

Robot Adoption and Labor Market Dynamics

Anders Humlum*

Princeton University

Job Market Paper

[[Link to latest version](#)]

November 14, 2019

Abstract

I use administrative data that link workers, firms, and robots in Denmark to study the distributional impact of industrial robots. I structurally estimate a dynamic model of the firm that rationalizes how firms select into and reorganize production around robot adoption. Using event studies, I find that firms expand output, lay off production workers, and hire tech workers when they adopt industrial robots. I embed the firm model into a dynamic general equilibrium framework that takes into account the ability of workers to reallocate across occupations in response to robots. To this end, I develop a fixed-point algorithm for solving the general equilibrium that features two-sided (firm and worker) heterogeneity and dynamics. I find that industrial robots have increased average real wages by 0.8 percent but have lowered real wages of production workers employed in manufacturing by 6 percent. Welfare losses from robots are concentrated on old production workers, as younger workers benefit from the option value of switching into tech and other occupations whose premiums rise as robots diffuse in the economy. Industrial robots can account for a quarter of the fall in the employment share of production workers and 8 percent of the rise in the employment share of tech workers since 1990. I use the estimated general equilibrium model to evaluate the dynamic incidence of a robot tax.

*E-mail: humlum@princeton.edu. I am extremely grateful to my advisor Stephen Redding for his guidance in this project, and to my committee Bo Honoré and Alex Mas for invaluable suggestions and encouragement. I am indebted to Jan De Loecker, Gene Grossman, Henrik Kleven, and Eduardo Morales for insightful comments. I thank University of Copenhagen and Jakob R. Munch for providing data access, and I acknowledge financial support from the International Economics Section at Princeton University.

1 Introduction

The arrival of industrial robots in modern manufacturing is one of the most salient technological changes in recent decades. Defined as “automatically controlled, reprogrammable, multipurpose manipulators programmable in three or more axes” (ISO 8373), industrial robots were developed for car assembly in the 1990s but have since diffused widely in manufacturing. Today, robot adopters account for half of manufacturing sales, and adoption rates are accelerating. The potential labor displacing effects of industrial robots have received much public attention, culminating when the European Parliament voted in 2017 on a proposal to tax the use of robotics (Delvaux, 2016).

This paper asks who gains and who loses when industrial robots are adopted. To answer this question, I use administrative data that link workers, firms, and robots in Denmark. My first contribution is to combine event studies with a structural model that rationalizes how firms select into and reorganize production around robot adoption. I find that firms expand output by 20 percent but shrink their wage bill on production workers, such as assemblers and welders, by 20 percent when they adopt industrial robots. Firms’ total wage bill increases 8 percent as labor demand shifts toward tech workers, such as skilled technicians, engineers, and researchers. I structurally estimate a dynamic model of the firm that matches these reduced-form responses to robot adoption, the observed size premium in the selection of firms into robot adoption, as well as the S-shape in robot diffusion over time.

To understand the macroeconomic implications of robot adoption, I embed the firm model into a general equilibrium framework that endogenizes the dynamic option for workers to reallocate across occupations. The estimated general equilibrium model captures several indirect effects of industrial robots that are not identified in micro-level diff-in-diff designs. These indirect effects include the extent to which the expansion of robot adopters crowds out non-adopter firms in product and labor markets, as well as the ability of workers to reallocate across occupations in response to equilibrium wage pressures from robot diffusion.

Using the general equilibrium model, I estimate that industrial robots have increased average real wages by 0.8 percent, but with substantial distributional consequences. At the opposite ends of the spectrum, I find that production workers employed in manufacturing have lost 6 percent in real wages, while tech workers have gained 2.3 percent. I find that welfare

losses from robots are concentrated on old production workers. Younger workers, with less specific skills and a long career ahead of them, benefit from the option value of switching into tech and other occupations whose premiums rise as robots diffuse in the economy.

Occupational reallocation in response to industrial robots can account for 25 percent of the fall in the employment share of production workers and 8 percent of the rise in the employment share of tech workers in Denmark since 1990. The adoption of industrial robots have thus been a driver of employment polarization (Autor and Dorn, 2013; Goos et al., 2014). Without these labor supply responses, I find that the real wage loss of production workers from robots would have been five times larger.

My findings highlight the importance of allowing for labor supply responses when evaluating the distributional impact of industrial robots. I use a dynamic occupational choice model that represents the state of the art for studying labor market dynamics in response to trade liberalizations (Dix-Carneiro, 2014; Traiberman, 2019), and I estimate the barriers to occupational switching using observed worker transitions together with a conditional choice probability (CCP) estimator that controls for the unobserved continuation values of workers.

As a final counterfactual exercise, I evaluate the dynamic incidence of a robot tax. The undistorted equilibrium of the model is efficient (except for markups in product markets), but I use the estimated model to quantify the distributional implications of a robot tax and to evaluate its impact on aggregate economic activity. I find that a temporary robot tax can be an effective way to slow the diffusion of industrial robots. However, compared to a permanent tax of similar magnitude, a temporary tax creates larger welfare losses per dollar of revenue collected and a larger fraction of its deadweight burden falls on workers. These larger adoption elasticities and relative efficiency losses reflect the forward-looking nature of adoption whereby firms foresee that the temporary tax will expire and postpone adoption until then. Based on the estimated responses, I conclude that a robot tax is an ineffective and costly way to redistribute income to production workers in manufacturing.

Evaluating the counterfactuals above requires solving the firm and worker problems jointly, and I develop a fixed-point algorithm for solving the dynamic general equilibrium of this class of models. A key property of the general equilibrium model is that the firm and worker problems are separable conditional on the path of wages. This separable structure is highly useful in estimation and in simulation. First, it allows me to estimate the firm (worker) model with-

out specifying the problem of the worker (firm) by simply conditioning on the observed path of wages. Second, it breaks the curse of dimensionality wherein firm variables become states for the worker, and worker variables become states for the firm. The separable structure enables me to incorporate the rich firm and worker heterogeneity estimated in the micro data, and still be able to compute the general equilibrium featuring joint firm and worker dynamics.

To measure robot adoption at the firm level, I leverage the fact that almost all industrial robots used in Denmark are not produced in the country. In particular, once an imported robot crosses the country border, it is recorded by the customs authorities under the 6-digit product code *847950 Industrial Robots*. The customs records, which contain information on the timing and value of firm robot imports, offer a unique opportunity to study what happens when firms adopt industrial robots. I supplement the customs records with a representative robot adoption survey conducted by Statistics Denmark, and I validate that these micro data sources on robot adoption align with industry-level measures used in the prior literature (Acemoglu and Restrepo, 2019b). By merging the firm robot adoptions to the Danish matched employer-employee data, I obtain a dataset with unusually rich information on both firms and workers that is ideally suited to studying the distributional impacts of industrial robots.

This paper is related to and builds on several literatures. The most immediately related work is a recent series of papers that have collected reduced-form evidence on how industrial robots affect firm performance and labor market outcomes (Acemoglu and Restrepo, 2019b; Bessen et al., 2019; Graetz and Michaels, 2018; Koch et al., 2019). I complement this work with two key structural contributions. First, I estimate a model of firm robot adoption that allows me to interpret the new reduced-form evidence in terms of structural primitives. Second, I embed the model into a general equilibrium framework, enabling me to extend the identified micro-level effects to quantify the macroeconomic impacts of industrial robots. The two-sided nature of the general equilibrium model allows me to connect evidence on firm (e.g., Koch et al. (2019)) and worker outcomes (e.g., Dauth et al. (2018)) of robotization.

The methodology developed in this paper builds heavily on the literature of dynamic discrete choice models. The robot adoption model draws on the Rust (1987) optimal stopping model, and the labor supply module follows closely a series of structural labor papers, including Dix-Carneiro (2014) and Traiberman (2019). In the structural estimation, I build on the work by Doraszelski and Jaumandreu (2018) on estimating production functions with

endogenous technical change, and I apply the methods of Arcidiacono and Miller (2011) on conditional choice probability (CCP) estimation of dynamic discrete choice models.

The remainder of the paper is structured as follows. Section 2 describes the Danish data and collects stylized facts on firm robot adoption. Sections 3 and 4 develop and estimate a partial equilibrium model of firm robot adoption. Section 5 estimates the labor supply module. Section 6 unites the firm and worker blocks, and then uses the general equilibrium model to estimate the distributional impact of industrial robots and to evaluate the incidence of a robot tax. Section 7 concludes.

2 Data

I use register data that link workers, firms, and robots in the Danish economy from 1995 to 2015. The dataset is the product of merging the Danish matched employer-employee data with two new micro data sources on firm robot adoption. This linked dataset contains unusually rich information on both firms and workers, making it ideally suited to studying the distributional impacts of industrial robots. The matched worker-firm-robot panel data offer a unique opportunity to study what happens, at the micro level, when industrial robots are adopted. The data contain detailed occupational codes of workers, allowing me to study how firms substitute between labor tasks in production. The universal coverage of the Danish data is essential for estimating the general equilibrium environment that robot adopters operate in.

The firm data come from the Firm Statistics (FirmStat) Register, which covers the universe of private-sector firms from 1995 to 2016. FirmStat associates each firm with a unique identifier, and provides annual data on many of the firm's activities, such as sales, number of full-time employees, and industry affiliation. The data on workers and establishments come from the Integrated Database for Labor Market Research (IDA), which covers the entire Danish population. IDA associates each person with her unique identifier, and provides annual data on many individual characteristics such as income, hours, hourly wage, detailed occupation, education, and other sociodemographics. To match the firm and worker data, I draw on the Firm-Integrated Database for Labor Market Research (FIDA), which links every firm in FirmStat with every worker in IDA who is employed by that firm in week 48 of each year. To study worker tasks, I build on the occupational classification developed by Bernard et al.

(2017); see Appendix A.6 for details.

I use two new and complementary micro data sources to measure robot adoption at the firm level. I first use a robot adoption firm survey conducted by Statistics Denmark in 2018. The survey asked a representative sample of Danish firms if they use industrial robots in production. Appendix A.1 provides details on the survey which had a response rate of 97 percent. Second, I leverage the fact that industrial robots are highly tradable goods to measure robot adoption from firm customs records. Almost all robots in the world are manufactured in Japan, South Korea, or Germany, and once such an imported robot crosses the country border, it is recorded by the customs authorities according to a 6-digit product code where one of the codes identifies “847950 Industrial Robots”. Acemoglu and Restrepo (2018a) show that a country’s imports of industrial robots correlate strongly with its total robot installments reported by the International Federation of Robotics (IFR). Appendix A.4 calculates that the share of imports in total robot investments in Denmark averaged 95 percent between 1993 and 2015. The Danish customs records are organized in the Foreign Trade Statistics Register (UHDI).

The main challenge in using the customs records is that a substantial share of machinery is imported through domestic distributors. In the case of industrial robots, there is an industry of robot integrators that specialize in importing robots and installing them at local production facilities. Appendix A.2 describes the robot supply chain and develops a procedure for identifying robot imports done by final adopters. I validate the sample selection procedure against the firm robot adoption survey in 2018 as well as a complete list of robot integrators and producers in Denmark. In total, I identify 454 robot adoption events through direct imports.

The existing literature on industrial robots has mostly relied on an industry-level dataset compiled by the International Federation of Robotics (IFR) (Acemoglu and Restrepo, 2019b; Dauth et al., 2018; Graetz and Michaels, 2018). Appendix A.3 shows that the micro data used in this paper align well with the IFR statistics both across industries and over time.

The customs records allow me to directly study what happens when firms adopt robots. However, when quantifying the aggregate effects of robots and for parts of the structural estimation, I also want to include the adoptions done through domestic distributors. Appendix A.5 describes how I supplement the customs records with the robot adoption survey (VITA) and the IFR statistics to measure robot adoptions that are sourced domestically.

The customs records (UHDI) and the robot adoption survey (VITA) use the same firm identifier as FirmStat and FIDA, allowing me to construct a matched employee-employer-robot dataset covering the Danish economy.

2.1 Stylized Facts on Firm Robot Adoption

In this section, I present two stylized facts that will inform the modeling choices in Section 3. The first fact concerns the observed lumpiness of firm robot expenditures, which motivates modeling robot adoption as a one-off decision. The second fact documents the non-random selection of firms into robot adoption, which informs the specification of a selection model for firm robot adoption.

Fact 1. Robot Adoption Is Lumpy

Table 1 reports summary statistics for the robot adoptions identified from firm customs records in Appendix A.2. The take-away from the table is that robot adoption is lumpy. Out of the sample adopters, 70.6 percent invest in a single year only, and the peak year of investment accounts on average for 90.7 percent of total firm robot expenditures. Adopting firms purchase robot machinery for an average of \$311,000. This discrete nature of robot adoption motivates the choice in Section 3 to model robot adoption as a discrete choice problem.

Table 1: Firm Robot Investments

Adoptions (count)	454
Share of adopters with investments in one year only (percent)	70.6
Share of robot expenditures in max year (percent)	90.7
Robot machinery expenditures (\$1,000)	311

Fact 2. Larger Firms Select into Robot Adoption

Table 2 shows firm outcomes for the robot adopters in the year prior to adoption. Column 2 (“Industry”) reports average outcomes for non-adopters within the same two-digit industry-year cells as the robot adopters. Robot adopters are different from non-adopters along several dimensions, but the key feature that sets robot adopters apart is that they are substantially larger. The model in Section 3 rationalizes the selection into robot adoption by combining firm

heterogeneity with fixed costs of adoption, such that it is the firms with the largest expected efficiency gains from industrial robots that will choose to adopt the technology.

Table 2: Firm Outcomes in the Year Before Robot Adoption

	<u>Adopters</u>	Industry	<u>Matches</u>	P-value A-M
log Sales	18.28 (0.07)	16.35 (0.07)	18.19 (0.07)	0.37
log Wage Bill	16.93 (0.07)	15.15 (0.07)	16.89 (0.07)	0.66
log Employment	4.06 (0.06)	2.4 (0.06)	4.02 (0.06)	0.66
Wage bill shares (percent)				
– Managers	12.5 (0.5)	9.1 (0.7)	11.0 (0.4)	0.02
– Tech	16.0 (0.9)	6.9 (0.6)	14.3 (0.8)	0.14
– Sales	12.2 (0.4)	10.5 (0.6)	12.5 (0.5)	0.64
– Support	7.5 (0.4)	4.9 (0.5)	7.8 (0.5)	0.69
– Transportation/warehousing	5.9 (0.5)	3.6 (0.5)	6.8 (0.5)	0.23
– Line workers (mostly production)	39.9 (1.1)	47 (1.4)	40.7 (1.0)	0.61
Joint orthogonality (F test)				0.25
Observations	454	454	454	908

Note: “Joint orthogonality (F test)” represents a test of the joint hypothesis that all coefficients equal zero when the adopter indicator is regressed on the nine outcome variables in Table 2. Column 1 (Adopters) shows mean outcomes for robot adopters in the year before adoption. Column 2 (Industry) shows averages for randomly chosen non-adopters within the same industry-year cell as the adopters (one-to-one). Column 3 (Matches) shows averages for match firms within the same industry-year cell. These matches each have the minimum distance to an adopter with respect to log sales and production wage bill share (levels and two-year changes); see Appendix A.7.1 for details. Column 4 (P-value A-M) shows p-values for the null hypotheses that Adopters (column 1) and Matches (column 3) have the same population mean.

Once I match on firm sales and line worker wage bill shares in column 3 (“Matches”), the adopters look similar to the match firms on employment, wages, and wage bill shares across occupations. An F test of the joint hypothesis that none of the covariates in Table 2 predict robot adoption has a p-value of 0.25. Put differently, I cannot reject that robot adopters and matches indeed are observationally identical before adoption. The fact that adopters and match firms are balanced on these non-targeted outcomes provides supportive evidence for a model assumption in Section 3 that robot adoption is driven by an adoption cost shock once selection based on observable firm heterogeneity is taken into account.

3 A Model of Firm Robot Adoption

In this section, I develop a partial-equilibrium model of a manufacturing firm's decision to adopt industrial robots. A firm in the model faces a dynamic choice of whether to adopt the robot technology and a sequence of static decisions to hire workers and use intermediate inputs for production. The optimal adoption decision trades off a sunk cost of robot adoption with gains in future profits from being able to operate the robot technology.

In Sections 3.1 and 3.2, I characterize the firm's static production problem taking the robot technology choice as given. In Section 3.3, I then characterize the firm's dynamic problem of adopting robot technology. The firm problem is linked to the worker's problem in general equilibrium but only through the path of wages. This separable structure allows me to study and estimate the firm model in isolation by conditioning on the observed path of wages, and postpone the specification of the worker's problem to Section 5.

3.1 Production Technology

A manufacturing firm j uses workers of different occupations $L \in \mathbb{R}_+^{|\mathcal{O}|}$ and intermediate inputs $M \in \mathbb{R}_+$ according to the CES production function

$$Y_{jt} = F(M_{jt}, L_{jt} | R_{jt}, \varphi_{jt}) = z_{Hjt} \left\{ M_{jt}^{\frac{\sigma-1}{\sigma}} + \sum_{o \in \mathcal{O}} z_{ojt}^{\frac{1}{\sigma}} L_{ojt}^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}} \quad \text{with} \quad (1)$$

$$z_{Hjt} = \exp(\varphi_{Hjt} + \gamma_H R_{jt}) \quad (2)$$

$$z_{ojt} = \exp(\varphi_{ojt} + \gamma_o R_{jt}) \quad (3)$$

Firms are heterogeneous with respect to a vector of exogenous baseline productivities $\varphi \in \mathbb{R}^{\mathcal{O}+1}$ and an endogenous robot technology state $R \in \{0, 1\}$. The parameter γ_H captures the effect of robot technology on firm Hicks-neutral productivity z_H , and the parameters γ_o govern how robot technology changes the relative productivities of worker occupations in production z_o (measured relative to intermediate inputs M).¹

¹Intermediate inputs M include all non-labor inputs including materials and conventional capital equipment. I measure payments to these intermediate inputs as the part of firm sales that are not paid to labor or profits. As Section 3.3 will make clear, industrial robots are different from other non-labor inputs in that their adoption involves a change of production technology that is subject to a sunk robot adoption cost.

In modeling robot adoption as a technology choice, I follow a growing literature arguing for task-based models to study automation (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018b). Appendix B.1 derives the specification in Equation (1) from a micro-founded model in which robots substitute for production tasks performed by workers. I model robot technology as a binary state $R \in \{0, 1\}$ to reflect the fact that most robot users invest in robots in a single year only (Fact 1 from Section 2.1).

3.2 Demand and Flow Profits

The firm faces an iso-elastic demand curve

$$Y_{jt} = Y_{Mt} \times (P_{jt}/P_{Mt})^{-\epsilon}, \quad (4)$$

where Y_{Mt} is the aggregate manufacturing demand and P_{Mt} is the manufacturing price index. The firm takes the vector of factor prices w_t as given, such that the flow profit function reads

$$\pi_t(R, \varphi) = \max_X \left\{ P_{Mt} Y_{Mt}^{\frac{1}{\epsilon}} F(X|R, \varphi)^{1-1/\epsilon} - w_t^T X \right\} = \Omega_t C_t(R, \varphi)^{1-\epsilon}, \quad (5)$$

where C_t denotes the unit cost function, Ω_t is a common profit shifter, and the static inputs are stacked into the vector $X = (M, L)$.² By lowering production costs C_t , the robot technology allows firms to scale up output and increase flow profits.

The key assumption in Equation (1) is that the production function admits a static factor demand system (satisfying Equation (5)) that is invertible in firm productivities. Invertibility allows me to control for unobserved firm productivities by matching on observed factor choices, similar to the proxy variable approach to production function estimation (Akerberg et al., 2015; Levinsohn and Petrin, 2003; Olley and Pakes, 1996). Berry et al. (2013) show that a demand system is invertible if and only if it satisfies a “connected substitutes” condition. The set of such production functions includes CES as in Equation (1), non-homothetic CES, nested

²The unit cost function and profit shifter are given by the CES expressions

$$C_t(R, \varphi) = \frac{1}{z_H(R, \varphi)} \left\{ \sum_{x \in X} (w_{xt}/z_x(R, \varphi))^{1-\sigma} \right\}^{\frac{1}{1-\sigma}}, \quad \Omega_{jt} = P_{Mt}^\epsilon Y_{Mt} (\epsilon - 1)^{(\epsilon-1)} \epsilon^{-\epsilon}. \quad (6)$$

CES, mixed CES, and translog. Appendix C.2.2 relaxes the robot technology effects in Equation (2)-(3) to a distributed lag model to account for any adjustment dynamics in the transition of firms to robot production. The demand curve in Equation (4) can be relaxed to an arbitrary downward-sloping function as considered in Doraszelski and Jaumandreu (2018). Appendix E derives an extension of the model where firms face upward-sloping labor supply curves and thus do not take wages as given in Equation (5).

3.3 Adoption of Robot Technology

The firm faces a dynamic decision about whether and when to adopt the robot technology R . The optimal adoption decision trades off a sunk cost of robot adoption with gains in future profits from being able to operate robot technology. The sunk adoption cost includes a common time-varying component c_t^R and an idiosyncratic component ε_{jt}^R . The adoption decision is essentially an optimal stopping problem that is reminiscent of the seminal work on bus engine replacement by Rust (1987). The value of a firm is represented by the Bellman equation

$$V_t(0, \varphi) = \max_{R \in \{0,1\}} \pi_t(0, \varphi) - (c_t^R + \varepsilon_{jt}^R) \times R + \beta \mathbb{E}_t V_{t+1}(R, \varphi') \quad (7)$$

$$V_t(1, \varphi) = \sum_{\tau=0}^{\infty} \beta^\tau \pi_{t+\tau}(1, \varphi_{t+\tau}). \quad (8)$$

Robot technology does not depreciate in the baseline specification of the model.³

Firm baseline productivities evolve according to the Markov process

$$\varphi_{jt+1} = g_t(\varphi_{jt}, \dots, \varphi_{jt-k}) + \xi_{jt+1}, \quad \xi_{jt+1} \perp (\varphi_{jt}, \dots, \varphi_{jt-k}), (\varepsilon_{jt'}^R, \dots, \varepsilon_{jt-1}^R). \quad (9)$$

The idiosyncratic adoption cost shocks ε_{jt}^R are drawn i.i.d. from a cumulative distribution function F such that the probability that a firm adopts robot technology is

$$P_t(\Delta R_{jt+1} = 1) = F\left(\beta (\mathbb{E}_t V_{t+1}(1, \varphi_{jt+1}) - \mathbb{E}_t V_{t+1}(0, \varphi_{jt+1})) - c_t^R\right) \quad (10)$$

The multiplicative productivity effects of robots in Equations (2) and (3) imply that firms that operate on a larger scale will be better able to reap the benefits of robot technology. Combined

³Appendix C.5 specifies and estimates a model extension in which robots deteriorate at a fixed rate.

with the fixed component of robot adoption costs c_t^R , this allows the model to rationalize the observed size premium in robot adoption (Fact 2 from Section 2.1). It is, however, worth noting that the model also allows for variable costs of robot adoption through the γ_o parameters. Robot production will, for example, be more intensive in intermediate inputs if γ_o is negative or require more tech workers if γ_T is positive. The adoption model also implies that larger firms will spend more on robots when they adopt because these firms will be willing to pay a higher idiosyncratic adoption cost ε_{jt}^R .

The robot adoption model features two key simplifying assumptions about robot investment behavior. First, robot adoption is treated as a one-off decision. This assumption is motivated by the observed lumpiness (Fact 1 in Section 2.1) whereby most robot users invest entirely in a single year. Appendix C.5 estimates a model extension in which robots deteriorate at a fixed rate, thereby leaving scope for replacement investments. Second, firms cannot receive larger relative robot production effects γ by spending more on robots. The structural estimation in Section 4 will provide empirical evidence in support of this homogeneity assumption on the treatment effects of robot adoption.

Equation (7) entails a key timing assumption that robot adoption is decided one year in advance. Combined with the Markovian structure on the productivity process in Equation (9), this timing assumption will be key to separating out the causal impact of robot adoption on firm productivities in Section 4.⁴

4 Structural Estimation of Firm Robot Adoption

In this section, I estimate the robot adoption model presented in Section 3. The structure of the model allows me to estimate its parameters in sequence. In Sections 4.1 to 4.3, I estimate the parameters of firm production technologies without having to specify other parameters of the adoption model, including robot adoption costs. In Section 4.4, I then estimate the cost parameters of robot adoption. I set the elasticity of demand and the time discount factor to conventional values from the literature ($\epsilon = 4, \beta = 0.96$).⁵

⁴The timing assumption on investment decisions (a one-year time-to-build) combined with a Markov process for firm productivities is a common assumption in the production function estimation literature, including Olley and Pakes (1996) and Doraszelski and Jaumandreu (2013).

⁵I follow Bloom (2009) and Asker et al. (2014), who calibrate the elasticity of demand ϵ to 4 to reflect a markup on output prices of 1/3 and calibrate the annual discount rate β to the data reported in King and Rebelo (1999).

4.1 Elasticity of Substitution between Production Tasks

In this section, I estimate the elasticity of substitution between production tasks, σ . I distinguish between labor tasks of production workers, tech workers, and other workers.⁶ To preview, I use the model structure to derive an instrumental variables strategy, and I estimate that tasks are complements in firm production.

The first-order conditions for cost minimization in Equation (5) imply that firm factor demands satisfy the following relationship

$$\log(L_{o'jt}) - \log(L_{ojt}) = -\sigma(\log(w_{o'jt}) - \log(w_{ojt})) + \log(z_{o'jt}) - \log(z_{ojt}) \quad (11)$$

The challenge in using Equation (11) to estimate σ is the classic simultaneity problem (Marschak and Andrews, 1944) that wages w_{jt} may be correlated with firm productivities z_{jt} , which constitute the regression error term in Equation (11). Appendix E derives a model in which firms face upward-sloping labor supply curves, thus creating an explicit link between firm productivities and wages.

I use the structure of the model in Section 3 to derive a rational expectations generalized method of moments (GMM) estimator that explicitly solves this simultaneity problem. The identification strategy builds on the insight of Doraszelski and Jaumandreu (2018) that the Markovian structure on firm productivities implies that past factor choices X_{jt-1} and prices w_{jt-1} must be uncorrelated with the current productivity innovations ξ_{jt} . This restriction allows me to estimate σ from the moment condition

$$\mathbb{E} \left[A_{oo'}(Q_{jt-1})(\xi_{ojt} - \xi_{o'jt}) \right] = 0, \quad (12)$$

where $A_{oo'}$ is a vector function of the instruments Q_{jt-1} , including $\log(X_{jt-1})$ and $\log(w_{jt-1})$. The derivation of this moment condition closely follows Doraszelski and Jaumandreu (2018), and I therefore relegate the derivations to Appendix C.1. The key idea is to, first, break the productivity error term z_{jt} in Equation (11) into the predictable component g_{jt} and the innovation ξ_{jt} . Since firms behave with rational expectations, the unforeseeable innovations ξ_{jt}

⁶The classification of worker tasks builds on the occupational grouping of Bernard et al. (2017); see Appendix A.6 for details.

must be uncorrelated with past decisions and prices of firms. To the extent that lagged factor prices and decisions correlate with current factor prices, they thus constitute valid and relevant instruments for estimating the substitution elasticity σ .

I estimate Equation (12) on the sample of non-adopters, which allows me to identify σ without specifying how robot technology affects firm productivities in Equations (2)-(3), which I separately estimate in Section 4.2. I estimate the moment conditions using a two-step GMM procedure with Appendix C.1 providing additional details on the estimation problem. The GMM estimate of the elasticity of task substitution σ is 0.49, which implies that tasks are complements in firm production. This estimate is based on the Danish matched employer-employee data from 1995 to 2015.

Table 3: Estimating the Elasticity of Substitution between Tasks in Production

	GMM
Elasticity of task substitution, σ	0.493 (0.092)

To place this estimate in the literature, Doraszelski and Jaumandreu (2018) estimate that the elasticity of substitution between labor and materials lies between 0.4 and 0.8, while Raval (2019) estimates that the elasticity of capital-labor substitution falls between 0.3 and 0.5. There is, to my knowledge, no estimate in the existing literature of the micro elasticity of substitution between worker tasks, and one contribution of this section is to provide such an estimate.⁷

4.2 Robot Technology

In this section, I estimate the parameters of robot technology γ , a key input for evaluating the distributional impact of industrial robots. In Section 4.2.1, I first use the model in Section 3 to derive an identification strategy that is based on event studies of firm robot adoption. In Section 4.2.2, I then present the estimation results, which show that industrial robots increase production efficiency but cause a substantial bias in technology away from production workers and toward tech workers and intermediate inputs.

⁷An important reason for the absence of such an estimate is the lack of micro data on the labor tasks employed in firms. The detailed occupational codes in the Danish data are unusually rich in this regard.

4.2.1 Identification of Robot Technology

This section describes my strategy for identifying the parameters of robot technology, γ . I first discuss the identification challenges that arise from the fact that firms endogenously select into robot adoption. I then use the adoption model developed in Section 3 to derive an identification strategy that deals with this selection problem.

First, from the invertibility of the factor demand system, I can recover firm productivities from the first-order conditions to Equation (5)

$$z_{ojt} = l_{ojt} - m_{jt} + \sigma(\log(w_{ojt}) - \log(w_{Mjt})) \quad (13)$$

$$z_{Hjt} = \frac{1}{\epsilon - 1} m_{jt} + \frac{\sigma}{\epsilon - 1} w_{Mjt} + \frac{(\sigma - \epsilon)}{(\sigma - 1)(\epsilon - 1)} \log \left\{ w_{Mjt}^{1-\sigma} + \sum_o z_{ojt} w_{ojt}^{1-\sigma} \right\} \quad (14)$$

where lower-case factor choices denote log transforms. With these productivities recovered, it is now tempting to use Equations (2)-(3) to run the regression

$$\log(z_{jt}) = \gamma R_{jt} + \varphi_{jt} \quad (15)$$

The issue with using Equation (15) as an estimating equation is that firms adopt robots R_{jt} based on their expected baseline productivities φ_{jt} (see Equation (22)), which exactly is the error term in Equation (15), thus creating selection bias. For example, simply comparing robot adopters to non-adopters in the cross-section will create bias because high baseline productivity firms are better able to overcome the fixed cost of robot adoption. Similarly, simply comparing a firm before and after robot adoption will be biased because firms tend to adopt robots when their baseline productivity is high or when they expect to face high demand for their products. Indeed, Fact 2 of Section 2.1 showed that robot adopters tend to be larger.

As I will show formally below, the dynamic adoption model of Section 3 gives me a way to confront this selection problem. The key idea is to match on observed firm factor choices leading up to adoption to control for selection into robot adoption based on heterogeneity in firm productivities. The reason why observably similar firms make different decisions about robot adoption is then due to heterogeneity in the sunk costs of robot adoption ε_{jt}^R , which satisfies the exclusion restriction for identification in the model. The key identifying assumption

is that observed factor choices are sufficient to control for firm productivities, and that there is no selection on unobservables that directly affect firm outcomes. The matching-based event study identification strategy reads as follows.

Identification Strategy (Parameters of Robot Technology γ).

1. Take two firms with similar output and occupational wage bills in some initial k years.
2. In the following year, one of the firms adopts robots.
3. The differential paths of firm output and occupational wage bills identify the parameters of robot technology, γ .

The firm model in Section 3 falls into a general class of potential outcomes models for robot adoption. In these potential outcomes models, the assumptions for non-parametric identification of average treatment effects are well-understood (Imbens and Wooldridge, 2007). I first remind the reader of these general requirements for identification, and then show that they are satisfied in my adoption model. Finally, I show that the average treatment effects estimated by the event studies identify the robot technology model parameters of interest.

Note first that, since payments to intermediate inputs M are defined as the part of firm sales that is not paid to labor or profits (a constant markup on firm sales), matching on firm sales and occupational wage bills is equivalent in the model to matching on the full vector of firm factor bills, $X = (M, L)$.

In the model, a firm's factor demands $x_{jt} = (m_{jt}, l_{jt})$ can take two potential values, $(x_{jt}(0), x_{jt}(1))$, according to whether or not the firm has adopted robot technology. In the language of Rubin (1990), the two identifying assumptions are *unconfoundedness*

$$\{\Delta R_{jt+1} \perp\!\!\!\perp (x_{jt+1}(1), x_{jt+1}(0)) \mid (x_{jt-1}(0), \dots, x_{jt-k}(0))\} \quad (\text{A1})$$

and *overlap* in robot adoption

$$0 < P(\Delta R_{jt+1} = 1 \mid x_{jt-1}(0), \dots, x_{jt-k}(0)) < 1 \quad (\text{A2})$$

Assumption (A1) requires that, once I condition on the path of factor choices that lead a firm to adopt robots in year t , the act of adoption must be independent of the firm's potential factor

choice outcomes going forward. On top of this, Assumption (A2) requires that I can find another firm that experienced the same initial sequence of factor choices but did not adopt robots in year t . Under Assumptions (A1) and (A2), the difference in sample means between adopter and match firms identifies the average treatment effect of robot adoption (see Imbens and Wooldridge (2007))

$$\bar{x}_{t+1}^T - \bar{x}_{t+1}^C \xrightarrow{p} \mathbb{E} [x_{jt+1}(1) - x_{jt+1}(0) \mid j \in T], \quad (16)$$

where \bar{x}^T and \bar{x}^C denote the sample means for adopter and match firms, respectively.

Let us now see how the general identifying assumptions (A1) and (A2) derive from the adoption model in Section 3. First, by the invertibility of the factor demand system, I am implicitly conditioning on $(\varphi_{jt-1}, \dots, \varphi_{jt-k})$ when I match on firm factor choices in the k years that lead up to robot adoption (see Equations (13) and (14)).⁸ Once I condition on $(\varphi_{jt-1}, \dots, \varphi_{jt-k})$, firm future factor outcomes $(x_{jt+1}(0), x_{jt+1}(1))$ are driven solely by the productivity innovations ζ_{t+1} in Equation (9). Since these future productivity innovations are unforeseeable when firms choose to adopt robots in year t , the adoption model satisfies the *unconfoundedness* condition (A1).

Second, the probability of robot adoption in the model is given by

$$P_t(\Delta R_{jt} = 1 \mid \varphi_{jt-1}, \dots, \varphi_{jt-k}) = F\left(\beta\left(\mathbb{E}_t V_{t+1}(1, \varphi_{jt+1}) - \mathbb{E}_t V_{t+1}(0, \varphi_{jt+1})\right) - c_t^R\right) \quad (17)$$

which lies strictly within the unit interval as long as the distribution of idiosyncratic adoption costs F has full support. The adoption model thus also satisfies the *overlap* condition (A2). Put into words, the identification strategy relies here on firm heterogeneity in the costs of robot adoption ε_{jt}^R driving otherwise similar firms to make different decisions about robot adoption.

Finally, from the model equations (2), (3), (13) and (14), we see that the treatment effects in

⁸If wages are firm-specific, the identification strategy also requires me to match on wages. In the analysis below, I match on factor choices, and then show that the firms also match on wages. The non-targeted match on wages is as an overidentification check of the model assumption that robot adopters do not pay wage premiums.

Equation (16) identify the parameters of the robot technology

$$\gamma_o = z_{ojt}(1) - z_{ojt}(0) = (l_{ojt}(1) - l_{ojt}(0)) - (m_{jt}(1) - m_{jt}(0)) \quad (18)$$

$$\gamma_H = z_{ojt}(1) - z_{ojt}(0) \quad (19)$$

$$= \frac{1}{\epsilon - 1} (m_{jt}(1) - m_{jt}(0)) + \frac{(\sigma - \epsilon)}{(\sigma - 1)(\epsilon - 1)} \log \left\{ \frac{w_{Mjt}^{1-\sigma} + \sum_o z_{ojt}(1) w_{ojt}^{1-\sigma}}{w_{Mjt}^{1-\sigma} + \sum_o z_{ojt}(0) w_{ojt}^{1-\sigma}} \right\} \quad (20)$$

The identification of γ_H requires the values of the factor augmenting productivities z_{ojt} which at this point can be readily recovered from Equation (13).

4.2.2 Estimation Results

The identification strategy presented above suggests matching robot adopters to comparison firms with a similar path of factor choices leading up to the adoption event. The match firms found in column 3 of Table 2 in Section 2.1 satisfy exactly these criteria. To recap, I found these firms by matching each robot adopter to a non-adopter firm that operated in the same two-digit industry and had a similar trajectory of firm sales and line worker wage bill shares in the three years that led up to adoption.⁹ I then showed that these firms were similar to the robot adopters on the full vector of factor choices as required by the identification strategy above.¹⁰

Once I have matched firms based on their factor choices leading up to robot adoption, the model in Section 3 implies that the act of adoption is driven by the idiosyncratic cost shock ε_{jt}^R that is independent of all other drivers of firm outcomes. The fact that the adopter and match firms are similar on several non-targeted outcomes in Table 2 provides evidence in support of this identifying assumption. The fact that the firms pay similar wages, in particular, provides an overidentification check of the model assumption that robot adopters do not pay wage premiums.

To ease the exposition, I presented the adoption model in Section 3 assuming that the productivity effects of robotization γ manifest fully within the first year of adoption; see Equa-

⁹I use an Exact-Mahalabonis matching procedure described in Appendix A.7.1. The three-year match window allows for firm productivities in Equation (9) to follow an arbitrary Markov chain of length three.

¹⁰A test of the joint hypothesis that none of the covariates in Table 2 (firm sales, employment, and occupational wage bills) predict robot adoption has a p-value of 0.25 (reported in the second-to-last row of the table).

tions (2)-(3). When taking the model to the data, I allow for the possibility that firms take a longer time to fully adjust to robot production. In practice, I track firm outcomes for four years after robot adoption. This, however, opens the possibility that some of the control firms may have also adopted robots in the post-event time window. Appendix Figure C.1 shows that around 10 percent of control firms adopted robots four years after the event year, which works against finding an effect of robot adoption in the reduced form of the event studies. I take this change in treatment status into account when estimating the model parameters.¹¹

Figures 1 and 2 show the main results from the estimation of robot technology. The figures display the differential paths of firm size and factor choices around robot adoption as prescribed by the identification strategy above. The blue lines represent raw data and the dashed orange lines show the model fit.¹² As I showed in Section 4.2.1, these reduced-form effects exactly identify the parameters of robot technology γ .

I estimate the parameters of robot technology to match the reduced-form moments four years after robot adoption. I choose the four-year horizon to account for the smoother transition path to robot production found in the data. This transition path likely reflects complementary investments that occur post adoption but that the model assumes are incurred immediately upon adoption. Appendix C.2.2 generalizes the model in Section 3 to account for these dynamic adjustments to robot production by allowing the productivity effects of robot adoption in Equations (2)-(3) to follow a distributed lag model. The appendix section estimates the full dynamic path of robot productivity effects. This generalization adds to the computational complexity of the model by requiring me to keep track of the years since robot adoption when solving the firm's dynamic programming problem. With the aim of keeping the firm's state space tractable when solving the general equilibrium model in Section 6, I abstract from these dynamic adjustment processes and instead match directly on the reduced-form effects four years after robot adoption.

The figures show that the model-simulated diff-in-diffs tend to drift back toward zero in the years following adoption. This post-event drift toward zero reflects the control firms that adopt robots in the post-event time window (orange line in Appendix Figure C.1).

¹¹The model-implied correction is the Wald estimator used in the treatment effects literature to convert intention-to-treat (ITT) effects into treatment-on-the-treated (TOT) estimates; see Angrist and Pischke (2008).

¹²Appendix C.2.1 describes the econometric specification that generates the point estimates and confidence intervals plotted in Figures 1 and 2.

Figure 1(a) shows that the average firm’s sales increase 20 percent around robot adoption. Through the lens of the structural model, this sales effect implies that robot technology increases firm production efficiency by around 7 percent, given the calibrated elasticity of firm demand ϵ . Figure 1(b) shows that the wage bill increases by 8 percent around robot adoption. The wage bill increase is less than the 20 percent sales effect in Panel (a), and implies that the substitution effects of robot adoption on labor γ_o on average are negative.

Figure 1: Firm Outcomes Around Robot Adoption (Matching Diff-in-Diff)

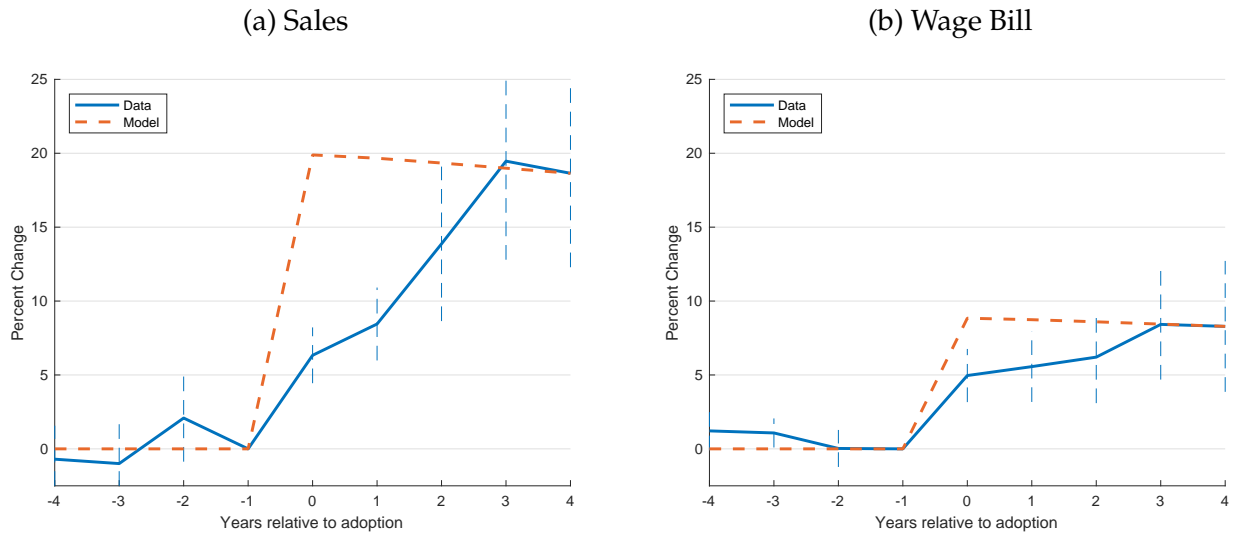


Figure 2 decomposes the wage bill effects in Figure 1(b) by occupations. Production workers include tasks from welding to assembly, while tech workers include engineers, researchers, and skilled technicians. Panel (a) of Figure 2 shows that the demand for production workers falls by around 20 percent around robot adoption, while Panel (b) shows that the demand for tech workers simultaneously increases by around 30 percent. This shift of labor demand away from the production line and toward the tech department implies that robot adoption lowers the relative productivity of production workers ($\hat{\gamma}_P = -0.461$) but increases the relative productivity of tech workers ($\hat{\gamma}_T = 0.043$).

Figure 2: Firm Wage Bills Around Robot Adoption (Matching Diff-in-Diff)

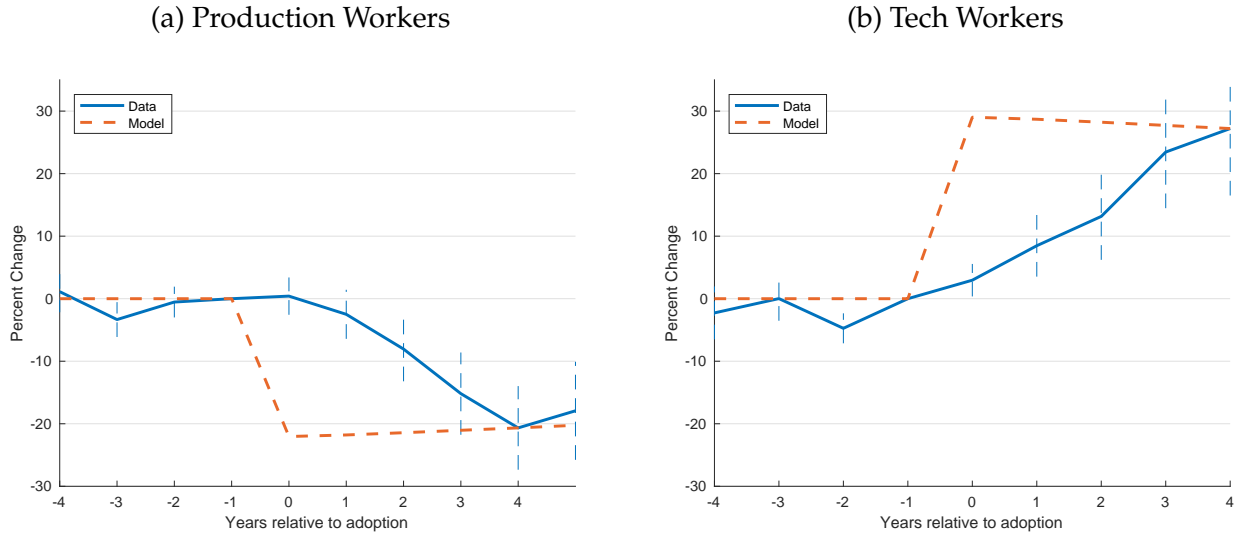


Table 4 summarizes the estimated parameters of robot technology.

Table 4: Estimated Parameters of Robot Technology

Parameter	Description	Estimated Value
γ_P	Production worker augmenting robot productivity	-0.461
γ_T	Tech worker augmenting robot productivity	0.043
γ_O	Other worker augmenting robot productivity	-0.115
γ_H	Hicks-neutral robot productivity (normalized)	0.066

Note: The relative productivity effects γ_o are measured relative to intermediate inputs. The parameter γ_H is normalized such that a zero sales effect of robot adoption would imply a value γ_H of zero.

The reduced-form effects in Figure 1 align well with Koch et al. (2019), who find that robot adoption increases output 20-25 percent and lowers labor costs per unit produced among Spanish manufacturing firms. It is worth keeping in mind that the reduced-form effects in Figures 1 and 2 only identify the partial effects of one firm adopting industrial robots, and that any general equilibrium effects of robotization are differenced out in the figures. The general equilibrium model in Section 6 will fit these partial effects but also take into account general equilibrium interactions in product and labor markets to be able to quantify what happens when many firms in the economy adopt industrial robots.

4.3 Baseline Technology

Baseline productivities φ_{jt} are structural residuals that capture changes in firm production technology that are not due to robot adoption. I can now recover these baseline productivities by inverting the model equations. To be precise, with the robot technology parameters γ estimated in Section 4.2.2 and firm productivities z_{jt} recovered from Equations (13) and (14), I can use Equations (2) and (3) to retrieve baseline productivities φ_{jt} .

To solve their forward-looking problem of robot adoption, firms must form expectations about their future productivities. To estimate this robot adoption problem, I specify that firm productivities (Equation (9)) follow a first-order vector autoregression VAR(1) with Gaussian innovations.

$$\varphi_{jt} = \mu_t + \Pi\varphi_{jt-1} + \zeta_{jt}, \quad \text{with} \quad \zeta_{jt} \stackrel{iid}{\sim} \mathcal{N}(0, \Sigma). \quad (21)$$

The unknown parameters (μ_t, Π, Σ) in Equation (21) can readily be estimated using either maximum likelihood or three-stage least squares.

The general equilibrium model in Section 6 restricts the labor-augmenting part of baseline productivities to a common time-varying parameter vector φ_{ot} . This simplification is done to keep the firm's state space tractable and to home in on the key size dimension that sets robot adopters apart from non-adopters (Fact 2 of Section 2.1).¹³ Appendix C.3.1 calibrates the path of common labor-augmenting productivities to match the path of manufacturing factor shares taking into account the diffusion of industrial robots. Appendix C.3.2 reports the results from estimating the productivity process in Equation (21). When solving the dynamic programming problem of robot adoption, I discretize the estimated baseline productivity process using the Tauchen (1986) method.

¹³The size premium in robot adoption is rationalized by the Hicks-neutral component of firm heterogeneity φ_{Hjt} which is left unrestricted. To be clear, the homogeneity restriction on firm baseline labor-augmenting productivities φ_{ot} is imposed solely for computational tractability: it does not alter the preceding analysis and can be relaxed without causing any conceptual or data complications.

4.4 Robot Adoption Costs

In this section, I estimate the costs of robot adoption. I first parameterize the path of common costs c_t^R and the distribution of idiosyncratic costs F , and then estimate their parameters to match the empirical robot diffusion curve and the observed firm size premium in robot adoption. To preview, I find that the model is able to generate the empirical S-shape in robot diffusion over time as well as the observed size premium of robot adopters, and that the estimated adoption costs align well with external cost measures.

I specify the idiosyncratic adoption cost shocks ε_{jt}^R to be drawn from a logistic distribution $F \sim \text{Logistic}(0, \nu)$ such that the probability that a firm adopts robot technology (Equation (10)) takes the form

$$P_t(\Delta R_{jt+1} = 1) = \frac{\exp(\frac{1}{\nu}(-c_t^R + \beta \mathbb{E}_t V_{t+1}(1, \varphi_{jt+1})))}{\exp(\frac{1}{\nu}(-c_t^R + \beta \mathbb{E}_t V_{t+1}(1, \varphi_{jt+1}))) + \exp(\frac{1}{\nu} \beta \mathbb{E}_t V_{t+1}(0, \varphi_{jt+1}))}. \quad (22)$$

To develop intuition for the estimation strategy that I adopt here, note that Equation (22) implies a linear relationship between the log odds ratio of robot adoption and the expected gain in future profits from operating industrial robots.

$$\log \left(\frac{P_t(\Delta R_{jt+1} = 1)}{1 - P_t(\Delta R_{jt+1} = 1)} \right) = -\frac{c_t^R}{\nu} + \frac{1}{\nu} \times (\beta \mathbb{E} V_{t+1}(1, \varphi_{jt+1}) - \beta \mathbb{E} V_{t+1}(0, \varphi_{jt+1})) \quad (23)$$

Equation (23) shows that the common cost c_t^R governs the rate of robot diffusion, while the sensitivity of robot adoption to future profit gains is inversely linked to the dispersion parameter ν .¹⁴ Since larger firms are the ones that can better scale up production to reap the benefits of robot technology, and thus enjoy larger profit gains when adopting robots, it follows that the size premium in robot adoption is also inversely tied to ν . Following on this intuition, I develop a simulation-based estimator that entails searching for the adoption cost parameters, c_t^R and ν , that bring the model as close as possible to the observed robot diffusion

¹⁴By inverting continuation values from choice probabilities as in Arcidiacono and Miller (2011), I can rewrite Equation (23) as follows

$$\beta \log P_{t+1} - \log \frac{P_t}{1 - P_t} = \frac{1}{\nu} (\beta c_{t+1}^R - c_t^R) - \frac{1}{\nu} \beta (\pi_{t+1}(1, \varphi') - \pi_{t+1}(0, \varphi')) \quad (24)$$

Equation (24) clarifies that the acceleration in robot diffusion pins down the change in robot adoption costs c_t^R over time, while $\frac{1}{\nu}$ measures the sensitivity of adoption to future profit flows.

curve and size premium in robot adoption.

I structure the exposition in two steps. In Section 4.4.1, I estimate the path of common adoption costs c_t^R to match the empirical robot diffusion curve, conditional on an estimate of ν . In Section 4.4.2, I then estimate the dispersion parameter ν to match the observed size premium in robot adoption. The final estimation procedure stacks the moments and estimates the parameters simultaneously using the method of simulated moments (MSM). Appendix C.4 provides details on the MSM estimation procedure.

4.4.1 Common Adoption Costs over Time

I estimate the path of common adoption costs $\{c_t^R\}_{t=0}^T$ to bring the model as close as possible to the observed robot diffusion curve. In particular, I parameterize the adoption cost schedule to be log-linear in time,

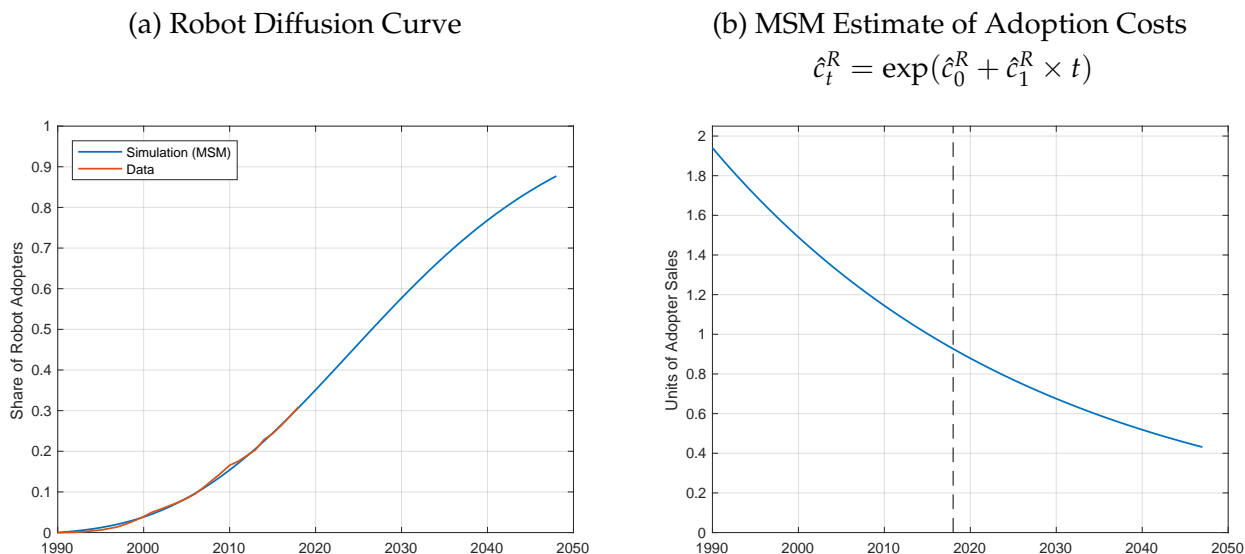
$$c_t^R = \exp(c_0^R + c_1^R \times t), \quad (25)$$

and then search over a grid of intercepts c_0^R and slopes c_1^R to minimize the distance between the simulated and empirical diffusion curve. That is, for each pair of intercept and slope (c_0^R, c_1^R) , I solve the dynamic programming problem of the firm, simulate the economy, and calculate the in-sample deviation to the empirical diffusion curve. The MSM estimator is the intercept-slope pair that brings the simulated diffusion curve the closest to the data. Appendix F.1 describes formally how to solve the dynamic programming problem of the firm. Put briefly, I first set a time horizon T sufficiently far in the future, such that robots are fully diffused by then. I then start at T , and solve the stationary, infinite horizon dynamic programming problem by iterating on the Bellman equation. I then solve for continuation values in $T - 1, T - 2, \dots$, back to the first period using backward induction. With the continuation values in hand, I can simulate firms forward using the adoption policy functions, and verify that industrial robots have actually diffused fully by time T .

Figure 3(a) compares the fit of the estimated adoption curve, and Figure 3(b) plots the MSM estimate for the path of adoption costs. The common component of robot adoption costs amounts to 0.9 times the adopter firms' sales in 2019. This is, however, not the average sunk cost $c_t^R + \varepsilon_{jt}^R$ borne by adopters because firms select into robot adoption based on their

idiosyncratic adoption cost ε_{jt}^R . Conditional on adoption, the total adoption cost amounts to around 10 percent of adopter sales.¹⁵ These are the costs needed to rationalize the fact that, despite enjoying substantial sales gains upon robotization, only 31 percent of manufacturing firms have adopted industrial robots almost 30 years after their arrival.

Figure 3: Estimating Adoption Costs on the Robot Diffusion Curve



Note: Firm sales (the units in Panel (b)) are an average of adopter sales measured over the full simulation period.

One notable feature of Figure 3 is that, despite the log-linear schedule for adoption costs, the model (blue line in Panel (a)) is able to generate the S-shaped diffusion curve commonly found in the literature on technology adoption (Griliches, 1957). This can be seen as an over-identification check of the estimated adoption model. The model-simulated S-shape reflects the combination of a Bell-shaped distribution for firm productivity and a model where robot adoption is driven by threshold crossing in firm productivity. The Gaussian cumulative distribution function for baseline Hicks-neutral productivity φ_H naturally gives rise to a tail of technology leaders, a bigger mass of followers, and a tail of laggards, as implied by an S-shaped diffusion curve.

The MSM adoption cost estimate is an inferred cost that not only includes the monetary

¹⁵Following Dubin and McFadden (1984), the expected cost borne by adopting firms may be calculated as

$$E(c_t^R + \varepsilon_{jt}^R | \Delta R_{jt+1} = 1) = c_t^R + v \left(\log P_t(\Delta R_{t+1} = 1) + \frac{P_t(\Delta R_{t+1} = 1)}{1 - P_t(\Delta R_{t+1} = 1)} \log P_t(\Delta R_{t+1} = 1) \right)$$

price of the robot machine but also expenditures for installation, the hassle of robot adoption and production reorganization, as well as changing accessibility of industrial robots. Still, we may ask how the inferred adoption cost from my estimation procedure compares to external measures of robot investment costs. Table 1 showed that robot adopters on average spend a total of \$311,000 on robot machinery. A rule of thumb is that machinery expenditures account for a third of the total cost of a robotic system that also includes expenditures for installation and integration (International Federation of Robotics, 2018). Taken together, this suggests that the monetary cost of robot adoption falls around \$1 million. This number is slightly smaller than, but in the ballpark of, the inferred cost for adopters ($c_t^R + \varepsilon_{jt}^R$) of around 10 percent of firm sales. Appendix C.4.3 shows further that the estimated rate of change in adoption costs c_t^R aligns well with the robot machine expenditures reported on customs records of adopting firms.

Importantly, the MSM estimation procedure also identifies the path of future adoption costs that are consistent with the observed adoption behavior. This future path of adoption costs will be key to evaluating the effects of imposing a robot tax in Section 6.3.

4.4.2 Variance of Idiosyncratic Adoption Costs


I estimate the dispersion in idiosyncratic adoption costs ν to match the observed size premium in robot adoption. Robot adopters were on average 2.61 times larger than non-adopter firms in 2018. The MSM procedure estimates ν to be 0.384, which delivers a simulated size premium of 2.61 in 2018. Figure 4 shows how the adopter size premium moment pins down the parameter ν by plotting the simulated size premium for varying values of ν .

To put this size premium into perspective, had selection into robot adoption been unrelated to firm size ($\nu \rightarrow \infty$), the adopter premium would only have reflected the 20 percent sales effect estimated in Section 4.2. At the other extreme, without heterogeneity in adoption costs ($\nu \rightarrow 0$), robot adopters would have been around 6 times larger than non-adopters in 2018.¹⁶ These estimates suggest that, while there is clear selection into robot adoption based on firm size (Fact 2 of Section 2.1), there is still ample heterogeneity in adoption costs ε_{jt}^R , leading

¹⁶The sales share of robot adopters in manufacturing was 53.9 percent in the data and in the model in 2018. In comparison, if firms did not select into robot adoption based on their size ($\nu \rightarrow \infty$) then the sales share of robot adopters would have been 34.5 percent. At the other extreme, without heterogeneity in adoption costs ($\nu \rightarrow 0$), the sales share would have been 72.8 percent.

observationally similar firms to make different decisions about robot adoption.

Figure 4: Size Premium of Robot Adopters for Varying Adoption Cost Dispersion ν



5 The Labor Supply Block

This section presents the labor supply block of the general equilibrium model. I incorporate this labor supply module into the general equilibrium model in Section 6 to allow for a labor supply response to industrial robots where workers move out of adversely affected occupations. I use here a dynamic occupational choice model that represents state-of-the-art for studying labor market dynamics in response to trade liberalizations (Dix-Carneiro, 2014; McLaren, 2017; Traiberman, 2019).

A key property of the general equilibrium is that the worker and firm problems are separable conditional on the path of wages. This block separable structure allows me to study and estimate the labor supply model now without reconsidering the firm’s problem from Section 3 by conditioning on the observed path of wages.

The labor force consists of overlapping generations of heterogeneous workers as in Lee and Wolpin (2006). Workers enter the labor market at age 25 with an educational skill level $s \in \{\text{Low, Mid, High}\}$ and retire at age 65. In each year before retirement, workers face a choice of which occupation o to work in. This labor supply decision is dynamic in two

ways. First, it is costly for workers to switch occupations. Second, workers may accumulate occupation-specific human capital on the job that is not transferable to other occupations. I allow labor markets to be segmented by occupation (production, tech, and other) and sector of employment (manufacturing and services).

A worker i of age a in occupation o in year t earns the product of a competitive occupational skill price, w_{ot} , and her human capital, H_{oit} . Her occupational human capital is given by

$$\log(H_{oit}) = \beta_s^o s_{it} + \beta_1^o a_{it} + \beta_2^o a_{it}^2 + \beta_3^o \text{ten}_{oit} + \zeta_{it} \quad (26)$$

where ten_o denotes tenure in occupation o , and $\zeta_{it} \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_h^2)$ is an ex-post productivity shock.

The worker's choice of occupation is an investment decision that trades off a sunk cost of switching occupations with future gains in wages and amenities of being employed in a new occupation. The occupational choice problem is represented by the Bellman equation

$$v_t(o, s, a, \text{ten}) = \max_{o' \in \mathcal{O}} \log(w_{ot} H_o(s, a, \text{ten})) + \eta_{ot} - (c_{oo'}(s, a) + \varepsilon_{o'}) \quad (27)$$

$$+ \mathbf{1}_{\{a < 65\}} \beta \mathbb{E}_t v_{t+1}(o', s, a + 1, \mathbf{1}_{\{o' = o\}}(\text{ten} + 1)) \quad (28)$$

where η_{ot} is a non-monetary amenity of working in occupation o , and $\varepsilon_o \stackrel{iid}{\sim} \text{GEV1}(\rho)$ is an idiosyncratic occupational switching cost shock. Income is implicitly assumed to be fully consumed in each period, and workers receive logarithmic flow utility of consumption. The occupational switching cost depends on the bilateral pair of current and prospective occupations, as well as the worker's age and skill

$$c_{oo'}(s, a) = c_{oo'} \exp \left\{ \alpha_s s + \alpha_1 \times a + \alpha_2 \times a^2 \right\} \quad (29)$$

I stack the worker state variables into the vector $\omega = (s, a, \text{ten}, o)'$.

5.1 Estimation of Labor Supply Parameters

I structurally estimate the labor supply model in Equations (26)-(28) using administrative data on the career paths of Danish workers. My approach to measurement and estimation follows

closely Traiberman (2019). I describe the estimation procedures below, and relegate the data description and estimation results to Appendices D.1 and D.2. To preview, the estimate show that production workers face steep barriers to switching into tech occupations, that it is easier for workers to switch sectors instead of occupations, that workers accumulate specific human capital on the job that is not transferable to other occupations, and that older workers find it more costly to reallocate in the labor market.

5.1.1 Human Capital Function

I estimate the human capital function in Equation (26) using a Mincer regression of log earnings on worker skill, age, and occupational tenure.

$$\log(\text{Earnings}_{it}) = \log(w_{ot}) + \beta_s^o s_{it} + \beta_1^o a_{it} + \beta_2^o a_{it}^2 + \beta_3^o \text{ten}_{oit} + \zeta_{it}, \quad (30)$$

where Earnings_{it} denotes labor earnings of worker i in year t , and w_{ot} is an occupation-time fixed effect. The key model assumption that enables me to identify the human capital parameters β in this regression is that workers cannot select on the productivity shock ζ when choosing occupation or education. Appendix Table D.1 provides the OLS estimation results, which align with estimates from the existing literature (Ashournia, 2017; Dix-Carneiro, 2014; Traiberman, 2019).

5.1.2 Occupational Switching Costs

I estimate the occupational switching costs $c_{oo'}$ on observed worker transition and a conditional choice probability (CCP) estimator adapted from Traiberman (2019). The estimator exploits the finite dependence in the labor supply model to difference out unobserved continuation values by comparing workers who start and end in the same states (Arcidiacono and Miller, 2011).

The occupational choice model in Equation (27) implies that the difference in the (discounted) probabilities of observing a worker in occupation o first switching into occupation o' and then transitioning into occupation o'' compared to observing the worker first staying in

occupation o and then transitioning into occupation o'' is

$$\log \frac{\pi_t(o o' | \omega)}{\pi_t(o o | \omega)} + \beta \log \frac{\pi_{t+1}(o' o'' | \omega')}{\pi_{t+1}(o o'' | \omega'')} = -\frac{1}{\rho} c_{o o'}(\omega) - \frac{\beta}{\rho} (c_{o' o''}(\omega') - c_{o o''}(\omega'')) \quad (31)$$

$$+ \frac{\beta}{\rho} (\log(w_{o' t+1} H_{o'}(\omega')) - \log(w_{o t+1} H_o(\omega''))) \quad (32)$$

$$+ \frac{\beta}{\rho} (\eta_{o'} - \eta_o) + \zeta_{o o' o'' t} \quad (33)$$

where $\pi_t(o o' | \omega)$ is the transition rate from occupation o to o' of workers with characteristics ω , H_o and w_{ot} are the human capital function and occupational skill prices estimated in Equation (30), and ζ is a mean-zero expectational error that is uncorrelated with the remaining RHS variables.

The occupational switching costs $c_{o o'}$ are identified off the excess likelihood of observing a worker staying in his own occupation from one year to the other, once his expected earnings differentials across occupations are controlled for. The occupational preference shock variance ρ is estimated as the inverse elasticity of occupational switching with respect to expected earnings differentials.

The key model assumption in Equations (31)-(33) is that occupational switching is a renewal action that clears past choices from a worker's state. Combining this assumption with the Hotz-Miller inversion of continuation values from choice probabilities (Hotz and Miller, 1993) allows me to cancel out continuation values.¹⁷

Equations (31)-(33) constitute a system of non-linear regressions that identify the switching cost function $c_{o o'}$ and the preference shock variance ρ . Appendix D.2.1 describes the computational implementation of the estimation procedure. Appendix Tables D.2 and D.3 present the non-linear least squares (NLLS) estimation results. The estimates show that production workers face steep barriers to switching into tech occupations, that workers find it easier to switch sector within the same occupation, and that older workers find it more costly to reallocate in the labor market. The estimated switching cost magnitudes are in the range of those found in the existing literature.

The NLLS procedure tightly estimate all the occupational choice parameters, except for the preference shock variance ρ . In the current setup, the estimate of ρ greatly exceeds estimates

¹⁷The derivation of Equations (31)-(33) closely follows Traiberman (2019), who estimates a richer model of labor supply that also accounts for unobserved (to the econometrician) types of workers.

in the existing literature. Since the labor supply responses to industrial robots are inversely related to this dispersion parameter, I choose to instead use a central estimate in the literature of ρ equal to 2. This value falls in between the estimates in Dix-Carneiro (2014), Ashournia (2017), Artuç et al. (2010), Caliendo et al. (2019), and Traiberman (2019).

5.1.3 Occupational Amenities

I estimate the path of occupational amenities η_{ot} to match the time series of employment shares across occupations. Appendix D.2.2 provides details on this estimation step.

6 Counterfactual Experiments

This section conducts counterfactual experiments to assess the general equilibrium impacts of industrial robots. I first present a general equilibrium model that unites the firm model from Section 3 with the worker model from Section 5. Section 6.1 defines the general equilibrium and develops a fixed-point algorithm for solving the equilibrium that features two-sided heterogeneity and dynamics. Section 6.2 uses the general equilibrium model to quantify how the arrival of industrial robots has affected the distribution of worker welfare. Section 6.3 evaluates the dynamic incidence of a robot tax.

6.1 Closing the General Equilibrium Model

The economy consists of a manufacturing sector and a service sector. The manufacturing sector consists of a mass $\mu_t^F(R, \varphi)$ of firms that are monopolistically competitive in product markets, pricetakers in factor markets, and otherwise operate as specified in Section 3.¹⁸ Services are produced with a Cobb-Douglas technology and supplied competitively,

$$Y_{st} = z_{st} M_{st}^{\alpha_M} \prod_{o \in \mathcal{O}} L_{ost}^{\alpha_o^S} \tag{34}$$

¹⁸The baseline mass of firms $\mu_t^F(\cdot, \varphi)$ is taken as given but its distribution over the robot technology state R evolves endogenously according to the equilibrium robot adoption model.

The economy is populated by a mass $\mu_t^W(\omega)$ of workers who supply labor as specified in Section 5, and consume the final output bundle

$$Y_t = Y_{Mt}^\mu Y_{St}^{1-\mu} \quad \text{with} \quad Y_{Mt} = \left[\int Y(R, \varphi)^{\frac{\epsilon-1}{\epsilon}} d\mu_t^F(R, \varphi) \right]^{\frac{\epsilon}{\epsilon-1}} \quad (35)$$

I model Denmark, a country of less than 6 million people located in the European free trade zone, as a small open economy. Intermediate inputs M are imported at world price w_{Mt} , which the Danish economy takes as given, and trade is balanced. The robot adoption cost c_t^R is determined on the world market for industrial robots and is thus exogenous to local conditions in Denmark. The general equilibrium of the economy is defined as follows.

Definition 1 (Dynamic General Equilibrium). A dynamic general equilibrium of the economy is a path of factor prices $\{w_t\}_t$, distributions of firm and worker states $\{\mu_t^F(R, \varphi), \mu_t^W(\omega)\}_t$, and policy functions $\{R_t(0, \varphi)\}_t, \{o'_t(\omega)\}_t$, such that taking the schedule of adoption costs $\{c_t^R\}_t$ and the price of intermediate inputs $\{w_{Mt}\}_t$ as given

1. Firms adopt robots to maximize expected discounted profits (Equation (7)) and demand static inputs to maximize profits period-by-period (Equation (5)).
2. Workers choose occupations to maximize expected present values (Equation (27)).
3. Labor markets clear (segmented by occupations and sectors)

$$\int L_{ot}(R, \varphi) d\mu_t^F(R, \varphi) = \int_{\omega} H_o(\omega) d\mu_t^W(\omega|M) \quad (36)$$

$$L_{ost} = \int_{\omega} H_o(\omega) d\mu_t^W(\omega|S), \quad (37)$$

where $L_{ot}(R, \varphi)$ is the static labor demand function satisfying Equation (5).

4. Firm output markets clear and trade is balanced.

$$Y_t = C_t + w_M M_t \quad (38)$$

where $M_t = \int M_t(R, \varphi) d\mu_t^F(R, \varphi) + M_{st}$ and $C_t = \sum_o w_{ot} L_{ot}^S + \Pi_t$. Equation (38) states that expenditures on intermediate input imports equal revenues from final goods exports.

5. The evolution of the distributions of firm and worker states $\{\mu_t^F, \mu_t^W\}_t$ is consistent with the policy functions $\{R_t(0, \varphi), o'_t(\omega)\}_t$.

A key property of the general equilibrium is that the firm and worker programs are separable conditional on the path of wages. This block separability breaks the curse of dimensionality where firm variables become states for the worker, and worker variables become states for the firm. The myriad of individual decisions taken by heterogeneous firms and workers is instead summarized into one aggregate state vector – the path of wages – which agents have perfect foresight about, up to unanticipated aggregate shocks to the economy. The block separable structure enables me to incorporate the rich firm and worker heterogeneity estimated in Sections 4 and 5, and still be able to compute the dynamic general equilibrium. In particular, the estimated general equilibrium model will fit the partial effects of firm robot adoption identified in Section 4 but also take into account how robotization affects non-adopter firms through product and labor markets, as well as the ability of workers to switch out of adversely impacted occupations.

I solve for the transitional dynamics of the economy where baseline productivities $\{\varphi_{jt}, z_{st}\}$, amenities $\{\eta_{ot}\}$, and robot adoption costs $\{c_t^R\}$ all have t-subscripts and are the time-varying fundamentals driving the system over time. The baseline estimated model perfectly matches the path of manufacturing factor bills (Appendix Figure C.3) and occupational employment shares (Appendix Figure C.3) observed in Denmark over time. I calibrate μ to match the manufacturing share in total output of the Danish economy and α_s to match the evolution of factor cost shares outside of manufacturing. Appendix Table G.1 provides a summary of the parameters of the general equilibrium model, as well as the moments used to estimate their values.

6.1.1 Solving the Dynamic General Equilibrium


The path of wages is the key endogenous variable that links the firm and worker decisions in general equilibrium. I solve for the general equilibrium wage schedule using a shooting algorithm adapted from Lee (2005). The procedure boils down to guessing a path of wages and manufacturing price indices, solving the dynamic programs related to the robot adoption decision of firms and the occupational choice problem of workers, simulating the economy

forward using the firm and worker policy functions, and then using the firm’s static labor demand functions to find the vector of wages that clear labor markets period-by-period. This algorithm iterates until convergence in the path of wages and the distributions of firm and worker states. Appendix F.3 details each step of the equilibrium solution algorithm.

6.2 The Distributional Impact of Industrial Robots

This section turns to the key question posed in this paper by asking how the distribution of worker earnings would have looked if industrial robots had not arrived. To evaluate this counterfactual, I solve the general equilibrium under a path of prohibitively high adoption costs ($c_t^R = \infty$). I then compare the results to the equilibrium under the baseline adoption cost schedule estimated in Section 4. The simulations assume that the arrival of industrial robot technology around 1990 came as a surprise to agents in the economy, but that firms and workers from that point on perfectly foresee the path of robot adoption costs. The robot diffusion curve in Figure 5 shows that if robot adoption had been infinitely costly (“No Robotization”), then robot technology would not have diffused at all.

Figure 5: Robot Diffusion Curve




The equilibrium effects of industrial robots depend not only on the direct impact of firm robot adoption estimated in Figures 1 and 2 but also on several indirect effects that are not

identified in micro-level diff-in-diff regressions. The indirect effects include the extent to which the expansion of robot adopters crowds out non-adopter firms in product and labor markets as well as the ability of workers to reallocate across occupations in response to equilibrium wage pressures from robot diffusion. The general equilibrium model captures these indirect effects by combining the structurally estimated behavior of firms and workers with internal consistency constraints imposed by equilibrium conditions on product and labor markets.

Figure 6 shows the impact of industrial robots on real wages in different occupations. Industrial robots have increased average real wages by 0.8 percent in Denmark but with substantial distributional consequences. Production workers employed in manufacturing are the big losers from industrial robots, as their real wages are 6 percent lower today due to robots. Tech workers employed in manufacturing earn 2.3 percent higher real wages today due to industrial robots, while the remaining occupations have gained between 0.3 and 1.2 percent from robots. While the real wage loss for production workers in manufacturing is substantial, it is important to keep in mind that the occupation only constitutes around 3 percent of total employment in Denmark.

Figure 6: Real Wage Effects of Industrial Robots
(Weighted Average in 2019: +0.76 percent)




To understand the general equilibrium forces driving the real wage outcomes, Figure 7 de-

composes the manufacturing production real wage effect into *labor demand* effects from robot adoption, *consumer price* effects from pass through of lower robot production costs, and *labor supply* effects from occupational reallocation of workers (changing the relative scarcity of labor across occupations).

As the decomposition shows, the real wage loss of manufacturing production workers would have been a half-order of magnitude larger than the estimated effect if workers could not reallocate across occupations in response to robots. Appendix Figure G.2 confirms this finding by evaluating the impact of industrial robots with exogenous labor supply, thus shutting off the occupational choice block estimated in Section 5.1. Real wages of production workers employed in manufacturing would in that world have been 30 percent lower today due to industrial robots.

Figure 7: Decomposition of the Production Wage Effect



Note: Labor demand effects are measured relative to the “Other Workers” occupation in the services sector.

Still, the labor supply and consumer price effects combined are not enough to overturn the negative labor demand effects of robot adoption from depressing real wages of production workers employed in manufacturing. The displacement effects identified in Figure 2(a) are in general equilibrium reinforced by two additional labor demand forces. First, the expansion of robot adopters crowds out activity in non-adopter firms through the stealing of

output markets. Second, the complementarity between occupations in manufacturing production (estimated in Section 4.1) means that firms spend a smaller fraction of their wage bill on production workers when they become less expensive.

Interestingly, among workers in the service sector, Figure 6 shows that production workers have experienced the largest real wage gain from robot adoption. This differential wage gain is a compensating differential for their excess risk of transitioning into production work in the manufacturing sector. In terms of expected lifetime earnings, production workers are the group of service workers with the lowest gain from industrial robots.


Finally, Figure 7 shows that more than half of the total consumer price gains from industrial robots have been realized already, even though only 30 percent of manufacturing firms have adopted robots. This finding reflects that the estimated model captures the fact that firms with larger efficiency gains from robot adoption (that is, firms that can better scale up production to take advantage of industrial robots) are the ones that adopt robots first.

Due to the possibility that workers can reallocate across occupations, the real wage effects in Figure 6 do not necessarily convert one-to-one into welfare effects for individual workers. The occupational reallocation margin opens an *option value* of being able to switch into occupations whose premiums rise as robots diffuse in the economy. As emphasized by Artuç et al. (2010), this option value source of worker welfare is not identified from static wage comparisons but is only captured once we factor in the dynamic occupational switching behavior observed over an individual's working life.

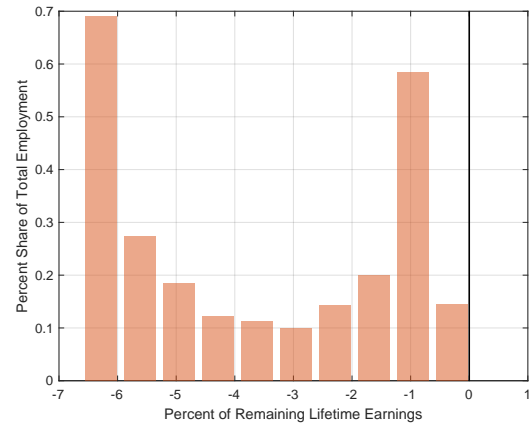
Figure 8 shows the impact of industrial robots on the welfare of workers in 2019. Panel (a) shows that 90 percent of workers have gained between 0.5 and 1 percent of lifetime earnings from the arrival of industrial robots. Yet, Panel (b) shows that the – considerably smaller – group of production workers employed in manufacturing have lost between 0 and 6 percent of lifetime earnings from robots.

Figure 8: Welfare Effects for Workers in 2019

(Average: +0.85 percent)



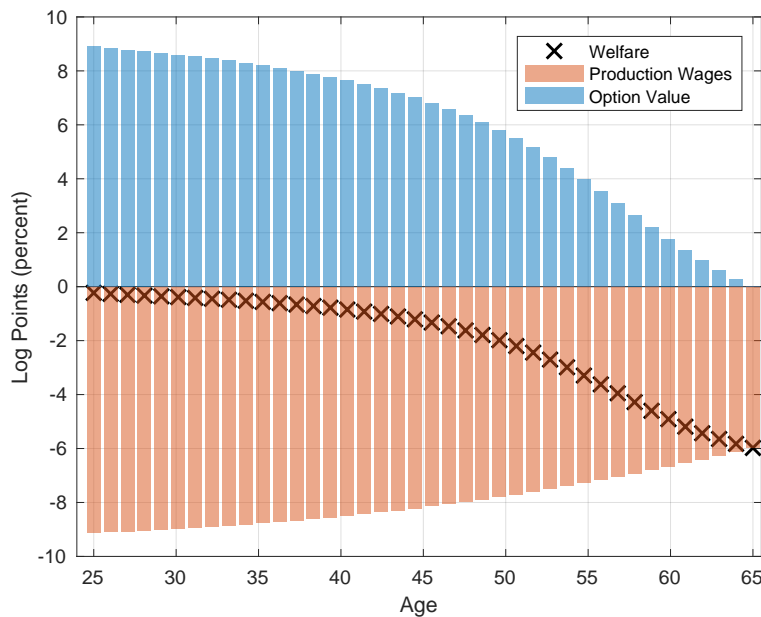
(a) All Workers Excl. Manufacturing Production



(b) Manufacturing Production Workers

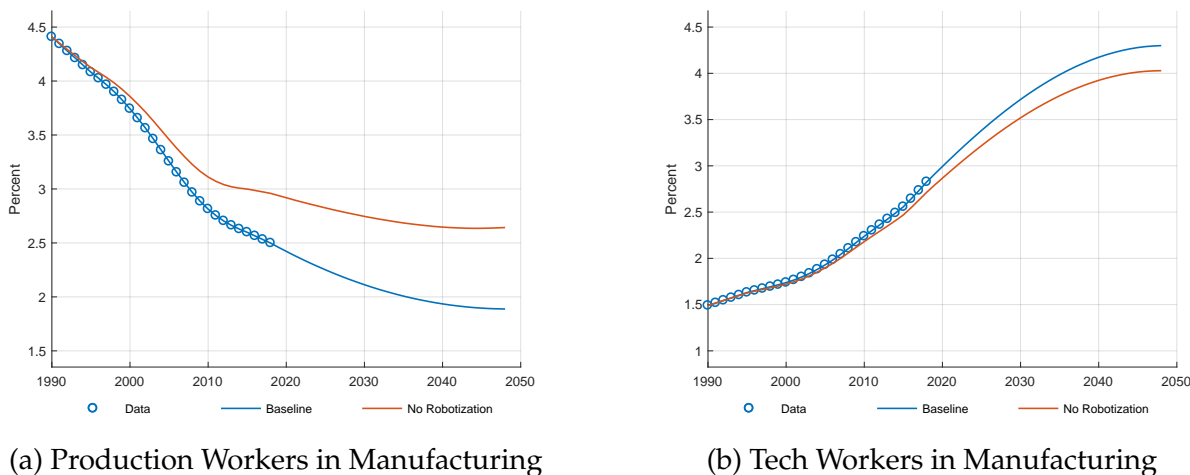
Figure 9 shows that the welfare losses in Figure 8(b) are concentrated on older workers. Younger production workers, with less specific skills and a long career ahead of them, are less affected by the arrival of industrial robots, as wage losses in their current occupation are offset by gains in the option value of switching into occupations whose premiums rise as robots diffuse in the economy.

Figure 9: Welfare Effects for Manufacturing Production Workers in 2019



The flip side of the labor supply responses found in Figure 7 is that industrial robots have contributed to employment polarization as documented in Autor and Dorn (2013) and Goos et al. (2014). Figure 10 shows that industrial robots can account for 25 percent of the fall in the employment share of manufacturing production workers and 8 percent of the rise in the employment share of tech workers in manufacturing since 1990.

Figure 10: The Effect of Industrial Robots on Employment Shares



To recapitulate, the estimates presented in this section are based on a general equilibrium model that has been validated on event studies of firm robot adoption, the observed diffusion of industrial robots, and worker transitions across occupations. The quantitative importance of the estimated general equilibrium responses to industrial robots warrants caution when comparing estimates from this section to findings in the reduced-form literature. For example, the conclusion from Figure 6 that industrial robots have increased average real wages may at first sight seem at odds with the finding in Acemoglu and Restrepo (2019b) that robots depress wages in local labor markets. Before drawing such a comparison, however, it is important to keep in mind that Figure 6 takes into account the general equilibrium consumer price and input-output linkage effects of industrial robots. For example, insofar as consumer price effects spill over across local labor markets, these contributions to real wages will be differenced out in the empirical strategy adopted in Acemoglu and Restrepo (2019b). In fact, the average real wage gain estimated in Figure 6 flips to a loss if I omit the consumer price effect of industrial robots. Furthermore, to the extent that some of the positive contributions to the service sector through input-output linkages extend beyond commuting zones, these effects

will also be differenced out in a diff-in-diff analysis of local labor markets. In fact, my estimate for the average wage effect in the manufacturing sector aligns well with the 0.4 percent loss estimated in Acemoglu and Restrepo (2019b).

In summary, I take the estimates presented in this section as complementary to existing reduced-form studies of industrial robots by highlighting the quantitative importance of general equilibrium effects that are not easily identified by reduced-form empirical strategies. In particular, I show the quantitative relevance of an occupational switching feedback mechanism that has been emphasized in the literature on international trade and labor market dynamics (Dix-Carneiro, 2014; McLaren, 2017; Traiberman, 2019). Although the labor supply responses are not strong enough to overturn the negative labor demand effects from depressing the real wages of manufacturing production workers, I find that the wage losses would have been a half-order of magnitude larger if workers could not reallocate across occupations. A speculative hypothesis is that the generous retraining subsidies offered in the Danish system of active labor market policies could be an underlying driver of the quantitative importance of the estimated occupational reallocation feedback response.

6.3 Policy Counterfactuals: The Dynamic Incidence of a Robot Tax

As a final counterfactual experiment, I now turn to evaluating the impact of a robot tax. The European Parliament voted in 2017 on a proposal to tax the use of robotics. The robot tax was motivated as a way to slow down the speed of robot adoption to give the economy more time to adjust to the new technology.¹⁹

I tax the schedule of robot adoption costs c_t^R to inform this policy counterfactual. To be clear, the undistorted equilibrium of the model is efficient (except for markups in product markets), but the robot tax could be motivated by distributional concerns.²⁰ In particular, Section 6.2 identified a group of production workers employed in manufacturing who have clearly lost from the use of industrial robots. A key policy question is how costly (in terms

¹⁹The proposal was ultimately voted down by the European Parliament but the idea of taxing robots to mitigate labor market polarization remains popular among public figures from Bill Gates (Quartz, 2017) to congresswoman Alexandria Ocasio-Cortez (Market Watch, 2019).

²⁰The *production efficiency* result of Diamond and Mirrlees (1971) establishes that it is always optimal to maintain production efficiency insofar as linear commodity taxes are available. Costinot and Werning (2018) derive sufficient-statistic formulas for optimal technology taxes when a non-linear income tax schedule is the only alternative policy instrument.

of lost economic efficiency) it is to insulate these production workers by taxing the further adoption of industrial robots. The answer to this question depends on several behavioral elasticities estimated from the micro data, including the sensitivity of firm robot adoption with respect to adoption costs (Section 4.4.2) as well as the ability of workers to switch occupations in response to robots (Section 5.1). I use the estimated general equilibrium model to quantify the distributional implications of a robot tax and to evaluate its impact on aggregate economic activity.

To map out the potential policies, I evaluate both a temporary and a permanent tax, each of 30 percent. The policies are announced and implemented in 2019, and the temporary tax is put in place for 10 years. Figure 11(a) shows the path of robot adoption costs under the tax policies. I assume that a robot tax in Denmark does not alter the pre-tax price for robots which is determined on world markets.

Panel (b) of Figure 11 shows the first key result from the robot tax counterfactuals: The temporary tax is more effective in slowing down the diffusion of industrial robots while it is put in place. With the temporary tax, only 43 percent of manufacturers will have adopted robots by 2029, compared to 48 percent with the permanent tax and 56 percent in the baseline scenario. The larger short-term effects of the temporary tax reflect the forward-looking nature of adoption, where firms foresee that the robot tax will expire and postpone adoption until then. The flip side of these delays is that the adoption of robots accelerates beyond its baseline speed after the temporary tax expires in 2030.

Figure 11: Robot Tax Counterfactuals

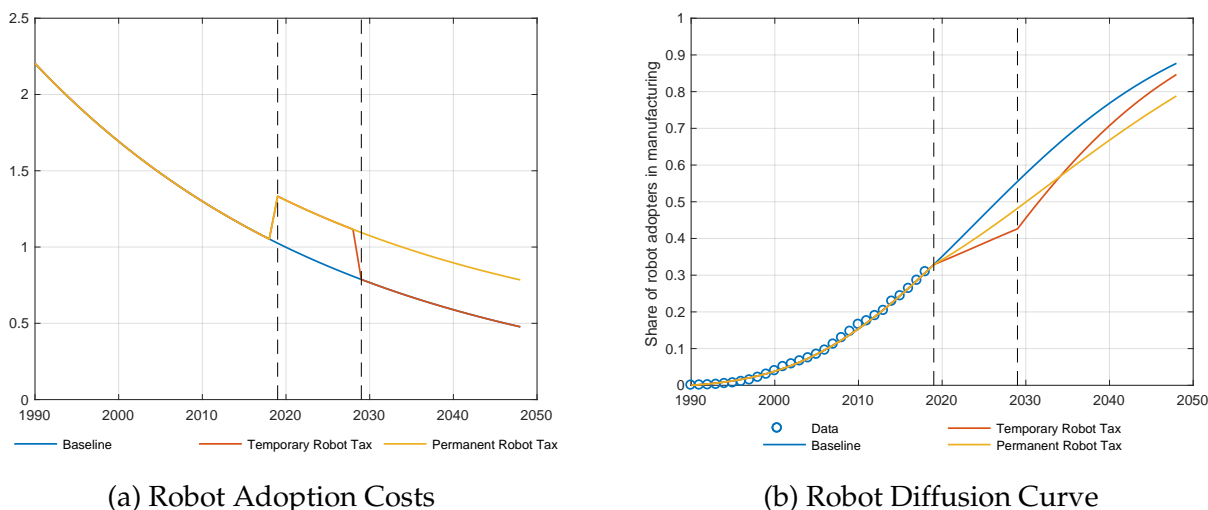


Figure 12 shows how the temporary robot tax affects the welfare of workers in 2019. The temporary tax lowers average welfare by 0.05 percent of lifetime earnings but benefits a group of older production workers employed in manufacturing by 0.2 to 0.3 percent.

Figure 12: The Impact of a Temporary Robot Tax on the Welfare of Workers in 2019
(Average: -0.054 percent)

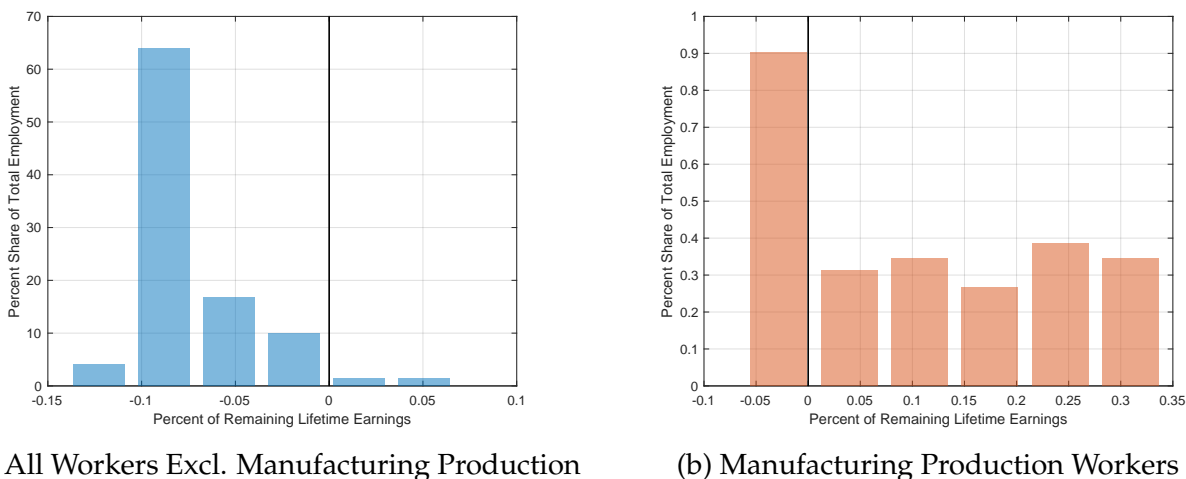


Table 5 shows how the burden of the robot taxes falls on workers and firms in the economy. Measured in present discounted terms, the robot taxes redistributes a total of 0.01 to 0.02 percent of GDP to production workers currently employed in manufacturing at the expense of a total welfare loss for workers of around 1 percent of GDP. These welfare losses reflect foregone efficiency gains from underinvestment in robot technology. Put differently, for the robot taxes to enhance social welfare, one needs to value production workers in manufacturing 50 to 100 times higher than the average worker.

The temporary robot tax creates welfare losses per dollar of tax revenue collected that are considerably larger than those of the permanent robot tax. These larger relative efficiency losses of the temporary tax is a direct consequence of the investment delays observed in Panel (b) of Figure 11: The intertemporal shifting of robot adoption out of the temporary policy window creates misallocation without raising tax revenues. In particular, if firm adoption behavior did not respond to the robot tax (“Mechanical Effect” in Table 5), the temporary robot tax would generate 133 percent more revenues, while revenues from the permanent tax would be only 17 percent higher.

The robot taxes do, however, generate substantial amounts of tax revenue, whose burdens

are primarily borne by manufacturing firms. As Table 5 shows, the tax revenues are sufficient to make all workers better off from the robot taxes, insofar as the revenues can be rebated appropriately and the planner does not care about firm profits. One should be cautious about drawing such a conclusions, however, as I do not model firms' entry decisions. If the robot taxes would cause some manufacturing firms to go out of business, these profit losses would be passed on to lower worker welfare.

Table 5: Robot Tax Incidence
(Discounted Present Values in Percent of GDP in 2019)

	Temporary Tax	Permanent Tax
Workers	-1.21	-1.00
Workers in 2019	-0.62	-0.47
– Manufacturing Production	0.02	0.01
Future Workers	-0.59	-0.53
Tax Revenues	2.39	9.41
Mechanical Effect	5.57	11.02
Behavioral Effect	-3.18	-1.61
Profits (excl. predatory externalities)	-4.14	-10.58

Note: *Workers* represent compensating variations; see Appendix G.1.1 for details. *Profits (excl. predatory externalities)* represent the effect on manufacturing firm values (Equations (7)-(8)) in 2019, holding constant pecuniary externalities of robot adoption in output markets; see Appendix G.2.1 for details. *Mechanical Effect* is the tax revenues collected if robot adoption did not respond to the tax.

In calculating the effects on firm profits in Table 5, I exclude so-called predatory investment externalities. Predatory investments refer to the pecuniary externality where robot adopters do not internalize that parts of the profit gain from robots come from stealing markets shares of competitor firms.²¹ By internalizing this predation effect, a robot tax has the possibility to increase aggregate profits of firms. To focus on the key equity-efficiency trade-off for workers, I hold the predatory externalities out of the baseline incidence calculations, and instead relegate their analysis to Appendix G.2.1.

To sum up, even though the temporary tax achieves the goal of delaying the diffusion of industrial robots, this analysis shows that the policy is an ineffective and relatively costly way to redistribute income to production workers employed in manufacturing.

²¹The implications of predatory investments have been studied extensively in the theoretical industrial organization literature, including Dixit (1980) and Spence (1986).

7 Conclusion

This paper makes two methodological contributions in order to study the distributional impact of industrial robots. First, I develop a dynamic firm model that can rationalize the selection into and reduced-form responses to robot adoption. Second, I model both firm and worker dynamics in general equilibrium. I use administrative data that link workers, firms, and robots in Denmark to structurally estimate a dynamic general equilibrium model that can account for event studies of firm robot adoption, the observed diffusion of industrial robots, and worker transitions in the labor market. The model fits the labor demand responses to robot adoption but also takes into account how production efficiency gains from robots are passed through to lower consumer prices as well as the ability of workers to reallocate between occupations in response to industrial robots.

Having validated the model using overidentification checks, I use it to estimate the distributional impacts of industrial robots. I find that industrial robots have increased average real wages by 0.8 percent but with substantial distributional consequences. At the ends of the spectrum, I find that production workers employed in manufacturing have lost 6 percent in real wages while tech workers have gained 2.3 percent.

The model captures worker heterogeneity in exposure to robot diffusion across occupation, industry, tenure, skill, and age of workers but abstracts from the possibility that robot adoption could differentially affect incumbent workers in the adopting firm. Using matched worker-firm-robot datasets to collect evidence on how firm robot adoption affects incumbent workers, as in Bessen et al. (2019), represents a promising avenue of future empirical research. Introducing such firm-specific wage or displacement effects into the general equilibrium framework developed in this paper without breaking the block separability that keeps the model computationally tractable is an important avenue of future theoretical research. I lay out one such model extension in the appendix of this paper.

I believe that the quantitative framework developed in this paper can be applied to studying the labor market impacts of other pressing technologies. For example, what will be the consequences when 1.3 million US truck drivers are expected to compete with self-driving vehicle technology by 2026 (Council of Economic Advisers, 2016)? The quantitative experiments conducted in this paper highlight that the ability of workers to switch occupations is

crucial for how a new technology can affect the distribution of earnings. In Humlum (2019), I find that retraining subsidies can be an effective tool to help workers transition across occupations in the labor market. These findings may help policymakers navigate in an era of rapid technological change.

References

- Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.
- Acemoglu, D. and P. Restrepo (2018a). Demographics and Automation. Technical report, National Bureau of Economic Research.
- Acemoglu, D. and P. Restrepo (2018b). Modeling Automation. *AEA Papers and Proceedings* 108, 48–53.
- Acemoglu, D. and P. Restrepo (2019a). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives* 33(2), 3–30.
- Acemoglu, D. and P. Restrepo (2019b). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy* (Forthcoming).
- Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification Properties of Recent Production Function Estimators. *Econometrica* 83(6), 2411–2451.
- Adda, J. and R. Cooper (2003). *Dynamic Economics: Quantitative Methods and Applications*. MIT press.
- Angrist, J. D. and J.-S. Pischke (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Arcidiacono, P. and R. A. Miller (2011). Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity. *Econometrica* 79(6), 1823–1867.
- Artuç, E., S. Chaudhuri, and J. McLaren (2010). Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *American Economic Review* 100(3), 1008–45.
- Ashournia, D. (2017). Labour Market Effects of International Trade When Mobility Is Costly. *The Economic Journal* 128(616), 3008–3038.
- Asker, J., A. Collard-Wexler, and J. De Loecker (2014). Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy* 122(5), 1013–1063.

- Autor, D. and D. Dorn (2013). The Growth of Low-Skill Service Jobs and the Polarization of the Us Labor Market. *American Economic Review* 103(5), 1553–97.
- Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992). International Real Business Cycles. *Journal of Political Economy* 100(4), 745–775.
- Bernard, A. B., T. C. Fort, V. Smeets, and F. Warzynski (2018). Heterogeneous Globalization: Offshoring and Reorganization. Technical report, Working Paper, Dartmouth College.
- Bernard, A. B., V. Smeets, and F. Warzynski (2017). Rethinking Deindustrialization. *Economic Policy* 32(89), 5–38.
- Berry, S., A. Gandhi, and P. Haile (2013). Connected Substitutes and Invertibility of Demand. *Econometrica* 81(5), 2087–2111.
- Bessen, J. E., M. Goos, A. Salomons, and W. Van den Berge (2019). Automatic Reaction – What Happens to Workers at Firms That Automate? *Boston Univ. School of Law, Law and Economics Research Paper*.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica* 77(3), 623–685.
- Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica* 87(3), 741–835.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics* 36(S1), S13–S70.
- Costinot, A. and I. Werning (2018). Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation. Technical report, National Bureau of Economic Research.
- Council of Economic Advisers (2016). Artificial Intelligence, Automation, and the Economy. *Executive Office of the President*, 18–19.
- Dauth, W., S. Findeisen, J. Suedekum, N. Woessner, et al. (2018). Adjusting to Robots: Worker-Level Evidence. Technical report.

- Delvaux, M. (2016). Draft Report with Recommendations to the Commission on Civil Law Rules on Robotics. *European Parliament Committee on Legal Affairs*.
- Diamond, P. A. and J. A. Mirrlees (1971). Optimal Taxation and Public Production I: Production Efficiency. *The American Economic Review* 61(1), 8–27.
- Dix-Carneiro, R. (2014). Trade Liberalization and Labor Market Dynamics. *Econometrica* 82(3), 825–885.
- Dixit, A. (1980). The Role of Investment in Entry-Deterrence. *The Economic Journal* 90(357), 95–106.
- Doraszelski, U. and J. Jaumandreu (2013). R&D and Productivity: Estimating Endogenous Productivity. *Review of Economic Studies* 80(4), 1338–1383.
- Doraszelski, U. and J. Jaumandreu (2018). Measuring the Bias of Technological Change. *Journal of Political Economy* 126(3), 1027–1084.
- Dubin, J. A. and D. L. McFadden (1984). An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica*, 345–362.
- Fort, T. C., J. R. Pierce, and P. K. Schott (2018). New Perspectives on the Decline of Us Manufacturing Employment. *Journal of Economic Perspectives* 32(2), 47–72.
- Goos, M., A. Manning, and A. Salomons (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review* 104(8), 2509–26.
- Graetz, G. and G. Michaels (2018). Robots at Work. *Review of Economics and Statistics* 100(5), 753–768.
- Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica*, 501–522.
- Hawkins, W., R. Michaels, and J. Oh (2015). The Joint Dynamics of Capital and Employment at the Plant Level. *Working Paper, Yale University*.
- Hotz, V. J. and R. A. Miller (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. *The Review of Economic Studies* 60(3), 497–529.

- Humlum, A. (2019). Retraining and Occupational Choice: Evidence from a Natural Experiment. *Working Paper, Princeton University*.
- Hummels, D., R. Jørgensen, J. Munch, and C. Xiang (2014). The Wage Effects of offshoring: Evidence from Danish Matched Worker-Firm Data. *American Economic Review* 104(6), 1597–1629.
- Imbens, G. and J. Wooldridge (2007). Estimation of Average Treatment Effects Under Unconfoundedness. *What's New in Econometrics, Lecture Notes 1*.
- International Federation of Robotics (2018). World Robotics Report 2018.
- King, R. G. and S. T. Rebelo (1999). Resuscitating Real Business Cycles. *Handbook of Macroeconomics* 1, 927–1007.
- Koch, M., I. Manuylov, and M. Smolka (2019). Robots and Firms. *CESifo Working Paper*.
- Lee, D. (2005). An Estimable Dynamic General Equilibrium Model of Work, Schooling, and Occupational Choice. *International Economic Review* 46(1), 1–34.
- Lee, D. and K. I. Wolpin (2006). Intersectoral Labor Mobility and the Growth of the Service Sector. *Econometrica* 74(1), 1–46.
- Levinsohn, J. and A. Petrin (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies* 70(2), 317–341.
- Market Watch (2019). Bill Gates Finds an Ally in Washington for His Idea to Tax Robots: Alexandria Ocasio-Cortez. *Market Watch*. March 11, 2019.
- Marschak, J. and W. H. Andrews (1944). Random Simultaneous Equations and the Theory of Production. *Econometrica*, 143–205.
- McLaren, J. (2017). Globalization and Labor Market Dynamics. *Annual Review of Economics* 9, 177–200.
- Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 1263–1297.

- Quartz (2017). The Robot That Takes Your Job Should Pay Taxes, Says Bill Gates. *Quartz Magazine*. February 17, 2017.
- Raval, D. R. (2019). The Micro Elasticity of Substitution and Non-Neutral Technology. *The RAND Journal of Economics* 50(1), 147–167.
- Region Syddanmark (2017). Robotter og Automatisering – Styrkepositioner, Udfordringer og Udviklingspotentiale. Technical report, Region Syddanmark.
- Rubin, D. B. (1990). Formal Mode of Statistical Inference for Causal Effects. *Journal of Statistical Planning and Inference* 25(3), 279–292.
- Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 999–1033.
- Spence, M. (1986). Cost Reduction, Competition and Industry Performance. In *New Developments In the Analysis of Market Structure*, pp. 475–518. Springer.
- Stokey, N. L. and R. Lucas (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters* 20(2), 177–181.
- Traiberman, S. (2019). Occupations and Import Competition. *American Economic Review* (Forthcoming).

A Data

A.1 Robot Adoption Firm Survey

Statistics Denmark conducts annually a technology adoption survey of firms in Denmark (IT usage in Danish enterprises, VITA). The survey is prepared in collaboration with the Danish Business Authority as a supplement to Eurostat’s technology survey. In 2018, the survey included a question on the use of industrial robots. The survey sampled 3,954 firms from the population of 16,465 private non-agricultural, non-financial firms with more than 10 employees. The response rate was 97 percent. Figure A.1 shows the questionnaire on industrial robot usage. Out of the survey respondents, a total of 473 firms answered ‘yes’ to using industrial robots in production.

Figure A.1: Firm Questionnaire on Firm Robot Adoption

Robot Technology

An industrial robot is an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications.

A service robot is a machine that has a degree of autonomy and is able to operate in complex and dynamic environment that may require interaction with persons, objects or other devices, excluding its use in industrial automation applications.

Software robots (computer programs) and 3D printers are out of the scope of the following questions.

18. Does your enterprise use any of the following types of robots?	Yes	No
- Industrial robots E.g. robotic welding, laser cutting, spray painting, etc.	<input type="radio"/>	<input type="radio"/>

A.2 Firm Customs Records

The firm customs records are organized in the Foreign Trade Statistics Register (UHDI) at Statistics Denmark. For each firm in each year 1993-2015, I have imports disaggregated by origin and 6-digit Harmonized System product code. One of these codes identifies “847950 Industrial Robots”.²² Industrial robots are heavily imported goods in Denmark (import share of 95 percent according to calculations in Section A.4), making customs records a valuable source of information on the adoption of industrial robots. The main challenge in using the customs records is that a substantial share of machinery is imported through domestic distrib-

²²Fort et al. (2018) use this product code to collect descriptive evidence on robot importers in the United States.

utors. Table A.3 develops a procedure for identifying robot imports done by final adopters.²³ Starting from the population of robot imports, I

1. *Pre-data coverage*: Restrict the sample to firms who are active three years before the import event. This condition is necessary for conducting the adoption event studies.
2. *Exclude wholesalers*: Exclude the one-digit industry code "5 Commerce".
3. *Exclude integrators*: Exclude 6-digit industry codes contained in a comprehensive list of robot integrators in 2018.²⁴
4. *Survey-validated adoptions*: Restrict the sample to 6-digit industries with a validation share in the robot adoption survey (VITA) of minimum 50 percent. The validation share is defined as the fraction of robot importers that in the robot adoption survey report that they use industrial robots.²⁵
5. *Single production establishment*: Restrict the sample to firms that only have a single establishment employing more than three workers in the year prior to robot adoption. This condition avoids dilution of the robot adoption effect in multi-plant firms (robot adoption happens at the plant level, but customs forms are filled out at the firm level).

The sample selection criteria exclude many of the robot import observations. For the sake of sustaining power in the statistical analysis, I use the 4-digit product code that includes industrial robots (HS 8479), as also done in Acemoglu and Restrepo (2018a).

²³I thank several industry experts for helpful inputs into developing this sample selection procedure, including Søren Peter Johansen (Technology Manager at the Danish Technological Institute, Robot Technology), Bo Hanfgarn Eriksen (Region Syddanmark), Per Rasmussen (BILA Robotics), and Martin Jespersen (Odense Robotics).

²⁴List of industry codes excluded: 51.60.00 Wholesale of machinery and equipment, 30.00.09 Manufacture of computer equipment, electric motors, etc., 29.40.09 Manufacturing of industrial machinery, 29.00.00 Manufacture of multi-purpose machines, 28.10.09 Manufacture of metal building materials. The list of robot integrators was developed by RoboCluster and Odense Robotics for the report Region Syddanmark (2017). I thank Bo Hanfgarn Eriksen at Region Syddanmark for providing the list.

²⁵List of industry codes included: 33.00.00 Manufacture of medical equipment, 30.00.09 Manufacture of computer equipment, electric motors, 29.30.00 Manufacture of agricultural machinery, 29.20.00 Manufacture of general purpose machinery, 29.10.00 Manufacture of ship engines, compressors, etc., 28.60.09 Manufacture of hand tools, metal packaging, etc., 28.10.09 Manufacture of metal building materials, 26.30.09 Brick, cement, and concrete industries, 25.00.00 Rubber and plastic industry, 24.40.00 Pharmaceutical Industry, 24.30.09 Manufacture of paints, soaps, cosmetics, etc., 24.10.09 Manufacture of chemical raw materials, 20.00.00 Wood industry, 15.89.09 Other food industry.

Table A.1: Identifying Robot Adoption in Customs Records

Step	Sample at End of Step		
	Imports (million USD)	Import events (firm-year)	Firms
Raw imports	3291.9	14355	4839
1. Pre-data coverage	1457.7	5936	2594
2. Exclude wholesalers	826.5	2016	1048
3. Exclude integrators	535.0	1375	754
4. Survey-validated industries	247.6	776	416
5. Single production establishment	91.1	454	293

A.3 Comparison of Data Sources on Robot Adoption

This section compares three data sources on robot adoption against each other: the robot adoption firm survey (Section A.1), the firm customs records (Section A.2), and the International Federation of Robotics (IFR) statistics. While the IFR statistics have been the main source of data for the existing papers on robotization (Acemoglu and Restrepo, 2019b; Dauth et al., 2018; Graetz and Michaels, 2018), the present paper is the first to use the adoption survey and the Danish customs records to study robot adoption. Section A.3.1 compares the industry representation across the three data sources, and Section A.3.2 examines how the time series of robot adoption compare in the different datasets.

A.3.1 Cross-Sectional Comparisons

Table A.2 shows that the industry composition of robot adoption in the micro data used in the present paper align well with the statistics compiled by the International Federation of Robotics. The data sources agree that industrial robots are a manufacturing technology, and that the metal, chemical, and plastic industries have been the main drivers of robot adoption.

Table A.2: Robot Adoption Across Industries: Comparison of Data Sources

	<i>Data Sources</i>		
	Robot Survey (StatDK)	Robot Stock (IFR)	Robot Imports (Customs)
<i>Share in Total Adoptions (%)</i>			
Manufacturing	79.1	85.9	83.5
<i>Share in Manufacturing Adoptions (%)</i>			
Food and beverages	7.2	18.3	7.2
Textiles	1.1	2.8	0.0
Wood and furniture	6.4	4.7	3.5
Paper	2.2	1.4	0.0
Plastic and Chemicals	14.0	22.0	32.3
Glass, stone, minerals	5.0	3.7	1.9
Metal	51.1	34.1	31.7
Electrical and Electronics	10.9	8.4	23.5
Automotives and vehicles	1.9	4.7	0.0

Note: “Robot Survey” indicates the share in total firm robot adopters. “Robot Stock” specifies the share in total robot stock. “Robot Imports” is the share in total firm robot import events (firm-year observations). Robot Imports represents the 454 adoption events identified in Table A.1.

A.3.2 Time Series Comparisons

The IFR statistics and the customs records each contain a time series dimension allowing me to compare how robot adoption has evolved according to the two data sources. Figure A.2 shows that total robot imports in Denmark (from custom records) have closely tracked the total number of robot installments (IFR statistics) since the 1990s.²⁶

²⁶The time series are normalized to 1 in 2010. This normalization implies a robot unit price of \$58,000, which falls within the range of common list prices for industrial robots (\$50,000 to \$100,000 according to the International Federation of Robotics (2018)).

Figure A.2: Robot Adoption Across Time: Comparison of Data Sources



A.4 Calculation of Robot Import Absorption Share

This section calculates the import share in robot adoption using micro data on re-exports and domestic producers. The robot import share is defined as

$$\text{Import Absorption Share}_t = \frac{\text{Imports Absorbed}_t}{\text{Imports Absorbed}_t + \text{Production for Domestic Absorption}_t}$$

I measure import absorption by summing over firms' robot imports, netting out their robot exports and transit trade. To measure the domestic supply (production for domestic absorption), I leverage the high export-orientation of robot producers to impute the domestic sales of robots. In particular, I use the customs records for the exports of robot producers (list provided by industry experts) to calculate the share of robots in total sales of the firm. I then multiply this robot share with the firm's domestic sales to impute the firm's domestic sales of robot.²⁷ Table A.3 shows that the robot import share has averaged 94.9 percent from 1993 to 2015.

²⁷Measuring the domestic supply (production for domestic absorption) is complicated by the fact that product code breakdowns of domestic sales in general do not exist.

Table A.3: Import Share in Robot Investments, Denmark 1993-2015 (percent)

Average	1993-2004	2005-2015
94.9	98.5	90.9

A.5 Measuring Domestically Sourced Robot Adoptions

This section describes how I supplement the customs records to measure robot adoption done through domestic distributors. I first use the representative robot adoption firm survey (VITA) conducted by Statistics Denmark; see Appendix A.1 for details. The survey provides a snapshot of which firms use industrial robots in 2018, regardless of whether the firms have imported their robots directly or have relied on a domestic distributor. From the adoption survey, I can directly calculate that 31 percent of manufacturing firms have adopted robots (last data point in Figure 3(a)) and that these adopters represent 54 percent of manufacturing sales (Figure 4).

For the time series of robot adoption, I use the International Federation of Robotics (IFR) statistics on the stock of industrial robots in Danish manufacturing over time (the data source of Acemoglu and Restrepo (2019b) and Graetz and Michaels (2018)). Assuming that the robot stock per adopter firm is constant over time, I can use the IFR time series to extend the number of robot adopters observed in 2018 back in time (Figure 3(a)). As a robustness check, I verify that the robot imports data imply the same evolution in total robot adoption over time.

A.6 Occupational Classification

I build on the occupational classification developed by Bernard et al. (2017) to study worker tasks. The classification groups detailed four-digit ISCO codes into six categories: managers, tech workers, sales workers, support workers, transportation/warehousing, and line workers (mostly production). The classification is used in Bernard et al. (2018).

Table 2 shows that the robot adopters and match firms are balanced on these occupational categories prior to adoption. In the main analysis, I focus on the three occupations that are most relevant to industrial robots: tech workers, production workers, and other workers. Tech workers is the second category of the Bernard et al. (2017) classification, and includes skilled technicians, engineers, and researchers. Production workers is the intersection of the sixth

category of Bernard et al. (2017) (line workers, mostly production) and the 1-digit ISCO88 code “7 Craft and Related Trades Workers.” Production workers consist of manual production tasks from welding to assembly.

A.7 Stylized Facts on Firm Robot Adoption

A.7.1 Matching Procedure

This section describes the matching algorithm used in column 3 of Table 2. The procedure is structured as follows.

1. Pick a vector X_e to match exactly on, and a vector $X_d \in \mathbb{R}^K$ to distance match on.
2. For each adopter firm f , find non-adopter match firm g that
 - (a) matches f exactly on X_e
 - (b) has minimal Mahalanobis distance to f in X_d

$$\text{Match}_f = \arg \min_{g \in \{X_e(f) \cap \text{na}\}} (X_{dg} - X_{df})' \Sigma (X_{dg} - X_{df}),$$

where Σ is the sample covariance matrix of X_d .

In my application, I match exactly (X_e) on industry (two-digit) and year $t - 1$. Within each industry-year bin, I then distance match (X_d) on firm sales and production line wage bill shares (levels at $t - 1$ and changes from $t - 3$).

B A Model of Firm Robot Adoption

B.1 Task-Based Micro Foundation for the Production Function

This section provides a task-based micro foundation for the production function used in Section 3.²⁸ Consider a firm j operating the task-based production technology

$$Y_{jt} = \left(\int_0^{\mathcal{I}_{jt}} Y_{jt}(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \quad Y_{jt}(i) = z_{ojt}(i) X_{ojt}(i) \mathbf{1}_{\{i \in \mathcal{A}_{ojt}\}}, \quad (39)$$

where production tasks are indexed by i and factors (production workers, tech workers, intermediate inputs, etc.) are indexed by o . Let $\mathcal{A}_{jt} = \{\mathcal{A}_{1jt}, \dots, \mathcal{A}_{\mathcal{O}jt}\}$ denote an assignment of tasks to factors (a partition of the interval $[0, \mathcal{I}_{jt}]$). Conditional on such a task assignment, the firm has to allocate the time of each factor across its assigned tasks. The first-order conditions to this time allocation problem are

$$X_{ojt}(i) = \frac{z_{ojt}(i)^{\sigma-1}}{\int_{i \in \mathcal{A}_{ojt}} z_{ojt}(i)^{\sigma-1} di} X_{ojt} \quad \text{for } i \in \mathcal{A}_{ojt}, \quad (40)$$

where $X_{ojt} = \int_{i \in \mathcal{A}_{ojt}} X_{ojt}(i) di$ is the total units of factor o employed at firm j in year t . By inserting Equation (40) into Equation (39), I can now represent the firm's technology with the production function

$$Y_{jt} = Y(X_{jt} | \mathcal{I}_{jt}, z_{jt}) = \left(\sum_{o \in \mathcal{O}} (z_{ojt} L_{ojt})^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \quad \text{with} \quad (41)$$

$$z_{ojt} = \frac{\int_{i \in \mathcal{A}_{ojt}} z_{ojt}(i)^{\sigma}_{ojt}}{\int_{i \in \mathcal{A}_{ojt}} z_{ojt}(i)^{\sigma-1}_{ojt}} \quad (42)$$

Following Acemoglu and Restrepo (2019a), suppose that firm robot adoption may

1. Affect each factor's productivity in a given task (the *productivity* effect), $z_{jt}(i, R_{jt})$
2. Require tasks to be reassigned between factors (the *substitution* effect), $\mathcal{A}_{jt}(R_{jt})$

²⁸The setup is inspired by Hawkins et al. (2015), who study the cost minimization problem of a plant operating a Ricardian task-based production technology where the assignment of productive factors to tasks is subject to a Calvo shock.

3. Create new tasks to be carried out in production (the *reinstatement* effect), $\mathcal{I}_{jt}(R_{jt})$.

I can then reformulate Equation (41) into a robot-contingent production function

$$Y_{jt} = Y(X_{jt} | R_{jt}, \omega_{jt}) = \left\{ \sum_{o \in \mathcal{O}} (z_{ojt} L_{ojt})^{\frac{\sigma-1}{\sigma}} di \right\}^{\frac{\sigma}{\sigma-1}}, \text{ with} \quad (43)$$

$$z_{ojt} = \exp(\varphi_{ojt} + \gamma_{ojt} R_{ojt}) \quad (44)$$

$$\varphi_{ojt} = \log \frac{\int_{i \in \mathcal{A}_{ojt}(0)} z_{ojt}(i, 0)^\sigma}{\int_{i \in \mathcal{A}_{ojt}(0)} z_{ojt}(i, 0)^{\sigma-1}} \quad (45)$$

$$\gamma_{ojt} = \log \frac{\int_{i \in \mathcal{A}_{ojt}(1)} z_{ojt}(i, 1)^\sigma}{\int_{i \in \mathcal{A}_{ojt}(1)} z_{ojt}(i, 1)^{\sigma-1}} - \log \frac{\int_{i \in \mathcal{A}_{ojt}(1)} z_{ojt}(i, 0)^\sigma}{\int_{i \in \mathcal{A}_{ojt}(0)} z_{ojt}(i, 0)^{\sigma-1}} \quad (46)$$

Equations (43)-(46) provide a direct micro foundation of the production function used in Equation (1). The only parametric restrictions imposed in Equation (1) are that of homogeneous robot productivity effects, $\gamma_{jt} = \gamma$.²⁹

C Structural Estimation of Firm Robot Adoption

C.1 Elasticity of Substitution Between Production Tasks

This section uses the model presented in Section 3 to derive the moment condition that I use to estimate the elasticity of substitution between production tasks σ in Section 4.1. The derivations follow closely those in Doraszelski and Jaumandreu (2018).

To derive the moment conditions, first insert Equation (13) into Equation (9) to express the deterministic component of firm productivities in terms of a non-parametric function of observables

$$\varphi_{ojt} = g_{ot}(\varphi_{ojt-1}, \dots, \varphi_{ojt-k}) + \xi_{ojt} \quad (47)$$

$$= g_{ot}(l_{ojt-1} - m_{jt-1} + \sigma(w_{ojt-1} - w_{Mjt-1}), \dots, l_{ojt-k} - m_{jt-k} + \sigma(w_{ojt-k} - w_{Mjt-k})) + \xi_{ojt} \quad (48)$$

$$= h_{ot}(l_{ojt-1} - m_{jt-1}, w_{ojt-1} - w_{Mjt-1}, \dots, l_{ojt-k} - m_{jt-k}, w_{ojt-k} - w_{Mjt-k}) + \xi_{ojt}, \quad (49)$$

²⁹The micro foundation in Equations (43)-(46) provides insights into the task-based sources of treatment effect heterogeneity in robot adoption. A promising avenue of further work is to use these expressions to empirically evaluate the task-based model predictions for heterogeneity in robot adoption treatment effects.

where lower-case letters denote log-transforms. Insert this function into Equation (11) to obtain

$$l_{o'jt} - l_{ojt} = -\sigma(w_{o'jt} - w_{ojt}) + (h_{o'jt} - h_{ojt}) + (\xi_{o'jt} - \xi_{ojt}), \quad (50)$$

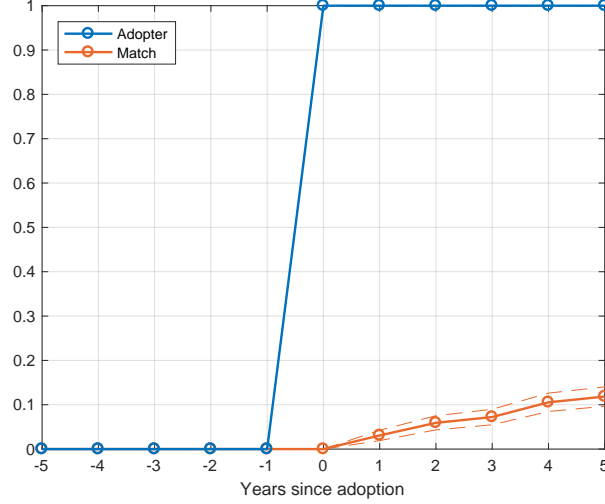
Equation (50) holds for firms that have not yet adopted robots. The Markovian structure on firm productivities implies that past factor choices l_{jt} and prices w_{jt} have to be uncorrelated with the current productivity innovations ξ_{jt} that constitute the error term in Equation (50). I can thus form a population moment condition that identifies σ , my parameter of interest

$$\mathbb{E}_t \left[A_{oo'}(Q_{jt-1}) \left(l_{o'jt} - l_{ojt} - \sigma(w_{o'jt} - w_{ojt}) + (h_{o't} - h_{ot}) \right) \right] = 0, \quad (51)$$

where $A_{oo'}$ is a vector function of the instruments Q_{t-1} including l_{jt-1}, w_{jt-1} . The instrument vector x_t consists of quadratic functions of $l_{jt-k} - m_{jt-k}$ and $w_{t-k} - w_{Mt-k}$ for $k = 1, 2, 3$, as well as quadratic functions of w_{jt-1} and l_{jt-1} (the excluded instruments). I set “Production Workers” and “Tech Workers” as o and o' , respectively, and I use “Other Workers” as the benchmark factor in production (M in the derivations above). I estimate (51) using a two-step GMM procedure (the `gmm` package in Stata).

C.2 Robot Technology

Figure C.1: Firm Robot Adoption Around the Event Year



Note: The figure shows separately the shares of firms in the treatment and control groups that have adopted robots around the event year.

C.2.1 Econometric Specification of the Event Studies

In this section, I describe the econometric specification that generates the matching-based event study estimates plotted in Figures 1 and 2. The estimates are differences-in-differences of outcomes y_{jt} for robot adopters versus match firms measured relative to the year prior to adoption.³⁰ Figures 1 and 2 plot OLS estimates of δ_k from the following specification

$$\frac{y_{jt}}{y_{jpre}^e} = \alpha \times \mathbb{R}_{je} + \sum_{k \in \mathcal{K}} \alpha_k \times \mathbb{1}_{\{t=e+k\}} + \sum_{k \in \mathcal{K} \setminus \{-1\}} \delta_k \times \mathbb{1}_{\{t=e+k\}} \times \mathbb{R}_{je} + u_{jt} \quad (52)$$

where e denotes event year, y_{jpre}^e are pre-event median outcomes, \mathbb{R}_{je} indicates that firm j adopted robots in year e , and $\mathbb{1}_{\{t=e+k\}}$ is an indicator that switches on iff event year e occurred k years ago. The event study window is denoted $\mathcal{K} = [-4, 4]$. Standard errors are clustered at the match level. I allow for zeros in occupational wage bills by calculating this relative difference as $(y_{jt}/w_{jpre})/(y_{pre}/w_{pre})$, where w_{jt} is the total wage bill of firm j in year t , and y_{pre} denotes mean pre-event outcomes.

³⁰The match firms are found using an Exact-Mahalanobis matching procedure described in Appendix A.7.1.

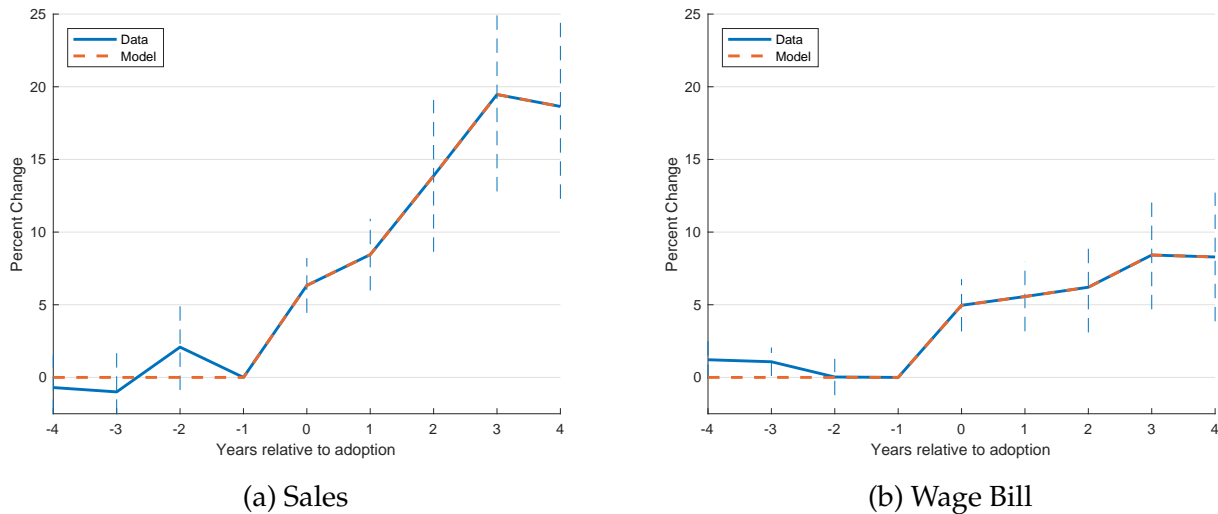
C.2.2 Robot Technology Distributed Lag Model

This section generalizes the robot technology equations (2)-(3) to account for the dynamic adjustments to robot production observed in Figures 1 and 2. I let robot technology follow a distributed lag model

$$\log(z_{jt}) = \varphi_{jt} + \sum_{\tau=0}^4 \gamma_{\tau} R_{jt-\tau} \quad (53)$$

Following the identification argument in Section 4.2.1, the adoption event study moments in Figures 1 and 2 exactly identify the dynamic robot technology parameters γ_{τ} . Figure C.2 shows the model fit for firm sales and wage bills.

Figure C.2: Distributed Lag Model for Robot Productivities




C.3 Baseline Technology

C.3.1 Labor-Augmenting Baseline Productivities


The general equilibrium model restricts the labor-augmenting baseline productivities to a common time-varying parameter vector. I calibrate this path of common productivities γ_{ot} to match the aggregate factor shares in manufacturing taking into account the diffusion of robot technology. Figure C.3 shows data (dots) and model simulations (line) from 1990 to 2018 together with out-of-sample forecasts from 2019 to 2049. The data have been HP-filtered to focus

on medium-run movements (smoothing parameter of 100 following Backus et al. (1992)). The forecasts extrapolate the growth rate from 2011 to 2018, assuming a linear reduction in rates of growth to zero by 2049.

Figure C.3: Aggregate Factor Shares in Manufacturing Production



(a) Wage Bill Shares



(b) Wage Bill Share in Manufacturing Sales

C.3.2 Hicks-Neutral Baseline Productivities

With the homogeneity restriction imposed on firm baseline labor-augmenting productivities, the productivity process in Equation (21) boils down to an AR(1) process for the Hicks-neutral term φ_{Hjt} .

$$\varphi_{Hjt} = \mu_{zt} + \rho_z \varphi_{Hjt-1} + \sigma_z \epsilon_{jt}, \tag{54}$$

where ρ_z is the persistence parameter for baseline productivity, and ψ_t is a time fixed effect.

Table C.1: Baseline Productivity Parameters

Parameter	Description	Estimated Value
$\hat{\rho}_z$	Persistence of firm productivity	0.901 (0.062)
$\hat{\sigma}_z$	Standard deviation of firm productivity innovations	0.140

I discretize the estimated AR(1) process using the Tauchen (1986) method.

C.4 Robot Adoption Costs

C.4.1 Method of Simulated Moments (MSM) Estimator

In this section, I describe the method of simulated moments (MSM) estimation procedure adopted in Section 4.4. Table C.2 reports the MSM parameter estimates.

1. Parameterize robot adoption costs to be log-linear in time: $c_t^R = \exp(c_0^R + c_1^R \times t)$.
2. Stack the robot adoption cost parameters into the parameter vector $\theta = (c_0^R, c_1^R, \nu)'$.
3. Stack the robot diffusion curve and the adopter size premium into the moment vector $\pi \in \mathbb{R}^N$ with $N = 2018 - 1990 + 2$
4. Define a grid on the parameter space Θ . For each point on the grid $\theta^{(j)} \in \Theta$,
 - (a) Solve for continuation values given $c_t^R = \exp(c_0^{(j)} + c_1^{(j)} \times t)$ and $\nu = \nu^{(j)}$. The solution algorithm is specified in Section F.1.
 - (b) Simulate firms forward using policy functions.
 - (c) Calculate the in-sample squared deviations between the simulated and observed moment vectors

$$(\pi_S(\theta^{(j)}) - \pi_D)'W(\pi_S(\theta^{(j)}) - \pi_D) \tag{55}$$

where W is the identity weighting matrix.

5. The MSM estimator, $\hat{\theta}$, attains the minimum in (55).

Table C.2: Robot Adoption Cost Parameters (MSM)

Parameter	Description	Estimate
c_0^R	Intercept of the common adoption cost schedule over time	1.155
c_1^R	Slope of the common adoption cost schedule over time	-0.026
ν	Dispersion in idiosyncratic adoption costs	0.384

C.4.2 Variance-Covariance Matrix of the MSM Estimator

I calculate the variance-covariance matrix of the MSM estimator $\hat{\theta}$ using the formula on pages 88 and 89 of Adda and Cooper (2003). The MSM estimator has the asymptotic distribution

$$\sqrt{N}(\theta - \theta_0) \xrightarrow{d} \mathcal{N}(0, V) \quad \text{with} \quad (56)$$

$$V = \left[\mathbb{E}_0 \frac{\partial \pi'}{\partial \theta} W^{-1} \frac{\partial \pi}{\partial \theta'} \right]^{-1} \mathbb{E}_0 \frac{\partial \pi'}{\partial \theta} W^{-1} \Sigma(\theta_0) W^{-1} \frac{\partial \pi}{\partial \theta'} \times \left[\mathbb{E}_0 \frac{\partial \pi'}{\partial \theta} W^{-1} \frac{\partial \pi}{\partial \theta'} \right]^{-1}, \quad (57)$$

I estimate $\mathbb{E}_0 \frac{\partial \mu'}{\partial \theta}$ using numerical derivatives of the simulated moments around $\hat{\theta}$. The confidence bands in Figure 3 are calculated using the delta method.

C.4.3 Comparison of Robot Adoption Cost Estimates

Table C.3 compares the MSM estimate for the rate of change in the common component of robot adoption costs c_1^R to the robot machine expenditures reported on customs forms of adopting firms. I report the time-slope estimates of log expenditures and log mean expenditures. As the table shows, the MSM estimate of c_1^R falls within the confidence bands of both specifications.

Table C.3: Rate of Change in Robot Adoption Costs: Model Estimates vs. External Measures

	Lower Bound (95% CI)	Point Estimate	Upper Bound (95% CI)
MSM Estimate (\hat{c}_1^R)		-0.0264	
Customs Expenditures 1	-0.1158	-0.0693	-0.0229
Customs Expenditures 2	-0.0661	0.0179	0.1019

Note: The first row is the MSM estimate of \hat{c}_1^R . The second row (Customs Expenditures 1) is the OLS estimate of β_1 in $\log(Y_{jt}) = \beta_0 + \beta_1 t$ (reweighted to the yearly level). The third row (Customs Expenditures 2) is the OLS estimate of β_1 in $\log(\tilde{Y}_t) = \beta_0 + \beta_1 t$. I deflate the customs expenditures with the consumer price index.

C.5 Depreciation of Robot Technology

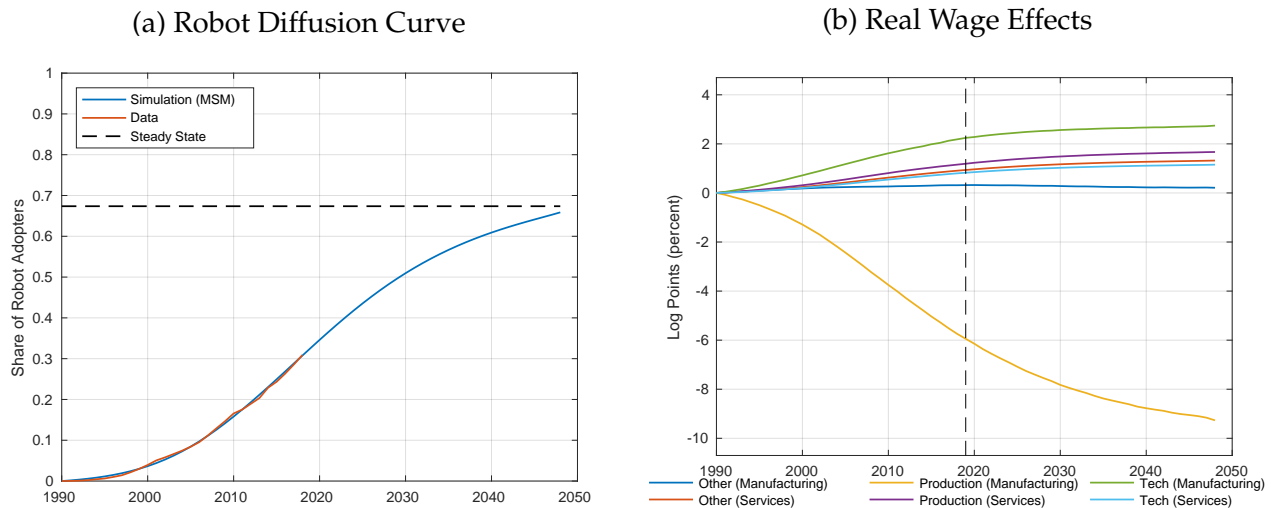
This section derives a model extension in which robot technology deteriorates with a probability θ . The Bellman equation for robot adoption now reads

$$V_t(0, \varphi) = \max_{R \in \{0,1\}} \pi_t(0, \varphi) - (c_t^R + \varepsilon_{jt}^R) \times R + \beta \mathbb{E}_t V_{t+1}(R, \varphi') \quad (58)$$

$$V_t(1, \varphi) = \pi_t(1, \varphi) + (1 - \theta) \mathbb{E}_t V_{t+1}(1, \varphi') + \theta \mathbb{E}_t V_{t+1}(0, \varphi') \quad (59)$$

Equations (58)-(59) collapse to the current setup in Equations (7)-(8) if $\theta = 0$. Figure C.4 shows the simulated robot diffusion curve and real wage effects on industrial robots under a robot depreciation rate θ of 10 percent (the depreciation rate used in Graetz and Michaels (2018)). Compared to baseline Figures 3a and 6, the model extension to robot depreciation does not affect the in-sample estimate of the real wage effects of industrial robots as the extended model is estimated to match the same observed robot diffusion curve. The model extension does alter the long-run predictions, however, as the robot diffusion curve asymptotes to a long-run steady-state level below full adoption when robot technology deteriorates (dashed line in Figure C.4a).

Figure C.4: Effect of Industrial Robots with Depreciation of Robot Technology



D The Labor Supply Block

D.1 Data on Worker Transitions

This section describes how I measure the worker transitions used to estimate the labor supply model. I follow Traiberman (2019) as closely as possible; please refer to his Appendices B.5 and D for additional details.

I use the Integrated Database for Labor Market Research (IDA), which links every worker to her employer in the month of November. The data contain information about the occupation, salary, and sociodemographics of workers. I follow the recommendation of Statistics Denmark, and use only the high-quality occupational codes that come from administrative registers or pension funds (DISCOTYP 1, 2, 4 and 10). If a worker is employed in the same occupation in year $t - 1$ and $t + 1$ but has missing data at t , I impute the occupation as that at $t - 1$. If a worker switches firms between year t and $t + 1$ and there are data on occupation in $t + 2$, I impute any missing occupational codes at $t + 1$. As the measure of salary $w_{ot}H_{oit}$, I use the variable `joblon` together with a correction for labor market pension contributions from Hummels et al. (2014). The fine level of disaggregation into worker types and occupations implies that the observed transition matrix $\pi_t(s, a, \text{ten}, o, o')$ has zero elements. To avoid bias from dropping these zero elements, I follow Traiberman (2019) and perform the following first-stage smoothing regression of the transition probabilities.

$$\pi_t(s, a, \text{ten}, o, o') = \beta_{soo't}^0 + \beta_{soo't}^a a + \beta_{soo't}^{a2} a^2 + \beta_{soo't}^{a3} a^3 \quad (60)$$

$$+ \beta_{soo't}^{\text{ten}} \text{ten} + \beta_{soo't}^{\text{ten}2} \text{ten}^2 + \beta_{soo't}^{\text{ten}3} \text{ten}^3 + \beta_{soo't}^{\text{ten},a} a \times \text{ten} \quad (61)$$

I then use the predicted transition rates $\hat{\pi}_t$ (bounded by 10^{-6} and $1 - 10^{-6}$) as input variable into the estimating regression equations (31)-(33).

D.2 Estimation of the Labor Supply Model

Table D.1: Human Capital Function

	Tech (services)	Tech (manuf)	Production (services)	Production (manuf)	Other (services)	Other (manuf)
Age β_1^o	0.0285 (0.0010)	0.0265 (0.0005)	0.0096 (0.0006)	0.0055 (0.0006)	0.0124 (0.0007)	0.0139 (0.0010)
Age-Squared β_2^o	-0.0590 (0.0016)	-0.0543 (0.0013)	-0.0236 (0.0011)	-0.0171 (0.0014)	-0.0266 (0.0014)	-0.0301 (0.0023)
Tenure β_3^o	0.0300 (0.0018)	0.0153 (0.0010)	0.0277 (0.0012)	0.0234 (0.0016)	0.0537 (0.0030)	0.0307 (0.0012)
Mid Skill β_M^o	-0.0428 (0.0015)	0.0028 (0.0028)	0.1025 (0.0015)	0.1168 (0.0025)	0.0537 (0.0012)	0.1165 (0.0018)
High Skill β_H^o	0.1671 (0.0016)	0.2958 (0.0022)	0.0997 (0.0103)	0.1629 (0.0061)	0.2502 (0.0037)	0.5108 (0.0053)
Observations	2147314	602741	1029836	681133	17176380	2780515

Note: SD of income shock: Tech (services): .118, Tech (manufacturing): .077, Production (services): .096, Production (manufacturing): .077 Others (services): .148, Others (services): .133. Standard errors are clustered at the occupation-year level. Coefficient on Age Squared is presented $\times 10^2$.

D.2.1 Occupational Switching Costs

The non-linear least squares objective function (NLLS) reads

$$\min_{\{c, \rho\}} \sum_{\omega, \rho, t} \left[\log \frac{\pi_t(oo'|\omega)}{\pi_t(oo|\omega)} + \beta \log \frac{\pi_{t+1}(o'o''|\omega')}{\pi_{t+1}(oo''|\omega'')} - \left(\frac{1}{\rho} c_{oo'}(\omega) - \frac{\beta}{\rho} (c_{o'o''}(\omega') - c_{oo''}(\omega'')) + \frac{\beta}{\rho} (w_{o't+1} H_{o'}(\omega') - w_{ot+1} H_o(\omega'')) + \frac{\beta}{\rho} (\eta_{o'} - \eta_o) \right) \right]^2 \quad (62)$$

I use the Matlab solver `lsqnonlin` to estimate $c(\omega)$ and ρ in Equation (62). Table D.2 presents the estimated bilateral occupational switching costs, and Table D.3 presents the remaining switching cost estimates.

Table D.2: Bilateral Switching Costs $c_{oo'}/\rho$

	Tech (serv)	Tech (manuf)	Production (serv)	Production (manuf)	Other (serv)	Other (manuf)
Tech (services)	0	4.71	1.7	5.04	1.17	4.64
Tech (manufacturing)	0.76	0	3.49	0.58	2.51	0.01
Production (services)	5.9	11.12	0	2.73	3.78	6.6
Production (manufacturing)	9.24	8.79	2.75	0	6.35	4.28
Other (services)	3.8	8.44	1.87	3.9	0	2.97
Other (manufacturing)	6.6	5.94	3.68	1.27	2.19	0

Table D.3: Switching Cost Parameters

Parameter	Description	Estimate
α_1	Semi-elasticity of switching costs with respect to age (linear term) [‡]	13.86
α_2	Semi-elasticity of switching costs with respect to age (quadratic term) [‡]	-0.14
α_M	Semi-elasticity of switching cost with respect to mid skill	0.01
α_H	Semi-elasticity of switching cost with respect to high skill	0.00
ρ	Occupational preference shock variance [†]	2.00

Note: [‡] Coefficients of age polynomial are presented $\times 10^3$. [†]Parameter value of ρ used in Section 6.

D.2.2 Occupational Amenities

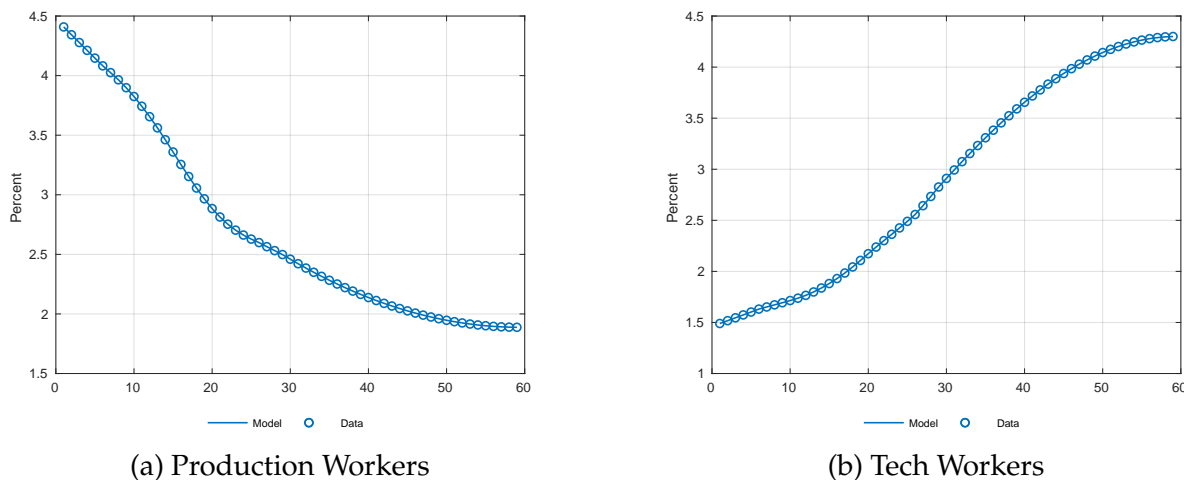
The employment shares across occupations are tightly connected to the vector of occupational amenities, η_{ot} . Conditional on the distribution of workers in $t - 1$, the relative change in employment shares, $\hat{s}_{ot} = \frac{s'_{ot}}{s_{ot}}$ from a different occupational amenity η'_{ot} is given by

$$\hat{s}_{o't} = \frac{\exp(\frac{1}{\rho}\hat{\eta}_{o't})}{\sum_{o''} s_{o''t} \times \exp(\frac{1}{\rho}\hat{\eta}_{o''t})}. \quad (63)$$

Of course, when evaluating the total effect of η_{ot} on s_{ot} , I cannot take the distribution of worker states, μ_{t-1}^W , as given because workers also choose occupations before $t - 1$ in anticipation of the amenities in year t . In practice, I estimate the path of occupational amenities η_{ot} by matching simulated occupational employment shares to the data. The estimation procedure is a fixed-point shooting algorithm.

Figure D.1 shows data (dots) and model simulations (line) for the share of employment across two example occupations from 1990 to 2018 together with out of sample forecasts from 2019 to 2049. The data have been HP-filtered to focus on medium-run movements (smoothing parameter of 100 following Backus et al. (1992)). The forecasts extrapolate growth rates from 2011 to 2018 by assuming a linear reduction in rates of growth to zero by 2049.

Figure D.1: Employment Shares Across Occupations (Manufacturing)



E Model Extension to Firm-Specific Wages

This section proposes an extension of the model in Sections 3 and 5 that accommodates firm-specific wages in equilibrium. On the worker side, I embed a random utility model for workplace environments as in Card et al. (2018) into the dynamic discrete occupational choice model of Section 5. The heterogeneity in worker preferences for employers implies that firms face upward-sloping labor supply functions. I derive the cost, profit, and factor demand functions of the firm, and I clarify how to implement these expressions into the remaining model structure.

E.1 The Worker's Problem

The worker first chooses which occupation o' to work in next period. Upon arriving in the chosen occupation next period, the worker then chooses which employer $j \in \mathcal{J}_{o'}$ to work for. This job choice model is a nested logit with occupations in the upper nest and firms constituting the lower nest. The Bellman equation of the worker reads

$$v_t(o, s, \text{age}, \text{ten}) = \max_{j \in \mathcal{J}_o} \{ \log(w_{ojt} H_o(s, \text{age}, \text{ten})) + \log(a_{ojt}) + \varepsilon_{ojt} \} \quad (64)$$

$$+ \max_{o' \in \mathcal{O}} \left\{ -(c_{oo'} + \varepsilon_{o'}) + \mathbf{1}_{\{\text{age} < 65\}} \beta \mathbb{E}_t v_{t+1}(o', s, \text{age} + 1, \mathbf{1}_{\{o'=o\}}(\text{ten} + 1)) \right\}, \quad (65)$$

where a_{ojt} is a firm-specific amenity common to all workers, and $\varepsilon_{ojt} \stackrel{iid}{\sim} \text{GEV1}(\alpha)$ are idiosyncratic workplace preference shocks that are drawn from a Gumbel distribution and only realized once a worker arrives in the occupational labor market. The parameter α measures the dispersion in these idiosyncratic preference shocks. The probability that a worker in occupation o with characteristics ω chooses to work for firm j is

$$P(j_t(\omega, o) = j | \omega) = \frac{(a_{ojt} w_{ojt} H_o(\omega))^{1/\alpha}}{\sum_{j' \in \mathcal{J}_o} (a_{oj't} w_{oj't} H_o(\omega))^{1/\alpha}} \quad (66)$$

E.2 The Firm's Problem

The firm faces the labor supply curve

$$L_{ojt}(w_{ojt}) = \int \frac{(a_{ojt} w_{ojt} H_o(\omega))^{1/\alpha}}{\sum_{j' \in \mathcal{J}_o} (a_{oj't} w_{oj't} H_o(\omega))^{1/\alpha}} H_o(\omega) dF_t^W(\omega) \quad (67)$$

If individual firms each constitute a negligible share of the total occupational labor market, the inverse labor supply curve to the firm becomes

$$w_{ojt} = \frac{w_{ot}}{a_{ojt}} \times L_{ojt}^\alpha \quad (68)$$

with $w_{ot} = \left\{ \int \frac{H_o(\omega)^{1/\alpha}}{\sum_{j' \in \mathcal{J}_o} (a_{oj't} w_{oj't} H_o(\omega))^{1/\alpha}} dF_t^W(\omega) \right\}^{-\alpha}$.³¹ To ease the exposition, stack the static inputs (including intermediate inputs) into the vector L , and reparameterize, without loss of generality, the CES production function as follows

$$Y_{jt} = F(L_{jt} | z_{jt}) = \left\{ \sum_o z_{ojt}^{\frac{1}{\sigma}} L_{ojt}^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}} \quad (69)$$

The firm's profit maximization problem reads

$$\pi_t(z_{jt}) = \max_{L_{jt}} \left\{ P_{Mt} Y_{Mt}^{\frac{1}{\epsilon}} F(L_{jt} | z_{jt})^{1-1/\epsilon} - \sum_o \frac{w_{ot}}{a_{ojt}} L_{ojt}^{1+\alpha} \right\}, \quad (70)$$

³¹If firms are price takers in intermediate input (M) markets, then the specification should be amended to allow for $\alpha \neq \alpha_M = 0$.

The first-order conditions for cost minimization imply that

$$\log(L_{o'jt}) - \log(L_{ojt}) = -\frac{\sigma}{1+\alpha\sigma}(\log(w_{o'jt}) - \log(w_{ojt})) \quad (71)$$

$$+ \frac{1}{1+\alpha\sigma} \{ \sigma \log(a_{ojt}/a_{o'jt}) + \log(z_{o'jt}/z_{ojt}) \}, \quad (72)$$

which clarifies that σ in Equation (11) is estimated off the variation in firms' relative amenities $\frac{a_{ojt}}{a_{o'jt}}$ once I in Section 4.1 control for the evolution in firm productivities $\frac{z_{ojt}}{z_{o'jt}}$. Equation (71) shows that the GMM estimate $\hat{\sigma}$ is attenuated toward zero if workers have heterogeneous preferences for workplaces ($\alpha > 0$).

The cost function of the firm is (after some algebra)

$$C_t(Y_{jt}, z_{jt}) = \min_{L_{jt}} \sum_o w_{ojt} L_{ojt} \quad \text{s/t} \quad F(L_{jt}|z_{jt}) \geq Y_{jt} \quad (73)$$

$$= Y_{jt}^{1+\alpha} \left\{ \sum_o (z_{ojt}^{1+\alpha} \left(\frac{w_{ojt}}{a_{ojt}} \right)^{1-\sigma})^{\frac{1}{1+\alpha\sigma}} \right\}^{\frac{1}{1-\sigma}} = Y^{1+\alpha} \Omega_{jt}, \quad (74)$$

where I denote $\Omega_{jt} = \{ \sum_o (z_{ojt}^{1+\alpha} \left(\frac{w_{ojt}}{a_{ojt}} \right)^{1-\sigma})^{\frac{1}{1+\alpha\sigma}} \}^{\frac{1}{1-\sigma}}$. The profit-maximizing output level is

$$Y_{jt} = \left\{ Y_{Mt} P_{Mt}^\epsilon \left((1+\alpha) \frac{\epsilon}{(\epsilon-1)} \Omega_{jt} \right)^{-\epsilon} \right\}^{\frac{1}{1+\alpha\epsilon}} \quad (75)$$

The firm's factor demands are

$$L_{ojt} = \frac{\{ z_{ojt} \left(\frac{w_{ojt}}{a_{ojt}} \right)^{-\sigma} \}^{\frac{1}{1+\alpha\sigma}}}{\Omega_{jt}^{-\sigma}} \times Y_{jt} \quad (76)$$

The profit function is

$$\pi_t(z_{jt}) = \left(\frac{1-\alpha\epsilon}{\epsilon-1} \right) \left\{ \Omega_{jt}^{1-\epsilon(1-\alpha)} Y_{Mt} P_{Mt}^\epsilon \left((1+\alpha) \frac{\epsilon}{\epsilon-1} \right)^{-\epsilon} \right\}^{\frac{1}{1+\alpha\epsilon}} \quad (77)$$

Equations (71)-(77) simplify to the usual CES expressions when worker preferences for workplaces vanish, $\alpha \rightarrow 0$.

The flow profit function (77) can be directly plugged into the Bellman equation (7) of the

adoption model in Section 3. The firm factor demands (76) can be readily aggregated up when clearing labor markets in the general equilibrium model of Section 6.1.

F Solution Algorithms

This section provides details on the solution algorithms used in Sections 4, 5, and 6.

F.1 Solving the Firm's Problem

This section details the algorithm for solving the firm's dynamic programming problem of robot adoption.

1. Set a time horizon, T , sufficiently far in the future such that robots are fully diffused and robot adoption costs are stationary by then (I set $T = 2050$ in practice).
2. Start at T . Solve the stationary, infinite horizon dynamic programming problem by iterating on the expected value functions until convergence.

$$\mathbb{E}V_T^{(j+1)}(1, \varphi) = \pi_T(1, \varphi) + \beta \sum_{z'} p(\varphi' | \varphi) \mathbb{E}V_T^{(j)}(1, \varphi') \quad (78)$$

$$\mathbb{E}V_T^{(j+1)}(0, \varphi) = \pi_T(0, \varphi) + \beta \sum_{z'} p(\varphi' | \varphi) \nu \log \left\{ \exp\left(\frac{1}{\nu}(-c_T^R + \beta \mathbb{E}V_T^{(j)}(1, \varphi'))\right) + \exp\left(\frac{1}{\nu} \beta \mathbb{E}V_T^{(j)}(0, \varphi')\right) \right\}, \quad (79)$$

where I used the log-sum expression for the expected maximum (EMAX) function.³² Convergence of Equation (79) in the unique fixed point $\mathbb{E}V_T(R, \varphi)$ is ensured from Blackwell's sufficient conditions for contraction mappings (Stokey and Lucas, 1989, Theorem 4.6).

3. Solve for $\{\mathbb{E}V_t(R, \varphi)\}_{t=t_0}^{T-1}$ using backward recursion from $T - 1$ to the initial period t_0 .

$$\mathbb{E}V_t(1, \varphi) = \pi_t(1, \varphi) + \beta \sum_{z'} p(\varphi' | \varphi) V_{t+1}(1, \varphi') \quad (80)$$

$$\mathbb{E}V_t(0, \varphi) = \pi_t(0, \varphi) + \beta \sum_{z'} p(\varphi' | \varphi) \nu \log \left\{ \exp\left(\frac{1}{\nu}(-c_t^R + \beta \mathbb{E}V_{t+1}(1, \varphi'))\right) + \exp\left(\frac{1}{\nu} \beta \mathbb{E}V_{t+1}(0, \varphi')\right) \right\} \quad (81)$$

4. From the initial year t_0 , use policy functions to simulate firms forward. Verify that the robot adoption share is 1 at time T .

³²Note that my setup with a logit shock for adoption (Equation (22)) is isomorphic to the setup in Rust (1987) with Gumbel shocks for both adoption and non-adoption (up to a recentering for the mean of a Gumbel). This is due to the well-known result that the difference between two Gumbels is logistically distributed.

In solving steps 3 and 4, I assume that firms have perfect foresight with respect to the in-sample path of wages, and use regressions to forecast these aggregate state variables out of sample.

F.2 Solving the Worker's Problem

This section details the algorithm for solving the worker's dynamic occupational choice problem.

1. Set a time horizon, T , sufficiently far in the future such that robots are fully diffused by then (I set $T = 2050$ in practice).
2. Start at T . Solve the stationary worker value functions
 - (a) Start at age of retirement. The value function is

$$\mathbb{E}_{\varepsilon, \zeta} V_T(o, 65, \omega) = \log(w_{oT} H_{oT}(65, \omega)) + a_{oT}. \quad (82)$$

- (b) Solve the value function for ages $a = 64, \dots, 25$ by backward recursion

$$\mathbb{E}_{\varepsilon, \zeta} V_T(o, a, \omega) = \log(w_{oT} H_{oT}(a, \omega)) + a_{oT} + \rho \left[\gamma + \log \left\{ \sum_{o'} \exp \left(\frac{1}{\rho} (c_{oo'}(\omega) + \beta \mathbb{E}_{\varepsilon, \zeta} V_T(o', a+1, \omega')) \right) \right\} \right], \quad (83)$$

where $\gamma = 0.577$ is Euler's constant.

3. Calculate the worker value functions for $t = T - 1, \dots, t_0$ using backward recursion

$$\mathbb{E}_{\varepsilon, \zeta} V_t(o, 65, \omega, \zeta) = \log(w_{ot} H_{ot}(65, \omega)) + a_{ot} \quad (84)$$

$$\mathbb{E}_{\varepsilon, \zeta} V_t(o, a, \omega, \zeta) = \log(w_{ot} H_{ot}(a, \omega)) + a_{ot} + \rho \left[\gamma + \log \left\{ \sum_{o'} \exp \left(\frac{1}{\rho} (c_{oo'}(\omega) + \beta \mathbb{E}_{\varepsilon, \zeta} V_{t+1}(o', a+1, \omega')) \right) \right\} \right], \quad (85)$$

In solving this dynamic program, I assume that workers have perfect foresight with respect to the in-sample path of wages, and use regressions to forecast these aggregate state variables out-of-sample.

F.3 Solving the Dynamic General Equilibrium

This section describes the algorithm for solving the general equilibrium featuring the two-sided dynamics defined in Section 6.1. A key property of the general equilibrium model is

that, despite the rich worker and firm heterogeneity, the only aggregate state variables that agents need to keep track of to solve their dynamic programming problem is the path of wages and the manufacturing price index.³³ I use a fixed-point shooting algorithm that solves for the wage path that clears labor markets given the optimal policy functions of workers and firms.

1. Guess a path of wages $w_t^{(0)}$ and manufacturing price index P_{Mt}^0 .
2. Solve for firm and worker continuation values (see Appendices F.1 and F.2).
3. Simulate firm and worker states forward using the policy functions from Step 2.
4. Find wages, $w_t^{(e)}$, that clear labor markets for each occupation period by period (using the firms' static labor demand conditions from Equation (5)). Calculate the implied manufacturing price index $P_{Mt}^{(e)}$.
5. Update wages and manufacturing price index

$$w_t^{(j+1)} = \lambda w_t^{(j)} + (1 - \lambda) w_t^{(e)} \quad (86)$$

$$P_{Mt}^{(j+1)} = \lambda P_{Mt}^{(j)} + (1 - \lambda) P_{Mt}^{(e)} \quad (87)$$

where $\lambda \in [0.8, 0.95]$ is the relaxation parameter in the Gauss-Seidel update.

6. Iterate until convergence in $\{w_t, P_{Mt}\}_t$.

³³The path of wages are sufficient to solve the worker's problem. Manufacturing firms also need to keep track of the manufacturing output price index as it summarizes the competitive pressures from robot adoption.

G Counterfactual Experiments

Table G.1: Parameters of the General Equilibrium Model

Description	Related Moments	Time varying
<i>Manufacturing Firms</i>		
c_t^R Common robot adoption costs	Robot diffusion curve (Figure 3)	✓
v Variance of idiosyncratic adoption costs	Size premium in robot adoption (Figure 4)	
γ_o Labor-augmenting robot productivity	Robot adoption event studies (Figures 1-2)	
γ_H Hicks-neutral robot productivity [†]	Robot adoption event studies (Figures 1-2)	
σ Elasticity of task substitution	Rational expectations GMM (Table 3)	
μ_H Mean of Hicks-neutral baseline productivity	Real wage index	✓
ρ_H Persistence of Hicks-Neutral productivity	Firm sales dynamics (Table C.1)	
σ_H Standard deviation of Hicks-Neutral innovations	Firm sales dynamics (Table C.1)	
φ_{ot} Baseline labor-augmenting productivities	Labor shares in manufacturing sales (Figure C.3)	✓
<i>Workers</i>		
β Human capital parameters	Mincer regression (Table D.1)	
$c_{oo'}$ Occupational switching costs	Occupational transition rates (Tables D.2-D.3)	
η_{ot} Occupational amenities	Employment shares across occupations and sectors (Figure D.1)	✓
<i>Services Production</i>		
α^s Cobb Douglas shares in services production	Wage bill shares in sales excl. manufacturing	
z_{st} Hicks-Neutral productivity in services	Real wage index	✓
<i>Common Parameters</i>		
β Discount factor	Interest rate of 4%	
μ Cobb-Douglas shares in final output	Share of manufacturing in total output	
ϵ Elasticity of manufacturing demand	Markup of 1/3 (Bloom, 2009)	

Notes: [†]I calibrate the path of γ_{Ht} to hold the sales elasticity with respect to robot adoption (Figure 1(a)) constant over time, given the estimated non-stationarity path of baseline labor-augmenting productivities γ_{ot} .

G.1 The Distributional Impacts of Industrial Robots

Figure G.1: The Effect of Industrial Robots on the Labor Share in Manufacturing Sales


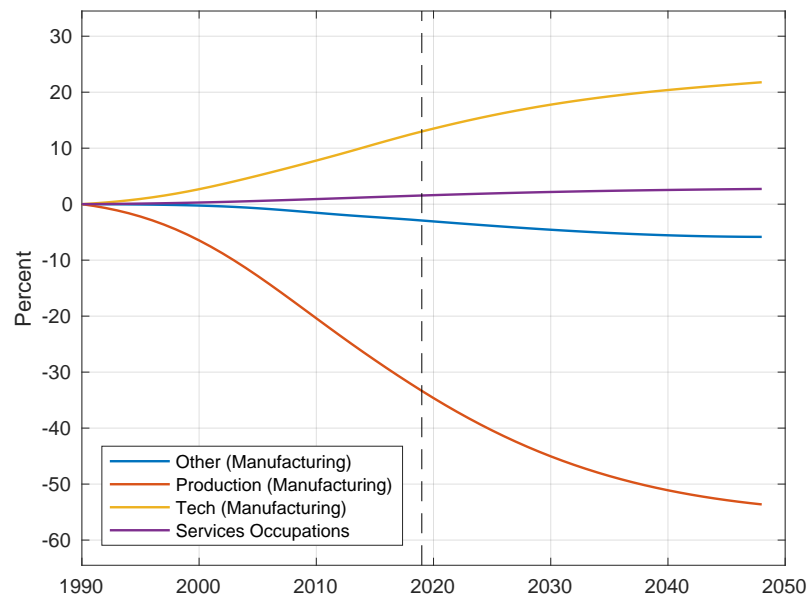


Figure G.2: Real Wage Effects of Industrial Robots with Exogenous Labor Supply



G.1.1 Compensating Variations

To measure the effects on worker welfare, I follow Caliendo et al. (2019) and calculate the percentage annual wage change δ needed to compensate a worker of characteristics ω and age a for a given change in policy. Let v^0 and v^1 denote the worker value functions in two policy scenarios whose welfare implications we would like to compare. Due to the logarithmic flow utility of workers in Equation (27), the compensating variations δ are simply given by

$$v_t^1(\omega, a) = v_t^0(\omega, a) + \sum_{\tau=0}^{\bar{A}-a} \beta^\tau \delta_t(\omega, a) \quad (88)$$

$$\delta_t(\omega, a) = (v_t^1(\omega, a) - v_t^0(\omega, a)) \frac{(1 - \beta)}{(1 - \beta^{\bar{A}-a+1})} \quad (89)$$

G.2 Policy Counterfactual: The Dynamic Incidence of a Robot Tax

G.2.1 Predatory Investment Externalities

This section incorporates predatory investment effects into the robot tax incidence analysis. Predatory investment effects refer to the pecuniary externality where parts of the profit gains from robot adoption come from crowding out competitors in output markets. If demand is sufficiently elastic, firms will be willing to undertake very costly fixed robot investments to obtain just an infinitesimal variable cost advantage over its competitors.

To analyze the effects of such predations, realize first that firm values in Equations (7)-(8) are driven by changes in flow profits π_t and robot adoption costs c_t^R . Flow profits depend in turn on firm unit costs C_t , manufacturing demand Y_{Mt} , and the manufacturing price P_{Mt} ; see Equations (5) and (6). The predatory investment externality works through the price index P_{Mt} . When tabulating the effects on firm values in Table 5, I hold this externality fixed by calculating

$$\tilde{v}_t^T - v_t^B = v(\{c_\tau^{RT}, C_\tau^T, Y_{M\tau}^T, P_{M\tau}^B\}_{\tau=t}^\infty) - v(\{c_\tau^{RB}, C_\tau^B, Y_{M\tau}^B, P_{M\tau}^B\}_{\tau=t}^\infty), \quad (90)$$

where superscripts T and B denote the robot tax counterfactual and baseline equilibrium, respectively.

Table G.2 now incorporates the predatory investment externalities by calculating

$$v_t^T - v_t^R = v(\{c_\tau^{RT}, C_\tau^T, Y_{M\tau}^T, P_{M\tau}^T\}_{\tau=t}^\infty) - v(\{c_\tau^{RB}, C_\tau^B, Y_{M\tau}^B, P_{M\tau}^B\}_{\tau=t}^\infty) \quad (91)$$

Table G.2 shows a stark finding: For baseline values of model parameters, the predatory externalities are large enough to make total tax revenues exceed total profit losses from the robot taxes. Put differently, if tax revenues can be rebated to firms appropriately, a robot tax has the potential to increase firm values by internalizing the predatory externalities of robot adoption.

Table G.2: Robot Tax Incidence with Predatory Investment Externalities
(Discounted Present Values in Percent of GDP in 2019)

	Temporary Tax	Permanent Tax
Profits	-1.65	-7.90
Predatory Investment Externalities	2.48	2.67
Tax Revenues	2.39	9.41

I hold these predatory externalities on firm profits out of the baseline analysis to focus on the key equity-efficiency trade-off for workers. That said, the analysis in this section suggests that studying predatory implications of recent automation technologies may be a fruitful avenue of future research.