

Automation and the Rise of Superstar Firms[★]

Hamid Firooz^a, Zheng Liu^b and Yajie Wang^c

^aFederal Reserve Bank of San Francisco, USA

^bFederal Reserve Bank of San Francisco, USA

^cUniversity of Missouri, USA

ARTICLE INFO

Keywords:

Automation
industry concentration
superstar firms
markup
productivity

ABSTRACT

We provide empirical evidence suggesting that the rise of superstar firms is linked to automation. We explain this empirical link in a general equilibrium framework with heterogeneous firms and variable markups. Firms can operate a labor-only technology or, by paying a per-period fixed cost, an automation technology that uses both workers and robots. The fixed costs lead to an economy-of-scale effect of automation, such that larger and more productive firms are more likely to automate. Automation boosts labor productivity, allowing those large firms to expand further, raising industry concentration. Since robots substitute for workers, increased automation raises sales concentration more than employment concentration, consistent with empirical evidence. Under our calibration, a modest robot subsidy mitigates markup distortions and improves welfare by stimulating automation investment, bringing aggregate output closer to the efficient level.


1. Introduction

Industries in the United States have become more concentrated over time, with each major sector increasingly dominated by a small number of superstar firms (Autor, Dorn, Katz, Patterson and Van Reenen, 2020). Based on empirical evidence and a theoretical framework, we argue that the rise in automation since the early 2000s has contributed significantly to the rise of superstar firms, particularly in the manufacturing sector.

Our study is motivated by the observation that the rise of industry concentration in U.S. manufacturing has been accompanied by steady increases in robot adoption over the past two decades. Figure 1 illustrates this correlation. Sales concentration in the U.S. manufacturing sector—measured by the sales share of the largest four firms—has increased from about 40.5 percent in the late 1990s to about 43.5 percent in 2012, while employment concentration has stayed relatively flat (Panel A, reproduced from Autor et al. (2020)), implying a widening gap between sales concentration and employment concentration. On the other hand, robot density—measured by the number of industrial robots per thousand manufacturing employees—has quadrupled since the early 2000s while the relative price of robots has declined by about 40 percent (Panel B).¹

We examine the economic forces that drive the link between automation and the rise of superstar firms. We first provide some additional stylized facts on the relationship between automation and industry concentration using Compustat data. We show that, at the industry level, robot density has a significantly positive correlation with sales

[★]We are grateful to the Associate Editor (Felipe Saffie) and an anonymous referee for constructive comments. We also thank David Autor, Yan Bai, Martin Beraja, Maryam Farboodi, Narayana Kocherlakota, Sylvain Leduc, Huiyu Li, Yueran Ma, Alan Moreira, Christina Patterson, Pascual Restrepo, Mathieu Taschereau-Dumouchel, Ivan Werning, and seminar participants at the NBER Economics of Artificial Intelligence conference, The Bank of Japan, Federal Reserve Bank of San Francisco, the 2024 EEA-ESEM Congress, the 2024 ESIF Economics and AI+ML Meeting, Midwest Macroeconomics Meetings (Logan), and China International Conference in Macroeconomics 2022 for helpful comments. We also thank Mollie Pepper for excellent research assistance and Anita Todd for editorial assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.


 hmdfirooz@gmail.com (H. Firooz); zheng.liu@sf.frb.org (Z. Liu); yajie.wang@missouri.edu (Y. Wang)

 <https://sites.google.com/view/hamidfirooz> (H. Firooz);

<https://www.frbsf.org/our-people/economists/zheng-liu> (Z. Liu); <https://www.yajiewang.info> (Y. Wang)

ORCID(s):

¹Sales concentration has also increased in Europe. As documented by Bajgar, Berlingieri, Calligaris, Criscuolo and Timmis (2019), manufacturing sales concentration in Europe started rising a few years ahead of that in North America (see their Figure 9). The adoption of industrial robots also started earlier in the European market than in the North American market (Acemoglu and Restrepo, 2020). The synchronized increases in industry concentration and robot adoptions in both Europe and North America suggest a potentially important link between the two salient trends.

Figure 1: Trends in Industry Concentration and Automation in Manufacturing

Note: Panel (a) is adapted from Autor et al. (2020) with permission from the Oxford University Press (License Number 5241431011126) and shows the industry concentration measured by both the sales share and the employment share of the top 4 firms (left scale) or the top 20 firms (right scale) across four-digit industries in the manufacturing sector. Panel (b) displays the unit value of newly shipped industrial robots deflated by the personal consumption expenditures price index (green line, left scale) and robot density measured by the operational stock of industrial robots per thousand manufacturing workers (blue line, right scale). Both the robot price and the operational stock of industrial robots are obtained from the International Federation of Robotics (IFR).

concentration but a small and insignificant correlation with employment concentration.² These empirical relations are robust when we estimate the correlations using an instrumental variable (IV) approach along the line of Acemoglu and Restrepo (2020). Specifically, we use lagged values of the average robot density at the industry level in five European economies as an IV for the robot density in the same industries in the United States.³ Our IV regressions suggest that the rise in sales concentration in the United States is associated with robot adoptions, whereas employment concentration is not. The estimated correlation between automation and sales concentration is economically important: a one standard deviation increase in robot density is associated with an increase in the sales share of the top 1% firms of about 10 percentage points, which is equivalent to an increase of one-third of its average value.

To understand the empirical link between automation and industry concentration, we construct a dynamic general equilibrium model featuring heterogeneous firms, endogenous automation decisions, and variable markups (with Kimball (1995) preferences). Firms have access to two types of technologies for producing differentiated intermediate goods: one is the traditional technology that uses labor as the sole input, and the other is an automation technology that uses both labor and robots with a constant elasticity of substitution. Operating the automation technology incurs a random per-period fixed cost, but it reduces the marginal cost of production relative to operating the labor-only technology. Firms also face idiosyncratic, persistent productivity shocks. A firm's automation decision (i.e., whether to use the labor-only technology or the automation technology) depends on the realization of the fixed cost relative to productivity. At a given fixed cost, a larger firm is more likely to automate because it has higher productivity, higher market power, and thus higher profits. Automation improves a firm's labor productivity, allowing large, robot-using firms to expand their sales share further. This economy-of-scale effect leads to a positive connection between automation and sales concentration. Since robots substitute for workers, the expansion of those large firms relies more on robots than on workers. Thus, a rise in automation raises sales concentration more than employment concentration, as we observe in the data.

In our model, a decline in the robot price drives the rise in automation, which in turn impacts industry concentration through two channels. First, a lower robot price and, therefore, a lower user cost of robots benefits large firms that operate the automation technology (an intensive-margin effect), enabling large firms to become even larger. Second, a lower robot price induces more firms to adopt robots (an extensive-margin effect), such that some smaller firms

²Throughout this paper, we focus on industrial robots, which is a specific type of automation technology. We use "automation" and "robots" interchangeably.

³The five European economies include Denmark, Finland, France, Italy, and Sweden, which all adopted robotics ahead of the United States.

that initially operate the labor-only technology would switch to the automation technology, reducing the sales share of the superstar firms and lowering industry concentration. The net effect of a decline in the robot price on industry concentration can be ambiguous, depending on the magnitude of the declines in the robot price and the calibration of the model parameters.

We calibrate the model parameters to match several moments observed in the U.S. manufacturing sector. Our calibrated model matches four key moments in the data, including the share of firms that use robots, their employment share, aggregate robot density, and the cumulative growth rate of robot density since the early 2000s. Under our calibration, the intensive-margin effect dominates, such that a decline in the robot price raises the sales concentration. The decline in the robot price also raises the employment concentration because automation boosts productivity, raising labor demand of automating firms. However, the increase in employment concentration is smaller than that in sales concentration because robots substitute for workers. These model predictions are consistent with our empirical evidence that robot adoptions are significantly correlated with the sales share of the top 1% firms, but not with the employment share of these firms. Under our calibration, the model predicts that a 40 percent decline in the relative price of robots—a magnitude observed during the past two decades—can explain about 49 percent of the rise in sales concentration in the U.S. manufacturing sector. It also explains about 25 percent of the divergence between sales and employment concentration.

Our calibrated model further predicts that the usage of automation technology is highly skewed toward a small fraction of large firms, in line with the cross-sectional evidence from the Annual Business Survey (ABS) conducted by the U.S. Census Bureau (Zolas, Kroff, Brynjolfsson, McElheran, Beede, Buffington, Goldschlag, Foster and Dinlersoz, 2020; Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo and Zolas, 2022).⁴ For example, according to the ABS, the adoption rate of robots in 2016-2018 by firms in the top percentile of the employment distribution within 6-digit industries was about 3 times of the adoption rate among firms in the 50th to 75th percentile (5.1% vs. 1.7%). Acemoglu et al. (2022) argue that this pattern supports the idea that the adoption of automation technologies involves large integration costs, consistent with our model's mechanism. Since larger firms have higher productivity, higher markups, and lower labor shares, our model suggests that the between-firm reallocation triggered by a decline in the robot price boosts aggregate productivity, increases the average markup, and reduces the average labor share. Such dynamics echo the reallocation channel documented by Autor et al. (2020), Acemoglu, Lelarge and Restrepo (2020a), Kehrig and Vincent (2021), and Hubmer and Restrepo (2022). Furthermore, a decline in robot prices raises equilibrium employment in automating firms, because it boosts those firms' labor productivity, leading to increased labor demand that dominates the labor-substituting effects of automation. This employment effect of automation is also in line with the evidence from the firm-level evidence documented by Zolas et al. (2020) and Aghion, Antonin, Bunel and Jaravel (2021).

The presence of monopolistic competition and variable markups in our model implies that the decentralized equilibrium allocation is inefficient, creating room for policy interventions to improve social welfare. A robot subsidy can alleviate the markup distortions. By lowering the user cost of robots, a subsidy reallocates production toward large automating firms that have high productivity, bringing aggregate output closer to the efficient level. However, since large firms also have high markups, a robot subsidy also raises the average markup through between-firm reallocation. Thus, our model implies an interior optimal level of robot subsidy. Under our calibration, a modest robot subsidy (about 1.41% of the value of robots) maximizes the steady-state welfare, yielding a welfare gain equivalent to about 4.23% of steady-state consumption compared to the laissez-faire benchmark. Although a permanent robot subsidy raises consumption in the steady state, consumption falls and employment rises in the transition process, reducing the overall welfare gains. Taking into account the transition dynamics, the optimal subsidy rate for automation is smaller (about 0.64%) and the ex ante welfare gains (relative to the case with no policy changes) are also smaller, at about 0.43% of consumption equivalent.

2. Related literature

Our work builds on the influential study of Autor et al. (2020), who document evidence of the steady rise of superstar firms in all major sectors of the U.S. economy since the early 2000s. Autor et al. (2020) discuss a few potential drivers of the rise of superstar firms (what they call a “winner takes most” mechanism), including greater market competition (e.g., through offshoring) or scale-biased technological changes driven by intangible capital investment and information technology. Other potential drivers of the rise in industry concentration have been studied in the literature, including

⁴The ABS covers a large and nationally representative sample of over 850,000 firms in all private, nonfarm business sectors.

uneven productivity growth across firms (Furman and Orszag, 2018; Yao, 2022; Choi, Levchenko, Ruzic and Shim, 2024), decline in knowledge diffusion between the frontier and laggard firms (Akcigit and Ates, 2019), a slowdown in radical innovations since the 1990s (Olmstead-Rumsey, 2019), the rise of specialized firms (Ekerdt and Wu, 2022), and digital advertising in customer accumulation (Shen, 2023). Our study focuses on the rise in automation technologies (and in particular, robot adoptions) as a driver of the rise of superstar firms.⁵

Our model underscores how fixed costs of automation disproportionately favor large, high-productivity firms, and thus raise sales concentration. The economy-of-scale feature of new technology adoptions has been explored in other studies, including, for example, Kwon, Ma and Zimmermann (2024), Aghion, Bergeaud, Boppart, Klenow and Li (2019), Hubmer and Restrepo (2022), Ridder (2023), Lashkari, Bauer and Boussard (2024), Tambe, Hitt, Rock and Brynjolfsson (2020), and Sui (2022). Relative to these studies, and in particular, the closely related parallel work of Hubmer and Restrepo (2022), our work makes three contributions.

First, we focus on explaining the link between robot adoptions and industry concentration. We show that the observed declines in robot prices can explain about half of the observed increase in sales concentration and about one-quarter of the divergence between sales and employment concentration in the U.S. manufacturing industry during the past two decades. Hubmer and Restrepo (2022) take a different approach by employing a task-based model in which firms incur fixed costs to automate new tasks. They show that a decline in the price of capital goods used for automation—including not only robots but also advanced technologies such as specialized software and dedicated machinery—can explain the observed changes in the labor share. Through a similar mechanism (economies of scale in automation), their model also generates increases in sales concentration. However, they do not address the divergence between sales and employment concentration—a salient empirical fact observed both in the time series (Autor et al., 2020; Hsieh and Rossi-Hansberg, 2019) and across industries (as we document in Section 3).

Second, our model highlights a non-monotonic relation between robot prices and industry concentration, reflecting the two opposing effects from the intensive margin versus the extensive margin of automation. When the robot price declines, some medium-sized firms switch technologies from labor-only to automation (i.e., the extensive margin). Although the drop in the robot price also benefits large and automating firms (intensive margin), when the robot price becomes sufficiently low, the usage of robots becomes widespread, such that the expansion of the medium-sized firms through automation erodes the market share of the top 1% firms, reducing sales concentration. This counterfactual illustrates a crucial difference between automation equipment and general capital equipment: the usage of robots is heavily skewed toward a small number of superstar firms, whereas the usage of general capital equipment is more widespread. When automation becomes widespread (e.g., when the robot price is sufficiently low), a further drop in the robot price may not increase industry concentration and may even reduce it. To our knowledge, this potential non-monotonic relation between robot prices and industry concentration is new to the literature.

Third, we use our calibrated model to examine the implications of automation policies, such as taxing (or subsidizing) robots, for macroeconomic allocations and welfare, taking into account transition dynamics under a given policy. In this aspect, our work contributes to the nascent but rapidly growing literature on automation policies. The rise in automation has raised an important policy question: Should robots be taxed? Some studies argue in favor of taxing robots because automation displaces routine workers and raises income inequality (Guerreiro, Rebelo and Teles, 2022; Acemoglu, Manera and Restrepo, 2020b; Beraja and Zorzi, 2022; Costinot and Werning, 2022).⁶ Our work contributes to this strand of literature by highlighting a new tradeoff facing automation policies: Taxing robots restrains the expansion of superstar firms, reducing both markup distortions and aggregate productivity. Under our calibration, a modest robot subsidy is welfare-improving relative to the laissez-faire equilibrium. Of course, our model abstracts from many other sources of frictions studied in the literature. Our results imply that, in a more general framework that incorporates those frictions along with the tradeoff between productivity gains and markup distortions highlighted in our model, the optimal size of robot taxes would likely be smaller than what is found in the literature.

⁵Existing studies show that automation can have important implications for employment, wages, and labor productivity (Acemoglu and Restrepo, 2018, 2020; Aghion et al., 2021; Aghion, Antonin, Bunel and Jaravel, 2023; Leduc and Liu, 2024; Firooz, Leduc and Liu, 2024). Automation has also contributed to wage inequality by displacing routine jobs in middle-skill occupations (Autor, Levy and Murnane, 2003; Autor, Dorn and Hanson, 2013; Jaimovich and Siu, 2020; Prettner and Strulik, 2020). There is also evidence that robot adoptions are associated with declines in the labor share (Autor and Salomons, 2018; Acemoglu et al., 2020a; Bergholt, Furlanetto and Maffei-Faccioli, 2022; Jeong, Baek and Peri, 2024).

⁶Thuemmel (2022) finds that a robot subsidy is optimal when robots are relatively expensive; when robots become sufficiently cheap, it would be optimal to tax them.

3. Automation and industry concentration: Some empirical evidence

Figure 1 shows the time-series correlations between industry concentration and automation during the past two decades. We now present some evidence that the correlations are also present in the cross section of manufacturing industries.

3.1. Data and measurement

We use firm-level data from Compustat to compute two measures of industry concentration: the sales share and the employment share of the top 1% of firms in a given industry.⁷

We construct a measure of robot density for each two-digit industry using data on manufacturing employment and operational stock of industrial robots from the International Federation of Robotics (IFR).⁸ We define robot density for industry j in year t as

$$robot_{jt} = \frac{\text{robot stock}_{jt}}{\text{thousands of employees}_{jt}}. \quad (1)$$

For robustness, we also consider an alternative measure of industry-level robot density, defined as the operational stock of robots per million labor hours. The data on industry-level employment (EMP) and labor hours (PRODH) are both obtained from the NBER-CES Manufacturing Industry Dataset.⁹ We obtain an unbalanced panel with 12 industries covering the 12 years from 2007 to 2018.¹⁰

Appendix Table A.2 reports the summary statistics of variables. Robot density varies widely in our sample. For example, the standard deviation of robot density (in log units) is 2.36, which is almost 5 times the mean (0.48). The variations in robot density reflect both within-industry changes in robot adoptions over time and across-industry heterogeneity in the adoption rates and heterogeneous growth rates of robot use. Industry concentration in our sample also displays large variations. For example, the sales share of the top 1% of firms averages about 30 percent, with a standard deviation of about 13 percent. The employment share of the top 1% of firms averages about 27 percent and varies less than the sales share, with a standard deviation of about 8 percent.

3.2. Cross-sectional correlations

In our sample, different industries experienced different growth rates in robot adoptions over time. This pattern can be visualized by the scatter plots in Figure 2, which shows the long differences in industry concentration against those in robot density during the sample periods from 2007 to 2018 in 12 two-digit manufacturing industries. The figure shows that, some industries (such as basic metals and machinery) experienced much larger increases in robot density than others (such as plastics and chemicals or electronics). There are large variations in the long-differences in both sales concentration and employment concentration across the 12 industries. Those industries with greater cumulative changes in robot density also experienced larger cumulative changes in sales concentration (Panel A). In contrast, changes in employment concentration are uncorrelated with those in robot density (Panel B).¹¹ These cross-sectional correlation patterns are in line with the time-series correlations shown in Figure 1.

3.3. Correlations from panel-data regressions

The cross-sectional correlations shown in Figure 2 do not capture industry-level variations in robot density and industry concentration over time and the linear trend line reflects the unconditional cross-sectional correlations. We now

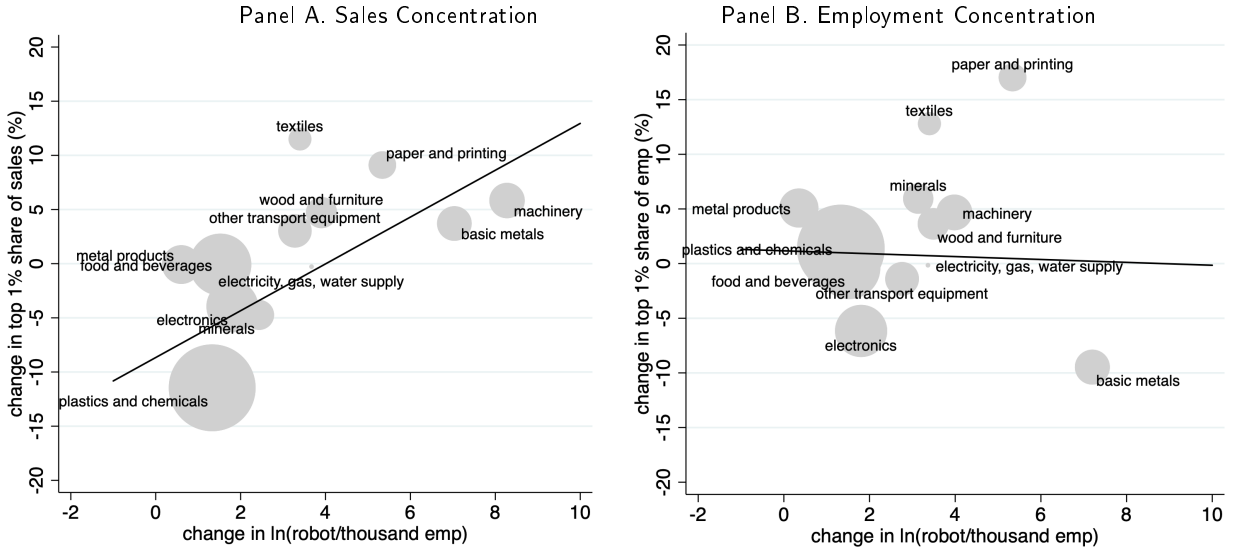
⁷Using a percentile is more appropriate than using a specific number of firms as the cutoff for our sample, given that the total number of public firms in Compustat changes greatly across time. The top 1% of firms is comparable to the top four firms analyzed by Autor et al. (2020), since an average four-digit manufacturing industry has around 364 firms and therefore the top four firms are approximately equivalent to the top 1% of firms.

⁸According to the IFR definition, industrial robots are automatically controlled, reprogrammable, and multipurpose manipulators with several axes.

⁹The IFR uses the International Standard Industrial Classification (ISIC, Rev. 4) for industry classification, while NBER-CES and Compustat use the NAICS classification. We match the ISIC Rev. 4 industry codes with the NAICS2017US codes using the concordance table from the U.S. Census Bureau.

¹⁰We selected 2007 as the starting point due to the limited availability of IFR data on U.S. industrial robots at the two-digit industry level prior to that year. Since we measure industry concentration by the share of the top 1% of firms within an industry, we exclude industries that have fewer than 10 firms in our sample (e.g., industry 29, automotive). Our sample includes 12 industries, identified by their ISIC rev4 codes: 10-12, 13-15, 16&31, 17-18, 19-22, 23, 24, 25, 26-27, 28, 30, D&E. Appendix Table A.1 reports the description of industries in our sample. Note that the sample sizes for some variables are smaller than $12 \times 12 = 144$ because of missing values in some industry-year cells.

¹¹The results are similar when we measure robot density using the operational stock of robots per million hours for each industry (see Figure A.1 in the appendix).

Figure 2: Industry Concentration and Robot Density

Note: This figure shows the cumulative changes in sales concentration (Panel A) and employment concentration (Panel B) against changes in robot density. The industry concentration is measured by the share of the top 1% of firms within an industry. Robot density is measured by the operational stock of industrial robots per thousand workers in each industry. The cumulative change is the long difference between the ending value and the starting value of each variable during the years from 2007 to 2018. Since we have an unbalanced panel, we use the first (last) year with non-missing values as the starting (ending) point for calculating the long differences. The circle size indicates an industry's sales share in the initial year (2007). The line shows the prediction from a linear regression weighted by industries' initial sales shares. The slope coefficient for sales concentration (Panel A) is 0.022 with a standard error of 0.008. The slope coefficient for employment concentration (Panel B) is -0.0013 with a standard error of 0.010.

Source: IFR, NBER-CES, Compustat, and authors' calculation

estimate the correlations between automation and industry concentration based on panel-data regressions, controlling for industry and year fixed effects. Specifically, we estimate the following OLS specification

$$Y_{jt} = \beta \log(robot_{jt}) + \gamma_j + \delta_t + \varepsilon_{jt}, \quad (2)$$

where the dependent variable Y_{jt} is a measure of industry concentration in industry j in year t (sales or employment share of the top 1% of firms), and γ_j and δ_t are industry and year fixed effects, respectively. The key independent variable is the log of robot density $robot_{jt}$. The term ε_{jt} denotes the regression residual. The coefficient of interest, β , measures the semi-elasticity of industry concentration with respect to robot density, controlling for aggregate conditions and other fixed industry characteristics.

Panel A of Table 1 reports the estimation results of the OLS regressions. Industries are weighted by their sales shares in the initial year (i.e., 2007), following the approach by Autor et al. (2020). Standard errors, shown in parentheses, are clustered at the industry level. Since we have a relatively small number of clusters (with 12 industries), we further report the wild cluster bootstrap robust p -value for the estimated coefficients, an approach proposed by Cameron, Gelbach and Miller (2008).¹²

Panel A of Table 1 shows that robot density is positively correlated with sales concentration (i.e., the sales share of the top 1% of firms), with the correlation being statistically significant at the 95 percent confidence level (Columns (1) and (2)). The point estimate in Column (1) implies that, in an industry with robot density (in log units) that is one standard deviation above the average, the sales share of the top 1% of firms is about 5 percentage points, or equivalently about 16 percent, above the sample mean (the average sales share of the top 1% of firms in our sample is about 31%).¹³

¹²We use the "bootstest" Stata command developed by Roodman, Nielsen, MacKinnon and Webb (2019).

¹³The standard deviation of logged robot density is 2.36. The point estimate in Column (1) indicates that a one standard deviation increase in logged robot density implies that the sales share of the top 1% firms increases by $0.021 \times 2.36 \approx 5$ percentage points, or about 16 percent of the mean of the sales share.

Table 1
Regressions of Industry Concentration on Robot Density

Panel A: OLS				
	top 1% share of sales		top 1% share of emp	
	(1)	(2)	(3)	(4)
ln(robot/thousand emp)	0.021** (0.007)		0.002 (0.015)	
ln(robot/million hours)		0.021** (0.007)		0.002 (0.015)
Observations	117	117	104	104
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Wild Bootstrap <i>p</i> -value	0.049	0.051	0.804	0.812
Panel B: IV (second stage)				
	top 1% share of sales		top 1% share of emp	
	(1)	(2)	(3)	(4)
ln(robot/thousand emp)	0.038** (0.019)		0.012 (0.016)	
ln(robot/million hours)		0.036* (0.020)		0.014 (0.016)
Observations	117	117	104	104
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Anderson-Rubin <i>p</i> -value	0.000	0.001	0.474	0.401
Wild Bootstrap <i>p</i> -value	0.019	0.047	0.385	0.306

Note: This table shows OLS and IV (second-stage) regression results from the empirical specification (2). Dependent variables are the sales share (first two columns) and employment share (last two columns) of the top 1% of firms. The industry-level robot density is measured by the operational stock of industrial robots per thousand workers or per million labor hours within the industry. The IV for the U.S. robot density is the one-year lag of the robot density averaged over five European countries (EURO5). In all regressions, industries are weighted by their sales shares in the initial year (2007), and the regressions also control for industry and year fixed effects. Standard errors in parentheses are clustered at the industry level. The last rows in both panels show the *p*-values of the wild cluster bootstrap inferences for the estimated coefficients, an approach proposed by [Cameron et al. \(2008\)](#) for linear regressions with a small number of clusters. The second from the last row in Panel B shows the *p*-values of Anderson-Rubin weak instrument robust tests adjusted for heteroskedasticity. Stars denote the statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The estimated correlation between the hours-based measure of robot density and sales concentration is very similar in magnitude and statistical significance (Column (2)).

In contrast, the correlation of robot density with employment concentration (i.e., the employment share of the top 1% of firms), although positive, is much smaller than that with sales concentration, and the estimated correlations are statistically insignificant (Columns (3) and (4)).¹⁴ These regression results corroborate well with the time-series and cross-sectional patterns illustrated in Figures 1 and 2, respectively.

3.4. An instrumental-variable approach

The correlations between robot density and industry concentration estimated from the OLS may suffer from an omitted variable bias that can arise when a time-varying industry-level factor (such as industry-specific productivity) affects both robot density and concentration in the industry.

To mitigate concerns about omitted-variable biases, we use an IV approach following [Acemoglu and Restrepo \(2020\)](#). Specifically, we use lags of industry-level robot adoptions in five European countries (EURO5) as an IV for robot adoptions in the United States in the same industries. The EURO5 economies include Denmark, Finland, France,

¹⁴There are fewer observations for employment concentration than sales concentration due to a higher occurrence of missing employment data in Compustat.

Italy, and Sweden, which all adopted robotics ahead of the United States, partly driven by their more rapidly aging population.¹⁵

Similar to our measure of robot density for the United States, we measure robot density in the EURO5 economies by the number of robots per thousand employees (or per million of labor hours) in each industry, with the employment (and hours) data taken from EUKLEMS. The average robot density of the five European economies in industry j at time t is calculated as

$$robot_{jt}^{EURO5} = \frac{1}{5} \sum_{k \in EURO5} \frac{\text{robot stock}_{kjt}}{\text{thousands of employees}_{kjt}}, \quad (3)$$

where k is an index of economies in the EURO5 group. We use the one-year-lagged EURO5 robot density as the IV for the U.S. robot density in our industry-level panel regression.

Our two-stage least squares (2SLS) regression specification has one endogenous regressor with one IV, and is thus just-identified. In the first stage, we regress robot density (in log units) at the two-digit industry level in the U.S. on lagged average robot density (also in log units) in the EURO5 group in the corresponding industries, controlling for industry and year fixed effects (see Table A.3 in the Appendix for the first-stage regression results). In the second stage, we regress our measures of U.S. industry concentration on the predicted robot density from the first stage.

Panel B of Table 1 displays the IV estimation results. The estimation shows that sales concentration in the U.S. manufacturing industries is positively and significantly correlated with the robot density predicted from the lagged robot density in the EURO5 economies (Columns (1) and (2)). A one standard deviation increase in the predicted robot density (in log units) is associated with an increase in the sales share of the top 1% firms of about 9 percentage points, or equivalently, about 29 percent relative to its sample average.¹⁶ This number is higher than the 5 percentage points (or 16 percent) obtained from the OLS estimation (Panel A of Table 1), suggesting that omitted variables lead to a downward bias of the coefficient in the OLS regressions. In comparison, the estimated correlations between employment concentration and the predicted robot density are small and statistically insignificant (Columns (3) and (4)). Our evidence therefore implies that automation is associated with the rise in sales concentration as well as the divergence between sales and employment concentration in the manufacturing sector.¹⁷

Although the F-statistics from the first-stage regression suggest that our instrument is weak (see Table A.3), our IV estimation and inferences are robust to potentially weak instruments, as indicated by the Anderson-Rubin (AR) test (Anderson and Rubin, 1949). In the second from the last row of Panel B of Table 1, we report the p -values of the AR test adjusted for heteroskedasticity. The p -values indicate that the estimated correlations of robot density with sales concentration are robust to weak instruments at the 99% confidence level, while those with employment concentration are not significant, with a p -value of the AR test larger than 0.40.¹⁸

The IV regressions help mitigate but not dispel concerns about omitted-variable biases. For example, some common, unobserved shocks might hit industries across the world and drive industry concentration. If so, an industry-level omitted variable for the U.S. might also be an industry-level omitted variable for Europe as well. Such unobserved shocks, if important, could potentially invalidate the exclusion restrictions of our IV.¹⁹ A related concern is that an increase in robot adoptions by European firms could drive out the least competitive firms in the European market, raising U.S. industry concentration through global sales, rather than through U.S. robot adoptions. If this global sales channel is important, then it might also invalidate the exclusion restrictions.²⁰ Given these concerns and the limitations

¹⁵Following Acemoglu and Restrepo (2020), we exclude Germany from our sample because it is far ahead of the other countries in robot adoptions, making it less informative for the U.S. adoption trends than those trends in the EURO5 economies.

¹⁶The estimation shown in Panel B of Table 1 implies that a one standard deviation increase in robot density (i.e., an increase of 2.36 log points) is associated with an increase in the sales share of the top 1% firms by $0.038 \times 2.36 \approx 9$ percentage points. In our sample, the average sales share of the top 1% firms is about 31%. Thus, our estimation suggests that a one standard deviation increase in robot density is associated with an increase in the sales share of the top 1% firms by about 29 percent relative to its sample average.

¹⁷With our small sample (covering 12 industries over 12 years), the estimation results are relatively noisy, and the estimated coefficients in Columns (1) and (3) (or those in Columns (2) and (4)) in Panel B of Table 1 are not statistically different from each other.

¹⁸The AR test is one of the most powerful tests for the null hypothesis in the second stage when the model is just-identified, regardless of the instrument's strength (Moreira, 2009; Andrews, Stock and Sun, 2019).

¹⁹Acemoglu and Restrepo (2020) provide a rebuttal to this concern. They argue that the trend of European robot adoptions is largely driven by its own demographic trends (aging population) and it is uncorrelated with other major global trends, such as import competition, offshoring, declines of routine jobs, and capital deepening.

²⁰The share of sales of U.S. affiliates in the European markets has been small relative to the total sales of the U.S. parent companies. For example, according to the Bureau of Economic Analysis, the sales of majority-owned U.S. affiliates in EURO5 were \$474 billion in 2020, about 3.4% of the total sales of their U.S. parent companies (\$13.85 trillion).

of the robot data, the IV estimation should not be interpreted as direct evidence of causal effects of robot adoptions on industry concentration.

Nonetheless, the empirical linkage between automation and industry concentration seems quite robust. The rise in automation has been accompanied by a rise of superstar firms, characterized by an increase in sales concentration and a divergence between sales and employment concentrations.

4. The Model

To understand the empirical link between automation and industry concentration, we construct a dynamic general equilibrium model featuring heterogeneous firms, variable markups, and endogenous automation decisions.

4.1. Households

The economy is populated by a continuum of identical, infinitely lived households of a unit measure. All agents have perfect foresight. The representative household has the utility function

$$\sum_{t=0}^{\infty} \beta^t \left[\ln C_t - \chi \frac{N_t^{1+\xi}}{1+\xi} \right], \quad (4)$$

where C_t denotes consumption at time t , N_t denotes labor supply, $\beta \in (0, 1)$ is a subjective discount factor, $\xi \geq 0$ is the inverse Frisch elasticity of labor supply, and $\chi > 0$ is the weight on the disutility from working.

The household faces the sequence of budget constraints

$$C_t + Q_{a,t} I_{a,t} \leq W_t N_t + r_{a,t} A_t + \pi_t, \quad (5)$$

where $I_{a,t}$ denotes the investment in robots, A_t denotes the beginning-of-period robot stock, $r_{a,t}$ is the real rental rate of robots, $Q_{a,t}$ denotes the relative price of robots, W_t denotes the real wage rate, and π_t denotes profits from the firms that the household owns. The stock of robots evolves according to the law of motion

$$A_{t+1} = (1 - \delta_a) A_t + I_{a,t}, \quad (6)$$

where $\delta_a \in [0, 1]$ denotes the robot depreciation rate.

The household takes the prices $Q_{a,t}$, W_t , and $r_{a,t}$ as given, and maximizes the utility function (4) subject to the budget constraints (5). The optimizing consumption-leisure choice implies the labor supply equation

$$W_t = \chi N_t^\xi C_t. \quad (7)$$

The optimizing choice of robot accumulation implies that

$$Q_{a,t} = \rho_{t,t+1} [r_{a,t+1} + Q_{a,t+1} (1 - \delta_a)], \quad (8)$$

where $\rho_{t,t+1} \equiv \beta \frac{C_t}{C_{t+1}}$ is the stochastic discount factor (SDF).

4.2. Final goods producers

There is a continuum of monopolistically competitive intermediate producers indexed by $j \in [0, 1]$. Final goods producers make a composite homogeneous good out of the intermediate varieties and sell it to consumers in a perfectly competitive market, with the final goods price normalized to one. The final good Y_t is produced using a bundle of intermediate goods $y_t(j)$, according to the Kimball aggregator

$$\int_0^1 \Lambda \left(\frac{y_t(j)}{Y_t} \right) dj = 1. \quad (9)$$

4.3. Demand for intermediate goods

Denote the relative output of firm j by $q_t(j) := \frac{y_t(j)}{Y_t}$. Taking the intermediate goods price $p_t(j)$ as given, the cost-minimizing decision of final good producers leads to the following demand schedule for intermediate good j

$$p_t(j) = \Lambda'(q_t(j))D_t, \quad (10)$$

where D_t is a demand shifter given by

$$D_t = \left(\int \Lambda'(q_t(j))q_t(j) dj \right)^{-1}. \quad (11)$$

We follow [Klenow and Willis \(2016\)](#) and assume that²¹

$$\Lambda(q_t) = 1 + (\sigma - 1) \exp\left(\frac{1}{\varepsilon} \varepsilon^{\frac{\sigma}{\varepsilon} - 1} \left[\Gamma\left(\frac{\sigma}{\varepsilon}, \frac{1}{\varepsilon}\right) - \Gamma\left(\frac{\sigma}{\varepsilon}, \frac{q_t^{\varepsilon/\sigma}}{\varepsilon}\right) \right]\right), \quad (12)$$

with $\sigma > 1$, $\varepsilon \geq 0$, and $\Gamma(s, x)$ denoting the upper incomplete Gamma function

$$\Gamma(s, x) = \int_x^\infty v^{s-1} e^{-v} dv. \quad (13)$$

Under the specification (12), we obtain

$$\Lambda'(q_t(j)) = \frac{\sigma - 1}{\sigma} \exp\left(\frac{1 - q_t(j)^{\frac{\varepsilon}{\sigma}}}{\varepsilon}\right), \quad (14)$$

which, using the demand schedule (10), implies that the demand elasticity (i.e., price elasticity of demand) faced by firm j is

$$\sigma(q_t(j)) = -\frac{\Lambda'(q_t(j))}{\Lambda''(q_t(j))q_t(j)} = \sigma q_t(j)^{-\frac{\varepsilon}{\sigma}}. \quad (15)$$

Given this demand elasticity, the firm with relative production $q_t(j)$ charges the optimal markup

$$\mu_t(j) = \frac{\sigma(q_t(j))}{\sigma(q_t(j)) - 1}. \quad (16)$$

As a result, larger firms face lower demand elasticities, have more market power, and charge higher markups.²²

4.4. Intermediate goods producers

Intermediate producers, from now on indexed by their idiosyncratic productivity ϕ_t , produce differentiated intermediate goods using one of two technologies: one with labor as the only input, and the other with both labor and robots as input factors. If the firm uses robots in production, it faces a per-period fixed cost which is realized after drawing the productivity ϕ_t , to be elaborated below. The production function takes the CES form

$$y_t(\phi_t) = \phi_t \left[\alpha_a A_t(\phi_t)^{\frac{\eta-1}{\eta}} + (1 - \alpha_a) N_t(\phi_t)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (17)$$

where $y_t(\phi_t)$ denotes the firm's output, $N_t(\phi_t)$ denotes the number of workers, and $A_t(\phi_t) \geq 0$ denotes the number of robots. The labor-only technology corresponds to the special case with $A_t(\phi_t) = 0$. The parameter $\eta > 1$ is the elasticity of substitution between robots and workers. The parameter α_a measures the relative importance of robot input in production.

²¹[Aruoba, Oue, Saffie and Willis \(2024\)](#) employ a different Kimball aggregator and show that their calibrated model is consistent with the markup distribution in the U.S.

²²We make the technical assumption that $q_t(j) < \sigma^{\frac{\sigma}{\varepsilon}}$ such that the effective demand elasticity is always greater than one. This assumption ensures a well-defined equilibrium under monopolistic competition. In our quantitative analysis, we find that this constraint is never binding.

The idiosyncratic productivity shock follows a stationary AR(1) process

$$\ln \phi_{t+1} = \gamma \ln \phi_t + \varepsilon_{t+1}, \quad \varepsilon_t \sim N(0, \sigma_\phi^2), \quad (18)$$

where $\gamma \in (0, 1)$ measures the persistence of the productivity shock and $\sigma_\phi > 0$ denotes the standard deviation of the innovation.

We assume that to use robots in production, firms face a per-period fixed cost that is proportional to their productivity. Specifically, a firm with productivity ϕ_t draws s_t from the *i.i.d.* distribution $F(\cdot)$ and needs to pay the per-period cost $s_t \phi_t$ if it uses robots in production.²³ We further assume that the distributions of s_t and ϕ_t are independent. A firm with the realized productivity ϕ_t that draws a fixed cost s_t chooses price p_t and quantity y_t of its differentiated product, labor input N_t , and robot input A_t to solve the dynamic programming problem

$$V_t(\phi_t; s_t) = \max_{p_t, y_t, N_t, A_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) - s_t \phi_t \mathbb{1}\{A_t(\phi_t) > 0\} + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \right], \quad (19)$$

where $\mathbb{1}\{x\}$ equals one if x holds and zero otherwise. The firm hires workers at the competitive real wage rate W_t and rents robots at the competitive rental rate $r_{a,t}$.

The firm solves the recursive problem (19) subject to the production function (17) and the demand schedule (10). Since robot operation incurs a fixed cost, a firm facing a sufficiently high s_t relative to its productivity would choose to use no robots by setting $A_t(\phi_t) = 0$.

Appendix B shows that the recursive problem (19) can be simplified to

$$V_t(\phi_t; s_t) = \max\{V_t^a(\phi_t) - s_t \phi_t, V_t^n(\phi_t)\}, \quad (20)$$

where the continuation value of operating the automation technology (i.e., having $A_t(\phi_t) > 0$) is given by

$$V_t^a(\phi_t) = \max_{p_t, y_t, N_t, A_t > 0} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) \right] + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}), \quad (21)$$

and the continuation value of operating the labor-only technology is given by

$$V_t^n(\phi_t) = \max_{p_t, y_t, N_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) \right] + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}). \quad (22)$$

Firms with automation technology in (21) optimally choose their production inputs $N_t(\phi_t)$ and $A_t(\phi_t)$ given their production $y_t(\phi_t)$. The first-order conditions for cost-minimizing imply the conditional factor demand functions

$$r_{a,t} = \alpha_a \lambda_t^a(\phi_t) \phi_t^{\frac{\eta-1}{\eta}} \left(\frac{y_t(\phi_t)}{A_t(\phi_t)} \right)^{\frac{1}{\eta}}, \quad (23)$$

$$W_t = (1 - \alpha_a) \lambda_t^a(\phi_t) \phi_t^{\frac{\eta-1}{\eta}} \left(\frac{y_t(\phi_t)}{N_t(\phi_t)} \right)^{\frac{1}{\eta}}, \quad (24)$$

where $\lambda_t^a(\phi_t)$ denotes the marginal cost of production for a firm with productivity ϕ_t operating the automation technology:

$$\lambda_t^a(\phi_t) = \frac{\left[\alpha_a^n r_{a,t}^{1-\eta} + (1 - \alpha_a)^\eta W_t^{1-\eta} \right]^{\frac{1}{1-\eta}}}{\phi_t}. \quad (25)$$

²³ Assuming that the fixed costs of automation are proportional to firm-level productivity captures the fact that large firms face higher fixed costs in production, which improves the model calibration as discussed later. However, our qualitative results remain valid even if fixed costs are not assumed to be proportional to productivity.

Moreover, firms operating the labor-only technology in (22) choose their labor input $N_t(\phi_t)$ given their production $y_t(\phi_t)$:

$$N_t(\phi_t) = \frac{(1 - \alpha_a)^{\frac{\eta}{1-\eta}} y_t(\phi_t)}{\phi_t}. \quad (26)$$

The marginal cost of production in this case would be

$$\lambda_t^n(\phi_t) = \frac{(1 - \alpha_a)^{\frac{\eta}{1-\eta}} W_t}{\phi_t}. \quad (27)$$

Notice that, given the productivity ϕ_t , the marginal cost of production using the labor-only technology is always larger than that using the automation technology, i.e., $\lambda_t^a(\phi_t) \leq \lambda_t^n(\phi_t)$.

The problem (20) implies that firms choose to operate the automation technology (i.e., to have $A_t(\phi_t) > 0$) if and only if their draw of the fixed automation cost is small enough:

$$s_t \leq s_t^*(\phi_t) \iff \mathbb{1}_t^a(\phi_t; s_t) = 1, \quad (28)$$

where $\mathbb{1}_t^a(\cdot)$ is an indicator of the automation decision, which is a function of the firm-level variables ϕ_t and s_t , and the cutoff fixed cost equals:

$$s_t^*(\phi_t) \equiv \frac{V_t^a(\phi_t) - V_t^n(\phi_t)}{\phi_t}. \quad (29)$$

It follows that, for a firm with productivity ϕ_t , the ex ante (i.e., before drawing the automation fixed cost) automation probability equals $F(s_t^*(\phi_t))$, which is the cumulative density of the fixed cost distribution evaluated at the indifference point.

Appendix B proves that the automation cutoff can be written as the difference between the flow profit from operating the automation technology versus that from employing the labor-only technology. In other words,

$$s_t^*(\phi_t) = \frac{\pi_t^a(\phi_t) - \pi_t^n(\phi_t)}{\phi_t}, \quad (30)$$

where

$$\pi_t^a(\phi_t) = \max_{p_t, y_t, N_t, A_t} \left[p_t(\phi_t) y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) \right], \quad (31)$$

subject to the demand schedule (10) and production function (17), and

$$\pi_t^n(\phi) = \max_{p_t, y_t, N_t} \left[p_t(\phi_t) y_t(\phi_t) - W_t N_t(\phi_t) \right], \quad (32)$$


subject to the same demand schedule and production function with $A_t(\phi_t) = 0$.

4.5. Equilibrium

The world robot price $Q_{a,t}$ is exogenously given. The equilibrium consists of aggregate allocations C_t , $I_{a,t}$, A_t , N_t , and Y_t , the wage rate W_t , the rental rate $r_{a,t}$, firm-level allocations $A_t(\phi_t)$, $N_t(\phi_t)$, and $y_t(\phi_t)$, and firm-level prices $p_t(\phi_t)$ for all $\phi_t \in G(\cdot)$, where $G(\cdot)$ denotes the ergodic distribution implied by the productivity process (18), such that (i) taking W_t and $r_{a,t}$ as given, the aggregate allocations C_t , N_t , A_t solve the representative household's optimization problem; (ii) taking W_t , $r_{a,t}$, and Y_t as given, the firm-level allocations and prices solve each individual firm's optimization problem; and (iii) the markets for the final good and labor clear.

The final goods market clearing condition is given by

$$C_t + Q_{a,t} I_{a,t} + \int_{\phi_t} \int_0^{s_t^*(\phi_t)} s_t \phi_t dF(s_t) dG(\phi_t) = Y_t. \quad (33)$$



Note: This figure shows the automation decisions as a function of firm-level productivity (ϕ) and the fixed cost of operating the automation technology (s). Firms with (ϕ, s) to the lower-right of the solid line choose to automate (the shaded area) and those to the upper-left of the line choose to use the labor-only technology. A decline in the robot price shifts the indifference line upward (from the solid to the dashed line), inducing more use of the automation technology.

The labor market clearing condition is given by

$$N_t = \int_{\phi_t} N_t(\phi_t) dG(\phi_t). \quad (34)$$

The stock of robots is given by

$$A_t = \int_{\phi_t} A_t(\phi_t) F(s_t^*(\phi_t)) dG(\phi_t). \quad (35)$$

Total investment in robots equals

$$I_{a,t} = A_{t+1} - (1 - \delta_a)A_t. \quad (36)$$

5. Model mechanism

In our model, firms are heterogeneous along two dimensions: they face idiosyncratic shocks to both productivity (ϕ) and the fixed cost of operating the automation technology (s). The automation decision depends on the combination of the realizations of ϕ and s . Firms face a tradeoff when deciding whether to automate. On the one hand, firms need to pay a fixed cost $s\phi$ to automate. On the other hand, the marginal cost of production using the automation technology (equation (25)) is always lower than that using the labor-only technology (equation (27)). Since higher-productivity firms are larger and charge higher markups, they earn higher profits and therefore are more likely to pay the fixed cost and automate.

Figure 3 illustrates the automation decision rule. For any given productivity ϕ , a firm will choose to automate if the realized fixed cost is sufficiently low. Similarly, for any given fixed cost s , a firm will automate if the realized productivity is sufficiently high. There is an upward-sloping line that separates the technology choices. To the right of the line (high ϕ or low s), firms use the automation technology and to the left of the line, they use the labor-only technology. Firms with combinations of ϕ and s on the upward-sloping line are indifferent between the two types of technologies, where the indifference line is given by the $s^*(\phi)$ function in Equation (30).

The location of the indifference line is endogenous, depending on aggregate economic conditions. A decline in the relative price of robots (Q_a), for example, will reduce the marginal cost of using the automation technology since

the robot rental rate declines (Equation (8) implies that r_a increases with Q_a in the steady state). This would shift the indifference line upward (from the solid to the dashed line), such that some new firms would adopt robots (the extensive margin) and those firms already using robots would increase their demand for robots (the intensive margin).

For a given technology choice (labor-only or automation), a high-productivity firm is also a large firm in terms of both employment and output. Moreover, high-productivity firms are also more likely to use robots at any given fixed cost, as illustrated in Figure 3. A decline in the relative price of robots improves labor productivity, enabling those robot-using firms to become even larger and increasing the share of top firms in the product market (through the intensive margin). However, the decline in robot price also induces some less-productive firms to switch from the labor-only technology to the automation technology (the extensive margin), partially offsetting the increase in the sales share of the top firms. The net effect of the decline in the robot price on sales concentration can be ambiguous, depending on the relative strength of the extensive versus the intensive margin effects. As we will show below, under our calibration, the intensive margin effect dominates, such that a lower robot price leads to a higher concentration of sales in large firms. This model prediction is consistent with the empirical evidence presented in Section 3.

An increase in the sales share of large firms following a decline in the robot price does not directly translate into an increase in the employment share of those firms. Since robots substitute for workers, large robot-using firms can increase production without proportional increases in labor input. Additionally, as these firms grow, they tend to charge higher markups. Thus, the share of employment of large firms increases by less than their sales share. This is the key model mechanism through which automation can contribute to the increase in sales concentration as well as the divergence between sales concentration and employment concentration.

6. Calibration

To assess the quantitative importance of the automation mechanism for explaining the observed rise in sales concentration and the divergence between sales and employment concentration in the U.S. manufacturing sector, we calibrate the steady state equilibrium of the model to match moments in the manufacturing data. Appendix C.1 outlines the computational algorithm to solve for the steady-state equilibrium. We focus on the manufacturing sector for two reasons. First, automation is more prevalent in the manufacturing sector than in the whole economy. According to the 2019 ABS, about 8.7% of manufacturing firms use robots and those firms employ about 45.1% of manufacturing workers. In comparison, in the whole economy, only about 2% of firms use robots and they employ about 15.7% of workers (Acemoglu et al., 2022).²⁴ Second, the increase in sales concentration in the manufacturing sector was accompanied by a divergence between sales and employment concentration in the past two decades (see Figure 1 and Autor et al., 2020).

Table 2 displays the calibrated parameters. We calibrate a subset of the parameters based on external sources in the literature (Panel A) and the remaining parameters by matching moments in the data (Panel B).

One period in the model corresponds to a quarter of a year. We set the subjective discount factor to $\beta = 0.99$, implying an annual real interest rate of 4%. We set the inverse Frisch elasticity to $\xi = 0.5$, following Rogerson and Wallenius (2009). We normalize the disutility from working to $\chi = 1$. We calibrate the quarterly robot depreciation rate to $\delta_a = 0.02$, implying an average robot lifespan of about 12 years, in line with the assumption made by the IFR in imputing the operational stock of industrial robots. We set the persistence of idiosyncratic productivity shocks to $\gamma = 0.95$ following Bloom, Floetotto, Jaimovich, Saporta-Eksten and Terry (2018). We set the standard deviation of productivity shocks to $\sigma_\phi = 0.1$, according to the estimation by Bloom et al. (2018).²⁵ To calibrate the elasticity parameters σ and ϵ in the Kimball aggregator, we follow Edmond, Midrigan and Xu (2021) and set $\sigma = 10.86$ and $\epsilon/\sigma = 0.16$.

We calibrate the remaining parameters to match several key moments in the micro-level data. We assume that the fixed cost of automation follows a log-normal distribution $\ln(s) \sim \mathcal{N}(0, \sigma_a^2)$, where σ_a is the standard deviation. The four parameters to be calibrated include the relative price of robots Q_a , the standard deviation of the fixed cost of

²⁴The 2019 ABS also shows that, within the manufacturing sector, the usage of robots was also more prevalent than that of other advanced technologies such as AI.

²⁵Bloom et al. (2018) estimate a two-state Markov switching process of firm-level volatility. They find that the low standard deviation is 0.051 and the high value is 0.209. Their estimated transition probabilities suggest that the unconditional probability of the low standard deviation is 68.7%. Therefore, the average standard deviation is 0.1 ($=0.051*68.7\%+0.209*(1-68.7\%)$).

Table 2
Calibrated Parameters and Matched Moments

Parameter	Notation	Value	Sources/Matched Moments
Panel A: Parameters calibrated to match external sources			
Discount factor	β	0.99	4% annual interest rate
Inverse Frisch elasticity	ξ	0.5	Rogerson and Wallenius (2009)
Working disutility weight	χ	1	Normalization
Robot depreciation rate	δ_a	0.02	8% annual depreciation rate
Productivity persistence	γ	0.95	Bloom et al. (2018)
Productivity standard dev.	σ_ϕ	0.1	Bloom et al. (2018)
Demand elasticity parameter	σ	10.86	Edmond et al. (2021)
Super elasticity	ϵ/σ	0.16	Edmond et al. (2021)
Panel B: Parameters calibrated to match moments in data			
Relative price of robots	Q_a	48.90	Fraction of automating firms
SD of log automation fixed costs	σ_a	3.38	Employment share of automating firms
Robot input weight	α_a	0.37	Robot density
Elasticity of substitution	η	2.03	Growth rate of robot density
Panel C: Matched Moments			
Moments	Data	Model	
Fraction of automating firms	8.7%	8.7%	
Employment share of automating firms	45.1%	45.1%	
Robot density	0.02	0.02	
Growth rate of robot density	300%	300%	

Note: This table presents the calibrated parameters and matched moments in the model. Panel A reports the externally calibrated parameters and their sources. Panel B presents the parameters calibrated by moment matching. Panel C reports the targeted data moments and the simulated moments by the model. The first two data moments are based on the ABS data (taken from Acemoglu et al., 2022) and the last two moments are authors' calculations using IFR and NBER-CES data.

automation σ_a , the robot input weight α_a , and the elasticity of substitution between robots and labor η . The calibrated parameters are shown in Panel B of Table 2.²⁶

We target four moments to jointly calibrate these four parameters. The four moments include (i) the share of manufacturing firms that use robotics was 8.7% during the period of 2016-2018, according to the 2019 ABS (Acemoglu et al., 2022); (ii) the employment share of the manufacturing firms that use robotics was 45.1% during the same period, also according to the 2019 ABS; (iii) the robot density measured by the aggregate operational stock of industrial robots per thousand manufacturing workers was about 20 in 2016, according to the data from the IFR and NBER-CES; and (iv) during the period from 2002 to 2016, the robot density increased by 300% while the relative price of robots declined by 40%.

These four moments in the data help pin down the four parameters in our model. Intuitively, the relative price of robots Q_a affects the fraction of firms that use the automation technology (i.e., the automation probability), which is given by $\int_\phi F(s^*(\phi)) dG(\phi)$. The parameter σ_a governs the skewness of the distribution of automation fixed costs, which in turn determines the skewness of automation decisions across the firm size distribution. Under a smaller σ_a ,

²⁶While we do not have an additional moment in the manufacturing sector to calibrate the mean of the automation fixed cost distribution, we examine the robustness of our quantitative results in Appendix D by calibrating this additional parameter to match a data moment in the whole economy. In particular, we calibrate the mean fixed cost of automation to match the ratio of the robot use rate among firms between the 50th and 75th percentile of the employment distribution (1.7%) to the average robot use rate among all firms in the whole economy (2%), taken from the 2019 ABS documented by Acemoglu et al. (2022). This data moment is available for the whole economy but not for the manufacturing sector. We find that our main quantitative results are robust to calibrating this additional parameter.

small firms would be less likely to cover the fixed cost of automation. As a result, the employment-weighted robot use rate would be larger. Therefore, to calibrate σ_a , we target the employment share of firms that use the automation technology, which in our model equals

$$\frac{\int_{\phi} F(s^*(\phi))N(\phi) dG(\phi)}{\int_{\phi} N(\phi) dG(\phi)}. \quad (37)$$

The robot input weight α_a in the production function of intermediate goods determines the steady-state level of robot density (i.e., A/N), which equals 0.02 in 2016 in our data (or equivalently, 20 robots per thousand workers). The elasticity of substitution η between robot input and labor input determines the changes in robot density in response to changes in the robot price. We calibrate the elasticity of substitution η by matching the cumulative increase of A/N of 300 percent associated with the cumulative decline in Q_a of 40 percent from 2002 to 2016 in our data.

Panel B of Table 2 reports these parameters that we internally calibrated. The calibrated model matches the targeted data moments exactly, as shown in Panel C of Table 2.²⁷

7. Model implications

We study the quantitative implications of our calibrated model, focusing on the steady state equilibrium. We also examine the model's transition dynamics following an exogenous decline in the robot price.

7.1. Firm-level implications in the steady state

Figure 4 shows firms' decision rules in the steady state as a function of the idiosyncratic productivity level ϕ , in both the benchmark model with calibrated parameters (black solid line) and a counterfactual scenario with a lower robot price (red dashed line). The figure shows that the automation probability increases with productivity (top left panel). Although revenues and fixed costs of automation both scale with productivity, revenues increase with productivity more than proportionately under our calibration, such that more productive firms are more likely to pay the fixed costs of operating the automation technology. In addition, firms with sufficiently low productivity do not use robots and operate the labor-only technology. A decline in the robot price boosts the automation probabilities, with a larger effect on more productive firms. It also reduces the productivity cutoff for accessing the automation technology.

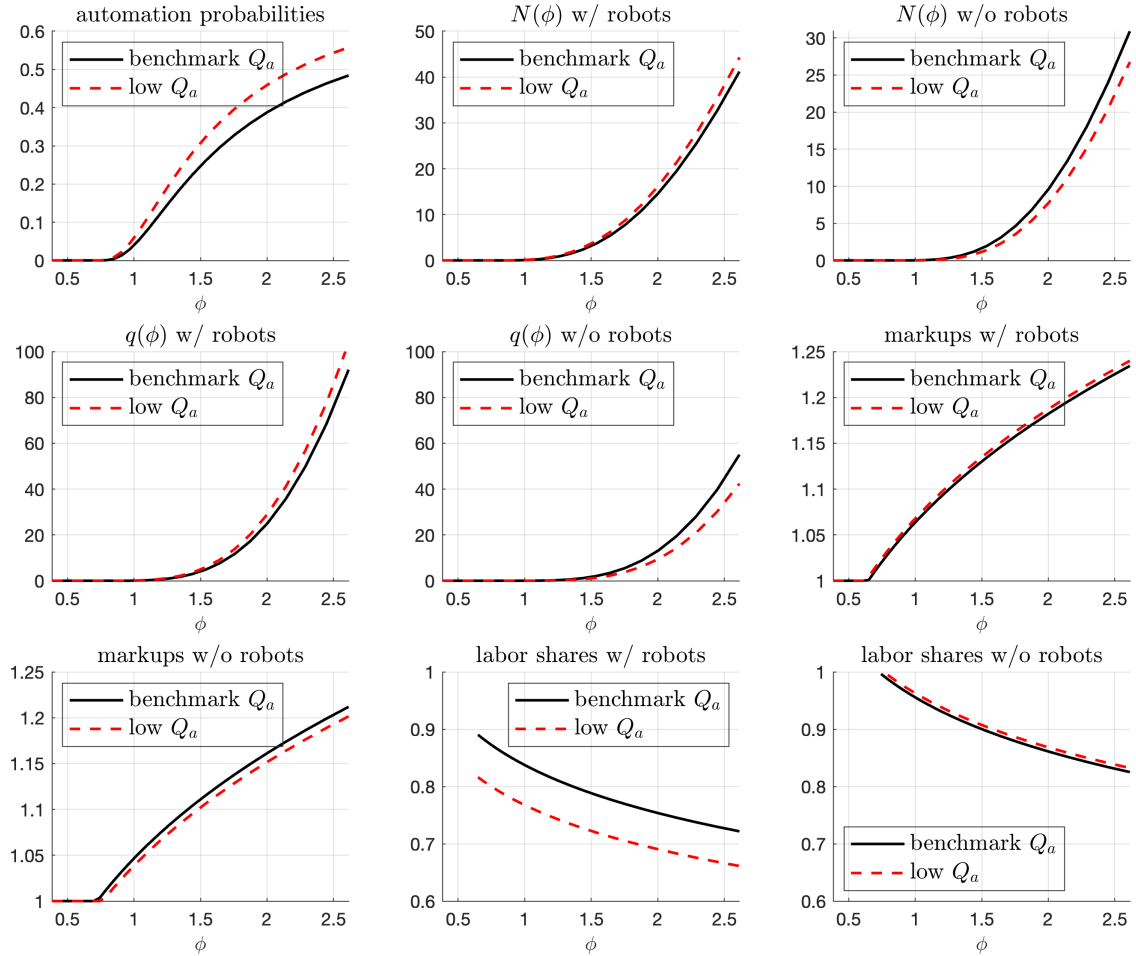
The figure also shows the decision rules for firms that use robots and those that don't at each level of productivity. In the benchmark model, the decision rules are qualitatively similar between the two types of firms. In particular, higher-productivity firms are larger, with higher employment ($N(\phi)$), higher relative output ($q(\phi)$), higher markups, and lower labor shares. Larger firms have lower labor shares for two reasons. First, these firms charge higher markups, reducing the share of labor compensation in value-added. This force is at play for all firms, regardless of whether they use robots. Second, larger firms are more likely to automate and, as a result, have lower labor shares. This effect works only for the firms that operate the automation technology.

The red dashed lines in Figure 4 further show that the impacts of a decline in the robot price on the firms' decision rules depend on whether firms use robots. For robot-using firms, a decline in the robot price raises employment, output, and markup at each level of productivity. A reduction in robot price activates two competing forces on the employment of the automating firms. On the one hand, robots substitute for workers, reducing employment by automating firms. On the other hand, however, robots boosts labor productivity, raising labor demand by those firms. Under our calibration, the latter effect dominates such that automating firms increase employment following a reduction in the robot price. This result is in line with firm-level evidence documented by [Zolas et al. \(2020\)](#) and [Aghion et al. \(2021\)](#). The labor shares of automating firms decline despite the increases in their employment, reflecting the substitution of robots for workers and also the increased markups.

For firms without robots, the decline in the robot price has the opposite effect on their decision rules. In particular, a decline in Q_a reduces employment, output, and markups, and increases the labor share at any given level of productivity. These changes in the decision rules reflect the reallocation of labor from non-automating firms to automating firms. As non-automating firms become smaller, their market power declines, resulting in lower markups and higher labor shares.

²⁷To put the calibrated elasticity η into context, we note that [Cheng, Drozd, Giri, Taschereau-Dumouchel and Xia \(2021\)](#) estimate the firm-level elasticity of substitution between labor and automation capital in China ranging from 3 to 4.5, with their preferred estimate being 3.8. Therefore, our calibrated elasticity of $\eta = 2.03$ is conservative relative to their estimates. We show that if we instead use a higher η in the range estimated by [Cheng et al. \(2021\)](#) the quantitative importance of automation in our model would be larger.

Figure 4: Firms' Decision Rules



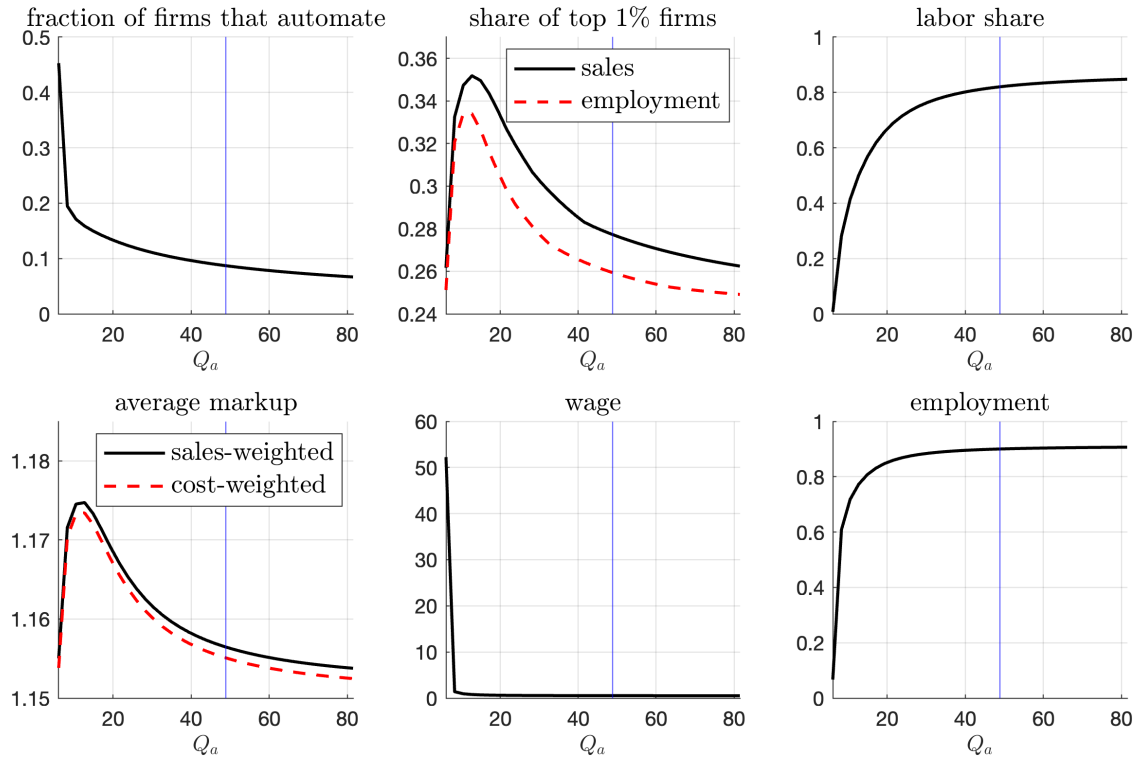
Note: This figure shows firms' decision rules for the firms that automate (w/ robots) and those that do not automate (w/o robots). The solid-black lines are associated with our benchmark calibration, whereas red-dashed lines show the results for a counterfactual in which robot price Q_a falls by 40%.

7.2. Aggregate implications in steady state

The heterogeneous automation decisions and the consequent between-firm reallocation have important implications for the steady-state relations between aggregate variables and the robot price, as shown in Figure 5. For illustration, we consider a wide range of the robot price around the calibrated value of $Q_a = 48.90$, indicated by the vertical blue line in the figure.

At a lower robot price, more firms find it profitable to automate, raising the fraction of automating firms. Given the fixed cost of operating the automation technology, larger firms are more likely to automate and thus they benefit more from the lower robot price.²⁸ As a result, the product market becomes more concentrated and the share of the top 1% of firms rises. Importantly, the sales share of the top firms rises more than their employment share as Q_a declines, because those top firms that use robots can expand production without proportional increases in their labor input, and also because they charge higher markups; while an increase in markups shows up in the sales share of top firms, it is not reflected in their employment share.

²⁸As discussed before, while automation fixed costs are proportional to firm-level productivity, revenues scale with productivity more than proportionately under our calibration, such that more productive firms are more likely to automate, as shown in the top-left panel in Figure 4.

Figure 5: Aggregate Variables in the Steady State

Note: This figure shows the effects of counterfactual changes in the robot price Q_a on the fraction of firms that automate, the share of the top 1% of firms, the labor share, the average markup, the wage rate, and employment in the steady state. The vertical blue line indicates the calibrated value of robot price Q_a .

As Q_a falls, large firms become even larger, raising the average markup in the economy (both sales- and cost-weighted).²⁹ As Figure 4 shows, a reduction in Q_a lowers labor shares in automating firms. It also reallocates production to automating firms that have on average lower labor shares than non-automating firms. Although labor shares in non-automating firms rise, the fall in labor shares in automating firms combined with the reallocation toward those firms contributed to the decline in the labor share in the aggregate economy. Our model thus incorporates the between-firm reallocation channel for explaining the declines in the aggregate labor share and the increases in the average markup, in line with the empirical evidence in Autor et al. (2020) and Acemoglu et al. (2020a).

Changes in robot prices affect employment and wages through various channels. A reduction in Q_a tends to reduce aggregate employment because production is reallocated to automating firms from the labor-intensive non-automating firms. The decline in Q_a raises equilibrium wages because it improves labor productivity in automating firms, raising aggregate labor demand and bidding up real wages. When automating firms expand production, however, they gain market power and their markups rise, thereby mitigating the increase in labor demand and dampening the increase in wages. The reduction in Q_a also creates a positive wealth effect: by raising consumption, the household is willing to supply less labor at each given wage level. In equilibrium, small reductions in Q_a have limited effects on employment and wages, while a large reduction in robot prices leads to an increase in wages and a decline in aggregate employment.³⁰

²⁹To derive the cost-weighted average markup, we use total variable costs at each firm, as in Edmond et al. (2021).

³⁰Our model's prediction that a reduction in the robot price raises worker wages seems to be at odds with the empirical evidence documented by Acemoglu and Restrepo (2021), who find substantial declines in the relative wages of workers specialized in routine tasks in industries experiencing rapid automation. This is perhaps not surprising because we focus on studying the relation between automation and industry concentration and abstract from labor market frictions in our model. In a model with elaborated labor market frictions, such as the business cycle model with labor search frictions and automation studied by Leduc and Liu (2024), an increase in automation threat effectively reduces workers' bargaining power in wage negotiations, and it can lower equilibrium wages. Incorporating labor market frictions into our framework is potentially important for

Automation and industry concentration. The top-middle panel in Figure 5 reports the relations between the robot price Q_a and industry concentration measured by the share of the top 1% of firms in sales (solid line) and in employment (dashed line). This graph helps us examine the quantitative importance of our automation mechanism in explaining the rise in sales concentration as well as the divergence between sales and employment concentration. In particular, we focus on a fall in the robot price from 81.49 to its calibrated value of 48.90, representing a 40% decline that captures the observed magnitude of changes in the relative price of robots in the data over the period from 2002 to 2016, as shown in Figure 1.³¹ We then examine the extent to which the resulting changes in industry concentration in the model can account for observed changes in the data.

As this figure shows, this decline in Q_a leads to the sales share of the top 1% of firms to rise by about 1.48 percentage points (from 26.24% to 27.72%). The employment share of the top 1% of firms also rises but with a smaller magnitude (1.02 percentage points). Thus, the gap between sales concentration and employment concentration widens by about 0.46 percentage points.

In the data, as documented by Autor et al. (2020), sales concentration in manufacturing measured by the sales share of the top four firms (i.e., CR4) rose from about 40.52% in 1997 to 43.32% in 2012, an increase of about three percentage points, while employment concentration rose from 33.26% to 34.51% during the same period, an increase of about 1.2 percentage points.³² The gap between sales and employment concentration during this period in the data therefore widens by about 1.8 percentage points. Our model can explain roughly 49.2% (1.48 out of 3 percentage points) of the increases in sales concentration as well as about 25.3% (0.46 out of the 1.8 percentage points) of the observed divergence between sales and employment concentration.

The top-middle panel in Figure 5 also illustrates that the relation between robot prices and industry concentration can be non-monotonic. If the economy starts with a small share of automating firms in the initial equilibrium, a reduction in the robot price would increase industry concentration, as we find in the calibrated model here. This is consistent with the positive correlation between automation and sales concentration in the U.S. that we documented in Section 3. However, in an economy with widespread automation (i.e., an economy with a sufficiently low level of the robot price), a further reduction in the robot price may not increase industry concentration as much, and it could even reduce concentration. As the automation technology becomes accessible to smaller firms, the share of top firms in the economy falls. These findings suggest that automation is different from general capital equipment. While equipment is widespread across firms in the economy, automation is highly skewed toward a small fraction of superstar firms.³³

7.3. Examining key features of the model

We now examine the importance of variable markups and fixed costs of automation—two important features of our model—for generating the model's main predictions. We show that the fixed cost of automation is crucial for generating the link between automation and industry concentration. However, the Kimball demand system with variable markups is not essential, although quantitatively helpful, since an alternative setup with a CES demand system and thus constant markups could generate qualitatively the same predictions, provided that we keep the fixed cost of automation in the model.

To examine the importance of the fixed cost, we consider a counterfactual economy with the Kimball demand system but without fixed costs of automation. Figure A.3 in the Appendix shows the steady-state relations between the robot price and the macroeconomic variables. In an economy with a lower robot price, firms use robots more intensively, raising labor productivity and the wage rate, while reducing employment and the labor income share, similar to the predictions of the benchmark model. However, without the fixed cost of automation, all firms would be using robots, such that changes in the robot price do not affect industry concentration.

To investigate the role of the Kimball demand system with variable markups, we consider a counterfactual model with a CES demand system, keeping the fixed cost of automation. This counterfactual is a special case of our benchmark

understanding the connection between automation and a broader set of labor market variables (including wages). We leave that important task for future research.

³¹The data on robot prices in the U.S. are available only after 2002. To have a comparable period with the concentration measures in Autor et al. (2020), we assume that the fall in robot prices from 1998 to 2012 is the same as that from 2002 to 2016 (i.e., 40%).

³²Notice that, as Figure 1 shows, sales concentration measured by the sales share of the top 20 firms (i.e., CR20) rose by a similar magnitude. We focus on CR4 since, as mentioned before, this is more comparable to the share of the top 1% firms.

³³Although we do not have firm-level data to show that general capital equipment is more evenly distributed across firms relative to robots, Figure A.2 in the Appendix shows that this is true across industries. For example, in 2018, the distribution of robot density is highly skewed toward a few industries such as electronics and basic metals. In comparison, capital equipment intensity is more evenly distributed across the same set of two-digit manufacturing industries.

model with the super-elasticity of demand set to $\epsilon = 0$. We recalibrate the constant demand elasticity σ to match the average markup in the benchmark model. We also recalibrate the four parameters Q_a , σ_a , α_a , and η to match the four moments in the data used for calibrating our benchmark model (see Panel C of Table 2). Table A.4 presents the calibrated parameters. The calibrated model with the CES demand system matches the four data moments exactly.

Figure A.4 in the Appendix shows the steady-state relations between the robot price and the macroeconomic variables in the CES model. The CES model generates qualitatively similar features as does the benchmark model, except that markups are constant by construction. Quantitatively, however, the CES model implies a steeper positive relation between automation and industry concentration. As shown in Figure A.4, the CES model predicts that a 40% decline in Q_a leads to the sales share of the top 1% of firms to rise by about 2.85 percentage points (from 33.66% to 36.51%), explaining about 95% of the observed increases in sales concentration, which seems implausible and is much larger than that predicted by the benchmark model (about 49%). The CES model also predicts a larger increase in employment concentration. Regardless of whether markups are constant, the presence of the fixed cost of automation implies that a decline in the robot price disproportionately benefits large firms that use robots, enabling them to expand production further, resulting in an increase in sales concentration. Under a Kimball demand system, however, the expansion of sales of the top firms is accompanied by an increase in their markups and, to maintain higher markups, the top firms restrain the expansion of production. Under the CES demand system, markup is constant by construction, such that the same decline in the robot price leads to larger increases in sales concentration.

7.4. Transitional dynamics

Our main analysis focuses on a stationary equilibrium. Since we have a dynamic model, we now solve for the transition dynamics following a sharp decline in the robot price. We assume that the economy starts from the calibrated initial steady state in period $t = 1$ and then the robot price declines unexpectedly by 40% in $t = 2$ and stays at that lower level. The agents have perfect foresight after period 2.³⁴ Figure 6 presents the transition paths for several key aggregate variables. The figure shows that the transition process in our model is relatively short: The economy reaches the new steady state in about 40 quarters, with most changes taking place in the first 20 quarters.

The transition paths are mostly smooth, except for the initial responses of the macro variables in the period when the shock hits (i.e., $t = 2$). The decline in Q_a in $t = 2$ boosts robot investment and aggregate output. Since the aggregate stock of robots is predetermined, the initial increase in aggregate output is driven by the initial increase in employment. As output increases, the robot rental rate (which equals the marginal product of robots) rises initially, raising the marginal cost of production for automating firms. Given the fixed costs of automation, larger and more productive firms use robots and they face an initial increase in the user cost of robots. The resulting reallocation away from those firms reduces aggregate productivity, industry concentration, and average markup while raising the labor share in the period of the shock. In subsequent periods, as new robots come online for production, labor productivity, concentration, and average markup all rise smoothly and the labor share also falls smoothly until reaching the new steady state.

7.5. Policy analysis

The source of inefficiencies in the model economy is the size-dependent markups, which in principle can be offset by size-dependent subsidies to firms. In reality, however, implementing such size-dependent subsidies requires knowledge about firm-specific markups, which can be challenging. Thus, we follow the literature and examine the effects of a robot tax on macroeconomic allocations and welfare.³⁵

For this purpose, we introduce a tax τ_t on the value of robots. Absent size-dependent subsidies that offset the markup distortions, the economy is in a second-best environment, where a robot tax policy can potentially improve welfare. We assume that the robot tax revenues are rebated to the representative household in the form of lump-sum transfers. The household budget constraint (5) now becomes

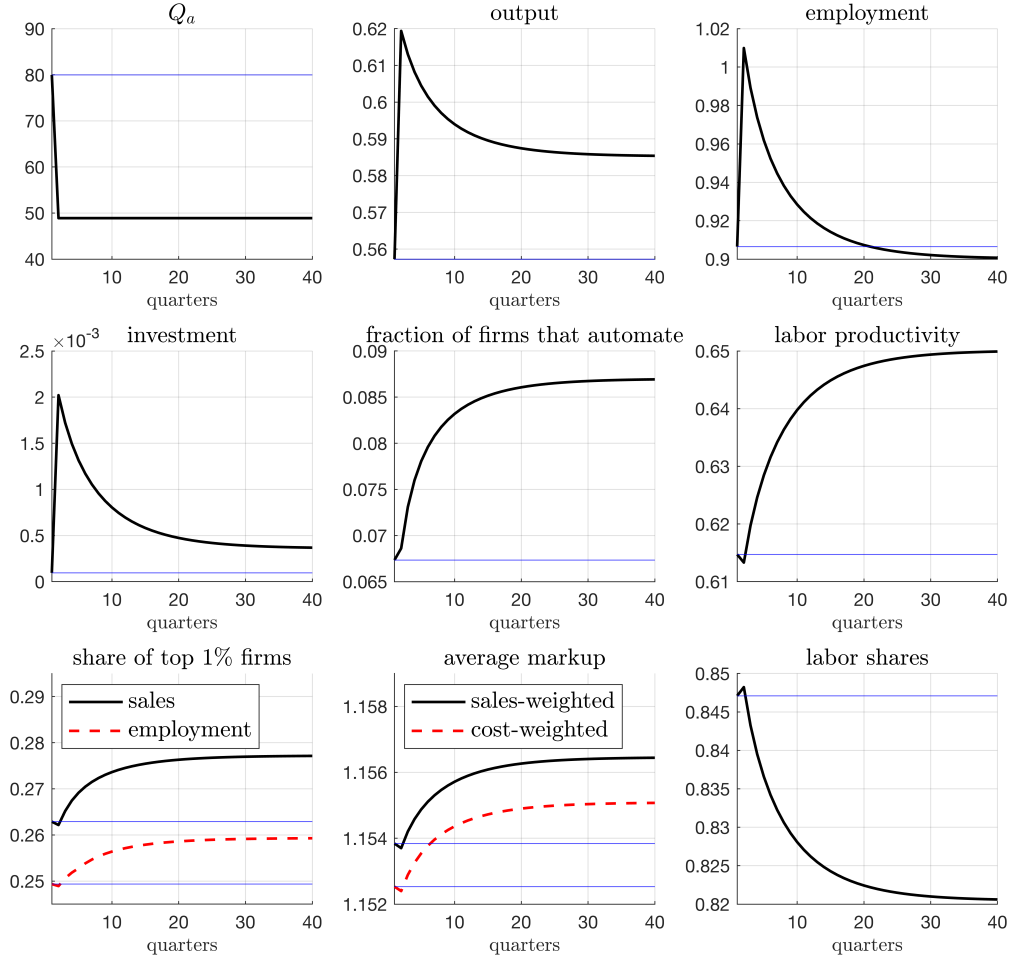
$$C_t + Q_{a,t}[A_{t+1} - (1 - \delta_a)A_t] \leq W_t N_t + r_{a,t} A_t + \pi_t - \tau_t Q_{a,t} A_{t+1} + T_t, \quad (38)$$

where T_t denotes the lump-sum transfer and we have substituted out robot investment using the law of motion (6). The optimizing choice of robot investment implies that

$$(1 + \tau_t)Q_{a,t} = \rho_{t,t+1} [r_{a,t+1} + Q_{a,t+1}(1 - \delta_a)]. \quad (39)$$

³⁴Appendix C.2 describes the algorithm for solving the transition dynamics under the assumption of perfect foresight, an approach similar to that in [Benguria, Saffie and Urzua \(2023\)](#).

³⁵See Section 2 for a discussion of the related literature.

Figure 6: Transition Dynamics: 40% One-Time Decrease in Q_a at $t = 2$ 

Note: This figure shows transition paths in response to an unexpected permanent decline in Q_a by 40% at $t = 2$ from the initial steady state.

To explore the welfare implications of the robot tax policy, we measure welfare losses (or gains) under a robot tax by the consumption equivalent variation relative to the laissez-faire benchmark. Specifically, we compute the percentage change in consumption in perpetuity that is required such that the representative household is indifferent between living in the economy with the robot tax and the benchmark economy without it.

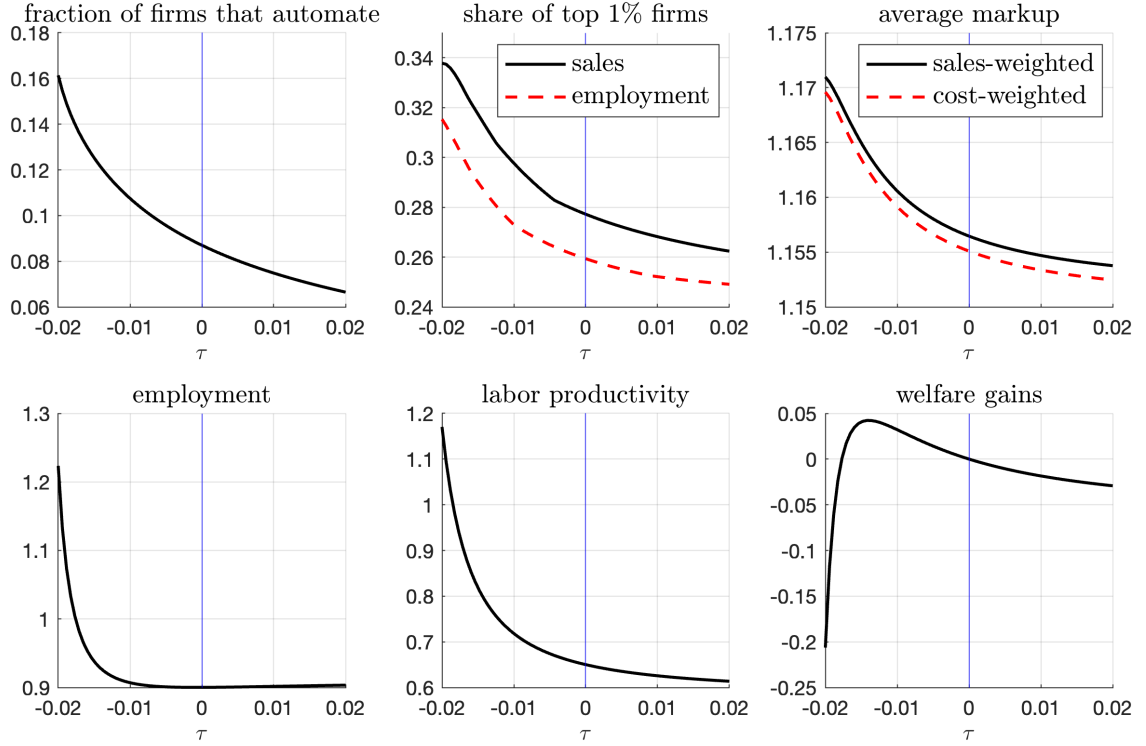
The welfare in the economy with the robot tax rate τ_t (denoted by $W(\tau)$) is given by

$$W(\tau) = \sum_{t=0}^{\infty} \beta^t \left[\ln C_t(\tau_t) - \chi \frac{N_t(\tau_t)^{1+\xi}}{1+\xi} \right], \quad (40)$$

where $C_t(\tau_t)$ and $N_t(\tau_t)$ denote, respectively, consumption and employment in the equilibrium under the tax policy. The welfare in the benchmark economy without the tax is given by

$$W(0) = \sum_{t=0}^{\infty} \beta^t \left[\ln C_t(0) - \chi \frac{N_t(0)^{1+\xi}}{1+\xi} \right], \quad (41)$$

where $C_t(0)$ and $N_t(0)$ denote, respectively, consumption and employment in the equilibrium of the benchmark economy without the tax (i.e., with $\tau_t = 0$ for all t). The welfare loss associated with the tax rate τ_t is given by

Figure 7: Steady-State Effects of Taxing Robots

Note: This figure shows the effects of imposing a tax τ on the stock value of robots on aggregate variables and welfare in the steady state. Welfare gains are measured by the consumption equivalent (percent) relative to the laissez-faire economy with $\tau = 0$.

the consumption equivalent μ , which is defined by the relation

$$\sum_{t=0}^{\infty} \beta^t \left[\ln C_t(0)(1 - \mu) - \chi \frac{N_t(0)^{1+\xi}}{1 + \xi} \right] = W(\tau), \quad (42)$$

Solving for μ from equation (42), we obtain

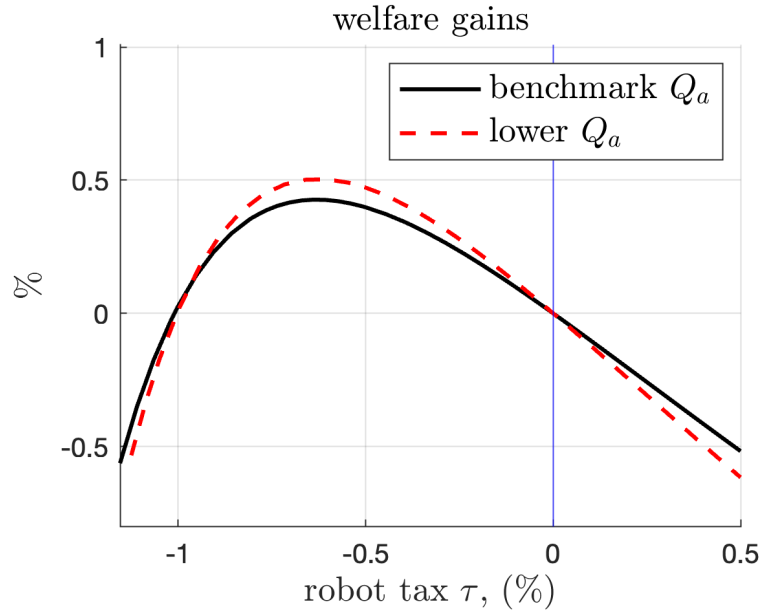
$$\mu = 1 - \exp[(1 - \beta)(W(\tau) - W(0))]. \quad (43)$$

A positive (negative) μ would imply a welfare loss (gain) in the economy with the robot tax relative to the laissez-faire benchmark.

7.5.1. Effects of a robot tax in the steady state

Figure 7 shows the steady-state effects of imposing a robot tax. In a steady state, the optimizing condition (39) implies that the rental rate is given by $r_a = Q_a[(1 + \tau)/\beta - 1 + \delta_a]$, which increases with the robot tax rate. Similar to a decline in the robot price, a reduction in the robot tax has two competing effects on the market share of the top firms. On the one hand, it reduces r_a and thus the cost of operating the automation technology, enabling large and automating firms to become even larger (the intensive margin). On the other hand, it raises the fraction of firms that automate (the extensive margin), such that some medium-sized firms that initially used the labor-only technology switch to the automation technology (the extensive margin), reducing the market share of the top firms. Under our calibration, the intensive margin effect dominates, such that a lower robot tax (or a higher robot subsidy) raises the sales concentration and the employment concentration, and it also enlarges the gap between the two measures of concentration because robots substitute for workers.

Since larger firms have higher markups and are more productive, the between-firm reallocation associated with a decline in the robot tax raises the average markup and aggregate labor productivity. The rise in labor productivity

Figure 8: Welfare Effects of Taxing Robots in the Economy with Transitional Dynamics

Note: This figure shows the welfare gains associated with different robot tax rates (τ), taking into account the transitional dynamics. Welfare gains are measured by the consumption equivalent (percent) relative to the laissez-faire economy with $\tau = 0$. The solid line indicates the welfare effects in the benchmark model with the calibrated value of the steady-state robot price Q_a . The dashed line indicates the welfare effects in the counterfactual economy with a robot price that is permanently 10% lower.

boosts labor demand, which tends to increase employment. At the same time, however, robots substitute for workers, reducing aggregate employment. Under our calibration, the net effect of the robot tax on employment is small, unless the subsidy rate is relatively large.

The presence of market power in our model renders the laissez-faire equilibrium inefficient. In particular, since firms have monopolistic markups, aggregate output is inefficiently low, such that a robot subsidy can improve welfare by stimulating robot investment and thereby bringing aggregate output closer to the efficient level. In this sense, a robot subsidy acts like a production subsidy that helps mitigate, but not completely undo, markup distortions.³⁶ A robot subsidy stimulates automation investment, resulting in more concentrated production in large firms that have high productivity and also high markups. By raising aggregate productivity, a robot subsidy can be effective in bringing aggregate output closer to the efficient level. However, a robot subsidy also raises the average markup, making it harder to undo markup distortions. With these effects taken into account, a modest robot subsidy is needed to maximize welfare. As shown in Figure 7, the optimal robot subsidy rate is about 1.42% under our calibration, with a maximum welfare gain of about 4.23% of steady-state consumption equivalent relative to the laissez-faire benchmark.

7.5.2. Effects of a robot tax with transition dynamics

We now examine the implications of the robot tax policy when transition dynamics are taken into account. We start from the steady state in the benchmark model without taxes (i.e., period one) and, in period two, we introduce an unexpected robot tax that remains in place permanently. We solve for the transition dynamics assuming that the agents have perfect foresight from period two onward.

Figure 8 plots the welfare gains (in consumption equivalent units) as a function of the robot tax rate in the model with transition dynamics. The figure shows that, similar to the steady-state analysis, the tax policy has a non-monotonic effect on welfare. Under our benchmark calibration (the black solid line), a permanent robot subsidy of 0.64% maximizes welfare, with a welfare gain of 0.43% of consumption equivalent relative to the no-tax economy.

³⁶As long as markups are positive, a robot subsidy (or a production subsidy) would help mitigate markup distortions. This implication does not rely on the Kimball demand system and continues to hold in, for example, an economy with a CES demand system, as shown in Figure A.5 in the Appendix. In the CES case, the optimal subsidy rate is 1.36%, resulting in consumption-equivalent welfare gains of 3.32%.

Compared to the steady-state effects of robot taxes shown in Figure 7, the optimal subsidy rate and the maximum welfare gains are both smaller in the dynamic model. This result reflects that the robot tax policy has different implications for welfare in the short run versus the long run. Relative to the original steady state with no taxes, imposing a modest, permanent robot subsidy boosts consumption and improves welfare in the new steady state, as we have discussed in Section 7.5.1. However, in the transition process, the subsidy stimulates robot investment, crowding out consumption (see Figure A.6 in the Appendix). The increase in robot investment also boosts labor productivity and raises labor demand, resulting in an increase in employment in the short run. The transitory declines in consumption and increases in employment reduce the utility flows for the representative household in the transition process, partly offsetting the steady-state welfare gains.

Figure 8 also shows the welfare gains under different robot tax rates in a counterfactual dynamic economy with a robot price that is 10% lower permanently (the dashed line). With a lower robot price, the optimal robot subsidy is smaller (0.60% vs. 0.64%) and the associated welfare gain is larger (0.50% vs. 0.43%). Thus, when the robot price is lower, a smaller subsidy is needed to maximize welfare. This is because, with a lower robot price, robot investment is larger and therefore production is closer to the efficient level, making robot subsidy less crucial.

8. Conclusion

We provide empirical evidence suggesting that automation is associated with the rise in sales concentration and the divergence between sales and employment concentration in the U.S. manufacturing sector since the early 2000s. We study the economic forces that drive the link between automation and industry concentration using a general equilibrium framework. The model highlights two important channels—fixed costs of operating the automation technology and labor-substituting effects of automation—that help explain the empirical relation between automation and concentration. Our calibrated model predicts that a decline in the robot price of a magnitude similar to that observed during the past two decades can account for about 49% of the rise in sales concentration and about 25% of the diverging trends between sales and employment concentration. Thus, the rise of automation is quantitatively important for driving the rise of superstar firms.

We use our general equilibrium model to evaluate the welfare implications of a robot tax policy. With the markup distortions, aggregate output is inefficiently low in the laissez-faire equilibrium. A robot subsidy induces reallocation toward large automating firms that have higher productivity, bringing aggregate output closer to the efficient level. However, large firms also have high markups, such that a robot subsidy raises the average markup through between-firm reallocation. Under calibrated parameters, a modest robot subsidy improves steady-state welfare relative to the laissez-faire benchmark. We further show that the robot tax policy has different welfare implications when the short-run transition dynamics are taken into account. During the transition process, an increase in robot investment stimulated by a robot subsidy crowds out consumption, reducing the flow value of the representative household's utility. Taking into account the dynamic effects leads to a smaller optimal robot subsidy and also a smaller welfare gain relative to those in the steady-state analysis. Furthermore, when the robot price is lower, a smaller robot subsidy is required to maximize welfare.

To highlight the key mechanism that connects automation with industry concentration, we focus on a stylized model that abstracts from several other mechanisms or sources of friction studied in the literature. For example, our benchmark model features a representative household with a homogeneous type of labor, and robots can substitute for workers. In reality, automation equipment substitutes for low-skill workers while complementing high-skill workers (Krusell, Ohanian, Ríos-Rull and Violante, 2000). In such an environment, a decline in the robot price would stimulate automation investment, which in turn boosts demand for high-skill labor, raising the skilled wage premium and also the cost of automation, partly offsetting the initial increase in automation investment (Leduc and Liu, 2024). Extending our model to incorporate skill heterogeneity is an important direction for future research. In that extended framework, changes in automation would affect not only the concentration of firms but also the distribution of income for households. Our framework can also be extended to an open-economy environment, which would allow us to examine potential interactions between automation, offshoring, and domestic production along the lines of Firooz et al. (2024). These extensions, in our view, are promising avenues for future research. Our work represents a small first step toward that direction.

References

- Acemoglu, D., Anderson, G.W., Beede, D.N., Buffington, C., Childress, E.E., Dinlersoz, E., Foster, L.S., Goldschlag, N., Haltiwanger, J.C., Kroff, Z., Restrepo, P., Zolas, N., 2022. Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey. Working Paper 30659. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w30659>, doi:10.3386/w30659.
- Acemoglu, D., Lelarge, C., Restrepo, P., 2020a. Competing with robots: Firm-level evidence from france. AEA Papers and Proceedings 110, 383–388. URL: <https://www.aeaweb.org/articles?id=10.1257/pandp.20201003>, doi:10.1257/pandp.20201003.
- Acemoglu, D., Manera, A., Restrepo, P., 2020b. Does the US Tax Code Favor Automation? Working Paper 27052. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w27052>, doi:10.3386/w27052.
- Acemoglu, D., Restrepo, P., 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108, 1488–1542.
- Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from us labor markets. *Journal of Political Economy* 128, 2188–2244.
- Acemoglu, D., Restrepo, P., 2021. Tasks, Automation, and the Rise in US Wage Inequality. Working Paper 28920. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w28920>, doi:10.3386/w28920.
- Aghion, P., Antonin, C., Bunel, S., Jaravel, X., 2021. What Are the Labor and Product Market Effects of Automation? New Evidence from France. Technical Report. Sciences Po. URL: <https://ideas.repec.org/p/spo/wpmain/infohdl2441-170cd4sul89ddpnfuomvfm0jco.html>.
- Aghion, P., Antonin, C., Bunel, S., Jaravel, X., 2023. The local labor market effects of modern manufacturing capital: Evidence from france. AEA Papers and Proceedings 113, 219–23. URL: <https://www.aeaweb.org/articles?id=10.1257/pandp.20231039>, doi:10.1257/pandp.20231039.
- Aghion, P., Bergeaud, A., Boppart, T., Klenow, P.J., Li, H., 2019. A Theory of Falling Growth and Rising Rents. Working Paper 26448. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w26448>, doi:10.3386/w26448.
- Akcigit, U., Ates, S.T., 2019. Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory. Working Paper 25755. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w25755>, doi:10.3386/w25755.
- Anderson, T.W., Rubin, H., 1949. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of mathematical statistics* 20, 46–63.
- Andrews, I., Stock, J.H., Sun, L., 2019. Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics* 11, 727–753. URL: <https://ideas.repec.org/a/anr/reveco/v11y2019p727-753.html>, doi:10.1146/annurev-economics.
- Aruoba, S.B., Oue, E., Saffie, F., Willis, J., 2024. Non-Constant Demand Elasticities, Firm Dynamics and Monetary Non-Neutrality: Role of Demand Shocks. Technical Report. National Bureau of Economic Research Working Paper No. 32518.
- Autor, D., Dorn, D., Katz, L.F., Patterson, C., Van Reenen, J., 2020. The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics* 135, 645–709. URL: <https://doi.org/10.1093/qje/qjaa004>, doi:10.1093/qje/qjaa004, arXiv:<https://academic.oup.com/qje/article-pdf/135/2/645/32994954/qjaa004.pdf>.
- Autor, D.H., Dorn, D., Hanson, G.H., 2013. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review* 103, 2121–2168.
- Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118, 1279–1333.
- Autor, D.H., Salomons, A., 2018. Is automation labor-displacing? productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity* 1, 1–87.
- Bajgar, M., Berlingieri, G., Calligaris, S., Crisculo, C., Timmis, J., 2019. Industry concentration in europe and north america. URL: <https://www.oecd-ilibrary.org/content/paper/2ff98246-en>, doi:<https://doi.org/https://doi.org/10.1787/2ff98246-en>. OECD Productivity Working Papers.
- Benguria, F., Saffie, F., Urzua, S., 2023. The Transmission of Commodity Price Super-Cycles. *The Review of Economic Studies* 91, 1923–1955. URL: <https://doi.org/10.1093/restud/rdad078>, doi:10.1093/restud/rdad078, arXiv:<https://academic.oup.com/restud/article-pdf/91/4/1923/58447542/rdad078.pdf>.
- Beraja, M., Zorzi, N., 2022. Inefficient Automation. Working Paper 30154. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w30154>, doi:10.3386/w30154.
- Bergholt, D., Furlanetto, F., Maffei-Faccioli, N., 2022. The decline of the labor share: New empirical evidence. *American Economic Journal: Macroeconomics* 14, 163–98. URL: <https://www.aeaweb.org/articles?id=10.1257/mac.20190365>, doi:10.1257/mac.20190365.
- Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S.J., 2018. Really uncertain business cycles. *Econometrica* 86, 1031–1065.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-Based Improvements for Inference with Clustered Errors. *The Review of Economics and Statistics* 90, 414–427. URL: <https://doi.org/10.1162/rest.90.3.414>, doi:10.1162/rest.90.3.414, arXiv:<https://direct.mit.edu/rest/article-pdf/90/3/414/1614600/rest.90.3.414.pdf>.
- Cheng, H., Drozd, L.A., Giri, R., Taschereau-Dumouchel, M., Xia, J., 2021. The Future of Labor: Automation and the Labor Share in the Second Machine Age. Working Paper WP 21-11. Federal Reserve Bank of Philadelphia.
- Choi, J., Levchenko, A.A., Ruzic, D., Shim, Y., 2024. Superstars or supervillains? large firms in the south korean growth miracle .
- Costinot, A., Werning, I., 2022. Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation. *The Review of Economic Studies* 90, 2261–2291. URL: <https://doi.org/10.1093/restud/rdac076>, doi:10.1093/restud/rdac076, arXiv:<https://academic.oup.com/restud/article-pdf/90/5/2261/51403015/rdac076.pdf>.
- Edmond, C., Midrigan, V., Xu, D.Y., 2021. How costly are markups? Unpublished Manuscript, New York University and Duke University.
- Ekerdt, L.K., Wu, K.J., 2022. The rise of specialized firms. Working Paper .
- Firooz, H., Leduc, S., Liu, Z., 2024. Reshoring, Automation, and Labor Markets Under Trade Uncertainty. Technical Report. URL: <https://doi.org/10.24148/wp2024-16>. federal Reserve Bank of San Francisco Working Paper 2024-16.

- Furman, J., Orszag, P., 2018. Slower Productivity and Higher Inequality: Are They Related? Working Paper Series WP18-4. Peterson Institute for International Economics. URL: <https://ideas.repec.org/p/iie/wpaper/wp18-4.html>.
- Guerreiro, J., Rebelo, S., Teles, P., 2022. Should robots be taxed? Review of Economic Studies 89, 279–311.
- Hsieh, C.T., Rossi-Hansberg, E., 2019. The Industrial Revolution in Services. NBER Working Papers 25968. National Bureau of Economic Research. URL: <https://ideas.repec.org/p/nbr/nberwo/25968.html>.
- Hubmer, J., Restrepo, P., 2022. Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution. Working Paper 28579. National Bureau of Economic Research.
- Jaimovich, N., Siu, H.E., 2020. Job polarization and jobless recoveries. The Review of Economics and Statistics 102, 129–147.
- Jeong, D., Baek, S., Peri, G., 2024. Automation, human task innovation, and labor share: Unveiling the role of elasticity of substitution .
- Judd, K.L., 1998. Numerical Methods in Economics. The MIT Press.
- Kehrig, M., Vincent, N., 2021. The Micro-Level Anatomy of the Labor Share Decline. The Quarterly Journal of Economics 136, 1031–1087. URL: <https://doi.org/10.1093/qje/qjab002>, doi:10.1093/qje/qjab002, arXiv:<https://academic.oup.com/qje/article-pdf/136/2/1031/36725272/qjab002.pdf>.
- Khan, A., Thomas, J.K., 2008. Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. Econometrica 76, 395–436.
- Kimball, M.S., 1995. The quantitative analytics of the basic neomonetarist model. Journal of Money, Credit and Banking 27, 1241–1277.
- Klenow, P.J., Willis, J.L., 2016. Real rigidities and nominal price changes. Economica 83, 443–472.
- Krusell, P., Ohanian, L.E., Ríos-Rull, J.V., Violante, G.L., 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. Econometrica 68, 1029–1053.
- Kwon, S.Y., Ma, Y., Zimmermann, K., 2024. 100 years of rising corporate concentration. American Economic Review 114, 2111–40. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20220621>, doi:10.1257/aer.20220621.
- Lashkari, D., Bauer, A., Boussard, J., 2024. Information technology and returns to scale. American Economic Review 114, 1769–1815. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20220522>, doi:10.1257/aer.20220522.
- Leduc, S., Liu, Z., 2024. Automation, bargaining power, and labor market fluctuations. American Economic Journal: Macroeconomics 16, 311–49. URL: <https://www.aeaweb.org/articles?id=10.1257/mac.20220181>, doi:10.1257/mac.20220181.
- Montiel Olea, J.L., Pflueger, C., 2013. A robust test for weak instruments. Journal of Business & Economic Statistics 31, 358–369.
- Moreira, M.J., 2009. Tests with correct size when instruments can be arbitrarily weak. Journal of Econometrics 152, 131–140.
- Olmstead-Rumsey, J., 2019. Market Concentration and the Productivity Slowdown. MPRA Paper 93260. University Library of Munich, Germany. URL: <https://ideas.repec.org/p/pra/mprapa/93260.html>.
- Prettner, K., Strulik, H., 2020. Innovation, automation, and inequality: Policy challenges in the race against the machine. Journal of Monetary Economics 116, 249–265. URL: <https://www.sciencedirect.com/science/article/pii/S0304393219301965>, doi:<https://doi.org/10.1016/j.jmoneco.2019.10.012>.
- Ridder, M.D., 2023. Market power and innovation in the intangible economy. American Economic Review .
- Rogerson, R., Wallenius, J., 2009. Micro and macro elasticities in a life cycle model with taxes. Journal of Economic Theory 144, 2277–2292.
- Roodman, D., Nielsen, M.Ø., MacKinnon, J.G., Webb, M.D., 2019. Fast and wild: Bootstrap inference in stata using boottest. The Stata Journal 19, 4–60.
- Shen, S., 2023. Customer acquisition, rising concentration and us productivity dynamics. Rising Concentration and US Productivity Dynamics (September 15, 2023) .
- Sui, X., 2022. Uneven firm growth in a globalized world. Working Paper .
- Tambe, P., Hitt, L., Rock, D., Brynjolfsson, E., 2020. Digital Capital and Superstar Firms. Working Paper 28285. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w28285>, doi:10.3386/w28285.
- Thuemmel, U., 2022. Optimal Taxation of Robots. Journal of the European Economic Association 21, 1154–1190. URL: <https://doi.org/10.1093/jeea/jvac062>, doi:10.1093/jeea/jvac062, arXiv:<https://academic.oup.com/jeea/article-pdf/21/3/1154/50525711/jvac062.pdf>.
- Yao, J., 2022. Knowledge is (market) power. Available at SSRN 4273942 .
- Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D.N., Buffington, C., Goldschlag, N., Foster, L., Dinlersoz, E., 2020. Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey. Working Paper 28290. National Bureau of Economic Research. URL: <http://www.nber.org/papers/w28290>, doi:10.3386/w28290.

Table A.1
Industries Included in the Sample

ISIC rev4		IFR	
Code	Label	Code	Label
10-12	Manufacture of food products, Manufacture of beverages, Manufacture of tobacco products	10-12	Food products and beverages; Tobacco products
13-15	Manufacture of textiles, Manufacture of wearing apparel, Manufacture of leather and related products	13-15	Textiles, leather, wearing apparel
16, 31	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials, Manufacture of furniture	16	Wood and wood products (incl. furniture)
17-18	Manufacture of paper and paper products, Printing and reproduction of recorded media	17-18	Paper and paper products, publishing & printing
19-22	Manufacture of coke and refined petroleum products, Manufacture of chemicals and chemical products, Manufacture of basic pharmaceutical products and pharmaceutical preparations, Manufacture of rubber and plastics products	19-22	Plastic and chemical products
23	Manufacture of other non-metallic mineral products	23	Glass, ceramics, stone, mineral products n.e.c. (without automotive parts)
24	Manufacture of basic metals	24	Basic metals (iron, steel, aluminum, copper, chrome)
25	Manufacture of fabricated metal products, except machinery and equipment	25	Metal products (without automotive parts), except machinery and equipment
26-27	Manufacture of computer, electronic and optical products, Manufacture of electrical equipment	26-27	Electrical/electronics
28	Manufacture of machinery and equipment n.e.c.	28	Industrial Machinery
29	Manufacture of motor vehicles, trailers and semi-trailers	29	automotive
30	Manufacture of other transport equipment	30	Other transport equipment
D, E	Electricity, gas, steam and air conditioning supply, Water supply; sewerage, waste management, and remediation activities	E	Electricity, gas, water supply

Note: This table shows the corresponding ISIC revision 4 and IFR codes and labels for the industries included in our sample.

Table A.2
Summary Statistics

	#obs	mean	min	p25	p50	p75	max	s.d.
ln(robot/thousand employees)	117	0.48	-6.57	-1.12	1.02	2.32	5.86	2.36
ln(robots/million hours)	117	0.20	-6.83	-1.34	0.79	1.94	5.30	2.43
top 1% share of sales	117	0.31	0.09	0.22	0.30	0.37	0.77	0.13
top 1% share of employment	104	0.27	0.11	0.21	0.28	0.32	0.46	0.08

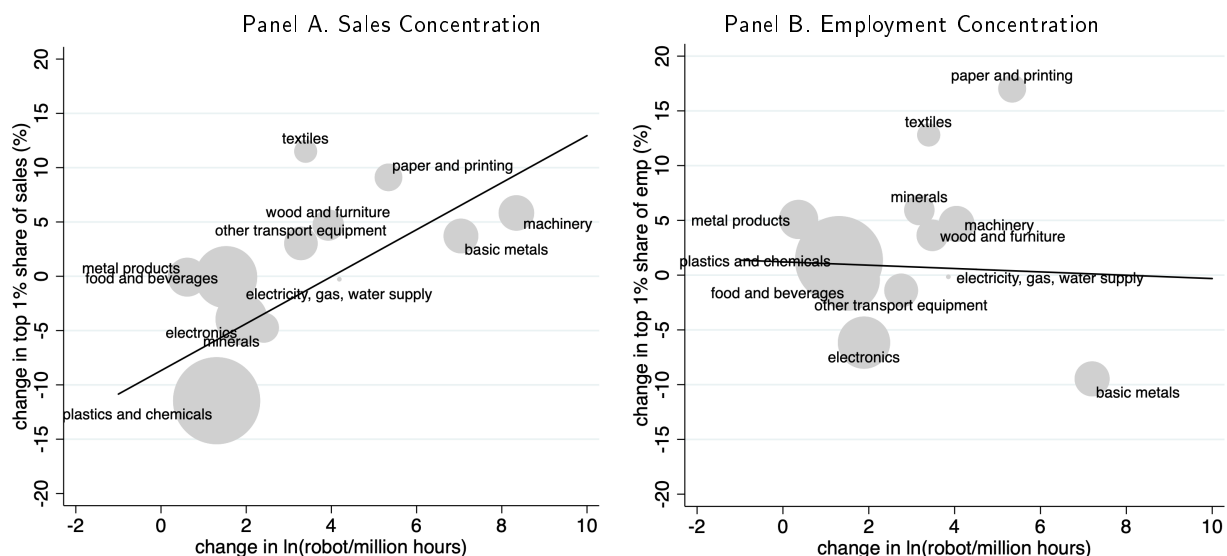
Note: This table shows the summary statistics of the data we use in the regressions. The industry-level robot density is measured as the operational stock of industrial robots per thousand employees or per million labor hours. We consider two measures of industry concentration: the sales share and the employment share of the top 1% of firms in the industry. For both measures of concentration, we restrict our sample to industry-year pairs with at least 10 firms.

Source: Authors' calculations using IFR, Compustat, and NBER-CES.

Appendices

A. Additional tables and figures

Figure A.1: Industry Concentration and Robot Density (per Million Hours)



Note: This figure shows the cumulative changes in sales concentration (Panel A) and employment concentration (Panel B) against changes in robot density. The industry concentration is measured by the share of the top 1% of firms within an industry. Robot density is measured by the operational stock of industrial robots per million hours in each industry. The cumulative change is the long difference between the ending value and the starting value of each variable during the years from 2007 to 2018. Since we have an unbalanced panel, we use the first (last) year with non-missing values as the starting (ending) point for calculating the long differences. The circle size indicates an industry's sales share in the initial year (2007). The line shows the prediction from a linear regression weighted by industries' initial sales shares. The slope coefficient for sales concentration (Panel A) is 0.022 with a standard error of 0.008. The slope coefficient for employment concentration (Panel B) is -0.0015 with a standard error of 0.010.

Source: IFR, NBER-CES, Compustat, and authors' calculation.

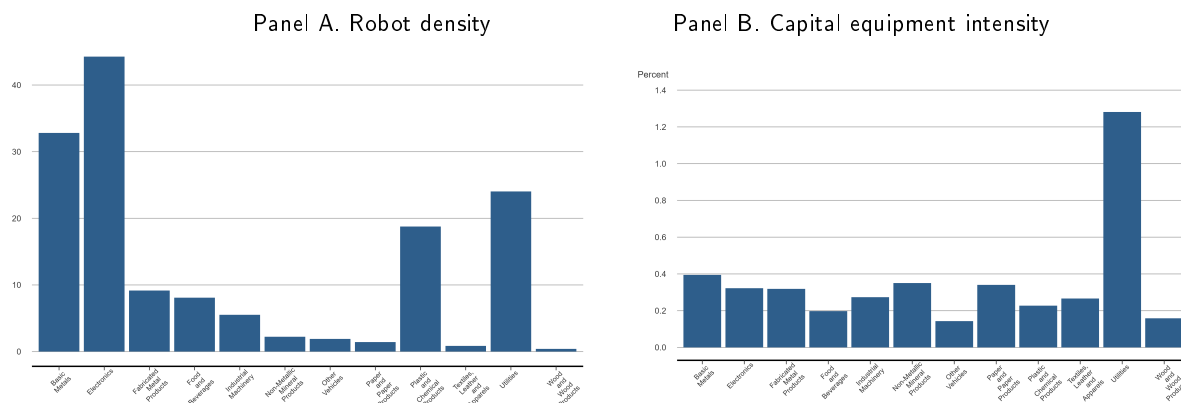
Table A.3

First-Stage of the IV Regressions for Robot Density and Industry Concentration

Second-stage dependent variable:	top 1% share of sales		top 1% share of emp	
First-stage dependent variable:	$\ln\left(\frac{\text{robot}}{\text{thousand emp}}\right)$	$\ln\left(\frac{\text{robot}}{\text{million hours}}\right)$	$\ln\left(\frac{\text{robot}}{\text{thousand emp}}\right)$	$\ln\left(\frac{\text{robot}}{\text{million hours}}\right)$
	(1)	(2)	(3)	(4)
EURO5 $\ln(\text{robot}/\text{thousand emp})$	1.815 (1.214)		1.404 (0.936)	
EURO5 $\ln(\text{robot}/\text{million hours})$		1.694 (1.240)		1.323 (0.904)
Observations	117	117	104	104
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
First-stage Effective F-statistic	2.235	1.866	2.251	2.141

Note: This table shows the first-stage results of the IV regression from the empirical specification (2). The second-stage dependent variables are the sales share (first two columns) and employment share (last two columns) of the top 1% of firms. The first-stage dependent variable is the U.S. robot density, measured as the operational stock of industrial robots per thousand workers or million labor hours within the industry. The IV for the U.S. robot density is the one-year lag of the robot density averaged over five European countries (EURO5). The last row shows the first-stage effective F-statistic of [Montiel Olea and Pflueger \(2013\)](#). In all regressions, the industries are weighted by their sales share in the initial year (2007), and the regressions also control for industry and year fixed effects. Standard errors in parentheses are clustered at the industry level. Stars denote the statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

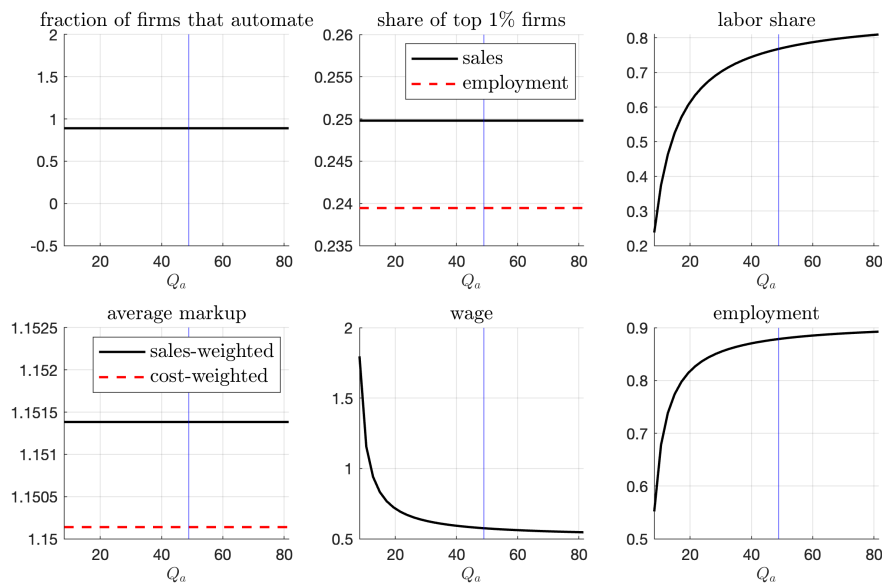
Figure A.2: Distribution of robot density and capital equipment intensity across industries



Note: This figure shows the distribution of robot density (Panel A) and of capital equipment intensity (Panel B) in the year 2018 across the 12 two-digit manufacturing industries in our sample. Robot density is measured by the operational stock of industrial robots per thousand workers in an industry. Capital equipment intensity is measured by the ratio of the nominal value of capital equipment to the nominal value added in an industry.

Source: IFR, NBER-CES, Bureau of Economic Analysis, and authors' calculation.

Figure A.3: Aggregate Variables (no fixed cost of automation)



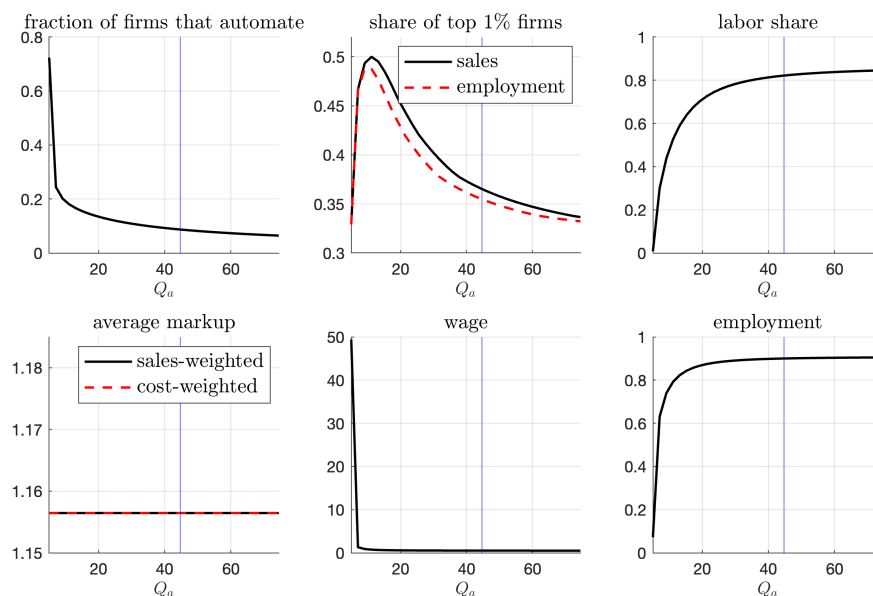
Note: This figure shows the steady-state effects of changes in the robot price Q_a on the fraction of firms that automate, the share of the top 1% of firms, the labor share, the average markup, the wage rate, and employment in the counterfactual model with no fixed cost of automation. The vertical blue line indicates the calibrated value of robot price Q_a .

Table A.4
Parameter calibration (CES)

Parameter	Notation	Value	Sources/Matched Moments
Panel A: Parameters calibrated to match external sources			
Discount factor	β	0.99	4% annual interest rate
Inverse Frisch elasticity	ξ	0.5	Rogerson and Wallenius (2009)
Working disutility weight	χ	1	Normalization
Robot depreciation rate	δ_a	0.02	8% annual depreciation rate
Productivity persistence	γ	0.95	Khan and Thomas (2008)
Productivity standard dev.	σ_ϕ	0.1	Bloom et al. (2018)
Demand elasticity parameter	σ	7.39	Matching a markup of 1.156 in the benchmark
Super elasticity	ϵ/σ	0	Imposing constant markups
Panel B: Parameters calibrated to match moments in data			
Relative price of robots	Q_a	44.70	Fraction of automating firms
SD of log automation fixed costs	σ_a	3.12	Employment share of automating firms
Robot input weight	α_a	0.36	Robot density
Elasticity of substitution	η	2.01	Growth rate of robot density

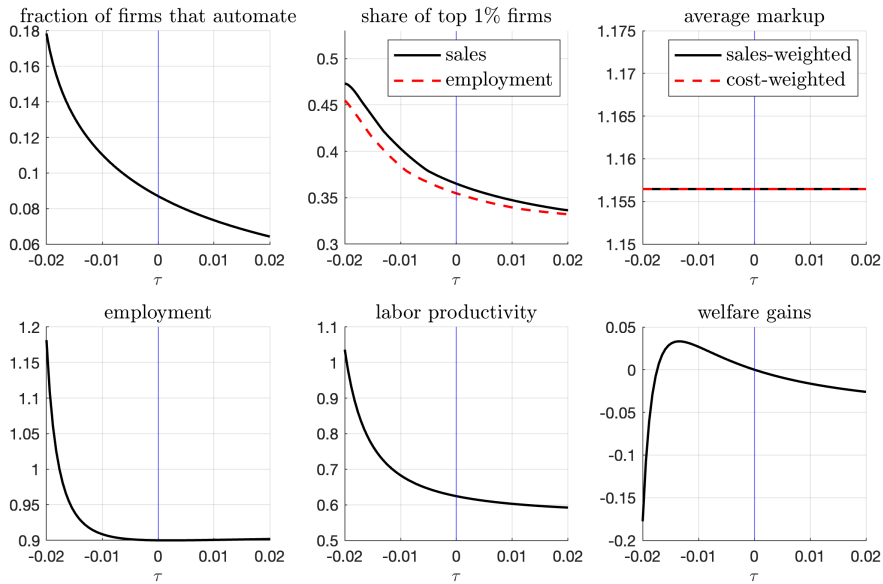
Note: This table shows the calibrated parameters in the counterfactual model with CES aggregation. Panel A reports the externally calibrated parameters and their sources. Panel B shows the parameters calibrated by moment matching.

Figure A.4: Aggregate Variables (CES)



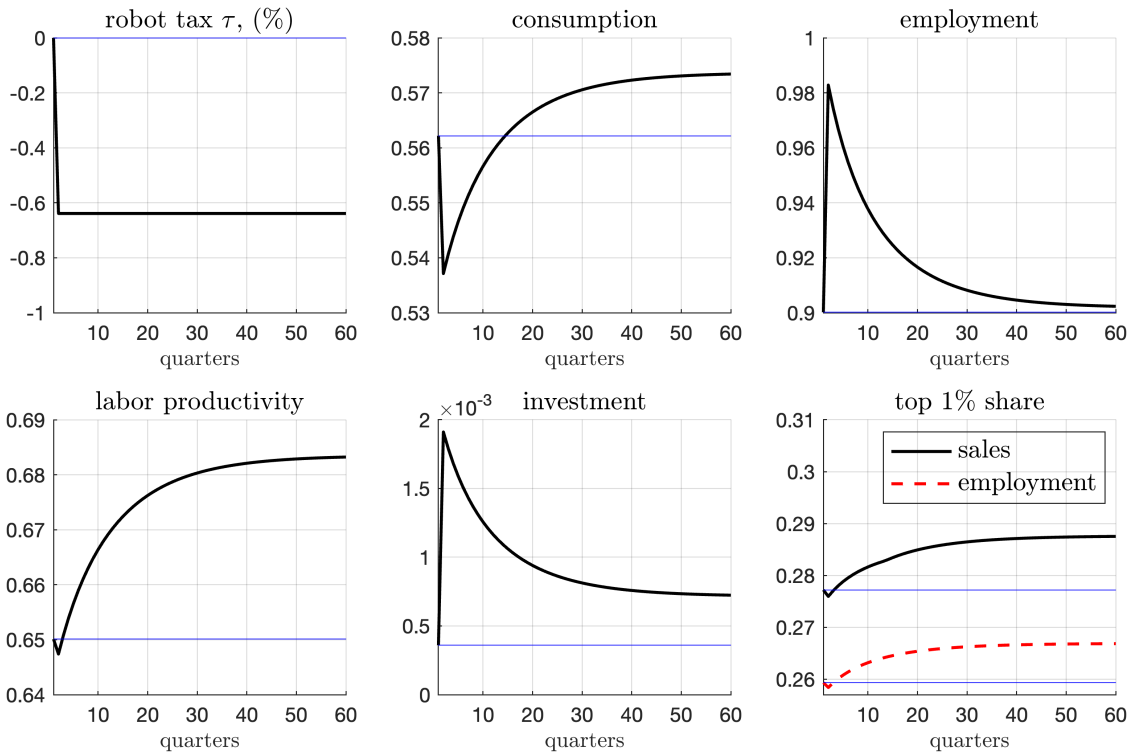
Note: This figure shows the effects of changes in the robot price Q_a on the fraction of firms that automate, the share of the top 1% of firms, the labor share, the average markup, the wage rate, and employment in the counterfactual model with a CES demand system. The vertical blue line indicates the calibrated value of robot price Q_a .

Figure A.5: Steady-State Effects of Taxing Robots (CES)



Note: This figure shows the effects of imposing a tax τ on the stock value of robots on aggregate variables and welfare in the steady state in the model with a CES demand system. Welfare gains are measured by the consumption equivalent (percent) relative to the laissez-faire economy with $\tau = 0$.

Figure A.6: Transition Paths Under the Optimal Robot Tax



Note: This figure shows the dynamic effects of imposing a permanent robot subsidy of 0.64%, corresponding to the optimal subsidy rate in the dynamic model under the benchmark calibration.

B. Derivations

To simplify the intermediate producers' problem in equation (19), rewrite the value function so that s is not a state variable:

$$\begin{aligned}
V_t(\phi_t; s_t) &= \max_{p_t, y_t, N_t, A_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) - s_t \phi_t \mathbb{1}\{A_t(\phi_t) > 0\} \right. \\
&\quad \left. + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \right] \\
&= \max \left\{ \underbrace{\max_{p_t, y_t, N_t, A_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \right]}_{\equiv V_t^a(\phi_t)} \right. \\
&\quad \left. - s_t \phi_t, \underbrace{\max_{p_t, y_t, N_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \right]}_{\equiv V_t^n(\phi_t)} \right\} \\
&= \max\{V_t^a(\phi_t) - s_t \phi_t, V_t^n(\phi_t)\}
\end{aligned} \tag{44}$$

The firm with productivity ϕ_t chooses $A_t(\phi_t) > 0$ if and only if $s_t \leq s_t^*(\phi_t) \equiv \frac{V_t^a(\phi_t) - V_t^n(\phi_t)}{\phi_t}$.

The value of an automating firm can be written as

$$V_t^a(\phi_t) = \max_{p_t, y_t, N_t, A_t > 0} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) \right] + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \tag{45}$$

The value of a non-automating firm can be written as

$$V_t^n(\phi_t) = \max_{p_t, y_t, N_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) \right] + \rho_{t,t+1} E_{\phi_{t+1}|\phi_t} \int_{s_{t+1}} V_{t+1}(\phi_{t+1}; s_{t+1}) dF(s_{t+1}) \tag{46}$$

To compute the automation cutoff $s_t^*(\phi_t)$, we can write:

$$s_t^*(\phi_t)\phi_t = V_t^a(\phi_t) - V_t^n(\phi_t) \tag{47}$$

$$= \max_{p_t, y_t, N_t, A_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) - r_{a,t} A_t(\phi_t) \right] - \max_{p_t, y_t, N_t} \left[p_t(\phi_t)y_t(\phi_t) - W_t N_t(\phi_t) \right], \tag{48}$$

which gives Eq. 30 in the text.

C. Solution Algorithm

C.1. Steady state

In the steady state, the rental rate of robots is

$$r_a = Q_a \left(\frac{1}{\beta} - 1 + \delta_a \right). \tag{49}$$

There are three loops to solve for the steady state. The Y loop is outside of the W loop and the W loop is outside of the q loop.

Y loop: Use bisection to determine the aggregate final goods and other aggregate variables.

1. Guess aggregate final goods Y .
2. Compute W and firms' relative production $q(\phi)$ in the W loop as explained below.

3. Given the equilibrium wage rate, compute other aggregate variables by finding Y using the bisection method:
 - (a) Given the solved relative production $q(\phi)$, we have $y(\phi) = q(\phi)Y$.
 - (b) Given the robot price Q_a and the wage rate W , compute the marginal costs $\lambda^n(\phi)$ and $\lambda^a(\phi)$ by eq. (25) and (27), and we can get $A(\phi)$ and $N(\phi)$ from eq. (23), (24), and (26).
 - (c) The aggregate employment and robot stock are determined by eq. (34) and eq. (35).
 - (d) Consumption C is determined by eq. (7).
 - (e) The steady state aggregate investment in robots I_a is from (36).
 - (f) Compute Y^{new} using the resource constraint (33). Stop if Y converges.
 - i. If $Y = Y^{\text{new}}$, Y and all other aggregate variables are found.
 - ii. If $Y > Y^{\text{new}}$, reduce Y . Go back to Step 1.
 - iii. If $Y < Y^{\text{new}}$, increase Y . Go back to 1.

W loop: Use bisection to determine the wage rate.

1. Guess a wage W .
2. Compute firms' relative production $q(\phi)$ in the q loop as explained below.
3. Check whether the Kimball aggregator (9) holds.
 - (a) If LHS = RHS, the wage rate is found and jump out of W loop to Y loop.
 - (b) If LHS > RHS, increase W to reduce $q(\phi)$ according to eq. (10). Go back to Step 2.
 - (c) If LHS < RHS, reduce W to raise $q(\phi)$ according to eq. (10). Go back to Step 2.

q loop: Find the relative production.

1. Given the prices Q_a and W , the marginal cost of production is determined by eq. (25) for the automation technology and by eq. (27) for the labor-only technology.
2. Guess a demand shifter D .
3. Use eq. (10) to solve for the relative output $q(\phi)$ for each ϕ , for firms with and without robots.
 - (a) The right-hand side of (10) is a function of $q(\phi)$ by plugging in (14).
 - (b) The price in the left-hand side is the marginal cost in (25) or (27) times the markup in (16), which is also a function of $q(\phi)$.
 - (c) Use the bisection method to solve for $q(\phi)$ in eq. (10).
4. Compute the automation decisions.
 - (a) Compute $y(\phi) = q(\phi)Y$ with and without robots.
 - (b) Compute the demand for $A(\phi)$ and $N(\phi)$ with and without robots from eq. (23), (24), and (26).
 - (c) For each productivity ϕ , compute the profits with and without robots and thus get the automation cutoffs $s^*(\phi)$ according to (30), and thus the automation probability $F(s^*(\phi))$.
5. Given the automation decisions, compute D^{new} by (11). Stop if D converges. Otherwise, go back to Step 2 and repeat until D converges.
 - (a) If $D = D^{\text{new}}$, D and $q(\phi)$ are found and jump out of q loop to W loop.
 - (b) If $D > D^{\text{new}}$, reduce D . Go back to Step 2.
 - (c) If $D < D^{\text{new}}$, increase D . Go back to Step 2.

C.2. Transitional dynamics

We assume that the economy is in the steady state at $t = 1$ and Q_a unexpectedly decreases by 40% in period 2 and remains deterministically constant afterward.

Given an exogenous path of $\{Q_{a,t}\}_{t=1}^T$, we solve the economy's transition path as follows:

1. Ensure that T is sufficiently large so that the economy reaches its new steady state by time T . For example, set $T = 300$. The economy begins at its initial steady state at $t = 1$ and reaches the new steady state at $t = T$, an unexpected change in robot prices.
2. Make initial guesses for the sequence of stochastic discount factors (SDFs), the sequence of aggregate output, and $r_{a,2}$. Set $\{\rho_{t,t+1}^{(init)}\}_{t=2}^{T-1} = \beta$, $\{Y_t^{(init)}\}_{t=2}^{T-1} = Y_T$, and $r_{a,2}^{(init)}$ in between $r_{a,1}$ and $r_{a,T}$.¹

¹Since aggregate investment depends on the next period's aggregate robot stock, another variable besides the SDF needs to be guessed. An alternative approach is to guess the sequence of robot stocks $\{A_t^{(init)}\}_{t=2}^{T-1}$ and solve for Y , W , and q loops, as described in the steady state solution algorithm. In practice, this approach is slower and does not improve convergence.

3. For each $t = 2, 3, \dots, T - 1$, given $Q_{a,t-1}$, $Q_{a,t}$, $r_{a,2}^{(init)}$, $\rho_{t-1,t}^{(init)}$, and $Y_t^{(init)}$, solve for the equilibrium as follows:
- (a) The rental rate of robots for $t = 3, 4, \dots, T - 1$ is given by

$$\begin{aligned} Q_{a,t} &= \rho_{t,t+1}^{(init)} [r_{a,t+1} + Q_{a,t+1}(1 - \delta_a)] \\ \Rightarrow r_{a,t+1} &= Q_{a,t} / \rho_{t,t+1}^{(init)} - Q_{a,t+1}(1 - \delta_a) \\ \Rightarrow r_{a,t} &= Q_{a,t-1} / \rho_{t-1,t}^{(init)} - Q_{a,t}(1 - \delta_a). \end{aligned} \quad (50)$$

- (b) Given $r_{a,t}$ and $Y_t^{(init)}$, solve the W and q loops outlined in the steady state solution algorithm. This yields W and firms' relative production $q(\phi_t)$.
- (c) Similar to the Y loop in the steady state solution algorithm, compute other variables as follows:
- Given $r_{a,t}$ and W , compute the marginal costs $\lambda_t^n(\phi_t)$ and $\lambda_t^a(\phi_t)$ given by eq. (25) and (27), and solve for $A_t(\phi_t)$ and $N_t(\phi_t)$ from eq. (23), (24), and (26).
 - The aggregate employment and robot stock are determined by eq. (34) and eq. (35).
 - Aggregate consumption C_t is determined by eq. (7).
- (d) Compute aggregate investment in robots, $I_{a,t} = A_{t+1} - (1 - \delta_a)A_t$.
- (e) Compute aggregate output for each t using the resource constraint (33):

$$Y_t^{(new)} = C_t + Q_{a,t}I_{a,t} + \int_{\phi_t} \int_0^{s_t^*(\phi_t)} s_t \phi_t dF(s_t) dG(\phi_t). \quad (51)$$

- (f) Compute stochastic discount factors for each t : $\rho_{t,t+1}^{(new)} = \beta \frac{C_t}{C_{t+1}}$.

- (g) Update $r_{a,2}^{(new)} = A_2 - A_1$. Notice that $r_{a,2}$ is not given by equation (50) because the shock at period 2 is unexpected. Instead, $r_{a,2}$ is determined such that robot demand equals the pre-determined robot supply at period 1, i.e., $A_2 = A_1$.

4. Continue iterating until the sequences of SDFs, aggregate output, and $r_{a,2}$ converge, i.e., $\text{dist}(\{\rho_{t,t+1}^{(new)}\}_{t=1}^T, \{\rho_{t,t+1}^{(init)}\}_{t=1}^T) < 10^{-6}$, $\text{dist}(\{Y_t^{(new)}\}_{t=1}^T, \{Y_t^{(init)}\}_{t=1}^T) < 10^{-6}$, and $|r_{a,2}^{(new)} - r_{a,2}^{(init)}| < 10^{-6}$. Here, the distance function is defined as $\text{dist}(f^{(new)}, f^{(init)}) = \frac{(\sum_t (f^{(new)}(t) - f^{(init)}(t))^2)^{1/2}}{1 + (\sum_t f^{(init)}(t)^2)^{1/2}}$, as in Judd (1998). If any of them does not converge, update our initial guess and start again from Step 3:

$$\begin{aligned} \rho_{t,t+1}^{(init)} &= \eta \rho_{t,t+1}^{(new)} + (1 - \eta) \rho_{t,t+1}^{(init)}, \\ Y_t^{(init)} &= \eta Y_t^{(new)} + (1 - \eta) Y_t^{(init)}, \\ r_{a,2}^{(init)} &= \eta r_{a,2}^{(new)} + (1 - \eta) r_{a,2}^{(init)}, \end{aligned}$$

with $\eta = 0.99$.

D. Calibrating the mean fixed cost of automation

In our benchmark calibration, we assume that the log fixed costs of automation have a mean of zero because we do not have an additional data moment in the manufacturing sector to calibrate this parameter. To examine the robustness of our results, we now calibrate the mean of the log-normal distribution of the fixed costs (denoted by μ_a) by targeting a data moment in the whole economy. The moment that we target is the ratio of the robot use rate among firms between the 50th and 75th percentile of the employment distribution (1.7%) to the average robot use rate among all firms in the whole economy (2%), taken from the 2019 ABS documented by Acemoglu et al. (2022). This moment ($\frac{1.7}{2} = 0.85$) also captures the skewness of using robotics across U.S. firms.²

²This moment is only available for the whole economy and not for the manufacturing sector. Hence, we report these results here as a robustness check.

Table D.1

Parameters (calibrating the mean fixed cost of automation)

Parameter	Notation	Value	Matched Moments
Relative price of robots	Q_a	46.47	Fraction of automating firms
Mean of log automation fixed costs	μ_a	-0.32	Skewness of robot use rate
SD of log automation fixed costs	σ_a	3.09	Employment share of automating firms
Robot input weight	α_a	0.34	Robot density
Elasticity of substitution	η	2.03	Growth rate of robot density

Note: This table shows the calibrated parameters by moment matching. Compared to the benchmark model, we calibrate an additional parameter, which is the mean of the log-normal distribution of the fixed cost of automation (μ_a) by matching the skewness of robot use rate measured by the ratio of the robot use rate among firms between the 50th and the 75th percentile of the employment distribution to the average robot use rate among all firms in the whole economy in the ABS data documented by [Acemoglu et al. \(2022\)](#).

Table D.2

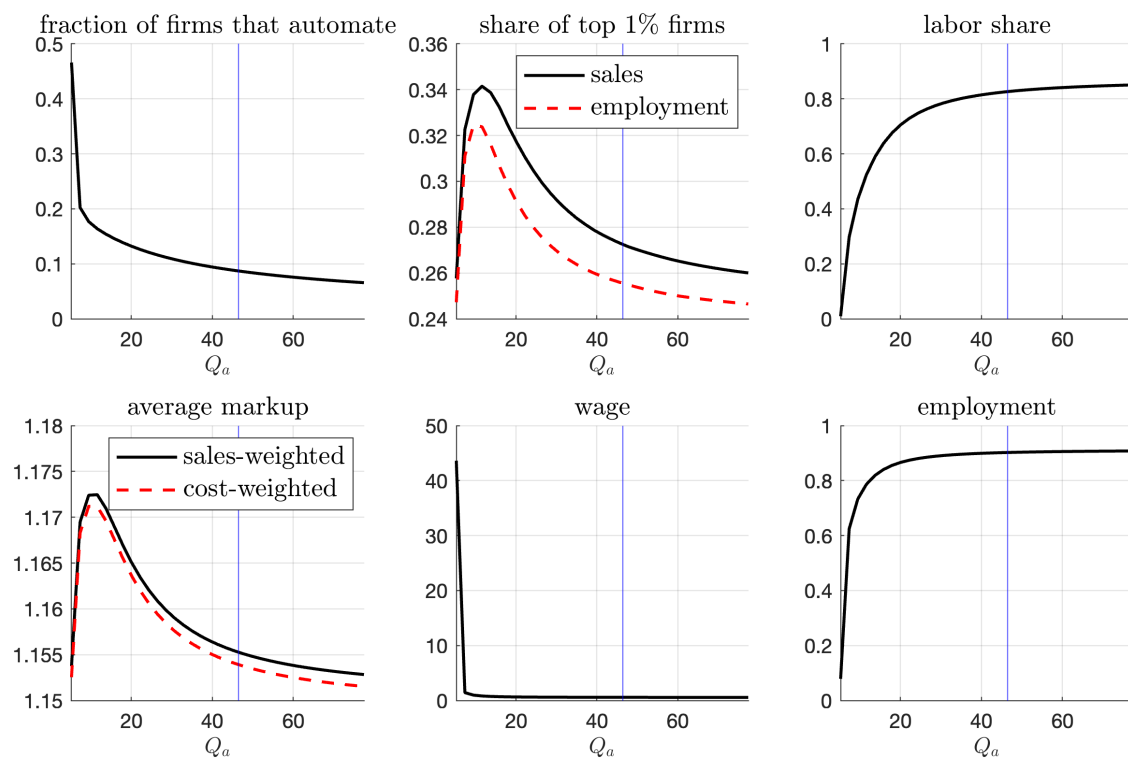
Matched Moments (calibrating the mean fixed cost of automation)

Moments	Data	Model
Fraction of automating firms	8.7%	8.7%
Skewness of robot use rate	0.85	0.85
Employment share of automating firms	45.1%	45.1%
Robot density	0.02	0.02
Growth rate of robot density	300%	300%

Note: This table shows the targeted data moments and the simulated moments by the model. The first three data moments are based on the ABS data (taken from [Acemoglu et al., 2022](#)), and the last two moments are authors' calculations using IFR and NBER-CES data. The skewness of robot use rate is measured by the ratio of the robot use rate among firms between the 50th and the 75th percentile of the employment distribution to the average robot use rate among all firms in the whole economy in the ABS data documented by [Acemoglu et al. \(2022\)](#).

Table D.1 presents the calibrated parameters. Table D.2 shows that the calibrated model exactly matches all the five moments in the data. In this calibrated model, as shown in Figure D.1, the predicted steady-state relations between the robot price and the macroeconomic variables are qualitatively similar to those in the benchmark model. Quantitatively, a 40% decline in the robot price raises the sales share of the top 1% of firms by about 1.23 percentage points (from 26% to 27.23%) and the employment share of the top 1% of firms by about 0.9 percentage points. Therefore, this model predicts that the decline in the robot price explains about 41% of the observed increases in sales concentration (1.23 out of 3 percentage points) and about 18% of the divergence between sales and employment concentration (0.32 out of the 1.8 percentage points). These magnitudes of the contributions from automation to industry concentration are slightly smaller than, but comparable to, those in the benchmark model. Thus, our main results are robust to calibrating the mean fixed cost of automation.

Figure D.1: Aggregate Variables (calibrated mean fixed cost of automation)



Note: This figure shows the effects of changes in the robot price Q_a on the fraction of firms that automate, the share of the top 1% of firms, the labor share, the average markup, the wage rate, and employment in the model with a calibrated value of the mean fixed cost of automation. The vertical blue line indicates the calibrated value of robot price Q_a .