

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Automation and the Rise of Superstar Firms

Hamid Firooz
University of Rochester

Zheng Liu
Federal Reserve Bank of San Francisco

Yajie Wang
University of Rochester

November 2023

Working Paper 2022-05

<https://www.frbsf.org/economic-research/publications/working-papers/2022/05/>

Suggested citation:

Firooz, Hamid, Zheng Liu, Yajie Wang. 2023 “Automation and the Rise of Superstar Firms,” Federal Reserve Bank of San Francisco Working Paper 2022-05.

<https://doi.org/10.24148/wp2022-05>

This paper was previously circulated under the title “Automation, Market Concentration, and the Labor Share.”

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Automation and the Rise of Superstar Firms*

Hamid Firooz[†] Zheng Liu[‡] Yajie Wang[§]

November 2023

Abstract

Using an instrumental variable approach, we document evidence that the rise in automation technology contributed to the rise of superstar firms. 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. Given the fixed cost, larger and more productive firms are more likely to automate. Automation boosts labor productivity, enabling those large automating firms to expand further, and thus raising industry concentration. Our calibrated model can replicate the highly skewed automation usage toward superstar firms observed in the Census data. Since robots can substitute for workers, increased automation raises sales concentration more than employment concentration, consistent with empirical evidence. In the model, automation raises aggregate productivity but exacerbates markup distortions. Our calibration suggests that a modest subsidy for automating firms improves welfare.

Keywords: Automation, industry concentration, superstar firms, markup, productivity.

JEL Codes: E24, L11, O33.

*We are grateful to David Autor, Yan Bai, Martin Beraja, 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, 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.

[†]University of Rochester; email address: hfirooz@ur.rochester.edu

[‡]Federal Reserve Bank of San Francisco; email address: Zheng.Liu@sf.frb.org

[§]University of Missouri; email address: yajie.wang@missouri.edu

1 Introduction

Industries in the United States have become increasingly concentrated, with each major sector increasingly dominated by a small number of superstar firms (Autor et al., 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.

The potential link between automation and industry concentration can be visualized from the time-series plots in Figure 1. The figure shows the shares of sales and employment of the largest firms within manufacturing industries (Panel A). The sales share of the top four firms (CR4) rose from about 40.5 percent in the late 1990s to about 43.5 percent in 2012, an increase of about 3 percentage points. The sales share of the top 20 firms (CR20) also increased during this period. The employment shares of the top firms, in contrast, stayed relatively flat. The rise in sales concentration coincides with the rise in automation, as Panel B of the figure shows. Since the early 2000s, the relative price of robots has declined by about 40 percent, and the number of industrial robots per thousand manufacturing employees has quadrupled.¹


Sales concentration has also increased in Europe. As documented by Bajgar et al. (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.

The correlations between automation and industry concentration are also present in cross-sectional data. We use firm-level data from Compustat to construct measures of industry concentration for two-digit manufacturing industries based on the North American Industry Classification System (NAICS). We find that robot density—measured by the operational stock of industrial robots per thousand workers in a given manufacturing industry—has a significantly positive correlation with sales concentration but a small and insignificant correlation with employment concentration.

We further establish causal effects of robot adoptions on industry concentration using an instrumental variable (IV) approach. Robot adoptions and industry concentration can both be endogenous. To address this endogeneity issue, we follow the approach of Acemoglu and Restrepo (2020) by using lagged values of the average robot density in five

¹Throughout this paper, we focus on industrial robots, which is a specific type of automation technology. We use “automation” and “robots” interchangeably.

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).

European economies as an IV for the U.S. robot density in our industry-level panel data regressions.² As documented by Acemoglu and Restrepo (2020), robot adoptions vary considerably across industries, with a common subset of industries in both the United States and Europe experiencing rapid robot adoptions in recent decades. Importantly, robot adoption trends in European markets have been ahead of the United States. Thus, those trends reflect global advancements in automation technologies, which in turn influence U.S. robot adoption patterns, indicating the relevance of the IV. Our identification is based on the assumption that the lagged robot density in those five European countries is related to U.S. industry concentration solely through the global automation technology progress. This methodology can filter out U.S.-specific factors that concurrently influence both robot adoptions and industry concentration, addressing potential omitted variable biases. Adding weight to this assumption, Acemoglu and Restrepo (2022) show that the faster robot adoption in Europe stems from its more rapidly aging population. This demographic trend, presumably unrelated to U.S. industry concentration, supports our IV’s exclusion restriction.

Our IV regressions show that robot adoptions have contributed significantly to the

²The five European economies include Denmark, Finland, France, Italy, and Sweden, which all adopted robotics ahead of the United States.

rise in sales concentration in the United States, but their effects on employment concentration have been small and statistically insignificant. Our evidence therefore implies that automation has contributed to the divergence between sales and employment concentration in the manufacturing sector. The effects of automation on sales concentration are economically important: a one standard deviation increase in robot density raises the sales share of the top 1% firms by about 10 percentage points, or about 34 percent relative to its average value.

To understand the economic mechanism that links automation to 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.

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 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 precisely 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 significantly raise the sales share of the top 1% firms, but the effects on the employment share of the top firms are small and insignificant. 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 et al., 2020; Acemoglu et al., 2022).³ 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 (2020), and Kehrig and Vincent (2021). 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 ABS survey documented by Zolas et al. (2020).

Our model implies that the relation between robot prices and industry concentration can be nonmonotonic. In a counterfactual scenario where the relative price of robots is sufficiently low, a further decline in the robot price could *reduce* the sales share of the top firms. When the robot price declines, some medium-sized firms would switch technologies from labor-only to automation. Although the drop in the robot price also benefits large and automating firms, when the robot price becomes sufficiently low, the usage of robots would become widespread, such that the expansion of the medium-sized firms through automation erodes the market share of the top 1% firms, reducing sales

³The ABS covers a large and nationally representative sample of over 850,000 firms in all private, nonfarm business sectors.

concentration. This counterfactual illustrates a crucial difference between automation equipment and general capital equipment: the usage of automation is heavily skewed towards a small number of superstar firms, whereas the usage of capital equipment is widespread. When automation becomes widespread (e.g., when the robot prices are sufficiently low), then a further drop in automation costs may not increase industry concentration and may even decrease it.

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. We examine the implications of taxing (or subsidizing) automation in our calibrated model for macroeconomic allocations and social welfare. On the one hand, taxing automating firms reallocates production from large, robot-using firms to smaller firms, reducing industry concentration, average markup, and markup dispersion. On the other hand, since larger firms are more productive, such reallocation reduces aggregate productivity. Overall, the tax policy faces a tradeoff between alleviating markup distortions and reducing aggregate productivity.

To quantify the economy-wide optimal size of robot taxes or subsidies, we re-calibrate the model to match the observed share of robot-using firms and the employment share of those firms in the whole economy, instead of the baseline calibration that matches those moments in the manufacturing sector. As an external validation, the calibrated model does well in matching the observed cross-sectional distribution of automation usage in the firm-level data (for the whole economy) documented by [Zolas et al. \(2020\)](#). In both the model and the data, automation usage is highly skewed towards a small number of very large firms. Given that our calibration does not target the entire distribution of automation usage, the model's ability to replicate this cross-sectional distribution lends credence to its underlying mechanism.

Finally, our calibrated model implies that a modest subsidy (about 2.5% of sales) for automating firms maximizes welfare, yielding a welfare gain equivalent to about 0.11% of steady-state consumption compared to the benchmark without policy interventions.

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 uneven productivity growth across firms (Furman and Orszag, 2018), declines in knowledge diffusion between the frontier and laggard firms (Akcigit and Ates, 2019), a slowdown in radical innovations since the 1990s (Olmstead-Rumsey, 2019), and the rise of specialized firms (Ekerdt and Wu, 2022).

Our study focuses on the rise in automation technologies as a driver of the rise of superstar firms. Existing studies show that automation can have important implications for employment, wages, and labor productivity (Acemoglu and Restrepo, 2018, 2020; Arnoud, 2018; Aghion et al., 2021; Graetz and Michaels, 2018; Leduc and Liu, 2019). 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, Lelarge and Restrepo, 2020; Bergholt, Furlanetto and Maffei-Faccioli, 2022). Our paper complements this literature by showing that automation plays a quantitatively important role in driving the rise of superstar firms.

Empirically, we document novel industry-level evidence that robot adoptions during the past two decades contributed significantly to the observed increases in sales concentration in the U.S. manufacturing sector. However, their effects on employment concentration are small and statistically insignificant. Our evidence suggests that the rise in automation can account for the divergence between sales and employment concentration observed in the manufacturing sector. This finding resonates with the findings of Hsieh and Rossi-Hansberg (2019), who find that employment concentration has remained flat or even declined in all but three broad sectors (services, wholesale, and retail) in the United States from 1977 to 2013, a period during which sales concentration in most sectors has steadily increased (Autor et al., 2020).

Theoretically, we highlight the economy-of-scale effect of robot adoptions through a quantitative general equilibrium framework. Our model underscores how fixed costs of automation disproportionately favor large, high-productivity firms, and thus raise sales concentration. Since robots substitute for workers, our model can also explain the divergence between sales and employment concentration observed in the manufacturing data. The economy-of-scale feature of new technology adoptions has been explored in existing studies, including, for example, Kwon, Ma and Zimmermann (2022), Aghion et al. (2019), Ridder (2023), Lashkari, Bauer and Boussard (2022), Tambe et al. (2020), and Sui (2022). In a closely related and parallel study, Hubmer and Restrepo (2022) present a task-based model of automation, featuring fixed costs of automating tasks, to examine how automation can contribute to the decline in the labor income share. Their model

implies that large firms automate more tasks following a decline in capital prices, while the median firm continues to operate a labor-intensive technology. This economy-of-scale effect reduces the average labor income share and raises sales concentration.

Our model shares the economy-of-scale perspective with this strand of literature. However, our goal is to provide a comprehensive, quantitative framework for understanding the empirical relations of automation with sales concentration and also with the divergence between sales and employment concentration. To our knowledge, our work is the first to document the causal empirical relation between automation and industry concentration, to quantify the contributions of the rise in automation to sales and employment concentration within a general equilibrium framework, and to use the same framework to study automation policies.

Our work also 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? For example, [Guerreiro, Rebelo and Teles \(2022\)](#) argue that steady declines in robot prices can lead to persistent increases in income inequality by displacing routine workers. To the extent that the current generation of routine workers cannot move to nonroutine occupations, optimal policy calls for taxing robots. [Acemoglu, Manera and Restrepo \(2020\)](#) argue that the U.S. tax system is biased against labor and in favor of capital and the current tax code has promoted inefficiently high levels of automation. Since the marginal automated tasks do not lead to large productivity gains while displacing workers, an automation tax helps reduce the level of automation towards the optimal level and also increases employment and welfare. [Beraja and Zorzi \(2022\)](#) study the implications of taxing automation in a heterogeneous agent framework with borrowing frictions. In their model, firms do not internalize that automation depresses the income of displaced workers who face borrowing constraints, especially during the early periods of the transition. Thus, slowing the speed of automation by tax policy can improve welfare. [Prettner and Strulik \(2020\)](#) examine how robot taxes can help redistribute income from high-skilled workers to low-skilled workers. [Costinot and Werning \(2022\)](#) argue that, while distributional concerns create a rationale for taxing robots and trade, the magnitude of these taxes may decrease as the process of automation and globalization deepens and inequality increases. Finally, [Thuemmel \(2022\)](#) finds that a robot subsidy is optimal when robots are relatively expensive; when robots become sufficiently cheap, it would be optimal to tax them.

Our model highlights a different tradeoff facing robot tax policies. Increases in automation disproportionately benefit large firms with high productivity and high markups. Following a decline in robot prices, the between-firm reallocation would raise both aggregate labor productivity and average markup. Taxing robots reduces the

markup distortions but also lowers the productivity gains from the reallocation induced by automation. 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 therefore 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.

3 Industry-level evidence

This section examines the empirical relation between automation and industry concentration for U.S. manufacturing industries. We first present ordinary least squares (OLS) evidence that automation (measured by robot density) positively correlates with industry concentration. The correlations of robot density with sales concentration are statistically significant and economically important, whereas the correlations with employment concentration are small and insignificant. We then use IV regressions to establish causal effects of robot adoptions on sales concentration. We further show that the effects of robot adoptions on employment concentration are small and statistically insignificant, suggesting that automation has also contributed to the divergence between sales and employment concentration in manufacturing.

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.⁴

We construct a measure of robot density for each two-digit industry using data on manufacturing employment and operation stocks of industrial robots from the International Federation of Robotics (IFR).⁵ We define robot density for industry j in year t as

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

⁴Using 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.

⁵According to the IFR definition, industrial robots are automatically controlled, reprogrammable, and multipurpose manipulators with several axes.

Table 1. Summary Statistics

	#obs	mean	min	p25	p50	p75	max	s.d.
robots/thousand employees	156	30.42	0.00	0.24	2.26	10.90	419.92	87.96
robots/million hours	156	19.58	0.00	0.18	1.72	7.72	243.54	52.42
top 1% share of sales	121	0.30	0.08	0.22	0.30	0.36	0.77	0.13
top 1% share of employment	106	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 operation 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 those industries with at least 10 firms.

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

For robustness, we also consider an alternative measure of industry-level robot density, defined as the operation 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.⁶ We obtain an unbalanced panel with 13 industries covering the 12 years from 2007 to 2018.⁷

Table 1 reports the summary statistics of variables. First, it shows that robot density varies widely in our sample. For example, the inter-quartile range (IQR) of robots per thousand workers is about 10, which is one-third of the sample mean. The standard deviation of robot density is also large—about three times the mean. These patterns reflect both within-industry changes in robot adoption over time and across-industry heterogeneity in robot adoption and the 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 an IQR of about 14 percent and a standard deviation of 13 percent. The employment share of the top 1% of firms averages about 27 percent and varies less than the sales share, with an IQR of about 11 percent and a standard deviation of about 8 percent.

⁶The 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.

⁷We 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. Our sample includes 13 industries, identified by their ISIC rev4 codes: 10-12, 13-15, 16&31, 17-18, 19-22, 23, 24, 25, 26-27, 28, 29, 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 $12 \times 13 = 156$ because of missing values in certain industry-year cells.

Table 2. OLS Regressions for Robot Density and Industry Concentration

	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)
Constant	0.295*** (0.022)	0.302*** (0.020)	0.267*** (0.039)	0.268*** (0.034)
Observations	117	117	104	104
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Note: This table shows the OLS regression results from the empirical specification (2) that projects the measures of industry concentration on robot density. 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 as the operation stock of industrial robots per thousand workers or million labor hours within the industry. 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.2 Correlations between automation and industry concentration

We calculate the correlations between automation and industry concentration, 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 at 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.

Table 2 reports the estimation results of the OLS regressions. Industries are weighted by their sales 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.

Table 2 shows that robot density is positively correlated with sales concentration (i.e., the sales share of the top 1% of firms), with the correlation 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 5.7 percentage points, or equivalently about 19 percent, above the sample mean (the average sales share of the top 1% of firms in our sample is 30%).⁸ 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 between robot density and 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)).⁹ These regression results from cross-sectional data corroborate well with the time-series correlations between automation and industry concentration illustrated in Figure 1.

3.3 Effects of automation on industry concentration

The correlations between robot density and industry concentration do not necessarily reflect causal effects, since both robot adoption and industry concentration are endogenous variables. An omitted variable bias can arise when a time-varying industry-level factor (such as industry-specific productivity) affects both robot density and concentration in the industry.

To examine how advancements in automation technology may have impacted industry concentration, we employ an IV approach. We follow [Acemoglu and Restrepo \(2020\)](#) and 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, Italy, and Sweden, which all adopted robotics ahead of the United States.¹⁰ 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

⁸The standard deviation of logged robot density is 2.71. 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 of 1% firms increases by $0.021 \times 2.71 \approx 5.69$ percentage points, or about 19 percent of the mean of the sales share.

⁹There are fewer observations for employment concentration than sales concentration due to a higher occurrence of missing employment data in Compustat.

¹⁰Following [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.

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.

As documented by [Acemoglu and Restrepo \(2020\)](#), robot adoptions vary considerably across industries, with a common subset of industries in both the United States and Europe experiencing rapid robot adoptions in recent decades. Importantly, robot adoption trends in those five European countries have been ahead of the United States. Thus, those trends reflect global advancements in automation technologies, which in turn influence U.S. robot adoption patterns, indicating the relevance of the IV. Our instrumental variable, derived from robot adoptions in foreign countries, is designed to estimate the causal impact of automation on U.S. industry concentration. Our identification is based on the assumption that the lagged robot density in EURO5 is related to U.S. industry concentration solely through the global automation technology progress. This methodology can filter out U.S.-specific factors that concurrently influence both robot adoptions and industry concentration, addressing potential omitted variable biases. Adding weight to this assumption, research by [Acemoglu and Restrepo \(2022\)](#) shows that the faster robot adoption in Europe stems from its more rapidly aging population. This demographic trend, presumably unrelated to U.S. industry concentration, supports our IV's exclusion restriction. We will further investigate several potential challenges to our identification assumption in [Section 3.4](#).

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. In the second stage, we regress our measures of U.S. industry concentration on the predicted logged robot density from the first stage.

[Table 3](#) displays the IV estimation results. The estimation shows that the quantitative effects of automation on sales concentration are statistically significant and economically important (Columns (1) and (2)). A one standard deviation increase in robot density raises the sales share of the top 1% firms by about 10 percentage points, or equivalently, about 34 percent relative to its sample average value (which is about 30%).¹¹ This

¹¹The logged robot density in the U.S. industries has a standard deviation of 2.71. Thus, the estimation shown in [Table 3](#) implies that a one standard deviation increase in robot exposure raises the sales share of

Table 3. IV Regressions for Robot Density and Industry Concentration

	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 p -value	0.000	0.001	0.474	0.401

Note: This table shows the second-stage results of the IV regression from the empirical specification (2). The dependent variables are the sales share (first two columns) and employment share (last two columns) of the top 1% of firms. The robot density is measured as the operation 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 p -values of Anderson-Rubin weak instrument robust tests adjusted for heteroskedasticity. 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

number is higher than the 6 percentage points (or 19 percent) obtained from the OLS estimation (Table 2), suggesting that omitted variables lead to a downward bias of the coefficient in the OLS regressions. In comparison, the estimated effects of automation on employment concentration are small and statistically insignificant (Columns (3) and (4)). Our evidence therefore implies that automation has contributed to the rise in sales concentration as well as the divergence between sales and employment concentration in the manufacturing sector.

Our estimation and inferences are robust to potentially weak instruments, as indicated by the Anderson-Rubin (AR) test (Anderson and Rubin, 1949). In the last row of Table 3, we report the p -values of the AR test adjusted for heteroskedasticity. The p -values indicate that the estimated effects of robot density on sales concentration are robust to weak instruments at the 99% confidence level, and the effects on employment concentration are not significant, with a p -value of the AR test larger than 0.40.¹²

the top 1% firms by $0.038 \times 2.71 \approx 10.30$ percentage points. In our sample, the average sales share of the top 1% firms is about 30%. Thus, our estimation suggests that a one standard deviation increase in robot density would raise the sales share of the top 1% firms by about 34 percent.

¹²The 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).

3.4 Potential challenges for identification

To identify the effects of automation on market concentration in the United States, we follow [Acemoglu and Restrepo \(2020\)](#) and use lagged robot density in five European countries as an IV for the robot density in the U.S. manufacturing industries. However, there might be concerns about the validity of this IV.

First, some unobserved shocks that are unrelated to the diffusion of automation technologies and are common to the U.S. and European countries might be influencing U.S. industry concentration. Such unobserved shocks, if important, could potentially invalidate the exclusion restrictions of our IV. Responding to this concern, [Acemoglu and Restrepo \(2020\)](#) show that the trend of European robot adoptions is largely uncorrelated with other major global trends, such as import competition, offshoring, declines of routine jobs, and capital deepening. Thus, the trend of European robot adoptions captures global advancement in automation technologies, which in turn drives U.S. robot adoptions as well.

Second, the development of other labor-substituting technologies such as AI can also potentially distort our IV regressions, especially if those technologies exhibit variations across industries and over time that are highly correlated with the observed variations in robotics. This concern, however, is not substantiated by the micro-level data. The evidence presented by [Acemoglu et al. \(2022\)](#) based on the 2019 ABS suggests that the conditional adoption rates of AI and robotics are weakly correlated. Specifically, conditional on adopting AI, the share of firms that also adopt robotics is under 20%; conditional on using robotics, the share of firms that use AI is also modest, at about 30% (see their Table 2).

Third, some large U.S. firms in our Compustat sample report global sales. An increase in robot adoptions in the European countries might increase sales of those U.S. firms in the European market, which in turn raise U.S. industry concentration through global sales, rather than through U.S. robot adoptions. If this global sales channel is important, then it might invalidate the exclusion restrictions of our IV. In the data, however, the share of sales of U.S. affiliates in the European market has been very 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). Thus, sales of foreign affiliates in EURO5 are unlikely an important force that drives U.S. industry concentration. To further isolate the global sales channel, we estimated a version of the IV regressions using domestic sales (i.e., total sales minus exports) instead of total sales

and obtained similar results (see Appendix Table A.2).¹³

Finally, there might have been an unobserved underlying trend in the U.S. industry concentration prior to the diffusion of robotics, such that the observed relation between robot density and industry concentration in our sample reflects the continuation of the pre-trends in industry concentration rather than the effects of robot adoptions. To address this concern, we re-estimate the baseline regression, using five-year lags of the concentration measures as the dependent variable. We obtain statistically insignificant estimates, as detailed in Appendix Table A.3. Thus, the increases in industry concentration in our sample were not a by-product of pre-existing trends before the expansion of robotics.

4 The Model

To understand the empirical connection 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, 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

$$C_t + v_t s_{t+1} \leq W_t N_t + (v_t + d_t) s_t, \quad (5)$$

where s_t denotes the equity share of firms held by the household, v_t denotes the equity price, d_t denotes the dividend flow, and W_t denotes the real wage rate. The household takes W_t and v_t as given and maximizes the utility function (4) subject to the budget

¹³A related concern about the exclusion restrictions in our IV regressions is that robot adoptions in EURO5 could drive out of the market the least competitive U.S. firms, leading to an increase in U.S. concentration. However, these effects are likely small since the average size of the EURO5 is very small relative to the U.S. economy.

constraints (5). The optimizing consumption-leisure choice implies the labor supply equation

$$W_t = \chi N_t^\xi C_t. \quad (6)$$

The optimizing decision for equity share holdings is given by

$$v_t = \rho_t(v_{t+1} + d_{t+1}), \quad (7)$$

where $\rho_t \equiv \beta \frac{C_t}{C_{t+1}}$ is the stochastic discount factor. We will be focusing on the steady state of the model and therefore $\rho_t = \beta$.

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 is produced using a bundle of intermediate goods $y(j)$, according to the Kimball aggregator

$$\int_0^1 \Lambda\left(\frac{y_t(j)}{Y_t}\right) dj = 1. \quad (8)$$

For ease of notation, we suppress the time subscript t in what follows.

4.3 Demand for intermediate goods

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

$$p(j) = \Lambda'(q(j))D, \quad (9)$$

where D is a demand shifter given by

$$D = \left(\int \Lambda'(q(j))q(j) dj \right)^{-1}. \quad (10)$$

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

$$\Lambda(q) = 1 + (\sigma - 1) \exp\left(\frac{1}{\varepsilon}\right) \varepsilon^{\frac{\sigma}{\varepsilon} - 1} \left[\Gamma\left(\frac{\sigma}{\varepsilon}, \frac{1}{\varepsilon}\right) - \Gamma\left(\frac{\sigma}{\varepsilon}, \frac{q^{\varepsilon/\sigma}}{\varepsilon}\right) \right], \quad (11)$$

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

$$\Gamma(s, x) = \int_x^\infty v^{s-1} e^{-v} dv. \quad (12)$$

Under the specification (11), we obtain

$$\Lambda'(q(j)) = \frac{\sigma - 1}{\sigma} \exp\left(\frac{1 - q(j)^{\frac{\varepsilon}{\sigma}}}{\varepsilon}\right), \quad (13)$$

which, using the demand schedule (9), implies that the demand elasticity (i.e., price elasticity of demand) faced by firm j is

$$\sigma(q(j)) = -\frac{\Lambda'(q(j))}{\Lambda''(q(j))q(j)} = \sigma q(j)^{-\frac{\varepsilon}{\sigma}}. \quad (14)$$

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

$$\mu(j) = \frac{\sigma(q(j))}{\sigma(q(j)) - 1}. \quad (15)$$

As a result, larger firms face lower demand elasticities, have more market power, and charge higher markups.¹⁴

4.4 Intermediate goods producers

Intermediate producers, from now on indexed by their idiosyncratic productivity ϕ , 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 , to be elaborated below. The production function takes the CES form

$$y = \phi \left[\alpha_a A'^{\frac{\eta-1}{\eta}} + (1 - \alpha_a) N^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (16)$$

where y denotes the firm's output; N denotes the inputs of workers; and $A' \geq 0$ denotes the end-of-period robot stock. The labor-only technology corresponds to the special case with $A' = 0$. 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.

¹⁴We make the technical assumption that $q(j) < \sigma^{\frac{\varepsilon}{\sigma}}$ 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

$$\ln \phi' = \gamma \ln \phi + \varepsilon, \quad \varepsilon \sim N(0, \sigma_\phi^2), \quad (17)$$

where ϕ' is next period productivity, $\gamma \in (0, 1)$ 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 draws s from the *i.i.d.* distribution $F(\cdot)$ and needs to pay the per-period cost $s\phi$ if it uses robots in production.¹⁵ We further assume that the distributions of s and ϕ are independent. A firm with the realized productivity ϕ and existing robot stock A that draws a fixed cost s chooses the price p and quantity y of its differentiated product, labor input N , and robot investment I_a to solve the dynamic programming problem

$$V(\phi, A; s) = \max_{p, y, N, I_a \geq (\delta_a - 1)A} \left[py - WN - Q_a I_a - s\phi \mathbb{1}\{A' > 0\} + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right], \quad (18)$$

where $\mathbb{1}\{x\}$ equals one if x holds and zero otherwise. The firm hires workers in a competitive labor market at a competitive real wage rate (W). The firm also chooses automation investment by purchasing I_a units of robots at the competitive price Q_a . Newly purchased robots add to the existing stock of robots, and robots depreciate at the constant rate $\delta_a \in (0, 1)$. The firm's stock of robots evolves according to the law of motion

$$A' = (1 - \delta_a)A + I_a. \quad (19)$$

Notice that we assume that the newly purchased robots can be used in the production process in the same period.

The firm solves the recursive problem (18) subject to the production function (16), the robot law of motion (19), and the demand schedule (9). Since robot operation incurs a fixed cost, a firm facing a sufficiently high s relative to its productivity would choose to sell its robots (i.e., by setting $A' = 0$) at the market price Q_a . In that case, we would have $I_a = (\delta_a - 1)A \leq 0$.

Appendix B shows that the recursive problem (18) can be simplified to

$$V(\phi, A; s) = Q_a(1 - \delta_a)A + \max\{V^a(\phi) - s\phi, V^n(\phi)\}, \quad (20)$$

¹⁵Assuming 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.

where the continuation value of operating the automation technology this period (i.e., having $A' > 0$) is given by

$$V^a(\phi) = \max_{p,y,N,A'>0} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right], \quad (21)$$

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

$$V^n(\phi) = \max_{p,y,N} \left[py - WN + \beta E_{\phi'|\phi} \int_{s'} V(\phi', 0; s') dF(s') \right]. \quad (22)$$

Firms with automation technology in (21) optimally choose their production inputs N and A' given their production y . As Appendix B shows, the first-order conditions imply

$$\gamma_a = \alpha_a \lambda_a(\phi) \phi^{\frac{\eta-1}{\eta}} \left(\frac{y}{A'} \right)^{\frac{1}{\eta}}, \quad (23)$$

$$W = (1 - \alpha_a) \lambda_a(\phi) \phi^{\frac{\eta-1}{\eta}} \left(\frac{y}{N} \right)^{\frac{1}{\eta}}, \quad (24)$$

where $\gamma_a \equiv Q_a[1 - \beta(1 - \delta_a)]$ denotes the effective user cost of robots, and $\lambda_a(\phi)$ denotes the marginal cost of production for a firm with productivity ϕ operating the automation technology:

$$\lambda_a(\phi) = \frac{\left[\alpha_a^\eta \gamma_a^{1-\eta} + (1 - \alpha_a)^\eta W^{1-\eta} \right]^{\frac{1}{1-\eta}}}{\phi}. \quad (25)$$

Moreover, firms operating the labor-only technology in (22) choose their labor input N given their production y :

$$N = \frac{y}{\phi} (1 - \alpha_a)^{\frac{\eta}{1-\eta}}, \quad (26)$$

The marginal cost of production in this case would be

$$\lambda_n(\phi) = \frac{(1 - \alpha_a)^{\frac{\eta}{1-\eta}} W}{\phi}. \quad (27)$$

Notice that, given the productivity ϕ , the marginal cost of production using the labor-only technology is always larger than that using the automation technology, i.e., $\lambda_a(\phi) \leq \lambda_n(\phi)$.

The problem (20) implies that firms choose to operate the automation technology (i.e.,

to have $A' > 0$) if and only if their draw of the fixed automation cost is small enough:

$$s \leq s^*(\phi) \iff \mathbb{I}_a(\phi, s) = 1, \quad (28)$$

where $\mathbb{I}_a(\cdot)$ is an indicator of the automation decision, which is a function of the firm-level variables ϕ and s , and the cutoff fixed cost equals:

$$s^*(\phi) \equiv \frac{V^a(\phi) - V^n(\phi)}{\phi}. \quad (29)$$

It follows that, for a firm with productivity ϕ , the ex ante (i.e., before drawing the automation fixed cost) automation probability equals $F(s^*(\phi))$, which is the cumulative density of the fixed costs 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,

$$s^*(\phi) = \frac{\pi^a(\phi) - \pi^n(\phi)}{\phi}, \quad (30)$$

where

$$\pi^a(\phi) = \max_{p, y, N, A'} \left[py - WN - Q_a[1 - \beta(1 - \delta_a)]A' \right], \quad (31)$$

subject to the demand schedule (9) and production function (16), and

$$\pi^n(\phi) = \max_{p, y, N} \left[py - WN \right]. \quad (32)$$

subject to the same demand schedule and production function with $A' = 0$.

4.5 Stationary equilibrium

We focus on the stationary equilibrium and thus drop the time subscript for all variables. The world robot price Q_a is exogenously given. The equilibrium consists of aggregate allocations C , I_a , A , N , and Y , wage rate W , firm-level allocations $A'(\phi)$, $I_a(\phi)$, $N(\phi)$, and $y(\phi)$, and firm-level prices $p(\phi)$ for all $\phi \in G(\cdot)$, where $G(\cdot)$ denotes the ergodic distribution implied by the productivity process (17), such that (i) taking W as given, the aggregate allocations C and N solve the representative household's optimization problem; (ii) taking W and Y 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

$$C + Q_a I_a + \int_{\phi} \int_0^{s^*(\phi)} s \phi dF(s) dG(\phi) = Y. \quad (33)$$

The labor market clearing condition is given by

$$N = \int_{\phi} N(\phi) dG(\phi). \quad (34)$$

The stock of robots is given by

$$A = A' = \int_{\phi} A'(\phi) F(s^*(\phi)) dG(\phi). \quad (35)$$

Total investment in robots equals

$$I_a = A' - (1 - \delta_a)A = \delta_a A = \delta_a \int_{\phi} A'(\phi) F(s^*(\phi)) dG(\phi). \quad (36)$$

Appendix C outlines the computational algorithm to solve the model.

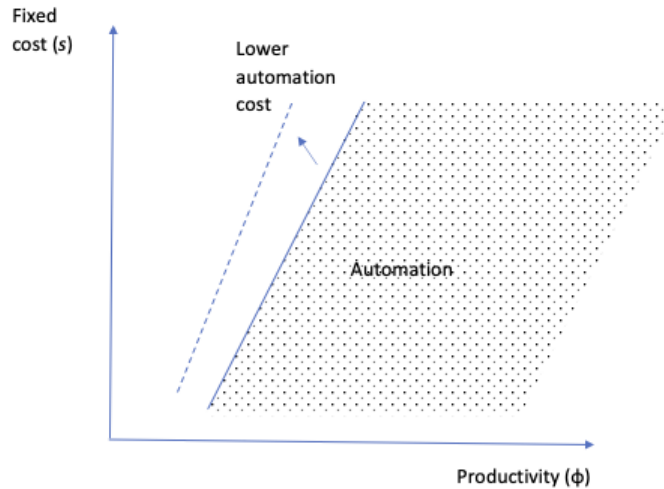
5 Model mechanism

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.¹⁶

Figure 2 illustrates the automation decision rules. 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

¹⁶As we will show in Section 7, while automation fixed costs are proportional to firm-level productivity, more productive firms are still more likely to automate.

Figure 2. Automation Decision Rules



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.

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.

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. This would shift the indifference curve up (from the solid to the dashed line), such that more firms would choose to automate (the extensive margin) and those firms already operating the automation technology would increase their use of 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 2. 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 (through 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 vs.

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 for explaining the observation that automation’s positive impact on sales concentration is stronger than its effects on 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 model parameters to match moments in the manufacturing data. 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).¹⁷ 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 4 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

¹⁷The usage of robotics in the U.S. manufacturing sector is more prevalent than that of other advanced technologies. For instance, according to the 2019 ABS survey, the fraction of manufacturing firms that use AI technologies and their employment share are 3.2% and 22.6%, respectively, much smaller than those for robotics (8.7% and 45.1%, respectively).

Table 4. Parameters

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	10.86	Edmond, Midrigan and Xu (2021)
Super-elasticity	ϵ/σ	0.16	Edmond, Midrigan and Xu (2021)
Panel B: Parameters calibrated to match moments in data			
Relative price of robots	Q_a	49.39	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

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

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 operation stocks of industrial robots. We set the persistence of idiosyncratic productivity shocks to $\gamma = 0.95$ following Khan and Thomas (2008). We set the standard deviation of productivity shocks to $\sigma_\phi = 0.1$, according to the estimation by Bloom et al. (2018).¹⁸ 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 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 4.

We target four moments to jointly calibrate these four parameters. The four moments

¹⁸Bloom 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 ($=0.051*68.7\%+0.209*(1-68.7\%)$).

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 , small firms would be less likely to cover the fixed cost of automation. As a result, the employment-weighted robot use rate would rise. 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 4 reports these parameters that we internally calibrated. The calibrated model matches the targeted data moments exactly, as shown in Table 5.¹⁹

¹⁹To put the calibrated elasticity η into context, we note that Cheng et al. (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. The results are available upon request.

Table 5. 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 shows 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.

7 Model implications

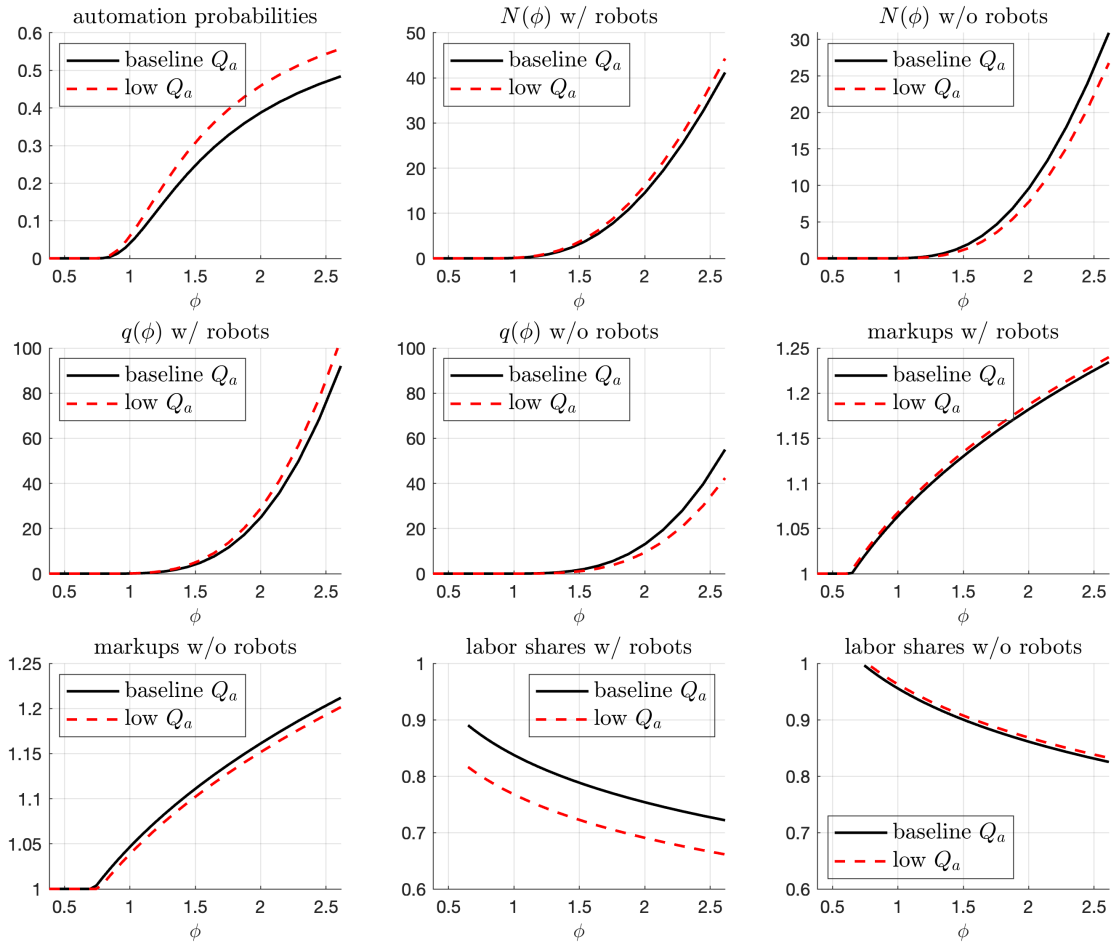
We solve the model’s steady-state equilibrium based on the calibrated parameters. We now report the model’s quantitative implications.

7.1 Firm-level implications

Figure 3 shows the firms’ decision rules as a function of the idiosyncratic productivity level ϕ , in both the baseline 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 since 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 baseline model, the decision rules are qualitatively similar between the two types of firms. In particular, higher-productivity firms are larger, with higher employment ($N(\phi)$) and relative output ($q(\phi)$), larger market power measured by 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.


Figure 3. 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 baseline calibration, whereas red-dashed lines show the results for a counterfactual in which robot price Q_a falls by 40%.

The red dashed lines in Figure 3 further show that the impacts of a decline in the robot price on the firms' decision rules depend on whether the firm uses 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, these firms substitute away from workers to robots, which tends to reduce employment at these firms. On the other hand, however, by adopting more robots, labor productivity at these firms rises, leading to an increase in labor demand to gain market share. Under our calibration, the latter effect dominates such that automating firms increase employment following the reduction in the robot price. This result is in line with the evidence documented in the ABS survey of Zolas et al. (2020). The labor shares of the automating firms decline despite the increases in their employment, reflecting the substitution of robots

Figure 4. Aggregate Variables



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. The vertical blue line indicates the calibrated value of robot price Q_a .

for workers and also the increase in markups as output increases.

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 the non-automating firms become smaller, their market power declines, resulting in lower markups and higher labor shares.

7.2 Aggregate implications

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 4. For illustration, we consider a wide range of the robot price around the calibrated value of $Q_a = 49.39$, indicated by the vertical blue line in the figure.

At a lower robot price, more firms would 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.²⁰ 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. We discuss the quantitative importance of automation in raising the concentration in the manufacturing sector below.

As Q_a falls, large firms become even larger, raising the average markup in the economy (both sales- and cost-weighted).²¹ As Figure 3 shows, a reduction in Q_a reallocates production and employment toward automating firms that have lower labor shares in the original steady state. Therefore, as Q_a falls, the labor share in the aggregate economy declines. Our model thus implies that declines in the aggregate labor share and increases in the average markup are mainly driven by the between-firm reallocation channel, in line with the empirical evidence in Autor et al. (2020) and Acemoglu, Lelarge and Restrepo (2020).

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, subsequently raising labor demand and bidding up real wages facing all firms. 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 substantial reduction in robot prices leads to an increase in wages and a decline in aggregate employment.²²

²⁰As discussed before, while automation fixed costs are proportional to firm-level productivity, more productive firms are still more likely to automate, as shown in the top-left panel in Figure 3.

²¹To derive the cost-weighted average markup, we use total variable costs at each firm, as in Edmond, Midrigan and Xu (2021).

²²Our 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

Automation and industry concentration. The top-middle panel in Figure 4 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 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 82.32 to its calibrated value of 49.39, 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 1.²³ We then examine the extent to which the resulting changes in industry concentration in the model can account for the actual changes observed 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 (see Figure 1), while employment concentration rose from 33.26% to 34.51% during the same period, an increase of about 1.2 percentage points.²⁴ 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 4 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 original 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 effects of automation on sales concentration in the U.S. that

Leduc and Liu (2019), 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 understanding the connection between automation and a broader set of labor market variables (including wages). We leave that important task for future research.

²³The 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%).

²⁴Notice 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.

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. Indeed, as illustrated by the top-middle panel in Figure 4, our model suggests that a decline in the price of general equipment that is widely used in the economy could decrease, rather than increase, industry concentration. In contrast, our model implies that a reduction in the price of any production input that is used by a small fraction of large firms would increase industry concentration.

7.3 Policy analysis

The rapid rise of the automation technology and the accompanying increase in industry concentration has stimulated ongoing policy debates on the efficacy of taxing automation. While it is argued that taxing automation might protect jobs for workers, it might also reduce labor productivity and lower labor demand. Moreover, taxing automation in a variable-markup world might seem attractive since it would reduce the market power of large, automating firms. In this section, we investigate the aggregate effects of taxing/subsidizing automation and examine welfare effects of such policies.

7.3.1 Calibrated model for the whole economy

For policy analysis, we re-calibrate the model to the whole economy. In particular, we calibrate Q_a and σ_a to match the fraction of firms using robots in the whole U.S. economy (2%) and the employment share of these firms (15.7%), both obtained from [Acemoglu et al. \(2022\)](#). Since the IFR data coverage beyond the manufacturing sector is limited, we do not have reliable data moments for robot density and the growth rate of robot density for the whole economy. We therefore set the parameters α_a and η to those calibrated in Section 6. We also keep the values of our externally calibrated parameters the same as in our baseline calibration in Panel A of Table 4.

Appendix Tables A.4 and A.5 report calibrated parameters and matched moments, respectively. The model is again able to match the targeted moments precisely.

As an external validation, the model calibrated to the whole economy does well in replicating the highly skewed distribution of firm-level automation usage in the ABS

data, a moment that we do not target in the calibration. Figure 5 plots the distribution of AI use rate (i.e., the fraction of firms that use AI in their production) in the ABS data documented by Zolas et al. (2020) (the bars), along with the model-predicted share of firms that use the automation technology (the line), as a function of firm size based on employment.²⁵ The figure shows that the model closely matches this non-targeted distribution of automation. In both the data and our model, automation usage is highly skewed toward the few largest firms. In this sense, automation is quite different from general capital equipment, the usage of which is widespread.²⁶

The ability of the model to correctly predict this non-targeted distribution lends credence to the model’s mechanism. Being consistent with the highly skewed distribution of automation usage, the model is capable of generating the observed sharper increases in sales concentration than in employment concentration when automation cost falls, as we showed above.

7.3.2 Taxing/Subsidizing automation.

We now use the model calibrated to the whole economy to examine the effects of taxing/subsidizing automation. To this end, we first introduce into the model a flat sales tax τ on firms that use the automation technology. The intermediate producers’ problem in equation (18) therefore becomes:

$$V(\phi, A; s) = \max_{p, y, N, I_a} \left[(1 - \tau \mathbb{1}\{A' > 0\}) p y - W N - Q_a I_a - s \phi \mathbb{1}\{A' > 0\} + \beta E_{\phi' | \phi} \int_{s'} V(\phi', A'; s') dF(s') \right]. \quad (38)$$

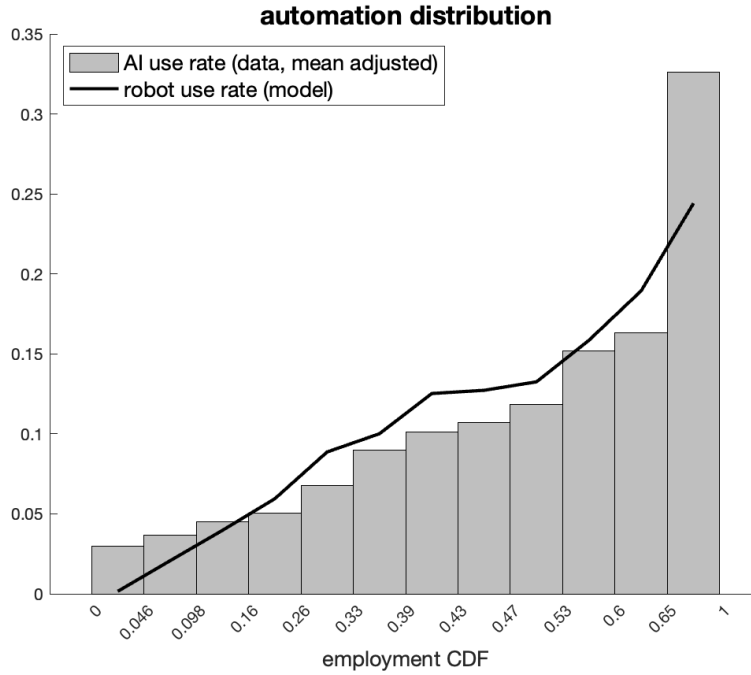
We assume that the tax revenue is rebated to consumers in a lump-sum fashion.

To explore the welfare implications of this policy, we compute the consumption equivalent variation as follows. Denote by $W(\tau)$ the social welfare in the economy with the automation tax rate τ . We measure the welfare losses (or gains) under the automation

²⁵Zolas et al. (2020) report the share of firms that use AI technologies across detailed size categories, e.g., 1-4 employees, 5-9 employees, or 10,000+ employees (see their Figure 8). To make this data comparable to our model, we convert the size bins into the cumulative density function (CDF) of employment, using the number of employees in each firm size category in the 2017 County Business Patterns and Economic Census. We then plot the AI usage rate across the employment CDF. Consistently, we calculate the robot usage rates across the employment CDF in the model using the same method. Note that AI is more commonly used than robots in the whole economy, and our focus is on the dispersion rather than the mean of these technologies. To ensure a fair comparison between the data and the model, Figure 5 scales the AI use rates in the data to have the same mean as that of the robot use rates in our model.

²⁶Acemoglu et al. (2022) document that the distribution of robot usage rates across firms is similar to that of AI usage rates, both skewed toward very large firms (see their Figure 3). However, they do not report more granular robot usage rates for firms within the top percentile. This is why we compare the distribution of robot usage rates in the model to the distribution of AI usage rates in the data documented by Zolas et al. (2020), who do report AI usage data within the top percentile of firms.

Figure 5. Automation Distribution



Note: This figure plots the distribution of AI usage rate (i.e., the fraction of firms that use AI in their production) in the ABS data (bars) and the fraction of firms using robots in the model (line), for firms of different sizes measured by the cumulative density of employment. The AI usage rates are scaled to have the same mean as that of the robot usage rates in the model.

tax by the percentage changes in consumption in perpetuity that are required such that the representative household is indifferent between living in the economy with the tax and the benchmark economy without the tax.

Specifically, the welfare in the economy with the tax rate τ is given by

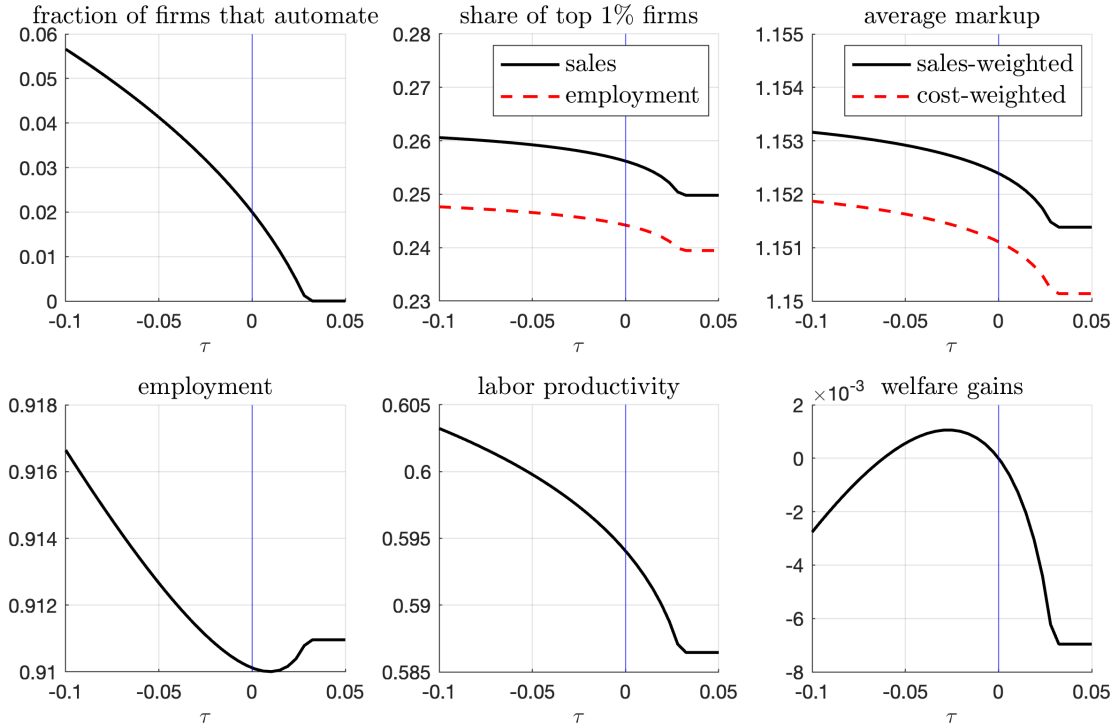
$$W(\tau) = \sum_{t=0}^{\infty} \beta^t \left[\ln C_t(\tau) - \chi \frac{N_t(\tau)^{1+\xi}}{1+\xi} \right], \quad (39)$$

where $C_t(\tau)$ and $N_t(\tau)$ are consumption and employment in the equilibrium with automation taxes. The welfare level in the benchmark economy without tax is given by

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

where $C_t(0)$ and $N_t(0)$ are consumption and employment in the equilibrium of the benchmark economy without automation taxes (i.e., with $\tau = 0$). The welfare losses associated with the tax rate τ is given by the consumption equivalent μ , which is defined

Figure 6. Effects of Taxing Automation



Note: This figure shows the aggregate effects of imposing a sales tax τ on firms that use automation technology.

by the relation

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

Solving for μ from Eq. (41), we obtain

$$\mu = 1 - \exp[(1 - \beta)(W(\tau) - W(0))]. \quad (42)$$

A positive (negative) μ would imply a welfare loss (gain) with automation taxes relative to the laissez-faire equilibrium without taxes.

Figure 6 shows the aggregate effects of imposing a sales tax on automating firms. Since only 2% of firms use automation technology in the calibrated economy, even a small automation tax of around 3% could drive the mass of automating firms down to zero, after which increasing the automation tax would have no effects.

A reduction in the robot tax τ works through a similar mechanism as a decline in the robot price in affecting industry concentration. A decline in τ raises the fraction of automating firms, with two competing effects on the market share of superstar firms (i.e., the top 1%). A lower τ disproportionately benefits large firms that use the automation

technology, enabling them to become even larger. At the same time, however, a lower τ also induces some medium-sized firms that initially use the labor-only technology to switch to the automation technology, reducing the market share of the top firms. Under our calibration, the former effect dominates, such that a lower robot tax (or a higher robot subsidy) raises the sales concentration and the employment concentration, although it also enlarges the gap between the two measures of concentrations because robots substitute for workers.

Since large firms have higher markups, the between-firm reallocation associated with a decline in robot taxes raises the average markup in the economy. At the same time, since larger firms have higher productivity, the reallocation also raises aggregate labor productivity.

In our model, changes in the robot tax have non-monotonic effects on employment. Starting from a high level of τ , lowering it would reduce employment, reflecting the labor-substituting effects of automation. Starting from a sufficiently low level of τ , lowering it further (or raising the robot subsidy) can raise aggregate employment because the increase in aggregate productivity boosts labor demand, which dominates the labor-substituting effects.

Reducing the robot tax has also non-monotonic effects on welfare, as shown in Figure 6, reflecting the tradeoff between labor productivity gains and markup distortions. Under our calibration, there is an interior optimum rate of robot subsidy, at about 2.5%, which maximizes welfare, with a maximum welfare gain of about 0.11 percent of steady-state consumption equivalent relative to the laissez-faire benchmark.

8 Conclusion

We have documented empirical evidence that automation has contributed to the rise in sales concentration and the divergence between sales and employment concentration since the early 2000s. We use a general equilibrium framework to show that this empirical relation between automation and industry concentration can be explained by an economy-of-scale effect stemming from fixed costs of operating the automation technology. In line with firm-level evidence, our calibrated model predicts a highly skewed distribution of automation usage toward a small number of superstar firms. Under our calibration, 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 in U.S. manufacturing 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.

In our model, taxing automating firms faces a tradeoff between productivity gains and markup distortions, because the tax policy induces reallocation from large automating firms with higher productivity and higher markups toward small firms. Given this tradeoff, our calibration suggests that a modest subsidy for automation improves welfare relative to the laissez-faire equilibrium.

References

- Acemoglu, Daron, and Pascual Restrepo.** 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6): 1488–1542.
- Acemoglu, Daron, and Pascual Restrepo.** 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6): 2188–2244.
- Acemoglu, Daron, and Pascual Restrepo.** 2021. "Tasks, Automation, and the Rise in US Wage Inequality." National Bureau of Economic Research Working Paper 28920.
- Acemoglu, Daron, and Pascual Restrepo.** 2022. "Demographics and Automation." *The Review of Economic Studies*, 89(1): 1–44.
- Acemoglu, Daron, Andrea Manera, and Pascual Restrepo.** 2020. "Does the US Tax Code Favor Automation?" National Bureau of Economic Research Working Paper 27052.
- Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo.** 2020. "Competing with Robots: Firm-Level Evidence from France." *AEA Papers and Proceedings*, 110: 383–88.
- Acemoglu, Daron, Gary W Anderson, David N Beede, Cathy Buffington, Eric E Childress, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas.** 2022. "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey." National Bureau of Economic Research Working Paper 30659.
- Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li.** 2019. "A Theory of Falling Growth and Rising Rents." National Bureau of Economic Research Working Paper 26448.
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel.** 2021. "What Are the Labor and Product Market Effects of Automation? New Evidence from France." Sciences Po.
- Akcigit, Ufuk, and Sina T Ates.** 2019. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." National Bureau of Economic Research Working Paper 25755.
- Anderson, Theodore W, and Herman Rubin.** 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *The Annals of mathe-*

matical statistics, 20(1): 46–63.

- Andrews, Isaiah, James H. Stock, and Liyang Sun.** 2019. “Weak Instruments in Instrumental Variables Regression: Theory and Practice.” *Annual Review of Economics*, 11(1): 727–753.
- Arnoud, Antoine.** 2018. “Automation Threat and Wage Bargaining.” Unpublished Manuscript, Yale University.
- Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen.** 2020. “The Fall of the Labor Share and the Rise of Superstar Firms.” *The Quarterly Journal of Economics*, 135(2): 645–709.
- Autor, David H., and Anna Salomons.** 2018. “Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share.” *Brookings Papers on Economic Activity*, 1: 1–87.
- Autor, David H., David Dorn, and Gordon H. Hanson.** 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review*, 103(6): 2121–2168.
- Autor, David H., Frank Levy, and Richard J. Murnane.** 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *The Quarterly Journal of Economics*, 118(4): 1279–1333.
- Bajgar, Matej, Giuseppe Berlingieri, Sara Calligaris, Chiara Criscuolo, and Jonathan Timmis.** 2019. “Industry Concentration in Europe and North America.” OECD Productivity Working Papers.
- Beraja, Martin, and Nathan Zorzi.** 2022. “Inefficient Automation.” National Bureau of Economic Research Working Paper 30154.
- Bergholt, Drago, Francesco Furlanetto, and Nicolò Maffei-Faccioli.** 2022. “The Decline of the Labor Share: New Empirical Evidence.” *American Economic Journal: Macroeconomics*, 14(3): 163–98.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry.** 2018. “Really Uncertain Business Cycles.” *Econometrica*, 86(3): 1031–1065.
- Cheng, Hong, Lukasz A. Drozd, Rahul Giri, Mathieu Taschereau-Dumouchel, and Junjie Xia.** 2021. “The Future of Labor: Automation and the Labor Share in the Second Machine Age.” Federal Reserve Bank of Philadelphia Working Paper WP 21-11.
- Costinot, Arnaud, and Iván Werning.** 2022. “Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation.” *The Review of Economic Studies*, 90(5): 2261–2291.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu.** 2021. “How Costly Are Markups?” Unpublished Manuscript, New York University and Duke University.
- Ekerdt, Lorenz K.F., and Kai-Jie Wu.** 2022. “The Rise of Specialized Firms.” *Working*

Paper.

- Furman, Jason, and Peter Orszag.** 2018. "Slower Productivity and Higher Inequality: Are They Related?" Peterson Institute for International Economics Working Paper Series WP18-4.
- Graetz, Georg, and Guy Michaels.** 2018. "Robots at Work." *The Review of Economics and Statistics*, 107(5): 753–768.
- Guerreiro, Joao, Sergio Rebelo, and Pedro Teles.** 2022. "Should Robots be Taxed?" *Review of Economic Studies*, 89: 279–311.
- Hsieh, Chang-Tai, and Esteban Rossi-Hansberg.** 2019. "The Industrial Revolution in Services." National Bureau of Economic Research NBER Working Papers 25968.
- Hubmer, Joachim, and Pascual Restrepo.** 2022. "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution." National Bureau of Economic Research Working Paper 28579.
- Jaimovich, Nir, and Henry E. Siu.** 2020. "Job Polarization and Jobless Recoveries." *The Review of Economics and Statistics*, 102(1): 129–147.
- Kehrig, Matthias, and Nicolas Vincent.** 2021. "The Micro-Level Anatomy of the Labor Share Decline." *The Quarterly Journal of Economics*, 136(2): 1031–1087.
- Khan, Aubhik, and Julia K Thomas.** 2008. "Idiosyncratic Shocks and the Role of Non-convexities in Plant and Aggregate Investment Dynamics." *Econometrica*, 76(2): 395–436.
- Kimball, Miles S.** 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking*, 27(4): 1241–1277.
- Klenow, Peter J., and Jonathan L. Willis.** 2016. "Real Rigidities and Nominal Price Changes." *Economica*, 83(331): 443–472.
- Kwon, Spencer Y., Yueran Ma, and Kaspar Zimmermann.** 2022. "100 Years of Rising Corporate Concentration." Unpublished Manuscript, University of Chicago Booth School of Business.
- Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard.** 2022. "Information Technology and Returns to Scale." Boston College.
- Leduc, Sylvain, and Zheng Liu.** 2019. "Robots or Workers? A Macro Analysis of Automation and Labor Markets." Federal Reserve Bank of San Francisco Working Paper No. 2019-17.
- Moreira, Marcelo J.** 2009. "Tests With Correct Size When Instruments Can Be Arbitrarily Weak." *Journal of Econometrics*, 152(2): 131–140.
- Olmstead-Rumsey, Jane.** 2019. "Market Concentration and the Productivity Slow-down." University Library of Munich, Germany MPRA Paper 93260.
- Prettner, Klaus, and Holger Strulik.** 2020. "Innovation, Automation, and Inequality:

- Policy Challenges in the Race Against the Machine." *Journal of Monetary Economics*, 116: 249–265.
- Ridder, Maarten De.** 2023. "Market Power and Innovation in the Intangible Economy." *American Economic Review*, , (forthcoming).
- Rogerson, Richard, and Johanna Wallenius.** 2009. "Micro and Macro Elasticities in a Life Cycle Model With Taxes." *Journal of Economic Theory*, 144(6): 2277–2292.
- Sui, Xiaomei.** 2022. "Uneven Firm Growth in a Globalized World." *Working Paper*.
- Tambe, Prasanna, Lorin Hitt, Daniel Rock, and Erik Brynjolfsson.** 2020. "Digital Capital and Superstar Firms." National Bureau of Economic Research Working Paper 28285.
- Thuemmel, Uwe.** 2022. "Optimal Taxation of Robots." *Journal of the European Economic Association*, 21(3): 1154–1190.
- Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz.** 2020. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." National Bureau of Economic Research Working Paper 28290.

Appendices

A Additional Tables

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. Regressions for Robot Density and Industry Concentration Based on Domestic Sales

	top 1% share of domestic sales			
	OLS		IV	
	(1)	(2)	(3)	(4)
ln(robot/thousand emp)	0.021** (0.007)		0.038** (0.020)	
ln(robot/million hours)		0.021** (0.007)		0.037* (0.020)
Observations	117	117	117	117
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Anderson-Rubin p -value			0.000	0.000

Note: The left-hand side variable in all columns is the domestic sales share of the top 1% of firms (in terms of domestic sales) in U.S. Compustat data. Domestic sales are defined as firm-level total sales minus exports, treating missing export data as zeros. The first two columns present OLS regression results analogous to specification (2), and the last two columns report the second-stage results of the instrumental variable (IV) regression. The industry-level robot density is measured as the operation stock of industrial robots per thousand workers or million labor hours within the industry. The instrumental variable for the U.S. robot density is the one-year lag of the robot density averaged over five European countries (EURO5), as described in the text. The last row shows the p -values of Anderson-Rubin weak instrument robust tests adjusted for non-homoskedasticity. All regressions weigh industries by their sales share in the initial year (2007), and all regressions control for industry and year fixed effects. Standard errors in parentheses are clustered at the industry level. Stars denote the statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3. Robot Density and Industry Concentration: Tests of Pre-trends

	five-year lagged top 1% share					
	sales		domestic sales		employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS regressions						
ln(robot/thousand emp)	0.006 (0.011)		0.007 (0.010)		0.002 (0.008)	
ln(robot/million hours)		0.005 (0.011)		0.006 (0.010)		0.002 (0.008)
Observations	122	122	122	122	102	102
Panel B: IV regressions						
ln(robot/thousand emp)	-0.024 (0.041)		-0.016 (0.034)		-0.010 (0.002)	
ln(robot/million hours)		-0.025 (0.045)		-0.017 (0.037)		-0.013 (0.021)
Observations	122	122	122	122	102	102
Covariates						
Industry FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Note: This table shows regression results of projecting the five-year lagged measures of industry concentration on robot density, testing for potential pre-trends. Panel A displays OLS regression results and Panel B reports the results of IV regressions. The dependent variables are the five-year lagged sales shares (1st column), domestic sales shares (2nd column), and employment shares (3rd column) of the top 1% of firms. Domestic sales are defined as firm-level total sales minus exports, treating missing export data as zero. The industry-level robot density is measured as the operation stock of industrial robots per thousand workers or million labor hours within the industry. 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. Note that the number of observations may exceed those in the baseline regressions presented in Tables 2 and 3, primarily due to a higher number of concentration observations from Compustat in earlier years. This arises from our sample restriction that each industry in Compustat must have at least 10 firms for concentration calculations to be performed. Since the number of firms is notably larger in earlier years, using the five-year lag of concentration measures results in more observations in our regression analysis. Stars denote the statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4. Parameters: The Whole Economy

Parameter	Notation	Value	Matched Moments
Relative price of robots	Q_a	139.84	Fraction of automating firms
SD of log automation fixed costs	σ_a	2.30	Employment share of automating firms

Note: This table shows the calibrated parameters by moment matching for the whole economy.

Table A.5. Matched Moments: The Whole Economy

Moments	Data	Model
Fraction of automating firms	2.0%	2.0%
Employment share of automating firms	15.7%	15.7%

Note: This table shows the targeted data moments and the simulated moments by the model for the whole economy. The data moments are based on the ABS data (taken from [Acemoglu et al., 2022](#)).

B Derivations

To simplify the intermediate producers' problem in equation (18), rewire the value function so that s is not a state variable:

$$\begin{aligned}
V(\phi, A; s) &= \max_{p, y, N, A'} \left[py - WN - Q_a[A' - (1 - \delta_a)A] - s\phi \mathbb{1}\{A' > 0\} + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right] \\
&= Q_a(1 - \delta_a)A + \max_{p, y, N, A'} \left[py - WN - Q_a A' - s\phi \mathbb{1}\{A' > 0\} + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right] \\
&= Q_a(1 - \delta_a)A + \max \left\{ \underbrace{\max_{p, y, N, A' > 0} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right]}_{\equiv V^a(\phi)} - s\phi, \right. \\
&\quad \left. \underbrace{\max_{p, y, N} \left[py - WN + \beta E_{\phi'|\phi} \int_{s'} V(\phi', 0; s') dF(s') \right]}_{\equiv V^n(\phi)} \right\} \\
&= Q_a(1 - \delta_a)A + \max\{V^a(\phi) - s\phi, V^n(\phi)\} \tag{43}
\end{aligned}$$

The firm with productivity ϕ chooses $A' > 0$ if and only if $s \leq s^*(\phi) \equiv \frac{V^a(\phi) - V^n(\phi)}{\phi}$.

We solve for the optimal decisions in $V^a(\phi)$ and $V^n(\phi)$ using the first-order conditions. Notice that the capital stock A is not a state variable since there is no friction on it. For automating firms, we have

$$V^a(\phi) = \max_{p, y, N, A' > 0} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right] \tag{44}$$

Conditional on paying the fixed cost of automation, $A' > 0$ would hold. Therefore, the value of an automating firm becomes:

$$\begin{aligned}
V^a(\phi) &= \max_{p, y, N, A'} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \int_{s'} V(\phi', A'; s') dF(s') \right] \\
&= \max_{p, y, N, A'} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \int_{s'} [Q_a(1 - \delta_a)A' + \max\{V^a(\phi') - s'\phi', V^n(\phi')\}] dF(s') \right] \\
&= \max_{p, y, N, A'} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \left[Q_a(1 - \delta_a)A' + \int_{s'} \max\{V^a(\phi') - s'\phi', V^n(\phi')\} dF(s') \right] \right] \\
&= \max_{p, y, N, A'} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \left[Q_a(1 - \delta_a)A' + \int_0^{s^*(\phi')} [V^a(\phi') - s'\phi'] dF(s') \right. \right. \\
&\quad \left. \left. + \int_{s^*(\phi')}^{\infty} V^n(\phi') dF(s') \right] \right]
\end{aligned}$$

$$\begin{aligned}
&= \max_{p,y,N,A'} \left[py - WN - Q_a A' + \beta E_{\phi'|\phi} \left[Q_a (1 - \delta_a) A' + F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' \phi' dF(s') \right. \right. \\
&\quad \left. \left. + [1 - F(s^*(\phi'))] V^n(\phi') \right] \right] \\
&= \max_{p,y,N,A'} \left[py - WN - Q_a A' + \beta Q_a (1 - \delta_a) A' \right] + \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' \phi' dF(s') \right. \\
&\quad \left. + [1 - F(s^*(\phi'))] V^n(\phi') \right] \\
&= \max_{p,y,N,A'} \left[py - WN - Q_a [1 - \beta(1 - \delta_a)] A' \right] + \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' \phi' dF(s') \right. \\
&\quad \left. + [1 - F(s^*(\phi'))] V^n(\phi') \right] \tag{45}
\end{aligned}$$

Let $\gamma_a \equiv Q_a [1 - \beta(1 - \delta_a)]$ denote the effective user cost of robots. Then the optimal choices of A' and N are those reported in equations (23) and (24) in the main text.

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

$$\begin{aligned}
V^n(\phi) &= \max_{p,y,N} \left[py - WN + \beta E_{\phi'|\phi} \int_{s'} V(\phi', 0; s') dF(s') \right] \\
&= \max_{p,y,N} \left[py - WN + \beta E_{\phi'|\phi} \int_{s'} [\max\{V^a(\phi') - s' \phi', V^n(\phi')\}] dF(s') \right] \\
&= \max_{p,y,N} \left[py - WN + \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' \phi' dF(s') + [1 - F(s^*(\phi'))] V^n(\phi') \right] \right] \\
&= \max_{p,y,N} \left[py - WN \right] + \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' \phi' dF(s') + [1 - F(s^*(\phi'))] V^n(\phi') \right] \tag{46}
\end{aligned}$$

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

$$\begin{aligned}
s^*(\phi)\phi &= V^a(\phi) - V^n(\phi) \\
&= \max_{p,y,N,A'} \left[py - WN - Q_a [1 - \beta(1 - \delta_a)] A' \right] \\
&\quad + \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' dF(s') + [1 - F(s^*(\phi'))] V^n(\phi') \right] \\
&\quad - \max_{p,y,N} \left[py - WN \right] - \beta E_{\phi'|\phi} \left[F(s^*(\phi')) V^a(\phi') - \int_0^{s^*(\phi')} s' dF(s') + [1 - F(s^*(\phi'))] V^n(\phi') \right] \\
&= \max_{p,y,N,A'} \left[py - WN - Q_a [1 - \beta(1 - \delta_a)] A' \right] - \max_{p,y,N} \left[py - WN \right], \tag{47}
\end{aligned}$$

and therefore

$$s^*(\phi) = \frac{\max_{p,y,N,A'} \left[py - WN - Q_a [1 - \beta(1 - \delta_a)] A' \right] - \max_{p,y,N} \left[py - WN \right]}{\phi}. \quad (48)$$

C Solution Algorithm

There are three loops to solve the problem. 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 $q(j)$ 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 $q(j)$, we have $y(j) = q(j)Y$.
 - (b) Given robot price Q_a and wage rate W , compute the marginal costs $\lambda(j)$ by eq. (25) and (27), and we can get $A'(j)$ and $N(j)$ 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. (6).
 - (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 $q(j)$ in the q loop as explained below.
3. Check whether the Kimball aggregator (8) 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 $q(j)$ according to eq. (9). Go back to Step 2.
- (c) If $LHS < RHS$, reduce W to raise $q(j)$ according to eq. (9). Go back to Step 2.

q loop: Find the relative production.

1. Given the factor 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. (9) to solve for the relative output $q(\phi)$ for each ϕ , for firms with and without robots.
 - (a) The right-hand side of (9) is a function of q by plugging in (13).
 - (b) The price in the left-hand side is the marginal cost in (25) or (27) times the markup in (15), which is also a function of q .
 - (c) Use the bisection method to solve for q in eq. (9).
4. Compute the automation decisions.
 - (a) Compute $y(j) = q(j)Y$ with and without robots.
 - (b) Compute the demand for $A'(j)$ and $N(j)$ 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 (10). Stop if D converges. Otherwise, go back to Step 2 and repeat until D converges.
 - (a) If $D = D^{\text{new}}$, D and $q(j)$ 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.