Quality-driven Automation*

Kyeongmin Park†

February 13, 2025

Abstract

This paper examines quality improvement as a key driver motivating firms to automate production processes. I present new evidence from the U.S. automotive industry on the relationship between automation and product quality improvement, as well as rising consumer demand for quality. Using microdata, I propose an explanation for the observed increase in automation, highlighting the role of shifting demand toward higher-quality products and changes in the input mix favoring automation to meet this demand. Additionally, microdata from the U.S. Census Bureau covering a broader set of manufacturing industries provide further evidence that greater robot use in establishments is associated with higher product quality. The model, calibrated to study the rise of automation in the U.S. automotive industry, suggests that quality-driven factors account for approximately one-third of the decline in the industry's labor share. A policy experiment demonstrates that slowing the pace of automation through a robot tax could enhance social welfare, while also revealing important distributional effects related to the consumption patterns of different consumer groups for goods produced via automation.

Keywords: product quality, automation, robot, robot tax, structural change, welfare

*I am deeply grateful to my advisor, Mark Bils, for his invaluable guidance and support throughout this project. I also thank Rafael Guntin, Christopher Sleet, Narayana Kocherlakota, George Alessandria, Yan Bai, Fernando Parro, Gaston Chaumont, Hamid Firooz, Matias Moretti, and numerous seminar participants for their helpful comments. Any views expressed are those of the author and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3131. (CBDRB-FY25-P3131-R12083)

†Department of Economics, University of Rochester. Email: kyeongmin.park@rochester.edu

1 Introduction

Rising automation[1] impacts not only the labor market but also the product market, shaping what people consume. However, relatively little is known about its impact on the product market, aside from the observation that firms adopting automation technologies often experience greater sales growth —a stylized fact referred to as the "productivity effect of automation" (e.g., see Acemoglu and Restrepo [2019]). But what drives this growth? Is it merely the result of lower prices from cost-cutting through automation, or does automation enable the production of higher-quality products that shift consumer demand toward these enhanced offerings?

Understanding this distinction is crucial for identifying the drivers of rising automation, evaluating its welfare implications, and informing regulatory policies, as price reductions and quality enhancements affect consumers and market dynamics in distinct ways. If consumer demand shifts toward higher-quality products facilitated by automation, it underscores the significant role of consumer demand in driving automation. This suggests that rising automation is driven not only by supply-side technological progress but also by a broader structural shift in the economy, where increasing demand for quality steers production toward higher-quality goods that are more efficiently produced through automation.

Moreover, if firms automate not just to lower prices but to produce higher-quality products, it raises questions about the inclusiveness of real consumption growth driven by these quality improvements. Since higher-quality products often come with higher prices, the benefits tend to accrue more to wealthier consumers, potentially widening consumption disparities across income groups. If regulatory policies aim to ensure that the gains from automation are shared inclusively, they should account for these distributional effects, extending beyond the traditional focus on labor market impacts.

The existing literature has emphasized the cost-saving effects of automation. As a result, it attributes the relative sales growth of automating firms to price reductions, based on demand models that abstract from shifts in consumer demand. However, this view is at odds with some empirical findings. First, firm surveys indicate that product quality improvement is one of the primary reasons for adopting automation technologies (see Figure A1 for the results from the 2023 Annual Business Survey). Second, while existing studies document that firms adopting robots experience sales growth, they do not provide evidence of a significant relative

[1] Automation generally refers to substituting capital for labor in a widening range of tasks. This "capital" encompasses technologies such as robots, AI, computer-numerically controlled machines (CNCs), and software. To ensure clarity and consistency with the data, this paper primarily uses the term "automation" to refer to the use of robots, representing the most prominent example of automation.

decline in product prices.

These facts suggest that consumer demand shifts toward firms that automate their production processes, with improved product quality likely driving much of this shift. This paper highlights this underexplored aspect of automation by presenting micro-level evidence that a firm's decision to automate is closely linked to its efforts to enhance product quality in response to increasing consumer demand for better products.

Providing evidence on the quality improvements associated with automation, alongside rising consumer demand for quality, requires detailed micro-level data to enable meaningful comparisons between product quality and the degree of automation. To this end, I first focus on the U.S. automotive industry, using data compiled from various sources. I then provide additional evidence from a broader range of U.S. manufacturing industries, drawing on U.S. Census Bureau microdata, including establishment-level data on robot usage and product unit prices.

The automotive industry is a natural starting point to study the effects of automation on product quality and consequent consumer welfare for three reasons. First, it is the most automated industry, extensively using industrial robots.[2] Second, vehicles are among the products for which consumers exhibit the strongest demand for quality (see Figure A2). Third, it provides feasible product quality measurements based on unit price and quality rating scores.

I present three main empirical findings from the U.S. automotive industry. First, plants adopting robots experience sales expansion, even though their unit prices show a relatively larger increase in the periods following adoption. Identifying robot adoption based on data collected on robot investments, I analyze its cumulative effects on plant-level variables using the local projections event study method. The results indicate that a demand shifter, potentially higher product quality, is needed to explain the sales growth observed in robot-adopting plants.

The second finding is that, within the sample of U.S. vehicle manufacturing plants, upgrading product quality is associated with a decline in labor share. By combining information on vehicle models produced in each plant, including unit prices and quality rating scores, I construct a plant-level quality measure. Panel regression analysis at the plant level reveals that plants with a relatively larger increase in this quality measure are associated with a

[2]In terms of robot density (robots installed per 10,000 employees), U.S. automotive industry (robot density 1,287) use robots about 10 times more intensively compared to general industries (robot density 139) as of 2019 according to the International Federation of Robotics.

relatively larger decrease in labor share. This suggests a shift in the input mix of production toward automation as product quality improves.

The third finding is that consumers show rising demand for quality as their consumption rises. Building on Bils and Klenow [2001]'s measure of the quality Engel curve,[3] which quantified the extent wealthier households consume higher-quality goods, we should expect consumer demand to shift towards higher-quality goods as overall consumption rises with economic growth. Using product-level U.S. vehicle market sales data from the past 30 years, I document a persistent shift in consumer demand toward higher-quality vehicle models, consistent with the rising demand for quality in a growing economy.

Extending beyond the automotive industry, I present evidence for a broader set of manufacturing industries, using microdata from the U.S. Census Bureau that includes establishment-level data on robot use and product unit prices. The findings show that establishments with more intensive robot use produce higher-quality (higher unit price) products compared to their peers producing the same product categories.

These empirical findings collectively indicate that automation and product quality are jointly determined. As consumer demand for quality increases, market share shifts toward higher-quality products. In industries where automation technology offers a comparative advantage over human labor—particularly for quality upgrades that require intricate production processes with minimal tolerance for human error—firms that enhance product quality are more likely to adopt higher levels of automation, leading to subsequent sales expansion. I refer to this dynamic as ***quality-driven automation***.

Based on this evidence, the paper proposes a framework in which rising automation is driven by the comparative advantage of robots over labor in producing higher-quality products, with consumers increasingly valuing higher quality as overall consumption grows.[4] As

[3]The slope of the quality Engel curve indicates the elasticity of the unit price of a consumer's purchase with respect to their total nondurable consumption.

[4]A central theoretical question is why previous analyses focusing on the cost-saving effects of automation are not isomorphic to the quality-driven framework proposed in this paper. In demand models that abstract from quality—or treat quality as equivalent to quantity in consumer utility, as in standard CES preferences with monopolistic competition—sales growth can be interpreted as resulting from reductions in the quality-adjusted price. From this perspective, even when firms expand sales by introducing higher-priced, higher-quality products, the outcome is treated analytically equivalent to price reductions driven by cost savings. (e.g., "Under the assumption of CES preferences and monopolistic competition, cost and appeal (quality) enter equilibrium firm revenue in the same way"(Hottman et al. [2016], p. 1296).)

However, a framework that solely emphasizes cost-saving effects fails to account for crucial data patterns: (a) market shares persistently shift toward higher-quality products, and (b) automation intensity tends to increase with the quality of the products firms produce. This subtle yet systematic distinction cannot be fully accounted for by the idea that automating firms expand sales merely by offering lower-cost goods (in terms of quality-adjusted price) through cost reductions.

the first paper to incorporate aspects of quality growth and evolving consumer demand into the analysis of automation, it examines how the welfare impacts of automation are distributed across households. Additionally, it explores how these effects might change if the government intervenes to curb the pace of automation due to concerns over labor displacement.

Given the distinct implications of a quality-driven framework, an important policy question arises. How should the government respond to the rise of quality-driven automation? While firms automate to efficiently produce higher-quality products, benefiting consumers at large, this process may also result in significant labor displacement among workers in sectors vulnerable to automation. This requires a reevaluation of the welfare effects of automation regulation policies, such as a robot tax. Unlike analyses that abstract from quality, my study suggests that policies aimed at slowing the pace of automation to address labor displacement may also affect quality growth, with distinct distributional effects across consumers.

I address this question by conducting a policy experiment within a calibrated model. There are two sectors in the model: a differentiated goods (automated) sector where quality-driven automation occurs[5], and an outside (non-automated) sector. Similar to Oberfield [2023], households have income differences and non-homothetic preferences over product quality, with quality a luxury good.

On the firm side, I extend the framework in Fieler et al. [2018] to incorporate changes in the production function based on product quality, along with a standard task based model (e.g., Acemoglu and Restrepo [2018]) in the automation literature. The parameters of the production function vary with output quality, assigning greater weight to robots as quality increases.

I embed rising automation within the context of economic growth. As consumption grows, the product space shifts toward higher-quality, leading to increased automation within the automated sector. Workers employed in the automated sector reallocate to the outside sector, with reallocation frictions. The model demonstrates structural change as the economy transitions to a state where most workers are employed in the outside sector, following reallocation amid labor market disruption.

I calibrate the model to the data by mapping the automated sector in the model to the

[5]The quality-driven framework may not explain all automation, as there are industries where automation is primarily not related to quality improvements. However, this article focuses on quality-driven automation and classify all other sectors, which do not involve quality-driven automation, as the outside sector.

U.S. automotive industry.[6] The calibrated model shows that a robot tax benefits workers initially employed in the automated sector by mitigating labor displacement. However, it raises the price index, disproportionately affecting households with consumption baskets focused on higher-quality goods. Nonetheless, due to reallocation frictions, the robot tax may enhance short-term social welfare by preserving wages in the automated sector.

In the long run, however, a tradeoff emerges: maintaining the robot tax ultimately disadvantages even workers initially employed in the automated sector, as it diminishes their option value—the potential benefits they might gain from reallocating to the non-automated sector. An optimal robot tax policy, therefore, suggests gradually phasing out the tax as labor reallocation advances, balancing short-term protection for displaced workers with long-term production efficiency.

Related Literature This paper relates to previous studies highlighting a stylized firm-level fact, known as the "productivity effect of automation": firms adopting automation technology often experience significant sales expansions, large enough to increase overall employment despite a decline in labor share (e.g., Acemoglu and Restrepo [2019]; Koch et al. [2021]; Restrepo [2023]). Abstracting from the quality aspect, previous studies have attributed this sales expansion to cost reductions and ensuing product price declines. However, evidence supporting this claim is limited,[7] mainly due to the lack of micro-level datasets that separate firm sales into unit price and quantity, while also providing data on the degree of automation.

This article contributes to the literature by providing a complementary explanation for the productivity effects of automation, highlighting quality improvement as a key demand shifter. This view relates to empirical findings that robot adoption is associated with firms' strategic priority in improving product quality among Canadian firms (Dixon et al. [2021]), and that robot diffusion leads to export quality upgrading based on cross-country and cross-

[6] It provides the lower bound for the share of the economy under quality-driven automation. Determining the extent to which this framework applies beyond the automotive industry remains an empirical question for future investigation.

[7] No such micro evidence exists in the U.S. context. Two earlier studies using French firm export data estimated the impact of robot adoption on product prices, but their results are mixed. Bonfiglioli et al. [2020] report that robot adoption is associated with a relative increase in export prices among French firms from 1996 to 2013, while Acemoglu et al. [2020] find that robot adopting firms between 2010-2015 show a relative decrease in export prices.

industry data (DeStefano and Timmis [2024]).[8] My view aligns with a strand of literature that quantifies the sources of firm/establishment size by distinguishing between technical efficiency (marginal cost) and quality (demand shocks), where true productivity is interpreted as a composite of both technical efficiency and quality (see Eslava et al. [2024]; Hottman et al. [2016]).

This paper explains the rise in automation through non-homotheticities in both demand and production. On the demand side, it builds on the insights of Bils and Klenow [2001], which documents that as households become richer, they demand higher quality goods. To model this, I use non-homothetic CES preferences, where households with larger consumption place more weight on higher-quality goods, reflecting a higher valuation of quality (e.g., Handbury [2021]; Faber and Fally [2022]; Nord [2022]). The model specifically builds on the setup from Oberfield [2023] to demonstrate how demand shifts toward higher-quality goods as the economy grows.

On the production side, the relationship between rise in automation and quality upgrading connects to the literature on how firms' input mix vary with output quality. The model draws from the work of Fieler et al. [2018] and Demir et al. [2024], where their production functions vary by quality. This perspective also relates to the O-ring theory in the seminal work of Kremer [1993], along with Schott [2004], Verhoogen [2008] and Kugler and Verhoogen [2012].

By incorporating these two non-homotheticities, which have not been previously explored in the context of automation, this paper offers fresh insights into the literature on the welfare effects of automation (e.g., Prettner and Strulik [2020]; Jaimovich et al. [2021]; Eden and Gaggl [2018]). Particularly, my approach emphasizes how automation interacts with consumer preferences as the product space evolves toward higher quality. This perspective provides a framework for analyzing the effects of automation on quality growth, going beyond the previous focus on cost reductions that overlook the quality dimension. It thereby offers a deeper understanding of the underlying drivers of rising automation and its distributional impacts. Emphasizing welfare implications in the product market relates to Aghion et al. [2020], while analyzing impacts across households with diverse consumption baskets connects with the findings of Arvai and Mann [2022].

The policy experiment and welfare analysis in this paper relate to the literature on automation regulation using robot taxes. Most related are Guerreiro et al. [2022] and Beraja

[8]My study, however, distinguishes itself from related papers that consider the quality aspect associated with automation by: (a) using direct measurements of product quality rather than relying on firms' reported strategic priorities; (b) conducting analysis at the product-establishment level, which allows for clearer comparisons across units of observation compared to cross-country or cross-industry analyses; and (c) exploring broader macroeconomic welfare implications through the application of a structural model.

and Zorzi [2024], which provide quantitative results from structural models with limited labor reallocation.[9] This paper provides novel insights into the discussion of optimal automation regulation policies by framing rising automation not merely as a result of declining relative factor prices between robots and labor, but as part of a broader shift driven by evolving consumer demand and concomitant changes in the production function in a growing economy.

The remainder of the paper is organized as follows. Section 2 describes the data used for empirical analysis. Section 3 presents empirical evidence of quality-driven automation in the U.S. automotive industry. Section 4 outlines the model. Section 5 describes the calibration and presents the results of quantitative evaluations. Section 6 explores policy experiments with a robot tax. Finally, Section 7 concludes.

2 Data Sources and Measurement

This section explains the construction of key variables and outlines the data sources used. The study primarily uses two sources of data: (1) a compiled dataset of the U.S. automotive industry, and (2) U.S. Census Bureau microdata covering a broader range of manufacturing industries.

2.1 U.S. Automotive Industry Data

Data have been collected from various sources, including printed hard copies of industrial journals[10] and online resources. The dataset covers the U.S. automotive industry, with sales including all vehicle model sales in the U.S. market, and production covering all vehicles produced by vehicle manufacturing plants located within the U.S. The constructed dataset spans the years 1992 to 2023, although some variables have shorter sample periods. The following subsections describe the key variables and their respective sources.

2.1.1 Product Data

Defining a Product. The vehicles sold in the market are highly differentiated products, where each make (e.g., Toyota) produces multiple models (e.g., Camry), and each model

[9] In Guerreiro et al. [2022], job reallocation occurs as new generations enter non-automated (non-routine) jobs, while previous generations in automated (routine) jobs can't reallocate. In Beraja and Zorzi [2024], reallocation frictions include random job change opportunities, temporary unemployment, and productivity losses in new jobs.

[10] Two industrial journals were extensively used to construct the dataset for this study by scanning hard copies. The Wards' Automotive Yearbook (for unit prices, sales in units for each vehicle model, and production in units at each plant) was accessed through the University of Rochester's Rush Rhees Library, while the Harbour Report (for employment data at each plant) was accessed through the Library of Congress.

offers various trims (e.g., LE, SE, XLE, XSE, etc.). The baseline definition of products in this study is at the make-model level (e.g., Toyota Camry). This choice was made because sales and production data are available and can be analyzed at this specific level of granularity.

Sales (in units). Unit sales data for vehicle models in the U.S. market is available on a calendar year basis and is sourced from the Ward's Automotive Yearbook. By combining this quantity information with model-year vehicle prices, I calculate market sales of each vehicle model. The terms "sales" and "revenue" are used interchangeably, with "unit sales" specifically referring to the number of vehicles sold.

Price. The price variable used in this study is the Manufacturer's Suggested Retail Price (MSRP) for the base model, available for each model-year. The primary source of price data is the Ward's Automotive Yearbook. For missing data, I supplemented the information using the Kelley Blue Book website, which provides the original MSRP for certain vehicles.

Quality. Quality is defined as any attribute that increases consumer's willingness to pay for a variety as defined in the literature (e.g., Khandelwal [2010]). This notion correspond to any shifter of demand conditional on price.[11] I use vehicle unit price as the primary measure for vehicle quality within specific vehicle types (cars, SUVs, pickups, and vans).

Vehicle is one of the products where unit price can effectively proxy for product quality, given its widely differentiated quality levels. However, I address a potential concern that unit price may not always perfectly correlate with quality. This discrepancy may arise from factors such as higher prices reflecting increased manufacturing costs or markups, rather than inherent quality. To address this potential concern, I incorporate vehicle quality rating scores as a supplementary measure for quality alongside the unit prices.[12]

Quality rating scores are sourced from the J.D. Power website, a company that provides vehicle quality ratings. The sample covers the years 2006 and onwards, where such ratings are available. Rating scores are provided as numbers relative to the vehicle's peers competing within the segment (e.g., compact car, midsize SUV, and large premium SUV), with higher scores indicating higher quality.

As vehicle "quality" can encompass multiple attributes, J.D. Power provides three separate rating categories: 1)Reliability; 2)Driving Experience; and 3)Resale Value. The "Relia-

[11] There are some nuanced differences in how this shifter is interpreted in the literature. Some studies use the concept of "appeal" to avoid taking a stance on whether the shifter of demand conditional on price arises from vertical quality differentiation or subjective differences in consumer taste (e.g., Hottman et al. [2016]). My perspective is to take a stand that it arises from vertical quality differentiation.

[12] Using quality ratings scores as direct measure for product quality relates to Bai et al. [2020] for vehicles, and Chen and Juvenal [2018] for wines.

bility" rating captures the dependability of a vehicle, accounting for defects and malfunctions. The "Driving Experience" captures how satisfying a vehicle is to own and drive, evaluating the overall advancement of features. "Resale Value" is based on the expected depreciation of the vehicle over the next three years.[13]

Here, I first validate the use of unit price as the primary measure of product quality by showing that variations in unit price are largely explained by variations in quality rating scores. To do this, I estimate the semi-elasticity of vehicle unit price with respect to the three quality rating categories, controlling for segment-year fixed effects. These fixed effects capture quality differences across segments (e.g., midsize car vs. midsize premium car). By partialling out these between-segment differences, the estimated coefficients show how much of the within-segment unit price differences can be explained by quality rating scores.

Table 1: Semi-elasticity of vehicle unit price w.r.t. quality rating scores

Dependent variable →	ln(unit price)
Reliability ($q^{(1)}$)	0.012 (0.003)
Driving Experience ($q^{(2)}$)	0.039 (0.003)
Resale Value ($q^{(3)}$)	0.014 (0.003)
Segment×Year FE	✓
Observations	1,821
R^2	0.90

Notes: Standard errors in parentheses. The sample covers the years 2006-2021. Quality rating scores in three categories ($q^{(1)}, q^{(2)}, q^{(3)}$) are standardized Z-scores within vehicle segment-year.

Table 1 presents the coefficient estimates. The results show that variations in vehicle unit price reflect quality differences. A vehicle model with a "Driving Experience" score one standard deviation above the average is about 4 percent relatively more expensive compared to its peers. Similarly, the unit price increases by approximately 1 percentage point in response to a one standard deviation increase in "Reliability" or "Resale Value" scores. The high R^2 of the regression supports interpreting differences in unit prices as reflecting quality differences, particularly in the context of vehicles.

In the following empirical analysis, where concerns arise about using unit price as a proxy for product quality, I address this by instrumenting vehicle unit price with three categories of quality rating scores and vehicle segment dummies. While this approach takes into account

[13]For fuller details on how quality ratings are determined, please refer to the subsection A.3.

quality differences across vehicle segments, as indicated by the average unit price differences between segments (e.g., midsize SUVs vs. midsize premium SUVs) captured by segment dummies, it also captures variations in quality within each segment based on quality rating scores. In addition, for better alignment with recent literature, I present some of the results using the quality measure introduced by Khandelwal et al. [2013], which has been widely used (e.g, DeStefano and Timmis [2024]; Bas and Paunov [2021]; Chen and Juvenal [2018]).[14]

2.1.2 Plant Data

Examining the association between automation and product quality upgrades requires plant-level variables, particularly measurements of automation and product quality. This paper uses labor share as a proxy for the degree of automation and exploits events of robot investments to capture plant-level robot adoption. For plant-level product quality, I construct a variable based on the sales-weighted average of product quality of all products produced at each plant, following a method similar to that used by Kehrig and Vincent [2021] to construct plant-level price premiums.

Product-level production (in units). Each vehicle manufacturing plant is a multi-product facility, producing multiple vehicle models simultaneously. I collect vehicle model-level production data by plant on a calendar-year basis from Ward's Automotive Yearbook. Combined with vehicle price data, this allows me to calculate revenue for each vehicle model produced at the plant. By aggregating across all vehicle models produced at each plant, I calculate total plant-level production (in units) and revenue.

Labor share. Plant-level labor share is defined as the total payroll divided by total revenue. I construct the labor share of plant j producing a set $M(j,t)$ of vehicle models, located in state $s(j)$, in year t as follows:

$$LS_{jt} = \frac{W_{s(j)t} L_{jt}}{\sum_{i \in M(j,t)} p_{it} x_{jit}}$$

$$\text{where } M(j,t) \equiv \{i : \text{vehicle model } i \text{ produced in plant } j \text{ at year } t\}$$

The denominator, total sales, is calculated by summing all vehicle model-level sales. These sales are derived by multiplying the number of units produced (x_{jit}) by the unit price of each

[14]The measure is derived from a CES demand system with an assumed elasticity of substitution, where products with higher market shares, conditional on their price, are regarded as having higher quality (acting as a demand shifter). This measure is estimated as a regression residual (see subsection A.4 for a derivation). I construct this quality measure using the vehicle model-level market share within each vehicle type (cars, SUVs, pickups, and vans).

vehicle model (p_{it}) produced in each plant, using data from the Ward's Automotive Yearbook.

The numerator, payroll, is derived from data obtained through multiple sources. Initially, information on the number of workers at each plant (L_{jt}) is sourced from the Harbour report, an industry journal that provides automotive manufacturing plant-level data on number of workers for specific years (1992, and 1994-2007). For missing years, additional employment data is supplemented through web searches on annual reports, plant homepages, and archived news articles. I combine this plant-level employment information with state-industry level[15] average annual wage data ($W_{s(j)t}$), sourced from the BLS Quarterly Census of Employment and Wages (QCEW), to calculate the total payroll.

Robot adoption. I define a robot adoption as the first observed robot investment at each plant during the sample period. This approach is necessary because comprehensive information on the intensity of robot use at the plant level, represented as a stock variable, is limited.

To collect data on robot investments, I first gather data on robot imports. When companies import products into the U.S. via sea freight, they are required to complete a document known as the Bill of Lading (BOL), which records information about the importer, supplier, and product description. I obtain data on the quantity and timing of industrial robot imports by each plant from this document, accessed through the data company ImportGenius, which covers historical U.S. ocean freight imports since 2006.

To supplement the robot investments not captured through the import database — such as robots procured through domestic channels — I also search online and collect information on plant investments explicitly related to robots. These additional data are gathered from publicly available sources, including archived news articles and company websites. By combining these sources, I construct plant-level data on robot investments starting from 2000.

Plant-level quality. Plant-level product quality q_{jt} is measured as the sales-weighted average log unit price of vehicle models produced in plant j in year t. In the empirical section, in the instrumental variable (IV) analysis aimed at validating the findings, this q_{jt} is instrumented using sales-weighted vehicle quality measures derived from either: 1) JD Power (incorporating quality rating scores and segment dummy variables); or 2) the regression residual as described by Khandelwal et al. [2013].

[15]NAICS 336 "Transportation Equipment Manufacturing". The choice of state-industry level analysis was made due to data availability, as more detailed-level data (e.g., county-industry) are often restricted by the Bureau of Labor Statistics (BLS) due to disclosure concerns.

2.2 U.S. Census Bureau Microdata

Combining datasets that include product-level unit prices and establishment-level robot usage, I test the hypothesis that establishments producing higher-quality (higher unit price) products use robots more intensively than their peers within the same product category. The dataset covers a set of manufacturing industries where unit price data is available, focusing on the years 2017–2022.

2.2.1 Product Data

Product is defined by narrowly defined product categories (10-digit NAPCS product code). In the Census of Manufacturers (CMF), the product trailer records the value of sales generated by individual products for each establishment. It also collects information on the physical quantity of products shipped for a sample of establishments whenever a meaningful metric is available. I calculate product-level unit prices by dividing revenue by the physical quantity of individual products. To make prices comparable across establishments, I normalize prices at the level of the 10-digit product category c, following Kehrig and Vincent [2021].

$$\tilde{p}_{jct} \equiv ln\, p_{jct} - \overline{ln\, p_{-j,ct}} \tag{1}$$

$$\text{where} \quad \overline{ln\, p_{-j,c,t}} \equiv \sum_{k \neq j} \frac{p_{kct} x_{kct}}{\sum_{k \neq j} p_{kct} x_{kct}} ln\, p_{kct}$$

Here, $\tilde{p}_{jct}$ denotes the log-point difference that the price of product c sold by establishment j at time t and the sales-weighted average price of the same product charged by all other establishments $k \neq j$ in the same year. I refer $\tilde{p}_{jct}$ as the product price premium that establishment j charges relative to its peers within a product category, using it as a proxy for product quality.

2.2.2 Establishment Data

I link the product data with establishment-level robot usage from the Annual Survey of Manufactures (ASM). The ASM (2018–2021) provides direct data on the degree of robot use, measured by the total number of industrial robots in operation at the establishment. It also includes data on capital expenditures for new and used industrial robots, as well as the number of robots purchased during the reference year.

I then match the robot usage data to the closest Census years where unit price data

is available. Specifically, I match unit prices from the 2017 and 2022 CMF to automation intensity data from the 2018 and 2021 ASM. The linkage is made using the establishment linkage variables in the Longitudinal Business Database (LBD). The resulting pooled sample consists of two waves, which I refer to as the 2018 and 2021 samples.

3 Empirical Results

The purpose of the empirical section is to provide evidence for the following industry dynamics: Consumer demand evolves through shifts toward higher-quality products as overall consumption grows. At the same time, the input mix of production shifts toward greater automation with improvements in product quality. As plants increase product quality, automation intensifies, and sales expand. The first three subsections present findings corresponding to each of these points, focusing on the U.S. automotive industry. In the final subsection, I provide additional evidence from a broader range of U.S. manufacturing industries, drawing on U.S. Census Bureau microdata, which includes establishment-level data on robot usage and product unit prices.

3.1 Plant responses to robot adoption events

First, I study the response of variables of interest at the plant level around the event of robot adoption, using robot adoption as a direct indicator of automation. I define a robot adoption event as the first observed robot investment at each plant during the sample period. To track this, I collect information on whether each plant made robot investments in year t. The collected data on plant-level robot investments covers the period from 2000 onward.

By combining this with other plant-level data I have compiled, I conduct local projections to estimate the cumulative effects of robot adoption on the outcome variables of interest. D_{jt} is defined as a binary absorbing treatment, such that if a plant j invests in robots at year t, $D_{jt} = 1$ and it stays as 1 afterwards (i.e., $D_{j,t+h} = 1$ for $h > 0$). Otherwise, $D_{jt} = 0$. I focus on the impacts of the first treatment ($\triangle D_{jt} = 1$), interpreting this as plant-level robot adoption in the sample.

This interpretation requires clarification. First, given the widespread use of robots in the vehicle industry, it is unrealistic to assume that non-treated plants have no robots at all. Rather than viewing the analysis as a comparison between plants with and without robots, it should be interpreted as capturing the response to increases in robot use versus no significant change. Second, due to data limitations on the stock of robots at each plant

from public sources,[16] I cannot rule out the possibility that the first observed treatment reflects replacement investment rather than an increase in automation. Acknowledging these limitations, this identification strategy focuses on events most indicative of increased robot use, capturing their impact as fully as possible with the available data.

To estimate the response of plants to the first treatment, I estimate the following equation:

$$y_{j,t+h} - y_{j,t-1} = \alpha_t + \beta_h \triangle D_{jt} + \varepsilon^h_{jt} \tag{2}$$

$$D_{jt} = \begin{cases} 1 & \text{if plant } j \text{ invests in robots in year } t \quad (D_{j,t+h} = 1 \text{ for } h > 0) \\ 0 & \text{o.w.} \end{cases}$$

Here, y denotes the outcome variable, and the coefficient of interest, β_h, captures the cumulative effects of robot adoption over a horizon of h years. α_t represents year fixed effects. I apply this specification to various outcome variables, including labor share, log sales, log payroll, and log unit price[17], to analyze plant-level dynamics associated with the event of robot adoption. I define the control group as plants that have not invested in robots by the period of comparison $t + h$ (clean control).

Figure 1 displays the estimates for the all four variables. Robot adoptions were observed in 37 out of 74 plants in the unbalanced panel during the sample period. While the estimates show relatively large standard errors, likely because the robot investments data collected from online sources may be incomplete, they still reveal dynamics with important implications for what occurs at the most granular level of production (i.e., the plant) around the event of robot adoption. Consistent with the stylized facts documented in the literature, sales expand following the period of first robot adoption. However, this sales expansion is not accompanied by a significant decline in unit prices; instead, unit prices actually increase.

The upper-right panel illustrates the response of unit price, constructed as the sales-weighted average of the product log(unit price) for all products produced at each plant. The estimates indicate that the plant-level unit price tends to increase following the event of robot adoption. This suggests that the sales expansion observed in plants adopting robots is not driven by price reductions but rather by a demand shock that shifts consumer demand toward these plants despite relatively larger increase in unit prices.

This suggests that the adoption of robots is driven not solely by the need to reduce

[16] I address this limitation using restricted-use U.S. Census Bureau microdata.

[17] Calculated as the sales-weighted average of the log unit prices of vehicle models produced on each plant's production lines, this measure serves as a proxy for plant-level quality.

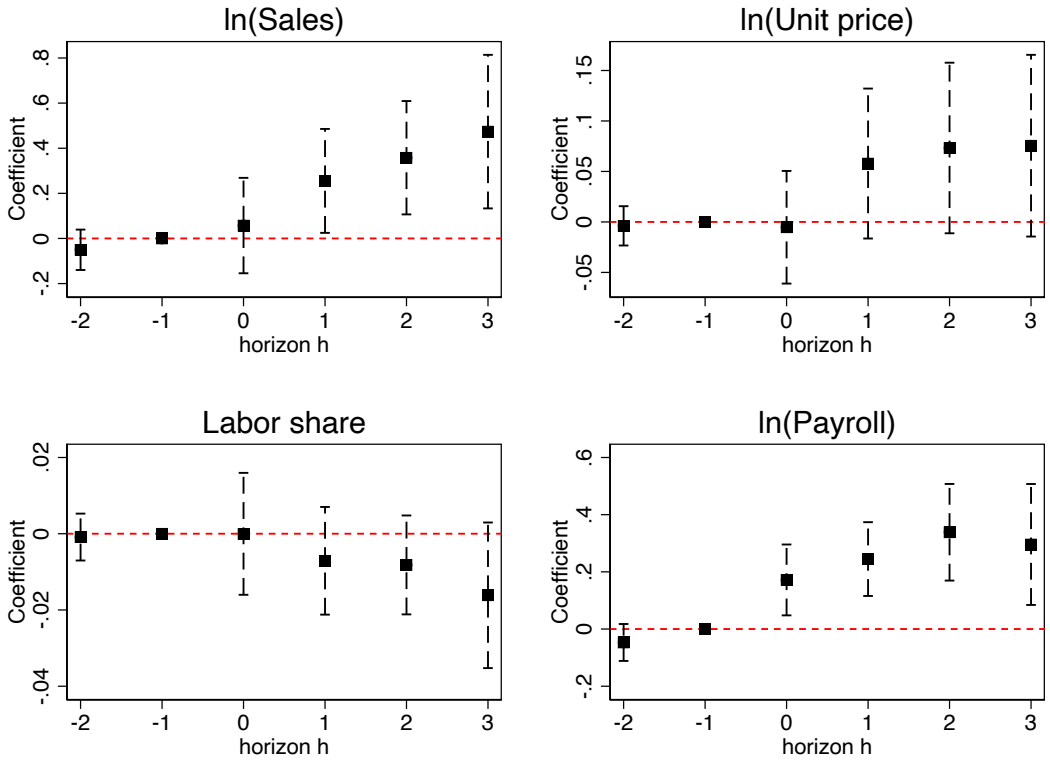


Figure 1: Responses to Robot Adoption Event

Notes: Each point represents the estimated coefficient β_h from the local projections in Equation 2. The dotted line indicates the 95% confidence interval. Standard errors clustered by plant. The coefficients capture the response of the outcome variable with respect to the robot adoption event (first treatment) over the horizon h. The regression controls for year fixed effects. In the upper right panel, ln(unit price) represents the sales-weighted average of the product ln(unit price) for all products produced at each plant.

production costs for existing products, but also by firms' choices regarding their product offerings. If firms shift their input mix toward greater automation when producing higher-quality goods, the observed demand shift toward robot-adopting plants—despite their higher relative unit prices—can be accounted for.

Labor share and payroll exhibit responses that qualitatively align with the patterns of firm-level robot adoption documented in the literature: an increase in payroll despite a decline in labor share. Consistent with the productivity effects of automation, the results suggest that while labor share declines as the input mix shifts toward greater use of robots, the expansion in sales is substantial enough to result in a relative increase in payroll.

There are some caveats to these results, as the data used to identify robot adoption across plants may be incomplete, having been compiled from multiple sources. Most notably,

missing instances of robot investments could bias the estimates by misclassifying treated plants as part of the control group. Despite these limitations, this analysis is one of the few to examine the impacts of robot adoption at the establishment level, providing rare and valuable insights into the response of unit prices—an aspect largely overlooked in previous research due to data constraints.

3.2 Quality upgrading associated with labor share decline

Firms continually upgrade product quality by either improving existing models with new versions or introducing new varieties. I investigate how this product quality upgrading is associated with input choices, particularly in the context of robotization within the vehicle industry.

3.2.1 Changes in labor share depend on plant-level quality changes

Since the direct measurement of robot use is limited in sample availability, I instead use the labor share as a proxy for automation and estimate how changes in labor share are associated with plant-level quality changes. Specifically, I interpret a relatively larger decrease in labor share as an indication that the input mix has shifted more toward automation and away from human labor. The labor share of plant j in year t, LS_{jt}, is expressed as:

$$LS_{jt} = \alpha_j + \beta q_{jt} + \gamma X_{jt} + FE + \varepsilon_{jt} \tag{3}$$

Here, α_j denotes a set of plant dummies, q_{jt} represents plant-level quality measured as the sales-weighted average log unit price of vehicle models produced. X_{jt} is a vector of additional controls that includes the sales share of each vehicle type (cars, SUVs, pickups, or vans) produced within the plant. This accounts for differences in production processes across vehicle types, ensuring that variations in labor share are not simply driven by shifts in the plant's production composition. FE represents a set of additional fixed effects, including year, state, and company fixed effects. These fixed effects account for unobserved factors that might influence labor share or quality changes, such as temporal trends, regional characteristics, or firm-specific attributes. ε_{jt} denotes the error term.

The coefficient of interest is β, which represents the semi-elasticity of the labor share with respect to plant-level quality, while controlling for the effects of other variables. If product quality upgrading is associated with greater automation, we expect β to be negative. Table 2 column (1) reports the estimate from the full sample. Column (2) reports the estimate from a subsample where IV estimation from JD Power quality ratings is feasible, ensuring

comparability with estimates in columns (2)-(3).

Column (1) indicates that, within a plant, a 10% relative increase in product quality is associated with a 0.2 percentage point relative decline in the revenue labor share. Given the small magnitude of the dependent variable by definition, I also report the standardized coefficient for easier interpretation. The standardized coefficient indicates that a one-standard-deviation increase in product quality is associated with approximately a 0.6-standard-deviation decrease in the labor share within a plant, underscoring that the effect is economically significant.

Table 2: Regression of Plant Labor Share on Quality

	(1)	(2)	(3)
Dependent variable →		LS_{jt}	
q_{jt}	-0.024	-0.027	-0.032
	(0.005)	(0.007)	(0.010)
standardized coef.	-0.57	-0.53	-0.64
Specification	OLS	OLS	2SLS
Instrument			JD Power
Sample Year	1993-2021	2006-2021	2006-2021
Observations	1,311	575	575
KP F-stat			82.66
R^2	0.547	0.595	-

Notes: Standard errors clustered by plant in parentheses. Plant-level quality proxy q_{jt} is measured as the sales-weighted average log unit price of vehicle models product in plant j in year t. Regression controls a set of fixed effects including year, state, company and plant. Observations are weighted by the sales share within a year (s_{jt}). The KP F-stat is the Kleibergen-Paap rk Wald F statistic for weak instrument.

IV estimation. In Equation 3, there is a potential concern about endogeneity, because the proxy for plant-level quality (q_{jt}) is constructed using product unit prices, and the dependent variable, labor share (LS_{jt}), has unit prices in its denominator. Factors such as markups, unrelated to inherent product quality, could inflate the denominator of the dependent variable and confound the estimates.

To address this concern, I employ an IV strategy that leverages external quality rating scores for product quality. This instrument is intended to capture the quality component reflected in the unit prices, ensuring that the estimated β is not confounded by other factors that affect unit prices but are unrelated to quality.

The instrument, denoted as Z_{jt}, is constructed by aggregating the quality of vehicle models

produced in each plant. First, I develop a vehicle model quality measure. In column (3), this measure is obtained from a linear projection of the vehicle log unit price on JD Power quality ratings and segment dummies.[18] Next, I aggregate these vehicle model quality measures by averaging them, weighted by sales, at the plant level.

The instrument is strongly correlated with q_{jt} (**relevance**), as higher quality rating scores and segment dummies are strongly associated with higher unit prices. For the instrument to be valid, it must not directly influence the labor share (LS_{jt}) except through its effect on plant-level quality proxy (q_{jt}) (**exclusion restriction**). The IV estimation relies on the identifying assumption that the instrument Z_{jt}, constructed using quality rating and segment dummies, is orthogonal to the error term such that: $Z_{jt} \perp \varepsilon_{jt}$. This assumption is plausible because confounding factors, such as an increase in pure markups unrelated to quality, would influence the proxy q_{jt} but are unlikely to affect Z_{jt}, as quality rating scores primarily reflect inherent quality differences.

Thus, the β estimate from the 2SLS regression more precisely captures the relationship between quality-related changes and labor share, effectively isolating the effect of product quality from other confounding factors influencing unit prices. In column (3), the IV estimate, though slightly larger than the OLS estimate in column (2), conveys a similar message. Overall, the estimates suggest that an increase in plant-level quality is significantly associated with a reduction in labor share. Interpreting this reduction in labor share as a change in the input mix, the estimates suggest that within plants, upgrading product quality involves shifting the input mix toward greater automation.[19]

3.2.2 Decomposition of labor share changes

I decompose the changes in labor share into the numerator (payroll) and denominator (sales) to discern how each component responds to quality upgrades, thus explaining the decline in labor share. For this purpose, I estimate Equation 3 with either the log of payroll or sales as the dependent variable. The estimates are presented in Table 3 alongside the estimates for the labor share itself. The second and third rows decompose the response of the labor

[18] In the appendix, I also report the results with IV based on the quality metric proposed by Khandelwal et al. [2013].

[19] If product quality and markups systematically co-move positively, interpreting the reduction in labor share as a change in the production function (input mix) requires more caution. This is because the reduction in labor share may partly stem from an increase in markups rather than from changes in the plant's input mix toward more automation.

While it seems reasonable, given the rise of robots in the vehicle sector, that a decline in labor share likely indicates changes in production function, it would be more precise to use either the labor cost share or a direct measure of automation intensity as the dependent variable. This robustness check will be conducted using the U.S. Census Bureau microdata, which offers the necessary variables.

share into the numerator and denominator, respectively.

The estimates indicate that much of the decline in the labor share due to quality upgrades comes from a larger increase in the denominator rather than a decline in the numerator. In fact, both the numerator and denominator increase, but the surge in sales relatively outweighs the increase in payroll, resulting in an overall decrease in the labor share. This pattern is pronounced with IV estimation, where the instrument is designed to capture the inherent quality component reflected in the unit price.

Table 3: Decomposition of the Labor Share Changes to Quality Upgrading

Dependent variable ↓	(1)	(2)	(3)
LS_{jt}	-0.024	-0.027	-0.032
	(0.005)	(0.007)	(0.010)
$log(Payroll_{jt})$	0.32	0.13	0.51
	(0.14)	(0.18)	(0.30)
$log(Sales_{jt})$	0.85	0.71	1.18
	(0.15)	(0.21)	(0.28)
Specification	OLS	OLS	2SLS
Instrument			JD Power
Sample Year	1993-2021	2006-2021	2006-2021
Observations	1,311	575	575

Notes: Each estimate represents the coefficient for the plant-level quality proxy q_{jt} from a separate regression. Standard errors clustered by plant in parentheses. Regression controls a set of fixed effects including year, state, company and plant. Observations are weighted by the sales share within a year (s_{jt}).

This finding aligns with the **productivity effect** of automation, a well-documented phenomenon in the literature such that firms adopting industrial robots experience significant sales expansions. These sales expansions are often so substantial that, despite the decrease in labor share as the firm's input mix shifts away from human labor (**displacement effect**), total payroll increases.

Results in Table 3 suggest that improvements in product quality can help explain the observed pattern of sales expansion and labor share decline. As plants increasingly rely on robots relative to labor to enhance product quality, they experience a decline in labor share. Nevertheless, the resulting sales expansion driven by improved product quality is so substantial that payroll increases, even as labor share declines.

In addition, this finding also relates to the establishment-level dynamics of labor share

documented in Kehrig and Vincent [2021] for the universe of U.S. manufacturing establishments. They found that establishments experiencing a decline in labor share also underwent significant growth in size within the U.S. manufacturing industries. Furthermore, the observation that these low labor share establishments enjoy a product price premium compared to their peers suggests that demand-side forces are responsible for these outcomes. However, the drivers behind the observed strong positive demand shift towards low labor share establishments are not fully elucidated and are instead left generally as "demand shocks" with some anecdotal examples in their study. In industries such as the vehicle industry, where consumers show rising demand for quality and production technology requires more automation for quality upgrades, the quality-driven automation framework provides a consistent explanation for the micro labor share dynamics.

3.2.3 Ruling out the adjustment costs scenario

In the vehicle industry, model changeovers refer to the transition of a manufacturing plant from producing one vehicle model to another.[20] This process involves extensive retooling of the plant to adjust production lines for a new variety. If adjustment costs are the primary determinant of firms' automation decisions, model changeovers provide an ideal opportunity to increase automation, as retooling already requires production line adjustments, minimizing additional downtime and costs associated with automation.

Hence, model changeovers provide a natural opportunity to examine whether firms increase robot use due to lower adjustment costs, independent of product quality upgrades, or if the quality level of the new model drives their automation decisions. To investigate this, I estimate the following regression equation:[21]

$$\triangle LS_{jt} = \alpha_t + \beta_1 D_{jt} + \beta_2 D_{jt} \triangle q_{jt} + \beta_3 (1 - D_{jt}) \triangle q_{jt} + \gamma X_{jt} + \varepsilon_{jt} \tag{4}$$

Here, the dependent variable is the one-year difference in labor share, and this change in labor share is used as a proxy for an increase in the degree of automation. $D_{jt} \in [0, 1]$ represents the degree of model changeovers, intended to capture the extent of plant retooling for plant j in year t. It is calculated as the sales share of a new variety at each plant in year

[20] To clarify, by "model changeovers," I'm specifically referring to the introduction of a new vehicle variety, such as when Tesla began producing the Model Y in 2020 at their Fremont plant in California. While "new model" may also indicate new versions (new model year) of existing models, the level of retooling required at the plant is typically greater when introducing a completely new variety. I focus on this scenario to identify significant plant retooling efforts.

[21] The equation is equivalent to $\triangle LS_{jt} = \alpha_t + \beta_1 D_{jt} + (\beta_2 - \beta_3) D_{jt} \triangle q_{jt} + \beta_3 \triangle q_{jt} + \gamma X_{jt} + \varepsilon_{jt}$. However, I'm using the specification in Equation 4 to highlight the estimate of β_3.

t. $\triangle q_{jt}$ denotes changes in plant-level quality. α_t is a set of year dummies. X_{jt} is a vector of additional controls, including changes in the sales shares of each vehicle type produced within plant j to account for differences in production processes by vehicle types. ε_{jt} is the error term.

Here, β_1 captures the impact of changeovers with no changes in plant-level quality ($\triangle q_{jt} = 0$). β_2 captures the impact of changes in plant-level quality with no changeovers. β_3, the coefficient of primary interest, measures the extent to which product quality explains changes in labor share during model changeovers. If adjustment costs are the primary factor determining robot adoption, firms are expected to increase robot use during plant retooling, regardless of changes in product quality. This scenario would imply insignificant estimate of β_3. However, if quality upgrading influences automation decisions during model changeovers, coefficients β_3 would be negative.

Table 4: Response of Labor Share to Model Changeovers

Dependent variable →	$\triangle LS_{jt}$
D_{jt}	0.007 (0.003)
$(1 - D_{jt})\triangle q_{jt}$	-0.020 (0.004)
$D_{jt}\triangle q_{jt}$	-0.052 (0.015)
Specification	GLS
Year FE	X
Observations	1,112

Notes: Standard errors are in parentheses. The sample covers the years 1993-2021. Observations are weighted using the average of sales shares (s_{jt}) over two years, $w_{jt} = 0.5(s_{jt} + s_{j,t-1})$.

Table 4 presents the estimates from Generalized Least Squares (GLS), addressing heteroskedasticity across panels. The estimates for β_1 and β_3 indicate that a decrease in labor share is associated with model changeovers only when there is an increase in plant-level product quality. Meanwhile, β_2 reaffirms the baseline findings in Table 2, showing that a decline in labor share is associated with improved plant-level quality even in the absence of model changeovers (i.e., quality changes among continuing models). Notably, the magnitudes of β_1 and β_2 suggest that the relationship between labor share decline and quality upgrades is particularly pronounced during model changeovers, a period when plants undertake more extensive updates to their production lines.

In sum, the main takeaway is that during model changeovers, plants experience a relative decline in labor share only if they introduce higher-quality products ($\beta_3 < 0$). While low

adjustment costs during model changeovers appear to amplify the degree of labor share changes, the fact that the relative decline in labor share occurs only when product quality improves underscores the critical role of product quality upgrades in driving automation decisions.

3.3 Rising consumer demand for quality

In the cross-sectional household consumption data, vehicles show strong non-homotheticity, quantified by a quality Engel curve slope[22] of approximately 0.94 (Bils and Klenow [2001]). This suggests that as a household's total nondurable consumption doubles, the unit price of vehicles it purchases nearly doubles as well. Such cross-sectional evidence points to a shift in consumer demand toward higher quality as the economy grows.[23]

I present evidence of this trend from U.S. vehicle market sales data, highlighting two key findings: First, among continuing models (available at both t and $t-1$), higher-quality vehicles experience faster market share growth. Second, newly introduced models tend to have systematically higher prices.

3.3.1 Market share grows faster for higher quality vehicles

First, I focus on continuing vehicle models (available at both t and $t-1$) and show that vehicles with higher quality experience faster growth in their market share over time. To demonstrate this, I estimate the following demand equation:

$$\triangle \ln s_{igt} = \alpha_{gt} + \beta_1 \ln \overline{p}^Z_{igt} + \beta_2 \triangle \ln p_{igt} + \varepsilon_{igt} \tag{5}$$

Here, the dependent variable is the one-year difference in the log market share of vehicle model i in vehicle type $g \in G = \{\text{Car, SUV, Pickup, Van}\}$ in year t. α_{gt} is a vehicle type-year fixed effects. p_{igt} is the unit price of a vehicle. $\overline{p}_{igt}$ is a proxy for the quality of the vehicle, and ε_{igt} is the error term.

[22] The elasticity of the unit price of a consumer's purchase with respect to the consumer's total nondurable consumption.

[23] This argument can be more clearly depicted with a model, similar to Oberfield [2023]. Suppose there are continuous quality levels $q \in R_+$ along the real line. Households' valuation of quality differs by their consumption level, with non-homothetic preferences for quality. Taking the supply of qualities as given, this generates a cross-sectional distribution of quality demanded by household income groups. Over time, as the economy grows (e.g., due to an increase in TFP), households will place more value on higher quality, causing demand to shift toward higher-quality levels. This empirical section aims to document evidence of this rightward shift.

$\overline{p}_{igt}$ is calculated using an approach similar to that in Bems and Di Giovanni [2016]. It is the average of each model's relative unit price within a vehicle type, $p_{igt}/\overline{p}_{gt}$ over two years.[24] When calculating the relative unit price, the denominator $\overline{p}_{gt}$ is the group's average unit price weighted by sales. $\overline{p}_{igt}$ is intended to capture the relative quality of model i among the continuing models (available at both time t and $t-1$) within vehicle type g. $\ln \overline{p}^{Z}_{igt}$ denotes the standardized Z-score of $\ln \overline{p}_{igt}$ within vehicle type-year. This standardization simplifies the interpretation of β_1, enabling comparisons across groups and years on a consistent scale.[25]

The main parameter of interest is β_1, which assesses whether higher-quality vehicle models, on average, experience faster market share growth. Meanwhile, the inclusion of $\triangle \ln p_{igt}$ controls for the influence of price changes on market share. Vehicle type-year fixed effects, α_{gt}, account for shocks common to each vehicle type-year, allowing the estimated coefficients to be identified from within vehicle type-year variations. Table 5 presents the estimation results of Equation 5. Columns (1) and (2) present the results without and with controls for changes in unit prices, $\triangle \ln p_{igt}$, respectively.

Table 5: Regression of Market Share Growth Rate on Product Quality

	(1)	(2)
Dependent variable $\rightarrow$	$\triangle \ln s_{igt}$	
$\ln \overline{p}^{Z}_{igt}$	0.015	0.017
	(0.005)	(0.006)
$\triangle \ln p_{igt}$	-	1.06
		(0.16)
consumption growth (%)	1.8	
Observations	6,799	
R^2	0.15	0.19

Notes: Standard errors in parentheses. Standard errors are clustered by vehicle model. Observations are weighted using the average of expenditure shares over two years, $w_{igt} = 0.5(s_{igt} + s_{ig,t-1})$. The sample covers years 1993-2023.

In column (1), the estimate of β_1 indicates that, on average, a vehicle model i with quality one standard deviation above the average experiences a relative market share growth rate approximately 1.5 percentage points higher. In column (2), even after controlling for changes

[24] $\overline{p}_{igt} \equiv 0.5 \left(\frac{p_{igt}}{\overline{p}_{gt}} + \frac{p_{ig,t-1}}{\overline{p}_{g,t-1}} \right)$.

[25] The standard deviation of $\ln \overline{p}_{igt}$ varies significantly across vehicle types: Car(0.56); SUV(0.39); Pickup(0.27); Van(0.13). Cars and SUVs have a greater variety of models, resulting in larger price dispersion. Pickups and vans, with a relatively smaller set of varieties, have significant quality differences across trims. However, the data source (Ward's Automotive Yearbook) does not provide breakdowns by trims regarding unit sales and prices. This leads to relatively smaller price dispersion in the sample. Due to the substantial differences across vehicle types, I present the results using a standardized variable.

in unit prices, the estimate for β_1 remains largely unchanged. These positive estimates of β_1 indicate that consumers have placed increasing value on higher quality products over time.

This estimate aligns with the prediction of the quality Engel curve in a growing economy. Over the sample years, overall consumption has grown at an average annual rate of approximately 1.8%. Along this growth path, the quality Engel curve predicts that as consumption increases, consumers place greater value on higher-quality products, shifting their demand toward them. Given the supply of available varieties, this shift translates into faster market share growth for models with relatively higher quality.

The positive estimate of β_2 in column (2) suggests that vehicle models with relatively larger increases in unit prices also experience relatively greater market share growth. This pattern likely arises because annual upgrades to vehicle models (e.g., 2020 Toyota Camry vs. 2021 Toyota Camry) often result in unit price increases that primarily reflect quality improvements.[26] This results in an increase in its market sales share despite the relative unit price increase. The β_2 estimate, together with the positive β_1, underscores the pivotal role of higher product quality in driving shifts in consumer demand.

3.3.2 Entering models tend to have systematically higher prices

If consumer demand is shifting toward higher quality over time as the economy's overall consumption grows, entering models (new varieties introduced at year t) are expected to feature higher quality to meet and capture this evolving demand. Therefore, I now focus on entering models and demonstrate that their prices are systematically higher compared to existing models. To show this, I estimate the following simple regression:

$$\ln p_{igt} = \alpha_{gt} + \beta \cdot Entry_{igt} + \varepsilon_{igt} \tag{6}$$

Here, the dependent variable is is the log of the unit price of vehicle model i in group g, and in year t. $Entry_{igt}$ is a binary variable indicating whether vehicle model i is a new variety introduced in year t. α_{gt} represents vehicle group-year effects, and ε_{igt} is the error term. The coefficient β captures the average price differential between entering and existing vehicle models. If entering models tend to be priced higher on average, we expect β to be positive and statistically significant.

[26] For a product i without changes in characteristics, $\triangle \ln p_{igt}$ indicates changes in pure prices. However, vehicle models have quality changes with new features each year as companies introduce new versions of each model annually. Therefore, for vehicles, $\triangle \ln p_{igt}$ reflects not only changes in pure prices but also quality changes.

Table 6: Average Price Differential Between Entering vs. Existing Models

	(1)	(2)
Dependent variable →	$\ln p_{igt}$	
$Entry_{igt}$	0.045	0.048
	(0.022)	(0.010)
Type×Year FE	✓	
Segment×Year FE		✓
Observations	7,212	7,199

Notes: Standard errors in parentheses. The sample covers years 1992-2023.

Table 6 presents the estimates of β. Column (1) suggests that, within a vehicle type and year, entering models have, on average, a 4.5% higher price compared to existing models. Column (2), which accounts for finer vehicle segment classifications within types, reveals that entering models are, on average, priced 4.8% higher than existing models within a vehicle segment and year. These findings suggest that firms introduce higher-quality models as new varieties over time to meet shifting consumer demand.

In sum, the past 30 years of U.S. vehicle market sales data reveal a significant demand shift toward higher-quality models through two main channels: (1) among continuing models, consumers increasingly shift their demand toward higher-quality models, and (2) new varieties introduced to the market are systematically of higher quality compared to existing models. This rising demand for quality creates strong incentives for firms to continuously enhance their product quality, as consumers value these higher-quality products and are willing to purchase more, even at relatively higher unit prices.

3.4 Evidence in U.S. Manufacturing Industries

Combining datasets that include product-level unit prices and establishment-level robot usage from the U.S. Census Bureau microdata, I present evidence that establishments producing higher-quality (higher unit price) products use robots more intensively than their peers within the same product category. To investigate this, I estimate the following regression equation for the establishment j producing product category c in year t:

$$a_{jt} = \lambda_t + \delta_c + \beta \tilde{p}_{jct} + \gamma X_{jt} + \varepsilon_{jt} \tag{7}$$

The dependent variable, a_{jt}, measures robot use at the establishment and is constructed

as robot density, defined as the number of industrial robots per 1,000 employees, or $\log(1 + Robot_{jt})$, where $Robot_{jt}$ indicates the number of industrial robots. This transformation allows the inclusion of samples with no robots at all. The regression includes year fixed effects (λ_t), and 10-digit product cateogry fix effects (δ_c). $\tilde{p}_{jct}$ represents the product price premium as defined in Equation 1. X_{jt} represents additional controls, with $\log(sales_{jt})$ included in column (3) to account for firm size differences. ε_{jt} denotes the error term. The parameter of interest is β, which captures compared to its peers producing the sampe product category, how establishment's produc price premium is associated with the degree of robot use.

Table 7: Regression of Robot Use Measure on Product Price Premium

	(1)	(2)	(3)
Dependent variable →	Robot density$_{j,t}$	$\log(1 + Robot_{j,t})$	$\log(1 + Robot_{j,t})$
$\tilde{p}_{jct}$	$+^{**}$	$+^{**}$	$+^{**}$
$\log(sales_{jt})$			$+^{**}$
10-digit product FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: $^{**}p < 0.01$. Robot density is defined as the number of industrial robots per 1,000 workers. In columns (2) and (3), $Robot_{i,t}$ represents the number of industrial robots in operation at establishment i in year t.

The parameter of interest is β, which captures how an establishment's product price premium is associated with the degree of robot use, compared to its peers producing the same product category. Table 7 reports the sign and significance of the estimates.[27] The results provide cross-sectional evidence that, compared to their peers within the same product category, establishments producing higher-priced (higher-quality) products exhibit more intensive robot use. Column (1) shows that these establishments have higher robot density, while columns (2) and (3) indicate that those producing higher-priced products use more robots, even after controlling for size. The results suggest that a firm's automation is associated with product quality across a broader set of U.S. manufacturing industries, extending beyond the automotive industry.

To summarize the empirical section, the findings demonstrate that quality-driven automation offers a consistent explanation for the shift in sales toward establishments adopting automation—a shift that cannot be fully explained by cost reduction alone. Building on these findings, I use the model to explore how incorporating the quality dimension enhances our

[27]Coefficient estimates and standard errors will be requested for disclosure going forward, along with additional estimates from longitudinal analyses.

understanding of the economic effects of rising automation.

4 Model

Using a model, I explore the implications of quality-driven automation in three aspects: (a) explaining the rise in automation, (b) assessing its welfare impacts across households, and (c) evaluating the welfare effects of regulating automation through a robot tax in the context of product quality growth. I highlight how my model provides distinct implications that are overlooked by previous models abstracting from product quality.

There are two sectors. The homogeneous (outside) good sector ($h = o$) operates under perfect competition and uses only labor in production. The homogeneous good serves as the numeraire. In the differentiated goods sector ($h = d$), for each quality level $q \in R_+$, there is a representative firm which supplies goods, producing with both robots and labor. Differentiated goods producers use a CES constant-returns-to-scale technology, but the share parameters for robots increase with product quality q relative to those for labor.[28]

Time is discrete, denoted by $t = 0, 1, 2, ...$ There is a set of exogenous aggregate state variables $S_t = (A_{o,t}, A_{d,t}, R_t, \tau_{k,t})$, which will be referred to as the *fundamentals*. Here, $A_{o,t}$ and $A_{d,t}$ represent labor-augmenting technical progress in each sector, R_t denotes the price of a robot, and $\tau_{k,t}$ indicates the size of the robot tax. The fundamentals of the economy are deterministic, with some variables varying over time. Similar to Boppart et al. [2023], the model considers a growing economy driven by increases in labor-augmenting technology.

Workers are characterized by their employment status, sector of work and efficiency units of labor. Individuals' valuation of quality varies with their overall consumption, leading to an increased demand for quality as their consumption grows. Workers are forward-looking and have perfect foresight, discounting the future at rate β. Given an initial distribution of workers across sectors, they choose which sector to move to, facing reallocation frictions such as idiosyncratic mobility costs and temporary unemployment spells, which make the reallocation costly. The sectoral reallocation decision is made based on the path of real wages over time and across sectors. This decision is closely related to the sectoral reallocation problem discussed in Beraja and Zorzi [2024].

Conditional on the supply of labor in a given sector, the model solves the static sub-problems of consumption and production to determine product prices and wages. A detailed

[28] In theory, this indicates that the marginal rate of technical substitution (MRTS) between robots and labor changes as the product quality level increases. In other words, the marginal product of robots relative to labor increases with product quality, resulting in changes in the isoquants.

description of each component of the model economy is provided below.

4.1 Workers

There is a continuum of workers characterized by their individual state $\mathbf{s} = (e, h, \ell)$. e denotes the employment status, where $e = \{E, U\}$, indicating whether an individual is employed E, or unemployed U. h represents the sector of employment, where $h = \{d, o\}$. ℓ represents the efficiency units of labor, and different endowments of ℓ generates an income distribution. $\mu_t(\mathbf{s})$ is the the probability measure indicating the cross-sectional distribution of workers at time t. At period $t = 0$, there is a mass $\mu_0(\mathbf{s})$ of workers at each state $\mathbf{s}$. $\mu_t(\mathbf{s})$ evolves according to the workers' dynamic sector reallocation given the fundamentals such that $\mu' = \mathbb{T}(S, \mu)$. $\mathbb{T}$ denotes a transition operator that define the law of motion for the distribution of workers.

When employed, workers supply a unit of labor inelastically, and earn wage income proportional to their ability ℓ. They face sector-specific competitive market wages, denoted as $W_{h,t}$, which represents the wage per unit of effective labor paid in sector h at time t. A worker with ability ℓ working in sector h earns labor income $W_{h,t}\ell$.

Workers may reallocate to another sector with some reallocation friction. I consider this reallocation primarily to capture the movement of workers from the differentiated goods sector ($h = d$), where robots increasingly displace workers as automation progresses, to the outside sector ($h = o$). To reallocate from sector h to h', individuals experience a spell of unemployment ($e = U$), during which they can exit to employment with a certain hazard rate. While unemployed, they receive unemployment benefits $bW_{h,t}\ell$, where b denotes the replacement ratio applied to the wage they would have earned in their previous sector of employment h. When characterizing the static subproblems, the time subscript is omitted for brevity.

Static consumption problem. Utility is derived from the consumption aggregator, $\mathcal{C}$. In each period, a worker of state $\mathbf{s} = (e, h, \ell)$ solves the static consumption problem.

$$v(\mathbf{s}; S, \mu) = \max_{\{x_j\}_{j \in J}, c_o} \log(\mathcal{C}) \tag{8}$$

$$\mathcal{C} = (c_d)^{\alpha}(c_o)^{1-\alpha} \tag{9}$$

$$c_d = \left[\int_J \Phi(q_j, \mathcal{C})^{\frac{1}{\eta}} x_j^{\frac{\eta-1}{\eta}} dj \right]^{\frac{\eta}{\eta-1}} \tag{10}$$

subject to the following budget constraint

$$c_o + \int_J p_j x_j dj \leq X(y(\mathbf{s})) \tag{11}$$

$$y(\mathbf{s}) = \begin{cases} W_h \ell & \text{for employed} \quad (e = E) \\ bW_h \ell & \text{for unemployed} \quad (e = U) \end{cases} \tag{12}$$

where $X(y(\mathbf{s})) = y(\mathbf{s}) - T(y(\mathbf{s})) + \zeta(\mathbf{s})$ denotes the disposable income after income tax $T(y(\mathbf{s}))$ and transfer $\zeta(\mathbf{s})$. The income tax schedule is linear with a tax rate τ_y, such that $T(y) = \tau_y y$. A tax system that includes a transfer financed by a proportional income tax governs the redistribution of income from pre-tax to post-tax levels.[29]

The consumption aggregator $\mathcal{C}$ is a Cobb-Douglas function defined over the consumption of differentiated goods c_d and the consumption of a homogeneous good c_o. The homogeneous good is the numeraire and its price is normalized to 1. With this Cobb-Douglas setting, individuals allocate a fraction $(1 - \alpha)$ of their expenditure to the numeraire good. By making this assumption, the model focuses on the within-sector expenditure shifts across different quality levels, while abstracting from shifts across sectors.

The remaining α fraction of expenditure is spent on c_d. c_d is itself a non-homothetic CES aggregator over varieties $j \in J$, each with quality q_j. η is the elasticity of substitution between varieties. In aggregation of c_d, individuals with larger consumption place greater value on higher-quality varieties. This is captured by $\Phi(q, \mathcal{C})$, which is a demand shifter that reflects how valuations of quality vary with household consumption levels.[30]

Building on Oberfield [2023] and Fieler et al. [2018], I assume

$$\Phi(q, \mathcal{C}) = \overline{\phi}(q) \cdot \phi\left(q - \nu \ln \mathcal{C}\right) \tag{13}$$

where $\phi\left(q - \nu \ln \mathcal{C}\right)$ is a weighting function that reflects varying valuations of quality based on an individual's total consumption. The weighting function $\phi(\cdot)$ is assumed to follow a standard normal distribution $N(0, \sigma_\phi^2)$. It indicates that for individuals with total consumption $\mathcal{C}$, the weighting function is centered at $q = \nu \ln \mathcal{C}$, with a standard deviation of σ_ϕ.

[29] I do not assume non-linear taxation in the spirit of Heathcote et al. [2017], such as $T(y) = \psi_0 y^{1-\psi_1}$, because in a growing economy, it would require the parameters ψ_0, ψ_0 to adjust over time to maintain a reasonable effective income tax rate, thereby introducing additional complexity to the estimation process.

[30] Compared to other non-homothetic preferences that restrict consumers to one unit of each differentiated good (e.g., Fajgelbaum et al. [2011] and Bils and Klenow [2001]), this framework has the advantage of simultaneously capturing both growth in the quantity of quality-differentiated goods and quality growth.

As overall consumption $\mathcal{C}$ rises, more weight is put on goods with higher quality q. The parameter $\nu > 0$ governs the strength of the non-homotheticity. $\overline{\phi}(q)$ captures the general relationship that higher quality leads to higher valuations across all consumers.

An individual's relative demand for any two varieties j and k with $q_j > q_k$,

$$\frac{x_j}{x_k} = \left(\frac{p_j}{p_k}\right)^{-\eta} \frac{\Phi(q_j, \mathcal{C})}{\Phi(q_k, \mathcal{C})} \tag{14}$$

is increasing in $\mathcal{C}$. Relative expenditures by consumers of total consumption $\mathcal{C}$ are given by

$$log\frac{\varphi_j}{\varphi_k} = log\frac{\Phi(q_j, \mathcal{C})}{\Phi(q_k, \mathcal{C})} - (\eta - 1)log\frac{p_j}{p_k} \tag{15}$$

The higher income consumer (with higher consumption $\mathcal{C}$) demands relatively more higher-quality varieties. This cross-sectional variation in consumption baskets across income groups can be mapped to the qualtiy Engel curve slope[31] estimated in Bils and Klenow [2001]. Assuming that prices increase with quality, given the product price schedule $\{p_j\}_{j \in J}$, the parameter ν determines the slope of the quality Engel curve in the model. A larger value of ν results in a steeper quality Engel curve.

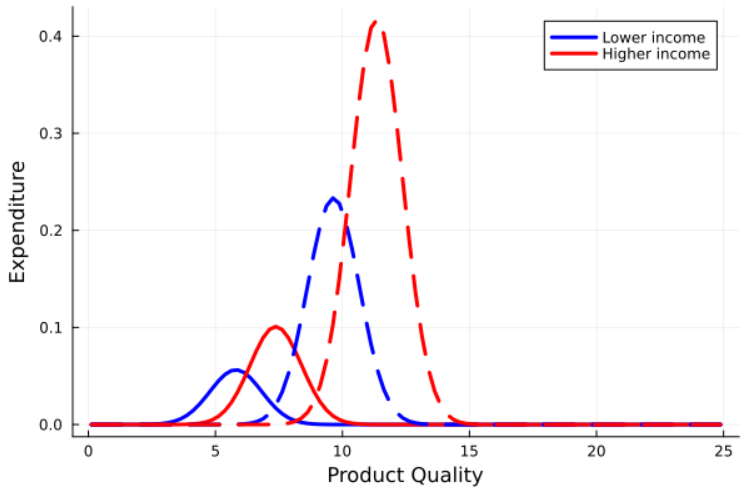


Figure 2: Expenditure Shift with Consumption Growth

Notes: The solid line illustrates the initial state, while the dashed line indicates the state after economic growth. Each color represents simplified income groups: one with lower income and the other with higher income. The figure illustrates not only the cross-sectional dispersion in quality demanded by different income groups but also shows that as overall consumption increases over time, expenditure shifts toward higher-quality products.

[31] The elasticity of the unit price of a consumer's purchase with respect to the consumer's total nondurable consumption.

A strand of recent papers primarily uses non-homothetic preferences for quality to explore the implications of cross-sectional income differences. For instance, studies have analyzed differences in inflation rates across income groups using the Nielsen dataset, which provides detailed information on unit prices and quantities for grocery products (e.g., Faber and Fally [2022]; Nord [2022]; Argente and Lee [2021]).

In contrast, I incorporate this preference within the context of economic growth, where income growth over time also leads to quality growth. As overall consumption grows with economic growth, consumers place increasing value on higher quality. Consequently, as illustrated in Figure 2, expenditure shifts toward the higher-quality products, given the price schedule $\{p_j\}_{j\in J}$.[32] This preference structure forms the foundation for analyzing an economy in which product quality evolves over time. It is closely related to the empirical findings in Bils and Klenow [2001] and the theoretical framework developed in Oberfield [2023].[33]

Dynamic sector reallocation problem. In the context of rising automation, the model focuses on the gradual reallocation of workers initially employed in sector d, where labor displacement by robots occurs, to the outside sector o. The timing for the worker's dynamic problem is as follows. If they begin the period as employed in a sector h, they work and earn the market wage. At the end of each period, workers observe the realization of their idiosyncratic migration cost $\delta_{m,t}$ and have the option to relocate to another sector.

Assumption 1 *The idiosyncratic migration cost δ_m is i.i.d. over time, following an exponential distribution with rate parameter ξ_m.*

If they decide to relocate to another sector, they experience a spell of unemployment with an exit hazard κ. The value function of a worker with individual state $\mathbf{s} = (e, h, \ell)$ at time t is expressed as

$$V_t(\mathbf{s}; S, \mu) = v(\mathbf{s}; S, \mu) + \beta \hat{V}_{t+1}(\mathbf{s}; S', \mu') \tag{16}$$

Here, $v(\mathbf{s}; S, \mu)$ represents the instantaneous utility, which is derived from solving the static consumption problem in Equation 8. Similar to the setup in Beraja and Zorzi [2024], $\hat{V}_{t+1}(\mathbf{s}; S, \mu)$ is the continuation value associated with the discrete sector choice for a worker with individual state $\mathbf{s} = (e, h, \ell)$ at period t. Formally, for workers employed in sector h at

[32] In the data, the market share shift toward higher-quality products reflects the dynamics of variety entry and exit. Rather than explicitly modeling this process, my model assumes a continuum of varieties spanning all quality levels. The product space evolves as expenditure reallocates across varieties of differing quality. This approach facilitates the long-run analysis of a growing economy.

[33] Oberfield [2023] uses the term "rank" to describe a concept analogous to quality in my paper. Higher-ranked goods have higher prices, yet consumers shift toward these goods as their consumption increases.

period t, their continuation value is

$$\hat{V}_{t+1}(E,h,\ell;S',\mu') = E\left[max\{V_{t+1}(E,h,\ell;S',\mu'), V_{t+1}(U,h,\ell;S',\mu') - \delta_m\}\right] \tag{17}$$

The continuation value reflects the worker's option value of either remaining in sector d or transitioning to temporary unemployment stage, followed by relocation to another sector h' after incurring the migration cost. The expectation is taken over the distribution of idiosyncratic migration cost shocks. Meanwhile, the continuation value for workers undergoing a temporary unemployment spell during the transition from sector h to h' is expressed as

$$\hat{V}_{t+1}(U,h,\ell;S',\mu') = (1-\kappa)V_{t+1}(U,h,\ell;S',\mu') + \kappa V_{t+1}(E,h',\ell;S',\mu') \tag{18}$$

This continuation value reflects the fact that unemployed workers can exit the unemployment stage and begin working in the destination sector h' with an exogenous hazard rate, κ. The probability of staying in sector h at the end of period t, denoted as $p_t^{h,h}$, becomes

$$\begin{aligned} p_t^{h,h}(\ell) &= Pr(\delta_{m,t} \geq V_{t+1}(U,h,\ell;S',\mu') - V_{t+1}(E,h,\ell;S',\mu')) \\ &= e^{-\xi_m[V_{t+1}(U,h,\ell;S',\mu') - V_{t+1}(E,h,\ell;S',\mu')]} \end{aligned} \tag{19}$$

It reflects that workers relocate to the sector with higher lifetime utility once the gap is large enough to compensate for the idiosyncratic mobility cost and earnings loss during the unemployment spell. Then, the law of motion for the mass of workers of individual state $\mathbf{s} = (e,h,\ell)$ is given by

$$\mu_{t+1}(E,d,\ell) = p_t^{d,d}(\ell)\cdot\mu_t(E,d,\ell) + \kappa\cdot\mu_t(U,o,\ell)$$

$$\mu_{t+1}(E,o,\ell) = p_t^{o,o}(\ell)\cdot\mu_t(E,o,\ell) + \kappa\cdot\mu_t(U,d,\ell)$$

$$\mu_{t+1}(U,d,\ell) = (1-\kappa)\cdot\mu_t(U,d,\ell) + (1-p_t^{d,d}(\ell))\cdot\mu_t(E,d,\ell)$$

$$\mu_{t+1}(U,o,\ell) = (1-\kappa)\cdot\mu_t(U,o,\ell) + (1-p_t^{o,o}(\ell))\cdot\mu_t(E,o,\ell)$$

The evolution of $\mu_{t+1}(E,h,\ell)$, mass of workers employed in sector h, reflects both the fraction of stayers in sector h, and the inflow of migrants who originated from another sector. These migrants have gone through an unemployment spell and exited it with an exogenous hazard κ. Meanwhile, the evolution of $\mu_{t+1}(U,h,\ell)$, mass of workers in a temporary unemployment spell during relocation from the origin sector h to the destination h', reflects both the fraction of unemployed workers who have not exited the spell and the inflow of new

migrants who have left the origin sector.

Together, these equations define the law of motion $\mu' = \mathbb{T}(S, \mu)$. Given the timing assumption, the supply of labor in each period is fully determined by the forward-looking decisions made at the end of the previous period. Conditional on the labor supplied in each sector, the static production problem can be solved for equilibrium prices at each time t such that markets clear.

4.2 Production

4.2.1 Homogeneous good producers

The representative firm producing the homogeneous good, which serves as the numeraire, uses labor as its input.

$$Y_{o,t} = A_{o,t} L_{o,t} \tag{20}$$

$A_{o,t}$ is the labor-augmenting technology. $A_{o,t}$ increases over time as a result of exogenous technical progress, $A_{o,t} = \tilde{A}_o \cdot e^{g_A \cdot t}$. This homogeneous good is competitively priced. It can be used for consumption or investment purposes (i.e., the purchase of robots). The wage per unit of effective labor in sector o is given by $W_{o,t} = A_{o,t}$.

4.2.2 Robot producers

Similar to Guerreiro et al. [2022], robots are produced by competitive firms at a cost of R_t per robot. The production function for robots is given by $K_t = \frac{1}{R_t} X_{o,t}$, where $X_{o,t}$ denotes the units of the numeraire good used. The producer chooses robot supply, K_t to maximize profits: $p_{k,t} K_t - (1 + \tau_{k,t}) R_t K_t$. $\tau_{k,t}$ represents the robot tax levied by the government for the purpose of regulating automation. In equilibrium, the robot price $p_{k,t} = (1 + \tau_{k,t}) R_t$ and profits are zero.[34] A decline in R_t represents relative technological progress in robots compared to labor-augmenting technologies, which can be interpreted as a reduction in the effective price of robots.

4.2.3 Differentiated goods producers

For each quality level $q \in R_+$, there exists a representative firm.[35] At each quality level q, the representative firm produces output (in physical units) by combining the services of a

[34] This is equivalent to the economy importing robots from abroad by paying the price $p_{k,t}$.

[35] This assumption abstracts from firms' choices of quality. For clearer boundaries of a firm, I could instead model heterogeneous firms selecting quality under monopolistic competition, as in Demir et al. [2024] and Fieler et al. [2018]. However, I assume representative firms at each quality level, as this simplification facilitates the analysis of long-run trends in rising automation within a growing (non-stationary) economy.

range of tasks. Indexing each representative firm by variety j, the production function of firm j, with quality level $q_j = q$ is given by:[36]

$$Y(q) = z(q) \left(\int_0^1 y_q(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} \tag{21}$$

where σ is the elasticity of substitution between tasks. Building on the task-based framework in the automation literature (e.g., Acemoglu and Restrepo [2018]), I model the production function to become more robot-intensive as product quality advances. The key difference from previous models is that the productivity schedule over tasks shifts, allowing robots to have a comparative advantage in a larger set of tasks at higher product quality levels.

There are two types of technologies that firms can use to produce each task $y_q(i)$. Without automation technology, the production function takes the form

$$y_q(i) = \psi_q^l(i) L(i) \tag{22}$$

where $\psi_q^l(i)$ denotes the productivity of labor in task i. In contrast, with automation technology, task i can be produced by either labor or robots:

$$y_q(i) = \psi_q^k(i) K(i) + \psi_q^l(i) L(i) \tag{23}$$

where $\psi_q^k(i)$ is the productivity of robots in task i. In technologically automated tasks, robots and labor are prefect substitutes. In line with the literature, I assume that labor has a comparative advantage in higher-indexed tasks. That is, $\psi_q^l(i)/\psi_q^k(i)$ is strictly increasing in i. Because labor has a strict comparative advantages in tasks with higher index, there is a threshold ${I_q}^*$ such that

$$\frac{W_d}{R} = \frac{\psi_q^l(I_q^*)}{\psi_q^k(I_q^*)} \tag{24}$$

where W_d represents the wage in sector d, and R represents the price of a robot, which becomes $(1 + \tau_k)R$ with the robot tax τ_k. This implies that there exists a unique optimal threshold task such that all tasks $i \leq I_q^*$ will be produced with robots, while all tasks $i \geq I_q^*$ will be produced with labor. I model firm's automation decision as a choice of I_q (i.e., the fraction of tasks executed by robots).

[36] Each variety j maps one-to-one with each quality q, making the index j interchangeable with q.

Then, the production function can be rewritten as

$$Y(q) = z(q) \left(\int_0^{I_q} (\psi_q^k(i) K(i))^{\frac{\sigma-1}{\sigma}} di + \int_{I_q}^1 (\psi_q^l(i) L(i))^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} \tag{25}$$

Given factor prices $\{R, W_d\}$, the degree of automation I_q depends on the productivity schedules of robots and labor across tasks in $[0, 1]$. To reflect greater automation with higher product quality, I assume that the schedule of comparative advantage $\psi_q^l(i)/\psi_q^k(i)$ over tasks varies across output quality $q \in R_+$. If $\psi_q^l(i)/\psi_q^k(i) < \frac{W_d}{R}$ for a larger share of tasks as q increases, the firm chooses to automate a larger set of tasks.

In the differentiated goods sector ($h = d$), product quality upgrading requires more intricate and refined production processes. Consequently, it becomes more challenging to execute these processes with human labor. Because robots exhibit greater accuracy and consistency, the comparative advantage shifts toward robots as the required output quality level q increases. To make the analysis more tractable, I introduce the following assumption.

Assumption 2 $\psi_q^l(i) = A_d e^{i-\lambda q}$ *and* $\psi_q^k(i) = 1\{i \in [0, \overline{a}]\}$.

where $1\{\cdot\}$ is a indicator function. This assumption implies that tasks $i > \overline{a}$ are technologically non-automatable, reflecting the upper bound of task automation. Within this task-based framework, given the factor prices, only the relative comparative advantage $\psi_q^l(i)/\psi_q^k(i)$ matters for determining the threshold task for automation. Therefore, I normalize $\psi_q^k(i) = 1$ for automatable tasks.

The function $\psi_q^l(i)$ reflects that, as product quality advances, human labor becomes less efficient, with the extent of this decline determined by the parameter $\lambda > 0$.[37] A_d represents exogenous labor-augmenting technical progress, which is common in the sector d and evolves according to $A_{d,t} = \tilde{A}_d \cdot e^{g_A \cdot t}$.

Given the productivity schedule for using robots versus labor, the threshold task for automation is given by $I_q^* = min\{\lambda q + log(\frac{W_d}{A_d R}), \overline{a}\}$. Here, the degree of automation is determined not only by relative factor prices in $\frac{W_d}{A_d R}$ but also by the term $\lambda q > 0$, which reflects changes in the marginal rate of technical substitution (MRTS) between robots and labor as product quality increases. The production function becomes (refer to subsection A.5

[37] While this assumption is drawn from evidence in the empirical section, it also relates to the cross-country evidence presented in Schott [2004], which shows that U.S. import data exhibits a positive association, within industries, between the unit value (price) of products and the exporter's capital intensity (capital per labor) in production.

for a detailed derivation)

$$
\begin{aligned}
Y(q) &= z(q)\left[\left(\int_0^{I_q}\psi_q^k(i)^{\sigma-1}di\right)^{\frac{1}{\sigma}}K_q^{\frac{\sigma-1}{\sigma}}+\left(\int_{I_q}^{1}\psi_q^l(i)^{\sigma-1}di\right)^{\frac{1}{\sigma}}L_q^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \\
&= z(q)\left[\Gamma^k(I_q)^{\frac{1}{\sigma}}K_q^{\frac{\sigma-1}{\sigma}}+\Gamma^l(I_q)^{\frac{1}{\sigma}}(A_dL_q)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}
\end{aligned}
\tag{26}
$$

Share parameters $\Gamma^k(I_q)=I_q$ and $\Gamma^l(I_q)=\left(\int_{I_q}^{1}(e^{i-\lambda q})^{\sigma-1}di\right)$ are endogenous and depend on the degree of automation, which is determined by the threshold task index $I_q=I_q^*$. I denote the production function for product quality q as $F_q(L_q,K_q)$.

Marginal cost and factor shares. The marginal cost of producing a product of quality $q\in R_+$ is $\frac{C(q,I_q)}{z(q)}$, where the function $C(q,I_q)$ is

$$
C(q,I_q)=\left[\Gamma^k(I_q)R^{1-\sigma}+\Gamma^l(I_q)\left(\frac{W_d}{A_d}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{27}
$$

$z(q)$ represents the economy's Hicks-neutral productivity for each quality level q. The representative firm at each quality level q sets price at its marginal cost $p(q)=C(q,I_q)/z(q)$. Together with $C(q,I_q)$, $z(q)$ determines the price schedule across quality levels $\{p(q)\}_{q\in R+}$. To reflect the idea that higher-quality goods are, on average, more costly to produce, I assume that $z(q)$ decreases with q, leading to $p(q)$ increasing with q.

The share of labor in cost $s^l(q)$ is given by

$$
s^l(q)=\frac{\Gamma^l(I_q)\left(\frac{W_d}{A_d}\right)^{1-\sigma}}{\Gamma^k(I_q)R^{1-\sigma}+\Gamma^l(I_q)\left(\frac{W_d}{A_d}\right)^{1-\sigma}}=\frac{1}{1+\frac{\Gamma^k(I_q)}{\Gamma^l(I_q)}\left(\frac{R}{W_d/A_d}\right)^{1-\sigma}} \tag{28}
$$

Given the schedule of comparative advantage and the factor prices, the ratio of share parameters $\frac{\Gamma^k(I_q)}{\Gamma^l(I_q)}$ increases in q and $s^l(q)$ decreases accordingly. This decline in the labor share for higher-quality products arises from changes in the share parameters, which reflect changes in the isoquants.

Revenue and price index. The revenue of the representative firm at product quality level q, charging price $p(q)$, is

$$
r(q)=p(q)^{1-\eta}D(q) \tag{29}
$$

where $D(q)$ is the economy's demand shifter for product quality q, integrated over different households with state $\mathbf{s} = (e, h, \ell)$.[38]

$$D(q) = \int \Phi(q, \mathcal{C}(\mathbf{s})) P_d(\mathbf{s})^{\eta-1} \alpha X(\mathbf{s}) d\mu \tag{30}$$

$$P_d(\mathbf{s}) = \left[\int_J p_j^{1-\eta} \Phi(q_j, \mathcal{C}(\mathbf{s})) dj \right]^{\frac{1}{1-\eta}} \tag{31}$$

In deriving $D(q)$, the spending of each consumer type $X(\mathbf{s})$ is allocated to each quality level q, based on a demand shifter $\Phi(q, \mathcal{C}(\mathbf{s}))$ and a price index $P_d(\mathbf{s})$, both of which are specific to the income group.[39] Then, the overall price index associated with the consumption aggregator $\mathcal{C}$ for the household of type $\mathbf{s}$ is given by

$$\mathcal{P}(\mathbf{s}) = \left(\frac{P_d(\mathbf{s})}{\alpha} \right)^{\alpha} \left(\frac{1}{1-\alpha} \right)^{1-\alpha} \tag{32}$$

It economically indicates the cost, in terms of the numeraire, required to obtain one unit of utility for worker of type $\mathbf{s}$. The unit price in the data corresponds to the product price p_j in the model, while the price index is utility-based, reflecting consumer preferences. Henceforth, I will explicitly use the term "price index" whenever it is necessary to distinguish it from the product price.

4.3 Government

The government collects taxes to fund government purchases, unemployment benefits, and transfers with revenues generated from proportional income taxes and a robot tax. The government is subject to the following budget constraints:

[38] This demand shifter is analogous to those used in firm-to-firm trade models in Fieler et al. [2018] and Demir et al. [2024]. I adapt their setup to consumer demand in order to generate rising consumer demand for quality as consumption increases.

[39] Note that firm j, a representative firm at quality level $q_j = q$, faces the downward sloping demand curve given by $p(q)^{-\eta} D(q)$. However, it does not charge a markup, and thus no profits accrue, as I assume perfect competition at every quality level. This setup is considered for two reasons: (1) to facilitate solving for the transition path with dynamic worker reallocation in a growing economy, and (2) to avoid complications from product market power in one sector, which would complicate the implications of a policy experiment involving a robot tax, as a social planner would instead prefer subsidizing goods production to enhance consumer welfare.

$$G_t + \underbrace{\int_{e=U} y(\mathbf{s})d\mu}_{\text{unemployment benefits}} + \underbrace{\int \zeta(\mathbf{s})d\mu}_{\text{transfer}} \leq \underbrace{\int T(y(\mathbf{s}))d\mu}_{\text{income tax}} + \underbrace{\tau_k R\left(\int_J K_j dj\right)}_{\text{robot tax}} \tag{33}$$

In the baseline laissez-faire model, where no robot tax is imposed, I assume that transfers are zero and that all tax revenue, aside from expenditures on unemployment benefits, is allocated to government purchases, G_t, which do not directly enter the worker's utility function. When analyzing tax reform in the policy experiment, I assume that government purchases remain fixed at the laissez-faire level in each period, with a balanced budget maintained through adjustments in transfers. I consider two types of transfers: (1) transfers proportional to tax payments and (2) lump-sum transfers, each of which differs in its redistributive effects.

4.4 Equilibrium

The endogenous state of the economy is the distribution of workers $\mu_t(\mathbf{s})$. The fundamentals of the economy $S_t = (A_{o,t}, A_{d,t}, R_t, \tau_{k,t})$ are deterministic. Given $(\mu_0, \{S_t\}_{t=0}^{\infty})$, a *sequential competitive equilibrium* of the model consists of a sequence of worker distribution $\{\mu_t(\mathbf{s})\}_{t=0}^{\infty}$, value function, $\{V_t(\mathbf{s})\}_{t=0}^{\infty}$, decision rules for consumption $\{c_{o,t}(\mathbf{s}), \{x_{j,t}(\mathbf{s})\}_{j\in J}\}_{t=0}^{\infty}$, firm inputs $\{L_{o,t}, \{L_{j,t}, K_{j,t}\}_{j\in J}\}_{t=0}^{\infty}$, product prices $\{\{p_{j,t}\}_{j\in J}\}_{t=0}^{\infty}$, factor prices $\{W_{o,t}, W_{d,t}, R_t\}_{t=0}^{\infty}$, and a law of motion for the distribution $\mu' = \mathbb{T}(S, \mu)$ such that for $\forall t$:

1. Individual optimization: given wages $\{W_{o,t}, W_{d,t}\}$, the individual decision rules $c_{o,t}(\mathbf{s})$, and $\{x_{j,t}(\mathbf{s})\}_{j\in J}$ solve $V_t(\mathbf{s})$

2. Firms maximize profits:

$$W_{o,t} = A_{o,t}$$

$$W_{d,t} = p_{j,t} \cdot \frac{\partial F_{q_j}(L_{j,t}, K_{j,t})}{\partial L_{j,t}}, \quad \text{for} \quad \forall j \in J$$

$$R_t = p_{j,t} \cdot \frac{\partial F_{q_j}(L_{j,t}, K_{j,t})}{\partial K_{j,t}}, \quad \text{for} \quad \forall j \in J$$

3. Goods markets clear:

$$Y_{o,t} = \int c_{o,t}(\mathbf{s})d\mu + G_t + R_t\left(\int_J K_{j,t}dj\right)$$

$$Y_{j,t} = \int x_{j,t}(\mathbf{s})d\mu, \quad \text{for} \quad \forall j \in J$$

4. Labor markets clear:

$$L_{o,t} = \int_{h=o} \ell(\mathbf{s}) d\mu$$

$$\int_J L_{j,t} dj = \int_{h=d} \ell(\mathbf{s}) d\mu$$

5. The government balances the budget:

$$G_t + \int_{e=U} y_t(\mathbf{s}) d\mu + \int \zeta_t(\mathbf{s}) d\mu = \int T(y_t(\mathbf{s})) d\mu + \tau_{k,t} R_t \left(\int_J K_{j,t} dj \right)$$

In each period, solving for equilibrium reduces to finding the wage $W_{d,t}$ that clears the labor market, which in turn determines all product prices $\{p_{j,t}\}_{j \in J}$. A stationary equilibrium in this economy is a state in which no aggregate variables (S, μ) change over time. This implies that fundamentals S_t remain constant, and the distribution of workers remains unchanged.

5 Quantitative Evaluation

I now proceed to explain how I apply this model to the data and simulate the economy. The differentiated goods sector in the model is specifically mapped to the U.S. vehicle manufacturing industry, as this is the sector where I have provided evidence of quality-driven automation. Consequently, the quantification of welfare effects should be interpreted within the context of this assumption.[40] To place this model analysis in the context of the data, I compare the simulated results to long-run trends in key data moments from the U.S. vehicle manufacturing industry over the past 40 years, starting in 1980, when the prominent trend of declining labor share began.

Each model period corresponds to one year in the data. Changes in fundamentals lead to an overall increase in consumption, driving a shift in consumer demand toward higher-quality goods, while firms increasingly use robots to produce these higher-quality products. The model aims to quantify how much quality growth can account for rising automation and assess the welfare implications of decelerating automation through a robot tax in a growing economy where the product space continues to evolve.

[40] The range of industries subject to quality-driven automation is an empirical question and may extend to a broader set of industries. In that case, the model can readily accommodate this by adjusting the parameter α, which represents the expenditure share of products subject to quality-driven automation, and the initial mass of workers in this broader set of industries to reflect greater economic impacts.

5.1 Calibration

Parameterization. Workers' ability ℓ follows a log-normal distribution, with $\log(\ell) \sim N(-\frac{\sigma_\ell^2}{2}, \sigma_\ell^2)$. Idiosyncratic migration cost δ_m follows a exponential distribution with rate parameter ξ_m, and its cumulative distribution function is given by $G(x) = 1 - e^{-\xi_m x}$. The economy's Hicks-neutral productivity for each quality level q, denoted by $z(q)$, is specified as $z(q) = 1/q$. This specification ensures that higher-quality products are associated with higher prices, as marginal costs increase with q.[41] Similar to Fieler et al. [2018], I set $\overline{\phi}(q) = z(q)^{1-\eta}$ so that in the consumer expenditure share, changes in productivity $z(q)$ are offset by changes in demand $\overline{\phi}(q)$. While the model's results remain robust as long as $\overline{\phi}(q)$ is increasing in q, this assumption ensures a more stable rate at which expenditure shares shift toward higher quality as total consumption $\mathcal{C}$ increases. This is because, under this assumption, the relative expenditure share in Equation 15 is not affected by the schedules of $\overline{\phi}(q)$ and $z(q)$, making it more directly governed by the weighting function $\phi(q - \nu \ln \mathcal{C})$.

Table 8: List of Parameters

	Parameter	Value	Target / Source
Externally calibrated			
Task substitution elasticity	σ	0.6	Oberfield and Raval [2021]
S.D. of worker ability	σ_ℓ	0.5	Battistin et al. [2009]
replacement rate	b	0.4	UI replacement rate
Discount factor	β	0.96	yearly interest rate 4%
Demand substitution elasticity	η	2	Faber and Fally [2022]
Expenditure share of c_d	α	0.037	CPI weight for new vehicles
S.D. of weighting function $\phi(\cdot)$	σ_ϕ	1	normalization
Unemployment exit hazard	κ	0.53	earnings loss from reallocation
Internally calibrated			
Preference nonhomotheticity	ν	4	β_1 estimate in Equation 5
Comparative advantage schedule	λ	0.055	semi-elasticity of LS to q
Migration cost distribution	ξ_m	0.004	relative wage in 1980-2020
			employment share in 1980-2020
Proportional income tax rate	τ_y	0.15	$G/C = 0.19$

External calibration. I assign eight parameters $\bar{\Theta} = (\sigma, \sigma_\ell, b, \beta, \eta, \alpha, \sigma_\phi, \kappa)$. Task substitution elasticity, σ, is set to the plant-level elasticity of substitution between capital and labor for the transportation equipment industry, as estimated in Oberfield and Raval

[41] The model's results remain robust under more general forms of $z(q)$, as long as it is decreasing in q. A key requirement for the model is the existence of a well-defined productivity frontier by quality, which defines the production possibilities frontier (PPF) of the economy. This structure allows product prices $\{p_{j,t}\}_{j \in J}$ corresponding to each quality q_j, to be expressed as a function of the equilibrium wage rate $W_{d,t}$.

[2021] using U.S. establishment-level microdata. The standard deviation of worker ability, σ_ℓ, determines the distribution of income and consumption expenditure in the model. It is set to the estimate in Battistin et al. [2009] for the empirical distribution of U.S. household consumption expenditure, based on the U.S. Consumer Expenditure (CEX) Interview Survey. The replacement rate b represents the unemployment insurance replacement rate and is set to align with the average across U.S. states, approximately 40% of previous income. The discount factor β is set to 0.96, implying an annual interest rate of approximately 4%. The elasticity of substitution between varieties, η, is set to 2, consistent with the estimates in Faber and Fally [2022], which are derived from the Nielsen dataset containing product price information.

α represents the consumer expenditure share for differentiated goods in the model, where quality-driven automation takes place. This parameter is crucial for determining the size of welfare effects, as it reflects the importance of categories in consumption that are exposed to quality-driven automation. I set this value to the U.S. BLS CPI weight for new vehicles, the category where I provide evidence of quality-driven automation. Therefore, the model results should be interpreted as a lower bound on the effects of quality-driven automation, as this phenomenon likely extends to a broader range of industries, particularly manufacturing sectors such as electronics.[42]

σ_ϕ is the standard deviation of the weighting function $\phi(\cdot)$. Since the weighting function is centered at $q = \nu \ln \mathcal{C}$, σ_ϕ governs the rate at which this weight diminishes as the quality level deviates from the center and determines the concentration of consumer expenditure around $q = \nu \ln \mathcal{C}$. Thus, together with the nonhomotheticity parameter ν, it determines the degree of overlap in consumption baskets across different income groups. I normalize this value to 1, as the slope of the quality Engel curve cannot separately identify σ_ϕ from the nonhomotheticity parameter ν.

$1/\kappa$ represents the average unemployment duration during the transition between the sectors. Hence, κ governs the cost of labor market reallocation, making it a key determinant of the welfare impact of labor market disruptions caused by automation. I calibrate this value to align with the upper bound of direct sector-switching cost estimates in Lee and Wolpin [2006], set at 75% of average annual earnings. This is derived by assuming individuals receive 40% of their previous sector wage as unemployment benefits. Given an average unemployment spell of $1/\kappa$, the direct switching cost can be approximated as $0.4W_{h,t} \cdot \frac{1}{\kappa} = 0.75W_{h,t}$, yielding $\kappa \approx 0.53$.

[42] Additionally, the impacts on new vehicles may propagate to other expenditure categories, such as used vehicles. Including used vehicles raises the CPI weight to 6.5%.

Given the yearly model, this implies a prolonged unemployment spell exceeding six months, which is the typical duration of unemployment insurance in real-world policies. However, the model's extended unemployment duration can be interpreted as broadly capturing other disadvantages of sectoral reallocation, such as lower earnings after transitioning to a different sector due to skill mismatches. Additionally, the tax burden associated with extended unemployment benefits can be understood as reflecting the broader social costs of labor market disruptions.

Internal calibration. I estimate four parameters, $\Theta = (\nu, \lambda, \xi_m, \tau_y)$, using the simulated method of moments (SMM). This method minimizes a weighted average distance between data moments and simulated moments with respect to the model's parameters. The target data moments come from several sources used in the empirical section. The parameters are jointly determined, but I explain which data moments are particularly informative for disciplining each parameter.

The preference non-homotheticity parameter ν, which increases consumers' weight on higher quality as total consumption $\mathcal{C}$ rises, governs the speed at which market share shifts toward higher-quality levels in the simulated growing model economy. I calibrate ν using simulated data by replicating the regression in Equation 5 and matching the coefficient to the empirical estimate reported in Table 5. Given this calibration, I compute the cross-sectional quality Engel slope in the model and compare it to the empirical estimate for cars, 0.94, documented in Bils and Klenow [2001], which serves as a non-targeted moment.

The parameter λ, which governs the comparative advantage of robots versus labor across different quality levels, is calibrated to match the semi-elasticity of plant labor share with respect to product quality. This semi-elasticity is estimated in Equation 3 in the empirical section using data from a sample of U.S. vehicle manufacturing plants.[43]

ξ_m, together with κ, determines the degree of reallocation friction in the model. ξ_m governs the speed of labor adjustment in response to wage gaps between sectors. A larger ξ_m means faster labor reallocation. This parameter is calibrated by targeting data moments reflecting the long-run trends in the relative wage of the vehicle manufacturing industry compared to the general private sector, as well as the industry's employment share, over the period from 1980 to 2020 (data moments sourced from the BEA, as shown in Figure A3).

The proportional income tax rate, τ_y, is calibrated to match the average U.S. government

[43] I use the IV estimate in column (3) of Table 2. Since the unweighted sample estimate is considerably larger (-0.076) than the reported weighted sample estimate, I target the midpoint of the two.

spending-to-private consumption ratio of 0.19, as reported in Guerreiro et al. [2022].

5.2 Model Computation

The model simulates the economy starting from an initial worker distribution, μ_0, and a deterministic path of fundamentals, $\{S_t\}_{t=0}^{\infty}$. The objective is to analyze the extent to which quality-driven automation contributes to the rise in automation and the decline in labor share, relative to the long-run trends observed in the U.S. vehicle manufacturing industry over the past 40 years, starting from 1980. Additionally, the model investigates the welfare effects of rising automation, capturing its impacts on both product and labor markets. In the baseline laissez-faire scenario, no robot tax is imposed ($\tau_{k,t} = 0, \forall t$). Counterfactual policy experiments introduce a robot tax to evaluate its impact on the economy.

The model begins with an initial distribution μ_0, where workers are employed in either sector according to this distribution. I assign a small share of workers (2.4 percent) to sector d. Although this exceeds the 1.1 percent of total private sector employment in the U.S. attributed to the vehicle manufacturing industry in 1980, this initial distribution ensures equal wages across sectors in the initial period—providing a useful benchmark for solving the sector reallocation problem—and aligns the sectoral labor share with the industry's share in 1980.[44]

The model feature structural change in growing economy. Labor-augmenting technologies $(A_{o,t}, A_{d,t})$ grow at an annual rate of 2%, which roughly aligns with the growth trend of U.S. real GDP per capita. This economic growth drives overall consumption growth, shifting demand toward higher product quality and prompting changes in the input mix that favor greater automation. To disentangle the additional explanatory power of my framework in explaining the rise in automation and the decline in labor share—relative to previous studies that attribute rising automation primarily to declining robot prices—the simulation also explores the impact of technical progress in robot production through the declining path of R_t.

The model economy is non-stationary, and solving the forward-looking worker reallocation problem requires defining the terminal state. Ideally, this terminal state would correspond to the final balanced growth path of the growing economy. However, the non-homotheticity

[44]In terms of payroll, workers in sector d earn 2.4% of total payroll in the initial period, compared to the 1.7% attributed to vehicle manufacturing workers in 1980. This discrepancy arises from several factors abstracted in my model, such as the role of vehicle exports and imports, which create a gap between domestic expenditure and production, as well as other production factors, including intermediates and input-output structures.

in both preferences and production makes it difficult to define the balanced growth path. In line with the Turnpike approach in Boppart et al. [2023], I set the terminal condition for the worker reallocation problem using a state at a sufficiently large time $t = T$ as the terminal state.

I then solve for the transitional dynamics, where the fundamentals evolve deterministically until time $t = \hat{T}$ and remain constant for $\hat{T} < t \leq T$.[45] This transitional dynamics is interpreted as an approximation of the evolution of a growing economy. In this growing economy, consumer quality choice and firms' input decisions regarding product quality are determined within the quality grid space $Q = (0, \bar{q}]$. The upper bound $\bar{q}$ is set as $\bar{q} = 20$, which is sufficiently high to ensure it is not binding during model simulations. The quality q is discretized into 150 grids, with $q \in (0, 20]$, and worker ability ℓ is discretized into 4 grids.

5.3 Simulation results: Laissez-faire

To highlight the key features of the economy that drive rising automation along the growth path, I begin by describing the case where only labor-augmenting technologies, $(A_{o,t}, A_{d,t})$, grow at an annual rate of 2%. As the economy grows, consumers not only consume larger quantities but also demand higher product quality, moving toward higher levels of $q \in Q = (0, \bar{q}]$. Figure 3 left panel illustrates that, in the simulation, over 40 years, the market share has shifted toward higher quality levels, with the strength of this shift governed by the parameter ν. Interpreting each quality level as a variety j and considering a variety to be continuing if it holds more than 1% of the market share for two consecutive periods, I run the model version of Equation 5 as follows.

$$\triangle \ln s_{jt} = \alpha_t + \beta_1 \ln \overline{p}^Z_{jt} + \varepsilon_{jt} \tag{1.a}$$

$\nu = 4$ gives the estimate of $\hat{\beta}_1^{Model} = 0.016$, which closely matches the target of $\hat{\beta}_1^{Data} = 0.015$ in Table 5 column (1). This calibrated value of $\nu = 4$ also generates a cross-sectional quality Engel curve slope of 0.88, which is close to the data estimate of 0.94 in Bils and Klenow [2001].

As sales shift toward higher-quality goods with the overall growth in consumption, the sec-

[45] I set $\hat{T} = 100$ and $T = 200$, which is sufficiently large to ensure the solution of backward induction. I solve the model conditional on an expected sequence of the mass of workers $\{\mu_t(\mathbf{s})\}_{t=1}^T$. Given this expected path, I solve the static household and firm problems to determine wages and product prices, updating the value functions and the sequence $\{\mu_t(\mathbf{s})\}_{t=1}^T$. This process iterates until the expected path $\{\mu_t(\mathbf{s})\}_{t=1}^T$ is close enough to the realized path. Once consistency is achieved, the result is a sequential competitive equilibrium.

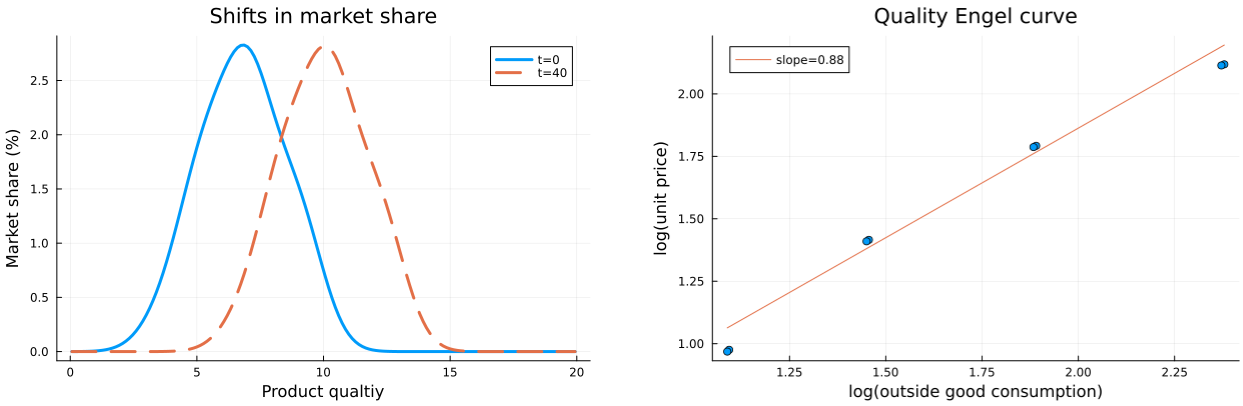


Figure 3: Nonhomothetic preference for qualtiy

Notes: The figure displays the market share over quality levels (left panel) and the quality Engel curve (right panel) in the model. The quality Engel curve slope is calculated as the elasticity of the average unit price of the consumer's purchase of differentiated goods to the consumer's outside good consumption, which corresponds to the empirical estimation in Bils and Klenow [2001].

tor's input mix changes. Firms that produce higher-quality goods rely more on robots than labor, as robots are more efficient for producing these higher-quality levels. Consequently, the industry-level decline in labor share is driven by the reallocation of sales toward firms producing higher-quality products, which use less labor and more robots. It is important to note that this channel is distinct from the decline in robot prices, R_t, which affects all firms across all quality levels. In contrast, the reallocation of sales toward higher-quality products results in a shift in market shares toward firms using less labor, leading to the industry-level decline in labor share.

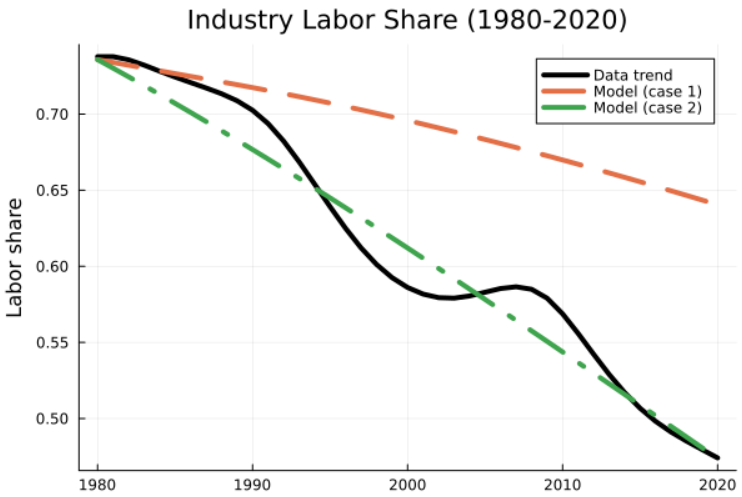


Figure 4: Trends of Motor Vehicle Manufacturing Labor Share

Notes: Sourced from the BEA-BLS KLEMS dataset, labor share is calculated as labor compensation divided by the value added for the industry category "Motor vehicles, bodies and trailers, and parts."

Figure 4 illustrates how well the model can explain the decline in labor share in the U.S. motor vehicle manufacturing industry between 1980 and 2020. In Case 1, where the decline in robot prices R_t is not considered, the model accounts for 9 percentage points of the 26 percentage point decline in labor share, or about one-third of the decline. To explain the remaining two-thirds of the decline, the model suggests that an annual 0.8% decline in robot prices R_t over $\hat{T}$ periods is required, as shown in Case 2 of the figure.

Over time, the decline in labor demand in sector d suppress the wage $W_{d,t}$, which, in turn, reduces the value of employment in that sector $V_t(E, d, \ell)$. Anticipating this worsening outlook for employment in sector d, workers decide to reallocate to the outside sector o, even if they must incur idiosyncratic mobility costs and experience a temporary unemployment spell, during which they face loss of earnings. In the end, the share of workers in sector d continuously declines during the transitional dynamics. This outflow of workers due to automation into other sectors can be understood in the context of structural change driven by economic growth, as well as the declining prices of robots.

Table 9: List of Moments

	Data	Model
Targeted moments		
β_1 estimate in Equation 5	0.015	0.016
semi-elasticity of LS to q	-0.055	-0.058
relative wage in 2020 (1980=1)	0.64	0.8
employment share in 2020 (1980=1)	0.72	0.81
change (%p) in labor share in 1980-2020	-26	-26
change (%) in income Gini	-23.5	-23.1
Untargeted moments		
quality Engel curve slope	0.94	0.88

Notes: Relative wage refers to the wage in the motor vehicle manufacturing industry (sector d) compared to the general private sector (sector o). Both relative wage and employment share are reported relative to 1980, which is normalized to 1. Refer Figure A3 for further details.

Table 9 summarizes the comparison of moments between the data and the model for case 2, where the annual decline in robot prices is included to match the size of the industry labor share decline. While the model matches the labor share decline, it cannot fully account for the extent of the declines in wages and employment share. This discrepancy could arise from other factors factors omitted in the model, such as the gap between expenditure and

production in an open economy, driven by outsourcing and imports.

However, even with the relatively small size of the robot price decline (an annual rate of 0.8%) compared to the larger rates assumed in the literature (e.g., a 2% annual decline in Guerreiro et al. [2022]), the model can explain significant labor displacement and ensuing labor reallocation in the context of economic growth. The model demonstrates that the quality-driven framework alone accounts for about one-third of the industry-level labor share decline, emphasizing the importance of rising demand for quality and the associated shift in the input mix toward automation. This provides an explanation for long-run trends in increasing automation and changes in factor shares within the industry.

6 Policy experiment

From an equity perspective, increasing automation may exacerbate inequality through two channels: the labor market, due to labor displacement, and the product market, as technological advances in robots disproportionately reduce the costs of higher-quality products, which are consumed more by wealthier households. Consequently, if the government values equity, it may choose to intervene. Taking this equity motive into account,[46] I analyze the welfare effects of government intervention. I begin by specifying the set of policy instruments available to the government.

Policy instruments. I assume the government has access to a restricted set of instruments: a linear tax on robots that can vary over time, $\{\tau_{k,t}\}_{t=0}^{\infty}$. This assumption aligns with the literature on automation regulation (e.g., Guerreiro et al. [2022]; Beraja and Zorzi [2024]), which excludes more flexible instruments, such as targeted transfers by worker type and time, due to their substantial information requirements, making them less practical in real-world applications.

Alternatively, policies that directly address labor reallocation frictions, such as subsidizing labor mobility, could be considered. In my model, such policies might increase the

[46] Alternatively, the government may intervene on efficiency grounds if excessive automation results in social costs that exceed the private costs faced by firms. For instance, Acemoglu and Restrepo [2024] demonstrates that when worker rents exist, the opportunity cost of labor—i.e., the social value generated by displaced workers—is lower than the wage rate, leading firms to over-automate and causing allocative inefficiency. Similarly, Beraja and Zorzi [2024] highlights inefficiencies arising when firms fail to internalize the impact of automation on displaced workers, particularly as it pushes them against their borrowing constraints. Moreover, if automation generates negative externalities in the form of labor market disruptions, leading to social costs not incorporated into firms' automation decisions, government intervention may be considered on efficiency grounds. I leave a detailed exploration of this motive for future research.

unemployment exit hazard rate, κ, or reduce migration costs, ξ_m, through measures like subsidies for retraining programs. While this is an important area for analysis, the current study focuses on the robot tax, which has received significant attention in both policy and academic discussions.

Social welfare function. To compute the optimal policy, I first specify the social welfare objective of the planner. Following the literature (e.g., Boar and Midrigan [2022]; Benabou [2002]), I calculate a measure of consumption-equivalent welfare for each worker type $\mathbf{s} = (e, h, \ell)$. Under the assumption log utility, worker welfare $\omega(\mathbf{s})$ is a solution to

$$V_0(\mathbf{s}) = \sum_{t=0}^{\infty} \beta^t \log(\omega(\mathbf{s}))$$

where $V_0(\mathbf{s})$ is the lifetime utility of the individual, evaluated at the initial period $t = 0$. The social welfare function, $\mathcal{W}$, is defined as an aggregation of individual welfare weighted by the initial distribution according to

$$\mathcal{W} = \left(\int \omega(\mathbf{s})^{1-\triangle} d\mu_0 \right)^{\frac{1}{1-\triangle}}$$

where $\triangle \geq 0$ is a parameter that captures the planner's preference for redistribution. If $\triangle = 0$, the planner maximizes average welfare, focusing solely on efficiency without valuing equity. If $\triangle \to 1$, the social welfare function reflects the preferences of a utilitarian planner, given by $\mathcal{W} = \exp(\int \log \omega(\mathbf{s}) d\mu_0)$, under the assumption of log utility. In the limit as $\triangle \to \infty$, the objective corresponds to that of a Rawlsian planner. I solve for the optimal policy chosen by a utilitarian planner as the baseline.

Optimal robot tax. The robot tax increases the price of robots by at a rate of $\tau_{k,t}$. The imposition of this tax raises the relative cost of robot use, leading to an increase in labor demand for each firm $j \in J$. This shift helps counteract the declining trend in wage in sector d and slows worker displacement by automation.

However, the tax also distorts the allocation of workers across sectors, thereby affecting the cost structure of production. As a result, higher input costs may negatively impact workers, potentially increasing the price index and slowing the trajectory of quality growth that consumers experience. I evaluate the welfare implications of this policy by analyzing its effects on heterogeneous workers and the aggregate economy, accounting for these trade-offs.

I solve the constrained Ramsey problem by evaluating social welfare across a set of considered robot tax paths during the transition, denoted by $\{\tau_{k,t}\}_{t=1}^{T}$, and referring to the path that maximizes social welfare as the optimal robot tax policy. For the set of alternative robot tax paths under consideration, I include the path derived from the parametric perturbation approach outlined in Beraja and Zorzi [2024], which uses the primal approach. In this framework, the government directly selects the wage gap between sectors through parametric perturbations.[47]

$$\hat{w}_t = \log(W_{o,t}) - \log(W_{d,t})$$

This parametric perturbation assumes temporary intervention, specifying the duration of intervention required for the controlled wage gap to gradually converge to the laissez-faire wage gap over time. This temporary intervention is assumed because maintaining the robot tax eventually disadvantages even workers initially employed in the automated sector, as it reduces their option value—the potential benefits they could gain from reallocating to the non-automated sector.

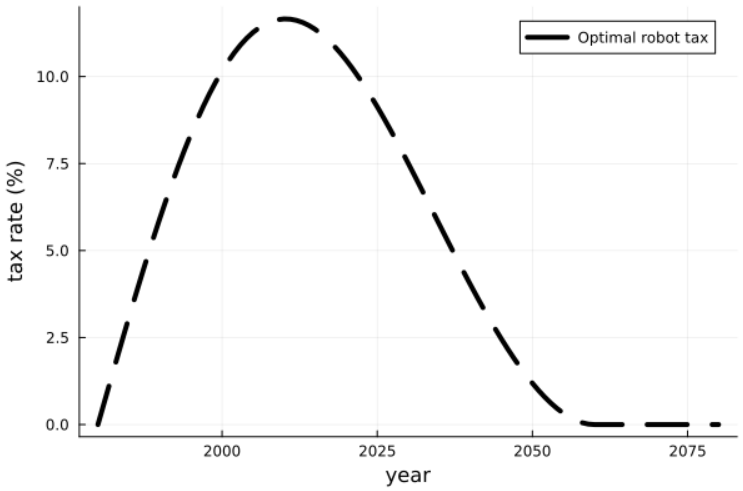


Figure 5: Optimal path of robot tax

Hence, the planner's problem considers a set of wage gap paths initially suppressed by government intervention, but which gradually converge to the laissez-faire wage gaps at varying speeds. For each considered path of the wage gap, the corresponding path of the robot tax $\{\tau_{k,t}\}_{t=1}^{T}$ that generates this wage gap can be determined. To explore a broader set of alternatives, I consider, in addition to the paths derived from parametric perturbations,

[47] For a detailed explanation of parametric perturbations, see subsection A.6.

paths where the tax rate is half of those in the parametric perturbation paths. The welfare effects of this tax policy are then evaluated, and the path that yields the highest social welfare among the considered alternatives is referred to as the optimal path of the robot tax.

Figure 5 illustrates the optimal robot tax trajectory determined within the model to maximize the utilitarian social welfare function. As the baseline, I assume transfer $\zeta(\mathbf{s})$ is proportional to tax payments, ensuring that the transfer itself does not have redistributive effects. The tax rate peaks at approximately 12% in 2010 ($t = 30$) and gradually decreases to 0 by 2060 ($t = 80$), as the model assumes a temporary intervention. The tax rate rises and falls because the model economy initially starts with equal wage rates across sectors, such that $W_{o,t} = W_{d,t}$ at $t = 0$. Therefore, there is no need to tax robots initially. Over time, as automation increases, the wage gap under laissez-faire $\hat{w}_t^{LF} = \log(W_{o,t}^{LF}) - \log(W_{d,t}^{LF})$ grows. The planner intervenes to manage this widening wage gap more gradually, and as automation increases, the robot tax size required to dampen the wage gap becomes larger. However, as the planner allows the wage gap to gradually approach the laissez-faire level, the intervention diminishes over time. Consequently, the tax rate rises and then declines.

Table 10: Welfare gains from optimal tax (%) by initial employment sector

Automated ($h = d$)	Non-automated ($h = o$)	Total
3.52	-0.07	0.011

Notes: Welfare gains in terms of CEV (Consumption Equivalent Variation).

Table 10 summarizes the welfare effects of the optimal robot tax, expressed as the consumption equivalent variation relative to the laissez-faire state (defined as Case 2 in the earlier section). The first column shows that workers initially employed in the automated sector d benefit significantly, with welfare gains of approximately 3.5%. However, the majority of workers, who were initially employed in the non-automated sector o, generally experience welfare losses. Since the subset of workers benefiting from the robot tax represents a small minority (initially only 2.4% of the workforce), the overall increase in social welfare is minimal, amounting to just about 0.01%.

Since robots are more intensively used in producing higher-quality products, which are consumed more by wealthier households, the robot tax creates additional distributional effects by increasing the price index for wealthier consumers more. Figure 6 displays the welfare effects of the robot tax on workers by ability ℓ for those initially employed in sector

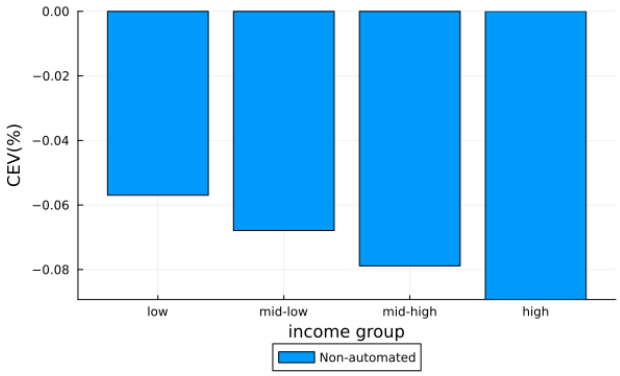


Figure 6: Welfare gains by worker ability initially employed in sector o

o, so that welfare changes reflect only the effects through the product market, highlighting this channel.

While the robot tax raises production costs and consequently increases the price index across all income groups, higher-income workers experience greater welfare losses, as their price index increases more due to their consumption of higher-quality goods, which are more affected by the robot tax. High-income (with the highest ability ℓ) workers experience a welfare loss of -0.09%, while the welfare loss for low-income workers decreases to approximately -0.056%.

While the difference in welfare effects between income groups may seem small, this quantification is based on the assumption that only 3.7% of consumer expenditure (the CPI weight for new vehicles) is exposed to quality-driven automation and the robot tax. As suggested by evidence from a broader set of manufacturing industries in the empirical section, if the categories of consumption exposed to quality-driven automation are larger, and the robot tax affects a greater set of consumption categories, the magnitude of welfare effects and the differential effects on the price index by income group explored in this policy experiment would increase accordingly. In this case, the current model results are likely to provide a lower bound for the welfare impacts.

7 Conclusion

This paper examines quality improvement as a key driver motivating firms to automate their production processes. It investigates the impacts of automation and the welfare effects of a robot tax, accounting for both consumer preferences and firms' input decisions concerning

output quality. Empirical findings from U.S. automotive industry data suggest that quality improvement is a primary motivation for firms to automate, providing a coherent explanation for the observed productivity effects of automation. This emphasizes the importance of understanding the evolution of firms' products and the relationship between automation and product quality, rather than focusing solely on the cost-saving effects of automation technologies.

The model suggests that quality-driven factors account for approximately one-third of the decline in labor share within the industry. Policy simulations further indicate that slowing the pace of automation through measures such as a robot tax could enhance social welfare. However, such policies also have important distributional consequences, depending on which consumer groups benefit most from goods produced through automation. The model results highlight the trade-offs associated with the robot tax and how welfare effects differ across households, including by sector of employment and income level. By incorporating differences in income and consumption baskets, as reflected in the quality Engel curve, the analysis provides insights into how regulating—and potentially slowing—automation can generate varied distributional effects through the product market. This aspect has received limited attention in the existing literature.

In conclusion, this paper underscores quality improvement as a central driver of automation in production processes. By presenting new evidence from the U.S. automotive industry, this study sheds light on this important, yet understudied, aspect of automation. Understanding this dynamic is crucial for explaining the rise of automation, its future trajectory, and its welfare implications.

References

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108 (6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30, 2019.

Daron Acemoglu and Pascual Restrepo. Automation and rent dissipation: Implications for wages, inequality, and productivity. Working Paper 32536, National Bureau of Economic Research, June 2024. URL http://www.nber.org/papers/w32536.

Daron Acemoglu, Claire Lelarge, and Pascual Restrepo. Competing with robots: Firm-level evidence from france. In *AEA papers and proceedings*, volume 110, pages 383–388. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2020.

Philippe Aghion, Céline Antonin, Simon Bunel, and Xavier Jaravel. What are the labor and product market effects of automation? new evidence from france. 2020.

David Argente and Munseob Lee. Cost of living inequality during the great recession. *Journal of the European Economic Association*, 19(2):913–952, 2021.

Kai Arvai and Katja Mann. Consumption inequality in the digital age. 2022.

Jie Bai, Panle Jia Barwick, Shengmao Cao, and Shanjun Li. Quid pro quo, knowledge spillover, and industrial quality upgrading: Evidence from the chinese auto industry. Technical report, National Bureau of Economic Research, 2020.

Maria Bas and Caroline Paunov. Input quality and skills are complementary and increase output quality: Causal evidence from ecuador's trade liberalization. *Journal of Development Economics*, 151:102668, 2021.

Erich Battistin, Richard Blundell, and Arthur Lewbel. Why is consumption more log normal than income? gibrat's law revisited. *Journal of Political Economy*, 117(6):1140–1154, 2009.

Rudolfs Bems and Julian Di Giovanni. Income-induced expenditure switching. *American Economic Review*, 106(12):3898–3931, 2016.

Roland Benabou. Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70(2):481–517, 2002.

Martin Beraja and Nathan Zorzi. Inefficient Automation. *The Review of Economic Studies*, page rdae019, 02 2024. ISSN 0034-6527. doi: 10.1093/restud/rdae019. URL https://doi.org/10.1093/restud/rdae019.

Mark Bils and Peter J Klenow. Quantifying quality growth. *American Economic Review*, 91(4):1006–1030, 2001.

Corina Boar and Virgiliu Midrigan. Efficient redistribution. *Journal of Monetary Economics*, 131:78–91, 2022.

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

Timo Boppart, Per Krusell, and Jonna Olsson. Labor supply when productivity keeps growing. *Review of Economic Dynamics*, 50:61–87, 2023.

Natalie Chen and Luciana Juvenal. Quality and the great trade collapse. *Journal of Development Economics*, 135:59–76, 2018.

Banu Demir, Ana Cecilia Fieler, Daniel Yi Xu, and Kelly Kaili Yang. O-ring production networks. *Journal of Political Economy*, 132(1):200–247, 2024.

Timothy DeStefano and Jonathan Timmis. Robots and export quality. *Journal of Development Economics*, 168:103248, 2024.

Jay Dixon, Bryan Hong, and Lynn Wu. The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605, 2021.

Maya Eden and Paul Gaggl. On the welfare implications of automation. *Review of Economic Dynamics*, 29:15–43, 2018.

Marcela Eslava, John Haltiwanger, and Nicolas Urdaneta. The size and life-cycle growth of plants: The role of productivity, demand, and wedges. *Review of Economic Studies*, 91 (1):259–300, 2024.

Benjamin Faber and Thibault Fally. Firm heterogeneity in consumption baskets: Evidence from home and store scanner data. *The Review of Economic Studies*, 89(3):1420–1459, 2022.

Pablo Fajgelbaum, Gene M Grossman, and Elhanan Helpman. Income distribution, product quality, and international trade. *Journal of political Economy*, 119(4):721–765, 2011.

Ana Cecília Fieler, Marcela Eslava, and Daniel Yi Xu. Trade, quality upgrading, and input linkages: Theory and evidence from colombia. *American Economic Review*, 108(1):109–146, 2018.

Joao Guerreiro, Sergio Rebelo, and Pedro Teles. Should robots be taxed? *The Review of Economic Studies*, 89(1):279–311, 2022.

Jessie Handbury. Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities. *Econometrica*, 89(6):2679–2715, 2021.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754, 2017.

Colin J Hottman, Stephen J Redding, and David E Weinstein. Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, 131(3):1291–1364, 2016.

Nir Jaimovich, Itay Saporta-Eksten, Henry Siu, and Yaniv Yedid-Levi. The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics*, 122: 1–16, 2021.

Matthias Kehrig and Nicolas Vincent. The micro-level anatomy of the labor share decline. *The Quarterly Journal of Economics*, 136(2):1031–1087, 2021.

Amit Khandelwal. The long and short (of) quality ladders. *The Review of Economic Studies*, 77(4):1450–1476, 2010.

Amit K Khandelwal, Peter K Schott, and Shang-Jin Wei. Trade liberalization and embedded institutional reform: Evidence from chinese exporters. *American Economic Review*, 103 (6):2169–2195, 2013.

Michael Koch, Ilya Manuylov, and Marcel Smolka. Robots and firms. *The Economic Journal*, 131(638):2553–2584, 2021.

Michael Kremer. The o-ring theory of economic development. *The quarterly journal of economics*, 108(3):551–575, 1993.

Maurice Kugler and Eric Verhoogen. Prices, plant size, and product quality. *The Review of Economic Studies*, 79(1):307–339, 2012.

Donghoon Lee and Kenneth I Wolpin. Intersectoral labor mobility and the growth of the service sector. *Econometrica*, 74(1):1–46, 2006.

Lukas Nord. Shopping, demand composition, and equilibrium prices. *Available at SSRN 4178271*, 2022.

Ezra Oberfield. Inequality and measured growth. Working Paper 31096, National Bureau of Economic Research, March 2023. URL http://www.nber.org/papers/w31096.

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

Klaus Prettner and Holger Strulik. Innovation, automation, and inequality: Policy challenges in the race against the machine. *Journal of Monetary Economics*, 116:249–265, 2020.

Pascual Restrepo. Automation: Theory, evidence, and outlook. *Annual Review of Economics*, 16, 2023.

Peter K Schott. Across-product versus within-product specialization in international trade. *The Quarterly Journal of Economics*, 119(2):647–678, 2004.

Eric A Verhoogen. Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530, 2008.

A Appendix

A.1 Tables

Table A1: Response of log(Sales) to Robot Adoption Event

horizon h	-2	-1	0	1	2	3
β_h	-0.05	0	0.06	0.26	0.36	0.47
	(0.05)	.	(0.11)	(0.12)	(0.13)	(0.17)
Observations	506	.	576	508	445	388

Notes: Standard errors are reported in parentheses and are clustered by plant. The sample period spans from 2000 to 2021.

Table A2: Regression of Plant Labor Share on Quality

	(1)	(2)	(3)	(4)
Dependent variable →		LS_{jt}		
q_{jt}	-0.024	-0.027	-0.032	-0.035
	(0.005)	(0.007)	(0.010)	(0.008)
standardized coef.	-0.57	-0.53	-0.64	-0.69
Specification	OLS	OLS	2SLS	2SLS
Instrument			JD Power	Khandelwal et al. [2013]
Sample Year	1993-2021	2006-2021	2006-2021	2006-2021
Observations	1,311	575	575	575
R^2	0.547	0.595	-	-

Notes: The tables reports the estimates from Equation 3. Standard errors clustered by plant in parentheses. Plant-level quality proxy q_{jt} is measured as the sales-weighted average log unit price of vehicle models product in plant j in year t. Regression controls a set of fixed effects including year, state, company and plant. Observations are weighted by the sales share within a year (s_{jt}).

Table A3: Decomposition of the Labor Share Changes to Quality Upgrading

Dependent variable ↓	(1)	(2)	(3)	(4)
LS_{jt}	-0.024	-0.027	-0.032	-0.035
	(0.005)	(0.007)	(0.010)	(0.008)
$log(Payroll_{jt})$	0.32	0.13	0.51	0.51
	(0.14)	(0.18)	(0.30)	(0.23)
$log(Sales_{jt})$	0.85	0.71	1.18	1.26
	(0.15)	(0.21)	(0.28)	(0.24)
Specification	OLS	OLS	2SLS	2SLS
Instrument			JD Power	Khandelwal et al. [2013]
Sample Year	1993-2021	2006-2021	2006-2021	2006-2021
Observations	1,311	575	575	575

Notes: The table reports the estimates from Equation 3 with either the labor share, log of payroll, or sales as the dependent variable. Each estimate represents the coefficient for the plant-level quality proxy q_{jt} from a separate regression. Standard errors clustered by plant in parentheses. Regression controls a set of fixed effects including year, state, company and plant. Observations are weighted by the sales share within a year (s_{jt}).

A.2 Figures

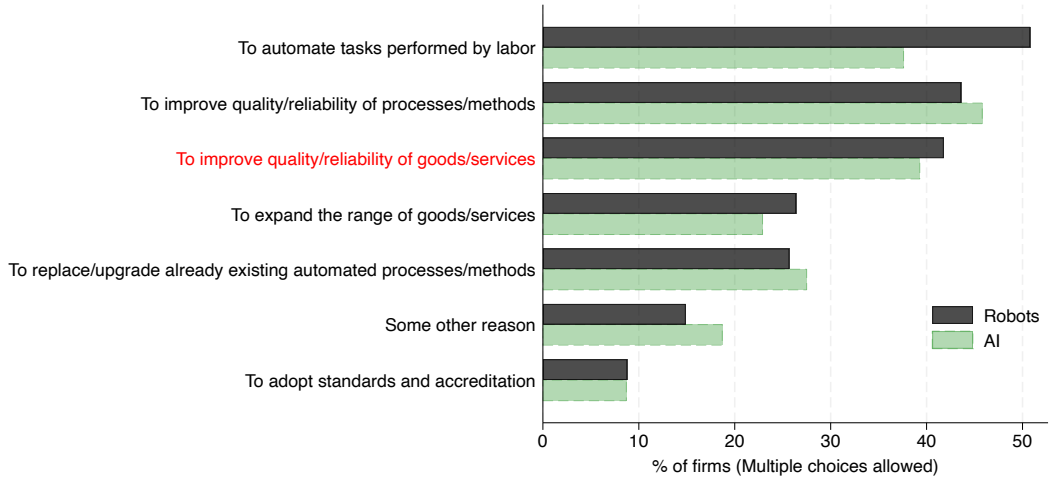


Figure A1: Motivation for technology use

Notes: Sourced from the U.S. Census Bureau's Annual Business Survey (ABS) 2023, the sample includes all sectors, with a reference year of 2022. The graph displays the responses to the survey question: "During the three years 2020 to 2022, why did this business adopt or use the technology?"

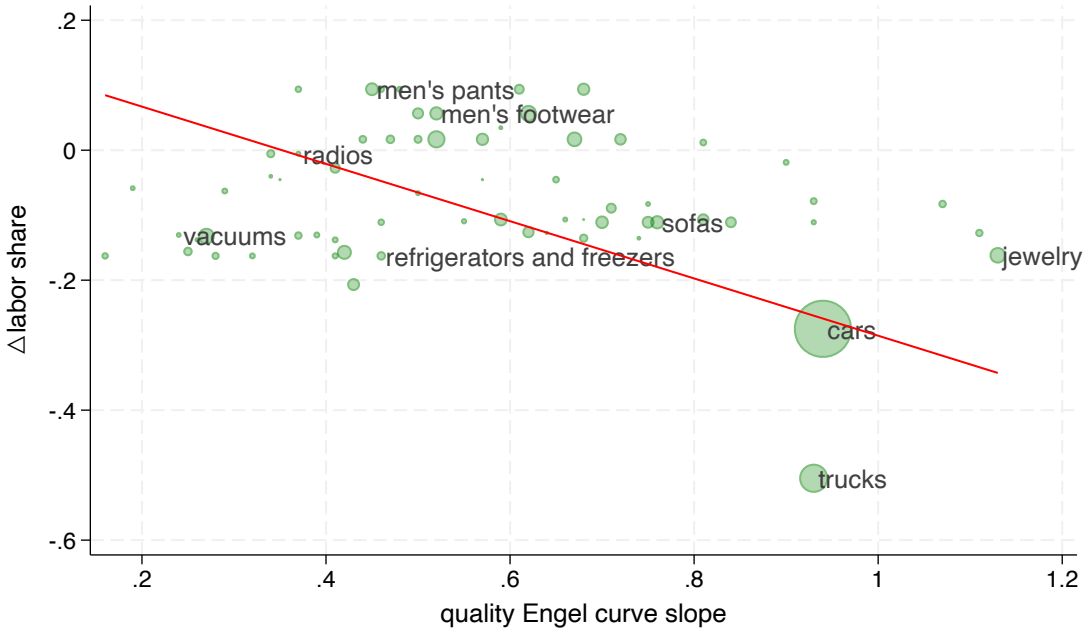


Figure A2: Quality Engel curve slope and changes in labor share (1980-2018)

Notes: The slope of the quality Engel curve refers to the elasticity of the unit price of a consumer's purchase with respect to the consumer's total nondurable consumption. It indicates how much consumers value the quality of a good as their consumption increases. The estimates for 66 durable goods are from Bils and Klenow [2001], using data from the U.S. Consumer Expenditure Surveys. The y-axis represents changes in labor share (payroll/value added) between 1980 and 2018, calculated from the NBER-CES manufacturing databases for industries mapped to each of these goods. Each point represents a good, and the size of the circle represents the weight in the CPI from Bils and Klenow [2001].

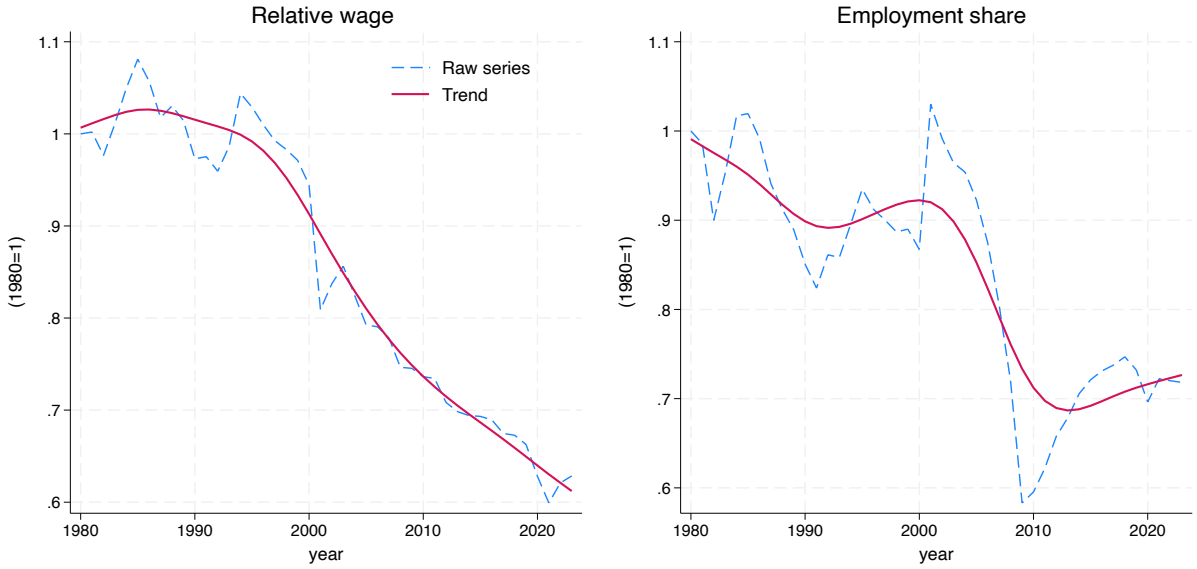


Figure A3: Trends of Motor Vehicle Manufacturing Relative to the General Private Sector

Notes: Sourced from the Bureau of Economic Analysis (BEA), the figure illustrates long-run trends in the motor vehicle manufacturing industry relative to the general private sector. The left panel shows wages in the motor vehicle manufacturing industry relative to the general private sector. The right panel displays the industry's employment share relative to the general private sector. Both series are normalized to 1 in 1980. The data for the left panel is sourced from BEA Table 6.6C, "Wages and Salaries Per Full-Time Equivalent Employee by Industry", while the right panel draws from BEA Table 6.5C, "Full-Time Equivalent Employees by Industry."

A.3 J.D. Power Quality Ratings

Quality ratings are based on initial owner response and feedback of their new vehicles during the first 90 days of ownership.

- Quality&Reliability: Rating focuses on problems experienced by verified owners. This rating is based on problems experienced including defects and malfunctions, or where controls or features are difficult to understand and operate.
- Driving Experience: It is based on verified owners' ratings of the performance, comfort, functionality, perceived safety and technology of their new vehicles.
- Resale Value: J.D. Power's forecast of depreciation (how much of its current value the vehicle will likely lose) over the next three years. Lower depreciation equals a higher rating

There is an another rating category - dealership experience. However, I focus on the previous three categories as they provide a more direct measure of a vehicle's quality.

A.4 Khandelwal et al. [2013] Quality Measure

For an item i in product group g, the following equation is derived from a CES demand system:

$$\ln x_{igt} = (\sigma_g - 1)\ln q_{igt} - \sigma_g \ln p_{igt} + (\sigma_g - 1)\ln P_{gt} + \ln E_{gt} \tag{34}$$

x_{igt} represents the quantity demanded, q_{igt} is the quality, p_{igt} is the price, P_{gt} is the price index, and E_{gt} is the expenditure on product group g. σ_g is the elasticity of substitution. With an assumed value of σ_g, we can rewrite the equation as:

$$\ln x_{igt} + \sigma_g \ln p_{igt} = FE + \varepsilon_{igt} \tag{35}$$

FE is a set of product group and year fixed effects which captures the $(\sigma_g - 1)\ln P_{gt} + \ln E_{gt}$. Then the residual ε_{igt} can be estimated from a regression, and the quality measure is derived from the residual estimates as:

$$\ln q_{igt} = \widehat{\varepsilon}_{igt}/(\sigma_g - 1) \tag{36}$$

In this paper, an item refers to a vehicle model, and a product group refers to a vehicle type. As a baseline, I assume $\sigma_g = 9$ for all $g \in \{\text{cars, SUVs, pickups, vans}\}$.

A.5 Production function

Suppress the quality subscript q for brevity. Set $N = 1$ to have the range of tasks between $[0, 1]$.

$$L = -\int_{N-1}^{N} p(i)y(i)di + \lambda \left[\left(\int_{N-1}^{N} y(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} - Y \right] \tag{37}$$

FOCs lead to:

$$y(i) = \left(\frac{p(i)}{P} \right)^{-\sigma} Y \tag{38}$$

where P is the price index

$$P := \left(\int_{N-1}^{N} p(i)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{39}$$

quantity of $K(i)$ required to produce $y(i),\ i \leq I^*$

$$\psi^k(i)K(i) = \left(\frac{R/\psi^k(i)}{P} \right)^{-\sigma} Y \Rightarrow K(i) = \psi^k(i)^{\sigma-1} \left(\frac{R}{P} \right)^{-\sigma} Y \tag{40}$$

$$K := \int_{N-1}^{I^*} K(i)di = \left(\frac{R}{P} \right)^{-\sigma} Y \left(\int_{N-1}^{I^*} \psi^k(i)^{\sigma-1} di \right) \tag{41}$$

$$\frac{R}{\psi^k(i)} = \left(\frac{Y}{K} \left(\int_{N-1}^{I^*} \psi^k(i)^{\sigma-1} di \right) \right)^{\frac{1}{\sigma}} \frac{P}{\psi^k(i)} \tag{42}$$

quantity of $L(i)$ required to produce $y(i),\ i > I^*$

$$\psi^l(i)L(i) = \left(\frac{W/\psi^l(i)}{P} \right)^{-\sigma} Y \Rightarrow L(i) = \psi^l(i)^{\sigma-1} \left(\frac{W}{P} \right)^{-\sigma} Y \tag{43}$$

$$L := \int_{I^*}^{N} L(i)di = \left(\frac{W}{P} \right)^{-\sigma} Y \left(\int_{I^*}^{N} \psi^l(i)^{\sigma-1} di \right) \tag{44}$$

$$\frac{W}{\psi^l(i)} = \left(\frac{Y}{L} \left(\int_{I^*}^{N} \psi^l(i)^{\sigma-1} di \right) \right)^{\frac{1}{\sigma}} \frac{P}{\psi^l(i)} \tag{45}$$

Plug Equation 42 and Equation 45 into Equation 39.

$$P^{1-\sigma} = \int_{N-1}^{I^*} \left[\left(\frac{Y}{K} \left(\int_{N-1}^{I^*} \psi^k(i)^{\sigma-1} di \right) \right)^{\frac{1}{\sigma}} \frac{P}{\psi^k(i)} \right]^{1-\sigma} di + \int_{I^*}^{N} \left[\left(\frac{Y}{L} \left(\int_{I^*}^{N} \psi^l(i)^{\sigma-1} di \right) \right)^{\frac{1}{\sigma}} \frac{P}{\psi^l(i)} \right]^{1-\sigma} di \tag{46}$$

$$Y = \left[\left(\int_{N-1}^{I^*} \psi^k(i)^{\sigma-1} di \right)^{\frac{1}{\sigma}} K^{\frac{\sigma-1}{\sigma}} + \left(\int_{I^*}^{N} \psi^l(i)^{\sigma-1} di \right)^{\frac{1}{\sigma}} L^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{47}$$

A.6 Primal approach with parametric perturbations

The solution for the optimal path of robot tax during the transition, $\{\tau_{k,t}\}_{t=1}^{T}$, draws on the framework developed by Beraja and Zorzi [2024]. The government is assumed to directly control the path of wage gap between sectors. This is modeled through parametric perturbations:

$$\hat{w}_t = \log(W_{o,t}) - \log(W_{d,t}) = S(t; \lambda_g)\hat{w}_t^{LF}$$

where $\hat{w}_t^{LF}$ denotes the wage gap under laissez-faire, and $S(t; \lambda_g)$ is a smoothstep function bounded between $[0, 1]$. The government selects λ_g from a grid Λ_g to maximize social welfare.

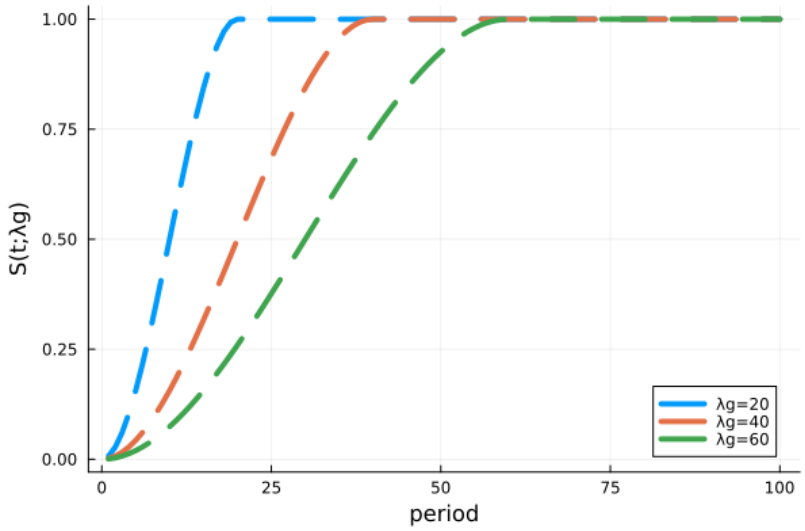


Figure A4: Smoothstep function $S(t; \lambda_g)$

Figure A4 illustrates that λ_g is the parameter determining the period at which the wage gap under government intervention converges to the laissez-faire wage gap, such that $S(t; \lambda_g) = 1$ at $t = \lambda_g$. This parametric perturbation assumes temporary intervention, specifying the duration of intervention required for the controlled wage gap to gradually converge to the laissez-faire wage gap over time.

The optimal λ_g^*, interpreted as the stopping time of the intervention, is obtained by searching over a grid Λ_g, where $\lambda_g^* \in \Lambda_g$. This parameter determines the speed at which the wage gap converges to its laissez-faire level over time. Once λ_g^* is selected, the path of the robot tax, $\{\tau_{k,t}\}_{t=1}^{T}$, that generates the wage gap $\{\hat{w}_t\}_{t=1}^{T}$ under λ_g^* is recovered. This resulting trajectory is referred to as the optimal robot tax path derived through the primal approach.