Elasticity of Substitution between Robots and Workers: Theory and Evidence from Japanese Robot Price Data[1]

Daisuke Adachi[a,b]

[a] *Department of Economics and Business Economics, Aarhus University, Fuglesangs Alle 4, Aarhus, 8210, Denmark*
[b] *Research Institute of Economy, Industry and Trade, 1-3-1, Kasumigaseki, Chiyoda-ku, 100-8901, Tokyo, Japan*

Abstract

This paper examines the wage effects of the increased use of industrial robots, focusing on their role in specific tasks and international trade. I construct a novel dataset by tracking shocks to the cost of acquiring robots from Japan, termed the Japan Robot Shock (JRS), and analyze these shocks across different occupations that have adopted robots. A general equilibrium model incorporating robot automation in a large open economy is developed, and a model-implied optimal instrumental variable of the JRS is constructed to address the identification challenges posed by the correlation between automation shocks and the JRS. The study finds that the elasticity of substitution (EoS) between robots and labor is heterogeneous across occupations, reaching up to 3 in production and material moving occupations, which is significantly higher than the EoS between other capital goods and labor.

JEL codes: J23, J24, J62, E24, F16, F66, O33

Keywords: Industrial Robots, Robot Prices, Elasticity of Substitution

[1] I acknowledge the financial support from the Japan Society for the Promotion of Science (No. 23K12498). All errors are my own.

1. Introduction

Industrial robots have rapidly transformed factory production. Over the past three decades, the global robot market has grown by 12% annually (IFR, 2021). Robotics has heterogeneous impacts on workers across occupations, raising concerns about its distributional effects. Therefore, policymakers have proposed various countermeasures for the potential harms of robotization, such as taxes on robot adoption. Motivated by these observations, a growing body of literature has estimated the effects of robot penetration on human employment (e.g., Acemoglu and Restrepo, 2020) and the potential impact of robot taxes (e.g., Humlum, 2021). However, few studies have explored the effect of factors such as the substitutability of robots for workers in each occupation that also determine the impact of robotization.

In this paper, I analyze the role of robots in wage inequality between occupations and welfare in the US. In contrast to previous research that reveals the substitutability between professions, I estimate the substitutability between robots and workers within an occupation using a novel dataset that tracks the cost of adopting Japanese robots. For this purpose, I construct a model-implied optimal instrumental variable (MOIV) and estimate the elasticity of substitution (EoS) between robots and workers that is heterogeneous across occupations. Finally, I conduct counterfactual exercises to analyze the distributional effects of robotization in the US from 1990 to 2007.

I use information on shipments of Japanese robots, accounting for approximately one-third of the world's robot supply, from the Japan Robot Association (JARA). A key feature of the JARA data is that sales quantity and total value are observed at the level of robot application or the specified task that robots perform. To obtain an occupation-level robot price measure, I combine the JARA data with the O*NET Code Connector match score. Ultimately, I extract a robot cost shock that controls for demand factors using leave-one-out regression, which I call the *Japan Robot Shock* (JRS).

I use an equilibrium model of robot automation in a large open economy. Occupations are bundles of tasks where tasks can be performed either by workers or robots (factors). I impose a Fréchet distribution for the task-specific productivity of each factor, enabling the aggregation of tasks to the occupational production function, featuring the constant EoS (CES) between robots and labor within each occupation. Using this formulation, I can interpret changes in robot quality in terms of changes in the robot expenditure share parameter, which I call the automation shock. In addition, I include the Armington-style robot trade to capture Japan's substantial robot exports.

An identification challenge in estimating robot–labor EoS is that the JRS may be correlated with the automation shock, which is unobserved. I overcome this challenge by using the general equilibrium restriction to obtain structural residuals of occupational wages, interpreted as the remaining variation in occupational wages after controlling for the effect of the automation shock. The identification assumption is that these structural residuals are not correlated with the JRS, implying a moment condition that provides consistent parameter estimates and an optimal instrumental variable to increase estimation precision.

Using this estimation method, I find that the average EoS between robots and workers is about 2. This estimate is higher than the typical values in the labor-capital EoS literature, highlighting a major difference between robots and other capital goods. Moreover, the EoS estimates are heterogeneous across occupations. In particular, for routine occupations that perform production tasks, the point estimates are as high as around 3, revealing the particular vulnerability of workers in these occupations to robots. These estimates are identified by a strong relationship between increased decline in robot price and lowered occupational wage growth rate in these occupations. In contrast, the estimates in other occupations are around 1, suggesting that robots and labor are less substitutable in such occupations.

The large EoS between robots and workers in occupations involving production and material moving implies that robotization significantly reduced the relative wage in these occupations over the sample period. In other words, the shock of robotization slowed the relative wage growth of occupations in the middle deciles, because robotized occupations tended to be in the middle of the occupational wage distribution in 1990. Moreover, the higher productivity in these occupations raised the marginal product of labor in other occupations, increasing labor demand.

This paper contributes to the literature on the economic impact of industrial robots by identifying the significant impact of robotization on wage inequality in the US. The closest papers to mine are Acemoglu and Restrepo (2020) and Humlum (2021). Acemoglu and Restrepo (2020) find that U.S. commuting zones with increased robot penetration in 1992–2007 experienced lower wage and employment growth.[2] Meanwhile, Humlum (2021) estimates

[2] Dauth et al. (2017) and Graetz and Michaels (2018) also use aggregate industry-level data on robot adoption to analyze its impact on labor markets. Galle and Lorentzen

a model of robot importers in a small open economy and the EoS between occupations using firm-level data on robot adoption, finding a positive average real wage effect with significant heterogeneity across occupations.[3] I complement the findings of these studies by providing a method to estimate the within-job EoS between robots and workers using occupation-level robot cost data. The estimations reveal the heterogeneous substitutability of robots and workers in the US.

Another strand of the literature focuses on occupations, aiming to clarify the potentially heterogeneous effects of automation (e.g., Cheng, 2018). In particular, Jaimovich et al. (2021) construct a general equilibrium model to study the impact of automation on the labor market of routine and nonroutine workers. I contribute to these efforts by providing a matching method for industrial robot applications and occupations that produces occupation-level robot cost data, allowing me to estimate the robot–labor EoS.

In addition, this paper is related to the vast literature on estimating the EoS between capital and labor (e.g., Arrow et al., 1961; Oberfield and Raval, 2014).[4] Although the literature provides numerous estimates with a

(2024) examine the interaction effects of trade and automation. In addition, Adachi et al. (2024) use JARA data to study the impact of robots on the Japanese labor market. In contrast, this paper studies the U.S. labor markets and examines the impact of robots on wages across occupations by estimating the EoS between robots and workers.

[3]Like Humlum (2021), a growing number of studies (including Koch et al., 2021) use firm-level data to study robots and workers.

[4]Caunedo et al. (2023) provide the EoS between labor and tools for each occupation by applying a natural language processing algorithm to tool descriptions, using data from the BEA fixed asset table. The exercise focuses on capital-embodied technological change

wide range, the upper limit appears to be around 1.5 (Karabarbounis and Neiman, 2014; Hubmer, 2023). By contrast, my EoS estimates of around 3 in occupations involving production and material moving are significantly higher than this upper limit. In this sense, the findings of this study highlight the particular vulnerability of workers to robots across occupations as one of the main differences between robots and other capital goods.

2. Model

The model adopts a task-based framework embedded in a multi-country Armington model. This framework has two main features: occupation-specific EoS between robots for workers and robot trade in a large open economy. In this study, I emphasize these features and discuss the other model elements based on later quantitative exercises in Appendix A.1. In the main text, we focus on the steady-state changes and omit the time subscript. I consider the dynamic problem in Appendix A.

2.1. Environment

There are N countries, O occupations, and two types of tradable goods (g): non-robot goods $g = G$ and robots $g = R$. Whenever possible, I use the country subscripts as follows:

- l: robot-exporting country,
- i: non-robot goods-exporting and robot-importing country,

(CETC), which is modeled as a reduction in tool prices. I treat the automation shock and robot price decline separately and address the resulting identification challenge.

- j: non-robot goods-importing countries country.

Each country has representative households and producers. As in the Armington model, non-robot goods are differentiated by country of origin while robots are differentiated by country of origin and occupation. non-robot goods can be consumed by households and invested to produce robots.[5]

In the main text, non-robot goods G are produced with two factors of production: labor $L_{i,o}$ and robot capital $K^R_{i,o}$ in each occupation o.[6] There is no international factor mobility. Producers own and accumulate robot capital. Households own the producers' shares in each country. All goods and factor markets are perfectly competitive. Workers are forward-looking, draw an idiosyncratic utility shock from a generalized extreme value distribution, pay a switching cost for changing occupations, and choose the occupation o that achieves the highest expected value $V_{i,o}$ among O occupations, following Caliendo et al. (2019). The discount rate is $\iota > 0$. The elasticity of the probability of changing occupation concerning the expected value is ϕ. The details of the worker problem are provided in Appendix A.1.

There are good-specific iceberg trade costs τ^g_{ij} for each $g = G, R$. There are no intra-country trade costs; therefore, $\tau^g_{ii} = 1$ for all i and g. Due to the iceberg costs, the bilateral price of good g that country j pays to i is $p^g_{ij} = p^g_i \tau^g_{ij}$.

Each country's government exogenously imposes a robot tax. Specifically,

[5] In the full model in Appendix A.1, non-robot goods are used as input for robot integration (Humlum, 2021).

[6] Appendix A.1 shows the model with intermediate goods and non-robot capital in the production function. The analytical results in our main analysis are unchanged.

buyer i of robot o from country l must pay an *ad valorem* robot tax u_{li} on top of the producer price of robots $p^R_{li,o}$ to buy from l. The tax revenue is uniformly rebated to households in the country.

2.2. Production Function, Tasks, and Automation

Production of Non-Robot Goods. In country i, the representative producer of non-robot good G uses the occupation-o service $T^O_{i,o}$ and produces with the following production function:

$$Y^G_i = A^G_i \left[\sum_o (b_{i,o})^{\frac{1}{\beta}} \left(T^O_{i,o}\right)^{\frac{\beta-1}{\beta}} \right]^{\frac{\beta}{\beta-1}}, \tag{1}$$

where A^G_i is a Hicks-neutral productivity, $b_{i,o}$ is the cost share parameter of each occupation o, and β is the EoS between each occupation in the production function. The parameters satisfy $b_{i,o} > 0$, $\sum_o b_{i,o} = 1$, and $\beta > 0$.

I adopt the canonical task-space framework at the occupation level (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020). The occupation service is a combination of tasks $\omega \in [0, 1]$ with CES technology

$$T^O_{i,o} = \left[\int_0^1 (t_{i,o}(\omega))^{\frac{\zeta-1}{\zeta}} d\omega \right]^{\frac{\zeta}{\zeta-1}}, \tag{2}$$

where $t_{i,o}(\omega)$ is the input of the task ω and $\zeta \geq 0$ is the EoS between tasks. A task is performed by robots or workers with perfect substitutability:

$$t_{i,o}(\omega) = Z^R_{i,o}(\omega)\, k^R_{i,o}(\omega) + Z^L_{i,o}(\omega)\, l_{i,o}(\omega)$$

where $Z^R_{i,o}(\omega)$ and $Z^L_{i,o}(\omega)$ are the task-specific productivity for robots and workers, respectively. Due to perfect competition, task prices are determined by marginal cost, the minimum of the efficiency prices of labor $w_{i,o}/Z^L_{i,o}(\omega)$

and robots $c^R_{i,o}/Z^R_{i,o}(\omega)$ for each task ω. The share of tasks performed by robots is denoted as $\xi_{i,o}$.

Following Artuc et al. (2023), I assume Fréchet-distributed productivity with scale parameter a^s_o $(s = R, L)$ and shape parameter θ_o, with the restriction $\theta_o \geq \zeta$. I assume that robot productivity is a common technical characteristic for all countries; thus, a^s_o does not vary across countries. I also normalize a^R_o and a^L_o so that they sum to one and write a_o as the normalized parameter for robots to make it easier to interpret the share parameter in the robot task space. The maximum stability property of the Fréchet distribution implies that $\xi_{i,o}$ is equal to the fraction of spending on robots (Eaton and Kortum, 2002), and

$$\xi_{i,o} = \frac{c^R_{i,o} K^R_{i,o}}{P^O_{i,o} T^O_{i,o}} = a_o \left(\frac{c^R_{i,o}}{P^O_{i,o}} \right)^{1-\theta_o}, \tag{3}$$

$$\text{where} \qquad P^O_{i,o} = \left(a_o (c^R_{i,o})^{1-\theta_o} + (1 - a_o)(w_{i,o})^{1-\theta_o} \right)^{1/(1-\theta_o)}, \tag{4}$$

where $c^R_{i,o}$ is the user cost of robot capital, formally given in Appendix A.2, and $P^O_{i,o}$ is the unit cost of occupation o. A key parameter is θ_o, which governs the EoS between labor and robots in each occupation o. Intuitively, the more dispersed the task productivities $Z^R_{i,o}(\omega)$ and $Z^L_{i,o}(\omega)$, the less sensitive the optimal allocation of labor and robots is to price changes because the unobserved productivity difference is more important.

Production of Robots. Robots for occupation o are produced by investing non-robot goods $I^R_{i,o}$ with productivity $A^R_{i,o}$ due to perfect competition:

$$Y^R_{i,o} = A^R_{i,o} I^R_{i,o}, \qquad \text{so} \qquad p^R_{i,o} = \frac{P^G_i}{A^R_{i,o}}, \tag{5}$$

where P_i^G is the price index for non-robot goods, as given below in (6). The robot price is inversely proportional to the productivity term $A_{i,o}^R$. Therefore, I refer to the change in $A_{i,o}^R$ for $i = JPN$ as the JRS throughout.

Trade in Goods and Robots. The elasticity of trade in non-robot goods (or robots) is denoted as ε (or ε^R). The import shares of goods and robots in j from i and their price indices are provided by

$$x_{ij}^G = \left(\frac{p_{ij}^G}{P_j^G}\right)^{1-\varepsilon} \text{ and } x_{ij,o}^R = \left(\frac{p_{ij,o}^R}{P_{j,o}^R}\right)^{1-\varepsilon^R}$$
$$\text{where } P_j^G = \left[\sum_i \left(p_{ij}^G\right)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}} \text{ and } P_{j,o}^R = \left[\sum_i \left(p_{ij,o}^R\right)^{1-\varepsilon^R}\right]^{\frac{1}{1-\varepsilon^R}}, \quad (6)$$

because of the Armington assumption.

2.3. Discussion of Model Assumptions

The robot technological efficiency parameter a_o in (3) plays a central role in estimations and counterfactuals and is discussed in detail here. Because the task-based framework developed in Section 2.2 includes the allocation of factors to tasks, I can interpret a_o as the shifter in the robots' share of tasks as opposed to labor by appropriately modifying the productivity term $b_{i,o}$, which is discussed in detail in Section 2.5. Thus, I call the change in a_o the *automation shock*.

The robot cost share, a_o, can also represent robot quality, as it is a non-pecuniary attribute whose value all agents agree on (Khandelwal, 2010). As (3) states that the increase in a_o implies an increase in the value of robots among production factors, the automation shock can be interpreted as a quality upgrade of robots relative to labor when combined with the productivity adjustment.

Therefore, my model does not distinguish between the automation shock and the quality upgrade: they have the same effect on equilibrium due to the restrictions of the Fréchet distribution assumption. To my knowledge, there has been no formal discussion of this point. Nevertheless, retaining this assumption is helpful to maintain complex technology improvements along with task automation and quality upgrades within a single parameter a_o.[7]

As comparative statics, I consider the JRS and the automation shock, which are together referred to as *robotization shocks*. It is likely that the JRS and the automation shocks are correlated with each other at the occupation level because innovations in robot technology improve the applicability of robots while reducing the cost of adoption.[8] This will be a source of the identification challenge discussed later.

2.4. Equilibrium

The rest of the model is standard in the dynamic general equilibrium literature and is presented in Appendix A.1. For the purpose of notation, I summarize the solution of the workers' dynamic discrete choice problem of occupations given occupational wages by the labor supply function $L_{i,o}(\boldsymbol{w}_i)$, suppressing its dependence on future values. The non-robot producer solves the dynamic robot capital investment problem under convex adjustment costs

[7] One of the reasons to impose this assumption is the lack of data on the set of tasks for each robot or the quality of the robots. Relaxing this assumption using rich data on this dimension will be addressed in future work.

[8] See Appendix B.1 for more concrete accounts of such a correlation.

(Cooper and Haltiwanger, 2006).[9] The prices of goods, labor, and robots equilibrate the respective markets in general equilibrium.

2.5. Solving the Model

I apply the first-order approximation to the steady state (Blanchard and Kahn, 1980). I chose this strategy over the exact solution method like Caliendo et al. (2019) because the trade literature has shown that the errors of the first-order approximation with respect to (unilateral) productivity shocks are considerably smaller than those due to bilateral trade shocks (Kleinman et al., forthcoming). This paper considers a unilateral robotization shock. For example, my model assumes that Japanese robots have become accessible to all countries (not just the US). This subsection focuses on the steady-state change and drops the subscript t. The complete characterization of the approximation and transition dynamics is provided in Appendix A.3.

I describe the log total derivative using the hat notation. The exogenous shocks are the shocks to a_o, $A^R_{l,o}$, and the adjustment to the occupational productivity term $b_{i,o}$. Throughout the paper, I only consider a type of automation shock that does not change labor productivity, reflecting the rapid growth of robotic technology relative to that of human capital in recent decades. Mathematically, this is equivalent to

$$\widehat{b_{i,o}}^{\frac{1}{\beta-1}} \left(\widehat{1-a_o}\right)^{\frac{1}{\theta_o-1}} = 0, \tag{7}$$

for each automation shock $\widehat{a_o}$, such that the effect of the change in a_o on labor productivity is offset by a corresponding adjustment in $\widehat{b_{i,o}}$. This approach

[9] de Souza and Li (2023) also apply the problem to the robot context.

still captures overall productivity growth due to the change in $\widehat{a_o}$. This is the typical approach in the literature for controlling labor productivity growth when modeling robot shocks. For example, the canonical setup in Acemoglu and Restrepo (2020) model automation by increasing the robot availability threshold across tasks. This does not change labor productivity; however, the overall productivity increases due to the threshold increase.

I provide several approximation expressions useful in the following sections when defining the estimator. First, I combine (5) and (6) to get the change in the robot price index $P^R_{i,o}$ in country i due to the change in robot production technology $A^R_{l,o}$ in country l:

$$\widehat{P^R_{i,o}} = -x^R_{li,o}\widehat{A^R_{l,o}} + \sum_{l'} x^R_{l'i,o}\widehat{P^G_{l'}}, \tag{8}$$

where the first term reflects the direct effects of the change in robot productivity in l mediated by the import share of robots from l in i. The second term summarizes the general equilibrium effects due to changes in the production cost of robots in other countries on the robot price index.

Second, from (3) and (4), the labor demand in dollar units in (i, o) is given by $(1 - \xi_{i,o})P^O_{i,o}T^O_{i,o}$. As a result, the approximated labor market equilibrium condition is as follows:

$$\widehat{w_{i,o}} + \sum_{o'} \frac{\partial \ln L_{i,o}}{\partial \ln w_{i,o'}}\widehat{w_{i,o'}} = (\widehat{1 - a_o}) + (1 - \theta_o)(\widehat{w_{i,o}} - \widehat{P^O_{i,o}}) + \widehat{P^O_{i,o}} + \widehat{T^O_{i,o}}, \tag{9}$$

where the LHS and RHS are the changes in supply and demand, respectively.

3. Estimation Strategy

Following Adao et al. (2023), I develop an estimation method using MOIV, which is applied to my novel measure of Japanese robot price reductions. Throughout this section, I consider the following identification challenges: (i) Robot prices may be driven by demand rather than cost, (ii) There is a correlation between the automation shock and robot price, (iii) Unilateral technical changes may drive robot prices, and (iv) Non-Japanese robot prices also change.

3.1. Parameterization

First, I set the sample period to 1992-2007 (or 1990-2007 for the labor data) and, given the data availability, write $t_0 \equiv 1992$ and $t_1 \equiv 2007$. I will relate the long difference to the steady-state changes of the model.

I account for the heterogeneity of EoS between robots and labor across occupations while maintaining estimation power by defining the following occupational groups. First, occupations are divided into three broad occupational groups: Abstract, Service (Manual), and Routine, following Acemoglu and Autor (2011). Given the trend of intensive robot adoption in production and transportation (material moving) occupations over the sample period, I further divide routine occupations into three subcategories: Production (e.g., welders), Transportation (indicating transportation and material moving, e.g., laborers), and Other (e.g., repairers). This leads to five occupational groups, the full list of which is presented in Appendix B.2. Within each group, I assume a constant EoS between robots and workers. Each occupation group is denoted by the subscript g; thus, the robot-labor EoS for group g is written as θ_g.

As I use Japanese robot prices and study the US labor market, I set $N = 3$ and aggregate the country groups to the US (USA, country index 1), Japan (JPN, index 2), and the Rest of the World (ROW, index 3). The annual discount rate is $\iota = 0.05$. Following Graetz and Michaels (2018), the robot depreciation rate is 10%. I take the trade elasticity of $\varepsilon = 4$ from the literature on trade elasticity estimation (e.g., Simonovska and Waugh, 2014) and $\varepsilon^R = 1.2$ derived by applying the estimation method developed by Caliendo and Parro (2015) to the robot trade data, which is discussed in detail in Appendix D.1. The remaining parameters $\mathbf{\Theta} \equiv \{\theta_g, \beta\}$ are the target of the following structural estimation.

The first-order approximation requires shares in the initial period, which are taken from the International Federation of Robotics (IFR), Integrated Public Use Microdata Series (IPUMS) USA, Current Population Survey (CPS), Database for International Trade Analysis (BACI), and the World Input-Output Table (WIOT). I set the initial robot share parameter a_{o,t_0} to the initial US occupation-specific expenditure share $c^R_{US,o,t_0} K^R_{US,o,t_0} / w_{US,o,t_0} L_{US,o,t_0}$ and the initial robot tax is zero in all countries. The remaining labor market outcomes are measured as standard and mentioned in Appendix B.2.

3.2. Data Source on Robots

Industrial robots are formally defined as multi-axis manipulators and measured by the number of manipulators or robot arms.[10]. The main data source for robots by occupation is the JARA, a general incorporated association comprising Japanese robot manufacturing companies. In its "Export Statistics of Manipulators, Robots and Applied Systems by Country and Application", JARA annually surveys major robot manufacturers regarding the units and monetary values of robots sold for each destination country and

[10]The full ISO-based definition is presented in Appendix B.1.

robot application. Robot applications are defined as the specified tasks that robots perform and are discussed in detail in Section 3.3.

To convert robot applications to occupations, I use the Occupational Information Network Online (O*NET) Code Connector. The O*NET Code Connector is an online database of occupations sponsored by the US Department of Labor, Employment, and Training Administration and provides an occupational search service. The algorithm used in the search service provides a match score indicating the relevance of each occupation to the search term, as discussed by Morris (2019) and Appendix B.2.

To integrate Japanese robot data from JARA and international trade data from BACI, I use HS code 847950 ("Industrial robots for multiple uses") as the robot definition in the trade data. I match the BACI robot trade data to JARA robot exports by aggregating applications in the JARA data. As I do not observe the occupation-level disaggregation of robot trade in other countries, I impose $x^R_{ij,o} = x^R_{ij}$ for all o in the estimation. See Appendix B.4 for the details of the robot measurement issues in JARA and BACI.

3.3. Data Construction

This subsection describes the construction of the robot price at the occupation level. Although Graetz and Michaels (2018) provide data on robot prices from IFR, their price data are aggregated but not distinguished by occupation. In contrast, I use variation at the occupation level to estimate substitutability between robots and workers.

Step 1. Application-Occupation Matching. The first step is to match robot applications and worker occupations. A heterogeneous mix of tasks in each occupation generates a difference in ease of automation across occupations, implying heterogeneous robot adoption across occupations (Manyika et al., 2017).[11] Formally, let a denote a robot application and o a labor occupation at the 4-digit level. The JARA data provide the number of robots sold and the total monetary transaction values for each application a. These robot measures are denoted as X_a^R, a generic notation indicating quantities and monetary values. The application-level robot measure X_a^R is converted to an occupation-level measure X_o^R using a weighted average. For this purpose, I search occupations in the O*NET Code Connector for the title of the robot application a and web-scrape the match score m_{oa} between a and o. Using m_{oa} as the weight, I compute[12]

$$X_o^R = \sum_a \omega_{oa} X_a^R \text{ where } \omega_{oa} \equiv \frac{m_{oa}}{\sum_{o'} m_{o'a}}. \tag{10}$$

where $\sum_o \omega_{oa} X_a^R = X_a^R$ because $\sum_o \omega_{oa} = 1$.

This matching method has low data requirements, which is useful given that I only observe the titles of robot applications, and not detailed descriptions such as patent texts. In this sense, this method complements the ones used in previous studies. For example, Webb (2019) provides a natural language processing method to match recent technological advances (e.g., robotics) embodied in patent titles and abstracts with occupations. Montobbio et al. (2020) extends this approach to analyzing full patent texts by applying the topic modeling method.

[11] Appendix B.1 provides examples of robot applications.

[12] More details on matching are described in Appendix B.5, including the use of hard-cut matching, which does not significantly affect the matching result.

Step 2. Constructing JRS. Using the occupation-level robot quantity $q^R_{i,o,t}$ and sales $(pq)^R_{i,o,t}$ in destination country i, occupation o, and year t, the cost shocks to robot users are constructed in each occupation as follows. First, I take the average export price $p^R_{i,o,t} \equiv (pq)^R_{i,o,t} / q^R_{i,o,t}$.[13] One concern with using unit value data is simultaneity, i.e., demand shocks and not cost shocks drive prices, as in point (i) at the beginning of this section.[14] My measure of export prices is based on external robot sales; thus, I am less concerned with the endogeneity from the use of domestic robot prices. Nevertheless, I exclude US robot import prices from the sample to mitigate simultaneity concerns. Here, the argument is consistent with Hausman et al. (1994), who argued that changes in demand shocks are uncorrelated between the US and other countries, but that price variations are primarily driven by robot production costs in producer countries. This leave-one-out idea is widely used in the automation literature (e.g., Acemoglu and Restrepo, 2020).

To further address cross-country correlation in demand shocks, I exploit the fact that the data are from bilateral trade flows and control for the destination country-specific demand effect. Formally, I fit the fixed-effects regression as follows:

$$\ln\left(p^R_{i,o,t}\right) - \ln\left(p^R_{i,o,t_0}\right) = \psi^D_{i,t} + \psi^J_{o,t} + \epsilon_{i,o,t},\ i \neq USA \tag{11}$$

[13] I also compute the chain-weighted robot price index, which is commonly used to measure the price of capital goods. The results using this index are not qualitatively different from the main results.

[14] Another concern is robot quality upgrading. A data-driven approach to address this issue is the hedonic and cost-estimation approaches, both of which are discussed in Appendix B.6.

where t_0 is the initial year, $\psi_{i,t}^{D}$ is the destination-year fixed effect (FE), $\psi_{o,t}^{J}$ is the occupation-year FE, and $\epsilon_{i,o,t}$ is the residual. This regression controls for any country-year-specific effect $\psi_{i,t}^{D}$ that includes the demand shock of country i or the trade shock between Japan and country i. I use the remaining variation across occupations $\psi_{o,t}^{J}$ as the cost shock of robot adoption and define $\psi_{o}^{J} \equiv \psi_{o,t_1}^{J}$ as the measured JRS.

Finally, I relate JRS to the model's robot productivity using the perfect competition assumption and the robot production function (5):

$$\psi_{o}^{J} = -\widehat{A_{2,o}^{R}}. \tag{12}$$

Appendix C presents stylized facts and reduced-form evidence about robots and workers at the occupation level, suggesting strong substitutability between robots and workers.

3.4. Estimation Procedure

The constructed data provide information about the robot price shock, a critical input for estimating the elasticity parameter in (3). The next identification threat is the unobserved automation shock, a_o, as pointed out in (ii) at the beginning of this section. I develop a moment condition using the model's restriction to address this concern.

First, I decompose the automation shock $\widehat{a_o}$ into an "implied" component $\widehat{a_o^{\text{imp}}}$ and an "unobserved residual" component $\widehat{a_o^{\text{res}}}$ such that $\widehat{a_o} = \widehat{a_o^{\text{imp}}} + \widehat{a_o^{\text{res}}}$ for all o. The steady-state change in the relative demand for robots and labor implicitly defines the implied component. Using (3), (8), and (12),

$$\left(\widehat{\frac{c_{US,o}^{R} K_{US,o}^{R}}{w_{US,o} L_{US,o}}}\right) = (1-\theta_g)\left(x_{JP,US}^{R}\psi_{o}^{J} - \widehat{w_{US,o}}\right) + \frac{\widehat{a_o^{\text{imp}}}}{1-a_{o,t_0}} + D, \tag{13}$$

where $x^R_{JP,US}$ is the base-year import share of robots from Japan in the US, and $D \equiv (1-\theta_g)\sum_l x^R_{l,US}\widehat{P^G_l}$ is the international spillover term due to changes in price indices in other countries. That is, $\widehat{a_o^{\text{imp}}}$ is the automation shock component explaining the shift in the expenditure share of robots. In contrast, the unobserved residual component $\widehat{a_o^{\text{res}}}$ is the residual term, which I consider as the measurement error.

The identification challenge is that the JRS ψ_o^J is potentially correlated with the implied automation shock $\widehat{a_o^{\text{imp}}}$. The literature estimates the capital-labor elasticity of substitution using the CES demand function of the form (3). However, this assumes that the technology shock is fixed or orthogonal to price changes.[15] As many task-based models yield a demand function where the technology shock $\hat{a}_o$ can be interpreted as the expansion of the task space for robots, the correlation of this shock with the decline in robot prices, another measure of technological progress, should be addressed formally.

A key observation is that the residual component $\widehat{a_o^{\text{res}}}$ can be inferred from the observed endogenous variables using the first-order solution and $\widehat{a_o^{\text{imp}}}$. Namely, the occupational labor market clearing condition (9) relates occupational wage changes and the automation shock. More specifically, combined with $\widehat{a_o^R} = \widehat{a_o^{R,\text{imp}}} + \widehat{a_o^{R,\text{res}}}$, I have

$$\widehat{a_o^{R,\text{res}}} = -\widehat{a_o^{R,\text{imp}}} - (1-a_o)\left[\widehat{w_{i,o}} + \sum_{o'}\frac{\ln L_{i,o}}{\ln w_{i,o'}}\widehat{w_{i,o'}} - (1-\theta_o)(\widehat{w_{i,o}} - \widehat{P^O_{i,o}}) - \widehat{P^O_{i,o}} - \widehat{T^O_{i,o}}\right] \tag{14}$$

where $\widehat{P^O_{i,o}}$ is implied by (4) and $\widehat{T^O_{i,o}}$ is given by (2). Equation (14) obtains

[15] See, for example, Herrendorf et al. (2015) and Eden and Gaggl (2018).

a structural residual after controlling for the automation shock measured from the expenditure share expression in (13). Thus, the following moment condition is imposed on this structural residual and the JRS $\boldsymbol{\psi}^J \equiv \left\{\psi_o^J\right\}_o$.

Assumption 1. *(Moment Condition) For any subset of occupations $G \subset \{1, \ldots, O\}$,*

$$\mathbb{E}\left[\widehat{a_o^{R,res}} | \boldsymbol{\psi}^J, o \in G\right] = 0. \tag{15}$$

The EoS between occupations, β, is estimated jointly with the EoS between robots and workers, θ_g; I take the sample analogs of (15) for each occupational group g and all occupations. Roughly, the conditional moment of each group g identifies θ_g conditional on the value of β, and the overall moment condition pins down β. Given (15), it is routine to construct the optimal GMM and implement it with the two-step estimator following Adao et al. (2023). Therefore, I leave a detailed explanation in Appendix D.3.

3.5. Discussion of the Identification Assumption

Assumption 1 restricts the structural residual $\widehat{a_o^{R,\text{res}}}$ such that it should not be predicted by the JRS. Note that it allows the automation shock $\widehat{a_o}$ to correlate with changes in robot producer productivity $\widehat{A_{2,o}^R}$. Intuitively, the structural residual $\widehat{a_o^{R,\text{res}}}$ refers to the remaining variation after controlling for the effects of the robotization shocks on wage changes, $\widehat{\boldsymbol{A}_2^R}$, and $\widehat{\boldsymbol{a}}$ (and the associated adjustment $\widehat{\boldsymbol{b}}$ in (7)). My restriction is that the remaining variation, as it is a measurement error, cannot be predicted by the JRS.

What breaks this assumption? First, the correlation of the structural residuals with other shocks, such as trade shocks, could. To this, a sensitivity analysis in Section 4 controls for the China shock at the occupation level and demonstrates the result's robustness. The robustness is further verified in Appendix C, which shows that the reduced-form linear regression coefficients do not change qualitatively after controlling for the China shocks.

The second threat is the directed technological changes raised in (iii) at the beginning of the section, where occupational labor demand drives changes in the cost of robots (e.g., Acemoglu and Restrepo, 2018). If a positive labor demand shock in occupation o induces research and development of robots in occupation o and drives down costs in the long run rather than exogenous technological change in the production function (5). In this case, the structural residual $\widehat{a_o^{R,\text{res}}}$ does not control for this effect and is negatively correlated with JRS ψ_o^J. However, Appendix Figure B.2 in the original manuscript (Figure Appendix C.1a in the revised manuscript) shows that there is no correlation between the baseline wage and the JRS after controls and occupation group fixed effects, suggesting the limited evidence on the directed change of Japanese robot technology due to the U.S. wages.

Another possibility that fails Assumption 1 is increasing returns to robot producers, implying that the unobserved increase in robot demand reduces robot costs. However, my estimation relies on the *foreign* robot price data, mitigating this concern. Moreover, even though these concerns bias the estimates, they imply a negative bias in the elasticity estimates, thus preserving my qualitative results of strong substitutability.

Finally, the estimation procedure assumes that unobserved reductions in the cost of robots from other countries are independent of the evolution of Japanese robot costs, as in (iv) at the beginning of the section. I discuss the plausibility of this assumption in Appendix C.3 by comparing the data from the JARA and the IFR.[16]

4. Estimation Results

Before showing the estimation results, I briefly review the JRS variation. Figure 1a plots the distribution (10th, 50th, and 90th percentile) of the growth rates of the nominal price of Japanese robots in the US each year relative to the initial year. The figure shows two patterns: (i) the robot prices follow an overall decreasing trend, and (ii) there is significant heterogeneity in the rate of price decline across occupations.

Figure 1b shows the distribution of the long-run trend (1992–2007) for each occupation group: the three routine occupation groups, service, and abstract. It confirms a significant price variation across occupations, even within occupation groups. Surprisingly, the average reduction in production robot prices is not as stark as other robots. This indicates that the robot technological change in production occupations is not reflected by the price decline but rather by the automation shock.

Figure 1b also shows the variation in JRS, or ψ^{J}_{i,t_1}, in (11). The large variation of the changes in prices by occupations persists. It also confirms that after controlling for US demand shocks, the Japanese robot cost significantly decreases, especially in the production occupation.

At the same time, there is a strikingly rapid growth in robot adoption in the US (Acemoglu and Restrepo, 2020). This includes increased robot

[16] Appendix B.4 shows the international robot flows, including Japan, the US, and the rest of the world.

Figure 1: Distribution of the Robot Prices and Japan Robot Shock

(a) Robot Prices in the US (b) Variation in the Japan Robot Shock

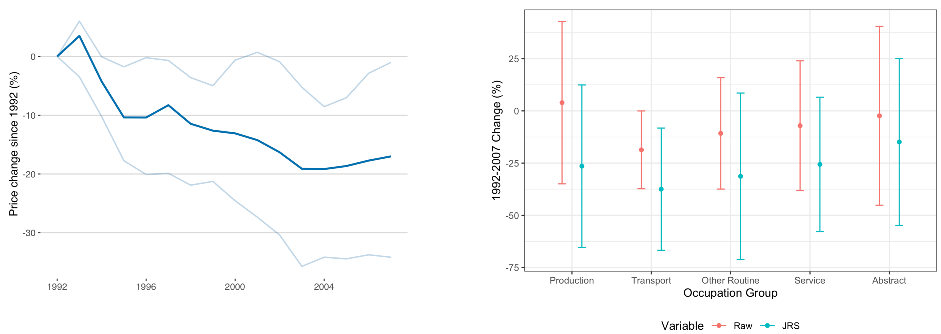

Note: The left panel shows the trend of nominal prices of robots in the US by occupations, $p^R_{USA,o,t}$. The bold and dark lines show the median price each year, and the two thin and light lines represent the 10th and 90th percentile. Three-year moving averages are shown to smooth out yearly noises. The right panel shows the mean and standard deviation of the long-run (1992–2007) raw price decline ("Raw") and the Japan Robot Shock measured by the fixed effect ψ^C_{o,t_1} in equation (11) ("JRS"). The occupation group is routine, service (manual), and abstract, where routine is further divided into production, transportation, and other.

imports from Japan, suggesting that robotization is a supply shock to the US economy. Figure Appendix B.3 shows the trends in robot stocks.

4.1. Parameter Estimates

Table 1 presents the estimates of the structural parameters. Column 1 shows the EoS parameter between robots and workers when constrained to be constant across occupation groups. The estimate of the within-occupation EoS between robots and labor θ is around 2, implying that robots and labor are substitutes within an occupation. The high estimate of EoS between labor and automation capital is also found in Eden and Gaggl (2018), who estimate the elasticity between ICT capital and labor. The point estimate of the EoS between occupations, β, is 0.83, indicating that the occupational

Table 1: Parameter estimates

	(1)			(2)		(3)		(4)	
	Constant θ			Main		Past wage		China shock	
			Production	2.71	(0.32)	2.95	(0.42)	3.03	(0.60)
			Transportation	1.76	(0.15)	2.90	(0.48)	2.01	(0.16)
θ_g	2.05	(0.19)	Other Routine	1.96	(0.17)	1.16	(0.32)	1.08	(0.28)
			Manual	1.01	(0.71)	1.23	(0.55)	1.16	(0.71)
			Abstract	1.01	(0.62)	0.64	(1.24)	1.00	(0.33)
β	0.83	(0.03)		0.73	(0.06)	0.73	(0.17)	1.18	(0.13)
F-stat	52.0			41.0		28.2		29.0	

Note: The structural parameter estimates based on the moment condition (15) and the two-stage optimal GMM estimates described in Appendix D.3 are shown. θ_g is the within-occupation elasticity of substitution (EoS) between robots and labor, while β is the EoS between occupations. Column (1) presents the results with the constraint that θ_g is constant across occupation groups. Column (2) presents the main results with θ_g allowed to be heterogeneous across five occupational groups. Column (3) presents the results of a sensitivity analysis using historical occupational wages. Column (4) presents the results of a sensitivity analysis using the China shock. Production, Transportation, and Other are the three subcategories of routine occupations. The plug-in optimal standard errors are presented in parentheses. F-statistics are calculated based on the post-estimation regression of the score functions evaluated at the estimated parameters on the Japan Robot Shock, described in detail in Appendix D.4

.

groups are complementary.[17] Appendix D.5 performs the model fit exercise and shows that it is critical to consider the automation shock when estimating the EoS between robots and labor.

[17]This estimate is higher than the central estimate of 0.49 in Humlum (2021), while it is lower than 1.78 in Burstein et al. (2019). I conduct the sensitivity analysis with respect to the values of β in Appendix E.4.

Column 2 presents the estimation result when heterogeneity is allowed across occupational groups. The EoS for routine-production occupations is 2.7. In contrast, those for other routine occupations (transportation and other routine) are close to 2, while those for other occupation groups are not significantly different from 1. Therefore, the routine-production occupation estimates indicate the particular vulnerability of workers in these occupations to robot capital. The estimate of EoS between occupations β does not change qualitatively between columns 1 and 2.

Consistent with the literature that estimates the capital-labor substitution elasticity, the source of identification of these large and heterogeneous EoS between robots and labor is the correlation between the JRS and the change in the labor market outcome. Intuitively, if θ_g is high, the steady-state relative demand for robots (or labor) responds strongly in the positive (or negative) direction to a unit decrease in the cost of using robots.[18]

Another potential cost shifter for occupational labor demand is the historical wage, which affects the contemporary incentive to adopt robots. To control for this effect, I consider an alternative measure of the JRS measured relative to the occupation wage in 1970. Column 3 of Table 1 shows the estimation result of this sensitivity analysis. In addition, I consider the role of the significant China trade shock during the sample period (Autor et al., 2013). To do so, I residualize the JRS by the measure of occupational exposure to Chinese imports before estimation.[19] The result is shown in Column 4. In both sensitivity analyses, the main message prevails: production work-

[18]This point is shown in a reduced-form analysis in Appendix C.1.

[19]See Appendix D.2 for how to construct the variable.

ers are particularly vulnerable to robots. I find an even larger estimate of the EoS for transportation occupations in column 3.

A related concern is that as the US is a large economy, its demand shock may affect robot prices in the international market, simultaneously driving US labor demand. To address this concern, I check the data from the Netherlands, a small open economy, in Appendix C.2, showing a similar empirical pattern to the US data.

4.2. Decomposing the Source of Task Automation

The estimated model's task allocation (3) allows me to recover the automation shock. Specifically, I obtain the implied automation shock by inverting (13), using the observed change in relative robot demand, the EoS estimates θ_g, and the change in the relative price of robots $x^R_{JP,US}\psi^J_o - \widehat{w_{US,o}}$.[20]

Figure 2a illustrates a scatterplot between the JRS and the automation shock, showing a slight positive relationship. This correlation is consistent with the example of robotic innovations discussed in Appendix B.1. Figure 2b summarizes the two shocks aggregated at the occupational group level. The figure shows 0.2-0.6 log points of the JRS, reflecting the decline in the price of robots from Japan. More importantly, the estimated automation shocks are positive and show greater variation across occupation groups. The two highly automated occupations, transportation and production, ob-

[20] The international price spillover term D in (13) is excluded because it is quantitatively small, as the contribution of robots to the national price index is small. This can be confirmed by substituting the implied shock in the model-implied price index change. Note that including D does not change the main results on distributional effects, because D is constant across occupations.

Figure 2: The Automation Shock, Japan Robot Shock, and the Total Automation

(a) Calibrated Automation Shock (b) Aggregates and Total Automation

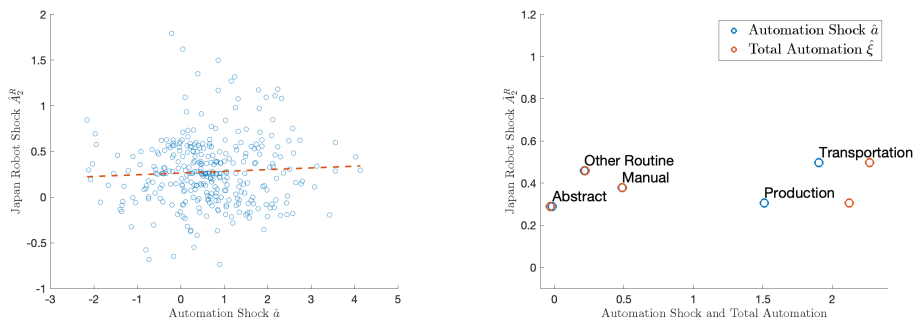

Note: The left panel shows the estimated automation shock (calibrated from Equation 13 and the estimated parameters in Table 1) on the horizontal axis and the Japan Robot Shock (obtained from the fixed effects in Equation 11) on the vertical axis. Each point is a 4-digit occupation, and the dashed line is the fitted line. The right panel adds total automation (implied by Equation 3) on the horizontal axis and shows the results at the occupation group level. Each occupation in the group is aggregated to the group level with the initial robot expenditure weight.

serve increases of 1.5-2 log points in robot task shares, whereas the other occupational groups experience a maximum of 0.5 log points.

Figure 2b also illustrates the total automation or change in the share of tasks performed by robots along the horizontal axis. Note that, according to (3), total automation can be driven by the exogenous change in the scale parameter of the Fréchet distribution a_o (the automation shock) and the endogenous reallocation of tasks due to the cheap robots caused by the JRS, $A^R_{2,o}$. In the two heavily robotized occupation groups, transportation and production, the total automation experiences as large as a 200% increase in the share of robotized tasks. This is driven by the automation shock and endogenous task allocation, although the former plays a more important role. There is no evidence of task allocation toward robots in other occupations

Figure 3: Robot Effects across Wage Distribution

(a) Estimated Automation Shock (b) Effect of the JRS on Wages

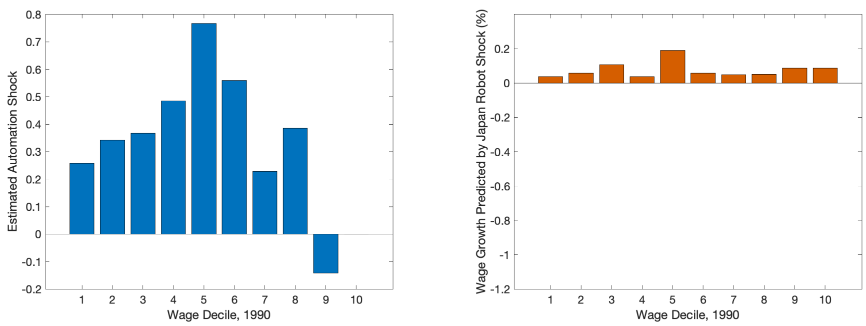

Note: The left panel shows the implied automation shock defined in Equation (13). The shocks are aggregated into 10 wage deciles in the base year 1990, weighted by initial employment levels. The right panel shows the annualized occupational wage growth rates for each wage decile predicted by the first-order approximated steady-state solution of the estimated model given in (A.32).

with less robotization.

5. The Effect of the JRS on Wage

First, I show the pattern of robot accumulation across the occupational wage distribution. Figure 3a shows the distribution of estimated automation shocks across baseline wage deciles. There is a pattern suggesting polarization–the automation shock hits the middle of the wage distribution harder compared with the bottom and top of the distribution.[21] Appendix E.1 summarizes the parameter values used in this quantitative exercise.

[21]Figure Appendix C.1a shows no correlation between the baseline wage and the JRS, contrasting with the result for the automation shock. These results suggest that it was the automation shock, not the JRS, that caused the dynamics of the wage distribution in the 1990s and 2000s.

Figure 3b illustrates the predicted steady-state wage growth per year due to JRS. Even though the automation shock falls in the middle of the wage distribution, the counterfactual wage growth rate is consistently positive across the initial wage distribution. There are two reasons. First, in the model, the reduction in the price of robots leads to an increase in the adoption of robots, which raises the marginal product of labor on average. Second, the elasticity of substitution between robots and workers shows large heterogeneity (Table 1), suggesting that workers in non-routine occupations are more complementary to robots. However, this aggregate effect masks important heterogeneity across different sources (occupations) of JRS. The decomposition exercises with different scenarios are carried out in Appendix E.5.

I also analyze the two robotization shocks (the automation shock $\widehat{\boldsymbol{a}}$ and the JRS $\widehat{\boldsymbol{A}}_2$) separately in another quantitative exercise. I find that the automation shock reduces labor demand by reallocating tasks from labor to robots, whereas the JRS increases the robot stock and the marginal product of labor. Appendix E.3 presents the detailed results.

Robot Tax Counterfactual Analysis. The estimated model provides insights into the short- and long-run effects of robot taxes on real wages across occupations and aggregate welfare losses. We find that introducing a robot tax generates redistributive effects to displaced production workers from complementary manual and abstract workers. In terms of aggregate income, there are small terms-of-trade gains from restricting the robot tax in the short run, which are more than offset by the loss in efficiency due to slower and lower robot accumulation. Appendix E.6 discusses these findings in detail.

6. Conclusion

This paper examines the wage effects of the increased use of industrial robots, considering that robots perform specific tasks and are traded internationally. I construct a measure of the cost reduction of buying robots from Japan (the JRS) across occupations in which robots are used. I then develop an open-economy general equilibrium model with automation within each occupation. To estimate the occupation-specific EoS between robots and labor of the model, I construct a MOIV of the JRS to address the correlation between the automation shock and the JRS, the key identification challenge.

The estimates of the within-occupation EoS between robots and labor are heterogeneous and are as high as 3 in production and material-moving occupations. These estimates are significantly larger than corresponding estimates in capital goods and labor, revealing the particular vulnerability of workers in production and material-moving occupations to robots. The model also implies that the JRS had little contribution to wage polarization across occupations in the US from 1990 to 2007. These results can be an important reference for policy discussions about industrial robots.

References

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings, in: Handbook of labor economics. Elsevier. volume 4, pp. 1043–1171.

Acemoglu, D., Restrepo, P., 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. American Economic Review 108, 1488–1542.

Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from us labor markets. Journal of Political Economy 128, 2188–2244.

Acemoglu, D., Restrepo, P., 2022. Demographics and automation. The Review of Economic Studies 89, 1–44.

Adachi, D., Kawaguchi, D., Saito, Y.U., 2024. Robots and employment: Evidence from japan, 1978-2017. Journal of Labor Economics .

Adao, R., Arkolakis, C., Esposito, F., 2023. General equilibrium effects in space: Theory and measurement .

Arrow, K.J., Chenery, H.B., Minhas, B.S., Solow, R.M., 1961. Capital-labor substitution and economic efficiency. The review of Economics and Statistics , 225–250.

Artuc, E., Bastos, P., Rijkers, B., 2023. Robots, tasks, and trade. Journal of International Economics 145, 103828.

Autor, D., Dorn, D., Hanson, G.H., 2013. The china syndrome: Local labor market effects of import competition in the united states. American Economic Review 103, 2121–68.

Blanchard, O.J., Kahn, C.M., 1980. The solution of linear difference models under rational expectations. Econometrica: Journal of the Econometric Society , 1305–1311.

Burstein, A., Morales, E., Vogel, J., 2019. Changes in between-group inequality: computers, occupations, and international trade. American Economic Journal: Macroeconomics 11, 348–400.

Caliendo, L., Dvorkin, M., Parro, F., 2019. Trade and labor market dynamics: General equilibrium analysis of the china trade shock. Econometrica 87, 741–835.

Caliendo, L., Parro, F., 2015. Estimates of the trade and welfare effects of nafta. The Review of Economic Studies 82, 1–44.

Caunedo, J., Jaume, D., Keller, E., 2023. Occupational exposure to capital-embodied technical change. American Economic Review 113, 1642–1685.

Cheng, W.J., 2018. Explaining job polariztion: The role of heterogeneity in capital intensity.

Cooper, R.W., Haltiwanger, J.C., 2006. On the nature of capital adjustment costs. The Review of Economic Studies 73, 611–633.

Dauth, W., Findeisen, S., Südekum, J., Woessner, N., 2017. German robots-the impact of industrial robots on workers .

Devol, J.G.C., 1961. Programmed article transfer. us 2988237a.

Eaton, J., Kortum, S., 2002. Technology, geography, and trade. Econometrica 70, 1741–1779.

Eden, M., Gaggl, P., 2018. On the welfare implications of automation. Review of Economic Dynamics 29, 15–43.

Galle, S., Lorentzen, L., 2024. The unequal effects of trade and automation across local labor markets. Journal of International Economics 150, 103912.

Graetz, G., Michaels, G., 2018. Robots at work. Review of Economics and Statistics 100, 753–768.

Hausman, J., Leonard, G., Zona, J.D., 1994. Competitive analysis with differenciated products. Annales d'Economie et de Statistique , 159–180.

Herrendorf, B., Herrington, C., Valentinyi, A., 2015. Sectoral technology and structural transformation. American Economic Journal: Macroeconomics 7, 104–133.

Hubmer, J., 2023. The race between preferences and technology. Econometrica 91, 227–261.

Humlum, A., 2021. Robot Adoption and Labor Market Dynamics. Technical Report.

Jaimovich, N., Saporta-Eksten, I., Siu, H., Yedid-Levi, Y., 2021. The macroeconomics of automation: Data, theory, and policy analysis. Journal of Monetary Economics 122, 1–16.

Karabarbounis, L., Neiman, B., 2014. The global decline of the labor share. The Quarterly journal of economics 129, 61–103.

Khandelwal, A., 2010. The long and short (of) quality ladders. The Review of Economic Studies 77, 1450–1476.

Kleinman, B., Liu, E., Redding, S.J., forthcoming. International friends and enemies. American Economic Journal: Macroeconomics .

Koch, M., Manuylov, I., Smolka, M., 2021. Robots and firms. The Economic Journal 131, 2553–2584.

Leigh, N.G., Kraft, B.R., 2018. Emerging robotic regions in the united states: insights for regional economic evolution. Regional Studies 52, 804–815.

Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., Dewhurst, M., et al., 2017. Harnessing automation for a future that works. McKinsey Global Institute .

Montobbio, F., Staccioli, J., Virgillito, M., Vivarelli, M., 2020. Robots and the origin of their labour-saving impact .

Morris, J., 2019. A weighted o*net keyword search (wws).

Oberfield, E., Raval, D., 2014. Micro data and macro technology. Technical Report. National Bureau of Economic Research.

Simonovska, I., Waugh, M.E., 2014. The elasticity of trade: Estimates and evidence. Journal of international Economics 92, 34–50.

de Souza, G., Li, H., 2023. Robots, tools, and jobs: Evidence from brazilian labor markets .

Tambe, P., Hitt, L.M., Rock, D., Brynjolfsson, E., 2019. It, ai and the growth of intangible capital. Available at SSRN 3416289 .

Traiberman, S., 2019. Occupations and import competition: Evidence from denmark. American Economic Review 109, 4260–4301.

Webb, M., 2019. The impact of artificial intelligence on the labor market. Available at SSRN 3482150 .

Online Appendix

Appendix A. Theory Appendix

Appendix A.1. The Full Model

The full model used for structural estimation extends that in the model section with worker dynamics, intermediate goods, and non-robot capital. Time is denoted by $t = 0, 1, \ldots$.

Workers' Problem. I formalize the assumptions behind the derivation and show (A.3) and (A.4). Workers are immobile across countries but choose occupations by solving a dynamic discrete choice problem (Humlum, 2021). Specifically, workers choose occupations that maximize lifetime utility based on switching costs and the draw of an idiosyncratic shock. The problem has a closed-form solution when the shock follows an extreme value distribution, which is a property used by the previous literature (e.g., Caliendo et al. (2019)).

Fix country i and period t. There is a mass $\overline{L}i, t$ of workers. At the beginning of each period, worker $\omega \in \left[0, \overline{L}i, t\right]$ draws a multiplicative idiosyncratic preference shock $\left\{Z_{i,o,t}\left(\omega\right)\right\} o$ that follows an independent Fr'echet distribution with scale parameter Ai, o, t^V and shape parameter $1/\phi$. To keep the expression simple, I focus on the case of an independent distribution. A worker ω works in the current occupation, earns income, consumes, and derives logarithmic utility, and then chooses the next period's occupation with the discount rate ι. When selecting the next period's occupation o, she pays an ad-valorem switching cost $\chi_{i,oo,t}$ in terms of a consumption unit that depends on the current occupation o. She consumes her income in each period. Thus, worker ω, who currently works in occupation o_t, maximizes the following objective function over the future stream of utilities by choosing occupations $\{o_s\}_{s=t+1}^{\infty}$:

$$E_t \sum_{s=t}^{\infty} \left(\frac{1}{1+\iota} \right)^{s-t} \left[\ln \left(C_{i,o_s,s} \right) + \ln \left(1 - \chi_{i,o_s o_{s+1},s} \right) + \ln \left(Z_{i,o_{s+1},s} \left(\omega \right) \right) \right] \quad \text{(A.1)}$$

where $C_{i,o,s}$ is a consumption bundle when working in occupation o in period $s \geq t$, and E_t is the expectation conditional on the value of $Z_{i,o_t,t}\left(\omega\right)$. Each worker owns an occupation-specific labor endowment $l_{i,o,t}$. Her income comprises labor income $w_{i,o,t}$ and an occupation-specific ad-valorem government transfer at the rate $T_{i,o,t}$. Given the consumption price $P_{i,t}^{G}$, the budget constraint is

$$P_{i,t}^{G} C_{i,o,t} = w_{i,o,t} l_{i,o,t} \left(1 + T_{i,o,t} \right) \quad \text{(A.2)}$$

for any worker, with $P_{i,t}^{G}$ denoting the price index of the non-robot good G.

Following a similar derivation as Caliendo et al. (2019), (A.1) and (A.2) imply the worker's optimization conditions, characterized by, for each country i and period t, the transition probability $\mu_{i,oo,t}$ from occupation o in period t to occupation o in period $t+1$, and the exponential expected value $V_{i,o,t}$ for occupation o, that satisfy

$$\mu_{i,oo',t} = \frac{\left(\left(1 - \chi_{i,oo',t} \right) \left(V_{i,o',t+1} \right)^{\frac{1}{1+\iota}} \right)^{\phi}}{\sum_{o''} \left(\left(1 - \chi_{i,oo'',t} \right) \left(V_{i,o'',t+1} \right)^{\frac{1}{1+\iota}} \right)^{\phi}}, \quad \text{(A.3)}$$

$$V_{i,o,t} = \widetilde{\Gamma} C_{i,o,t} \left[\sum_{o'} \left(\left(1 - \chi_{i,oo',t} \right) \left(V_{i,o',t+1} \right)^{\frac{1}{1+\iota}} \right)^{\phi} \right]^{\frac{1}{\phi}}, \quad \text{(A.4)}$$

respectively, where $C_{i,o,t+1}$ is the real consumption, $\chi_{i,oo,t}$ is an ad-valorem switching cost from occupation o to o, ϕ is the occupation-switch elasticity, and $\widetilde{\Gamma} \equiv \Gamma(1 - 1/\phi)$ is a constant that depends on the Gamma function $\Gamma(\cdot)$. For each i and t, the employment level satisfies the law of motion

$$L_{i,o,t+1} = \sum_{o'} \mu_{i,o'o,t} L_{i,o',t}. \tag{A.5}$$

Non-robot Good Producers' Problem. The producer's problem consists of two tiers—static optimization of labor input in each occupation and dynamic optimization of robot investment. The static part chooses employment conditional on market prices and the current stock of robot capital. Namely, for each i and t, conditional on the o-vector of the stock of robot capital $\left\{K^R_{i,o,t}\right\}_o$, producers solve

$$\pi_{i,t}\left(\left\{K^R_{i,o,t}\right\}_o\right) \equiv \max_{\{L_{i,o,t}\}_o} p^G_{i,t} Y^G_{i,t} - \sum_o w_{i,o,t} L_{i,o,t}, \tag{A.6}$$

where $Y^G_{i,t}$ is presented by the production function (1).

The dynamic optimization problem involves choosing the size of the robot investment, given the current stock of robot capital. It is derived from the following three assumptions. First, for each i, o, and t, robot capital $K^R_{i,o,t}$ accumulates according to

$$K^R_{i,o,t+1} = (1-\delta) K^R_{i,o,t} + Q^R_{i,o,t}, \tag{A.7}$$

where $Q^R_{i,o,t}$ is the amount of new robot investment and δ is the depreciation rate of robots. Second, new investment is represented by a CES aggregation of robot hardware from country l, $Q^R_{li,o,t}$, and non-robot goods input $I^{int}_{i,o,t}$, which is the input of software and integration or

$$Q^R_{i,o,t} = \left[\sum_l \left(Q^R_{li,o,t}\right)^{\frac{\varepsilon^R - 1}{\varepsilon^R}}\right]^{\frac{\varepsilon^R}{\varepsilon^R - 1}\alpha^R} \left(I^{int}_{i,o,t}\right)^{1-\alpha^R} \tag{A.8}$$

where l denotes the origin of the newly purchased robots, and α^R is the share of robot arms in investment costs. Discussion of the choice of the functional form of (A.8) is given in Appendix B.1. Third, robot installation is costly and requires a convex per-unit adjustment cost $\gamma Q^R_{i,o,t}/K^R_{i,o,t}$ measured in robot units, where γ governs the magnitude of the adjustment cost (e.g., Cooper and Haltiwanger, 2006), reflecting the sluggishness of robot adoption.

Given these assumptions, a producer of non-robot good G in a country i solves the dynamic optimization problem

$$\begin{aligned} \max_{\left\{\left\{Q^R_{li,o,t}\right\}_l, I^{int}_{i,o,t}\right\}_o} \sum_{t=0}^{\infty} \left(\frac{1}{1+\iota}\right)^t \Bigg[\pi_{i,t}\left(\left\{K^R_{i,o,t}\right\}_o\right) \\ - \sum_o \left(\sum_l p^R_{li,o,t}\left(1+u_{li,t}\right) Q^R_{li,o,t} + P^G_{i,t} I^{int}_{i,o,t} + \gamma P^R_{i,o,t} Q^R_{i,o,t} \frac{Q^R_{i,o,t}}{K^R_{i,o,t}}\right)\Bigg], \end{aligned} \tag{A.9}$$

subject to accumulation (A.7) and (A.8), and given $\left\{K^R_{i,o,0}\right\}_o$. A standard Lagrange multiplier method yields Euler equations for the investment, which are derived in Appendix A.2. Note that the Lagrangian multiplier $\lambda^R_{i,o,t}$ represents the equilibrium marginal value of robot capital.

The Cost of Using Robots and Robot Aggregation Function. I briefly mention the background for the functional form choice in the previous paragraph. A modern industrial robot typically does not have stand-alone hardware (e.g., robot joints and arms); instead, its ecosystem includes hardware and controllers driven by software (e.g., computers and robot programming language). Due to its complexity, the installation of robots in the production environment usually requires the hiring of expensive system integrators with the necessary engineering knowledge for integration. Therefore, the relevant costs of robots for users include hardware, software, and integration costs.

The average price measure of robots used in this paper should be interpreted as reflecting a portion of the total cost of robots. Following the literature convention due to data limitations on robot software and integration, I address this point in the model section by separately defining the observable hardware costs using my data and the unobserved components of the costs. Namely, (A.7) explicitly includes software and integration, reflecting the feature of modern industrial robots, which are typically not stand-alone hardware, but rather an ecosystem of controllers driven by software, requiring a significant amount of resources for integration.

Relatedly, (A.7) follows the capital goods trade formulation in the sense that robots are traded because they are differentiated by country of origin l. Note that (A.8) implies that the origin-differentiated investment good is first aggregated and then added to the capital stock after (A.7). This trick helps to reduce the number of capital stock variables and is also used in the international macroeconomic literature.

Intermediate Good Producers' Problem. The intermediate goods are the same goods as the non-robot goods, but are an input to the production function. The stock of non-robot capital is given exogenously in each period for each country, and producers rent non-robot capital from the rental market. The production function of the non-robot goods is represented by

$$Y_{i,t}^{G} = A_{i,t}^{G} \left\{ \alpha_{i,L} \left(T_{i,t}^{O} \right)^{\frac{\vartheta - 1}{\vartheta}} + \alpha_{i,M} \left(M_{i,t} \right)^{\frac{\vartheta - 1}{\vartheta}} + \alpha_{i,K} \left(K_{i,t} \right)^{\frac{\vartheta - 1}{\vartheta}} \right\}^{\frac{\vartheta}{\vartheta - 1}},$$

where ϑ is the EoS between occupation aggregates, intermediate goods, and non-robot capital, and $\alpha_{i,L}$, $\alpha_{i,M}$, and $\alpha_{i,K} \equiv 1 - \alpha_{i,L} - \alpha_{i,M}$ are cost share parameters for occupation aggregates, intermediate goods, and non-robot capital, respectively. The parameters satisfy $\vartheta > 0$ and $\alpha_{i,L}, \alpha_{i,M}, \alpha_{i,K} > 0$, and in the structural estimation I set $\vartheta = 1$ and compute each country's cost share parameters from the data. Intermediate goods are aggregated by

$$M_{i,t} = \left[\sum_{l} \left(M_{li,t}\right)^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{A.10}$$

where $\varepsilon > 0$ is the EoS between source countries. Since intermediate goods are traded between countries and aggregated by (A.10), the elasticity parameter ε plays the role of trade elasticity. The static decision of producers now includes the amount of non-robot capital to rent and the amount of intermediate goods to purchase from each source country.

Equilibrium. To close the model, the employment level must satisfy an adding-up constraint

$$\sum_{o} L_{i,o,t} = \overline{L}_{i,t}, \tag{A.11}$$

and market clearing conditions for robot and non-robot goods must hold. There is a numeraire good to fix the pricing system. The following defines a temporary equilibrium in each period, followed by a sequential equilibrium, leading to the definition of a steady state. The complete expressions are given in Appendix A.2.

I define the bold symbols as column vectors of robot capital $\boldsymbol{K}_t^R \equiv \left[K_{i,o,t}^R\right]_{i,o}$, marginal values of robot capital $\boldsymbol{\lambda}_t^R \equiv \left[\lambda_{i,o,t}^R\right]_{i,o}$, employment $\boldsymbol{L}_t \equiv \left[L_{i,o,t}\right]_{i,o}$, workers' value functions $\boldsymbol{V}_t \equiv \left[V_{i,o,t}\right]_{i,o}$, non-robot goods prices $\boldsymbol{p}_t^G \equiv \left[p_{i,t}^G\right]_i$, robot prices $\boldsymbol{p}_t^R \equiv \left[p_{i,o,t}^R\right]_{i,o}$, wages, $\boldsymbol{w}_t \equiv \left[w_{i,o,t}\right]_{i,o}$, bilateral non-robot goods trade levels $\boldsymbol{Q}_t^G \equiv \left[Q_{ij,t}^G\right]_{i,j}$, bilateral non-robot goods trade levels $\boldsymbol{Q}_t^R \equiv \left[Q_{ij,o,t}^R\right]_{i,j,o}$, and occupation transition shares $\boldsymbol{\mu}_t \equiv \left[\mu_{i,oo',t}\right]_{i,oo'}$, where $\boldsymbol{V}_t$ and $\boldsymbol{\mu}_t$ are explained in detail in Appendix A.1. I write $\boldsymbol{S}_t \equiv \left[\boldsymbol{K}_t^{R'}, \boldsymbol{\lambda}_t^{R'}, \boldsymbol{L}_t', \boldsymbol{V}_t'\right]'$ as state variables.

Definition 1. In each period t, given state variables $\boldsymbol{S}_t$, a *temporary equilibrium* (TE) $\boldsymbol{x}_t$ is the set of prices $\boldsymbol{p}_t \equiv \left[\boldsymbol{p}_t^{G'}, \boldsymbol{p}_t^{R'}, \boldsymbol{w}_t'\right]'$ and flow quantities $\boldsymbol{Q}_t \equiv \left[\boldsymbol{Q}_t^{G'}, \boldsymbol{Q}_t^{R'}, \boldsymbol{\mu}_t'\right]$ that satisfy: (i) given $\boldsymbol{p}_t$, workers choose occupations optimally by (A.3), (ii) given $\boldsymbol{p}_t$, producers maximize flow profit by (A.6) and demand robots by (A.17), and (iii) markets clear: labor adds up as in (A.11), and goods markets clear with trade balances as in (A.25) and (A.27).

In other words, the inputs to the TE are all state variables, while the outputs are all endogenous variables determined in each period. Adding the conditions on state variable transitions, a sequential equilibrium determines all state variables given initial conditions as follows.

Definition 2. Given initial robot capital stocks and employment $\left[\boldsymbol{K}_0^{R'}, \boldsymbol{L}_0'\right]'$, a *sequential equilibrium* (SE) is a sequence of vectors $\boldsymbol{y}_t \equiv [\boldsymbol{x}_t', \boldsymbol{S}_t']_t'$ that satisfies the TE conditions and the employment law of motion (**??**), the value function condition (A.4), capital accumulation (A.7), producer dynamic optimization (A.21) and (A.20).

Finally, I define the steady state as an SE $\boldsymbol{y}$ that does not change over time.

Appendix A.2. Equilibrium Characterization

To characterize the producer problem, I first show the static optimization conditions and then the dynamic ones. For simplicity, I focus on the case with $\vartheta = 1$, or Cobb-Douglas in the mix of occupation aggregates, intermediate goods, and non-robot capital. To solve the static problem of labor, intermediate goods, and non-robot capital, consider the first-order conditions (FOCs) of (A.6)

$$p_{i,t}^{G}\alpha_{i,L}\frac{Y_{i,t}^{G}}{T_{i,t}^{O}}\left(b_{i,o,t}\frac{T_{i,t}^{O}}{T_{i,o,t}^{O}}\right)^{\frac{1}{\beta}}\left(\left(1-a_{o,t}\right)\frac{T_{i,o,t}^{O}}{L_{i,o,t}}\right)^{\frac{1}{\theta_{o}}}=w_{i,o,t}, \tag{A.12}$$

where $T_{i,t}^{O}$ is the aggregated occupations $T_{i,t}^{O}\equiv\left[\sum_{o}\left(T_{i,o,t}^{O}\right)^{(\beta-1)/\beta}\right]^{\beta/(\beta-1)}$,

$$p_{i,t}^{G}\alpha_{i,M}\frac{Y_{i,t}^{G}}{M_{i,t}}\left(\frac{M_{i,t}}{M_{li,t}}\right)^{\frac{1}{\varepsilon}}=p_{li,t}^{G}, \tag{A.13}$$

and

$$p_{i,t}^{G}\alpha_{i,K}\frac{Y_{i,t}^{G}}{K_{i,t}}=r_{i,t}, \tag{A.14}$$

where $\alpha_{i,K}\equiv 1-\alpha_{i,L}-\alpha_{i,M}$. Note also that by the envelope theorem,

$$\frac{\partial\pi_{i,t}\left(\left\{K_{i,o,t}^{R}\right\}\right)}{\partial K_{i,o,t}^{R}}=p_{i,t}^{G}\frac{\partial Y_{i,t}}{\partial K_{i,o,t}^{R}}=p_{i,t}^{G}\left(\alpha_{L}\frac{Y_{i,t}^{G}}{T_{i,t}^{O}}\left(b_{i,o,t}\frac{T_{i,t}^{O}}{T_{i,o,t}^{O}}\right)^{\frac{1}{\beta}}\left(a_{o,t}\frac{T_{i,o,t}^{O}}{K_{i,o,t}^{R}}\right)^{\frac{1}{\theta}}\right). \tag{A.15}$$

Another static problem for producers is robot purchase. Define the "before-integration" robot aggregate

$$Q_{i,o,t}^{R,BI}\equiv\left[\sum_{l}\left(Q_{li,o,t}^{R}\right)^{\frac{\varepsilon^{R}-1}{\varepsilon^{R}}}\right]^{\frac{\varepsilon^{R}}{\varepsilon^{R}-1}}$$

and the corresponding price index $P_{i,o,t}^{R,BI}$. By the first order condition with respect to $Q_{li,o,t}^{R}$ for (A.8), I have

$$p_{li,o,t}^{R}Q_{li,o,t}^{R}=\left(\frac{p_{li,o,t}^{R}}{P_{i,o,t}^{R,BI}}\right)^{1-\varepsilon^{R}}P_{i,o,t}^{R,BI}Q_{i,o,t}^{R,BI}$$

and $P^{R,BI}_{i,o,t} Q^{R,BI}_{i,o,t} = \alpha P^R_{i,o,t} Q^R_{i,o,t}$. Thus,

$$p^R_{li,o,t} Q^R_{li,o,t} = \alpha \left(\frac{p^R_{li,o,t}}{P^{R,BI}_{i,o,t}} \right)^{1-\varepsilon^R} P^R_{i,o,t} Q^R_{i,o,t}.$$

Hence, I have

$$Q^R_{li,o,t} = \alpha \left(p^R_{li,o,t} \right)^{-\varepsilon^R} \left(P^{R,BI}_{i,o,t} \right)^{\varepsilon^R - 1} P^R_{i,o,t} Q^R_{i,o,t}.$$

Using $P^R_{i,o,t} = \left(P^{R,BI}_{i,o,t} \right)^{\alpha^R} (P_{i,t})^{1-\alpha^R}$, I have

$$Q^R_{li,o,t} = \alpha \left(\frac{p^R_{li,o,t}}{P^{R,BI}_{i,o,t}} \right)^{-\varepsilon^R} \left(\frac{P^{R,BI}_{i,o,t}}{P_{i,t}} \right)^{-\left(1-\alpha^R\right)} Q^R_{i,o,t}.$$

Alternatively, one can define the robot price index by

$$\widetilde{P}^R_{i,o,t} = \alpha^{\frac{1}{\varepsilon^R}} \left(P^{R,BI}_{i,o,t} \right)^{\frac{\varepsilon^R - \left(1-\alpha^R\right)}{\varepsilon^R}} P_{i,t}^{\frac{1-\alpha^R}{\varepsilon^R}}$$

and show

$$Q^R_{li,o,t} = \left(\frac{p^R_{li,o,t}}{\widetilde{P}^R_{i,o,t}} \right)^{-\varepsilon^R} Q^R_{i,o,t}, \tag{A.16}$$

which is a standard gravity representation of robot trade.

To solve the dynamic problem, set up the (current-value) Lagrangian function for non-robot goods producers

$$\mathcal{L}_{i,t} = \sum_{t=0}^{\infty} \left\{ \left(\frac{1}{1+\iota} \right)^t \left[\pi_{i,t} \left(\left\{ K^R_{i,o,t} \right\}_o \right) - \sum_{l,o} \left(p^R_{li,o,t} \left(1 + u_{li,t}\right) Q^R_{li,o,t} + P^G_{i,t} I^{int}_{i,o,t} + \gamma P^R_{i,o,t} Q^R_{i,o,t} \frac{Q^R_{i,o,t}}{K^R_{i,o,t}} \right) \right] \right.$$
$$\left. - \lambda^R_{i,o,t} \left\{ K^R_{i,o,t+1} - (1-\delta) K^R_{i,o,t} - Q^R_{i,o,t} \right\} \right.$$

Taking the FOC with respect to the hardware from country l, $Q^R_{li,o,t}$, I have

$$p^R_{li,o,t} \left(1 + u_{li,t}\right) + 2\gamma P^R_{i,o,t} \left(\frac{Q^R_{i,o,t}}{K^R_{i,o,t}} \right) \frac{\partial Q^R_{i,o,t}}{\partial Q^R_{li,o,t}} = \lambda^R_{i,o,t} \frac{\partial Q^R_{i,o,t}}{\partial Q^R_{li,o,t}}. \tag{A.17}$$

Taking the FOC with respect to the integration input $I^{int}_{i,o,t}$, I have

$$P^G_{i,t} + 2\gamma P^R_{i,o,t}\left(\frac{Q^R_{i,o,t}}{K^R_{i,o,t}}\right)\frac{\partial Q^R_{i,o,t}}{\partial I^{int}_{i,o,t}} = \lambda^R_{i,o,t}\frac{\partial Q^R_{i,o,t}}{\partial I^{int}_{i,o,t}}, \tag{A.18}$$

Taking the FOC with respect to $K^R_{i,o,t+1}$, I have

$$\left(\frac{1}{1+\iota}\right)^{t+1}\left[\frac{\partial \pi_{i,t+1}\left(\left\{K^R_{i,o,t+1}\right\}_o\right)}{\partial K^R_{i,o,t+1}} + \gamma P^R_{i,o,t+1}\left(\frac{Q^R_{i,o,t+1}}{K^R_{i,o,t+1}}\right)^2 + (1-\delta)\,\lambda^R_{i,o,t+1}\right] - \left(\frac{1}{1+\iota}\right)^t \lambda^R_{i,o,t} = 0, \tag{A.19}$$

and the transversality condition: for any j and o,

$$\lim_{t\to\infty} e^{-\iota t}\lambda^R_{j,o,t}K^R_{j,o,t+1} = 0. \tag{A.20}$$

Rearranging equation (A.19), I obtain the following Euler equation.

$$\lambda^R_{i,o,t} = \frac{1}{1+\iota}\left[(1-\delta)\,\lambda^R_{i,o,t+1} + \frac{\partial}{\partial K^R_{i,o,t+1}}\pi_{i,t+1}\left(\left\{K^R_{i,o,t+1}\right\}\right) + \gamma p^R_{i,o,t+1}\left(\frac{Q^R_{i,o,t+1}}{K^R_{i,o,t+1}}\right)^2\right]. \tag{A.21}$$

Turning to the demand for non-robot goods, in the following I characterize bilateral intermediate goods trade demand and total expenditure. Let $X^G_{j,t}$ be the total quantity (but not the value) purchased of good G in country j in period t. By (A.10), the bilateral trade demand is given by

$$p^G_{ij,t}Q^G_{ij,t} = \left(\frac{p^G_{ij,t}}{P^G_{j,t}}\right)^{1-\varepsilon} P^G_{j,t}X^G_{j,t}, \tag{A.22}$$

for all i, j, and t. In this equation, $P^G_{j,t}X^G_{j,t}$ is the total expenditure on non-robot goods. The total expenditure is the sum of final consumption $I_{j,t}$, payment for intermediate goods $\alpha_M p^G_{j,t}Y^G_{j,t}$, input for robot production $\sum_o P^G_{j,t}I^R_{j,o,t} = \sum_{o,k} p^R_{jk,o,t}Q^R_{jk,o,t}$, and payment to robot integration $\sum_o P^G_{j,t}I^{int}_{j,o,t} =$

$\left(1-\alpha^R\right)\sum_o P^R_{j,o,t}Q^R_{j,o,t}$. Therefore, I have

$$P^G_{j,t}X^G_{j,t} = I_{j,t} + \alpha_M p^G_{j,t}Y^G_{j,t} + \sum_{o,k} p^R_{jk,o,t}Q^R_{jk,o,t} + \left(1-\alpha^R\right)\sum_o P^R_{j,o,t}Q^R_{j,o,t}.$$

For country j and period t, by substituting into income $I_{j,t}$ the period cash flow of the non-robot good producer that satisfies

$$\Pi_{j,t} \equiv \pi_{j,t}\left(\left\{K^R_{j,o,t}\right\}_o\right) - \sum_{i,o}\left(p^R_{ij,o,t}\left(1+u_{ij,t}\right)Q^R_{ij,o,t} + \sum_o P^G_{j,t}I^{int}_{j,o,t} + \gamma P^R_{j,o,t}Q^R_{j,o,t}\left(\frac{Q^R_{j,o,t}}{K^R_{j,o,t}}\right)\right)$$

and robot tax revenue $T_{j,t} = \sum_{i,o} u_{ij,t}p^R_{ij,o,t}Q^R_{ij,o,t}$, I have

$$I_{j,t} = \left(1-\alpha_M\right)\sum_k p^G_{jk,t}Q^G_{jk,t} - \left(\sum_{i,o} p^R_{ij,o,t}Q^R_{ij,o,t} + \left(1-\alpha^R\right)\sum_o P^R_{j,o,t}Q^R_{j,o,t}\right). \tag{A.23}$$

I can rewrite this in terms of variables in the definition of equilibrium to have

$$I_{j,t} = \left(1-\alpha_M\right)\sum_k p^G_{jk,t}Q^G_{jk,t} - \frac{1}{\alpha^R}\sum_{i,o} p^R_{ij,o,t}Q^R_{ij,o,t}.$$

Thus, the total expenditure measured on the production side, as opposed to the income side, is

$$P^G_{j,t}X^G_{j,t} = \sum_k p^G_{jk,t}Q^G_{jk,t} - \sum_{i,o} p^R_{ij,o,t}Q^R_{ij,o,t}\left(1+\gamma\frac{Q^R_{ij,o,t}}{K^R_{j,o,t}}\right). \tag{A.24}$$

Note that this equation embeds the balanced trade condition. By substituting (A.24) into the (A.22), I have

$$p^G_{ij,t}Q^G_{ij,t} = \left(\frac{p^G_{ij,t}}{P^G_{j,t}}\right)^{1-\varepsilon^G}\left(\sum_k p^G_{jk,t}Q^G_{jk,t} + \sum_{k,o} p^R_{jk,o,t}Q^R_{jk,o,t} - \sum_{i,o} p^R_{ij,o,t}Q^R_{ij,o,t}\right). \tag{A.25}$$

The good and robot-o market-clearing conditions are given by,

$$Y^R_{i,t} = \sum_j Q^G_{ij,t}\tau^G_{ij,t}, \tag{A.26}$$

for all i and t, and

$$p_{i,o,t}^{R}=\frac{P_{i,t}^{G}}{A_{i,o,t}^{R}} \tag{A.27}$$

for all i,o, and t, respectively.

Conditional on state variables $\boldsymbol{S}_t=\left\{\boldsymbol{K}_t^R,\boldsymbol{\lambda}_t^R,\boldsymbol{L}_t,\boldsymbol{V}_t\right\}$, (A.3), (A.12), (A.17), (A.25), (A.26), and (A.27) characterize the TE $\left\{\boldsymbol{p}_t^G,\boldsymbol{p}_t^R,\boldsymbol{w}_t,\boldsymbol{Q}_t^G,\boldsymbol{Q}_t^R,\boldsymbol{L}_t\right\}$. In addition, conditional on initial conditions $\left\{\boldsymbol{K}_0^R,\boldsymbol{L}_0\right\}$, (A.7), (A.21), and (A.20) characterize the SE.

Finally, the steady-state conditions are provided by imposing the time-invariance condition on (A.7) and (A.21):

$$Q_{i,o}^{R}=\delta K_{i,o}^{R}, \tag{A.28}$$

$$\frac{\partial}{\partial K_{i,o}^{R}}\pi_i\left(\left\{K_{i,o}^{R}\right\}\right)=\left(\iota+\delta\right)\lambda_{i,o}^{R}-\sum_{l}\gamma p_{li,o}^{R}\left(\frac{Q_{li,o}^{R}}{K_{i,o}^{R}}\right)^{2}\equiv c_{i,o}^{R}. \tag{A.29}$$

Note that (A.29) can be interpreted as the flow marginal profit of capital equal to the flow marginal cost. Thus, I define the steady state marginal cost of robot capital $c_{i,o}^{R}$ from the right hand side of (A.29). Note that if there are no adjustment costs $\gamma=0$, the steady state Euler equation (A.29) implies

$$\frac{\partial}{\partial K_{i,o}^{R}}\pi_i\left(\left\{K_{i,o}^{R}\right\}\right)=c_{i,o}^{R}=\left(\iota+\delta\right)\lambda_{i,o}^{R},$$

which states that the marginal profit of capital is equal to the user cost of robots in the steady state.

Appendix A.3. The First-Order Approximation of the General Equilibrium

Since the GE system is highly nonlinear and does not have a closed-form solution due to the flexible robot-labor substitution, the equilibrium system of equations is log-linearized around the initial steady state. Consider the growth of the robot task space $a_{o,t}$ and the productivity of the robot production $A^R_{i,o,t}$ in the initial period t_0, and combine all these changes into a column vector $\boldsymbol{\Delta}$. Write state variables $\boldsymbol{S}_t = \left[\boldsymbol{K}^{R\prime}_t, \boldsymbol{\lambda}^{R\prime}_t, \boldsymbol{L}'_t, \boldsymbol{V}'_t\right]'$, and use "hat" notation to denote changes from t_0, or $\widehat{z}_t \equiv \ln(z_t) - \ln(z_{t_0})$ for any variable z_t,. I take the following three steps to solve the model.

Step 1. In a given period t, I combine the vector of shocks $\boldsymbol{\Delta}$ and (given) changes in state variables $\widehat{\boldsymbol{S}_t}$ into a column vector $\widehat{\boldsymbol{A}_t} = \left[\boldsymbol{\Delta}', \widehat{\boldsymbol{S}_t}'\right]'$. Log-linearizing the TE conditions, I solve for matrices $\overline{\boldsymbol{D}^x}$ and $\overline{\boldsymbol{D}^A}$ such that the log-difference of the TE $\widehat{\boldsymbol{x}}_t$ satisfies

$$\overline{\boldsymbol{D}^x}\widehat{\boldsymbol{x}}_t = \overline{\boldsymbol{D}^A}\widehat{\boldsymbol{A}_t}. \tag{A.30}$$

In this equation, $\overline{\boldsymbol{D}^x}$ is a substitution matrix, and $\overline{\boldsymbol{D}^A}\widehat{\boldsymbol{A}_t}$ is a vector of partial equilibrium shifts in period t Adao et al. (2023).[22]

Step 2. Log-linearizing the laws of motion and the Euler equations around the initial steady state, I solve for the matrices $\overline{\boldsymbol{D}^{y,SS}}$ and $\overline{\boldsymbol{D}^{\Delta,SS}}$ such that $\overline{\boldsymbol{D}^{y,SS}}\widehat{\boldsymbol{y}} = \overline{\boldsymbol{D}^{\Delta,SS}}\boldsymbol{\Delta}$, where the superscript SS denotes the steady state. Note that there is a block separation $\overline{\boldsymbol{D}^A} = \left[\overline{\boldsymbol{D}^{A,\Delta}}|\overline{\boldsymbol{D}^{A,S}}\right]$, so the equation (A.30) can be written as

$$\overline{\boldsymbol{D}^x}\widehat{\boldsymbol{x}}_t - \overline{\boldsymbol{D}^{A,S}}\widehat{\boldsymbol{S}_t} = \overline{\boldsymbol{D}^{A,\Delta}}\boldsymbol{\Delta}. \tag{A.31}$$

Combined with this equation evaluated at the steady state, I have

$$\overline{\boldsymbol{E}^y}\widehat{\boldsymbol{y}} = \overline{\boldsymbol{E}^\Delta}\boldsymbol{\Delta}, \tag{A.32}$$

[22] Because the TE vector $\widehat{\boldsymbol{x}_t}$ includes wages $\widehat{\boldsymbol{w}_t}$, (A.30) generalizes the general equilibrium comparative statics formulation in Adao et al. (2023), who consider the variant of (A.30) with $\widehat{\boldsymbol{x}_t} = \widehat{\boldsymbol{w}_t}$.

where

$$\overline{\boldsymbol{E}^{y}} \equiv \begin{bmatrix} \overline{\boldsymbol{D}^{x}} & -\overline{\boldsymbol{D}^{A,T}} \\ \multicolumn{2}{c}{\overline{\boldsymbol{D}^{y,SS}}} \end{bmatrix}, \text{ and } \overline{\boldsymbol{E}^{\Delta}} \equiv \begin{bmatrix} \overline{\boldsymbol{D}^{A,\Delta}} \\ \overline{\boldsymbol{D}^{\Delta,SS}} \end{bmatrix},$$

which implies $\widehat{\boldsymbol{y}} = \overline{\boldsymbol{E}}\boldsymbol{\Delta}$, where the matrix $\overline{\boldsymbol{E}} = \left(\overline{\boldsymbol{E}^{y}}\right)^{-1}\overline{\boldsymbol{E}^{\Delta}}$ represents the approximate first-order steady-state effect of the shock $\boldsymbol{\Delta}$. This steady-state matrix $\overline{\boldsymbol{E}}$ will be a key object in the estimation of the model in Section 3.2.

Step 3. Log-linearizing the laws of motion and the Euler equations around the new steady state, I solve for the matrices $\overline{\boldsymbol{D}_{t+1}^{y,TD}}$ and $\overline{\boldsymbol{D}_{t}^{y,TD}}$ such that $\overline{\boldsymbol{D}_{t+1}^{y,TD}}\check{\boldsymbol{y}}_{t+1} = \overline{\boldsymbol{D}_{t}^{y,TD}}\check{\boldsymbol{y}}_{t}$, where the superscript TD stands for transition dynamics and $\check{z}_{t+1} \equiv \ln z_{t+1} - \ln z'$ and z' is the new steady state value for each variable z. Log-linearized SE satisfies the following first-order difference equation

$$\overline{\boldsymbol{F}_{t+1}^{y}}\widehat{\boldsymbol{y}_{t+1}} = \overline{\boldsymbol{F}_{t}^{y}}\widehat{\boldsymbol{y}}_{t} + \overline{\boldsymbol{F}_{t+1}^{\Delta}}\boldsymbol{\Delta}. \tag{A.33}$$

Following the insights in Blanchard and Kahn (1980), there is a convergent matrix representing the first-order transition dynamics $\overline{\boldsymbol{F}_{t}}$ such that

$$\widehat{\boldsymbol{y}}_{t} = \overline{\boldsymbol{F}_{t}}\boldsymbol{\Delta} \text{ and } \overline{\boldsymbol{F}_{t}} \to \overline{\boldsymbol{E}}. \tag{A.34}$$

The matrix $\overline{\boldsymbol{F}_{t}}$ characterizes the transition dynamics after robotization shocks and is used to study the effect of policy changes in counterfactual analyses.

Appendix B. Data Appendix

Appendix B.1. Details about Industrial Robots

Industrial robots are defined as multi-axis manipulators. Following the International Organization for Standardization (ISO), this paper defines robots

Figure Appendix B.1: Examples of Industrial Robots

(a) Spot Welding (b) Material Handling

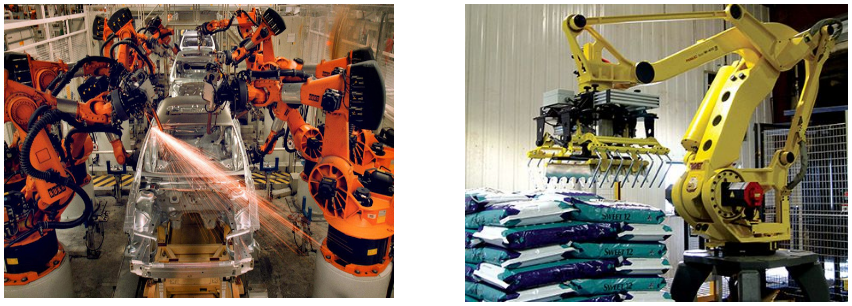

Sources: Autobot Systems and Automation (https://www.autobotsystems.com) and PaR Systems (https://www.par.com)

as "automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications" (ISO 8373:2012). This section provides a detailed discussion of such industrial robots. This definition excludes from the scope of this paper any automation devices that do not have multiple axes, even though some of them are often referred to as "robots" (e.g., Roomba, an autonomous home vacuum cleaner manufactured by iRobot Corporation). Figure Appendix B.1 shows examples of industrial robots that are used extensively in the manufacturing process and are considered in this paper. Spot welding and material handling robots are shown in the left and right panels, respectively.

Japan is a major innovator, manufacturer, and exporter of robots. As of 2017, the US had imported 5 billion dollars worth of Japanese robots, accounting for about one-third of the robots used in the US. Thus, the cost reduction of Japanese robots has a significant impact on the adoption of robots in the US and around the world.

JARA Robot Applications. The robot applications available in the JARA data are the following: Die casting; Forging; Resin molding; Pressing; Arc welding; Spot welding; Laser welding; Painting; Loading/unloading; Mechanical cutting; Polishing and deburring; Gas cutting; Laser cutting; Water jet cutting; General assembly; Inserting; Mounting; Bonding; Soldering; Sealing and gluing; Screwdriving; Picking, alignment and packaging; Palletizing; Measurement/inspection/testing; and Material handling.

Since robots are characterized by their versatility, unlike older specified industrial machines, the question arises as to whether robots can be classified as one of these applications (Kawasaki Heavy Industry, 2018). Although a robot can be reprogrammed to perform more than one task, I argue that robots can be classified as one of the above applications because the level of dexterity is different. Robots may be able to adapt to a model change of products, but they are not designed to perform other tasks beyond the 4-digit occupation level. Since small and medium enterprises are mostly high-mix, low-volume producers, robots are still too rigid to be transferred from one job to another at a reasonable cost. Because of this technological bottleneck, a versatile robot capable of replacing a wide range of workers at the 4-digit occupation level is not feasible for the sample period of this study.

Examples of Robotics Innovation. In section 2.2, the automation shock is defined as the change in the robot task space $a_{o,t}$, and the cost shock of producing robots is defined as the robot producer's total factor productivity (TFP) shock $A^R_{l,o,t}$. This paragraph provides examples of changes in robot technology and new patents to help understand these interpretations. An example of task space expansion is the introduction of the *Programmed Article Transfer* (PAT, Devol, 1961). The PAT is a machine that moves objects by the "teaching and playback" method, which requires a one-time teaching of how to move, after which the machine repeatedly and automatically plays back the movement. This feature frees workers from repetitive tasks. PAT has been intensively applied to spot welding tasks. Kawasaki Heavy Industry (2018) reports that of 4,000 spot welding points, 30% were previously performed by humans, which the PAT then took over. Therefore, I interpret the adoption of PAT as an example of expanding the task space of robots by increasing $a_{o,t}$.

An example of cost reduction is the introduction of the *Programmable Universal Manipulator for Assembly* (PUMA). The PUMA was designed to quickly and accurately transport, handle, and assemble automotive accessories. This was made possible by a new computer language, *Variable Assembly Language (VAL)*, which made the teaching process less complicated and more sophisticated. In other words, PUMA performed tasks previously performed by other robots, but at a lower unit cost per task unit.

It is also worth noting that the introduction of a new robot brand typically includes both components of innovation (task space expansion and cost reduction). For example, PUMA also expanded the task space of robots because VAL enabled the use of sensors and "expanded the range of applications to include assembly, inspection, palletizing, resin casting, arc welding, sealing, and research" (Kawasaki Heavy Industry, 2018).

Appendix B.2. More on Data Sources

*Details on the O*NET Code Connector Search.* From the O*NET Code Connector search, we use the match score generated by the *weighted search algorithm* used by the O*NET Code Connector. The weighted search algorithm is an internal search algorithm developed and used by O*NET since September 2005. Since then, O*NET has continually updated the algorithm and improved the quality of the search results. Morris (2019) reports that the latest weighted search algorithm scores 95.9% based on the position and score of the target best 4-digit occupation for a given query, a significant improvement over the previous search algorithm.

Additional Data Sources. In addition to JARA and O*NET data, I use data from IFR, BACI, WIOT, IPUMS USA, and CPS. IFR is a standard data source for industrial robot adoption in several countries (e.g., Graetz and Michaels (2018); Acemoglu and Restrepo, 2020, hereafter AR) to which JARA contributes Japan's robot data.[23] I use IFR's data to show total robot adoption in each destination country. BACI provides disaggregated trade flow data for more than 5000 products and 200 countries, from which the measures of international trade in industrial robots and baseline trade shares are derived. I used data from WIOT for the year closest to the starting year, 1992, to obtain intermediate input shares. IPUMS USA collects and harmonizes U.S. census microdata. I use censuses (1970, 1980, 1990, and 2000) and American Community Surveys (ACS, 2006-2008 3-year sam-

[23]As of August 2020, JARA includes 381 member companies, with 54 full members, 205 associate members, and 112 supporting members.

ple and 2012-2016 5-year sample). Occupational wages, employment, and labor cost shares are obtained from these data sources.

I focus on occupational codes that existed between the 1970 census and the 2007 ACS, which covers the sample period and the period before the trend analysis, in order to obtain consistent data across time periods. Thus, this paper focuses on intensive substitution in occupations, as opposed to the extensive effect of automation that creates new labor-intensive tasks and occupations, as in Acemoglu and Restrepo (2018). My dataset shows that 88.7% of workers in 2007 worked in the same occupations that existed in 1990. How to attribute the creation of new occupations to different types of automation goods, such as occupational robots in my case, remains an open question.

I follow Autor et al. (2013) for the census/ACS data cleaning procedure. Namely, I extract the 1970, 1980, 1990, and 2000 censuses, the 2006-2008 3-year ACS file, and the 2012-2016 5-year ACS file from the Integrated Public Use Micro Samples. For each file, I select all workers with occupation code OCC2010 between the ages of 16 and 64 who are not institutionalized. I compute the education share for each occupation by the share of workers with more than "any year of college," and the foreign-born share by the share of workers whose birthplace is neither in the U.S. nor in the U.S. extraterritorial areas/territories. I calculate hours worked by multiplying the usual weeks worked by hours worked per week. For 1970, I use the median in each bin of the usual weeks worked variable and assume that all workers worked 40 hours per week because the hours variable does not exist. I compute the hourly wage by first imputing the top-coded values for each state year by multiplying by 1.5 and dividing by the hours worked. To remove outliers, I take wages below the first percentile of the distribution in each year and set the maximum wage as the top-coded wage divided by 1,500. I compute the real wage in 2000 dollars by multiplying the CPI99 variable prepared by IPUMS. The person weight variable is used to aggregate all of these variables to the occupation level.

Occupational groups are formally defined as follows: Routine occupations include occupations such as production, transportation, sales, clerical, and administrative support. Abstract occupations include professional, managerial, and technical occupations. Service occupations include protective services, food preparation, cleaning, personal care, and personal services. Routine occupations are further divided into production, transportation, and other. This results in the following five categories in terms of US Census OCC2010 codes: Routine Production Occupations are in $[7700, 8965]$, Routine Transportation Occupations are in $[9000, 9750]$, Routine Other Occupations are in $[4700, 6130]$, Service (Manual) Occupations are in $[3700, 4650]$, and Abstract Occupations are in $[10, 3540]$.

Furthermore, following the idea of Caliendo et al. (2019), I use the bilateral occupational flow data to estimate the model with a dynamic discrete choice of occupation by workers. Specifically, I obtain the 1976 Annual Social and Economic Supplement (ASEC) of the CPS. For each year, I select all workers with occupation code 2010 for the current (OCC2010) and previous year (OCC10LY), aged 16 to 64, who are not institutionalized. I then construct variables using the same method as for the Census/ACS above. I assume that workers do not move between 4-digit occupations within the 5

Table Appendix B.1: List of Data Sources

Variable	Description	blackSource
$\widetilde{y}^G_{ij,t_0}, \widetilde{x}^G_{ij,t_0}, \widetilde{y}^R_{ij,t_0}, \widetilde{x}^R_{ij,t_0}$	Trade shares of goods and robots	BACI, IFR
$\widetilde{x}^O_{i,o,t_0}$	Occupation cost shares	IPUMS
l_{i,o,t_0}	Labor shares within occupation	JARA, IFR, IPUMS
$s^G_{i,t_0}, s^V_{i,t_0}, s^R_{i,t_0}$	Robot expenditure shares	BACI, IFR, WIOT
$\alpha_{i,M}$	Intermediate input share	WIOT

occupational groups defined in section 3.2, but rather between the 5 groups. I also assume that workers draw a target 4-digit occupation from the initial year's occupational employment distribution within the target group when switching occupations. With these data, I compute the probability of changing occupation by year.

Data on Initial Shares Used in Simulations. I need the data baseline share because the log-linearized sequential equilibrium solution depends on the initial steady-state shares. I define $t_0 = 1992$ and use annual frequency data. I consider the world consisting of three countries $\{USA, JPN, ROW\}$. Table Appendix B.1 summarizes the variable notations, descriptions, and data sources. I take the matrices of trade in goods and robots using BACI data. As in Acemoglu and Restrepo (2022), I measure robots by HS code 847950 ("Industrial Robots For Multiple Uses") and approximate the starting year value by the year 1998, when the robot HS code was first available.

The domestic robot adoption data are obtained by taking the flow quantity variable and the aggregate price variable from the IFR data for a selected set of countries. I then multiply these to obtain the U.S. and Japanese robot adoption values. For robot prices in ROW, I take the simple average of prices among the set of countries (France, Germany, Italy, South Korea, and the UK, as well as Japan and the US) for which prices are available in 1999, the earliest year for which price data are available. Graetz and Michaels (2018) discusses robot prices using the same data source. Figure Appendix B.2 shows the comparison of the available US price index measure between JARA and IFR. The JARA measures are disaggregated by 4-digit occupations. The figure shows the 10th, 50th (median), and 90th percentiles for each year, as in figure 1a. All measures are normalized to 1999, the year the first price measure is available in the IFR data. Overall, the JARA price trend variation tracks the IFR price trend quite well: The long-term trends from 1999 to the late 2010s are similar between the JARA median price and the IFR price index. During the 2000s, the IFR price index fell faster than the JARA data median price. This may be due to technological changes in robotics in countries other than Japan during the corresponding period.

I construct occupation cost shares $\widetilde{x}^{O}_{i,o,t_0}$ and labor shares within occupation l_{i,o,t_0} as follows. To measure $\widetilde{x}^{O}_{i,o,t_0}$, I aggregate the total wage income of workers primarily employed in each occupation o in year 1990, the census year closest to t_0. I then take the share of this total labor compensation measure for each occupation. Total labor compensation as a share of total labor cost and robot user cost is then used to measure l_{i,o,t_0} for each occupation. The user cost of robots is calculated using the occupation-level robot price data available in IFR and the set of calibrated parameters in section 3.1. Table Appendix B.2 summarizes these statistics for the aggregated 5 occupational groups in the US. The costs for production occupations and transportation

Figure Appendix B.2: Comparison of US Price Indices between JARA and IFR

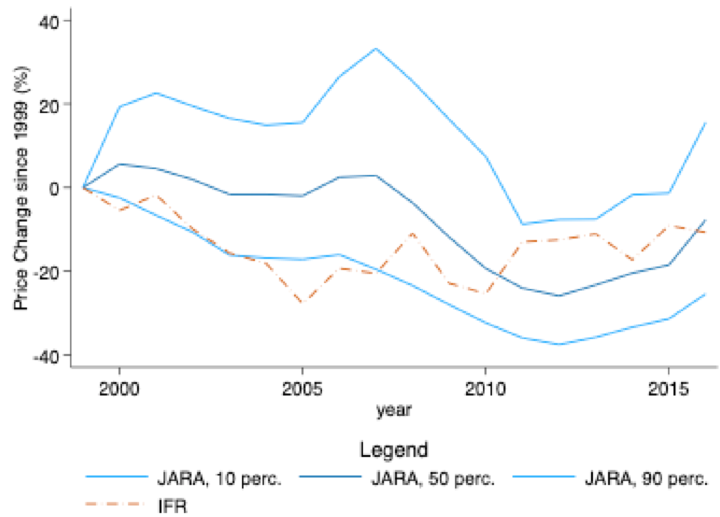

Note: The author's calculation of US robot price measures in JARA and IFR. The JARA measures are disaggregated by 4-digit occupation, and the figure shows the 10th, 50th (median), and 90th percentiles for each year. All measures are normalized to 1999, the year in which the first price measure was available in the IFR data.

occupations account for 18% and 8% of the US economy, respectively, which together represent more than a quarter of the US economy. Furthermore, the share of robot costs in all occupations is still quite low, with the highest share of 0.19% in production occupations, revealing the overall low adoption of robots in the US economy.

To calculate the effect on total income, I also need to calculate the sales share of robots by occupation $y^R_{i,o,t_0} \equiv Y^R_{i,o,t_0} / \sum_o Y^R_{i,o,t_0}$ and the absorption share $x^R_{i,o,t_0} \equiv X^R_{i,o,t_0} / \sum_o X^R_{i,o,t_0}$. To obtain y^R_{i,o,t_0}, I compute the share of robots by occupation produced in Japan $y^R_{2,o,t_0} = Y^R_{2,o,t_0} / \sum_o Y^R_{2,o,t_0}$ and assume the same distribution for other countries due to data limitation: $y^R_{i,o,t_0} = y^R_{2,o,t_0}$ for all i. To get x^R_{i,o,t_0}, I compute the occupational robot adoption in each

Table Appendix B.2: Baseline Shares by 5 Occupation Group

Occupation Group	$\widetilde{x}^{O}_{1,o,t_0}$	l^{O}_{1,o,t_0}	y^{R}_{2,o,t_0}	x^{R}_{1,o,t_0}	x^{R}_{2,o,t_0}	x^{R}_{3,o,t_0}
Routine, Production	17.58%	99.81%	64.59%	67.49%	62.45%	67.06%
Routine, Transportation	7.82%	99.93%	12.23%	11.17%	13.09%	11.04%
Routine, Others	28.78%	99.99%	10.88%	9.52%	11.68%	10.40%
Service	39.50%	99.99%	8.87%	8.58%	9.17%	8.32%
Abstract	6.32%	99.97%	3.43%	3.24%	3.60%	3.18%

Note: The author's calculation of the initial year share variables is shown based on the US Census, IFR, and JARA. As in the main text, country 1 denotes the United States, country 2 denotes Japan, and country 3 denotes the rest of the world. See the main text for the construction of each variable.

country by $X^{R}_{i,o,t_0} = P^{R}_{i,t_0} Q^{R}_{i,o,t_0}$, where Q^{R}_{i,o,t_0} is the occupation-level robot set obtained by the O*NET concordance generated in section 3.3 to the IFR application classification. As mentioned above, the robot price index P^{R}_{i,t_0} is available for a selected set of countries. To compute the rest of the world price index P^{R}_{3,t_0}, I use the average of all available countries, weighted by the occupational robot values each year. The summary table for these variables y^{R}_{i,o,t_0} and x^{R}_{i,o,t_0} at 5 occupation groups is shown in table Appendix B.2. All values in table Appendix B.2 are obtained by aggregating occupations at the 4-digit level.

I take the intermediate input share $\alpha_{i,M}$, from the WIOT. I combine the trade matrix generated above and the WIOT to construct the good and robot expenditure shares s^{G}_{i,t_0}, s^{V}_{i,t_0} and s^{R}_{i,t_0}. Specifically, with the robot trade matrix, I take the total sales value by summing across importers for each exporter and the total absorption value by summing across exporters for each importer. I also obtain the total absorption of goods by the WIOT. I

Table Appendix B.3: 1990 Occupation Group Switching Probability

		Routine			Service	Abstract
		Production	Transportation	Others		
Routine	Production	0.961	0.011	0.010	0.006	0.012
	Transportation	0.020	0.926	0.020	0.008	0.025
	Others	0.005	0.006	0.955	0.020	0.014
Service		0.003	0.002	0.020	0.967	0.007
Abstract		0.014	0.014	0.036	0.015	0.922

Note: The table shows the transition rates between occupational groups calculated from the 1990 CPS-ASEC data. The probability is the probability of moving to the column occupation group conditional on being in each row occupation.

compute the expenditure shares from these values.

I take the 1990 flow Markov transition matrix from the cleaned CPS-ASEC data. Table Appendix B.3 shows this conditional switching probability for the first year. The matrix for the other years is available upon request. Because occupation employment data are difficult to obtain worldwide, I assign the same transition probabilities for other countries in my estimation strategy.

Appendix B.3. Trends of Robot Stocks and Prices

Figure Appendix B.3 shows the US robot trends at the occupation level. The left panel shows the raw stock trend, which shows that overall robot stocks increased rapidly over the period, as found in the previous literature, and that the increase occurred at different rates across occupations. To highlight such a difference, I plot the normalized trend at 100 in the first year in the right panel. There is significant heterogeneity in the growth

Figure Appendix B.3: Trends of Japanese Robot Use at the US Occupation Level

(a) Stock (b) Prices

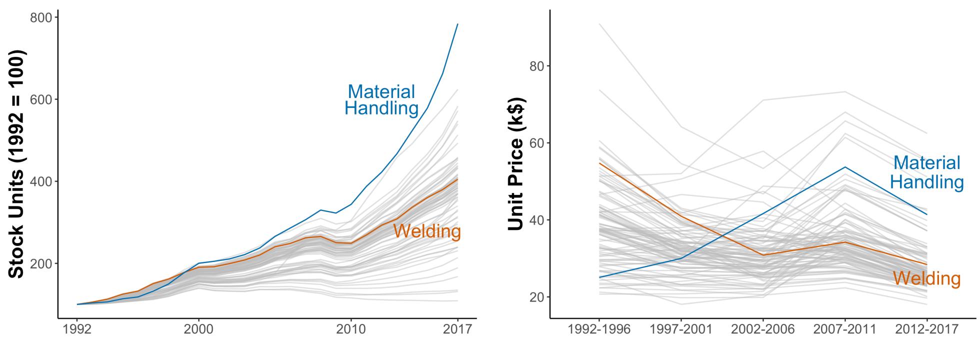

Note: The left panel shows the trend in U.S. robot stocks for each occupation, normalized to 100 in 1992. The right panel shows the trend in robot prices in the US for each occupation. Two occupations are highlighted in both panels: "Welding" corresponds to the occupation code in IPUMS USA, OCC2010 = 8140 "Welding, soldering, and brazing workers." Material handling" corresponds to the occupation code OCC2010 = 9620 "Laborers and Freight, Stock, and Material Movers, Hand". Years are aggregated into five-year bins (with the last bin, 2012-2017, being a six-year bin) to smooth out yearly noise.

rates, ranging from a factor of one to eight. Next, figure Appendix B.3b shows the trend of robot prices in the US for each occupation. In addition to the overall downward trend, there is considerable heterogeneity in the pattern of price declines across occupations. The price patterns are strongly correlated across countries, with a correlation coefficient of 0.968 between US and non-US prices at the occupation-year level. Motivated by this finding, I use the prices of non-US countries as the JRS to the US in the reduced-form analysis.

To further emphasize the heterogeneity of the trends, the following two occupations are colored: "Welders, Solderers, and Brazers" (or "Welding") and "Laborers and Freight, Stock, and Material Movers, Manual" (or "Material Handling") in these two figures. A spot welding robot is used in routine production occupations, while a material handling robot is used in transportation (material moving) occupations. On the one hand, the stock of welding robots grew throughout the period in the left panel, and their average price fell during the 1990s. On the other hand, the stock of material handling robots grew rapidly and its price increased over the sample period. These results indicate the difference in automation shocks; robots such as welding robots followed a standard pattern of expansion along the demand curve, while other robots such as material handling robots expanded their adoption even as the average price increased, indicating the impact of the automation shock described in the model section.

Figure Appendix B.3b suggests an anomaly in the increasing trend from 2007 to 2011. This pattern emerges because, during the Great Recession, total units declined more than total sales. After the Great Recession, the growth in the value and number of robots accelerated. These observations suggest a structural break in the robot industry during the Great Recession, which is beyond the scope of this paper.

Appendix B.4. Trade of Industrial Robots

I combine BACI and IFR data to calculate the trade shares of industrial robots. In particular, I use the HS code 847950 ("Industrial robots for multiple uses") to measure robots, following (Humlum, 2021; Acemoglu and Restrepo, 2022), using 1998 as the starting year value, the first year in which the HS code 847950 was available. To calculate the total absorption value of robots in each country, I use the robot units (quantities) of the IFR data combined with the robot price indices published in the IFR annual reports for selected countries (Graetz and Michaels, 2018). Note that these price indices do not provide a disaggregation by robot tasks or occupations, highlighting the value-added of the JARA data. Figure Appendix B.4 illustrates the international trade pattern of robots. In the left panel, I compute the import-absorption ratio. To remove the noise from annual observations and focus on long-term trends, I aggregate the data into five-year bins: 2001-2005 and 2011-2015. The result shows that many countries are importing robots rather than producing them. Japan's low import ratio stands out, revealing its comparative advantage in this area. Notably, China has gradually domesticated the production of robots over the study period. Another way to capture the comparative advantage of the robot industry is to examine the export share, as shown in the right panel of Figure Appendix B.4. In 2001-2005, the EU dominated half of the world robot market and Japan a third. The remaining 20% is shared by the rest of the world, mainly the US and South Korea.

Figure Appendix B.5 shows the trend of robot export and import shares for the US, Japan, and the rest of the world (RoW). The trends are fairly stable for the three regions, except that the US import share has decreased relative to RoW.

Robots from Japan in the US, Europe, and the Rest of the World. To compare the pattern of robot adoption internationally, I generate growth rates of the robot stock between 1992 and 2017 at the occupation level for each group of destination countries. The groups are the US, the non-US countries (all countries except the US and Japan), and the five European countries (or "EU-5") of Denmark, Finland, France, Italy, and Sweden used in AR.

Figure Appendix B.4: Trade of Industrial Robots

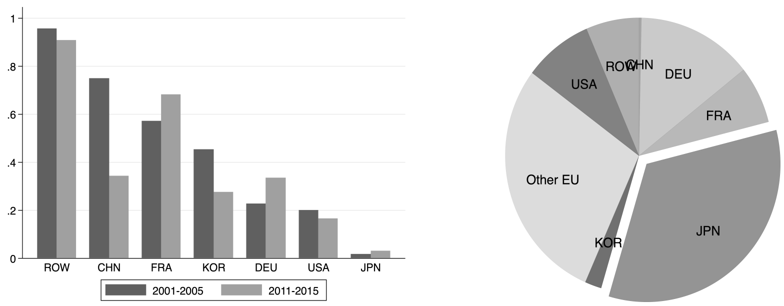

(a) Robot Import-Absorption Ratio (b) World Robot Export Share, 2001-2005

Note: The author's calculation from IFR and BACI data. The left panel shows the share of imports in total absorption value. The import value is calculated by aggregating trade values across countries of origin in the BACI data (HS-1996 code 847950), and the absorption value is calculated using the price index and the quantity variable available for selected countries in the IFR data. The data are aggregated by 5 years in 2001-2005 and 2011-2015, and countries are sorted in descending order by import shares in 2001-2005. The right panel shows the export share for the 2001-2005 aggregates obtained from the BACI data.

Figure Appendix B.5: Robot Trade Share Trends

(a) Exports (b) Imports

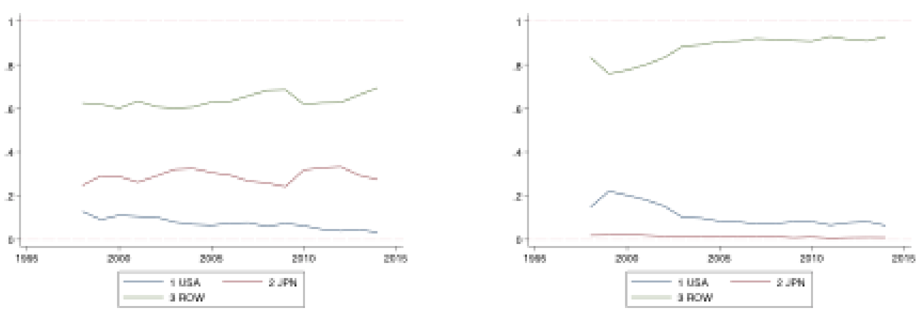

Note: The author's calculation of world trade shares is shown based on the BACI data. Industrial robots are measured by HS code 847950 (Industrial robots for multiple uses).

According to Graetz and Michaels (2018), the perpetual inventory method with a depreciation rate of $\delta = 0.1$ is used to calculate the stock of robots.

Figure Appendix B.6 shows scatter plots of the growth rates at the occupation level. The left panel shows the growth rates in the US on the horizontal axis, while the vertical axis shows the non-US countries. The right panel shows the same measures on the horizontal axis, but the growth rates in the set of EU-5 countries on the vertical axis. These panels show that robot stocks at the occupation level (1992-2017) have grown proportionately between the US and non-US, compared to the US and EU-5. This finding contrasts with AR, which found that aggregate robot stocks in the US grew at roughly the same rate as those in the EU-5. It also suggests that the non-US growth patterns reflect the growth of robot technology at the occupational level available in the US. These patterns are used as a proxy for robot technology available in the US. In the model section, I take a further step and solve for the quantity and value of robot adoption in non-US countries in a general equilibrium including the US and non-US countries.

One possible reason for the difference between my results and those of AR is the difference in data sources. In contrast to the JARA data I use, AR use IFR data, which includes all robot-producing countries. Because the EU-5 is closer to major robot-producing countries other than Japan, including Germany, the pattern of robot adoption across occupations may be influenced by their presence. If these nearby producers have a comparative advantage in producing robots for a particular occupation, then EU-5 may adopt robots for such occupations intensively from nearby producers. In contrast, non-EU-5 countries, including the United States, may not benefit from the proximity of

Figure Appendix B.6: Growth Rates of Robots at the Occupation Level

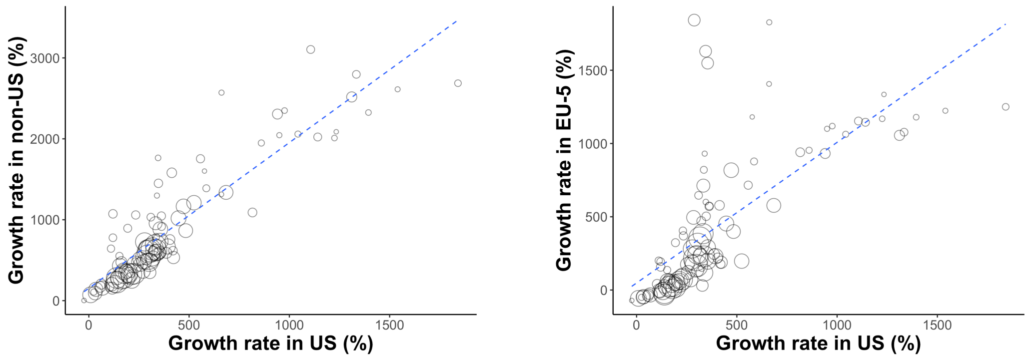

(a) Comparison between the US and non-US

(b) Comparison between the US and EU-5

Note: The growth rates of robot stocks based on JARA and O*NET are shown. The left panel shows the correlation between the occupation-level growth rates of robot stocks from Japan to the US and the growth rates of stocks in non-US countries. The right panel shows the correlation between the growth rates of quantities in the U.S. and EU-5 countries. Non-US is the aggregate of all countries except the US and Japan. EU-5 is the aggregate of Denmark, France, Finland, Italy, and Sweden used in Acemoglu and Restrepo (2020). Each bubble represents an occupation. The bubble size reflects the stock of robots in the US in the base year 1992.

these producers; thus, they are more likely to purchase robots from producers located far from the EU-5, such as Japan.

Appendix B.5. Details in Application-Occupation Matching

Details of the application-occupation matching are discussed. First, I accessed the O*NET Code Connector (https://www.onetcodeconnector.org/) and web-scraped search results as follows. For each robot application title listed in the Appendix B.1 section, I search for matches on the web page and record all occupation codes, names, and match scores. Then I append the result files across all applications, which is called the match score file. At this stage, I drop the assembly and measurement/inspection/test robots from the data due to poor match quality. Second, I merge the match score file and the JARA data at the application level and take the weighted average of the robot sales values and quantities with the weight of the score, as in (10).

For example, consider spot welding and material handling robots. First, spot welding joins two or more sheets of metal into one by applying heat and pressure to a small area called a spot. O*NET-SOC code 51-4121.06 is titled ”Welders, Cutters, and Welder Fitters” (”Welders” below). This suggests that spot welding robots and welders perform the same welding task. Second, material handling involves moving heavy materials a short distance, another primary robot application. ONET-SOC code 53-7062.00 is titled ”Laborers and Freight, Stock, and Material Movers, Hand” (”material handlers” below). Again, both material handling robots and material handlers perform the material handling task. Figure Appendix B.7 shows the top five match scores for spot welding and material handling, with these two occupations at the top of the match score rankings.

Hard-cut Matching of Applications and Occupations. Although matching between applications and occupations based on the equation (10) is transparent and performed automatically rather than using the researcher’s judgment, there may be a concern that such a matching method may potentially contain errors due to noise in the textual descriptions of the Occupation Dictionary. For example, Figure Appendix B.7 shows a case where spot-welding robots are matched with ”Laundry and dry-cleaning workers” with a high score. This is primarily because the textual description for these workers includes ”Apply bleaching powders to spots and spray them with steam to remove

Figure Appendix B.7: Examples of Match Scores

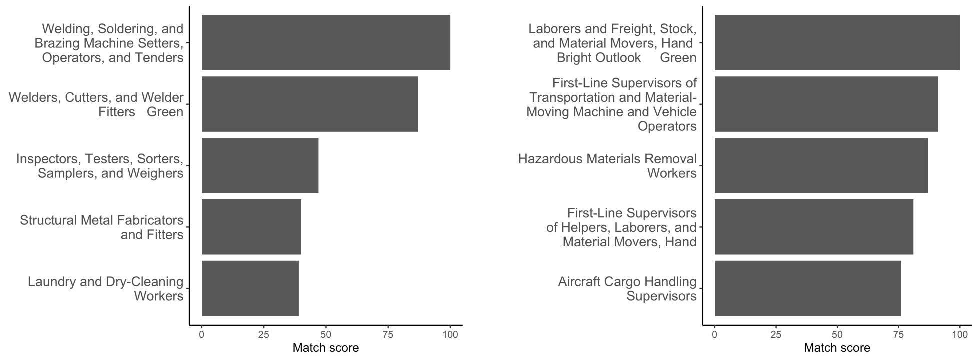

(a) Spot Welding

(b) Material Handling

Note: The author's calculation from the O*NET Code Connector search result. The bars indicate the match scores for the query term "Spot Welding" in panel (a) and "Material Handling" in panel (b). Occupation codes are 2010 O*NET-SOC codes. In each panel, the occupations are sorted descending by their relative relevance scores, and the top five occupations are shown.

stains from fabrics...", which has a high matching score with the term "spot".

To mitigate this concern, I explore manual hard-cut matching between applications and occupations by dropping potentially problematic application-occupation matches with a matching score of 75 or less while including sufficient data variation. I then construct the matching score according to (10) conditional on the remaining pairs and compute the robot quantity and price variables. Figure Appendix B.8 shows the result of the regression specification (C.1) using these measures. The estimated coefficients are somewhat larger than with the preferred data matching procedure, mainly because the hard-cut matching removes erroneous matches that may contain noise. Statistical significance is maintained in all columns.

Figure Appendix B.8: Wage and Robot Prices with a Hard-cut Matching Method

Note: The figure shows the relationship between the Japan Robot Shock, based on application-level robot measures matched to occupations using the hard-cut method described in the main text (horizontal axis), and changes in log wages (vertical axis). The sample includes all occupations that existed between 1970 and 2007. Bubble sizes reflect employment in the base year, and the number of observations is 324. All variables are residualized by control variables (female share of occupation, college share, age distribution, foreign-born share, and the China shock in equation D.2).

Appendix B.6. Other Potential Methods for Adjusting the Robot Prices

In this paper, I use the general equilibrium model to control for the quality component of robot prices. However, other methods have been proposed in the literature to measure the price of capital goods. In this subsection, I briefly describe these methods and their limitations.

The first approach is to control for quality change using the hedonic approach used by Tambe et al. (2019), among others, in their application to digital capital. The hedonic approach requires information about the detailed specifications of each robot. Unfortunately, it is difficult to keep track of the detailed specifications of commonly used robots as the robotics industry changes rapidly.

The second method is more data-driven. The Bank of Japan (BoJ) provides the quality-controlled price index. Unfortunately, the method is not clearly explained in the BoJ's technical documentation. It is said to be a "cost evaluation method" in which the BoJ asks producer firms to measure the quality improvement component of price changes between periods. Obtaining the quality measures is challenging as I do not know the surveyed firms or the quality components.

Appendix C. Reduced-form Analysis

Combining all of the data in **??**, I present in this section several facts and data patterns about robots, JRS, and their relationship to labor market outcomes in the United States.

Appendix C.1. The Effects of the Japan Robot Shock on US Occupations

Because labor demand may be affected by trade liberalization, especially the China shock in my sample period, I control for the occupational China shock using the method developed by Autor et al. (2013). Specifically, I compute the occupational China shock by (D.2). For the list of non-US countries, I follow Autor et al. (2013) and take eight countries: Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland. Appendix B.2 shows the distribution of occupational employment l_{s,o,t_0} for each sector. Intuitively, an occupation receives a large trade shock if sectors facing increased import competition from China employ the occupation intensively. With this measure of the trade shock in the control variable, I run the following regression:

$$\Delta \ln\left(\ln w_o\right) = \alpha_0 + \alpha_1 \times \left(-\psi_o^J\right) + \alpha_2 \times IPW_{o,t_1} + \boldsymbol{X}_o \cdot \boldsymbol{\alpha} + \varepsilon_o, \qquad \text{(C.1)}$$

where w_o is the log hourly wage and $\boldsymbol{X}_o$ is the vector of baseline demographic control variables. The controls are female share, college graduate share, age distribution, and foreign-born share.

First, I examine the correlation between various robot measures and wage measures. In figure Appendix C.1a, the left panel shows the correlation between the JRS and baseline US wages in 1990 at the occupation level. No systematic relationships are found between these variables, suggesting that the JRS did not necessarily trigger an increase in wage inequality in the 1990s and 2000s. Next, the middle panel shows the result of the estimated (C.1) in a scatterplot, which shows that a 10% reduction in Japanese robot prices reduces US occupational wages by 1.2%. Thus, the JRS adversely affected US occupations, suggesting a substitution of robots for labor. Finally, total spending on robots quantitatively affects the demand for labor in each occupation, conditional on robot prices. The right panel shows the relationship between the change in robot spending and wages, suggesting that the negative impact on wages also operates through the margin of spending, indicating the substitutability of labor due to robot penetration at the occupation level.

Next, table Appendix C.1 shows the result of the regression (C.1) to vary across the occupational groups defined above. I find the negative effects in routine production and routine transportation occupations, demonstrating the heterogeneity of the effects across occupation groups. This finding motivates me to consider the group-specific EoS between robots and workers in the model section.

Again, the novelty of these results lies in the use of robot cost reductions

Table Appendix C.1: The heterogeneous effects of the Japan robot shock on US occupations

	(1)
VARIABLES	$\Delta \ln(wage)$
$(-\psi^J) \times$ Routine, production	-0.627***
	(0.112)
$(-\psi^J) \times$ Routine, transportation	-0.738***
	(0.0624)
$(-\psi^J) \times$ Routine, others	0.00770
	(0.0536)
$(-\psi^J) \times$ Service	-0.0639
	(0.107)
$(-\psi^J) \times$ Abstract	0.00693
	(0.0789)
Observations	324
R-squared	0.462

Note: The table shows the coefficients of the regression (C.1), which allows the coefficient α_1 to vary across occupation groups. Observations are occupations at the 4-digit level, and the sample includes all occupation codes that existed consistently between 1970 and 2007. ψ^J stands for the Japan robot shock from equation (11). Control variables are included for the female share, the college graduate share, the age distribution (shares of ages 16-34, 35-49, and 50-64 among workers aged 16-64), the foreign-born share since 1990, and the China shock from equation (D.2). Standard errors are clustered at the 2-digit occupation level. *** p<0.01, ** p<0.05, * p<0.1.

at the occupation level. Therefore, I will present additional results that complement the results. Table Appendix C.2 shows the results of the regression (C.1) using several alternative outcome periods and robot measures on the right-hand side. Panel A takes the wage change between 1990 and 2007, the

Figure Appendix C.1: The Japan Robot Shock and US Occupational Wages

(a) Japan Robot Shock and Baseline Wage

(b) Changes in Wage and Robot Prices

(c) Changes in Wage and Robot Expenditure

Note: The left panel shows the scatterplot, weighted line of best fit, and 95% confidence interval of the baseline (1990) US log wage (horizontal axis) and the Japan Robot Shock (JRS) in equation (11) (vertical axis) at the 4-digit occupation level. The middle panel shows the relationship between the JRS (horizontal axis) and changes in log wages (vertical axis). The right panel shows the relationship between log total expenditure on Japanese robots in non-US countries (horizontal axis) and changes in log wage (vertical axis). In all panels, the sample includes all occupations between 1970 and 2007, bubble sizes reflect employment in the base year, and the number of observations is 324. In the middle and right panels, the variables are residualized with control variables (female occupational share, college share, age distribution, foreign-born share, and the China shock in equation D.2) and occupational group fixed effects.

Table Appendix C.2: Regression of Wages on Robot Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	dln_wage	dln_wage	dln_wage	dln_wage	dln_wage	dln_wage	dln_wage	dln_wage
					A. 1990-2007			
Robot Measure	-0.169***	-0.196***	-0.180***	-0.171***	-0.0399	-0.0798**	-0.210***	-0.206***
	(0.0395)	(0.0398)	(0.0460)	(0.0463)	(0.0399)	(0.0346)	(0.0601)	(0.0458)
R-squared	0.066	0.283	0.055	0.245	0.005	0.214	0.093	0.284
					B. 1970-1990			
Robot Measure	0.00691	0.00772	-0.00388	0.00142	0.00699	-0.00480	0.00866	0.0189
	(0.0262)	(0.0233)	(0.0306)	(0.0269)	(0.0236)	(0.0244)	(0.0286)	(0.0240)
R-squared	0.000	0.079	0.000	0.079	0.000	0.079	0.000	0.081
Robot Measure	US Stock	US Stock	- US Price	- US Price	Non-US Stock	Non-US Stock	- Non-US Price	- Non-US Price
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	324	324	324	324	324	324	324	324

Note: Regressions of log wage changes on robot measures are shown. Panel A takes the wage change between 1990 and 2007, the main period, while Panel B takes the change between 1970 and 1990, the pre-sample period. The regressors are robot stock in the US (columns 1 and 2), robot stock in non-US countries (columns 3 and 4), robot price in the US (columns 5 and 6), or robot price in non-US countries (columns 7 and 8). The control variables are demographic variables (the share of women, the share of college graduates, the share of 16-34, 35-49, and 50-64 years old workers, and the share of foreign-born workers since 1990) and the China trade shock defined in equation (D.2). Bootstrapped standard errors are shown in parentheses.

main period, while Panel B takes the change between 1970 and 1990, the pre-sample period. In each panel, the columns differ along two dimensions: (i) the robot measure, from the robot stock in the U.S. and other countries (non-US) and the robot price in the U.S. and other countries, and (ii) whether the regressions include control variables for demographic variables and the China trade shock.

Table Appendix C.3 shows the regression result of C.1 with employment as the outcome variable. A qualitatively similar pattern is found in the sense that employment in a subset of the routine occupation group (production workers) is reduced in occupations that experienced the JRS; in contrast, there is no statistically significant point estimate for transportation workers.

Furthermore, to address the concern that the U.S. is a large country that affects robot prices more directly, I confirm that the effect of robot price reductions on labor demand is also observed in a small open economy in Appendix C.2.

Although these data patterns and regressions are informative about the substitutability of robots, they do not provide definitive answers about the value of the substitution parameter or the distributional and aggregate effects of robotization. First, the observed JRS may reflect the quality improvement of robots, implying that the quality-adjusted cost reduction of robots may be even larger. Second, changes in labor demand for one occupation after the shock may affect wages and employment in other occupations by changing their marginal products. Third, the coefficients in equation (C.1) reveal the relative impact of the JRS but not the real wage impact. I develop and

Table Appendix C.3: The heterogeneous effects of the Japan Robot Shock on US occupations

VARIABLES	(1) $\Delta \ln(emp)$
$(-\psi^J) \times$ Routine, others	-0.657***
	(0.229)
$(-\psi^J) \times$ Routine, transportation	-0.258
	(0.180)
$(-\psi^J) \times$ Routine, production	-0.0651
	(0.143)
$(-\psi^J) \times$ Service	-0.126
	(0.227)
$(-\psi^J) \times$ Abstract	-0.342
	(0.256)
Observations	324
R-squared	0.126

Note: The table shows the coefficients in a regression (C.1) that allows the coefficient α_1 to vary across occupation groups, with the outcome variable being the long difference of log employment from 1990 to 2007. Observations are occupations at the 4-digit level, and the sample includes all occupation codes that existed consistently between 1970 and 2007. ψ^J stands for the Japan robot shock from equation (11). Control variables are included for the female share, the college graduate share, the age distribution (shares of ages 16-34, 35-49, and 50-64 among workers aged 16-64), the foreign-born share since 1990, and the China shock from equation (D.2). Standard errors are clustered at the 2-digit occupation level. *** p<0.01, ** p<0.05, * p<0.1.

estimate a general equilibrium model to address these issues in the main text.

Appendix C.2. Validation Exercise in a Small Country

One concern with my reduced-form analysis is that US demand, as a large buyer of robots, may influence the price. To mitigate this, I perform a robustness exercise using data from a small country that is unlikely to influence the world price of robots. I use data from the Netherlands because it is the largest export destination for Japanese robots in Europe, after Germany, the UK, Italy, and France, and it is a small open economy. The data come from the international IPUMS and provide the ISCO 1-digit level occupation indicator for 2001 and 2011. The prices of occupational robots are aggregated at the same level and the relationship between the JRS and occupational employment growth is examined. Since the wage variable is not available in the international IPUMS, I use the employment variable as a proxy for changes in labor demand. Figure Appendix C.2 summarizes the results. Despite a significant difference in context and level of data aggregation, I find a significant negative relationship between these two variables. This exercise suggests that a reduction in the price of Japanese robots, which is likely to exogenously affect small open economies, reduces labor demand in the Netherlands.

Appendix C.3. The Effect of Robots from Japan and Other Countries

A potential concern in my empirical setting is the selection problem regarding the robot's country of origin. In particular, because robots from Japan may differ from those from other countries, the labor market effects may also differ. Unfortunately, data limitations make it difficult to directly compare the effects of these two different groups of robots, so I focus on the best comparable measure of robotization between robots from Japan and

Figure Appendix C.2: The Effect of Japan Robot Shock in the Netherlands

Note: The bubble plot and the fitted line between the Netherlands' occupational growth and the Japan robot shock are shown. The period is 2001 to 2011. The size of the bubble reflects the size of employment in the initial period. Occupations are aggregated to the 1-digit ISCO level. The shading indicates the 95% confidence interval.

those from all countries, which is the quantity of robot stock. I take the total stock of robots in the U.S. from the IFR data. The IFR data gives only the total number and does not specify the country of origin. I then convert the IFR application codes to JARA application codes to match the JARA application codes to the occupation codes using the allocation rule. As a result, I have the robots used in the U.S. that are sourced from any country at the occupation level. I then run the following regression using the obtained robot measures and my preferred measure from JARA:

$$\Delta Y_o = \beta^Q \Delta K_o^{R,Q} + X_o \gamma^Q + \varepsilon_o^Q, \tag{C.2}$$

where ΔY_o is the change in wages at the occupation-o level, ΔK_o^Q is the measure of the number of robots taken from either JARA (i.e., robots from

Japan) or IFR (i.e., robots from the world), and ε_o^Q is the error term. The coefficient of interest is β^Q, which provides insight into the correlation between the changes in labor market outcomes and the number of robots depending on whether the robots are from Japan. Specifically, if robots from Japan substitute workers more than robots from other countries, the coefficient β^Q is expected to be larger when I use the JARA robot measure than when I use the IFR.

Table Appendix C.4 shows the regression result of (C.2). The result of the IFR data is consistent with the previous results of Acemoglu and Restrepo (2020). Table Appendix C.4 shows that both the JARA- and IFR-based robot measures capture the substitution of workers by robots, although the coefficient is slightly larger for the JARA robot measures than for IFR.

Appendix D. Estimation Appendix

Following the convention in the literature, I assume $\alpha^R = 2/3$, meaning that the robot integration cost is two-thirds of the total robot-related expenditure. As in Cooper and Haltiwanger (2006), the adaptation cost parameter is set to $\gamma = 0.295$. Estimates from the literature on dynamic discrete choice of occupations are used, and the elasticity of switching occupations is set to $\phi = 1.4$.

Appendix D.1. Robot Trade Elasticity

To estimate the robot trade elasticity ε^R, I apply and extend the trilateral method of Caliendo and Parro (2015). First, I decompose the robot trade cost $\tau_{li,t}^R$ into $\ln \tau_{li,t}^R = \ln \tau_{li,t}^{R,T} + \ln \tau_{li,t}^{R,D}$, where $\tau_{li,t}^{R,T}$ is the tariff on robots from the UNCTAD-TRAINS database and $\tau_{li,t}^{R,D}$ is the asymmetric non-tariff

Table Appendix C.4: Regression Results of Labor Market Outcome on JARA and IFR Robot Stocks

	(1)	(2)	(3)	(4)
VARIABLES	$\Delta \ln(w)$	$\Delta \ln(w)$	$\Delta \ln(w)$	$\Delta \ln(w)$
$\Delta \ln(K^{R,Q}_{JPN \to USA})$	-0.372		-0.271	
	(0.0466)		(0.0304)	
$\Delta \ln(K^{R,Q}_{USA})$		-0.144		-0.111
		(0.0300)		(0.0185)
Observations	324	324	324	324
R-squared	0.307	0.200	0.349	0.262
Controls			✓	✓

Note: Regression results of occupational wage changes are shown. The observations are occupations at the 4-digit level, and the regression is between 1990 and 2007 with the sample of all occupations that existed between 1970 and 2007. Columns 1 and 3 include robot measures from Japan using JARA data, while columns 2 and 4 include robot measures from the world using IFR data as explained in the main text. Columns 1 and 2 do not include the control variables of demographic variables (female share, age distribution, college graduate share, and foreign-born share) and the China trade shock in (D.2), while columns 3 and 4 do. Heteroskedasticity-robust standard errors are reported in parentheses.

trade cost. The latter term is assumed to be $\ln \tau^{R,D}_{li,t} = \ln \tau^{R,D,S}_{li,t} + \ln \tau^{R,D,O}_{l,t} + \ln \tau^{R,D,D}_{i,t} + \ln \tau^{R,D,E}_{li,t}$, where $\tau^{R,D,S}_{li,t}$ captures symmetric bilateral trade costs such as distance, common border, language, and free trade agreement (FTA) status and satisfies $\tau^{R,D,S}_{li,t} = \tau^{R,D,S}_{il,t}$, $\tau^{R,D,O}_{l,t}$ and $\tau^{R,D,D}_{i,t}$ are the origin and destination FEs such as non-tariff barriers, respectively, and $\tau^{R,D,E}_{li,t}$ is the random error orthogonal to the tariffs. By (A.16), I have

$$\ln\left(\frac{X^R_{li,t}X^R_{ij,t}X^R_{jl,t}}{X^R_{lj,t}X^R_{ji,t}X^R_{il,t}}\right) = \left(1-\varepsilon^R\right)\ln\left(\frac{\tau^{R,T}_{li,t}\tau^{R,T}_{ij,t}\tau^{R,T}_{jl,t}}{\tau^{R,T}_{lj,t}\tau^{R,T}_{ji,t}\tau^{R,T}_{il,t}}\right) + e_{lij,t}, \tag{D.1}$$

where $X^R_{li,t}$ is the bilateral sales of robots from l to i in year t and $e_{lij,t} \equiv \ln\tau^{R,D,E}_{li,t} + \ln\tau^{R,D,E}_{ij,t} + \ln\tau^{R,D,E}_{jl,t} - \ln\tau^{R,D,E}_{lj,t} - \ln\tau^{R,D,E}_{ji,t} - \ln\tau^{R,D,E}_{il,t}$. The advantage of this approach is that it does not require symmetry for the non-tariff trade costs $\tau^{R,D}_{li}$, but only the orthogonality condition for the asymmetric component of the trade costs. My method also extends Caliendo and Parro (2015) by using time series variation as well as trilateral country-level variation to complement the relatively small number of observations in the robot trade data.

When regressing (D.1), I also consider controlling for two separate sets of FEs. The first is the unilateral FE, which indicates whether a country is included in the trilateral country pair, and the second is the bilateral FE for the country pair. These FEs are relevant in my setting because only a few countries export robots, and controlling for the unobserved characteristics of these exporters is crucial.

Table Appendix D.1 shows the result of the regression of (D.1). The first two columns show the result for HS code 847950 ("Industrial multi-purpose robots", the definition of robots used in Acemoglu and Restrepo, 2022) and the last two columns for HS code 8479 ("Machines and mechanical appliances having individual functions, not specified or included elsewhere in this chapter", used by Humlum, 2021). The first and third columns control for unilateral FE, while the second and fourth control for bilateral FE. The implied trade elasticity of robots ε^R is estimated quite tightly, ranging between

Table Appendix D.1: Coefficient of equation (D.1)

	(1)	(2)	(3)	(4)
	HS 847950	HS 847950	HS 8479	HS 8479
Tariff	-0.272	-0.236	-0.146	-0.157
	(0.0718)	(0.0807)	(0.0127)	(0.0131)
Constant	-0.917	-0.893	-1.170	-1.170
	(0.0415)	(0.0381)	(0.00905)	(0.00853)
FEs	h-i-j-t	ht-it-jt	h-i-j-t	ht-it-jt
N	4610	4521	88520	88441
r2	0.494	0.662	0.602	0.658

Note: The author's calculation, based on BACI data from 1996 to 2018 and the equation (D.1), is shown. The first two columns show the result for HS code 847950 ("Industrial robots for multiple uses"), while the last two columns show the result for HS code 8479 ("Machines and mechanical appliances having individual functions, not specified or included elsewhere in this chapter"). The first and third columns control for unilateral fixed effects, while the second and fourth columns control for bilateral fixed effects.

1.13 and 1.34. Given these estimation results, I use $\varepsilon^R = 1.2$ in the estimation and in the counterfactuals.

To put my estimation result in context, note that Caliendo and Parro (2015) showed in Table 1 that the regression coefficient of equation (D.1) is 1.52, with the standard error of 1.81, for "Machinery n.e.c.," which corresponds to HS 84. Thus, my estimate for industrial robots is within one standard deviation of their estimate for a broader category of goods.

Note that the average trade elasticity across sectors is estimated to be significantly higher than these values, such as 4 in Simonovska and Waugh (2014). The low trade elasticity for robots ε^R reflects that robots are highly heterogeneous and hardly substitutable. This low elasticity implies small gains from robot taxes, with the incidence of the robot tax almost on the side of the US (robot buyer) rather than the robot-selling country.

Appendix D.2. Constructing the China Shock Across Occupations

Inspired by Autor et al. (2013), I use the import penetration measure defined at the occupation level:

$$IPW_{o,t} \equiv \sum_{s} l_{s,o,t_0} \Delta m^{C}_{s,t}, \tag{D.2}$$

where l_{s,o,t_0} is the sector-s share of employment in occupation o, and $\Delta m^{C}_{s,t}$ is the per-worker growth of Chinese exports to non-US developed countries. This method is in the spirit of Autor et al. (2013), while I measure the occupational variation in exposure.

Appendix D.3. Detailed Discussion of the Estimator

Using the assumption 1, I develop a consistent and asymptotically efficient two-stage estimator. Specifically, I follow the method developed by Adao et al. (2023), who extended the classical two-stage GMM estimator to the general equilibrium setting and defined the MOIV. The key idea is that the optimal GMM estimator is based on the instrumental variable that depends on unknown structural parameters. The two-step estimator solves this unknown dependence problem and achieves consistency and asymptotic efficiency. Specifically, I define the IVs $Z_{o,n}$ where $n = 0, 1$ as follows:

$$Z_{o,n} \equiv H_{o,n}\left(\boldsymbol{\psi}^{J}\right) = \mathbb{E}\left[\nabla_{\boldsymbol{\Theta}} \nu_o\left(\boldsymbol{\Theta}_n\right) |\boldsymbol{\psi}^{J}\right] \mathbb{E}\left[\nu_o\left(\boldsymbol{\Theta}_n\right)\left(\nu_o\left(\boldsymbol{\Theta}_n\right)\right)^{\top} |\boldsymbol{\psi}^{J}\right]^{-1}, \tag{D.3}$$

where ν_o is the function of the structural residual satisfying

$$\boldsymbol{\nu}_{\boldsymbol{w}} = \boldsymbol{\nu}_{\boldsymbol{w}}(\Theta) = \widehat{\boldsymbol{w}} - \bar{\boldsymbol{E}}_{\boldsymbol{w}_1,\boldsymbol{a}}\widehat{\boldsymbol{a}^{\mathrm{obs}}} - \bar{\boldsymbol{E}}_{\boldsymbol{w}_1,\boldsymbol{A}_2^R}\widehat{\boldsymbol{A}_2^R} - \bar{\boldsymbol{E}}_{\boldsymbol{w}_1,\boldsymbol{b}}\widehat{\boldsymbol{b}},$$

in a matrix notation. The formal statement requires the following additional assumption.

Assumption 2. *(i) A function of* $\widetilde{\boldsymbol{\Theta}}$, $\mathbb{E}\left[H_o\left(\boldsymbol{\psi}_{t_1}^J\right)\nu_o\left(\widetilde{\boldsymbol{\Theta}}\right)\right] \neq 0$ *for any* $\widetilde{\boldsymbol{\Theta}} \neq \boldsymbol{\Theta}$. *(ii)* $\underline{\theta} \leq \theta_o \leq \overline{\theta}$ *for any* o, $\underline{\beta} \leq \beta \leq \overline{\beta}$, $\underline{\gamma} \leq \gamma \leq \overline{\gamma}$, *and* $\underline{\phi} \leq \phi \leq \overline{\phi}$ *for some positive values* $\underline{\theta}, \underline{\beta}, \underline{\gamma}, \underline{\phi}, \overline{\theta}, \overline{\beta}, \overline{\gamma}, \overline{\phi}$. *(iii)* $\mathbb{E}\left[\sup_{\boldsymbol{\Theta}} \parallel H_o\left(\boldsymbol{\psi}_{t_1}^J\right)\nu_o\left(\widetilde{\boldsymbol{\Theta}}\right) \parallel\right] < \infty$. *(iv)* $\mathbb{E}\left[\parallel H_o\left(\boldsymbol{\psi}_{t_1}^J\right)\nu_o\left(\widetilde{\boldsymbol{\Theta}}\right) \parallel^2\right] < \infty$ *(v)* $\mathbb{E}\left[\sup_{\boldsymbol{\Theta}} \parallel H_o\left(\boldsymbol{\psi}_{t_1}^J\right)\nabla_{\widetilde{\boldsymbol{\Theta}}}\nu_o\left(\widetilde{\boldsymbol{\Theta}}\right) \parallel\right] < \infty$.

Under the assumptions 1 and 2, Adao et al. (2023) showed that the estimator $\boldsymbol{\Theta}_2$ obtained in the following procedure is consistent, asymptotically normal, and optimal: Step 1: With a guess $\boldsymbol{\Theta}_0$, estimate $\boldsymbol{\Theta}_1 = \boldsymbol{\Theta}_{\boldsymbol{H}_0}$ using $Z_{o,0}$ defined in (D.3); and Step 2: With $\boldsymbol{\Theta}_1$, estimate $\boldsymbol{\Theta}_2$ by $\boldsymbol{\Theta}_2 = \boldsymbol{\Theta}_{\boldsymbol{H}_1}$ using $Z_{o,1}$ defined in (D.3).

Appendix D.4. Calculating the F-statistics

Write the set of structural parameters $\boldsymbol{\Theta} \equiv (\theta_g, \beta)$. The moment condition is

$$E\left[\hat{a}_o^{err}\left(\boldsymbol{\Theta}\right)|\boldsymbol{\psi}\right] = 0,$$

where $\hat{a}_o^{err}\left(\cdot\right)$ is a non-linear function of the structural parameter $\boldsymbol{\Theta}$. If we impose technical conditions to guarantee that the approximation error vanishes in the asymptote, we can still have the first-order approximation around the true value $\boldsymbol{\Theta} = \boldsymbol{\Theta}_0$, so that

$$\begin{aligned} 0 &= E\left[\hat{a}_o^{err}\left(\boldsymbol{\Theta}\right)|\boldsymbol{\psi}\right] \\ &\approx E\left[\hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right) + \nabla_{\boldsymbol{\Theta}}\hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right)\left(\boldsymbol{\Theta} - \boldsymbol{\Theta}_0\right)|\boldsymbol{\psi}\right], \end{aligned}$$

where $\nabla_{\boldsymbol{\Theta}}$ is the partial derivative operator. This implies that

$$0 \approx C + E\left[\psi_o \nabla_{\boldsymbol{\Theta}}\hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right)\right]\boldsymbol{\Theta} \tag{D.4}$$

where $C \equiv E\left[\psi_o \hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right)\right] - E\left[\psi_o \nabla_{\boldsymbol{\Theta}}\hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right)\boldsymbol{\Theta}_0\right]$ is the constant that is a counterpart of the reduced correlation term in the linear IV-2SLS case. Thus, the counterpart of the first-order F-statistic would be the strength of the correlation between the IV (in my case, the JRS variable) and the score value $\nabla_{\boldsymbol{\Theta}}\hat{a}_o^{err}\left(\boldsymbol{\Theta}_0\right)$, which indicates how much the structural error varies with each element of the structural parameter evaluated at the true parameter value.

To get the intuition, note that the structural residual is the element of the change in the relative demand for robots (to workers) after controlling for the observable component of the automation shock. If it correlates well with the JRS, the robot price shock, then the JRS is "relevant" for estimating the EoS between robots and workers. I implement this idea using the plug-in estimator. That is, I use the sample analog of (D.4) and take the F-statistic of the regression of the score value evaluated at the estimated parameter value on the JRS since I do not know the true parameter value.

Appendix D.5. Model Fit

I apply the simulated data to the linear regression model (C.1). First, I apply the JRS and the implied automation shock, calling this counterfactual wage change the “targeted change.” The predicted wage change is consistent with the moment condition (15), and thus the linear regression coefficient α_1 of (C.1) is expected to match that obtained from the data. Second, I apply only the JRS but not the automation shock, calling this counterfactual wage change an “unintended change.” In this case, the moment condition (15) is violated because the structural residual does not include the unobserved automation shock, introducing a bias into the regression. The difference in estimates from the one using the targeted wage change reveals the magnitude of this bias. Thus, this exercise demonstrates the importance of including the automation shock in the estimation. The details of the method used to simulate the data are given in Appendix E.2.

Table Appendix D.2 shows the results of these exercises. The first column shows the estimates of (C.1) using the data, the second column is the estimate based on the targeted wage change, and the third column is the estimate based on the untargeted wage change. As expected, comparing the first and second columns confirms that the targeted moments are consistent. Furthermore, comparing the third column with these two columns reveals a stronger negative correlation between the simulated wage and the JRS. This is due to the positive correlation between the JRS $-\psi_o^J$ and the implied automation shock $\widehat{a_o^{\text{imp}}}$, which is consistent with the fact that robotic innovations that save costs (hence $\widehat{A_{2,o}^R} > 0$ or $-\widehat{\psi_o^J} > 0$) and quality improvements (thus $\widehat{a_o^{\text{imp}}} > 0$) are likely to occur at the same time. More specifically, with the real data, the regression specification (C.1) contains a positive bias due to this positive correlation. In contrast, the untargeted wage is free of this bias, since its data-generating process includes only the JRS, not the automation

Table Appendix D.2: Model Fit: Linear Regression with Observed and Simulated Data

	(1)	(2)	(3)
VARIABLES	$\widehat{\boldsymbol{w}}_{data}$	$\widehat{\boldsymbol{w}}_{\psi^J \widehat{\boldsymbol{a}^{imp}}}$	$\widehat{\boldsymbol{w}}_{\psi^J}$
$-\psi^J$	-0.118	-0.107	-0.536
	(0.0569)	(0.0711)	(0.175)
Observations	324	324	324

Note: Exercises to examine the model fit using different simulations based on the estimated model are shown. Column (1) is the coefficient on the JRS ψ^J in the reduced-form regression with the China shock control. Column (2) takes the change in US wages predicted by the model with ψ^J and the implied automation shock $\widehat{\boldsymbol{a}^{\text{imp}}}$. Column (3) shows the US wage change predicted by the model with only the JRS (but not the automation shock). Heteroskedasticity-robust standard errors are in parentheses.

shock. Thus, the linear regression coefficient α_1 is higher than that obtained from the real data. In other words, if I had mistakenly assumed that the economy did not experience the automation shock, and if I had believed that the coefficient obtained in Figure Appendix C.1 was free of bias, I would have estimated a higher EoS by ignoring the actual positive correlation between $-\psi_o^J$ and $\widehat{a_o^{\text{imp}}}$. This thought experiment shows that it is critical to account for the automation shock when estimating the EoS between robots and labor using the JRS and that the large EoS in my structural estimates is robust after accounting for this point.

Appendix E. Quantitative Exercise Appendix

Appendix E.1. Parameter Values Used in the Main Quantitative Exercise

Appendix E.2. Details in the Simulation Method

The simulation for the counterfactual analysis consists of three steps. First, I back out the observed shocks from the estimated model for each year

Table Appendix E.3: Parameter Values in the Quantitative Exercise

Parameter	Description	Values	Source
	Calibrated parameters		
ι	Annual discount rate	0.05	-
ε	Goods trade elasticity	4	Simonovska and Waugh (2014)
ε^R	Robot trade elasticity	1.2	Caliendo and Parro (2015)*
ϕ	Workers' occupation switch elasticity	1.4	Traiberman (2019)
γ	Robot capital adjustment cost	0.45	Cooper and Haltiwanger (2006)
α^R	Robot production expenditure share	2/3	Leigh and Kraft (2018)
$\alpha_{i,L}$, $\alpha_{i,M}$, $\alpha_{i,K}$	Goods production expenditure share		WIOD
	Estimated parameters		
Parameter	Description	Values	Source
$\theta_{\text{production}}$	Robot-worker EoS in production	2.71	Table 1, column 2
$\theta_{\text{transportation}}$	Robot-worker EoS in transportation	1.76	Table 1, column 2
$\theta_{\text{other routine}}$	Robot-worker EoS in other routine	1.96	Table 1, column 2
θ_{service}	Robot-worker EoS in service	1.01	Table 1, column 2
θ_{abstract}	Robot-worker EoS in abstract	1.01	Table 1, column 2
β	Occupation demand substitution	0.73	Table 1, column 2

*: The triad strategy of Caliendo and Parro (2015) applied to the robot trade data from BACI.

between 1992 and 2007. Namely, I obtain the efficiency increase of Japanese robots $\widehat{A^R_{2,o,t}}$ using (12). With the point estimates in table 1, the implied automation shock $\widehat{a^{\text{imp}}_{o,t}}$ using (13). To get the efficiency shock of robots in other countries, I assume $\widehat{A^R_{i,o,t}} = \widehat{A^R_{i,t}}$ for $i = 1, 3$. Then, using the robot trade prices $p^R_{ij,t}$ from BACI, I fit a fixed effects regression $\Delta \ln \left(p^R_{ij,t}\right) = \widetilde{\psi}^D_{j,t} + \widetilde{\psi}^C_{i,t} + \widetilde{e}_{ij,t}$, and use $\widehat{A^R_{i,t}} = -\widetilde{\psi}^C_{i,t_1}$. The idea to back out the negative efficiency shock $\widetilde{\psi}^C_{i,t_1}$ is similar to the fixed effects regression in section 3.2, but without the occupational variation that is not observed in the BACI data. Second, applying the backed-out shocks $\widehat{A^R_{i,o,t}}$ and $\widehat{a^{\text{imp}}_{o,t}}$ to the first-order solution of the GE in (A.34) yields the prediction of the changes in the endogenous variables to these first-order shocks. Finally, the predicted level of endogenous variables is obtained by applying the predicted changes to the initial data in $t_0 = 1992$.

Appendix E.3. The Effect of Robotization and the Sources of Shocks

Figure Appendix E.1a shows the effect of two robotization shocks in a sum: the automation shock $\widehat{\boldsymbol{a}}$ and the JRS $\widehat{\boldsymbol{A}}_2$. Although both are relevant shocks to robotics technology during the sample period, the result on the wage distribution combines these two effects, making it difficult to assess the contribution of each shock. To address this concern, figure Appendix E.1 shows the decomposition of the main exercise. The right panel shows the same result as figure 3b. In contrast, the middle panel shows the predicted wage changes with only the automation shock, and the left panel shows the effect of both the automation shock and JRS. In particular, the automation shock reduces the demand for labor and thus the wage for many occupations. In contrast, the JRS lowers the price of robots and raises the marginal product of labor, which raises occupational wages on average.

Appendix E.4. Sensitivity Analysis about the Elasticity of Substitution between Occupations

Figure Appendix E.2 shows the sensitivity analysis of the main counterfactual analysis across different β estimates. The high β value uses $\beta = 1.78$ from Burstein et al. (2019), and the low β value uses $\beta = 0.49$ from Humlum (2021).

Appendix E.5. Underlying Mechanisms for the Impact of the JRS

To understand the underlying mechanisms for the positive effect of the JRS, I conduct additional analyses with three scenarios. The first scenario is the baseline, using the heterogeneous elasticity of substitution and the observed trade shares of robots. The second scenario uses the constant EoS across occupations. The third scenario sets robot trade across countries to zero.

Figure Appendix E.3 shows the effect of JRS in each of the five occupational groups. In the "baseline estimates" panel, I used the estimated heterogeneous elasticities of substitution (EoS) and the observed robot trade shares. The total effect shows the effect of all JRS and is thus consistent with figure 3b of the revised manuscript (note the difference in scale). There are five different bars for each decile, indicating the sub-effects of JRS in each occupational group, holding other occupational JRS constant. For example, the "production" sub-effect shows the effect of JRS in production occupations across wage deciles. These five sub-effects add up to the total first-order effect. The result shows that there is clear heterogeneity across

Figure Appendix E.1: The Effect on Occupational Wages by Sources of Shocks

(a) Both Shocks

(b) Automation Shock

(c) Japan Robot Shock

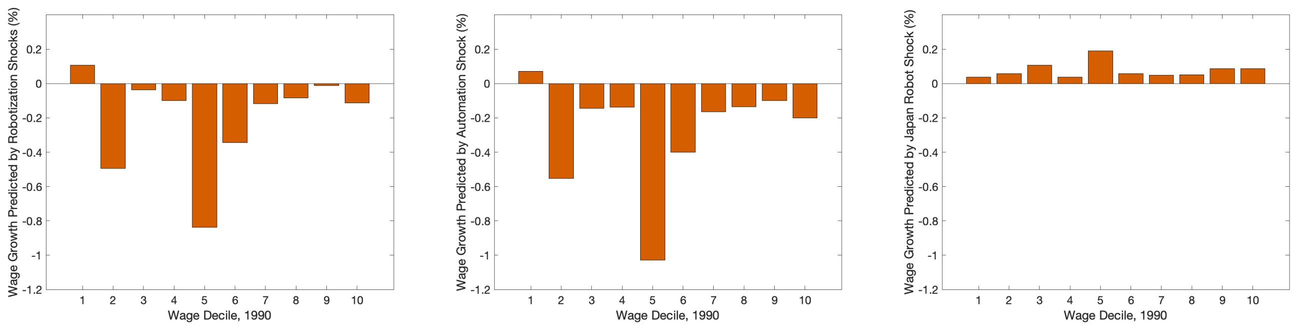

Note: The left panel shows the annualized occupational wage growth rates for each wage decile predicted by the steady-state first-order solution of the estimated model, given by equation (A.32), for each of the ten deciles of the occupational wage distribution in 1990. The middle and right panels distinguish the effect of the automation shock (middle) and the Japan robot shock (right). The right panel corresponds to figure 3b.

Figure Appendix E.2: The Effect on Occupational Wages by Sources of Shocks with Different β Values

(a) High β, Both Shocks

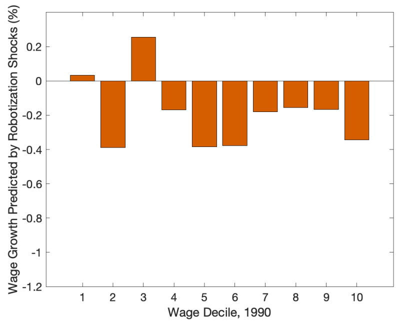

(b) High β, Automation Shock

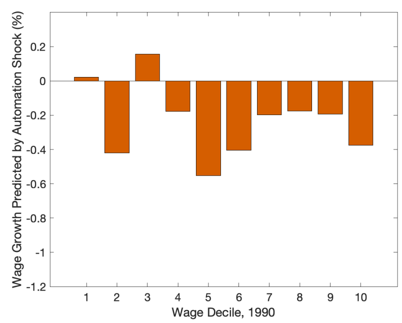

(c) High β, Japan Robot Shock

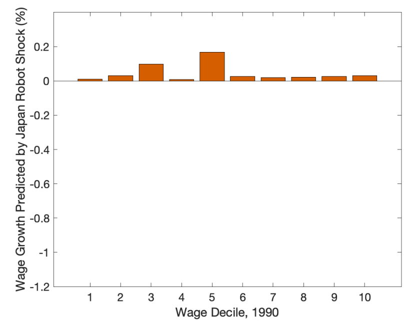

(d) Low β, Both Shocks

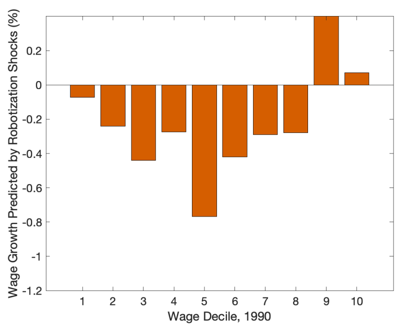

(e) Low β, Automation Shock

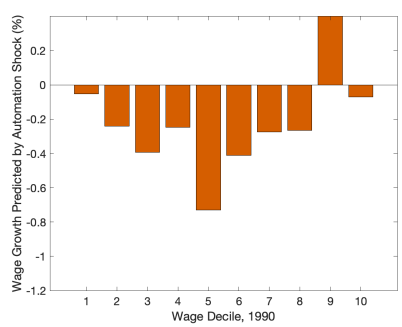

(f) Low β, Japan Robot Shock

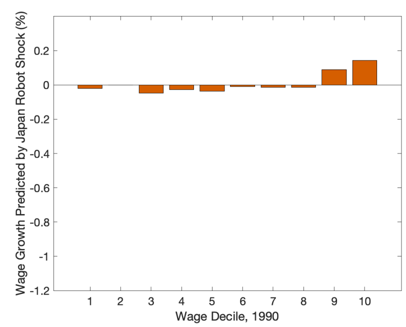

Note: The result of the sensitivity analysis to different values of β is shown. All panels show the annualized occupational wage growth rates for each wage decile. Compared to the baseline elasticity of substitution between occupations, $\beta = 0.73$, the top panels analyze the case with high $\beta = 1.78$, the central estimate of Burstein et al. (2019), while the bottom panels analyze the case with low $\beta = 0.49$, the central estimate of Humlum (2021). The left panels show the wage changes due to the robotization shock (the sum of the automation shock and the Japan Robot Shock, JRS). The middle and right panels distinguish the effect of the automation shock (middle) and the Japan robot shock (right).

the sub-effects. Overall, the Production sub-effect contributes significantly to the total effect, while for some deciles (3 and 5) the Abstract sub-effect is important. Although small, the transportation sub-effect has a negative impact on wages in the lowest wage deciles. Therefore, the model generates negative wage changes due to factor substitution.

The "constant elasticity" panel analyzes the counterfactual with constant EoS across occupations and shows more cases with negative robotization effects. There are some deciles with a negative total effect, and the positive effects are overall weaker than in the first panel. Where does this effect come from? Looking at each sub-effect, the "other routine" JRS consistently shows negative effects across occupations, so this is an important contributor to the smaller wage effect. This is probably due to the constant substitution elasticity. In column 1 of table 1, I estimated the EoS between robots and workers to be 2.05 across all occupations. Using this value, we have an unrealistically strong substitution effect on labor demand for non-production and non-transportation occupations. This finding highlights the importance of accounting for heterogeneity when estimating the EoS.

Finally, the "robot autarky" panel shows the case where countries do not trade robots. In this case, the JRS is not felt as a direct substitution effect on U.S. labor demand, but only as an indirect effect through the input-output linkages between Japanese and U.S. industries. Compared to the baseline estimation panel, the positive wage effect is even stronger because of the I/O link and productive inputs (and stronger demand) in Japan. This result suggests that robot trade has a negative direct substitution effect that is more than offset by the positive I/O linkage effect in an open economy.

Figure Appendix E.3: Understanding the Wage Impact of the JRS

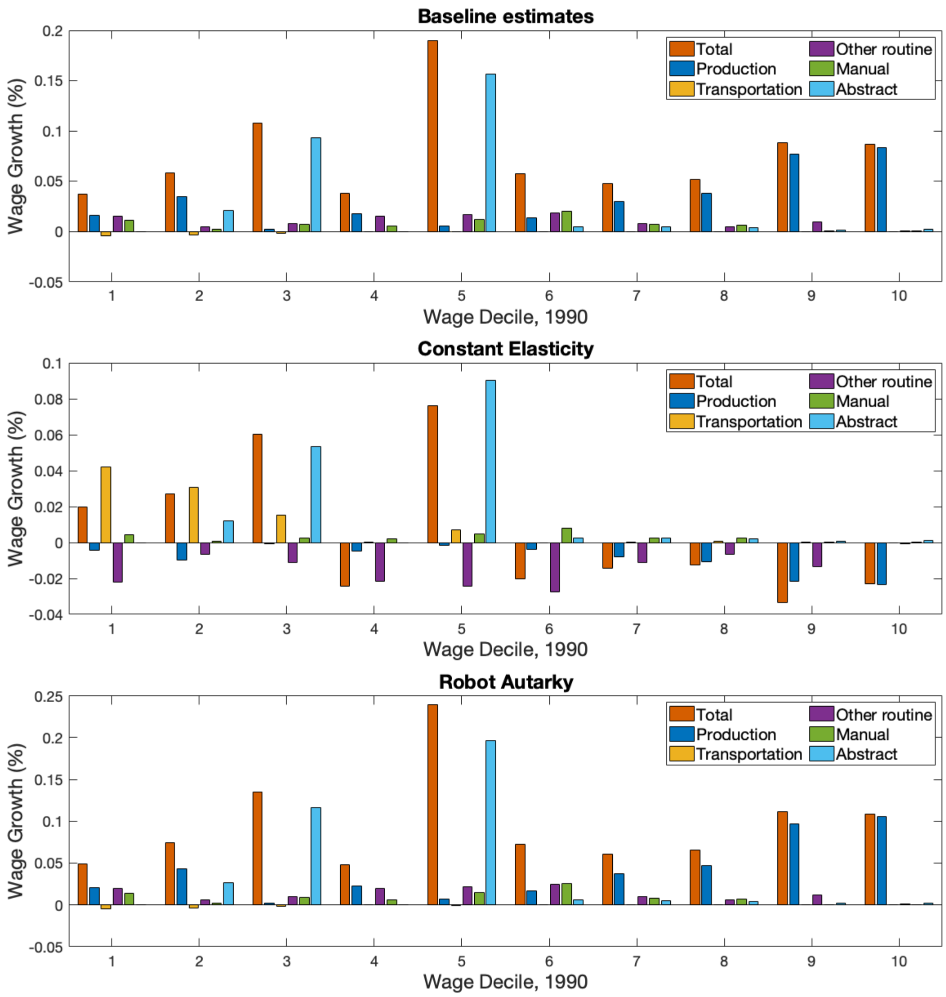

Note: The figure shows the impact of the Japan Robot Shock (JRS) on wages across base year wage deciles. Each panel analyzes different scenarios: The first panel uses the baseline estimates with heterogeneous elasticities of substitution between robots and workers and baseline robot trade shares. The second panel uses the constant elasticity of substitution between robots and workers. The third panel assumes that countries do not trade in robots. In each panel, the total effect shows the effect of all JRSs observed in the data. There are five different sub-effects (Production, Transportation, Other Routine, Manual, and Abstract), each of which shows the effect of JRS in each occupation group, holding constant the JRS of other occupation groups. These five sub-effects add up to the total effect.

.

This mechanism is reminiscent of Galle and Lorentzen (2024), who analyze the interplay between robot automation and the implications of quantitative trade models.

Taken together, the overall positive effects of JRS on US wages are attributed to several factors: the variety of JRS shocks across occupations, the estimated heterogeneous EoS, and the open economy setup.

Appendix E.6. Counterfactual Analysis on Robot Taxes

The Effect of Robot Tax on Occupations. To study the effect of the counterfactual introduction of a robot tax, I consider an unexpected, unilateral, and permanent 6% increase in the robot tax in the U.S., which I call the general tax scenario. I also consider the 33.6% tax only on imported robots and call it the import tax scenario, which implies the same amount of tax revenue as in the general tax scenario and makes the comparison of the two scenarios straightforward.[24] I first examine the effect of the general robot tax on occupational inequality.

Figure Appendix E.4a shows two scenarios of steady-state changes in real occupational wages. In one scenario, the economy is hit only by the automation shock. In the other scenario, the economy is hit by both the automation shock and the robot tax. The results show heterogeneous effects of the robot tax on real occupational wages. The tax mitigates the negative impact of automation on routine production and transport workers, while the tax also reduces the small gains that workers in other occupations would

[24] The 6% rate of the general tax is more modest than the 30% rate considered by Humlum (2021) for the Danish case.

Figure Appendix E.4: The Effects of the Robot Tax on Real Occupational Wages

(a) Steady-state Comparison (b) Transitional Effect of Tax

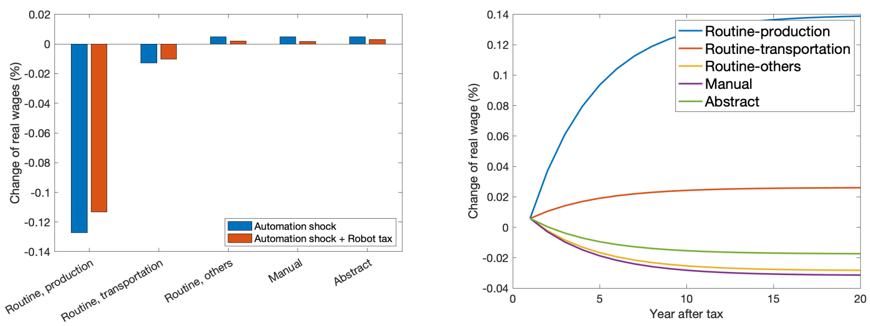

have enjoyed. Overall, the robot tax mitigates the large heterogeneous effects of the automation shock, which could be negative or positive depending on the occupation group, and compresses the effects toward zero.

Figure Appendix E.4b illustrates the dynamics of the effects of the robot tax alone. Although the steady-state effects of the robot tax were heterogeneous, as shown in figure Appendix E.4a, the effect is not immediate, but materializes after about 10 years due to the sluggish adjustment in the accumulation of the robot capital stock. Overall, I find that the robot tax rolls back the real wage effect of automation because the robot tax hinders the adoption of robots. In other words, workers in occupations that have experienced a significant negative automation shock (e.g., production and transportation in the routine occupations group) benefit from the tax, while others lose. Appendix E.7 discusses the effect of the robot tax on the welfare of workers in each occupation.

Figure Appendix E.5: Effects of the Robot Tax

(a) US Real Income

(b) US Robot Stock and Import Robot Price

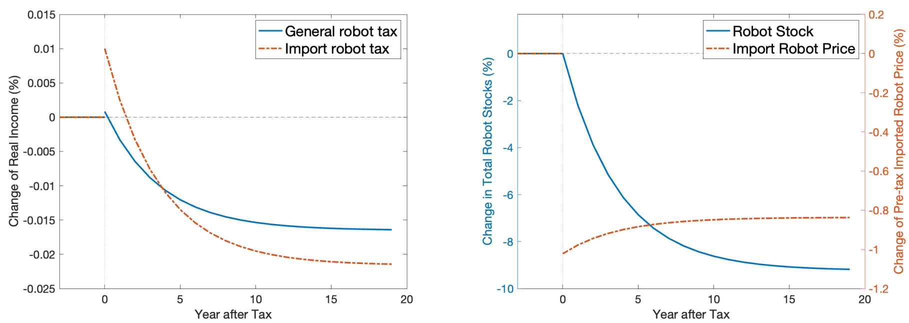

Note: The left panel shows the counterfactual effect on the U.S. real income of the two robot tax scenarios described in the main text over a 20-year time horizon. The right panel shows the effect of the import robot tax on total U.S. robot stocks (solid line) and the pre-tax robot price from Japan (dashed line) over the same time horizon.

Robot Tax and Aggregate Income. I examine how the two robot tax schemes affect US real income. In figure Appendix E.5a, the solid line tracks the real income effect of the general robot tax over a 20-year time horizon after the tax is introduced. First, the magnitude of the effect is small because the cost of purchasing robots is small relative to total production costs. Second, there is a positive effect in the short run, but this effect quickly turns negative and remains negative in the long run.

To understand why there is a short-run positive effect on real income, it is useful to distinguish the source of national income in the model. The total income of a country consists of the wage income of workers, the profit of producers of non-robot goods, and the tax rebate. Since robots are traded, and the US is a large economy that can influence the price of robots produced in other countries, there is a terms-of-trade effect of the robot tax in the US.

Namely, the robot tax reduces the demand for robots traded in the global market and causes the equilibrium robot price to fall along the supply curve. This reduction in the robot price compresses the cost of robot investment, which increases the firm's profit and raises real income. This positive effect is stronger in the Import Robot Tax scenario because the higher tax rate induces a larger decrease in the price of imported robots.

The reason for the different effects on real income in the long run is as follows. The solid line in figure Appendix E.5b shows the dynamic effect of the import robot tax on the accumulation of robot stocks. The robot tax significantly slows the accumulation of robot stocks, reducing the steady-state robot stock by 9.7% relative to the no-tax case. The small robot stock reduces firm profits, which contributes to low real incomes.[25] These results highlight the role that costly robot capital (de-)accumulation plays in the effect of the robot tax on aggregate income. Figure Appendix E.5b also illustrates the dynamic effect on the price of imported robots in the dotted line. In the short run, the price falls due to the reduced demand from the US as explained above. As the SE reaches the new steady state where the US stock of robots is lower, the marginal value of the robots is higher. The effect of this increased marginal value partially more than offsets the short-run effect of the reduced price of robots in the long run.

[25] For each occupation, the counterfactual evolution of robot stocks is similar in percentage terms to each other, and thus similar in percentage terms to the aggregate trend. This is not surprising because the robot tax is ad valorem and uniform across occupations.

Appendix E.7. Robot Tax and Workers' Welfare

This subsection examines how the robot tax affects workers in different occupations. First, I define the equivalent variation (EV) as follows. Consider the unilateral (no reaction in other countries), unanticipated, and permanent U.S. tax on robot purchases as in section Appendix E.6. Denote the consumption flow under the robotized economy with tax as $C'_{i,o,t}$ and that under the robotized but not taxed economy as $C_{i,o,t}$, where the robotization shock is backed out in Appendix D.5. For each country i and occupation o, $EV_{i,o}$ is implicitly defined as follows:

$$\sum_{t=t_0}^{\infty}\left(\frac{1}{1+\iota}\right)^t \ln\left(\left[C'_{i,o,t}\right]\right) = \sum_{t=t_0}^{\infty}\left(\frac{1}{1+\iota}\right)^t \ln\left(C_{i,o,t}\left[1+EV_{i,o}\right]\right). \tag{E.1}$$

Namely, EV is the fraction of the occupation-specific subsidy that would make the present discounted value (PDV) of utility in the robotized and taxed economy equal to the PDV of utility if the occupation-specific subsidy were exogenously given each period in a non-taxed economy. Workers in country i and occupation o prefer the taxed economy if and only if $EV_{i,o}$ is positive.

Figure Appendix E.6a shows this occupation-specific EV as a function of the tax rate. The far left side of the figure is the zero robot tax case, that is, a case with only the robotization shock. Consistent with the occupational wage effects (see figure Appendix E.4a), workers in production and transportation occupations lose significantly due to robotization. In contrast, other workers are roughly indifferent between the robotized world and the non-robotized initial steady state or slightly prefer the former. Moving through the figure, the EV of production and transportation workers improves as the robot tax

Figure Appendix E.6: Robot Tax and Workers' Welfare

(a) Occupational Equivalent Variation (b) Total EV and Revenue

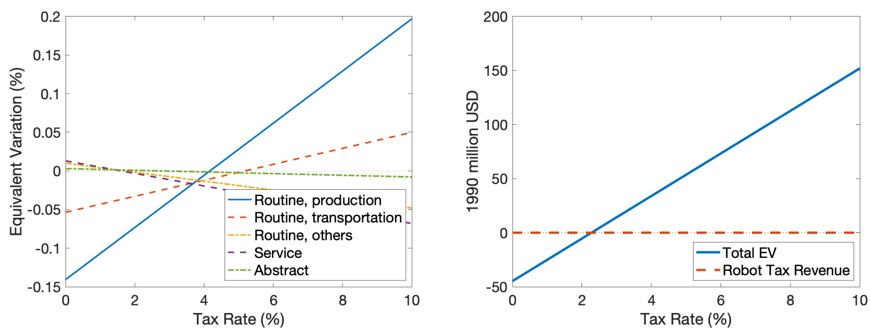

Note: The left panel shows the equivalent variation of US workers, defined in equation (E.1), as a function of the US robot tax rate. The right panel shows the monetary values of equivalent variation aggregated over workers and robot tax revenue as a function of the robot tax rate, measured in 1990 million USD.

reduces the adoption of robots that replace their jobs. The EV of production workers turns positive when the tax rate is about 6%, and the EV of transportation workers turns positive when the tax rate is about 7%. However, these tax rates are too high and would negatively affect the EVs of other occupations. This is because such a high tax rate would significantly reduce robot accumulation in production and transportation occupations, which would negatively affect labor demand in other occupations.

To study whether the robot tax reallocation policy can work, I also compute the equivalent change in monetary value aggregated by occupational group (total EV) and compare it to the robot tax revenue, both as a function of robot tax. Figure Appendix E.6b shows the result. One can confirm that the marginal robot tax revenue is far from sufficient to compensate for the loss of workers, which is concentrated on production and transportation workers in the initial steady state with a zero robot tax rate. The robot tax revenue at this margin is negligible compared to the worker loss due to robotization. As the robot tax rate increases, the total EV increases: When the rate is as high as 2-3%, the sum of the total EV and the robot tax revenue is positive.