 and A.4 in Appendix A.3. The results are striking and indicate that firms in the top quartile of the productivity/output distribution have the highest probability of adopting robots. For example, firms in the top quartile of the output distribution are 15 percentage points more likely than firms in the bottom quartile to subsequently adopt robots even after controlling for factor intensity and globalization variables; see column (4) in Table A.3.

5 Firm-level effects of robot adoption

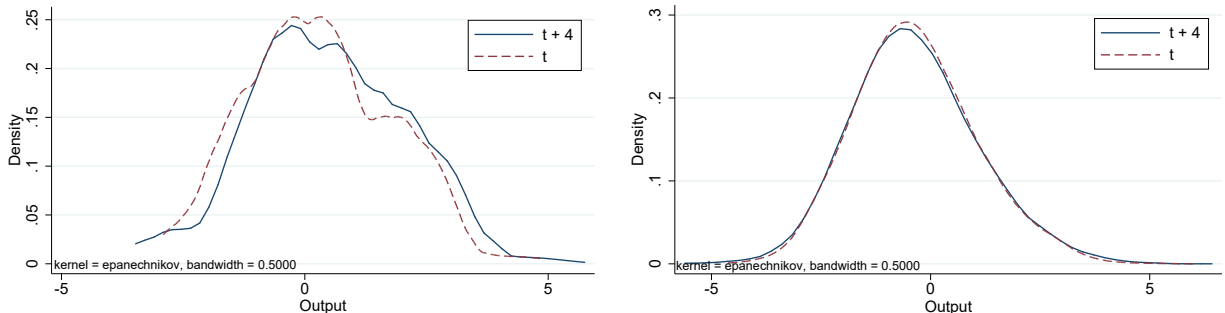
Heaving established novel facts on the selection of firms into robot adoption, we now aim to identify the consequences of robot adoption at the firm level. Our focus is on the effects on output, as well as on employment, labor costs, and average wages.

5.1 Output effects

We first present graphical evidence on the output distribution of robot adopters before and after the adoption, and benchmark it against changes in the output distribution of non-adopters. Figure 5 provides a first indication that, in contrast to non-adopters, robot adopters were able to significantly expand the scale of their operations. The left panel makes a before-after comparison among robot adopters. It reveals that the distribution of output (deflated and in logs) when firms enter the sample in t (dashed red line) is clearly dominated by the distribution of output four years later at $t + 4$ (solid blue line) when the same firms have adopted robots. The right panel makes the same comparison for firms that do not adopt robots and reveals almost no differences in the output distribution.

²³We report estimates obtained with the non-linear Probit model in Table A.2 in Appendix A.3.

Figure 5: Before-after comparison of the output distribution for robot adopters (left panel) vs. non-adopters (right panel)



Notes: The left panel makes a before-after comparison of the output distribution of robot adopters, i.e., firms that do not use robots when they first appear in the sample at time t , but will have adopted robots four years later (i.e. at time $t + 4$). The red dashed line and the solid blue line show the empirical probability density function of output at time t and at time $t + 4$, respectively. The right panel makes the same comparison for non-adopters, i.e., firms that do not use robots when they first appear in the sample at time t and will not have adopted robots four years later (i.e. at time $t + 4$). Output is given in logs, deflated, and demeaned by industry.

To identify the effect of robot adoption on firm-level output more formally, we estimate the following equation:

$$\text{Output}_{it} = \gamma_1 \text{Robots}_{it} + \gamma_2 \text{Robots}_{it-4} + \beta \mathbf{X}_{it-4} + \mu_i + \mu_{st} + \varepsilon_{it}, \quad (8)$$

where the dependent variable is deflated output of firm i in year t (in logs), \mathbf{X}_{it-4} is a vector of time-varying firm-level controls lagged by four years, with a corresponding vector of parameters β to be estimated, μ_i and μ_{st} are firm and industry-year fixed effects, respectively, and ε_{it} is an error term with zero conditional mean. The parameter μ_{st} captures general time trends and industry shocks affecting firms equally within industries. The parameters of interest in (8) are γ_1 and γ_2 , both capturing the impact of robot adoption on firm-level output. These parameters indicate the percentage change in output after firms start using robots in their production process.

By including fixed effects for individual firms, we identify the output effects of robot adoption only through within-firm variation, i.e., firms switching from non-robot use to robot use over time. The firm fixed effects control for robot adoption based on time-invariant factors, like the firm's baseline productivity $\phi(\omega)$ in our theoretical framework. To control for robot adoption based on not just time-invariant but also time-varying firm-level variables, we include labor productivity, capital intensity, skill intensity, R&D intensity (all in logs), as well as indicator variables for exporting, importing, and foreign ownership in \mathbf{X}_{it-4} .²⁴ We also estimate specifications including the four-year forward of our robot indicator variable (Robots_{it+4}). This allows us to see whether our model is reasonably successful at controlling for positive selection into robot adoption as identified in the

²⁴We let the firm-level control variables enter with a four-year lag in order to control for selection into robot adoption in $t - 4$ and t . However, we have also used a one-year lag instead of a four-year lag, to find that this does not alter our estimates in any significant way.

previous section.

To make further progress in establishing a causal effect of robot adoption on output, we closely follow the empirical methodology proposed by Guadalupe et al. (2012) and combine the firm fixed effects approach with a propensity score reweighting estimator in the spirit of DiNardo et al. (1996). Specifically, we construct propensity scores and reweigh each firm in order to generate a similar distribution of key observable characteristics across robot adopters and non-adopters. By matching along observable firm characteristics, we hope to also match the distribution of important unobservable characteristics. To estimate the propensity scores, we consider the years 1991, 1994, 1998, ..., 2014 in our panel and sort those firms that adopt robots in that year into the treatment group and those that never use robots into the control group. We then pool observations in the treatment and in the control group across all these years and obtain the propensity scores for all firms by running industry-specific probit regressions for robot adoption (the treatment) on one-year lags of sales, sales growth, labor productivity, labor productivity growth, capital-, skill- and R&D-intensity, indicators for exporter, importer and foreign ownership, and year dummies. The growth rates of both labor productivity and sales control for recent performance differences among firms. We then use the estimated propensity scores and reweigh each treated firm by $1/\hat{p}$ and each control firm by $1/(1 - \hat{p})$, where \hat{p} is the estimated propensity score.²⁵

Table 4: Output effects of robot adoption

	Output (in logs)					
	(1)	(2)	(3)	(4)	(5)	(6)
Robots _t	0.157*** (0.0289)	0.106*** (0.0344)	0.162*** (0.0315)	0.120*** (0.0370)	0.126*** (0.0385)	0.119** (0.0495)
Robots _{t-4}	0.121*** (0.0325)	0.126*** (0.0446)	0.119*** (0.0337)	0.111** (0.0468)	0.121*** (0.0415)	0.0815 (0.0545)
Robots _{t+4}		0.0743** (0.0348)		0.0471 (0.0383)		0.0724 (0.0478)
Observations	4977	2813	4570	2574	4633	2634
R-squared	0.240	0.295	0.249	0.294	0.264	0.284
Selection controls	No	No	Yes	Yes	No	No
Propensity scores	No	No	No	No	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Robots is a 0/1 indicator variable equal to one if the firm uses robots in the specified period. The dependent variable in all columns is the log of the firm's deflated output value. Selection controls (in $t - 4$) are the firm's deflated labor productivity (in logs), deflated capital intensity (in logs), skill intensity (in logs), deflated R&D intensity (in logs), as well as exporter, importer, and foreign ownership dummies. We add one to all factor intensity variables before taking logs in order to keep zero observations. For details on the propensity score reweighting estimator see the text. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Robust standard errors clustered by firm are given in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

Table 4 shows our estimates of Eq. (8). We first estimate the equation with firm fixed effects, but without selection controls (columns (1) and (2)); we then add time-varying firm-specific variables

²⁵We only keep those observations in the analysis that are in the region of common support, and we have checked that the balancing property is supported by the data in all industries, i.e., all observed characteristics of robot adopters and non-adopters are balanced. More output corresponding to the propensity score estimation can be found in Table A.5 in Appendix A.4.

as selection controls (columns (3) and (4)); and we finally use the propensity score reweighting estimator as described above (columns (5) and (6)). We estimate each of the three variants with and without the four-year forward of the robot indicator variable. Throughout all specifications employed, we find positive and significant output effects of robot adoption. We also see that, once we include selection control variables or use the propensity score reweighting estimator to control for positive selection into robot adoption, the four-year forward of the robot indicator variable is not significantly different from zero. This makes us confident that, for our purposes, we are modelling the selection decision reasonably well. To get a sense of the magnitude of the effects, consider the estimates of γ_1 and γ_2 in column (5), which are equal to +0.126 and +0.121, respectively. These estimates imply that the adoption of robots in the production process raise output by almost 25% within four years.^{26,27,28}

5.2 Labor market effects

We now turn to the labor market effects of robot adoption at the firm level. Specifically, we consider the effects on the firm’s employment (for specific groups of workers and overall), the labor cost share, and the average wage. Our theoretical considerations in the previous section imply that robot adopters will reduce their labor cost share, while the impact on total employment is ambiguous and depends on the relative strength of the displacement effect and the productivity effect. The employment effects might also be specific to certain groups of workers, especially to those performing automatable tasks (low-skilled workers as well as workers in the firm’s manufacturing rather than service-oriented establishments). As for the wage effects, our theoretical framework implies that the average wage in firms adopting robots increases if the firm changes the composition of its workforce by hiring relatively more high-skilled workers (and given a positive exogenous skill premium). To shed light on these effects, we estimate an equation akin to Eq. (8), where we use a variety of different labor market variables as dependent variables. Table 5 reports the results. In Panel A we control for selection into robot adoption by including the same set of time-varying selection controls as before. In Panel B we combine the firm fixed effects estimator with our propensity score weighting approach. All models include firm and industry-year fixed effects.

²⁶Since we have robot information available in our data, not every year, but every four years, there is some uncertainty regarding the precise timing of first time robot adoption. A firm that reports robot use in $t - 4$, but no robot use in $t - 8$, can have adopted robots for the first time in either $t - 4$, $t - 5$, $t - 6$, or $t - 7$. Hence, the most conservative interpretation is that the adoption of robots raises output by almost 25% within seven years.

²⁷Since we use labor productivity in $t - 4$ as a selection control and the level and growth of labor productivity in the propensity score estimates, we restrict our attention here to the effects of robots on physical output of firms. In the next section, we will also analyze the productivity effects of robot adoption.

²⁸In an additional set of estimates, we investigate whether the output gains from robot adoption are more pronounced in firms that are more integrated into the global economy. We do not find robust evidence that exporters (or importers or foreign-owned firms) experience stronger output gains from adopting robots.

Table 5: Labor market effects of robot adoption

	Employment	Labor cost share	Low-skilled	High-skilled	Manufacturing employment	Share of manuf. employment	Average wage
	(1a)	(2a)	(3a)	(4a)	(5a)	(6a)	(7a)
<i>PANEL A: Selection Controls</i>							
Robots _t	0.0582** (0.0252)	-0.0362*** (0.00883)	0.0595** (0.0261)	0.0820** (0.0384)	0.0410 (0.0280)	-0.00668 (0.00569)	-0.00259 (0.0118)
Robots _{t-4}	0.0532** (0.0246)	-0.0317*** (0.0109)	0.0413 (0.0252)	0.106*** (0.0379)	0.0499* (0.0260)	-0.00484 (0.00516)	-0.0151 (0.0161)
Observations	4572	4541	4549	4549	4565	4565	4532
R-squared	0.201	0.158	0.209	0.140	0.203	0.062	0.615
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>PANEL B: Propensity Score</i>							
	(1b)	(2b)	(3b)	(4b)	(5b)	(6b)	(7b)
Robots _t	0.0594** (0.0271)	-0.0253** (0.0114)	0.0703** (0.0279)	0.0571 (0.0437)	0.0512* (0.0285)	-0.00333 (0.00478)	0.0103 (0.0166)
Robots _{t-4}	0.0646* (0.0350)	-0.0289** (0.0138)	0.0637* (0.0347)	0.0562 (0.0485)	0.0633* (0.0347)	-0.00467 (0.00503)	-0.0144 (0.0183)
Observations	4632	4595	4611	4611	4624	4624	4585
R-squared	0.208	0.202	0.222	0.156	0.237	0.121	0.662
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: All dependent variables are expressed in logs except for the labor cost share and the share of manufacturing employment. Robots is a 0/1 indicator variable equal to one if the firm uses robots in the specified period. Selection controls (in $t - 4$) are the firm's deflated labor productivity (in logs), deflated capital intensity (in logs), skill intensity (in logs), deflated R&D intensity (in logs), as well as exporter, importer, and foreign ownership dummies. We add one to all factor intensity variables before taking logs in order to keep zero observations. For details on the propensity score reweighting estimator see the text. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Robust standard errors clustered by firm are given in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

A striking result in Panel A in Table 5 is that within four years robot adopters raise overall employment by around 10 percent. Importantly, while positive employment effects turn out to be especially pronounced for high-skilled workers, they also apply to other types of workers, namely low-skilled workers as well as workers employed in the firm’s manufacturing establishments. Moreover, while the labor cost share decreases significantly due to robot adoption, by almost 7 percentage points, we find no significant effect on the firm’s average wage, although the coefficient is estimated with a negative sign. The results based on the propensity score estimates in Panel B, by and large, confirm these results, although the positive employment effects for the group of high-skilled workers are smaller in magnitude and lose significance.

5.3 Effects of alternative systems in the production process

As detailed in Section 2, firms are not only asked whether they use robots in their production process, but they also report alternative systems they use, namely computer-digital machine tools (CAM), computer-assisted design (CAD), and a combination of some of the systems through a central computer (FLEX).²⁹ We use this information to investigate differences between the effects of robots and those of alternative technologies in the production process. Moreover, we want to see whether the estimated output and labor market effects of robot adoption are robust to controlling for alternative technologies in production. To save on space, we briefly summarize the main findings from this analysis here and relegate detailed regression output and discussion to the online supplement of this paper.

The main findings are as follows. First and foremost, the output and labor market effects of robot adoption reported above are fully robust to controlling for alternative systems in the production process. Secondly, both the adoption of computer-digital machine tools and of flexible manufacturing systems through a central computer raise firm output. However, the effects are smaller in magnitude than in the case of robot adoption, and they are not robust to including all technologies (including robots) simultaneously in the estimation. Adopting computer-assisted designs has no statistically significant effect on output. Thirdly, there is evidence for positive employment effects across all three technologies (CAD, CAM, and FLEX) and in all skill groups. In addition, we find robust evidence that the use of computer-digital machine tools disproportionately benefits high-skilled and non-manufacturing workers. Finally, we identify a striking difference between robot technology and other technologies used in the firm’s production process: only robots lead to a significant reduction in the firm’s labor cost share.

6 Robot adoption and intra-industry reallocations

In the previous sections we have documented that robot adoption is much more likely in ex-ante larger and more productive firms and that firms experience substantial output and employment

²⁹Descriptive analyses show that robots are used less frequently in the production process than these other system, and that there is a slight positive correlation between the use of robots and these other systems.

gains following robot adoption. In this section we investigate how non-adopting firms are affected by the use of robot technology within their industries. We also shed light on the contribution of robots to gains in aggregate productivity, and we decompose this contribution into two components: (1) direct efficiency gains within robot adopters and (2) indirect gains through reallocation of labor from non-adopters to adopters.

6.1 Robot density and its impact on non-adopting firms

We first estimate the effects of robot diffusion within industries on non-adopting firms. To do so, we estimate variants of the following equation:

$$\text{Outcome}_{it} = \gamma_1 \text{Robot-density}_{st} + \gamma_2 \text{Robot-density}_{st} \times \text{Robot-use}_i + \beta_1 \mathbf{X}_{it-4} + \mu_i + \mu_t + \varepsilon_{it}, \quad (9)$$

where we use firm-level employment, output, and market exit as different outcomes of firm i at time t , and where we interact a time-varying industry-specific measure of robot density with a time-constant firm-level dummy variable for the use of robots. The variable $\text{Robot-density}_{st}$ is constructed in two different ways using two different data-sets. First, we use our ESEE data-set and define the variable as the share of sales attributable to robot-using firms in total industry sales.^{30,31} This measure is only available in those years in which we have information on robot use in the survey (i.e. every four years). In an alternative approach, we use data from the International Federation of Robotics (IFR) and more specifically the industry-specific stock of robots over the period 1993 to 2016.³² This measure of robot density is available on an annual basis and features yearly variation. The variable Robot-use_i in Eq. (9) is a 0/1 indicator variable equal to one if the firm uses robots at least once during our sample period, and zero otherwise.

The coefficients of interest are γ_1 and γ_2 . The first coefficient tells us the effect of rising robot density in an industry on non-adopting firms, while the second coefficient tells us the difference in the effect of robot density on robot adopters vs. non-adopters. We report the two coefficients in Table 6, where Panel A, B, and C focus on the effects on employment, output, and market exit, respectively. In columns (1) to (3) and columns (4) to (6), we use our robot density measure from the ESEE data and the IFR data, respectively. All specifications include both firm fixed effects (μ_i) and year fixed effects (μ_t). In columns (2), (3), (5), and (6), we include our selection controls for robot adoption in the vector \mathbf{X}_{it-4} (see Section 4). To make sure that our results are indeed due to differences in robot density across industries, and not other important industry-specific factors, we also augment the model by including time-varying industry-specific factor intensity variables

³⁰We have verified that our results are robust to alternative definitions of this variable using the ESEE data, viz. the share of robot-using firms in the total number of firms, the share of output attributable to robot-using firms in industry output, and the share of employment in robot-using firms in total industry employment.

³¹To construct meaningful measures of robot density, when computing this variable, we do not restrict the sample to firms that do not use robots in the first year they appear in the sample. We also use the full sample of firms in the estimation. However, our results do not change when restricting the estimation sample to firms that do not use robots in the first year they appear in the sample.

³²Table A.6 in Appendix A.5 describes the concordance between the different industry classifications in the ESEE and the IFR data-sets.

(namely annual industry averages of capital, skill, and R&D intensity), and we also interact these variables with our firm-level indicator variable for robot use; see columns (3) and (6).³³

Table 6: Robot adoption and intra-industry reallocations

	ESEE			IFR		
	(1a)	(2a)	(3a)	(4a)	(5a)	(6a)
<i>PANEL A: Employment in t</i>						
Robot-density _t	-0.193*** (0.0488)	-0.222*** (0.0716)	-0.210*** (0.0737)	-0.00630 (0.00963)	-0.00389 (0.0130)	-0.00288 (0.0133)
Robot-density _t × Robot-use	0.190*** (0.0732)	0.313*** (0.0968)	0.255*** (0.0988)	0.0266** (0.0122)	0.0380** (0.0150)	0.0360** (0.0163)
Observations	13372	6104	6104	34426	21590	21590
R-squared	0.080	0.125	0.128	0.114	0.142	0.143
<i>PANEL B: Output in t</i>						
Robot-density _t	-0.356*** (0.0645)	-0.401*** (0.0896)	-0.333*** (0.0957)	-0.0354*** (0.0122)	-0.0312* (0.0160)	-0.0257 (0.0160)
Robot-density _t × Robot-use	0.463*** (0.0911)	0.601*** (0.123)	0.401*** (0.131)	0.0807*** (0.0150)	0.0970*** (0.0185)	0.0846*** (0.0197)
Observations	13417	6103	6103	34358	21557	21557
R-squared	0.145	0.163	0.170	0.143	0.139	0.140
<i>PANEL C: Exit in t + 1</i>						
Robot-density _t	0.0475** (0.0215)	0.0581* (0.0330)	0.0417 (0.0329)	0.00468* (0.00260)	0.00938** (0.00366)	0.00726* (0.00384)
Robot-density _t × Robot-use	-0.0872*** (0.0274)	-0.0816** (0.0406)	-0.0565 (0.0404)	-0.0143*** (0.00261)	-0.0128*** (0.00367)	-0.00831** (0.00419)
Observations	12958	5558	5558	33643	19223	19223
R-squared	0.027	0.033	0.035	0.033	0.034	0.035
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Selection controls	No	Yes	Yes	No	Yes	Yes
Industry controls + interact.	No	No	Yes	No	No	Yes

Notes: In columns (1) to (3) we define robot density as the share of sales attributable to robot-using firms in total industry sales constructed from the ESEE data. In columns (4) to (6) we use the stock of robots in an industry (in logs) constructed from the IFR data. The variable robot-use is a 0/1 indicator variable equal to one if the firm uses robots (at least once) during our sample period, and zero otherwise. In Panels A and B we use employment and deflated output (both in logs) as the dependent variables, respectively. In Panel C, we use a 0/1 indicator variable as the dependent variable; this variable is equal to one if the firm exits the market in the next period, and zero if it continues its operations. Selection controls include the firm's deflated labor productivity (in logs), deflated capital intensity (in logs), skill intensity (in logs), deflated R&D intensity (in logs), exporter status, importer status, and foreign ownership status (all in $t-4$). We add one to all factor intensity variables before taking logs in order to keep zero observations. Industry controls are annual industry averages of capital, skill, and R&D intensity; these variable are also interacted with the time-constant firm-specific robot-use dummy variable. Robust standard errors clustered by firm are given in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

The negative estimates of γ_1 in Panel A of Table 6 show that an increase in robot density has a significant negative impact on employment in firms that do not adopt robot technology. The estimates in the first three columns imply that 10% of jobs in non-adopting firms are destroyed when the share of sales attributable to robot-using firms increases from zero to one-half. Importantly, the positive and significant estimates of γ_2 indicate that these effects are exclusive to non-adopters.

³³Following the analysis in Section 5.3, we have also run an additional robustness analysis (not reported), where we augment the model to include density measures for the use of computer-digital machine tools (CAM), computer-assisted design (CAD), or a combination of some of the systems through a central computer (FLEX). The results reported here are robust to this augmented specification.

Looking at Panel B, we see the same pattern of effects in terms of output, but the implied magnitude is even more pronounced. In Panel C, we document higher exit rates among non-adopters due to an increase in robot density, which is consistent with the predicted increase in the survival cut-off productivity in our theoretical framework. Importantly, we find similar results on employment, output, and exit rates when using the stock of robots within industries from the IFR data. This is remarkable because the IFR measure captures the intensive margin of robot diffusion, regardless of how many firms use this technology, whereas the ESEE measure reflects the share of firms using robots and thus the extensive margin of robot use.

Taking stock, we provide strong support for the idea that robot adopters expand their scale of operations and create jobs, while non-adopters experience negative output and employment effects in the face of tougher competition with high-technology firms. Our results thus imply substantial intra-industry reallocation of market shares and resources as a result of more widespread diffusion of robot technology and a polarization between high-productivity robot adopters and low-productivity non-adopters.


6.2 Decomposing productivity gains

We now investigate how the diffusion of robot technology affected aggregate productivity from 1990 to 2016. To do so, we first estimate time-varying and firm-specific total factor productivity (TFP) using a standard production function approach as described in Olley and Pakes (1996). To obtain a measure of *aggregate* TFP, we compute the weighted average of all firm-specific TFPs using employment shares as weights. The solid black line in Figure 6 depicts the thus constructed measure of aggregate TFP. To simplify the analysis, we restrict the analysis to a balanced sample of firms between 1990 to 2016. We see a steady rise in TFP with a stagnation during the crisis in the early 1990s and a strong negative shock around the time of the global financial crisis. TFP tripled over the 27-year period of our sample (we normalize the value of TFP in 1990 to 100), which corresponds to a plausible annual growth rate of about 4%.

To evaluate the extent to which the diffusion of robot technology contributed to this development, we consider two factors: (1) *direct* technical efficiency gains arising from robot adoption at the firm level, and (2) *indirect* gains arising from the reallocation of labor towards robot adopters and away from non-adopters. To quantify the first component, we estimate the effect of robot adoption on firm-level TFP following our analysis described in Section 5.1. Specifically, adopting the specification with selection controls, we obtain estimated coefficients of Robots_t and Robots_{t-4} equal to 0.135 and 0.084, respectively.³⁴ Both coefficients are statistically significant at the 1 percent level. These estimates are used to correct the TFP growth in robot adopters, which allows us to construct a counterfactual evolution of aggregate TFP eliminating firm-specific TFP gains due to robot adoption. This is represented by the long-dashed dark grey line in Figure 6. To quantify the second component, the gains stemming from labor reallocation, we use our estimates in column (1a)

³⁴These are estimates based on a specification without the forward variable Robots_{t+4} . When including this forward variable, its coefficient is not significantly different from zero.

Figure 6: Actual vs. counterfactual evolution of aggregate TFP (1990-2016)



Notes: The figure depicts the evolution of aggregate TFP (constructed as the weighted average of firm-level TFP using employment shares as weights) in a balanced sample of firms from 1990-2016. The solid black line depicts the actual evolution; the long-dashed dark grey line depicts the counterfactual evolution eliminating firm-specific TFP gains due to robot adoption; the short-dashed light grey line depicts the counterfactual evolution eliminating firm-specific TFP gains due to robot adoption, as well as labor reallocations caused by firm-level robot adoption and rising robot density at the industry level.

Source: Authors' computations based on ESEE data.

of Table 5 and column (3a) of Table 6 to adjust the employment weights for both robot adopters and non-adopters. This allows us to construct a counterfactual evolution of aggregate TFP eliminating not just firm-specific TFP gains due to robot adoption, but also changes in aggregate productivity due to labor reallocations caused by robot adoption of individual firms as well as rising robot density at the industry level. This is represented by the short-dashed light grey line in Figure 6.

Our analysis reveals that, without the availability and adoption of robot technology, TFP would have doubled over the period from 1990 to 2016 rather than tripled. Moreover, we see that the first component, direct technical efficiency gains, explains about two thirds of the total TFP gains attributable to robots, while the second component, the gains due to labor reallocation, explains the remaining one third. However, note that our analysis here is based on a balanced sample of firms and thus abstracts from market exit of low-productivity non-adopters. Hence, we provide a lower bound for the total TFP gains due to robots, as well as for the contribution of the second component to these gains.

7 Conclusion

This paper provides novel evidence on how automation in the form of robot adoption affects firm-level outcomes. We use detailed firm level information from a survey conducted on Spanish manufacturing firms over a 27-year period (1990-2016). We focus on three central questions: (1) Which firms adopt robots? (2) What are the labor market effects of robot adoption at the firm level? (3) How does firm heterogeneity in robot adoption affect the industry equilibrium? As for the first question, we establish robust evidence that ex-ante larger and more productive firms and exporters are more likely to adopt robots, while ex-ante more skill-intensive firms are less likely to do so. As for the second question, we find that robot adoption generates substantial output gains in the vicinity of 20-25% within four years, reduces the labor cost share by 5-7%-points, and leads to net job creation at a rate of 10%. Finally, we reveal substantial job losses in firms that do not adopt robots, and a productivity-enhancing reallocation of labor across firms, away from non-adopters, and toward adopters.

By focusing attention on heterogeneity in robot adoption within narrowly defined industries, our results provide novel evidence how robots affect industry heterogeneity. Importantly, we do not find any negative employment effects in those firms that start to adopt robots, even if we focus on specific skills or groups of workers. On the contrary, our results robustly show that robot adopters create jobs in the subsequent years, relative to the control group, i.e., their competitors that do not adopt robots. In other words, negative employment effects materialize where they are ex-ante the least expected, namely in those firms that do not automate their production process. Hence, our study points to the importance of reallocation of resources across firms within industries, as robots are creating new opportunities for some firms, while simultaneously leading to job losses in non-adopting firms.

References

- Acemoglu, Daron and Autor, David H.** (2011), Chapter 12 - Skills, tasks and technologies: Implications for employment and earnings, *in Orley Ashenfelter and David Card.*, eds, 'Handbooks in Economics Handbook of Labor Economics', Vol. 4, Part B, Elsevier, pp. 1043–1171.
- Acemoglu, Daron and Restrepo, Pascual.** (2017), Robots and jobs: Evidence from US labor markets, Working Paper 23285, National Bureau of Economic Research.
- Acemoglu, Daron and Restrepo, Pascual.** (2018*a*), Modeling automation, *in* 'AEA Papers and Proceedings', Vol. 108, pp. 48–53.
- Acemoglu, Daron and Restrepo, Pascual.** (2018*b*). 'The race between man and machine: Implications of technology for growth, factor shares, and employment', *American Economic Review* 108(6), 1488–1542.
- Antràs, Pol, Garicano, Luis and Rossi-Hansberg, Esteban.** (2006). 'Offshoring in a knowledge economy', *Quarterly Journal of Economics* 121(1), 31–77.
- Armenter, Roc and Koren, Miklos.** (2015). 'Economies of scale and the size of exporters.', *Journal of the European Economic Association* 13(3), 482 – 511.
- Autor, David H, Levy, Frank and Murnane, Richard J.** (2003). 'The skill content of recent technological change: An empirical exploration', *Quarterly Journal of Economics* 118(4), 1279–1333.
- Bustos, Paula.** (2011). 'Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms', *American Economic Review* 101(1), 304–40.
- Capuano, Stella, Egger, Hartmut, Koch, Michael and Schmerer, Hans-Jörg.** (2017), Offshoring and firm overlap, CESifo Working Paper Series 6361, CESifo Group Munich.
- Dauth, Wolfgang, Findeisen, Sebastian, Suedekum, Jens and Woessner, Nicole.** (2018), Adjusting to robots: Worker-level evidence, Working Papers 13, Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute.
- Davis, Donald R. and Harrigan, James.** (2011). 'Good jobs, bad jobs, and trade liberalization', *Journal of International Economics* 84(1), 26–36.
- DiNardo, John, Fortin, Nicole M and Lemieux, Thomas.** (1996). 'Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach', *Econometrica* 64(5), 1001–1044.
- Eaton, Jonathan, Kortum, Samuel and Kramarz, Francis.** (2011). 'An anatomy of international trade: Evidence from French firms.', *Econometrica* 79(5), 1453 – 1498.


- Egger, Hartmut, Kreickemeier, Udo and Wrona, Jens.** (2015). ‘Offshoring domestic jobs’, *Journal of International Economics* 97, 112–125.
- Frey, Carl Benedikt and Osborne, Michael A.** (2017). ‘The future of employment: How susceptible are jobs to computerisation?’, *Technological Forecasting and Social Change* 114, 254 – 280.
- Graetz, Georg and Michaels, Guy.** (2018). ‘Robots at work’, *Review of Economics and Statistics* 100(5), 753–768.
- Groizard, Jose L., Ranjan, Priya and Rodriguez-Lopez, Antonio.** (2014). ‘Offshoring and jobs: The myriad channels of influence’, *European Economic Review* 72, 221–239.
- Grossman, Gene M. and Rossi-Hansberg, Esteban.** (2008). ‘Trading tasks: A simple theory of offshoring’, *American Economic Review* 98(5), 1978–1997.
- Guadalupe, Maria, Kuzmina, Olga and Thomas, Catherine.** (2012). ‘Innovation and foreign ownership’, *American Economic Review* 102(7), 3594–3627.
- Hallak, Juan Carlos and Sivadasan, Jagadeesh.** (2013). ‘Product and process productivity: Implications for quality choice and conditional exporter premia.’, *Journal of International Economics* 91(1), 53 – 67.
- Harrigan, James and Reshef, Ariel.** (2015). ‘Skill biased heterogeneous firms, trade liberalization, and the skill premium’, *Canadian Journal of Economics* 48(3), 1024–1066.
- Helpman, Elhanan, Itskhoki, Oleg, Muendler, Marc-Andreas and Redding, Stephen J.** (2017). ‘Trade and inequality: From theory to estimation’, *Review of Economic Studies* 84(1), 357–405.
- Jäger, Angela, Moll, Cornelius and Lerch, Christian.** (2015). ‘Analysis of the impact of robotic systems on employment in the European Union - 2012 data update. Update of Final Report’, *Publications Office of the European Union: Luxembourg, Luxembourg* .
- Koch, Michael.** (2016). ‘Skills, tasks and the scarcity of talent in a global economy’, *Review of International Economics* 24(3), 536–563.
- Lileeva, Alla and Trefler, Daniel.** (2010). ‘Improved access to foreign markets raises plant-level productivity ... for some plants’, *Quarterly Journal of Economics* 125(3), 1051–1099.
- Melitz, Marc J.** (2003). ‘The impact of trade on intra-industry reallocations and aggregate industry productivity’, *Econometrica* 71(6), 1695–1725.
- Olley, G. Steven and Pakes, Ariel.** (1996). ‘The dynamics of productivity in the telecommunications equipment industry’, *Econometrica* 64(6), 1263–1297.

- Ornaghi, Carmine.** (2006). ‘Assessing the effects of measurement errors on the estimation of production functions’, *Journal of Applied Econometrics* 21(6), 879–891.
- Raj, Manav and Seamans, Robert.** (2018), AI, labor, productivity, and the need for firm-level data, *in* ‘The economics of artificial intelligence: an agenda’, University of Chicago Press.
- Syverson, Chad.** (2011). ‘What determines productivity?’, *Journal of Economic Literature* 49(2), 326–65.
- UNCTAD.** (2017), *Trade and Development Report 2017. Beyond austerity: Towards a global new deal*, United Nations, New York and Geneva.

A Appendix

A.1 Labor cost share for robot adopters vs. non-adopters

Figure A.1: Evolution of firm-level labor cost share (1990-2016)



Notes: The figure depicts the evolution of average firm labor cost share (defined as labor costs divided by the total production) in a balanced sample of firms from 1990-2016, separately for robot adopters (solid black line) and non-adopters (dashed grey line). Robot adopters are defined as firms that entered the sample in 1990 and had adopted robots by 1998. Non-adopters are firms that never use robots over the whole sample period.

Source: Authors' computations based on ESEE data.

A.2 Detailed derivations corresponding to the theory in Section 3

We use this section to provide analytical details corresponding to the results presented in Sections 3.2 and 3.3 in the main text of the paper. We proceed as follows. First, we discuss the main results of the basic model with one source of firm heterogeneity, viz. across-firm differences in $\phi(\omega)$ but no differences in the degree of complexity, $N(\omega) = N$. These results form the basis of our discussion in Section 3.3. Secondly, we introduce the possibility of exporting and discuss its implications. This corresponds to Section 3.2.2. Finally, we allow for firm heterogeneity in $N(\omega)$ and consider two skill types of labor, low- and high-skilled workers. This is relevant for Section 3.2.3.

Basic model. In the model, firms differ in their baseline productivity $\phi(\omega)$ and the complexity of their production process $N(\omega)$. In a first step, we focus on just one-dimensional heterogeneity by assuming that all firms have to perform the same set of tasks, given by $N(\omega) = N$. Hence, firms are fully described by their productivity ϕ and we can omit the firm index ω to save on notation. We can write firm profits for robot adopters and non-adopters, respectively, as

$$\pi(\phi) = (1 - \beta)A \left\{ \frac{1}{\beta} [\gamma(N, N - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F \quad \text{and} \quad (\text{A.1})$$

$$\pi^a(\phi) = (1 - \beta)A \left\{ \frac{1}{\beta} [\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F - F^a. \quad (\text{A.2})$$

Given that robots have a comparative advantage in all tasks $i \leq I$, we know that $[\gamma(N, N - 1)w^{1-\sigma}] < [\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]$. Without loss of generality, we normalize the left-hand side by setting $[\gamma(N, N - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} = 1$ and we define $[\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]^{\frac{1}{1-\sigma}} = 1/\bar{\eta}$ with $\bar{\eta} > 1$. Furthermore, we choose $F^a = (\alpha - 1)F$ with $\alpha > 1$. We can thus rewrite profits as

$$\pi(\phi) = (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi} \right)^{-\frac{\beta}{1-\beta}} - F, \quad (\text{A.3})$$

$$\pi^a(\phi) = (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi\bar{\eta}} \right)^{-\frac{\beta}{1-\beta}} - \alpha F. \quad (\text{A.4})$$

To determine the domestic cut-off productivity, denoted by ϕ^* , we can use $\pi(\phi^*) = 0$. The cut-off productivity for robot adoption ϕ^r can be determined by using the indifference condition $\pi(\phi^r) = \pi^a(\phi^r)$ along with $\pi(\phi^*) = 0$ to compute

$$\phi^r = \phi^* \left(\frac{\alpha - 1}{\frac{\beta}{\bar{\eta}^{\frac{1-\beta}{1-\beta}} - 1}} \right)^{\frac{1-\beta}{\beta}}. \quad (\text{A.5})$$

Using these cut-off productivities, we can define the share of firms that use robots as

$$s_r \equiv \frac{1 - G(\phi^r)}{1 - G(\phi^*)}, \quad (\text{A.6})$$

where $G(\cdot)$ denotes the cumulative distribution function of productivity. From inspection of Equation (A.5) we can conclude that a lower fixed cost for robot adoption or a higher share of automatable tasks (and thus $\bar{\eta}$) raises the share of robot adopters, i.e. $\partial s_r / \partial \alpha < 0$ and $\partial s_r / \partial I > 0$.

Discussing the implications for the composition of firms within industries requires to also specify the details on the entry (and exit) process of firms. As this is standard in the literature on heterogeneous firms, we refer the interested reader for details to Melitz (2003). Here, we briefly outline how the endogenous cut-off productivity ϕ^* can be determined. Specifically, it is determined by two conditions. The first condition uses the relation between the average profit per firm and the cut-off productivity level, the so-called zero-cutoff productivity. It can be computed as the average profits over all active firms, that is

$$\bar{\pi} = (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\tilde{\phi}} \right)^{-\frac{\beta}{1-\beta}} - F - F(\alpha - 1) \frac{1 - G(\phi^r)}{1 - G(\phi^*)}, \quad (\text{A.7})$$

where $\bar{\pi}$ denotes the average profits over all active firms and $\tilde{\phi}$ is the average (expected) productivity level, defined as

$$\tilde{\phi} \equiv \left(\int_{\phi^*}^{\phi^r} \phi^{\frac{\beta}{1-\beta}} \frac{g(\phi)}{1 - G(\phi^*)} d\phi + \int_{\phi^r}^{\infty} (\eta\phi)^{\frac{\beta}{1-\beta}} \frac{g(\phi)}{1 - G(\phi^*)} d\phi \right)^{\frac{1-\beta}{\beta}}. \quad (\text{A.8})$$

The second condition, called the free entry condition, requires that the net value of entry is zero, i.e. the sunk market entry costs (f_e) are equal to the expected profits (discounted by δ). Formally, this condition reads as:

$$\bar{\pi} = \frac{\delta f_e}{1 - G(\phi^*)}. \quad (\text{A.9})$$

Both equations can be used to determine a unique cut-off productivity level and show that a lower fixed cost of robot adoption or a higher level of robot technology affects the composition of firms within industries. Following Melitz (2003), we know that ex-ante more productive firms gain market share by reducing marginal costs due to robot adoption. This raises the cut-off productivity at which firms are able to survive in the market. Put differently, increasing robot exposure raises the exit rate among non-robot firms and reduces their output and employment. This proves the results described in Section 3.3.

Exporting. When allowing for trade with a symmetric partner country, we can sort firms into four groups, namely combinations of robot adopters vs. non-adopters (indicating robot adopters by a superscript a) and exporters vs. non-exporters (indicated by subscripts x and d , respectively). Specifically, we can write firm profits for the different types as

$$\pi_d(\phi) = (1 - \beta)A \left\{ \frac{1}{\beta} [\gamma(N, N - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F, \quad (\text{A.10})$$

$$\pi_d^a(\phi) = (1 - \beta)A \left\{ \frac{1}{\beta} [\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F - F^a, \quad (\text{A.11})$$

$$\pi_x(\phi) = \left(1 + \tau^{-\frac{\beta}{1-\beta}}\right) (1 - \beta)A \left\{ \frac{1}{\beta} [\gamma(N, N - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F - F^x, \quad (\text{A.12})$$

$$\pi_x^a(\phi) = \left(1 + \tau^{-\frac{\beta}{1-\beta}}\right) (1 - \beta)A \left\{ \frac{1}{\beta} [\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]^{\frac{1}{1-\sigma}} \right\}^{-\frac{\beta}{1-\beta}} - F - F^a - F^x. \quad (\text{A.13})$$

Again, setting $[\gamma(N, N - 1)w^{1-\sigma}]^{\frac{1}{1-\sigma}} = 1$, defining $[\eta(N, I)r^{1-\sigma} + \gamma(N, I)w^{1-\sigma}]^{\frac{1}{1-\sigma}} = 1/\bar{\eta}$ with $\bar{\eta} > 1$, and setting $F^a = (\alpha - 1)F$, we can rewrite profits as

$$\pi_d(\phi) = (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi} \right)^{-\frac{\beta}{1-\beta}} - F, \quad (\text{A.14})$$

$$\pi_d^a(\phi) = (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi\bar{\eta}} \right)^{-\frac{\beta}{1-\beta}} - \alpha F, \quad (\text{A.15})$$

$$\pi_x(\phi) = \left(1 + \tau^{-\frac{\beta}{1-\beta}}\right) (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi} \right)^{-\frac{\beta}{1-\beta}} - F - F^x, \quad (\text{A.16})$$

$$\pi_x^r(\phi) = \left(1 + \tau^{-\frac{\beta}{1-\beta}}\right) (1 - \beta)A \left(\frac{1}{\beta} \frac{1}{\phi\bar{\eta}} \right)^{-\frac{\beta}{1-\beta}} - \alpha F - F^x. \quad (\text{A.17})$$

Except for different variable labels, this system is identical to the one described in Bustos (2011) (on page 310). We can thus build on her insights and follow the same steps. Accordingly, we focus on cost and parameter conditions that guarantee that the least productive firms serve only the domestic market and do not adopt robots, while more productive firms export and only the most productive exporters find it attractive to adopt robots. Importantly, the descriptive statistics obtained from our data and described in the main text reveal that the share of robot adopters is considerably lower than the share of exporting firms. It is therefore plausible to assume that the marginal exporter is a non-adopter, i.e., a firm that does not use robots. As shown in Bustos (2011), this is the case with a sufficiently high fixed cost of robot adoption relative to exporting. The exporter cut-off ϕ^x is determined by the indifference condition $\pi_d(\phi^x) = \pi_x(\phi^x)$. Combining this condition with $\pi_d(\phi^*) = 0$ entails

$$\phi^x = \phi^* \tau \left(\frac{F^x}{F} \right)^{\frac{1-\beta}{\beta}}. \quad (\text{A.18})$$

To determine the cut-off productivity for robot adoption in the open economy ϕ^r , we use $\pi_x(\phi^r) = \pi_x^a(\phi^r)$. Using the zero cut-off profit condition for the least productive firm, this allows us to

compute:

$$\phi^r = \phi^* \frac{1}{\left(1 + \tau^{-\frac{\beta}{1-\beta}}\right)^{\frac{\beta}{1-\beta}}} \left(\frac{\alpha - 1}{\bar{\eta}^{\frac{\beta}{1-\beta}} - 1}\right)^{\frac{1-\beta}{\beta}}. \quad (\text{A.19})$$

Using Equation (A.19), we can conclude that a reduction in variable trade costs τ raises the share of robot adopters, i.e. $\partial s_r / \partial \tau < 0$. As discussed in detail in Bustos (2011), we know that the incentives for adopting robots are higher for exporting firms, as the gains from doing so—the reduction in variable production costs—can be scaled up to a larger customer base in home and foreign. This completes the discussion corresponding to Section 3.2.2.

Skill heterogeneity. In the main text we briefly discuss an extension with two types of workers, namely low-skilled and high-skilled workers, indexed by subscripts l and h , respectively. Accordingly, we have

$$x(\omega, i) = \mathbb{1}[i \leq I] \eta(i) k(\omega, i) + \gamma_l(i) l_l(\omega, i) + \gamma_h(i) l_h(\omega, i). \quad (\text{A.20})$$

In such an environment, firms will not only compare the production costs of robots and human labor across tasks, but also consider the skill-specific effective labor costs in each task, i.e., the firm will benchmark $w_l / \gamma_l(i)$ against $w_h / \gamma_h(i)$. The task-level production function in (A.20) implies that low-skilled and high-skilled workers are substitutes in the performance of tasks. Following Acemoglu and Autor (2011), we impose a comparative advantage of high-skilled workers over their low-skilled coworkers that is increasing in the complexity of tasks. As discussed in detailed in Koch (2016), we can define a unique threshold task $z \in (0, 1)$ for which the firm is exactly indifferent between hiring low-skilled and hiring high-skilled workers, at prevailing skill premium $s \equiv w_h / w_l$. Put differently, the unit costs of performing task z are the same irrespective of the assigned skill type $k = l, h$. This establishes


$$w_l / \gamma_l(z) = w_h / \gamma_h(z). \quad (\text{A.21})$$

Koch (2016) discusses parameter constraints (on the comparative advantage schedule, factor endowments, etc.) within a general equilibrium framework that guarantee the existence of an interior solution, $z \in (N - 1, N)$. Intuitively, we need a skill premium that exceeds the productivity advantage of high-skilled workers in some tasks. Under this constraint, we can establish that low-skilled workers will be assigned to all tasks $i < z$, while high-skilled workers will be assigned to all tasks $i \geq z$. Under the additional constraint that robots cannot automate all tasks performed by low-skilled workers, $I < z$, we obtain for the unit production costs

$$c^a(\phi, N, I) = \frac{1}{\phi} \left[\eta(N, I) r^{1-\sigma} + \gamma_l(I, z) w^{1-\sigma} + \gamma_h(N, z) w^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (\text{A.22})$$

where $\eta(N, I) \equiv \left(\int_{N-1}^I \eta(i)^{\sigma-1} di \right)^{\frac{1}{\sigma}}$, $\gamma_l(I, z) \equiv \left(\int_I^z \gamma_l(i)^{\sigma-1} di \right)^{\frac{1}{\sigma}}$ and $\gamma_h(N, z) \equiv \left(\int_z^N \gamma_h(i)^{\sigma-1} di \right)^{\frac{1}{\sigma}}$. In the main text, we use this extension with two skill types to conclude that firms with a higher skill intensity are less likely to adopt robots. Therefore, we now also consider heterogeneity of firms in the complexity of the production process, $N(\omega)$. For ease of exposition, we assume that some firms operate with a complexity equal to N , while others operate with $N^c > N$.³⁵ Suppose that $I > N^c - 1$, so that there is always some tasks that are automatable. Figure A.2 illustrates this situation. It is evident that more complex firms have (i) a higher share of tasks that are performed by high-skilled workers and (ii) that in these firms only a smaller fraction of tasks can be performed by robots. It follows that firms with a lower skill intensity are more likely to adopt robots. This completes the discussion corresponding to Section 3.2.3.

Figure A.2: Skill allocation and automatable tasks for different complexities of the production process



³⁵Different studies in the field of international economics have extended the Melitz (2003) framework to allow for heterogeneity in more than one dimension. Prominent examples include Davis and Harrigan (2011), Eaton et al. (2011), Hallak and Sivadasan (2013), Armenter and Koren (2015), Harrigan and Reshef (2015), and Helpman et al. (2017). For instance, Harrigan and Reshef (2015) also consider two types of labor with firms differing in both the baseline productivity φ and the Cobb-Douglas share parameter α which governs the skill intensity of the firm. They characterize firms by their “competitiveness”, determined by both φ and α , and they apply the theory of copulas from mathematical statistics to determine the distribution of firms’ competitiveness allowing for flexible correlations between φ and α . Another example is Capuano et al. (2017), who allow for two-dimensional heterogeneity in the context of offshoring. In their framework, firms differ in the range of tasks to be performed as well as the share of offshorable tasks.

A.3 Further results on robot adoption

Table A.1: Robot adoption A.I: Probit cross-sectional and panel specification

	Robot adoption (0/1 indicator)			
<i>PANEL A: Cross-sectional specification</i>	(1a)	(2a)	(3a)	(4a)
Base year labor productivity	0.197*** (0.0425)	0.0863* (0.0499)	0.0931** (0.0448)	0.0356 (0.0516)
Base year skill intensity		-2.373*** (0.718)		-2.838*** (0.793)
Base year exporter status			0.350*** (0.0646)	0.289*** (0.0726)
Observations	3272	2732	3208	2689
Pseudo R-squared	0.061	0.100	0.089	0.114
		Robot adoption (0/1 indicator)		
<i>PANEL B: Panel specification</i>	(1b)	(2b)	(3b)	(4b)
Lagged labor productivity	0.270*** (0.0390)	0.116*** (0.0404)	0.147*** (0.0394)	0.0564 (0.0403)
Lagged skill intensity		-0.890** (0.442)		-1.213** (0.485)
Lagged exporter status			0.231*** (0.0537)	0.178*** (0.0570)
Observations	7225	6738	7157	6683
R-squared	0.060	0.082	0.080	0.093
Industry(-base)-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	No	Yes	Yes

Notes: In Panel A the dependent variable in all columns is a 0/1 indicator variable equal to one if the firm adopts robots during our sample period and zero otherwise, while in Panel B it is equal to one if the firm uses robots in a specific year and zero otherwise. Labor productivity is the firm's deflated value added per worker (in logs). Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. All estimates in Panel A (B) include industry-base-year (industry-year) fixed effects. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. In Panel A all explanatory variables are measured in the base year defined as the first year the firm appears in the sample. In Panel B all explanatory variables are lagged by one year. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Moreover, in Panel B we condition on the firm not using robots in the previous year (or the most recent year for which robot data are available for that firm). Robust standard errors are given in parentheses and clustered by firm in Panel B. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

Table A.2: Robot adoption A.II: Probit cross-sectional and panel specification

	Robot adoption (0/1 indicator)			
<i>PANEL A: Cross-sectional specification</i>	(1a)	(2a)	(3a)	(4a)
Base year output	0.219*** (0.0171)	0.184*** (0.0220)	0.185*** (0.0220)	0.153*** (0.0267)
Base year skill intensity		-2.596*** (0.693)		-2.658*** (0.726)
Base year exporter status			0.190*** (0.0661)	0.180** (0.0739)
Observations	3436	2839	3368	2793
Pseudo R-squared	0.105	0.123	0.111	0.127
		Robot adoption (0/1 indicator)		
<i>PANEL B: Panel specification</i>	(1b)	(2b)	(3b)	(4b)
Lagged output	0.231*** (0.0141)	0.211*** (0.0171)	0.227*** (0.0179)	0.204*** (0.0203)
Lagged skill intensity		-1.498*** (0.486)		-1.454*** (0.491)
Lagged exporter status			0.0356 (0.0562)	0.0244 (0.0595)
Observations	7424	6891	7350	6831
Pseudo R-squared	0.112	0.115	0.113	0.115
Industry(-base)-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	No	Yes	Yes

Notes: In Panel A the dependent variable in all columns is a 0/1 indicator variable equal to one if the firm adopts robots during our sample period and zero otherwise, while in Panel B it is equal to one if the firm uses robots in a specific year and zero otherwise. Output is the firm's deflated output value (in logs). Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. All estimates in Panel A (B) include industry-base-year (industry-year) fixed effects. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. In Panel A all explanatory variables are measured in the base year defined as the first year the firm appears in the sample. In Panel B all explanatory variables are lagged by one year. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Moreover, in Panel B we condition on the firm not using robots in the previous year (or the most recent year for which robot data are available for that firm). Robust standard errors are given in parentheses and clustered by firm in Panel B. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

Table A.3: Robot adoption based on productivity quartiles: linear cross-sectional specification

	Robot adoption (0/1 indicator)			
	— <i>Labor productivity</i> —		— <i>Output</i> —	
	(1)	(2)	(3)	(4)
<i>Productivity</i>				
Base year 2nd quartile	0.0443*** (0.0167)	0.0185 (0.0184)	0.0378*** (0.0139)	0.0205 (0.0164)
Base year 3rd quartile	0.0531*** (0.0166)	0.0213 (0.0189)	0.0792*** (0.0148)	0.0507*** (0.0186)
Base year 4th quartile	0.0780*** (0.0170)	0.0179 (0.0207)	0.205*** (0.0173)	0.146*** (0.0245)
Base year skill intensity		-0.400*** (0.100)		-0.407*** (0.101)
Base year exporter status		0.0573*** (0.0161)		0.0346** (0.0160)
Observations	4053	3443	4221	3551
R-squared	0.111	0.161	0.141	0.170
Industry-base-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	Yes	No	Yes

Notes: The dependent variable in all columns is a 0/1 indicator variable equal to one if the firm adopts robots during our sample period and zero otherwise. The regressions include a full set of dummy variables indicating the firm's (quartile) position in the productivity distribution of the industry in which it is active. Columns (1) and (2) do this based on labor productivity, i.e., the firm's deflated value added per worker, while columns (3) and (4) use output, i.e., the firm's deflated output value. Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. All explanatory variables are measured in the base year defined as the first year the firm appears in the sample. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Robust standard errors are given in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

Table A.4: Robot adoption based on productivity quartiles: linear panel specification

	Robot adoption (0/1 indicator)			
	— <i>Labor productivity</i> —		— <i>Output</i> —	
	(1)	(2)	(3)	(4)
<i>Productivity</i>				
Lagged 2nd quartile	0.0166** (0.00827)	0.00126 (0.00855)	0.0242*** (0.00725)	0.0161** (0.00800)
Lagged 3rd quartile	0.0464*** (0.00915)	0.0121 (0.00972)	0.0615*** (0.00862)	0.0410*** (0.0102)
Lagged 4th quartile	0.0719*** (0.00983)	0.0169 (0.0110)	0.149*** (0.0102)	0.121*** (0.0137)
Lagged skill intensity		-0.151** (0.0633)		-0.161** (0.0636)
Lagged exporter status		0.0250*** (0.00820)		0.00612 (0.00826)
Observations	7368	6879	7535	6997
R-squared	0.041	0.059	0.068	0.072
Industry-year fixed effects	Yes	Yes	Yes	Yes
Factor intensity controls	No	Yes	No	Yes
Globalization controls	No	Yes	No	Yes

Notes: The dependent variable in all columns is a 0/1 indicator variable equal to one if the firm uses robots and zero otherwise. The regressions include a full set of dummy variables indicating the firm's (quartile) position in the productivity distribution of the industry in which it is active. Capital intensity is the firm's deflated capital stock per worker. Skill intensity is the firm's share of workers with a five-year university degree (in logs). Exporter status is a dummy variable for positive exports. Factor intensity controls are a firm's capital intensity, defined as the firm's deflated capital stock per worker, and R&D intensity as the firm's deflated R&D expenditures relative to its deflated total sales (both in logs). Globalization controls are importer status, defined as a dummy variable for positive imports, and a dummy variable for foreign ownership (equal to one if the firm is foreign owned by more than 50 percent and zero otherwise). We add one to all factor intensity variables before taking logs in order to keep zero observations. All explanatory variables are lagged by one year. The sample is restricted to firms that do not use robots in the first year they appear in the sample. Moreover, we condition on the firm not using robots in the previous year (or the most recent year for which robot data are available for that firm). Robust standard errors (clustered by firm) are given in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels, respectively.

A.4 Propensity score estimates

In column (1) of Table A.5 we present univariate probit regressions where we regress the robot indicator variable on a set of lagged variables we use in the propensity score estimation. In column (2) we present the multivariate probit regression using the same variables. To construct the table we pool across all industries, while for the results shown in the paper, we estimate the propensity score by industry. All regressions include industry dummies.

Table A.5: Propensity scores estimation equation (probit specification)

	Robots multivariate (1)	Robots univariate (2)
Sales	0.284*** (0.0311)	0.304*** (0.0209)
Sales growth	-0.0145 (0.126)	0.221** (0.101)
Labor productivity	-0.114* (0.0684)	0.361*** (0.0545)
Labor productivity growth	0.0226 (0.0664)	-0.0144 (0.0458)
Capital intensity	0.127*** (0.0381)	0.320*** (0.0323)
Skill intensity	-1.806*** (0.649)	1.014** (0.459)
R&D intensity	0.165*** (0.0613)	0.359*** (0.0556)
Exporter status	0.0567 (0.0780)	0.554*** (0.0615)
Importer status	0.0331 (0.0803)	0.579*** (0.0619)
Foreign ownership status	-0.0529 (0.109)	0.475*** (0.0984)
Observations	4053	4053
Pseudo R-squared	0.157	

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.5 Details to data from the International Federation of Robotics (IFR)

Table A.6: Sector mapping IFR to ESEE

Description SEPI website	Corresponding IFR industries
1. Meat products	10-12 - Food and beverages
2. Food and tobacco	10-12 - Food and beverages
3. Beverage	10-12 - Food and beverages
4. Textiles and clothing	13-15 - Textiles
5. Leather, fur and footwear	13-15 - Textiles
6. Timber	16 - Wood and furniture
7. Paper	17-18 - Paper
8. Printing	17-18 - Paper
9. Chemicals and pharmaceuticals	19 - Pharmaceuticals, cosmetics & 20-21 - other chemical products n.e.c. & 229 - Chemical products, unspecified
10. Plastic and rubber products	22 - Rubber and plastic products (non-automotive)
11. Nonmetal mineral products	23 - Glass, ceramics, stone, mineral products
12. Basic metal products	24 - Basic metals & 289 - Metal, unspecified
13. Fabricated metal products	25 - Metal products
14. Machinery and equipment	28 - Industrial machinery
15. Computer products, electronics and optical	275 - Household/domestic appliances & 262 - Computers and peripheral equipment & 263 - Info communication equipment, domestic and prof. & 265 - Medical, precision, optical instruments & 279 - Electrical/electronics unspecified
16. Electric materials and accessories	271 - Electrical machinery n.e.c. & 260 - Electronic components/devices & 261 - Semiconductors, LCD, LED
17. Vehicles and accessories	29 - Automotive
18. Other transport equipment	30 - Other vehicles
19. Furniture	16 - Wood and furniture
20. Other manufacturing	91 - All other manufacturing branches

Notes: This table shows our mapping of industries between the official classification in the ESEE data according to the SEPI website (left column) and the official sector definition in the IFR data (right column).