Automation and the Rise of Superstar Firms*

Hamid Firooz† Zheng Liu‡ Yajie Wang§

April 2023

Abstract

We document evidence that the rise in automation technology contributed to the rise of superstar firms in the past two decades. We explain the empirical link between automation and industry concentration in a general equilibrium framework with heterogeneous firms and variable markups. A firm can operate a labor-only technology or, by paying a per-period fixed cost, an automation technology that uses both workers and robots as inputs. Given the fixed cost, more productive, larger firms are more likely to automate. Increased automation boosts labor productivity, enabling large, robot-using firms to expand further, which raises industry concentration. Our calibrated model does well in matching the highly skewed automation usage toward a few superstar firms observed in the Census data. Since robots substitute for labor, increased automation raises sales concentration more than employment concentration, also consistent with empirical evidence. A modest subsidy for automating firms improves welfare since productivity gains outweigh increased markup distortions.

Keywords: Automation, industry concentration, markup, robots, reallocation, heterogeneous firms.

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 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 Rochester; email address: yajie.wang@rochester.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 steady increases in the usage of automation technology since the early 2000s have contributed significantly to the rise of superstar firms, particularly in the manufacturing sector.

The link between automation and industry concentration can be visualized from the time-series plots in Figure 1. The figure shows the average shares of sales and employment of the largest firms within manufacturing industries (Panel A).[1] The sales share of the top four firms (CR4) rose from about 40.5% in the late 1990s to about 43.5% in 2012, an increase of about 3 percentage points. During the same period, the sales share of the top 20 firms (CR20) also increased substantially. The employment shares of the top firms, in comparison, stayed relatively flat. The rise in industry 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%, while the number of industrial robots per thousand manufacturing employees has quadrupled.

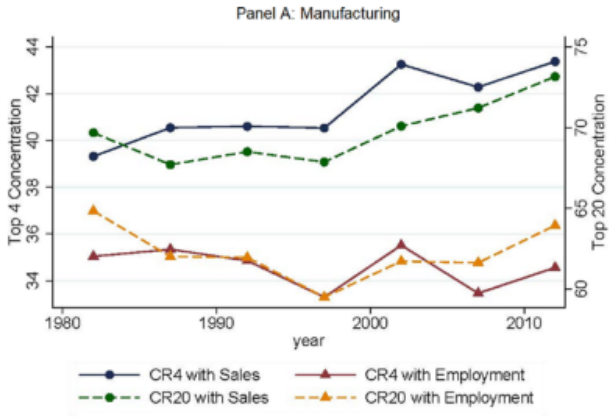
Sales concentration has also increased in Europe (Bajgar et al., 2019). For the manufacturing sector, sales concentration in Europe started rising a few years ahead of that in North America (see Figure 9 of Bajgar et al. (2019)). This timing of the increases in industry concentration aligns with the timing of automation adoptions: adoptions of automation technologies (in particular, industrial robots) started earlier in the European market than in the North American market (Acemoglu and Restrepo, 2020).

Our empirical findings suggest that the correlations between automation and concentration observed in the aggregate time series are also present in the industry-level data. We use Compustat firm-level data to construct industry concentration measures at the North American Industry Classification System (NAICS) two-digit industry level. We

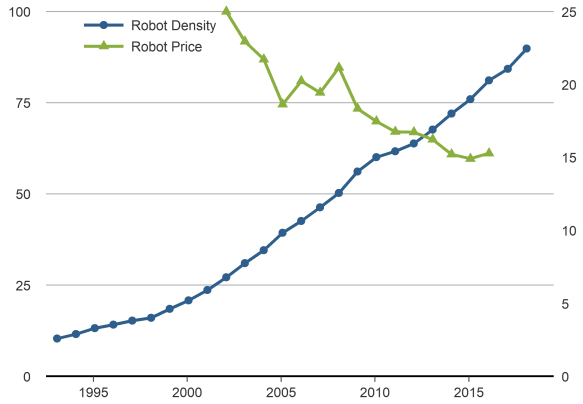
[1]The figure is taken from Figure IV in Autor et al. (2020) with permissions from the Oxford University Press (License Number 5241431011126).

Figure 1. Trends in Industry Concentration and Automation in Manufacturing

(a) Industry concentration



(b) Robot prices and robot density



Note: Panel (a) is taken from Autor et al. (2020) 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) shows the unit value of newly shipped industrial robots deflated by the personal consumption expenditures price index (red line, left scale) and robot density measured by the operation stock of robots per thousand manufacturing workers (blue line, right scale). Both series of robot price and the operation stock of industrial robots are taken from the International Federation of Robotics (IFR).

also construct an industry-level measure of robot density, which is defined as the ratio of the operation stock of industrial robots from the International Federation of Robotics (IFR) to thousands of manufacturing employment from the Bureau of Labor Statistics (BLS). We find that robot density is positively correlated with a sales-based measure of industry concentration (i.e., the sales share of the top 1% firms) and the correlation is economically important and statistically significant. In contrast, the correlation of robot density with employment-based concentration (measured by the employment share of the top 1% firms) is positive but statistically insignificant, with a magnitude much smaller than the correlation with sales concentration.

The observed correlations, however, do not necessarily reflect causal relations because both industry concentration and robot density are endogenous. To study the potential causal effects of automation on industry concentration, we estimate an instrumental variable (IV) panel specification. As documented by Acemoglu and Restrepo (2020), robot adoptions vary considerably across industries, and a common set of industries

in both the United States and Europe are more automatable and therefore experienced rapid robot adoptions in recent decades. More importantly, robot adoption trends in many European economies have been ahead of the United States. Thus, we use the lagged average robot density in five European economies as an instrumental variable for the U.S. robot density in our industry-level panel data regressions.[2] The IV regressions provide robust evidence that automation has significantly contributed to the rise of sales concentration in the United States, but its effect on employment concentration is small and insignificant. Our results are economically important: a one standard deviation increase in robot density raises the sales share of the top 1% firms by about 31 percent.

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. Firms also face idiosyncratic, persistent productivity shocks. A firm's automation decision (i.e., whether to operate the traditional 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. Increased 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 industry 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.

To assess the plausibility of the model's quantitative predictions, we calibrate the

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

model parameters to match several moments in the data and in the empirical macro literature. Specifically, we calibrate two key, non-standard parameters: the fixed cost of operating the automation technology and the cost of robot adoption. To this end, we target the observed share of firms that use robots for production and the employment share of those firms, both taken from the 2019 Annual Business Survey (ABS) conducted by the U.S. Census Bureau (Acemoglu et al., 2022). The ABS survey shows that approximately 2% of U.S. firms used robotics during 2016-2018, while 15.7% of U.S. workers were employed at these firms. These two moments indicate that robot adoptions are heavily skewed toward a few very large firms (see also Zolas et al., 2020).

As an external validation, we find that the calibrated model does well in predicting the cross-sectional distribution of automation usage observed in the firm-level data from the ABS. In particular, the usage of automation technologies is highly skewed toward large, high-productivity firms, both in the model and in the data. Given that we do not target the distribution of automation usage in our calibration, the ability of the model to correctly predict this cross-sectional distribution lends credence to the model's mechanism.

We use the calibrated model to examine the implications of an exogenous change in the relative price of robots for industry concentration. A decline in the robot price raises the probability of automation through two channels. First, it reduces the user cost of robots, benefiting large firms that operate the automation technology (an intensive-margin effect). Second, it induces more firms to adopt robots (an extensive-margin effect). Through the intensive-margin effect, the decline in the robot price enables large firms to become even larger, raising industry concentration. Through the extensive-margin effect, however, some smaller firms that initially operate the worker-only technology would switch to the automation technology when the robot price declines. This would reduce the sales share of the superstar firms and lower industry concentration.

Under our calibration, the intensive-margin effect dominates, such that a decline in the robot price leads to an increase in industry concentration. This is because the

calibrated model captures the fact that the usage of automation technology is highly skewed toward a small fraction of large firms, in line with the micro-level data (Zolas et al., 2020; Acemoglu et al., 2022). A modest decline in the robot price would not induce a sufficiently large share of smaller firms to switch to the automation technology, while it enables large firms that already use the automation technology to expand further, raising sales concentration.

A decline in the relative price of robots also increases the employment concentration, although the increase is smaller than that in sales concentration because the expansion of large firms relies more on robots than on workers. This model prediction is 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% decline in the relative price of robots—a magnitude similar to that observed during the past two decades—can explain about 80% of the rise in sales concentration in the U.S. manufacturing and about 61% of the diverging trends between sales concentration and employment concentration.

Although automation in our model is a labor-substituting technology, employment in automating firms increases following a decline in the robot price, in line with the ABS survey of Zolas et al. (2020). This is because increased robot usage raises labor productivity, which in turn boosts the labor demand of automating firms. Furthermore, since larger firms have higher markups and lower labor shares, the between-firm reallocation triggered by a decline in automation costs reduces the average labor share and increases the average markup, consistent with the reallocation channel documented by Autor et al. (2020), Acemoglu, Lelarge and Restrepo (2020), and Kehrig and Vincent (2021).

Our model implies that the relation between robot prices and industry concentration can be non-monotonic. We show that, in a counterfactual exercise with a sufficiently large decline in the relative price of robots, the extensive-margin effect would become dominant, such that a sufficiently large number of (medium-sized) firms would switch

technologies and expand production, reducing the sales share of the top 1% firms. Thus, in an economy with widely spread automation technologies, a decline in automation costs may not increase (and even reduce) industry concentration.[3]

The rapid rise of automation in recent years has raised an important policy question: Should robots be taxed?[4] We use our calibrated model to examine the implication of taxing (or subsidizing) robot-using firms for allocations and welfare. Because of monopolistic competition and variable markups in the product markets, equilibrium allocations in the model are inefficient. Taxing automating firms reallocates production from large, robot-using firms to smaller firms, reducing industry concentration and also reducing both the average markup and the markup dispersion. However, since larger firms are more productive, this reallocation lowers aggregate productivity. The optimal flat tax policy faces a tradeoff between alleviating markup distortions and reducing aggregate productivity. Under our calibration, a modest subsidy (about 4.7% of sales) for automating firms maximizes welfare, yielding a welfare gain equivalent to about 0.3% of steady-state consumption compared to the benchmark without policy interventions.

2 Related literature

Our work is motivated by the empirical evidence on industry concentration documented by Autor et al. (2020). Their study highlights an important between-firm reallocation channel that connects the rise in industry concentration with the fall in the labor share.

[3]Robots in our model are different from general capital equipment. Although both types of capital can substitute for workers, they differ in the sense that robot usage is highly concentrated in large firms, whereas equipment usage is much more widely spread. Our counterfactual simulation shows that, if the use rate of robots (i.e., the fraction of firms that operate automation technologies) is sufficiently high (e.g., bringing it to a level similar to the use rate of equipment), a decline in the relative price of robots would *reduce* industry concentration because the extensive margin would dominate the intensive margin. Based on this finding, we conjecture that a decline in the relative price of capital equipment, which is more widely used than robots, could reduce industry concentration.

[4]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 non-routine occupations, optimal policy calls for taxing robots. See also Prettner and Strulik (2020) for an analysis of how robot taxes can help redistribute income from high-skilled workers to low-skilled workers.

They discuss a few potential drivers of the rise of superstar firms (what they call a "winner takes most" mechanism), such as greater market competition (e.g., through globalization) or scale-biased technological change driven by intangible capital investment and information technology. Other potential drivers of the rise of concentration studied in the literature include 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 evidence suggests that the increased use of automation technology has also contributed to the observed increases in industry concentration.

Our evidence further indicates that although automation has contributed to increases in sales concentration since the late 1990s, it has not raised employment concentration in the U.S. manufacturing sector. This finding is in line with Hsieh and Rossi-Hansberg (2019), who document evidence 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).

Our theoretical model suggests that an important driver of the empirical link between automation and industry concentration is an economies-of-scale effect associated with high fixed costs of automation and low marginal costs in the production process enabled by the automation technology. This finding aligns with several other studies that highlight the importance of economies of scale for driving the increase in industry concentration (Kwon, Ma and Zimmermann, 2022; Aghion et al., 2019; Lashkari, Bauer and Boussard, 2022; Tambe et al., 2020; Sui, 2022).

Our work is closely related to Hubmer and Restrepo (2022), who present a model featuring heterogeneous firms with fixed costs of automating tasks. Their focus is on the role of automation in driving the observed decline in the labor share. A decline in capital prices reduces the aggregate labor share because large firms automate more tasks, while

the median firm continues to operate a labor-intensive technology with a rising labor share. Although our model's implications for the relation between automation and the labor share align with those found by them, our study, in contrast, focuses on how automation drives industry concentration. Moreover, a key contribution of our work is that our quantitative model successfully generates the highly skewed distribution of automation usage toward a few superstar firms observed in the U.S. firm-level data. To our knowledge, no other models in the literature have been able to replicate the observed distribution of automation usage across firms.

More broadly, our work contributes to the burgeoning literature on automation and labor markets. Automation has 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). Empirical evidence suggests that, at the firm level, robot adoptions are associated with declines in the labor share (Autor and Salomons, 2018; Acemoglu, Lelarge and Restrepo, 2020). Our paper complements this literature by showing that automation also has important implications for the rise of superstar firms, and increases sales concentration more than employment concentration.

3 Industry-level Evidence

This section examines the empirical relation between automation and industry concentration for U.S. manufacturing industries. We first present 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 present causal evidence demonstrating that automation

has significant impacts on sales concentration, but it has no such effects on employment concentration.

3.1 Data and measurement

We use firm-level data from Compustat to compute two measures of industry concentration: the sales share as well as the employment share of the top 1% of firms in a given industry.[5] The top 1% of firms is comparable to the top 4 firms analyzed by Autor et al. (2020), since an average four-digit manufacturing industry has around 364 firms and therefore the top 4 firms are approximately equivalent to the top 1% of firms.

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).[6] We define robot density for industry j in year t as

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

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 of industry-level employment (EMP) and labor hours (PRODH) are both obtained from the NBER-CES Manufacturing Industry Dataset.[7] We obtain an unbalanced panel with 13 industries covering the 12 years from 2007 to 2018.[8]

Table 1 reports the summary statistics of variables. First, it shows that robot density

[5]Using a percentile is more appropriate than using a specific number of firms as the cutoff, given that the total number of public firms in Compustat changes greatly across time.

[6]According to the definition of IFR, industrial robots are automatically controlled, reprogrammable, and multipurpose manipulators with several axes.

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

[8]We 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. The sample sizes for some variables are smaller than $12 \times 13 = 156$ because of missing values in certain industry-year cells.

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 that 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 as well as 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.

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%, with an IQR of about 14% and a standard deviation of 13%. The employment share of the top 1% of firms averages about 27% and varies less than the sales share, with an IQR of about 11% and a standard deviation of about 8%.

3.2 Correlations between automation and industry concentration

We calculate the correlations of automation with industry concentration, controlling for industry and year fixed effects. Specifically, we estimate the following ordinary least squares (OLS) specification

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

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**		0.002	
	(0.007)		(0.015)	
ln(robot/million hours)		0.021**		0.002
		(0.007)		(0.015)
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. The industry-level robot density is measured as the operation stock of industrial robots per thousand workers or million labor hours within the industry. All regressions weigh the industries 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

where the dependent variable Y_{jt} is a measure of industry concentration in industry j and 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%).[9] The estimated correlation between the hours-based measure of robot density and sales concentration is similar in magnitude and statistical significance (Column (2)).

The correlation of robot density with employment concentration (i.e., the employment share of the top 1% of firms), although positive, is much smaller than that with sales concentration, and the estimated correlations are statistically insignificant (Columns (3) and (4)). 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 estimate an instrumental-variable (IV) panel specification. We follow Acemoglu and Restrepo (2020) and use the lag of robot adoptions in European countries as our instrument for robot adoption in the U.S. As documented by Acemoglu and Restrepo (2020), robot adoptions vary considerably across industries, and a common set of industries in both the U.S. and Europe are rapidly adopting robots. Furthermore, partly due to Europe's more rapidly aging population, the robot adoption trends in Europe are ahead of those in the U.S. (Acemoglu and Restrepo, 2022). Therefore, the rise in robot adoption in European countries can indicate the development of automation technologies, which would then be correlated with U.S. robot adoption.

[9]The 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 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.

Moreover, using an instrument constructed from foreign countries' robot adoptions can help isolate from the U.S.-specific forces that drive both robot adoptions and industry concentration. As for other shocks that may be common to the U.S. and European countries, Acemoglu and Restrepo (2020) show that U.S. robot adoption is not strongly correlated with several major trends, including import competition, offshoring, and capital deepening.

We use the lag of the average robot density of five European economies (EURO5) for each industry as an instrumental variable for U.S. robot density in our regressions. 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 U.S., we measure robot density in the EURO5 economies by the number of robots per thousand employees (or per million of labor hours) in each industry, with the employment (and hours) data taken from EUKLEMS. The average robot density of the five European economies in industry j at time t is calculated as

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

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

Our two-stage least squares (2SLS) regressions are just identified, with one endogenous regressor and one instrumental variable. Specifically, 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.

[10] Following 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 US adoption trends

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**		0.012	
	(0.019)		(0.016)	
ln(robot/million hours)		0.036*		0.014
		(0.020)		(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 instrumental variable (IV) regression from the empirical specification (2). The 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). The last row shows the p-values of Anderson-Rubin weak instrument robust tests adjusted for non-homoskedasticity. All regressions weigh the industries 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 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)).

To address concerns about the robustness of our estimation and inferences in the

than those trends in the EURO5 economies.

[11]The logged robot density in the US 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 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.

presence of potentially weak instruments, we perform an Anderson-Rubin (AR) test (Anderson and Rubin, 1949), which 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). In the last row of Table 3, we report the AR test's p-values adjusted for non-homoskedasticity. The p-values indicate that the estimated effects of robot density on sales concentration are robust to weak instruments at around the 99% confidence level, although the effects on employment concentration are not robust, with a p-value of the AR test larger than 0.40.

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 large number of infinitely lived identical households with a unit measure. All agents have perfect foresight. The representative household has the utility function

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

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

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

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

$$W_t = \chi N_t^{\xi} C_t. \tag{6}$$

The optimizing decision for equity share holdings is given by

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

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 good producers

There is a large number of monopolistically competitive intermediate producers with a unit measure indexed by $j \in [0, 1]$. Final good 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

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

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

$$p(j) = \Lambda'(q(j))D, \tag{9}$$

where D is a demand shifter given by

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

We follow Klenow and Willis (2016) and assume that

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

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

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

Under the specification (11) for Λ, we obtain

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

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

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

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

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

As a result, larger firms face lower demand elasticities, have more market power, and charge higher markups.[12]

4.4 Intermediate goods producers

Intermediate producers, from now on indexed by their productivities ϕ, produce differentiated intermediate goods using two alternative technologies: one with labor as the sole 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

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

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.

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

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

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

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

production. We further assume that the distributions of s and ϕ are independent.[13] 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

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

where $\mathbb{1}\{x\}$ equals one if x holds and zero otherwise. The firm hires workers in a competitive labor market with the 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

$$A' = (1 - \delta_a)A + I_a. \tag{19}$$

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

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

where the continuation value of operating the automation technology this period (i.e.,

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

having $A' > 0$) is given by

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

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

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

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

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

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

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:

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

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

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

The marginal cost of production in this case would be

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

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:

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

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:

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

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.

As Appendix A shows, 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,

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

where

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

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

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

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 optimizing problem; (ii) taking W and Y as given, the firm-level allocations and prices solve each individual firm's optimizing problem; and (iii) the markets for the final good and labor clear.

The final goods market clearing condition is given by

$$C + Q_a I_a + \int_{\phi} \int_{0}^{s^*(\phi)} s\phi \, dF(s) \, dG(\phi) = Y. \tag{33}$$

The labor market clearing condition is given by

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

The stock of robots is given by

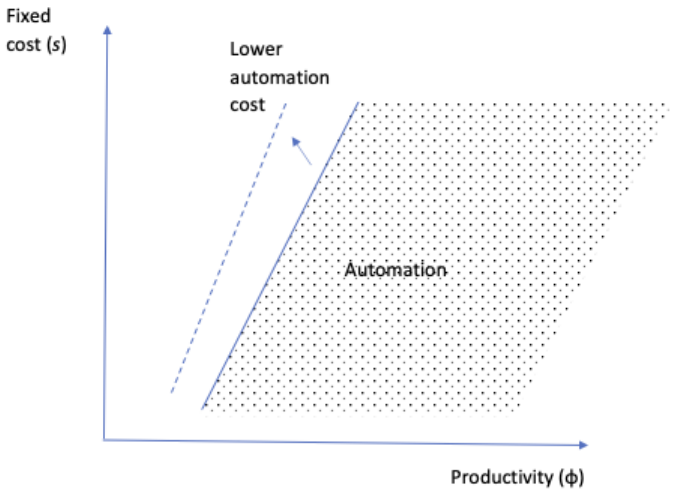
$$A = A' = \int_{\phi} A'(\phi) F(s^*(\phi)) dG(\phi). \tag{35}$$

Total investment in robots equals

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

Appendix B outlines the computational algorithm to solve the model.

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.

5 Model mechanism

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 trade-off 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.[14]

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

[14] As 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.

any given fixed cost s, a firm will automate if the realized productivity is sufficiently high. There is an upward-sloping line that separates the technology choices. To the right of the line (high ϕ or low s), firms use the automation technology and to the left of the line, they use the labor-only technology. Firms with combinations of ϕ and s on the upward-sloping line are indifferent between the two types of technologies.

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 share of sales 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 model's quantitative implications, we parameterize the model and solve the steady-state equilibrium. Table 4 displays the calibrated parameters. We externally calibrate a subset of parameters based on the literature (Panel A). 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 set the elasticity of substitution between robots and workers in the automation technology to $\eta = 3$, and the input weight of robots to $\alpha_a = 0.465$ following the study of Eden and Gaggl (2018).[15] We calibrate the quarterly robot depreciation rate to $\delta_a = 0.02$, implying an average robot lifespan of about 12 years, in line with the assumption made by the IFR in imputing the 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).[16] To calibrate the elasticity parameters σ and ϵ in the Kimball preferences, we follow Edmond, Midrigan and Xu (2021) and set $\sigma = 10.86$ and $\epsilon/\sigma = 0.16$.

[15] 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, the elasticity of $\eta = 3$ is conservative relative to their benchmark estimate.

[16] Bloom 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. Moreover, their estimated transition probabilities suggest that the unconditional probability of the low standard deviation is 68.7%. Therefore, the average standard deviation is 0.1 (=0.051*68.7%+0.209*(1-68.7%)).

Table 4. Parameters

Parameter	Notation	Value	Sources/Matched Moments
Panel A: Assigned Parameters			
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
Elasticity of substitution	η	3	Eden and Gaggl (2018)
Robot input weight	α_a	0.465	Eden and Gaggl (2018)
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 from Moment Matching			
Relative price of robots	Q_a	45.9	Fraction of automating firms
SD of log of automation fixed costs	σ_a	2.2	Employment share of automating firms

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.

We calibrate the remaining parameters to match some key moments in the micro-level data. These parameters include the relative price of robots Q_a and the parameters in the distribution of the fixed cost of automation. We assume that the fixed cost of automation follows a log-normal distribution $\ln(s) \sim \mathcal{N}(0, \sigma_a^2)$, and we calibrate the standard deviation parameter σ_a. The calibrated values are shown in Panel B of Table 4.

Our calibration strategy is as follows. 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

$$\int_\phi F(s^*(\phi))\, dG(\phi).$$

We calibrate Q_a to target the observed fraction of firms that use robots in the micro-level data. In particular, we target this moment to match that in the 2019 ABS conducted by the U.S. Census Bureau, which shows that the fraction of firms that use robots is about

Table 5. Matched Moments

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. The data moments are based on the ABS data (taken from Acemoglu et al., 2022).

2% during 2016-2018 (Acemoglu et al., 2022).

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

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

In the ABS survey, the employment share of automating firms is about 15.7% during 2016-2018 (Acemoglu et al., 2022).

By matching the fraction of automating firms and the employment share of those firms in the ABS data, we obtain $Q_a = 45.9$ and $\sigma_a = 2.2$, as shown in Panel B of Table 4. The calibrated model matches the targeted moments closely, as shown in Table 5.

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 Automation distribution

The calibrated model does well in replicating the observed distribution of firm-level automation usage in the ABS data, a moment that we do not target in the calibration.

Figure 3 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.[17] 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.[18]

The ability of the model to correctly predict this non-targeted distribution therefore lends credence to the model's mechanism. By matching 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 show below.

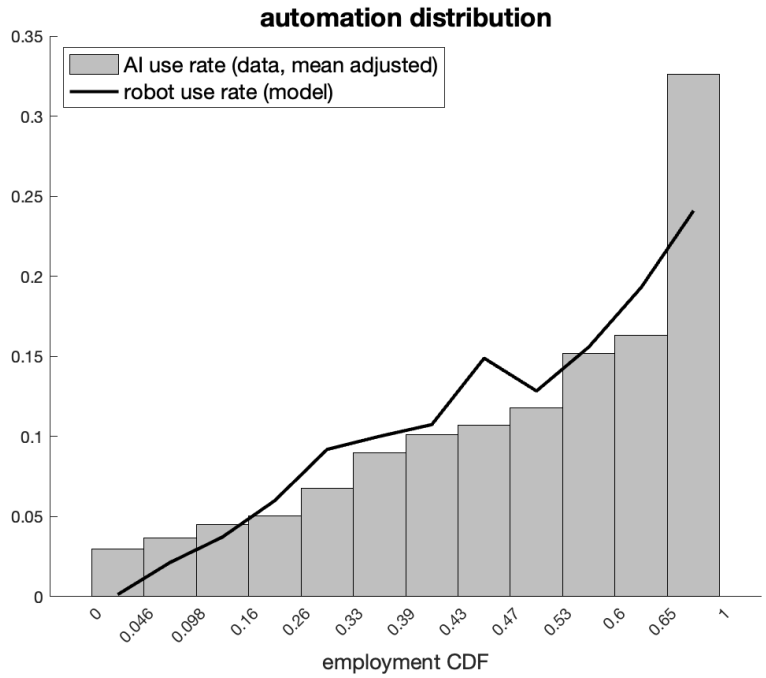
7.2 Firm-level implications

To further examine the automation mechanism, we plot in Figure 4 the firms' decision rules as a function of the idiosyncratic productivity level ϕ. In each panel, we show two

[17] Zolas 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 use rate across the employment CDF. Consistently, we calculate the robot use rates across the employment CDF in the model using the same method. Note that AI is more commonly used than robots, 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 3 scales the AI use rates in the data to have the same mean as that of the robot use rates in our model.

[18] Acemoglu et al. (2022) document that the distribution of robot use rates across firms is also skewed toward large firms, although they do not report more granular robot use rates for firms within the top percentile.

Figure 3. Automation Distribution

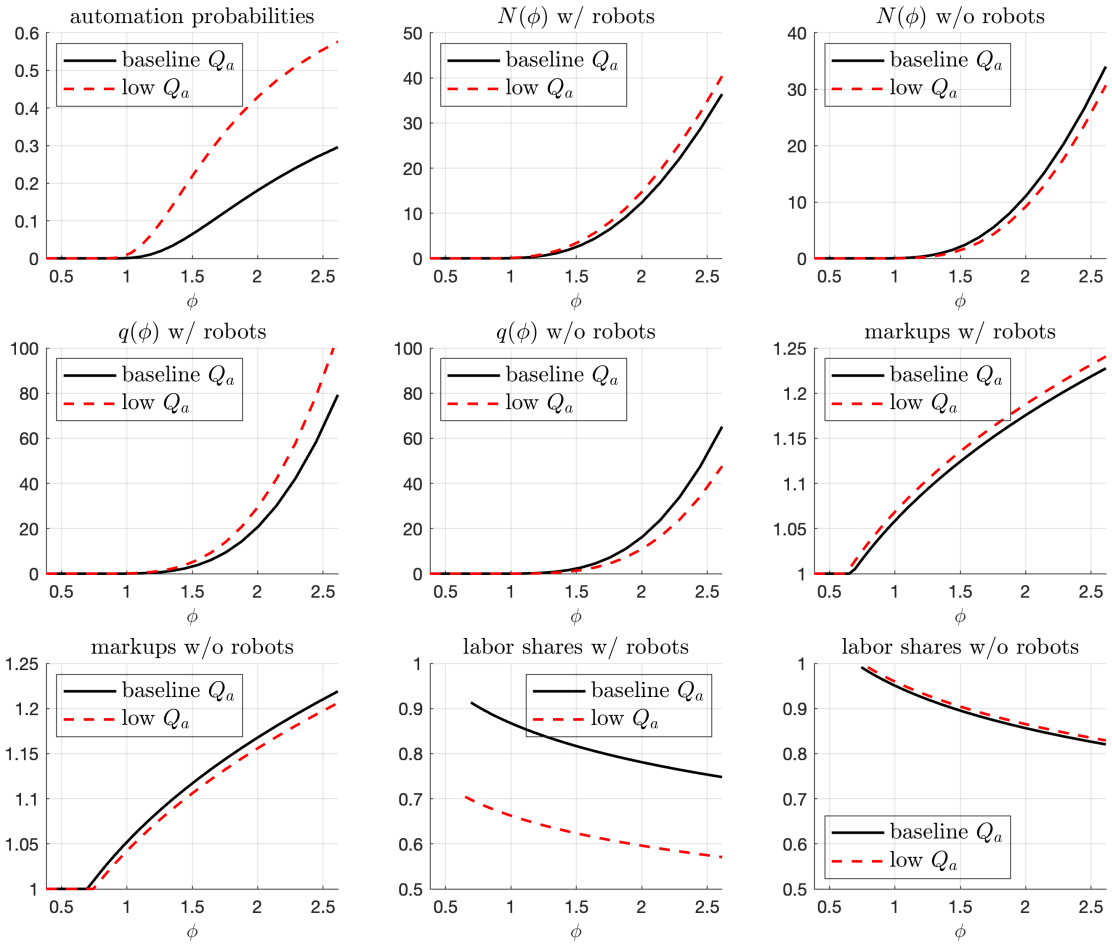


Note: This figure plots the distribution of AI use rate (i.e., the fraction of firms that use AI in their production) in the ABS data (bars) and the fraction of firms using robots in the model (line), for firms of different sizes measured by the cumulative density of employment. The AI use rates are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

lines, one in the baseline model with calibrated parameters (black solid line) and the other in a counterfactual scenario with a lower robot price (low Q_a, red dashed line). The figure shows that the automation probability increases with productivity, since more productive firms are more likely to be able to cover the fixed costs to access the automation technology. In addition, a non-degenerate set of firms with sufficiently low productivity do not use robots and operate the worker-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 ($N(\phi)$), higher relative output ($q(\phi)$), have larger market power

Figure 4. Firms' Decision Rules



Note: This figure shows firms' decision rules for the firms that automate (w/ robots) and those that do not automate (w/o robots). The solid-black lines are associated with our baseline calibration, while red-dashed lines show the results for a counterfactual in which robot price Q_a falls by 50%.

measured by markups, and have 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 4 further shows 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 and to gain market share. Under our calibration, the latter effect dominates such that automating firms increase employment following the reduction in the robot price. The labor shares of the automating firms decline despite the increases in employment, reflecting the substitution of robots 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.

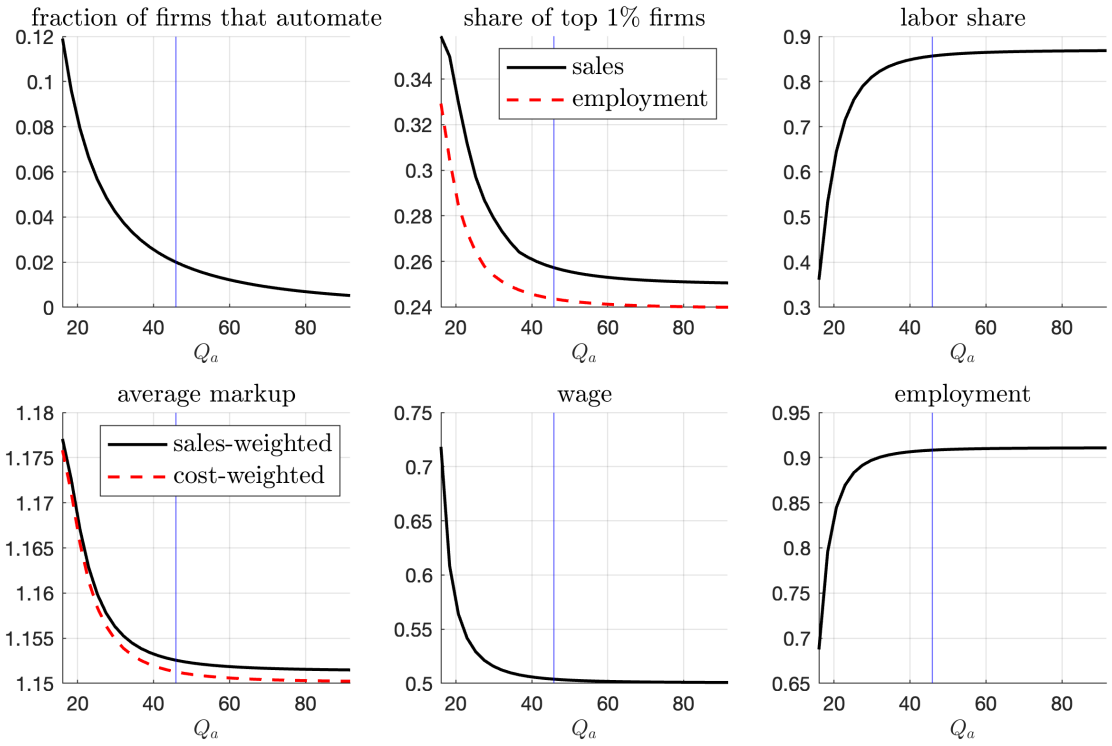
7.3 Aggregate implications

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 5. To illustrate, we consider a wide range of the robot price that covers the calibrated value of $Q_a = 45.9$, 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.[19] As a result, the product market becomes more concentrated

[19] As discussed before, while automation fixed costs are proportional to firm-level productivity, more

Figure 5. 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.

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).[20] As Figure 4 shows, a reduction in Q_a reallocates

productive firms are still more likely to automate, as shown in the top-right figure.

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

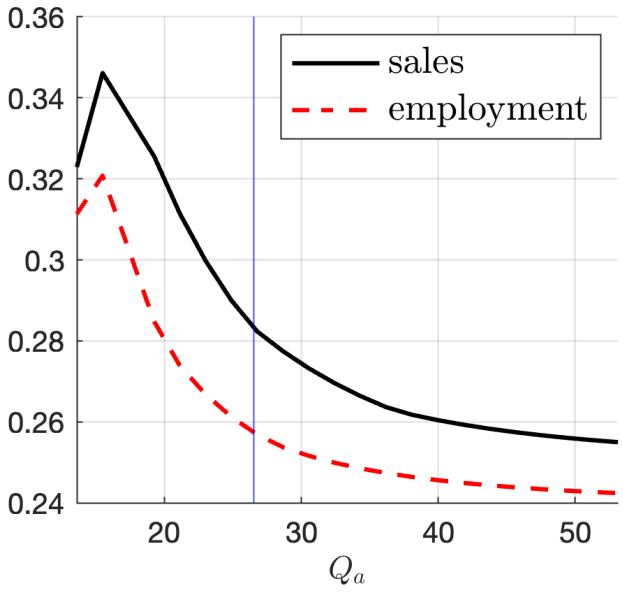
production and employment toward automating firms that have lower labor shares in the original steady state. Therefore, as Q_a falls, the labor share in the aggregate economy declines. Our model thus implies that declines in the aggregate labor share and increases in the average markup are mainly driven by the between-firm reallocation channel, in line with the empirical evidence in Autor et al. (2020) and Acemoglu, Lelarge and Restrepo (2020).

A reduction in Q_a also reduces 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, a reduction in Q_a leads to an increase in wages and a decline in aggregate employment.[21]

Concentration in manufacturing. We now shift our focus to the manufacturing sector to examine the quantitative importance of the automation mechanism in our model, specifically in explaining the observed rise in sales concentration as well as the divergence between sales and employment concentration in this sector. 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%

[21]Our model's prediction that a reduction in the robot price raises worker wages seems to be at odds with the empirical evidence documented by Acemoglu and Restrepo (2021), who find substantial declines in the relative wages of workers specialized in routine tasks in industries experiencing rapid automation. This is perhaps not surprising because we focus on studying the relation between automation and industry concentration and abstract from labor market frictions in our model. In a model with elaborated labor market frictions, such as the business cycle model with labor search frictions and automation studied by Leduc and Liu (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.

Figure 6. Robot Prices and Industry Concentration in Manufacturing



Note: This figure shows the effects of counterfactual changes in the robot price Q_a on the share of the top 1% of firms when the model is calibrated to match the robot usage in the U.S. manufacturing sector. The vertical blue line indicates the calibrated value of robot price Q_a.

of manufacturing firms use robots and those firms employ about 45.3% of manufacturing workers. In comparison, in the whole economy, about 2% of firms use robots and they employ about 15.7% of workers (Acemoglu et al., 2022). Second, one of the largest increases in sales concentration accompanied by one of the largest rises in the gap between sales and employment concentration in the past two decades has happened in the manufacturing sector (see Figure 1 and Autor et al. (2020)).

To this end, we re-calibrate the robot price Q_a and the standard deviation of the fixed cost of automation σ_a to match the fraction of firms using robots in the manufacturing sector (8.7%) and the employment share of these firms (45.3%). We keep the other parameter values the same as in our baseline calibration. Based on this calibration, we perform counterfactual exercises in our model by varying the robot price Q_a. In particular, we reduce the robot price from 44.3 to its calibrated value of 26.6, representing a 40% decrease 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.[22] 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.

Figure 6 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) or employment (dashed line). The vertical blue line shows the calibrated value of $Q_a = 26.6$. When Q_a declines from 44.3 to the calibrated value of 26.6 (a 40% decline, as in the data), the sales share of the top 1% of firms rises by about 2.4 percentage points (from 25.8% to 28.2%). The employment share of the top 1% of firms also rises but with a smaller magnitude (1.3 percentage points). Thus, the gap between sales concentration and employment concentration widens by about 1.1 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.[23] Thus, the model can explain roughly 80% (2.4 out of 3 percentage points) of the increases in sales concentration in the data. The model also explains about 61% (1.1 out of the 1.8 percentage points) of the observed diverging trends between sales concentration and employment concentration.

Figure 6 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 steady state, a reduction in the robot price would increase industry concentration, as we find in the model here. This is consistent with the positive effects of automation on sales concentration in the U.S. that we documented in Section 3. However, in an economy with widespread automation (i.e., an economy with a sufficiently low

[22]The 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 between 1998 to 2012 is the same as that between 2002 to 2016 (i.e., 40%).

[23]Notice 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.

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 in Figure 6, our model implies that a decline in the price of general equipment that is widely used in the economy could decrease, rather than increase, industry concentration.

7.4 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 create jobs for workers, it might also reduce labor productivity and put downward pressure on 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 whether such policies are welfare-improving.

We introduce a flat sales tax τ on firms that use the automation technology. The intermediate producers' problem in equation (18) therefore becomes:

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

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

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

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

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

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

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 by the relation

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

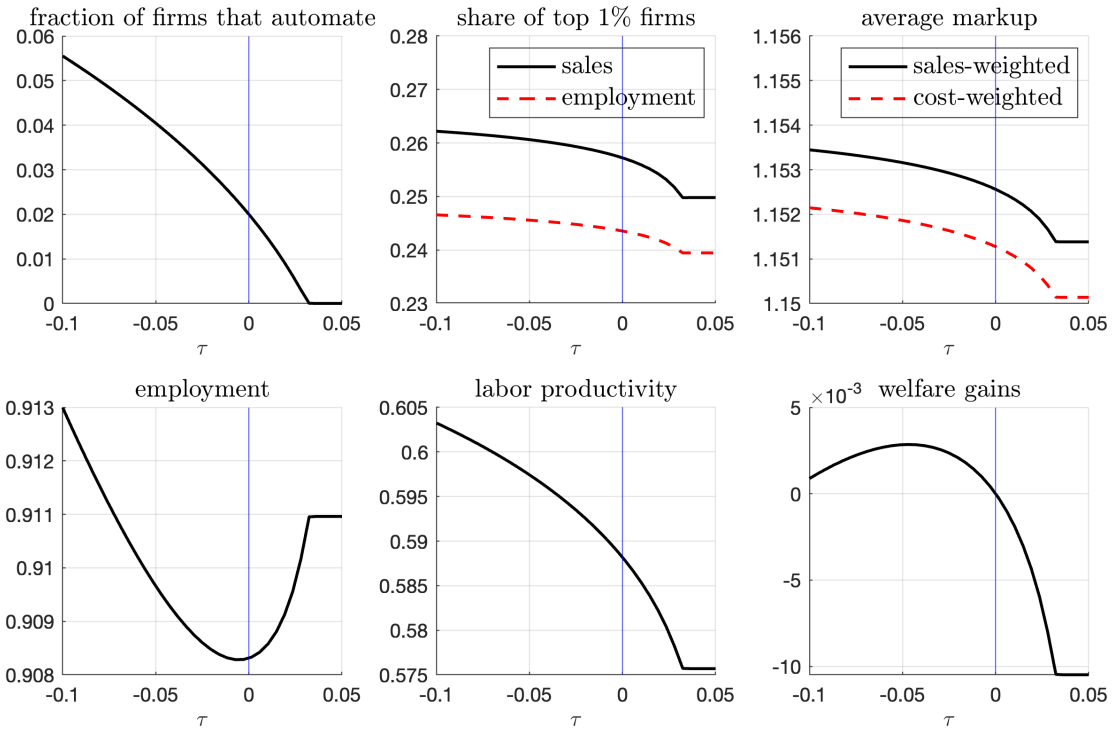
Solving for μ from Eq. (41), we obtain

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

A positive μ implies that the economy with automation taxes has lower welfare than the one without, that is, automation taxes would lead to welfare losses. On the other hand, a negative μ would imply welfare gains from taxing automation.

Figure 7 shows the aggregate effects of imposing a sales tax on automating firms, where we use our calibrated model for the whole economy. First of all, notice that since only 2% of firms use automation technology in the calibrated economy, even a small automation tax of around 3% would drive the mass of automating firms down to zero,

Figure 7. Effects of Taxing Automation



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

after which increasing the automation tax would have no effects. As most of the actions are for the case of $\tau < 0$, we will focus on the effects of an automation subsidy in what follows.

An automation subsidy (i.e., $\tau < 0$) reduces the marginal cost of using robots and therefore increases the fraction of automating firms. Automation subsidy has two competing effects on the market share of superstar (i.e., the top 1%) firms. On the intensive margin, since larger firms are more likely to be able to cover the fixed cost of automation, an automation subsidy will favor large, automating firms and make them even larger, leading to an increase in the share of superstar firms. On the extensive margin, however, an automation subsidy incentivizes some non-automating firms to pay the fixed cost and automate production, which would make these firms larger and in turn reduces the market share of superstar firms. Moreover, since an automation subsidy will favor large,

automating firms that charge higher markups, it also increases the average markup in the economy.

More robot usage in the economy that is induced by an automation subsidy increases labor productivity, while reducing the aggregate labor share because a larger fraction of output in the economy will be produced using the automation technology.[24] Starting from a high level of automation taxes, reducing τ would initially reduce employment, reflecting the labor-substituting effects of automation. With a sufficiently high automation subsidy, however, increasing subsidies further would boost employment, because automation raises labor productivity, shifting up labor demand and raising aggregate employment, despite its labor-substituting effects.

Automation subsidies have non-linear effects on welfare, as shown in Figure 7. An automation subsidy can raise welfare by improving labor productivity through increased automation. On the other hand, an automation subsidy can reduce welfare by raising industry concentration and the average markup. Figure 7 shows that, under our calibration, there is an interior optimum rate of automation subsidy at about 4.7%, which maximizes welfare, with a maximum welfare gain of about 0.3 percent of steady-state consumption equivalent relative to the benchmark without automation taxes or subsidies.

8 Conclusion

We have presented empirical evidence suggesting that automation has contributed to the rise in industrial concentration since the early 2000s. We explain the link between automation and industry concentration by an economy-of-scale effect stemming from fixed costs of operating the automation technology in a general equilibrium model. Our calibrated model predicts a highly skewed distribution of automation usage toward a small number of superstar firms, and this prediction aligns well with the firm-level

[24] Notice that the rise in average markups also contributes to the fall in the labor share.

data. Our model predicts that a decline in the robot price of a magnitude similar to that observed during the past two decades can account for about 80% of the rise in sales concentration in U.S. manufacturing and about 61% of the diverging trends between sales concentration and employment concentration. Thus, the rise of automation is quantitatively important for driving the rise of superstar firms.

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, 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 mathematical 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.

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.

Eden, Maya, and Paul Gaggl. 2018. "On the Welfare Implications of Automation." *Review of Economic Dynamics*, 29: 15–43.

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 Slowdown." 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.

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.

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 Derivations

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

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

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

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

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

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

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

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:

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

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

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

and therefore

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

B 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.