

Robot Adoption, Organizational Capital, and the Productivity Paradox

Rodimiro Rodrigo*

George Washington University

December 2023

Abstract

Major technological changes have come with an adjustment period of stagnant productivity before the economy operates at its full potential. The mechanism of this adoption process is still not well understood. Combining Brazil employer-employee matched data with a novel measure of robot adoption, I document that the short-term effect of robots is on within-firm organization of labor across occupations rather than on productivity. Only when labor reorganization stabilizes -five years from adoption- productivity gains kick in. I estimate a general equilibrium model with heterogeneous firms, endogenous robot adoption, and organizational capital accumulation. The model shows that the organizational capital mechanism explains 20% of Brazil's aggregate skills demand change over the last decade.

*Email: rodimiro.rodriago@gwu.edu. I thank Toshi Mukoyama, Adriana Kugler, John Rust and Mark Huggett for their support and guidance in this project. I thank Pol Antras, Luis Baldomero, Adrien Bilal, Nick Bloom, Minsu Chang, Jasmina Chauvin, Behzad Diba, Mayara Felix, Ariel Gomez, Juan Hernandez, Henry Hyatt, Pedro Juarros, Bill Kerr, Ernesto Lopez-Cordova, Rodney Ludema, Anna Maria Mayda, Ferdinando Monte, Pascual Restrepo, Rich Reyes-Heroles, Magaly Saenz, Daniela Scur, Daniel Valderrama, Cesar Uribe, and seminar and conference participants at AOM, Boston U-TPRI, EEA, Georgetown, GWSB, Harvard, HBS, HKS, IDB, ITAM, LACEA-LAMES, LACEA-RIDGE, NASMES, NBER SI, PUC Chile, SED, U Delaware, U Los Andes, U Queensland, U Toronto, Wharton-WINDS, and the World Bank for helpful comments. I thank Ministerio do Trabalho of Brazil for granting access to RAIS.

1 Introduction

Technological advancement has long been considered a driving force behind growth and productivity enhancement. Nevertheless, previous episodes of major technological progress, such as the Second Industrial Revolution, characterized by new technology based on electricity around 1870 and the Third Industrial Revolution, marked by the rise of information and communication technologies (IT) in the 1980s, were followed by a surprisingly long delay of productivity gains (see [Greenwood and Yorukoglu \(1997\)](#)). This puzzling dynamic of productivity has been called the productivity paradox.¹

This paper revisits the question of why productivity has a slow reaction after a major technological change by studying the latest episode associated with automation, also known as “the Fourth Industrial Revolution.” I focus on the adoption of industrial robots, defined by the International Federation of Robotics (IFR) as “automatically controlled, re-programmable, and multi-purpose machines” in Brazil. I document that meaningful productivity gains are only realized five years after the adoption of robots, what I refer to as the productivity paradox of robots. These productivity gains are economically significant. A \$3 per worker increase in robot imports by a local labor market results in a 2.66% cumulative increase in productivity over the following 10-year period.²

These three industrial revolutions not only have in common a short-term slow productivity growth; they also implied a major reconfiguration of the workplace or a redesign of the production line, as understood by Henry Ford,³ in terms of the division of tasks among machines and labor, and the new distribution of workers across the tasks reserved for labor. Motivated by this observation, this paper hypothesis is that a firm’s organizational capital, defined as the “expertise to enhance its own production efficiency”⁴ is technology-specific, and hence, a major technological change will depreciate the stock accumulated by the firm. This depreciation requires that firms

¹It is also known as the Solow paradox, in reference to [Solow \(1987\)](#): “you can see the computer age everywhere but in the productivity statistics.”

²I show this result is not driven by declines in employment as documented by [Acemoglu et al. \(2014\)](#) for the ITs productivity paradox in the US, also known as the Solow paradox.

³Ford’s production line was a manufacturing process in which parts are added as a semi-finished product moves from workstation to workstation until the final product is achieved. On December 1, 1913, Henry Ford installed the first assembly line for the mass production of an entire automobile, reducing the time to build a car from 12 hours to one hour and 33 minutes.

⁴As defined by [Prescott and Visscher \(1980\)](#).

undergo a process of manager experimentation to learn the new optimal distribution of workers across tasks associated with the new technology.⁵ The process can be slow and costly for the adopting firms and the aggregate.

By combining a Brazilian-matched employer-employee dataset and granular administrative robot adoption data, I provide new establishment-level evidence on labor reorganization and the consequences for organizational capital associated with robots. In particular, the occupation switching rate of workers increases between 18% and 27% within firms during the five years after adopting robots. During this period, I observed increased worker movement from production to support activities. Using the average of their workers' occupation-specific tenure as a measure of a firm's organizational capital, I document the immediate depreciation of the firm's organizational capital and a slow re-accumulation after the new technology is adopted.

These empirical results motivate and allow me to estimate a dynamic general equilibrium model with heterogeneous firms, where firms accumulate organizational capital, imperfectly transferable across technologies, creating a trade-off between two dynamic decisions: adopting robot technology and accumulating organizational capital. Firms also face a sequence of static decisions on hiring production and support (administrative) workers. Therefore, robot adoption induces an adjustment period for firms to re-organize or accumulate organizational capital to re-optimize their production line given the new technology, thus explaining the late appearance of the productivity gains. To the best of my knowledge, this work is the first to use granular data on robot adoption and labor dynamics to estimate a model of endogenous technology adoption and organizational capital to explain the transitional path of productivity.

This paper contributes to different strands of research. First, it aligns with the expanding literature examining the impacts of industrial robots on labor markets. Prior studies have explored this impact through cross-country analysis (like [Graetz and Michaels \(2018\)](#)); specific-country labor-market analysis (as [Acemoglu and Restrepo \(2020\)](#) for the U.S.);⁶ and recent studies focusing on firm and worker levels using detailed labor data and measures of robot adoption, with signs and magnitudes varying across studies.⁷ I add to this literature by estimating the short- and long-term effects of robot adoption for firms and workers in Brazil using high-quality data, one of the few

⁵[Autor et al. \(2002\)](#) analyze the reorganization of labor in the check processing of a Bank after the introduction of digital check imaging technology. They describe the reorganization of labor as a phase of manager experimentation.

⁶These studies include [Chiacchio et al. \(2018\)](#), [Dauth et al. \(2018\)](#), [Mann and Puttmann \(2018\)](#), and [Webb \(2019\)](#).

⁷See [Bessen et al. \(2019\)](#), [Koch et al. \(2019\)](#), [Acemoglu et al. \(2020\)](#), [Humlum \(2020\)](#), and [Aghion et al. \(2020\)](#).

estimates for a developing economy.⁸

Second, this paper delve into the productivity paradox literature along two dimensions. One, I contribute by empirically characterizing the productivity paradox of robots, building on past work documenting the paradox for previous technology episodes such as [David \(1990\)](#) did for both the electricity and the information technology revolutions.⁹ Additionally, I contribute to investigating the causes behind this paradox. [Greenwood and Yorukoglu \(1997\)](#) linked the stagnation of labor-productivity growth observed after the rise in IT adoption to the high cost of learning to use the new technology, a cost that worker skills could mitigate. [Parente \(1994\)](#) and [Violante \(2002\)](#) micro-found this fact using vintage capital models, where expertise in the old technology is not fully transferable to the new one, and a period of firm-specific learning-by-doing occurs. As a result, productivity growth may fall in the short run and increase later. [Atkeson and Kehoe \(2007\)](#) build a quantitative model of technology diffusion and growth, where learning needs to be substantial and protracted to explain the productivity paradox. The authors infer the learning parameters from U.S. plants' life-cycle patterns. Nevertheless, none of these papers provide direct empirical evidence of such learning, and that their length can last for so long. Moreover, in their models, the idea of learning is governed by exogenous parameters. More recently, [Acemoglu et al. \(2014\)](#) analyze the productivity paradox in the IT sector, concluding that more direct evidence of the transformations induced by new technology is needed. I complement the existing literature using microdata to provide evidence of the workplace reorganization and organizational capital depreciation associated with adopting robots.

Third, I estimate a quantitative model that accounts for organizational capital and explains productivity and occupational dynamics, robot diffusion, and changes in aggregate skill demand seen in the data. This model allows for evaluating counterfactual policies such as tariffs or taxes and provides insights into challenges faced by managers and entrepreneurs, such as the “innovator’s dilemma.”

Finally, I contribute to the literature on organizational capital that emphasizes its role in firms’ growth over the life cycle ([Prescott and Visscher \(1980\)](#)), the slow economic recovery after a re-

⁸Previous literature analyzing developing countries, instead investigate the indirect effect of robots adopted by developed countries in the labor markets of developing countries (see [Kugler et al. \(2020\)](#)) through the trade channel. [Giuntella et al. \(2022\)](#) is an exception analyzing the local effects of robot adoption in China.

⁹Also see [Acemoglu et al. \(2014\)](#), who use event studies to document the fact more rigorously, and [Huggett and Ospina \(2000\)](#), who provide plant-level evidence from the Colombian manufacturing sector that productivity growth falls when a plant goes into the adoption of new technology embodied in new equipment.

cession (Koenders and Rogerson (2005)), and the productivity paradox of electricity (Atkeson and Kehoe (2007)). Existing papers in this area are often lack direct measures of organizational capital or rely on small-sample firm surveys to measure this intangible asset (see Bloom et al. (2010) for a literature review). This paper provides establishment-level evidence on the effects of robot adoption on worker reallocation across occupations within firms and on the evolution of organizational capital. My results are consistent with evidence from Bresnahan et al. (2002) on the relevant complementarities between technology investments and organizational investments while emphasizing that the latter take place slowly and are associated with productivity's slow responses in the short run.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 presents the macro evidence, namely, findings of the analysis at the local-labor-market level. Section 4 presents the micro evidence, namely, worker-level results regarding the organizational capital mechanism. Section 5 develops the general equilibrium model, and Section 6 develops the strategy for the identification of the model parameters. Section 7 presents the results of the quantitative estimation of the model, and Section 8 uses the model to quantify the implications of the "innovator's dilemma" and to analyze the incidence of different policies to overcome it.

2 Data

The main data source for the labor market information is *Relação Anual de Informações Sociais* (RAIS), a matched employer-employee dataset compiled annually since 1976 by the Brazilian Ministry of Labor. This paper uses the period between 2000 and 2013 to avoid the economic crisis beginning in 2014. RAIS is well known as a high-quality Census of Brazilian formal firms and workers (see Dix-Carneiro (2014)). Unique identifiers link workers to the establishment and firm where they are employed, which allows me track workers over time and across firms and occupations. The dataset provides detailed workers' information on gender, age, race, education level, monthly wage, number of hours in the contract, month of access, and separation from a job and occupation at the 5-digit level CBO-94 classification before 2002 and at the 6-digit level CBO-2002 classification compatible with ISCO after 2002. For firms and establishments, RAIS reports the location, industry at the 5-digit level in CNAE classification compatible with ISIC, and the number of employees. From the panel structure of the data, I can construct measures of occupation-firm

tenure. The data appendix contains more details on the definitions and construction of workers' characteristics. I limit the sample to include working-age individuals (18–64 years old) working full time by December 31 (more than 30 hours per week). For the worker-level analysis, I omit those working in public administration, those without valid information on industry and occupation, and those with missing or zero earnings. If multiple jobs are reported for one individual, I include only the highest-paying job each year.

This data has two main limitations. First, it provides no information on workers' activities once they leave the formal labor market. This lack of information is relevant for Brazil, where informality rates exceed 40% of employment during my sample period.¹⁰ Second, it lacks important information such as sales, value-added, and investment in technology.¹¹

Labor productivity is measured at the local-labor-market level as GDP divided by total employment at the same level of aggregation. I employ the GDP at the municipality level constructed by the Institute of Applied Economic Research (IPEA) from 1997 to 2014 based on Brazilian National Accounts from IBGE and aggregate into local labor markets. Following the literature analyzing labor dynamics in Brazil, I define local labor markets according to the Brazilian Institute of Geography and Statistics (IBGE) 557 microregions, which are defined as groups of economically integrated contiguous municipalities with similar geographic and productive characteristics, slightly larger groupings than commuting zones in the U.S.

To measure local robot adoption, I take advantage of the fact that the world's total manufacture of robots is concentrated in a few countries. In 2019, 15 countries shipped almost 90% of globally exported industrial robots.¹² Given that Brazil does not produce industrial robots domestically, every new robot adopted in the country is first imported and therefore recorded by the customs authorities with the 6-digit product code corresponding to industrial robots and compiled by the Secretariat of Foreign Trade in Brazil. The data are available at the municipality level yearly from 1997 to 2019.

¹⁰RAIS lack of information on informal labor should not be a major concern for the analysis at the local-labor-market level given that informality did not experience any change in trends over my sample period. The micro facts rely on the occupational dynamics inside firms, so we restrict our attention to firms larger than five workers, which is the case for 98% of informal firms in Brazil.

¹¹Information on sales, value-added, technology investments, and other inputs are at PIA, a manufacturing survey that can be merged with RAIS. Access to PIA was already granted for a continuation project, but it is still restricted due to COVID-19.

¹²These countries are Japan, Germany, South Korea, Singapore, Italy, France, Denmark, United States, Netherlands, Austria, Sweden, Taiwan, Belgium, United Kingdom, and China, the only developing country on the list.

In Appendix A.1, I display two data-validation exercises. First, I verify that the definition of industrial robot from the administrative trade records and the one established by the IFR are consistent at the national level. The IFR defines an industrial robot as “an automatically controlled, reprogrammable, and multipurpose [machine]” (IFR, 2014) and provides a world census of industrial robots.¹³ Appendix A.1.1 shows a strong correlation (0.94) of Brazil’s total robot installments reported between these two information sources from 2000 to 2016. I then verify the distribution of robots across municipalities within Brazil. Customs records’ geographical distribution could be biased for two reasons. First, one might be worried that import records are huddled in cities where firms’ headquarters are located and not where their plants are. Second, with the same effect, this machinery is frequently imported through domestic robot integrators that specialize in importing and installing them at local plants. To address these issues, I implement a second validation of the data focused on the automotive industry, which accounts for more than two-thirds of robot adoption in Brazil. I then identify all openings and expansions of car factories over the sample period –60 events– and verify those events coincide with spikes of imported robots in the municipalities where plants are located. Appendix A.1.2 describes in more detail the algorithm of this procedure.¹⁴

The variable of robot adoption considered in the empirical analysis, both at the local-labor-market and at the worker level, is robots per worker, defined as the value of imported robots in constant prices weighted by the local labor market employment in an old period, to avoid introducing a source of endogeneity through our measure of robot adoption. Appendix A.2 describes the evolution over time of my robots-per-worker measure for a typical local labor market. Adoption occurs in spikes, that is, large amount of adoptions in a given year preceded and followed by years of low levels of adoption. This observation motivates the implementation of event studies as my main empirical specification.

¹³The data are available at the country-year-industry level. I thank Adriana Kugler for the access to the IFR data.

¹⁴Acemoglu and Restrepo (2019), Dauth et. al. (2018), and Graetz and Michaels (2018) use data on robot adoption from IFR at the country-industry level to analyze local effects. Aghion et. al. (2020) use firm adoption of machinery and equipment, where the idea of technology change is less clear. Humlum (2019) follows the same strategy as mine for Denmark. Note that the U.S. trade data fail the second validation.

3 Macro facts: What are the long-term dynamic effects of robot adoption on labor productivity?

I first estimate the evolution of the effect of robot adoption by Brazilian local labor markets on labor productivity over a 10-year period. The baseline specification is a distributed-lag model, which includes contemporaneous robot adoption along with 10 lags as independent variables, to observe the long-run effects of the new technology, and two leads as a falsification test. Indexing local labor markets by m and years by t , the distributed-lag model is specified as follows:

$$y_{m,t} = \beta_0 + \sum_{k=-4}^{10} \beta_1^k RPW_{m,t-k} + \alpha X_{m,t} + \gamma_m + \theta_t + \epsilon_{m,t} \quad (1)$$

$$RPW_{m,t-k} = \frac{\$BRL.Imported.Robots_{m,t-k}}{L_{m,t_0}},$$

where $y_{m,t}$ is the logarithm of labor productivity at local labor market m and year t , $RPW_{m,t-k}$ is the value of imported robots (in thousand constant Brazilian reals) per worker at the same local labor market in 1997 (the oldest year available in the data), $X_{m,t}$ are covariates controlling for the proportion of high-skill workers, average age, industry and occupation diversity, and the gini coefficient as a measure of inequality, and γ_m and θ_t are micro-region and time fixed effects. This empirical strategy allows to control for time-invariant as well as local-labor-market-invariant unobservables.¹⁵

The lead-lag coefficient β_1^k is the marginal dynamic response of labor-productivity outcome ($y_{m,t}$) at time $t+k$ to an increase in robot technology adopted at time t , controlling for adoption at all other time periods. The length of the lagged polynomial is usually determined by when the effect dies and whether it passes the tests for multicollinearity between $RPW_{m,t-k}$ and $RPW_{m,t-s}$.¹⁶

Given that labor productivity is defined as the ratio of GDP over employment, observed positive effects of robots in productivity could be driven by either positive effects of robots in output or by negative effects of robots in employment. Ex ante, we cannot rule out this possibility, because previous work such as [Acemoglu and Restrepo \(2020\)](#) documents negative effects of robots

¹⁵[Schmidheiny and Sieglöck \(2021\)](#) show the equivalence between a distributed-lag model and an event study strategy.

¹⁶The multicollinearity test allows the inclusion of less than 12 periods; the effect vanishes differently for each variable but always between the 10th and 11th. Therefore I include 10 lags in all exercises.

on employment. Therefore, I look at the long-term dynamic effects of robots on both the logarithm of GDP and the logarithm of employment of the local labor markets, respectively, by running a distributed-lag model with each of these two dependent variables. An additional robustness check to my productivity measure is running the distributed-lag model with average wages as another proxy of labor productivity.

One implicit assumption of my estimates is the time-invariant treatment effects, ignoring technology improvements. Therefore, the lagged reaction of productivity and wages could be mechanically driven by new technologies being more productive than older technologies. Because year-, lagged-, and technology cohort- effects cannot be estimated simultaneously, due to multicollinearity, I address this concern by splitting my time series into before and after 2007.

3.1 Identification discussion

The identification assumption in the estimation of equation (1) can be formalized as follows:

$$E[RPW_{m,t-k} \times \epsilon_{m,t} | X_{m,t}, \gamma_m, \theta_t] = 0, \forall (t, k). \quad (2)$$

The leads estimated coefficients, specifically β_1^{-4} to β_1^{-1} , allow me to control for pre-trends. If the identification assumption holds, we would expect these four parameters to be zero. Even in the case of no pre-trends, the identification condition could be violated. Correlated demand and supply shocks may occur at the same time that the firm or plant automates the production process. For example, increased demand or increased competition could lead to increased automation with a simultaneous direct impact on labor productivity.

A major concern with this empirical strategy is that the adoption of robots in a given local labor market could be due to other trends affecting either the adopters, the non-adopters, or the overall labor market. To address this endogeneity concerns, I implement a shift-share instrumental variable (IV) design. Given that 60% of robots in Brazil are adopted in the automotive sector, I use the sectoral adoption of robots in countries were the headquarters (HQ) of the different brands manufacturing cars in Brazil are located, which are France, Germany, Italy, Japan, Korea, Sweden and the United States as the shift of my instrument. The exclusion restriction is that automation in headquarters' countries influences the automation of subsidiary plants abroad included of those in Brazil by pushing the technology frontier without directly affecting its labor force. Finally, I

interact this shift with the share of employment across the different sectors in Brazil in a previous period as follows:

$$Exp.to.robots_{m,(t-l,t-k)}^{HQ} = \sum_j \frac{L_{m,j,1997}^{BR}}{L_{m,1997}^{BR}} \times \frac{\Delta Robots_{j,(t-l,t-k)}^{HQ}}{L_{j,1997}^{HQ}} \quad (3)$$

If robot adoption is endogenous in equation (1), we have 10 endogenous variables to correct for. Unfortunately, we only know the distribution needed to make inference of the two-stage least-squares (2SLS) estimator for up to three endogenous variables, which is less than in this case. To overcome this limitation, I transform the distributed-lag model specification into a comparable static version with two endogenous variables suitable to be instrumented:

$$y_{m,t} = \beta_0 + \beta_1 \sum_{k=0}^4 RPW_{m,t-k} + \beta_2 \sum_{k=5}^{10} RPW_{m,t-k} + \alpha X_{m,t} + \gamma_m + \theta_t + \epsilon_{m,t} \quad (4)$$

where $y_{m,t}$ is the log in labor productivity, employment, and wages at time t and local labor market m , and $\sum_{k=0}^4 RPW_{m,t-k}$ is the value of the robots adopted over the last 5-year period and $\sum_{k=5}^{10} RPW_{m,t-k}$ the robots adopted from the previous five year, so β_1 would capture the short run effects of robots and β_2 would capture the long run effects of robots. While with linear algebra we can show this specification is equivalent to equation (1), the estimated β from equation (4) will be equal to the cumulative effect obtained by adding all the marginal effects estimated by equation (1), only if the effect of robots is independent of the timing of adoption, that is, if the marginal effect of a robot in $t + 1$ is the same in $t + 10$. If, as expected, the dynamic effect of robots is strictly increasing in time, the β from equation (4) will be strictly lower than $\sum_{k=0}^{10} \beta_1^k$ from (1).

3.2 Estimation results

Figure 1 reports the marginal effects of robot adoption on labor productivity over a 10-year period. Note the figure shows no signs of pre-trends: before some local labor markets adopted more robots, the labor-productivity path was comparable in both adopters and non-adopters afterwards. The results of the local-labor-market effects show labor productivity has a delayed positive reaction after the adoption of the new technology. The effect becomes statistically significant greater than zero only five years after the adoption and reaches its peak after six years. The increase in the value of robots per worker is associated with a labor-productivity gain of 12%, six

years after the adoption. I refer to this weak short-run productivity growth followed by a strong productivity growth as the *productivity paradox of robots*.

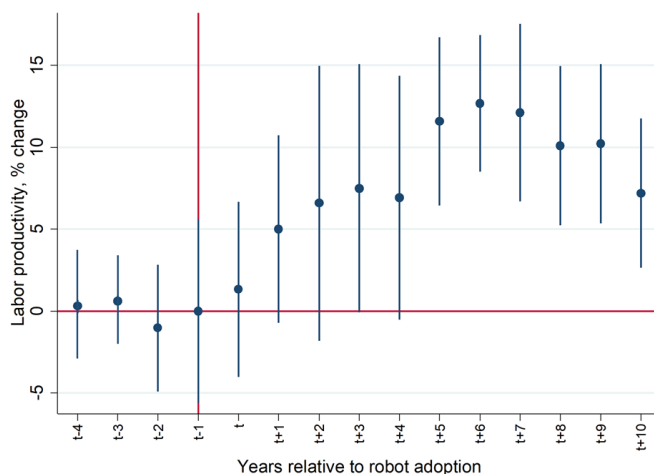
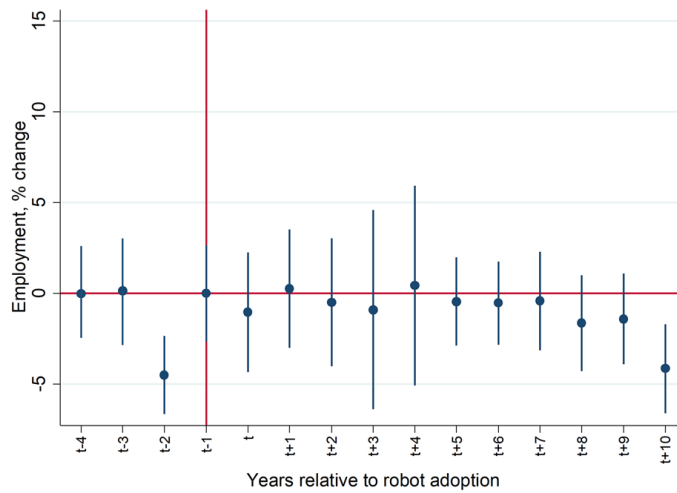


Figure 1: Robot-adoption effects on labor productivity

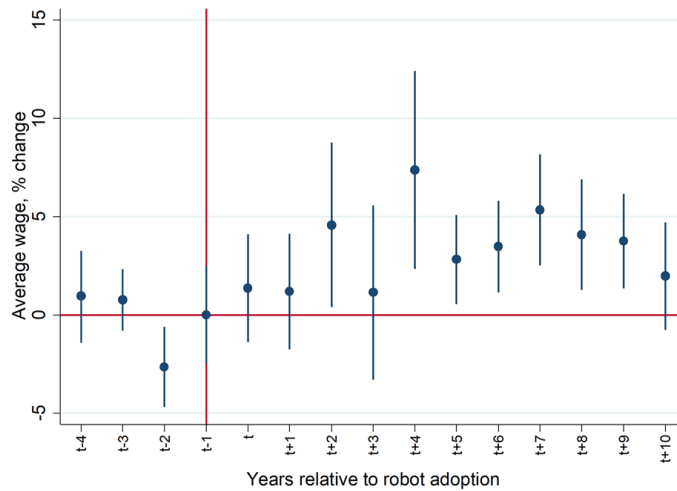
Figure 2 reports the robustness checks regarding our productivity measure. First, we check the effects of robot adoption on aggregate employment, which is the denominator of our labor-productivity measure. Figure 27(a) shows no effect of robots on aggregate employment, which confirms the productivity result we observe in Figure 1 is not mechanically driven by potentially less productive workers leaving the labor market. Figure 27(b) reports the dynamic effects of robots on average wages, which is a well-established proxy of labor productivity.

The results show the same pattern as that of labor productivity, no effects in the short term, and with positive and statistically significant effects four years after adoption. Note that although labor productivity and wages report the same pattern, wages have a lower magnitude. Wages grow more slowly than labor productivity after robots, which would lead to a decrease in labor share over time. Appendix B.1 verifies this conjecture.

To provide an illustration of the relevance of the estimated effects, I compute the 10-year cumulative effects of robot adoption reported in Table 1. To grasp the size of these effects, consider a local labor market that increases the value of imported robots by the interquartile range of the variable (\$3 per worker in 1997); the cumulative effect over a 10-year period on productivity would be a 2.66% increase. This magnitude is similar to Brazil’s labor-productivity growth, which is around 2.67% in the same period analyzed here. Average wages would have increased by 1.18%, while



(a) Robot-adoption effects on employment



(b) Robot-adoption effects on wages

Figure 2: Labor-productivity robustness checks

aggregate employment would have no effect over a 10-year period.

The evidence for Brazil suggests that in the long run, the productivity effect dominates the displacement effect of robots. By contrast, [Acemoglu and Restrepo \(2020\)](#) find small negative effects of robots on employment and negative effects on wages for the U.S. This evidence suggests in the U.S., the displacement effect dominates the productivity effect. Note they only look at the short-run effects of robots, and thus do not directly contradict my results. This result motivates

Table 1: Long-run (10-year) cumulative effects of robot adoption

	ln(Labor Prod.)	ln(Av. Wages)	ln(Empl.)
Imported robots value	0.89*** (0.19)	0.39*** (0.09)	-0.10 (0.08)
Observations	4,464	4,464	4,464
R-squared	0.34	0.40	0.43

Notes: *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Clustered standard errors are in parentheses.

the need to evaluate both short- and long-run effects for policy recommendations.

Finally, Table 2 reports the results of estimating equation (4), the cumulative version of the distributed-lag model, estimating the short and long run effects of robots. Panel A presents the results of running equation (4) by ordinary least squares (OLS) and Panel B through a two stage least squares (2SLS) procedure using the *exposure to Headquarter robots* to instrument for the *value of robots per worker* by local labor markets. As expected, the OLS results are qualitatively consistent with the cumulative effects estimated by the distributed lag model although with smaller magnitude. The results of the IV procedure show relevant first-stage and 2SLS estimates that confirm no statistically significant short-run effects and relevant long-run effects of robots on labor productivity, wages and employment.

3.3 Composition of labor across occupations

A second macro fact, which is more descriptive but informative about the mechanism that explains the productivity paradox, is the strong connection between the timing of Brazil doubling its average units of imported robots after 2007 (left panel of Figure 3) and the beginning of a change in the occupational distribution of labor (right panel of Figure 3). In particular, the number of production workers decreases and the number of workers in supporting activities increases after the same year. This fact suggests a labor reallocation across occupations induced by the arrival of new machine technology. In the next section, I provide a formal analysis to establish the relation between the adoption of a new technology and the observed change in the composition of labor across occupations.

Table 2: Long-run (10-year) IV effects of robot adoption

OLS	ln(Labor Prod.)	ln(Av. Wages)	ln(Empl.)
Δ Imported robots value _(t-1,t-5)	0.04 (0.03)	0.03* (0.02)	0.01 (0.02)
Δ Imported robots value _(t-6,t-10)	0.08*** (0.02)	0.05*** (0.01)	0.01 (0.01)
Adj R-squared	0.95	0.98	0.97
2SLS	ln(Labor Prod.)	ln(Av. Wages)	ln(Empl.)
Δ Imported robots value _(t-1,t-5)	-0.16 (0.53)	-0.06 (0.22)	-0.03 (0.20)
Δ Imported robots value _(t-6,t-10)	0.19** (0.09)	0.28*** (0.11)	0.07 (0.06)
R-squared	0.88	0.91	0.90
Observations	4,464	4,464	4,464

Notes: *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Clustered standard errors are in parentheses.

4 Micro facts: What are the effects of robots in firms' organizational capital?

Consider [Prescott and Vissscher's \(1980\)](#) definition of organizational capital as the firm's accumulated expertise to enhance its productivity, which includes the ability to match workers to tasks in the right proportion and then allowing them to master on it. The hypothesis on the effect of a technological change in firms' organizational capital is that firms' accumulated expertise is not perfectly transferable across technologies. Therefore, the new technology might induce firms' reorganization of labor across tasks, and the depreciation of firms' organizational capital or expertise developed under the old technology.

The process of reorganizing workers and gaining new expertise across tasks might be slow and costly not only for the workers, as documented by [Traiberman \(2019\)](#) after a trade shock and [Humlum \(2019\)](#) after a technology shock, but also for the firms. Previous literature proposes the

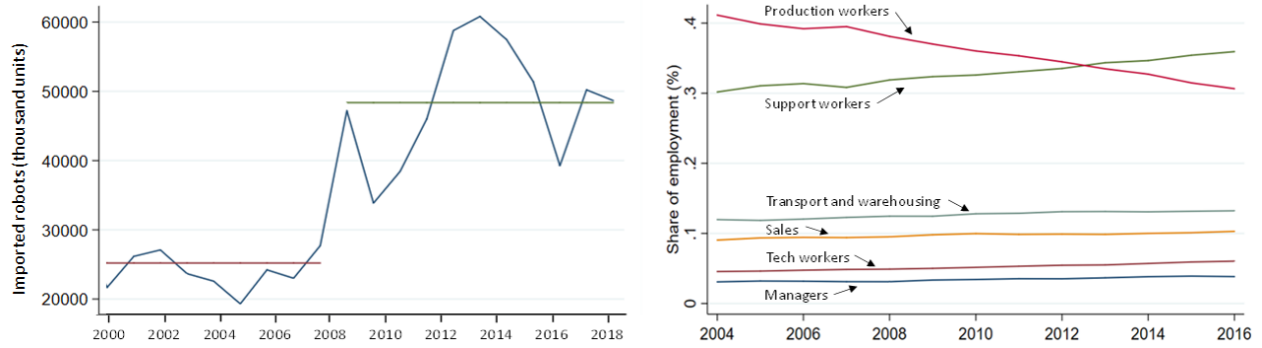


Figure 3: Composition of labor across occupations

organizational capital as the mechanism to understand firm growth over the life cycle (Prescott and Visscher (1980)), the economy’s slow recovery after a recession (Koenders-Rogerson (2005)), or the productivity paradox after the Second Industrial Revolution (Atkeson-Kehoe (2005)). To provide empirical evidence on the suggested mechanism, I exploit matched employer-employee data, which allows me to track workers over time and across occupation, firm, industry, and location, as long as workers remain in the formal sector.¹⁷

To investigate the effects of robot adoption on firms’ organizational capital, I implement two empirical exercises. First, I investigate whether the adoption of a robot induced a within-firm reallocation of workers across tasks, by running a distributed-lag model at the worker level where the dependent variable is $1\{\textit{occupation switching}_{i,f,t}\}$, which is an indicator function that takes the value of one if the worker switched occupation within the same firm. I restrict the sample to those workers who remain with the same firm over my entire sample period. I control for observable worker and firm characteristics, as well as several fixed effects including year, local labor market, and firm and individual fixed effects. I also include lags to control for pre-trends:

$$y_{i,f,j,m,t} = \beta_0 + \sum_{k=-4}^{10} \beta_1^k RPW_{m,t-k} + \alpha X_{i,f,j,m,t} + \xi_i + \kappa_f + \phi_j + \gamma_m + \theta_t + \epsilon_{i,f,j,m,t}. \quad (5)$$

Defining organizational capital according to Prescott and Visscher (1980) as the accumulated expertise of the firm, and taking advantage of the fact that the employer-employee data allow us to track workers moving across occupations even within firms, I measure it as the average of the

¹⁷An important limitation of RAIS data is that they do not include informal labor, which account for more the 40% of Brazilian labor force during my sample period.

worker's occupation-specific tenure ($tenure_{it}^{o,f}$) with a firm in a given year,¹⁸ which implies adding all the expertise of all of the workers in their current occupation, measured by their tenure, and divide this amount by the number of workers in the firm to make this measure comparable across firms of different size as follows:

$$OrgCapital_{f,m,t} = \frac{\sum_i tenure_{it}^{o,f}}{N_{ft}}$$

Then, I run a distributed lag-model at the firm level, with the organizational capital measure as the dependent variable. Controlling for firms' characteristics varying over time, and including firm, local-labor-market, and year fixed effects according to equation (5),

$$y_{f,m,t} = \beta_0 + \sum_{k=-4}^{10} \beta_1^k RPW_{m,t-k} + \alpha X_{f,m,t} + \kappa_f + \gamma_m + \theta_t + \epsilon_{f,m,t}. \quad (6)$$

It is worth noting that even after limiting the sample as described in section 2 for the worker-level analysis, the full sample is still very large (over 200 million observations). For computational tractability, I construct a random sample of 20% of Brazilian local labor markets or 119 microregions shown in Figure 13 in Appendix A.3. The main reason for sampling at this level is that the high-dimensional fixed effects used in my regressions rely on individual movements across occupations and establishments. Given the simultaneous estimation of individual and establishment level fixed effects, a sample either at the individual- or establishment-level only would reduce the statistical power of my regressions and therefore the precision of my estimates. This way I am able to analyze the effects of the technological shock, observed at the local-labor-market level, on the worker and firm dynamics over time. Appendix A.3 presents summary statistics comparing the full sample and final sample for 2004 (to be completed). Finally, I use the occupational classification developed by Bernard et al. (2017) to study workers' tasks; see Appendix A.3.1 for details.

¹⁸This definition is consistent with the notion of organizational forgiveness defined in Benkard (2000), who show the learning curve of the aircraft firms -decreasing in the number of produced aircrafts- was stopped by a strike that implied many workers laid off and the loss for the firm of the accumulated expertise by those workers.

4.1 Estimation results

The results for equation (4) regarding within-firm reorganization of labor across tasks are reported in Figure 4(a). I find the likelihood of workers switching across occupations within their same firm increases immediately after the adoption of a robot around three percentage points, and returns to the pre-robot levels that is around 12%, after five years. This result means the occupational mobility within firms increases on average 25%, when they are being reorganized around the new technology. Note the period of time during which this reorganization (5 years) takes place coincides with the period of stagnated productivity.

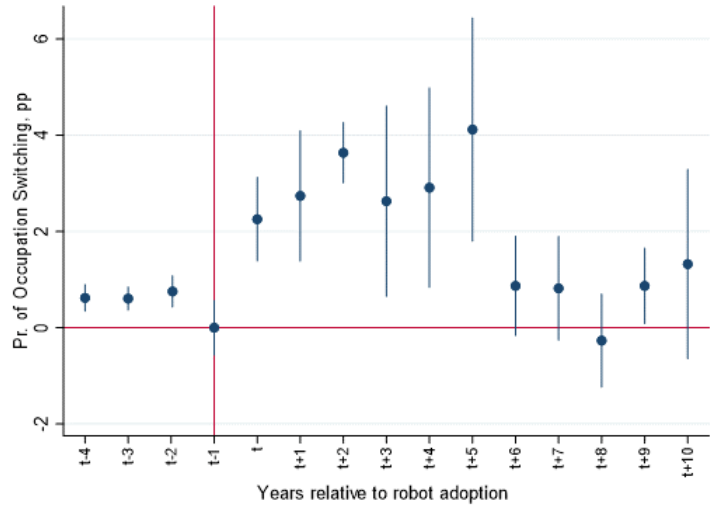
When looking at the effect of robot adoption on the likelihood of switching across specific pairs of occupation groups, the pattern observed across occupations corresponds to the switches from production to support activities, as reported in Figure 4(b). All the possible transitions across the six occupation groups are reported in Appendix B.3 in Figure 15.

Finally, the results of estimating equation (5) for the effects of robot adoption on firms' organizational capital are reported in Figure 5. I find that organizational capital suffers an immediate depreciation after the adoption of the new technology. Recovery is slow, becoming positive after six years of adoption to get back to a steady state. I use both results to inform the construction and estimation of my model presented in the next section.

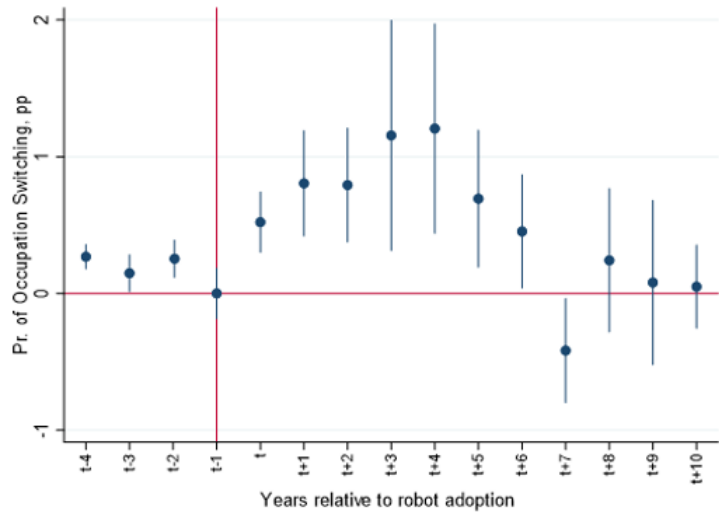
4.2 Robustness checks

4.2.1 Churning

A reasonable concern with the previous section's results, is that the observed occupational mobility was driven by churning. To address this concern, I run two versions of equation (4). In the first, the dependent variable is a dummy that takes the value of 1 if a worker leaves the firm and 0 otherwise. In the second, the dependent variable is a dummy variable that takes the value of 1 if a worker gets hired, and 0 otherwise. Appendix B.3, Figure 17, presents the results of both event studies. After the adoption of robots by a local labor market, no statistically significant increase occurs in the probability of a worker's separation. By contrast, a decrease occurs in the probability of workers in that local labor market being hired. These results not only mitigate the possibility of churning, but also strengthen the hypothesis of a re-optimization process within firms.



(a) Labor switching across occupations



(b) Labor switching from production to support activities

Figure 4: Robot-adoption and within firm occupation mobility

4.2.2 Poaching

Another threat to the evidence regarding the reorganization of labor within firms being so costly that it generates the observed delay in productivity reaction, is the possibility of firms poaching good workers. Firms adopting robots could in fact be poaching workers from previously adopted firms. Therefore, those workers with experience dealing with robots, far from arriving to the new

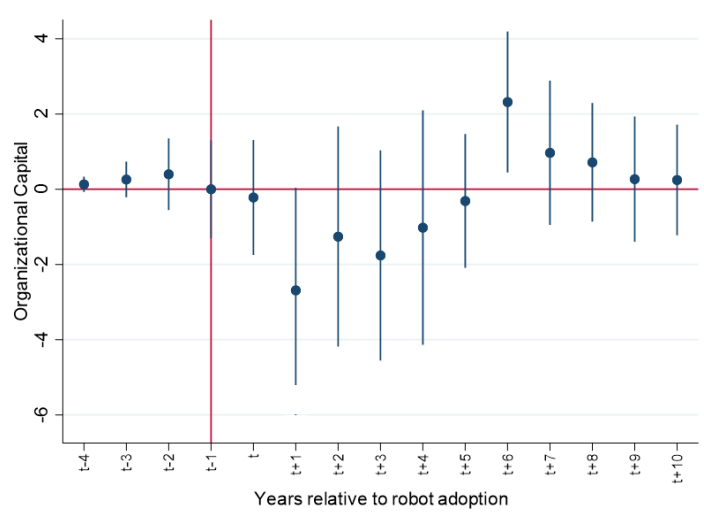


Figure 5: Robot effects on firms' organizational capital

firm to gain expertise, would be contributing to the firm's organizational capital with knowledge on best practices with the new technology and the potential learning from her peers. To investigate that possibility, I estimate the dynamic effect of robots on workers switching across firms. Specifically, I again run the distributed-lag model described by equation (4), this time with a dummy variable taking the value of 1 if a worker switched firms at a given period, and 0 otherwise. The result presented in Appendix B.3, Figure 18, shows no increase in the rate of workers moving across firms associated with the adoption of robots, supporting the hypothesis of this paper.

4.2.3 Supply-chain reorganization

Lastly, an alternative story behind the observed workplace reorganization, which could also be consistent with the lag reaction of productivity, is the reorganization of the supply chain. More precisely, the robots adopted by certain industries change their demand for inputs for the upstream industries, inducing a rearrangement of the value chain that is potentially slow and costly enough to generate the result we observe in the productivity dynamics.

To dig into this hypothesis, I look at the heterogeneous effect of robot adoption on the probability of workers' occupational switching for different industries, according to equation (4). Figure 19 in Appendix B.3 shows the results for firms in adopter industries in Brazil that, according to the International Federation of Robotics in Brazil, are Food and Beverages, Electronics, Autom-

tive, Other Vehicles, and Other Manufacturing, and firms in non-adopter industries such as Wood and Furniture, Paper and Plastics, Metal Products, Metal Machinery, and Textiles. Note the non-adopter industries represent an important component of the supply chain of most if not of all adopter industries. If the supply-chain reorganization hypothesis were true, we would expect to see robot adoption inducing occupational switching both at adopter and non-adopter firms. Whereas the adopter industries show the pattern presented in previous section, that is, an increase in workers occupational switching for the first periods after adoption, the non-adopter industries do not report any abnormal behaviour after the technology shock, leading us to reject the hypothesis.

4.3 External validity: Developing vs. developed countries

This paper presents the first estimates of the labor market consequences of robot adoption for an emerging country. Two results contrast with previous estimations for developed countries. Unlike [Acemoglu and Restrepo \(2020\)](#) for the U.S. and [Humlum \(2019\)](#) for Denmark, who both document a decrease in aggregate employment associated with robots, in Brazil, it remains unchanged both in the short and long-run. Moreover, [Humlum \(2019\)](#) also implements event studies to estimate the effect of robots in labor productivity and find it takes only three years for productivity to reach the maximum.

One possible explanation for these difference in results is the differences in the firing costs, which are known to be higher in Brazil compared with those in the U.S. and Denmark. A high firing cost would amplify the effects of the mechanism proposed here, because it forces the firm to redeploy workers within the firm instead of firing them and hiring new ones to bring the new skill required by the new technology.

I explore the role of the firing cost for the proposed mechanism by exploiting the variation in firing cost in Brazil induced by a major labor constitutional reform in Brazil in 1988. Among other inclusion, the reform implied a reduction in the maximum working hours per week from 48 to 44, an increase in the power of unions (to constitute and call for strikes), and an increase in the firing cost, namely, a fourfold increase in the penalty levied on employers for unjustified dismissal, going from 10% to 40% of the accumulated separation account. [Bosch, Goñi-Pacchioni, and Maloney \(2012\)](#) show how this national policy had differentiated effects depending on the

labor characteristics before the reform, such as the overtime of formal workers, union enrollment, and the average tenure of fired workers. I use the average tenure of fired workers in 1987, right before the labor reform, to classify local labor markets into low and high firing costs, according to whether if the average tenure was below or above the mean.

The results, reported in Appendix B.4, show the local labor markets with low firing costs experienced positive and statistically significant effects in labor productivity after the adoption earlier than local labor markets with high firing costs, where the effects of robots in labor productivity become positive and statistically significant only sixth year after robot adoption. These results are in line with the conjecture that firing costs reduces the margins of the firms to adjust to the new-skills demand given the new technology, amplifying the negative consequences for the organizational capital of the firm, which could explain the differences between the results for Brazil and Denmark.

4.4 The role of Managers

Following the same argument of this paper that recognize the tenure or accumulated expertise, firms with longer tenure managers should be able to have a smoother reorganization process compared with those firms with managers of shorter tenure.

To test this implication of our main hypothesis, I split firms based on their average manager tenure, below or above the median. Three results reported in Figure 21 on Appendix B.5, confirm this conjecture. Specifically, the results show that firms with managers tenure below the average are the ones that experienced a larger reorganization during the first years after the technological shock, while their counterpart firms with managers with more experience in their current firms had a less disruptive reorganization after the receiving the same shock.

5 Model

I build a dynamic general equilibrium model with heterogeneous firms. Every period, the firms decide the demand for production and support (or service) workers and how much to invest in organizational capital. Organizational capital will be imperfectly transferable across technologies, creating a trade-off between their two dynamic decisions: adopting robot technology and accu-

mulating organizational capital.

Firms There is a fixed mass of firms indexed by $j \in [0, 1]$. Firm j produces output y_{jt} by combining workers specialized in two tasks or occupations, namely, production activities l_{1jt} and support activities l_{2jt} , and a stock of organizational capital h_{jt} . The firm has a production function that exhibits decreasing returns to scale i.e. $\alpha \in (0, 1)$, and has elasticity of substitution across occupations σ , with θ_1 and θ_2 intensity-factor shares of each occupation's labor in total output as follows:

$$y_{jt} = F(l_{1jt}, l_{2jt}; h_{jt}, \epsilon_{jt}, R_{jt})$$

$$F(l_{1jt}, l_{2jt}; h_{jt}, \epsilon_{jt}, R_{jt}) = z_{Hjt} \left\{ (\theta_1 (h_{jt} z_{1jt} l_{1jt})^{\frac{\sigma-1}{\sigma}} + \theta_2 (z_{2jt} l_{2jt})^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \right\}^\alpha \quad (7)$$

Assumption 1: $z_{Hjt} = \exp\{\phi_H R_{jt}\} + \epsilon_{jt}$, $\epsilon_{jt} \sim \pi(\epsilon' | \epsilon) > 0$

Assumption 2: $z_{1jt} = \exp\{\phi_1 R_{jt}\}$ and $z_{2jt} = \exp\{\phi_2 R_{jt}\}$

Firms are heterogeneous with respect to an idiosyncratic productivity shocks ϵ_{jt} , which will be independently distributed across firms and will follow a Markov process within firms; endogenous organizational capital h and robot technology $R_{jt} \in \{0, 1\}$ states. I model robot technology as a binary state to reflect the fact that most robot users adopt in robots in a single year (see Figure A.2), and assume that once a firm gets a robot, it keep it forever. The parameter ϕ_H captures the effect of robot adoption on a firm's Hicks-neutral productivity z_{Hjt} , and the parameter ϕ_1 captures how the robot technology affects the productivity of workers in production activities, whereas ϕ_2 the productivity of workers in support activities. It will only be possible to identify $(\phi_1 - \phi_2)$ that captures the effect of robots on the relative productivity of production workers versus support workers.

The specification chosen in equation (7) can be micro-founded or derived from a task-based model with automation like the one proposed by [Acemoglu and Autor \(2011\)](#) or [Acemoglu and Restrepo \(2020\)](#), in which robots substitute for worker tasks in production.¹⁹

Finally, firms' organizational capital h_{jt} depreciates at the rate δ . Firms invest in organizational capital in units of the final good, according to the following quadratic function that will work as an adjustment cost:

¹⁹See [Humlum \(2019\)](#) for the derivation for a similar specification.

Assumption 3: $i_{jt} = (h_{jt+1} - h_{jt}(1 - \delta))^2$.

I model the imperfect transferability of organizational capital across technologies observed in Figure 5 as an immediate depreciation ϕ_3 on the period of adoption of the new technology, as shown below in the firms' problem. This additional depreciation, along with the implicit adjustment cost considered in the investment function of organizational capital, will drive the slow reaction of productivity over the transition after the adoption of a robot.

Skills demand. Every period, firms take the factor prices $\{w_1, w_2\}$ and the state variables $x = (h, \epsilon, R)$ as given, and choose their demand of workers for each occupation to maximize the static profit function defined as follows:

$$\pi(l_1, l_2; x) = z_H F(l_1, l_2; x) - w_1 l_1 - w_2 l_2. \quad (8)$$

where $F(l_1, l_2; x)$ is defined according to equation (7). The intra-temporal equilibrium condition derived from this problem shows the way the workplace reorganization mechanism will operate when a firm adopts the robot technology, formalized in Proposition 1:

$$\frac{l_2}{l_1} = \left(\frac{w_1 \theta_2}{w_2 \theta_1} \right)^\sigma \left(e^{R(\phi_1 - \phi_2)} \cdot h \right)^{(1-\sigma)}. \quad (9)$$

Proposition 1. *Let $\sigma \in (0, 1)$ and suppose $(\phi_1 - \phi_2) > 0$ and $h(\phi_1 - \phi_2) - \phi_3 < 0$, then robot-adopter firms will experience:*

- *A decrease in (l_2/l_1) , during the period of adoption ($\tau = t$). Therefore, they will increase the demand for production workers versus support workers.*
- *And an increase in (l_2/l_1) , for the later periods, ($\tau \geq t + 1$). Therefore, firms will reduce the demand for production workers versus support workers.*

In other words, the depreciation of organizational capital will induce in the short run an inefficient workplace reorganization or reallocation of labor across occupations. This reorganization will prevent the productivity gains from increasing during the first periods after adoption, which will be overcome as the organizational capital is re-accumulated. See Appendix C for the proof of Proposition 1.

Firm optimization. The dynamic problem of the infinitely lived firm j that at the beginning of the period has not adopt robots, that is, in the state without robots ($R = 0$), is directly integrated in the following Bellman equation:

$$V(h, \epsilon, 0) = \max \left\{ \begin{array}{l} \max_{l_1, l_2, h'} \left\{ \pi(l_1, l_2; h, \epsilon, 0) - (h' - h(1 - \delta))^2 + \beta \mathbb{E}[V(h', \epsilon', 0)|x] \right\}, \\ \max_{l_1, l_2, h'} \left\{ \pi(l_1, l_2; h, \epsilon, 1) - (h' - h(1 - \delta) - \phi_3)^2 - P_R + \beta \mathbb{E}[V(h', \epsilon', 1)|x] \right\} \end{array} \right\} \quad (10)$$

Each period, firm j (the sub-index is suppressed to simplify notation) chooses its labor demand l_1 and l_2 , as well as organizational capital accumulation h' , and whether to adopt a robot $R \in \{0, 1\}$. Adopting firms pay the price of robots P_R -exogenously determined by the international market-, and lose a fixed amount ϕ_3 of organizational capital or expertise achieved with the previous technology, both only during the period in they decided to adopt.

The Bellman equation of firm j that at the beginning of a period already possesses a robot, i.e. that already is in the absorbing state ($R = 1$), is defined as follows:

$$V(h, \epsilon, 1) = \max_{l_1, l_2, h'} \left\{ \pi(l_1, l_2, h, \epsilon, 1) - (h' - h(1 - \delta))^2 + \beta \mathbb{E}[V(h', \epsilon', 1)|x] \right\}. \quad (11)$$

Therefore, the firms that adopt the robots will experience permanently the productivity gains from z_{Hjt} governed by ϕ_H and a change in the relative productivity across occupations amplified by $(\phi_1 - \phi_2)$ biased towards supporting occupations complementing robots.

Households' optimization The representative household chooses consumption C_t and total labor supply for occupations L_{1t} and L_{2t} to maximize the following problem:

$$\max \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left\{ C_t - \gamma_1 \frac{L_{1t}^{1+\eta}}{1+\eta} - \gamma_2 \frac{L_{2t}^{1+\eta}}{1+\eta} \right\} \quad (12)$$

$$s.t. C_t = w_{1t}L_{1t} + w_{2t}L_{2t} + \pi_t.$$

where β is the discount factor, and η the Frisch elasticity of labor supply. The household owns

all the firms in the economy, and receives their profits $\pi_t = \int \pi_{jt} d\lambda_t(h, \epsilon, R)$.

5.1 Equilibrium

Definition. A dynamic general equilibrium of the economy is a path of factor prices $\{w_{1t}, w_{2t}\}_t$, and distribution of firms over three states: the organizational capital, the productivity shock, and the robot adoption decision $\{\lambda_t(h, \epsilon, R)\}_t$; firms' policy functions $\{R_{jt}, h'_{jt}, l_{1jt}, l_{2jt}\}_{j,t}$; and representative household allocations $\{C_t, L_{1t}, L_{2t}\}_t$; such that taking the price of robots $\{P_R\}_t$ as exogenously given by the international market:

1. Firms maximize their expected discounted profits by solving equations (10) and (11). 2. The representative household maximizes expected discounted utility by solving equation (12), satisfying:

$$w_{1t} = \gamma_1 L_{1t}^\eta \text{ and } w_{2t} = \gamma_2 L_{2t}^\eta.$$

3. Labor markets clear:

$$L_{1t} = \int l_{1jt} d\lambda_t(h, \epsilon, R), L_{2t} = \int l_{2jt} d\lambda_t(h, \epsilon, R).$$

4. Goods market clears:

$$C_t + \int i_{jt} d\lambda_t(h, \epsilon, R) = \int z_{Hjt} F(l_{1jt}, l_{2jt}; h, \epsilon, R) d\lambda_t(h, \epsilon, R).$$

It worth emphasizing that the main focus of the paper is to understand the transition path following the technology adoption, regarding two main elements: the productivity dynamics and the change in the skills demand. Accounting in the model for the organizational capital accumulation according to assumption (3) and its depreciation modeled by a stock loss of ϕ_3 after adoption, helps to generate the results observed in the data, with two main challenges: The computational challenge is to carry on a transition distribution determined by three state variables, the idiosyncratic productivity shock, the stock of organizational capital, and a discrete variable that will keeps track of whether the firm has adopted a robot. And the identification challenge, for which unlike previous literature invoking the abstract concept, organizational capital h dynamics, is disciplined

by its analogous occupational switching observed in the data. In the next section, I describe the work done for this two dimensions.

6 Quantitative analysis

6.1 Solution algorithm

I assume the economy is initially in a steady-state equilibrium, where the international price of robots is high, so no firms are adopting robots. This economy experiences a technology shock materialized by a reduction in the exogenously determined price of robots, inducing the adoption of the new technology by some firms. Given the new price, in the long run, the economy converges to a new steady-state equilibrium.

6.1.1 Steady-state equilibrium

I describe now the steps to compute the initial and final steady-state equilibrium. The computation method is similar to that of a Bewley-Huggett-Aiyagari model:

- After initializing parameters, guess (w_1, w_2) .
- Derive firms' labor demands $l_1(x)$ and $l_2(x)$ by maximizing the static profits in equation (7).
- Given (w_1, w_2) and $(l_1(x), l_2(x))$, iterate on the *value function* from equations (8) and (9) until convergence to find organizational capital accumulation $h'(x)$ and robot-adoption $R(x)$ rules.
- Iterate over the *distribution function* to find the invariant distribution $\lambda(h, \epsilon, R)$.
- Compute the aggregate labor demand (L_1, L_2) from the invariant distribution. Then, compute the wages (w_1, w_2) implied by equation (11) and compare them with the guess.
- Update the guess until convergence.

6.1.2 Transition dynamics

For the transition-dynamics analysis, consider a permanent change in the price of robots from P_R to \hat{P}_R , determined exogenously (Figure 22 shows the evolution of the international price of robots

from 1996 to 2015). The computation of the transition path goes as follows:

- Import the value functions and decision rules, for both the initial and terminal state, from the steady-state calculations. Also import the invariant distribution at the initial state.
- Guess the time series of $\{w_{1,t}, w_{2,t}\}_{t=0}^T$, as the linear combination of both equilibrium wages, where $t = 1$ is the period in which the unexpected change to \hat{P}_R happened, and $t = T$ is sufficiently far in future so we can safely assume that by time T , the economy is sufficiently close to the new steady state.
- Assume that at $t = T + 1$, the economy is in the steady state with \hat{P}_R . Then, use the terminal value function from the previous step on the right-hand side of the period- T Bellman equation. Similarly, solve the value functions and the decision rules for $t = 1, \dots, T$ by backward induction.
- Using the decision rule above, the economy can be simulated (again, using the density function) forward, starting from the uploaded density function at the initial steady state. Then, T_t for $t = 1, \dots, T$ can be computed. Compare the result with the guess, and modify and iterate until convergence.

6.2 Identification of model parameters

To derive the model quantitative implications according to the main observations on the transition to the equilibrium with robots, I follow different strategies for the calibration of different groups of parameters as summarized in Table 3. A first group of parameters relatively standard in the literature, including $\{\beta, \alpha, \sigma, \eta, \delta\}$, were externally calibrated. The intertemporal discount factor β is set to 0.94 according to Cooley and Prescott (1995). The decreasing-returns-to-scale parameter α is set to 0.6 considering Baus and Fernald (1997).²⁰ The elasticity of substitution across occupations σ is set to 0.5 following Humlum (2020) estimations that imply tasks are complements in firm production. The Frisch labor-supply elasticity η is set to 0.7 following Chetty et al. (2011). So far, the depreciation rate of organization capital δ is set to 0.1.

²⁰Given that Baus and Fernald (1997) show that different levels of aggregation of the data use to estimate this parameter leads to different results, I solved the model for different α values. The model is not sensitive to different values of this parameter.

Table 3: Parameters of the model

Parameter	Description	Identification	Value
α	Returns to scale	External calibration	0.6
σ	Elasticity of substitution	External calibration	0.5
β	Discount factor	External calibration	0.94
δ	Depreciation	External calibration	0.08
η	Labor supply elasticity	External calibration	0.7
θ_i	Factor shares of occupation $i = \{1,2\}$	Match labor composition 2007 - pre-robots	$\theta_1 = 1.2$ $\theta_2 = 0.8$
ϕ_H	Hicks productivity shifter	Indirect inference	0.9
ϕ_1	Relative productivity labor	Indirect inference	0.208
ϕ_2	Org Capital loss	Indirect inference	0.208

Additionally, I assign parameter values to the share of production ($\theta_1 = 1.2$) and support-activities workers ($\theta_2 = 0.8$) in order to match the pre-robot-adoption composition of labor across occupations reported in Figure 3.

The remaining four parameters $\{\phi_H, \phi_1, \phi_2, \phi_3\}$ are identified using an indirect-inference approach based on [Akcigit and Kerr \(2018\)](#), and targeting the transition dynamics induced by the adoption of the robot technology. Although we cannot identify ϕ_1 and ϕ_2 separately, we will identify $(\phi_1 - \phi_2)$, that is, the relative productivity of labor in production versus in supporting activities. To estimate ϕ_H , $(\phi_1 - \phi_2)$ and ϕ_3 , I target as moment conditions three 10-year series: (i) the marginal effects of robots on labor productivity, (ii) the diffusion of robots, and (iii) the occupation switching within firms, all over the same time period. I compute various model-implied moments over the transition path and compare them with the data-generated moments to minimize the following objective function:

$$\min \sum_{i=0}^3 \sum_{t=0}^{10} \frac{|model_{it} - data_{it}|}{1/2|model_{it}| + 1/2|data_{it}|}. \quad (13)$$

Note that these three parameters play a role only when a firm decides to adopts a robot. Therefore, I fully identify the model at the pre-robot stage, and identify these parameters by targeting the transition path.

I consider a change in the international price of robots P_R as the exogenous shock inducing

robot adoption and the associated effects implied by the model. I discipline the shock on the robot prices to generate the diffusion of robots across firms, specifically so that the share of adopters goes from 0.05 to 0.35, after the 10-year period of the adoption of robot technology.

7 Results

The estimated parameter values are for the neutral productivity gains associated to robots $\phi_H=0.9$, for the change in the relative efficiency of production vs. support labor induced by robots ($\phi_1 - \phi_2$)=0.208, and for the organizational cost of adoption $\phi_3=0.208$. The fit of the model to the data is remarkably good, as Figure 6 shows.

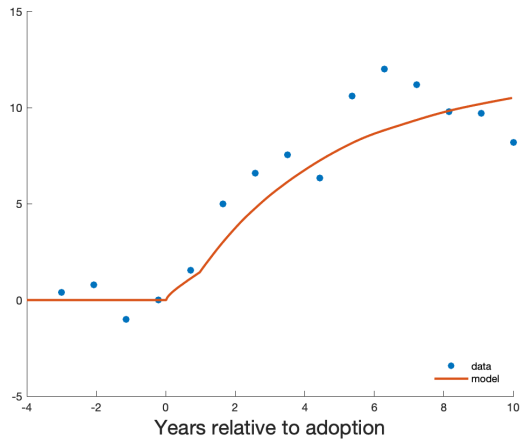
The prediction of the model for the evolution of labor productivity over the 10-years after the adoption of the new technology, always lies within the confidence interval of the effects estimated in Section 3. The diffusion of robot adoption over the 10-year period is followed almost exactly in every year over the transition path. Finally, the occupation switching at the firm level in the model goes up immediately as in the data, but by between the third and the six year begins to decrease gradually, whereas in the data, the abnormal switching stays at the same level until the fifth year, and by the sixth year, it returns to the pre-robots level.

Finally, even when the aggregate composition of labor across occupations time series was an untargeted moment, the model predicts 20% of the decrease in production workers and an increase in support workers observed in the data. Figure 7 compares the data and results from the calibrated model.

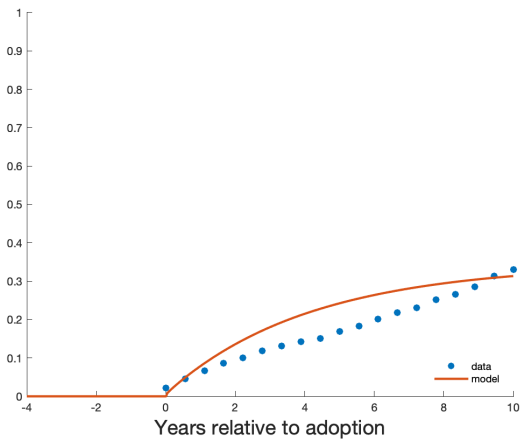
8 Policy Counterfactuals

8.1 Tariffs or taxes to robots

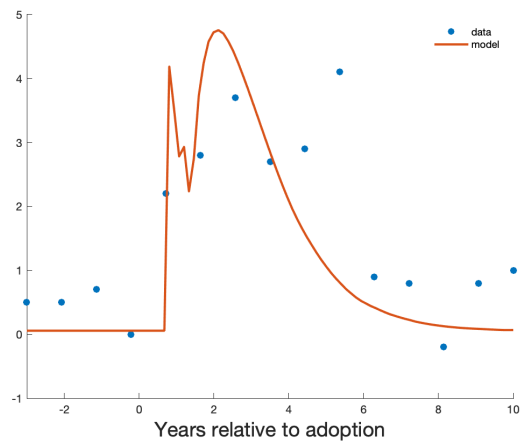
One of the main debates around robots, was started by Bill Gates in 2017 when he suggested robots should be taxed –in the same way labor is– as a way to at least temporarily slow the spread of automation and to fund other types of employment. At the same time there are arguments that, in particular for developing countries like Brazil, suggest the use of subsidies to technology to boost the adoption and close the technology gap with developed countries. Nevertheless, it remains to answer what would be the consequences for aggregate productive of this kind of measures.



(a) The productivity paradox



(b) Robot diffusion



(c) Occupational switching

Figure 6: Transition path model fit
30

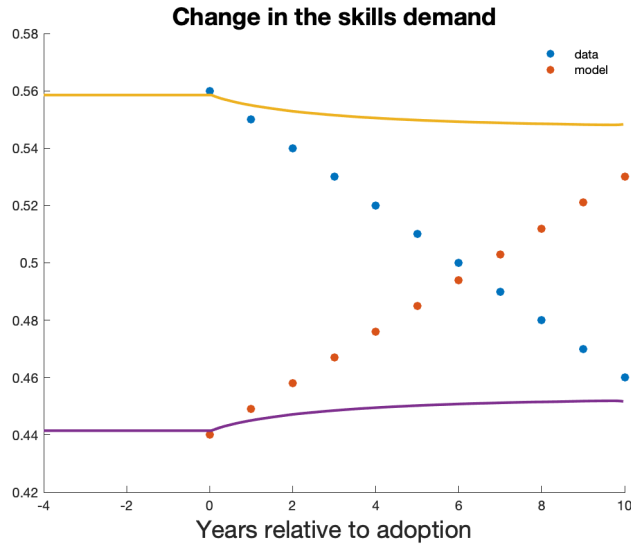


Figure 7: Change in the skills demand

The model presented in this paper allow us to provide one more argument to this debate. In Brazil, the tariffs to importing industrial robots remained fixed in 14% since 2002 and through 2022 in Brazil, practically all the period of analysis of this paper.

A straightforward policy experiment to try is removing this tariff. For that, we consider our estimated model and trajectories as the result of the current prices and tariffs. I then reduce the price of robots in 12.18%, the equivalent to removing the 14% tariff in the first place. As a result, the model predicts that in a decade the increase in productivity would have been 1.2% larger.

8.2 Innovator's dilemma

To illustrate the role of ϕ_3 , the parameter that captures the organizational cost of adopting a new technology, this section analyzes the implications of the model for the innovator's dilemma. Named by [Christensen \(1997\)](#) famous book, the dilemma arises when managers, trying to do their best to satisfy customers' current needs and minimize the likely short-run increase in waste and the slowdown of production associated with changes in business practices, end up taking sub-optimal decisions, such as avoiding the adoption of new technologies that might lead to large productivity gains for the company in the long-run.

In terms of the model, suppose managers are pessimistic regarding their expected value of the organizational cost for the firm to adopt robot technology. Specifically, suppose that when

the managers solve the optimization problem of the firm described by equations (8) and (9), she believes ϕ_3 is twice as large as it actually is. The manager will take the decisions on robot adoption R' for $\hat{\phi}_3 = 2\phi_3^*$.

The result of this exercise reported in Figure 8 shows the manager's misperception leads to about 40% decrease of the number of firms adopting robots and a decrease, similar in magnitude, in the productivity gains associated with the new technology.

In Appendix C.3, I perform two possible interventions to compensate for this manager's over-estimation of the organizational cost. The first intervention, reported in Figure 25, considers a 20% subsidy toward the price of robots, which leads to a 25% recovery versus the case with the conservative $\hat{\phi}_3$, but still 15% below the case with the true ϕ_3^* . The second intervention, reported in Figure 26, is a case in which the government without any financial incentive can achieve the previous improvement, consist in allowing non-adopter firms to learn from the experience of early adopters.

For the second case, in the baseline quantitative exercise I allow firms to learn the organizational cost parameter from early adopters by assuming a learning function ϕ_3^* , whose initial value is $\hat{\phi}_3$ but as the share of adopters goes to 1, it goes to the true value ϕ_3 . This second exercise emphasizes the role of sharing information among firms about the true cost and benefits of new technologies, and is consistent with the formation of clusters of firms of the same industry where they share expertise and best practices with very successful outcomes.

9 Conclusion

This paper documents the productivity paradox of robots by estimating the dynamic long-term effects of robot adoption on productivity for Brazilian local labor markets, and find that productivity gains come five years after adoption.

To investigate the mechanism behind this empirical regularity, I exploit rich employer-employee data and document the effects of robot adoption in firms' organizational capital. The findings provide new establishment-level evidence of (i) within-firm labor reorganization (mainly of workers moving from production to support activities) that last for the same amount of time productivity gains take to arise and (ii) an immediate drop in organizational capital associated a major technology change, followed by a slow recovery. To my knowledge, this finding is the first micro evidence

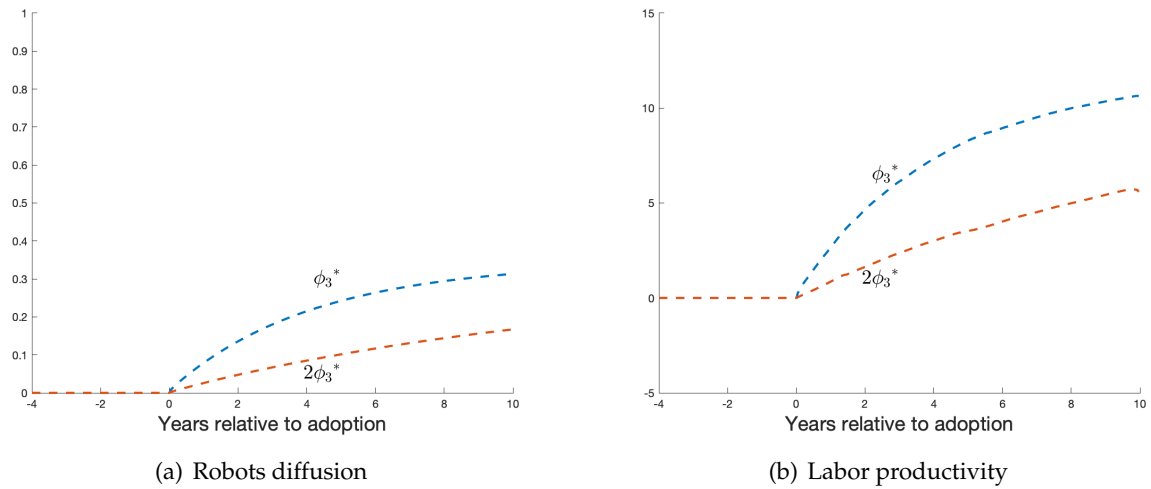


Figure 8: The innovator’s dilemma: Changes in organizational cost ϕ_3

of this powerful mechanism since it was defined in [Prescott and Visscher \(1980\)](#). My results go along with evidence in [Bresnahan et al. \(2002\)](#) on the relevant complementarities between technology investments and organizational investments, while emphasizing that the latter take place slowly and are associated with the productivity’s slow responses in the short run.

Finally, this paper’s dynamic general equilibrium model with endogenous robot adoption and organizational capital, fits well three main facts: the diffusion of robot technology, the slow increase in labor productivity, and the within-firm redeployment of workers across occupations. The model also reproduces consistently, the untargeted structural change in labor composition across occupations observed over the 10 years after robot adoption. The model emphasizes the role of organizational capital accumulation for productivity. I provide an example of the utility of the model built, by using it to understand a common manager or entrepreneur problem such as the *innovator’s dilemma*.

References

- Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*. 4, 1043-1171.
- Akcigit, U. and Kerr, W. (2018). Growth through Heterogeneous Innovations. *Journal of Political Economy*. 126(4) 1374-1443.
- Acemoglu, D., Autor, D., Dorn, D., Hanson, G., Price, B. (2014). Return of the Solow Paradox? IT, Productivity, and Employment in the US Manufacturing. *American Economic Review: Papers & Proceedings*. 104(5), 394-399.
- Acemoglu, D. and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*. 128(6), 2188-2244.
- Aghion, P., Antonin, C., Bunel, S., Jaravel, X. (2020). What are the Labor and Product Market Effects of Automation? New Evidence from France. *Working Paper*.
- Autor, D., Levy, F., Murnane, R. (2002). Upstairs, downstairs: computers and skills on two floors of a large bank. *Industrial and Labor Relations Review*. 55(3), 432-447.
- Atkeson, A. and Kehoe, P. (2007). Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions. *American Economics Review*. 97(1), 64-88.
- Basu, S. and Fernald, J. (1997). Returns to Scale in U.S. Production: Estimates and Implications. *Journal of Political Economy*. 105(2), 249-283.
- Benkard, L. (2000). Learning and Forgetting: The Dynamics of Aircraft Production. *American Economic Review*. 90(4), 1034-1054.
- Bernard, A., Smeets, V. and Warzynski, F. (2017). Rethinking Deindustrialization. *Economic Policy*. 32(89), 5-38.
- Bessen, J., Goos, M., Salomons, A. and Van Den Berge, W. (2019), Automatic Reaction: What Happens to Workers at Firms that Automate. *CPB Discussion Paper*.
- Bloom, N., Sadun, R., Van Reenen, J. (2010). Recent Advances in the Empirics of Organizational Economics. *Annual Review of Economics*. 2, 105-137.
- Bresnahan, T., Brynjolfsson, E., Hitt, L. (2002), Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *Quarterly Journal of Economics*. 117(1), 339-376.
- Chetty, R., Guren, A., Manoli, D., and Weber A. (2011), Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins, *American Economic Review*, 101(3), 471-75.
- Christensen, C. (1997), The Innovator's Dilemma: When New Technologies Cause Great Firms

to Fail. *Harvard Business School Press*.

Cooley, T. and Prescott, E. (1995). *Economic Growth and Business Cycles*. Frontiers of business cycle research. *Princeton University Press*.

David, P. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review*, 80(2), 355-361.

Dix-Carneiro, R. (2014). Trade Liberalization and Labor Market Dynamics. *Econometrica*. 82(3), 825–885.

Kugler, A., Kugler, M., Ripani, L., Rodrigo, R. (2020). U.S. Robots and their Impacts in the Tropics: Evidence from Colombian Labor Markets. *NBER Working Paper No. 28034*.

Giuntella, O., Lu, Y., and Wang, T. (2022). How Do Workers and Households Adjust to Robots? Evidence from China, *National Bureau of Economic Research*, 30707.

Graetz, G. and Michaels, G. (2018). Robots at Work, *Review of Economics and Statistics*, 100(5), 753-768.

Greenwood, J. and Yorukoglu, M. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy*. 46(1), 49-95.

Huggett, M. and Ospina, S. (2000). Does productivity growth fall after the adoption of new technology? *Journal of Monetary Economics*. 48, 173–195.

Humlum, A. (2019). Robot Adoption and Labor Market Dynamics. *Working paper*.

Koenders, K. and Rogerson, R. (2005). Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries. *Federal Reserve Bank of St. Louis Review*. 87(4), 555-79.

Mukoyama, T. (2006), Rosenberg's "learning by using" and technology diffusion, *Journal of Economic Behavior & Organization*, 61, 123-144.

Parente, S. (1994). Technology Adoption, Learning by Doing, and Economic Growth. *Journal of Economic Theory*. 63, 346-369.

Prescott and Visscher 1980. Organization Capital. *Journal of Political Economy*. 88(3), 446-461.

Solow, R. (1987). We'd better watch out. *New York Times Book Review*.

Schmidheiny, K. and Siegloch S. (2022). On Event Studies and Distributed-lags in Two-way Fixed Effects Models: Identification, Equivalence, and Generalization. *Journal of Applied Econometrics*. 1-19.

Traiberman, S. (2019). Occupations and Import Competition: Evidence from Denmark. *American Economic Review*. 109(12), 4260-4301.

Violante, G. (2002). Technological Acceleration, Skill Transferability and the Rise in Residual Inequality. *Quarterly Journal of Economics*. 117(1), 297-338.

A Data

A.1 Robots-data validation

A.1.1 Time-series validation: HS and IFR data comparison

To validate the customs records at the country level, I compare its time series with that of the IFR with a census of robots. Figure 9 shows that the total units of robots imported in Brazil is strongly correlated with the reported number of installments by the IFR from 2000 to 2016.

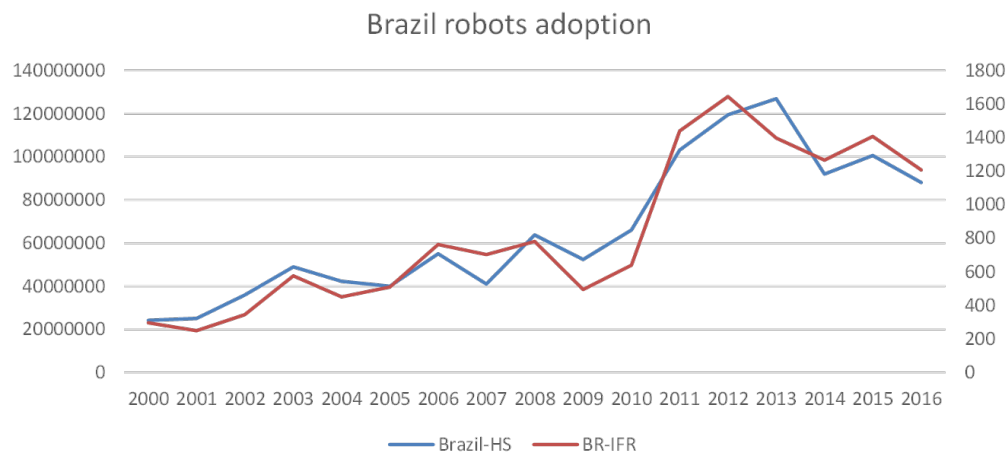


Figure 9: Robot adoption over time: Trade vs. IFR data

A.1.2 Geographical distribution validation: Car industry

The automotive industry in Brazil accounts for almost 70% of robots adopted in Brazil. To validate the distribution of robots across Brazilian municipalities, I conduct the following steps:

1. I build a database with information from Anfavea, the National Association of Vehicle Manufacturers in Brazil, identifying all the openings, closings, and expansions of car plants in Brazil.
2. I describe the time series of robot adoption for each municipality. Figure 10 shows an example with the local labor market Minas Gerais, where there are four plants in different municipalities. Additionally, I verify that the spikes of robots adopted in that municipality coincides with an opening or an expansion of a car plant.

Figure 11 shows in blue the intensity of my measure of robot adoption for Southeast Brazil,

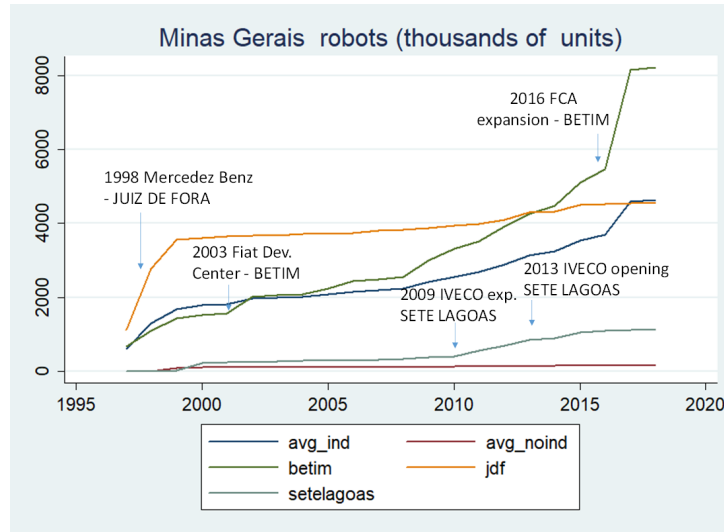


Figure 10: Robots in the automotive industry: Minas Gerais example

the most industrialized region of the country. The red dots are car plants operating in 1997 and 2017, respectively. By inspection, we can verify a correspondence between robot adoption and car plants opening or expanding.

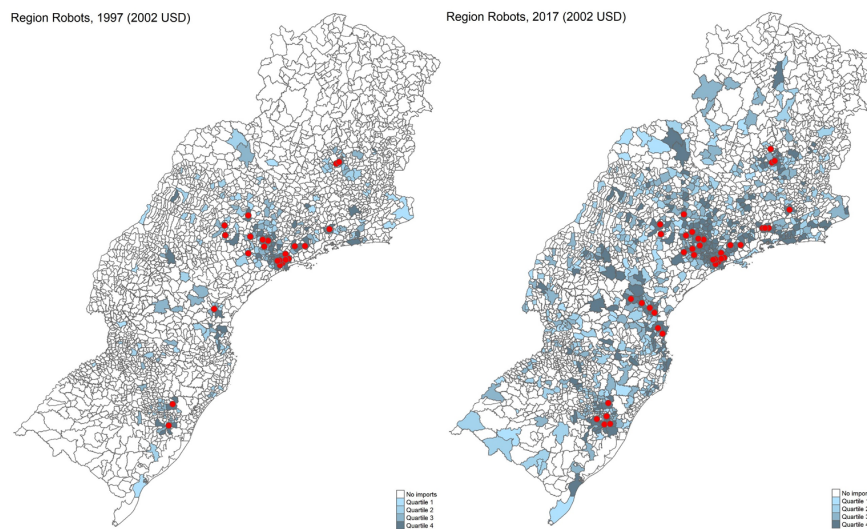


Figure 11: Robot adoption across municipalities: 1997 vs. 2017

A.2 Descriptive statistics of robot adoption

I document that robot adoption at the local labor market level occurs in spikes. Figure 12 shows the time series of imported robots of a typical local labor market. Note the level of adoption in years before and after the spike is much smaller than the adoption the year of the spike. Moreover, as in the example, 92% of the time, the adoption during the year of the spike is more than twice the mean adoption over the period of analysis (and 89% more than three times the mean).

This observation is consistent with the literature on lumpy investment and previous empirical work on robot adoption (Bessen et al. (2019), Humlum (2019)) and motivates the event-study empirical strategy used in the empirical section of this paper, as well as the model assumption of the discrete choice of robots.

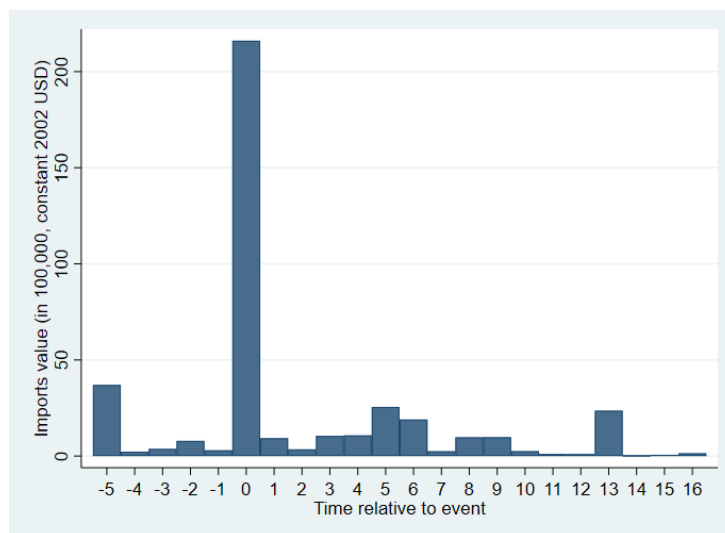


Figure 12: Time relative to largest robot adoption spike

I build the measure of value of imported robots per worker as described by equation (1), using the value of industrial robots in constant reals, imported by a local-labor-market, and divide it by the total employment in the same location in the oldest year in my sample, 1997. Table 4 shows the distribution of the constructed variable. Note that even when the median of the value of imported robots per worker is 6.582, the third quartile is 31.720, that is pulled by super adopters locations. Finally, note that by the end of the analyzed period, 290 of the 558 local-labor-markets in Brazil had never adopted robots, which guarantees that we will always have a control group²¹.

²¹This has been stressed by the new econometric literature analyzing diff-in-diff exercises with multiple time periods.

Table 4: Value of Robots per Worker

First Quartile	1.140
Second Quartile	6.582
Third Quartile	31.720
Interquartile Range	30.580
Adopter Labor Markets by 2014	268
Never Adopters by 2014	290
Total Labor Markets	558

A.3 Employment data, RAIS

The sample is limited to working-age individuals ages 18–64, who reported to work more than 30 hours per week on December 31 of each year, excluding around 10 percent of the observations. Workers with multiple jobs represent around 5% percent of the workers, in which case, I include only the highest-paying job.

Besides this filters, the full sample has above 200 million observations. Therefore, for computational tractability, I construct a random sample of 20% of local labor markets or microregions.

Equation (5) includes controls for the following: gender; age computed from the year of birth and age squared; education level, introduced as a dummy for workers with at least a high school diploma, who account for more than 60% of Brazil’s formal workers; race, classifying workers as white (48%), brown (27%), and others (24%); and firm size, classified as small (40%; less than 50 workers), medium (37% 51 - 250 workers), and large firms (17%; more than 250 workers).

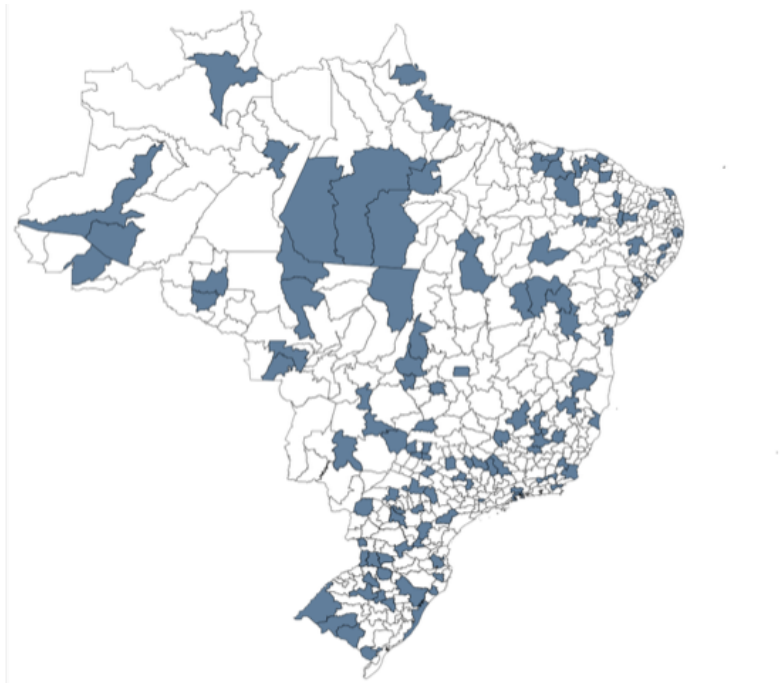


Figure 13: Random sample of microregions for panel analysis

A.3.1 Occupational Classification

To analyze workers' switches across occupations, and to avoid identifying renamings or reclassification as switches, I group the 6-digit CBO-2002 occupation codes available at RAIS into the six categories proposed by [Bernard et. al. \(2017\)](#): managers, R&D/tech, production, support, sales, transport and warehousing, and production workers.

Figure 3 shows that production and support workers are not only those occupation groups with a larger share of workers, but are also the ones who face most of the dynamics observed in the second part of the 2000, when robot adoption became more prevalent.

Table 5: Macro outcome variables

ln(Productivity)	3.839 (0.378)
ln(Av. Wage)	6.185 (0.325)
ln(Empl.)	9.906 (1.366)
ln(Labor Share)	2.870 (0.543)
Value of Robots per Worker	22.688 (121.324)
No. Observations	558

B Stylized Facts

B.1 Effect of robots in labor share

The effects of robot adoption on labor share are estimated by running equation (1) with this dependent variable. The results, reported in Figure 14, are consistent with the literature on labor market polarization as well as with the results from Figures 1 and 27(b) that show that even when they have the same patterns, labor-productivity growth increases more than wages.

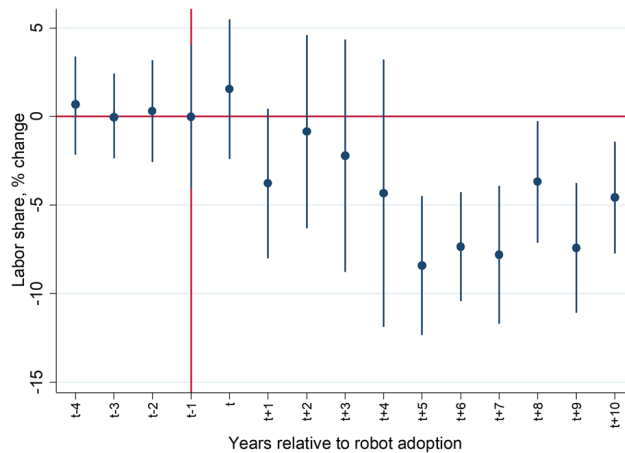


Figure 14: Robot-adoption effects on labor share

B.2 Occupation switching

Table 6 contain the mean and standard deviation of the occupation switching in 2000, as a reference year before the larger adoption is observed in Brazil according with Figure (9), that is observed at the worker level. Note that the occupation switching of workers within firm is on average of a-

Table 6: Micro outcome variables

Workers occupation switching	0.119 (0.325)	
Firms organizational capital		2.391 (3.069)
No. Observations	35,305,038	4,785,634

round 11%. The same table reports the mean and standard deviation of the organizational capital measure at the firm level, which are 2.39 and 3.07 respectively.

Figure 15 reports the result from estimating equation (4) with directed occupational switching moving out of or into production activities. Figure 16 is the analogous report but moving out of or into support activities. The relevant dynamics are mainly associated with workers who are moving out of the two occupations likely being substituted by robots, such as those workers in *production* and *transport* (or logistics in production) activities, in both cases when workers move towards supporting activities.

B.3 Robustness checks

Figure 17 presents the results of two sanity checks regarding the possibility of the labor mobility documented in section 4, to be driven by churning. The graph on the left of Figure 17 presents the estimated dynamic effects of robots on the probability of a separation. The graph on the right Figure 17 presents the results for the probability of a worker being hired. Both support the proposed mechanism.

Figure 18 presents the results of checking for the possibility that the labor mobility induced by robots was generated by good workers moving across firms injecting their experience in the firm of arrival, which would go against labor mobility being a cost from the firm perspective. The figure shows no effects of robots on labor mobility across firms, but a slight reduction in years five and six after the adoption, also supporting the proposed mechanism.

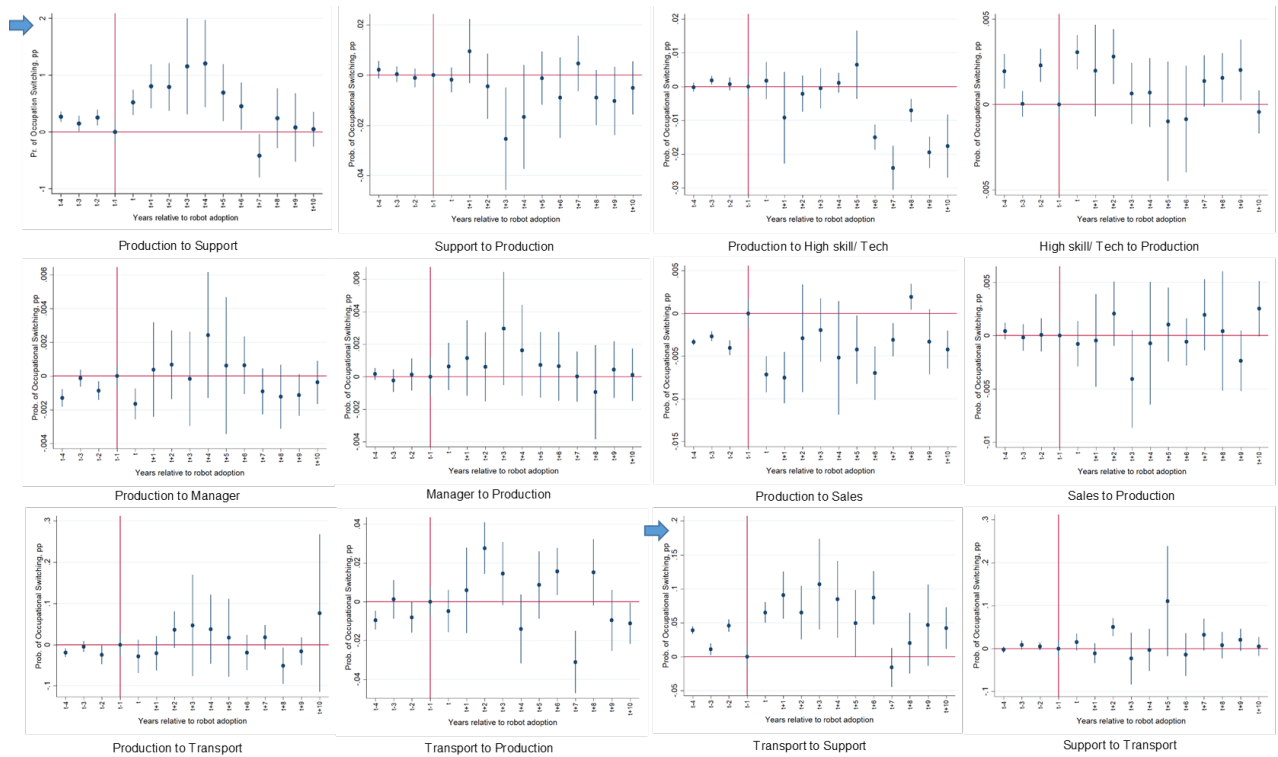


Figure 15: Robot-adoption effects on occupation switching from or to production

Figure 19 shows the effects of robot adoption on workers' occupational switching for two groups of manufacturing industries, adopter (Food and Beverages, Electronics, Automotive, Other vehicles, and Other manufacturing) and non-adopter (Wood and Furniture, Paper and Plastics, Metal Products, Metal Machinery, and Textiles) industries according to the IFR.

B.4 Firing cost

Figure 20 presents the results of running equation (1) for local labor markets with low firing costs (on the left) and high firing costs (on the right), defined as detailed in the main text of the paper.

B.5 The role of managers

Figure 21 presents the results of running equation (5) for firms with low or high average managers tenure.

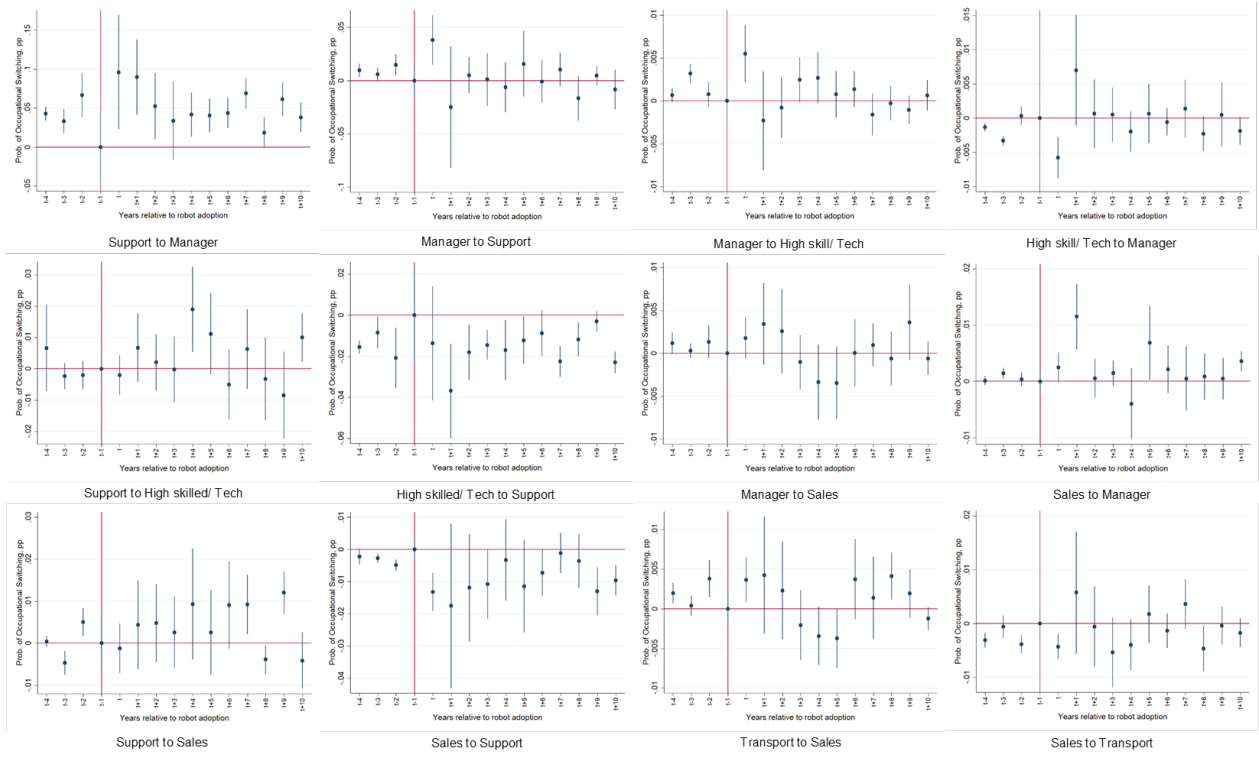


Figure 16: Robot-adoption effects on occupation switching from or to support activities

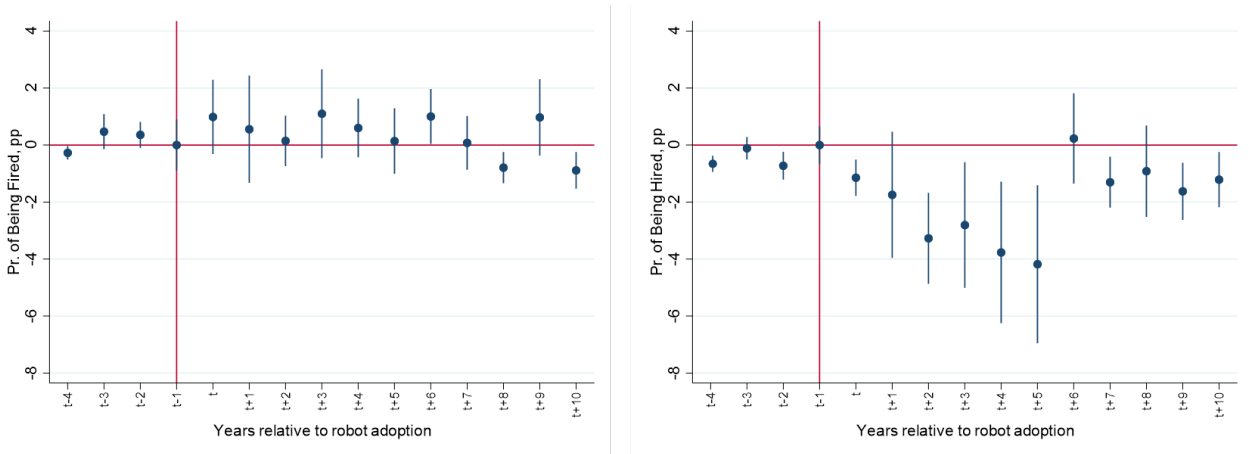


Figure 17: Robot-adoption effects on separations and firings

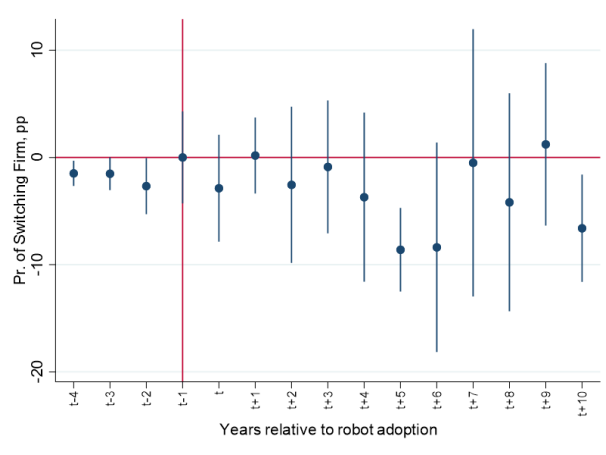
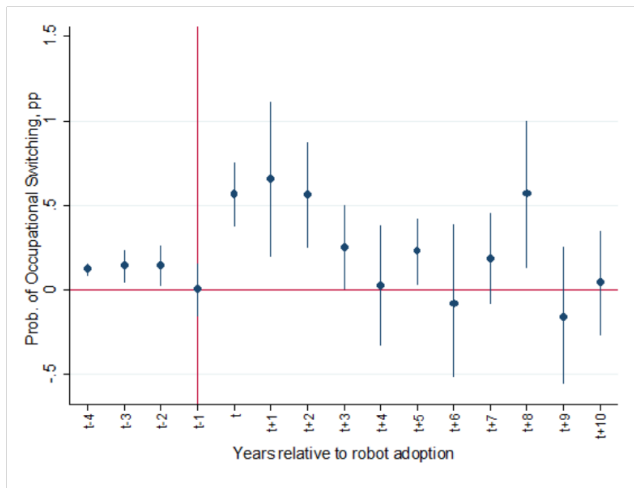
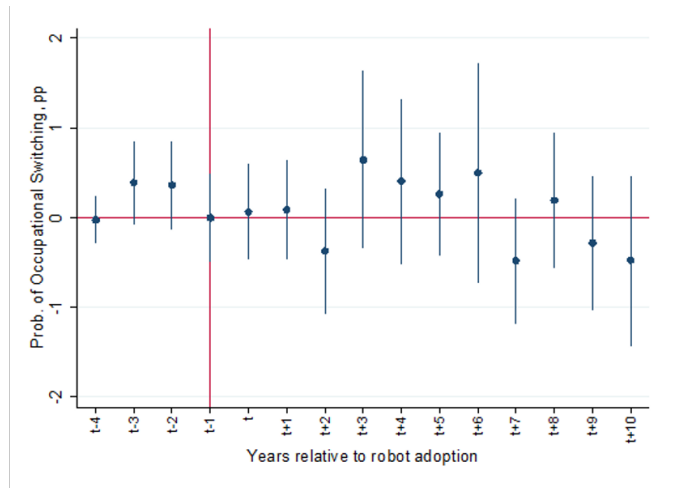


Figure 18: Robot-adoption effects on firm switching



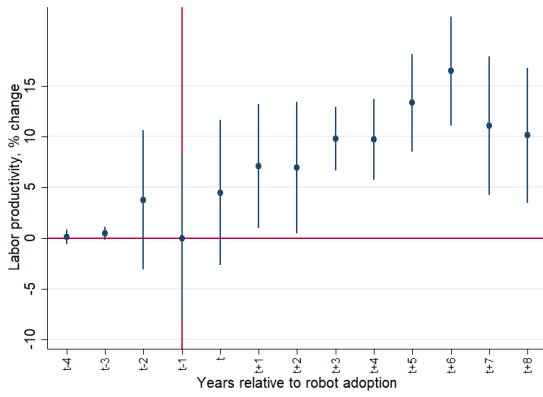
Manufacturing adopter sectors



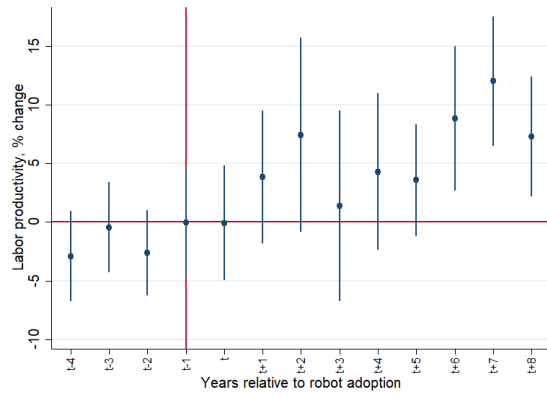
Manufacturing non-adopter sectors

*Based on the IFR: Food and Beverages, Electronics, Automotive, Other vehicles, Other manufacturings

Figure 19: Robot-adoption effects on occupational switching, by industry

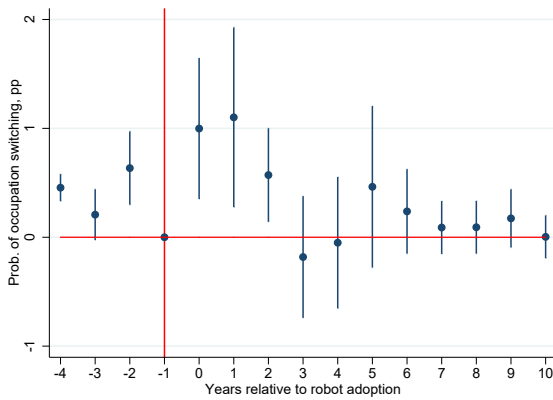


(a) Low firing cost

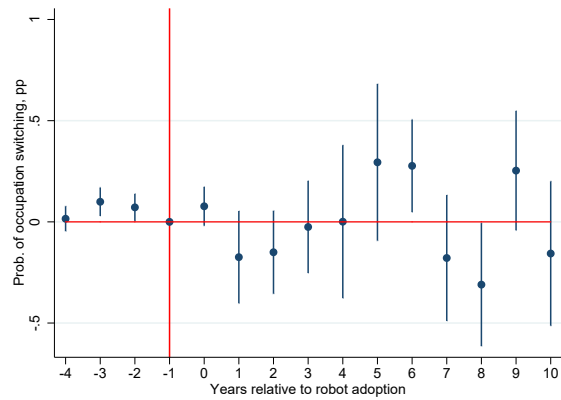


(b) High firing cost

Figure 20: Labor productivity for low and high firing cost



(a) Low manager tenure



(b) High manager tenure

Figure 21: Occupation switching for firms by manager tenure

C Model

Proposition 1. Let $\sigma \in (0, 1)$ and suppose $(\phi_1 - \phi_2) > 0$ and $h(\phi_1 - \phi_2) - \phi_3 < 0$, then robot adopter firms:

- In the period of adoption ($\tau = t$), (l_2/l_1) decreases. Therefore, firms will increase the demand for production workers versus support workers.
- For later periods, ($\tau \geq t + 1$), (l_2/l_1) increases. Therefore, firms will reduce the demand for production workers and increase the demand for support workers.

Proof. Consider the firms' intra-temporal equilibrium condition described in equation (9). The right-hand side of that expression is the marginal rate of transformation (MRT) for the production function (7). To prove the first bullet, take the derivative of the MRT with respect to robots in the short run, when the $(\partial h/\partial R)$ is given by the organizational cost of adoption ϕ_3 , as follows:

$$\frac{\partial MRT}{\partial R} = \left(\frac{w_1 \theta_2}{w_2 \theta_1} \right)^\sigma \left(e^{R(\phi_1 - \phi_2)} \right)^{(1-\sigma)} h^{-\sigma} (1 - \sigma) [h(\phi_1 - \phi_2) - \phi_3].$$

The first terms are positive; therefore, the sign of the derivative will be determined by the sign of the term in the brackets. Therefore,

$$\text{if } [h(\phi_1 - \phi_2) - \phi_3] < 0 \Rightarrow \frac{\partial MRT}{\partial R} < 0.$$

For the equilibrium condition (9) to continue holding, firms' will have to reduce l_2 and/or increase l_1 .

For the second bullet, we rely on the derivative of the MRT with respect to robots in the long run, when $(\partial h/\partial R = 0)$, leading to the following expression:

$$\frac{\partial MRT}{\partial R} = \left(\frac{w_1 \theta_2}{w_2 \theta_1} \right)^\sigma \left(e^{R(\phi_1 - \phi_2)} \right)^{(1-\sigma)} h^{-\sigma} (1 - \sigma) [h(\phi_1 - \phi_2)]$$

Therefore, as long as $[h(\phi_1 - \phi_2)] > 0$, it will be the case that $\partial MRT/\partial R > 0$. In this case, the equilibrium condition (9) will push the firms to reduce the demand for production workers (l_1) and increase the demand for support workers (l_2).

□

C.1 Transition dynamics

As described in section 6.1.2, both in the data and in the model, the transition dynamics are induced by a permanent change in the price of robots from P_R to \hat{P}_R , determined exogenously. Figure 22 shows the evolution of the international price of robots from 1996 to 2015.

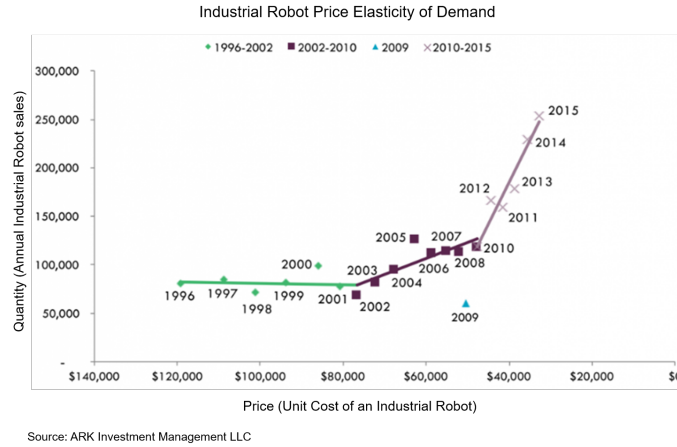


Figure 22: Price of industrial robot

Figure 23 shows the transition distribution after a change in price in $t=0$. Recall that the distribution of firms will have three state variables: productivity ϵ , stock of organizational capital, and whether the firm has a robot. Three observations (pending to formalize in the main body) are relevant for our policy experiments:

1. Adoption occurs gradually over time depending on ϵ , h , and R .
2. Adopter firms accumulate more organizational capital than non-adopters; a threshold of organizational capital separate these two groups;
3. The overall levels of organizational capital are lower for low-productivity firms than for productive ones.

C.2 Calibration

The model is fully calibrated at the initial steady state, and parameters ϕ_1 , ϕ_2 & ϕ_3 are calibrated to minimize the difference between the model and data moments displayed in the following table. I discipline the exogenous change in price in the model (observed in Figure 11), by the observed geographical diffusion of robot adoption, shown in Figure 12.

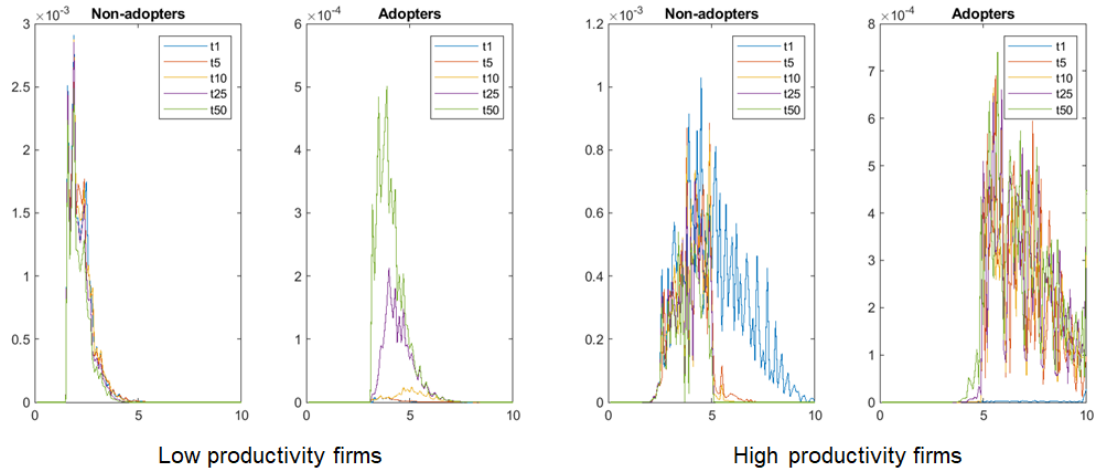


Figure 23: Industrial Robot Price

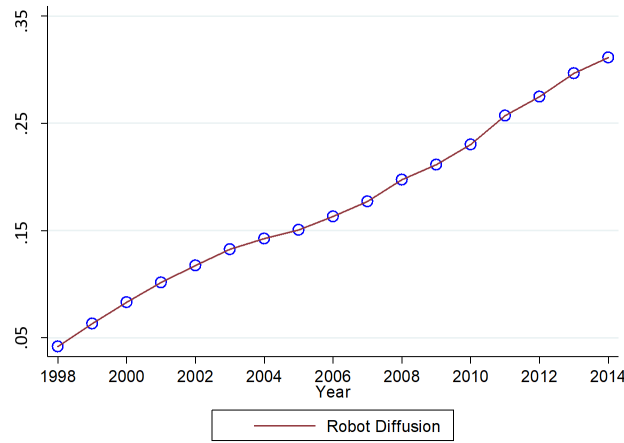
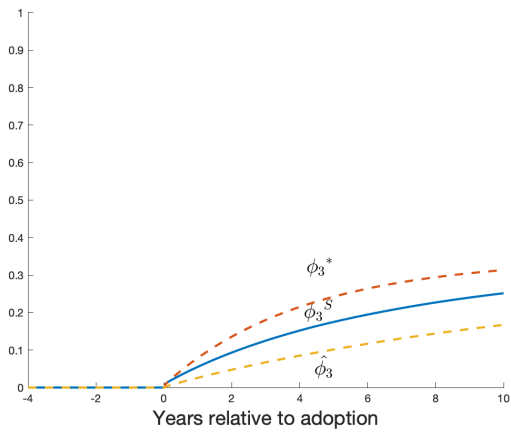


Figure 24: Share of adopter municipalities by local labor market

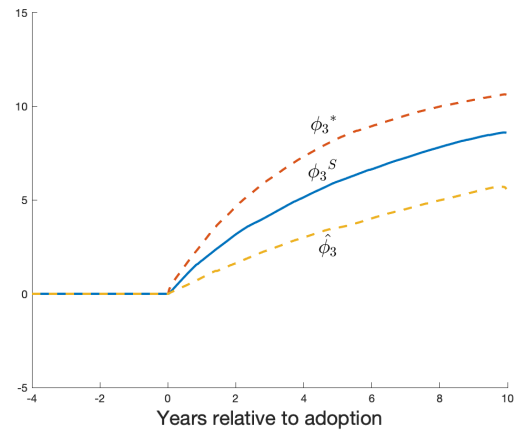
C.3 Counterfactual experiment: the innovator's dilemma

This section of the Appendix presents the results of a counterfactual experiment regarding the organizational cost of adoption ϕ_3 , described in section 8.1, and what is known as the innovator's dilemma. I perform two possible interventions to compensate for the manager overestimation's of the organizational cost:

1. A 20% subsidy toward P_R that leads to a recovery of 25% caused by considering $\hat{\phi}_3$ but still 15% below the case with the true ϕ_3^* .
2. A case in which the government can achieve the previous improvement, by disseminating



(a) Robots diffusion

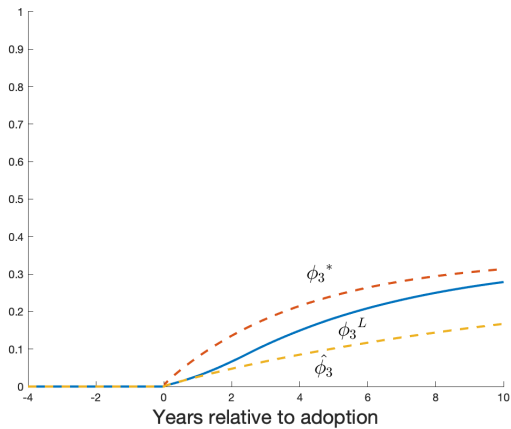


(b) Labor productivity

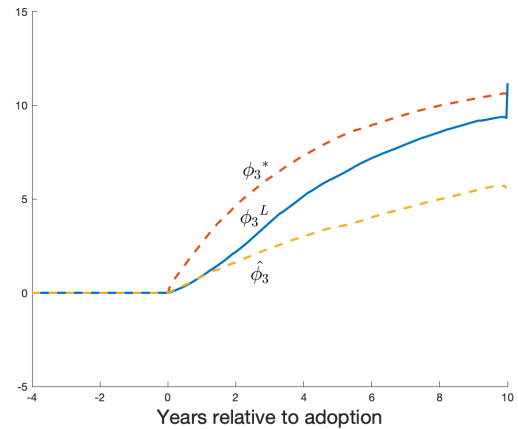
Figure 25: Effect of a subsidy on diffusion and productivity

the experience of adopter firms among non-adopters. I assume a learning function, in the spirit of Mukoyama (2005), defined to be an exponential and that initially will take values close to $\hat{\phi}_3$, but as the share of adopters goes to 1, they converge to the true ϕ_3^* , and produce a more pronounced S-shape diffusion curve.

$$\phi_3 = \phi_3^* \times \exp(\xi(1 - Shareadopters))$$

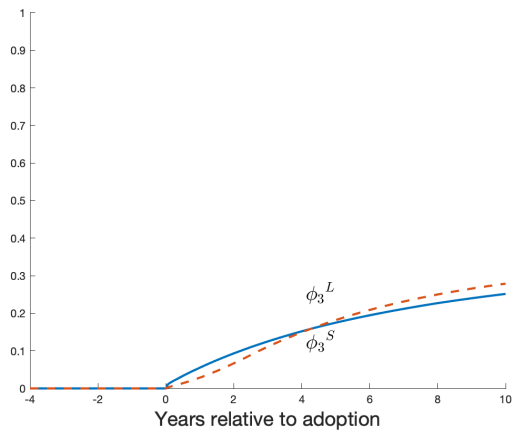


(a) Robots diffusion

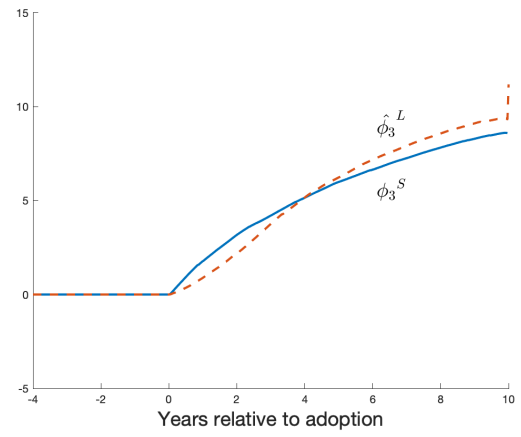


(b) Labor productivity

Figure 26: Effect of learning on diffusion and productivity



(a) Robots diffusion



(b) Labor productivity

Figure 27: Comparison of subsidy vs. learning