

The Future of Labor: Automation and the Labor Share in the Second Machine Age^{*}

Hong Cheng

Wuhan University

Lukasz A. Drozd

Federal Reserve Bank of Philadelphia

Rahul Giri

International Monetary Fund

Mathieu Taschereau-Dumouchel

Cornell University

Junjie Xia

Peking University

January 19, 2021

Preliminary — comments welcome

Abstract

We study the effect of modern automation on firm-level labor shares using a 2018 survey of 1618 manufacturing firms in China. We exploit geographic and industry variation built into the design of subsidies for automation paid under a vast government industrialization program, “Made In China 2025”, to construct an instrument for automation investment. We use a canonical CES framework of automation and develop a novel methodology to structurally estimate the elasticity of substitution between labor and automation capital among automating firms, which for our preferred specification is 3.8. We calibrate the model and show that the general equilibrium implications of this elasticity are consistent with the aggregate trends during our sample period.

Key words: labor share, labor’s share in income, automation, labor demand, industrial robots.

JEL Classifications: D33, E25, O33, J23, J24, E24, O25.

^{*}We thank Jonas Arias, Mark Bilal, Julieta Caunedo, Jonathan Eaton, Loukas Karabarbounis, Samuel Kortum, Alex Monje-Naranjo, and Marina Mendes Tavares for their insightful comments. We are particularly grateful to Jonas Arias for his technical insights. All remaining errors are ours. The views expressed in this paper are solely those of the authors and do not necessarily reflect those of the International Monetary Fund, the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Cheng: 919637855@qq.com; Drozd (corresponding author): lukasz.drozd@phil.frb.org, Federal Reserve Bank of Philadelphia Ten Independence Mall Philadelphia, PA 19106-1574; Giri: rgiri@imf.org; Taschereau-Dumouchel: mt763@cornell.edu; Xia: junjiexia@nsd.pku.edu.cn.

The share of labor in national income has been shrinking globally since the early 1980s, undermining economists’ confidence that a constant labor share is indeed an immutable fact of growth (Kaldor, 1961). Karabarbounis and Neiman (2013) measure that decline in multiple countries and find that the labor share has fallen by about five percentage points since then—a large movement by historical standards.¹ One hypothesis is that this decline is driven by a modern wave of automation, fueled by recent advances in artificial intelligence, machine learning and dexterous robotics, which have eroded humans’ advantage in cognitive tasks. As a result, “robots” are now eating into labor’s share of national income, as opposed to merely affecting the distribution of that income across laborers, as historically was the case. A growing number of studies seem to confirm that technology is indeed rapidly changing how capital can be substituted for labor.² But while these developments might have captured the imagination of intellectuals, popular writers and politicians, concrete evidence for how modern automation is affecting the labor share is still scant in the economic literature.³

Measuring the effect of modern automation on the labor share is a challenging task, and correlation can be easily confused with causation. The issue is that investment in automation is oftentimes driven by unrelated trends or shocks that difficult to measure and affect the labor share. For example, demand shocks that increase markups—and thereby lower the labor share of a firm or an entire industry—may push firms to expand production and invest in automation to satisfy the simultaneously growing demand for their products. An analogous effect can be driven by other shocks when a firm’s demand features a nonconstant elasticity. Similarly, offshoring of labor intensive activities to lower income countries—or any change in consumption patterns over time—may push existing firms to specialize in goods that are more capital intensive, potentially creating a non-causal link between the labor share and automation.⁴ To address such or similar concerns, a careful identification strategy is required to determine whether the link between automation and labor share is indeed a technological one.

¹See Dao et al. (2017) for updated evidence. See also Autor et al. (2003), Goos and Manning (2007), Acemoglu and Autor (2011), David and Dorn (2013), Michaels et al. (2014) and Dvorkin and Monge-Naranjo (2019) for evidence on how automation is impacting jobs and job polarization.

²For example, according to the McKinsey Institute (2019), a quarter of labor hours is estimated to be lost to automation by 2030, with highly skilled occupations expected to be at stake as much as low skilled occupations. The Brookings Institute (2019) similarly estimates that “approximately 25% of U.S. employment will face high exposure to automation in the coming decades—with greater than 70% of current task content at risk of substitution.” See also Frey and Osborne (2017). For a contrarian view, see the work by Arntz et al. (2016). While their estimates are considerably smaller, they are still sizable.

³See, for example, Brynjolfsson and McAfee (2014), Ford (2009) or Frey (2020).

⁴For example, Hubmer (2020) shows that changing consumption pattern may partly account for labor share declines in the US.

In this paper, we attack this problem by developing a novel approach that uses policy-induced local variation in the price of automation capital as an instrument for automation investment at the firm level. To implement our approach, we use new data from a survey of 1618 Chinese manufacturing firms conducted by the Wuhan University Institute of Quality Development Strategy and designed in collaboration with a number of researchers from the Hong Kong University of Science and Technology, Stanford University and the Chinese Academy of Social Science.⁵ This data provides detailed information on firms’ operations between 2015 and 2017, including the type of equipment they purchased and, most importantly, the subsidies that they received from the government for the purchases of automation capital (industrial robots and numerically controlled machines). These subsidies were paid under an unprecedented government industrialization program, “Made in China” (MIC, hereafter), but they were implemented by local municipalities. As a result, subsidy rates varied considerably across cities and industries. We exploit this variation in the effective price of automation capital to identify the causal impact of automation on the labor share.

To develop our empirical strategy, we extend a canonical model of automation—along the lines of [Graetz and Michaels \(2018\)](#) and [Acemoglu and Restrepo \(2018\)](#)—and lay out how, in that theory, after adding a large amount of firm and industry-level heterogeneity in parameters and shocks, the presence of orthogonal variation in subsidies for automation capital can identify the causal impact of automation on the labor share. We state the assumptions under which this identification scheme is valid, and discuss its validity and limitations in the context of our analysis. We then show how to structurally estimate the key parameter responsible for that relationship in the model: the elasticity of substitution between labor and automation capital. While we use this analysis in the context of a particular dataset, our approach is general and can be applied to identify the impact of automation on the labor share whenever (exogenous) variation in prices is available.

Substantively, we find that the firm-level elasticity of substitution between automation capital and labor falls between 3 and 4.5 across different empirical specifications. Our preferred point estimate is 3.8. The fact that this elasticity is significantly larger than one indicates that automation capital and labor are strong substitutes among actively automating firms, and so cheaper automation technologies do have a large negative impact on the labor share of income at the firm level.

⁵China is a good case to study the impact of automation as it has been one of the most aggressive adopters of robots in the world in the past decade. The staggering pace of robot adoption in China is evident from the fact that China’s share in the global market for robots went from 3.7 percent in 2005 to around 30 percent in 2016. China’s MIC aims to raise the global market share of Chinese-made robots to over 50 percent by 2020. At the same time, China’s robot density was below the global average, with only 68 units per 10,000 workers in 2016, comparing to the US stock of almost 200. Source: National Federation of Robotics, Statista.com.

We provide an assessment of whether this fairly high elasticity is consistent with the aggregate trends in our dataset and with the overall Chinese economy during this time period. To do so, we calibrate our model to match key moments characterizing our sample, such as the growth of value added among automation firms, the growth in the stock of automation capital, the growth of wages etc. We find that the firm-level microeconomic elasticity of substitution between automation and labor that we estimate is consistent with the observed decline in the aggregate labor share. Moreover, we show that a lower elasticity would counterfactually imply too much growth in the value added of automating firms relative to the industry average. We also find that the aggregate impact of automation ultimately depends on how spread out automation is across firms, which we refer to as the extensive margin of automation.

Finally, while our analysis pertains to Chinese data, the well-documented decades-long trend of declining quality-adjusted price of automation equipment sheds light on the declining labor share in the U.S. manufacturing sector during the last several decades—which fell from an average of .61 to .41 between 1960 and 2000s.⁶

Related literature. Our paper is one of a few that estimate the causal impact of modern automation on the labor share. The most closely related studies are [Acemoglu and Restrepo \(2019\)](#) and [Graetz and Michaels \(2018\)](#). Both papers use different variants of a shift-share type of identification and do not focus on the labor share. The effects they measure capture general equilibrium adjustments that may take place due to changes in prices and factor mobility over the long-run. In contrast, our identification focuses on micro-level elasticities that are embedded in the production technologies. Both approaches have their strengths, with ours being particularly useful as an input into modeling.⁷

The paper by [Autor and Salomons \(2018\)](#) is also related to our work due to its focus on the link

⁶See [Kehrig and Vincent \(2018\)](#) for a comprehensive analysis of the trends in the U.S. manufacturing sector, which these numbers are taken from. While their analysis does not explicitly relate observed changes to automation, it does eliminate several other possibilities. Developed countries were ahead of China in terms of automation by at least a decade as of 2015, but since then Chinese automation has progressed at an extremely rapid pace. For quality-adjusted series of the price of industrial robots, see [Graetz and Michaels \(2018\)](#) (Figure 1). Their evidence points to a decline between 1990 and 2004 by a factor of five (based on 1990 dollars). For comparison, nominal wages grew on average 105 percent in the six reported countries. We are not aware of more recent quality-adjusted series but nonadjusted series are available from various industry sources.

⁷The shift-share identification does not automatically remove the effects of offshoring. Import competition from low-wage countries such as China (or within NAFTA) might lead firms in rich countries to shift focus to capital-intensive products, and this, rather than automation, might reduce employment and wages in the most exposed areas in all developed countries. [Acemoglu and Restrepo \(2019\)](#) discuss this issue and devise supplementary ways of addressing it. They find that automation and import competition are not as strongly correlated and conclude this is unlikely to be a concern for their study.

between the labor share and productivity growth. It does not, however, provide causal identification of the impact of automation. Finally, in a related and complementary work to ours, [Humlum \(2019\)](#) provides a comprehensive analysis of robot adoption in Denmark and the distributional impact of robots on employment.⁸

The rest of the paper is organized as follows. Section 1 discusses the theory of automation that we use to develop our identification strategy. Section 2 discusses the data, our empirical strategy, and presents the results. Section 3 discusses the aggregate implications of our calibrated model and provides robustness analysis of our estimation procedure using model-generated data.

1 Model

We start by laying out the theory that forms the basis for both our empirical and quantitative analysis. Our model is fairly standard and links the quality-adjusted price of automation capital to automation investment and the labor share.

1.1 Environment

Time is discrete and the horizon is infinite. There are n cities indexed by $c \in \{1, 2, \dots, n\}$ and m industries indexed by $i \in \{1, 2, \dots, m\}$. For convenience, we refer to the tuple (c, i) as an island. Islands are populated by a continuum of firms of a fixed measure $\Omega_{ci} > 0$, each producing a differentiated good that is aggregated into a composite homogenous consumption good sold at home and abroad at a unitary (global) price. Firms and goods are indexed by a unique identifier $\omega \in \Omega_{ci} \subset \mathbb{R}$, and we denote by $c(\omega)$ and $i(\omega)$ the city and industry of firm/good ω . There is an economy-wide labor market such that the wage rate is common to all firms. The economy is sufficiently small to take world prices as given and so the cost of funds are exogenous. To model shocks to firm-level markups—a major concern in the measurement of the causal link between automation and the labor share—we assume that a firm ω operates in one of k distinct markets, indexed by $j \in \{1, 2, \dots, k\}$, on each island. The transition between these markets follows a Markov process independent of any firm or island characteristic, and a law of large number yields constant shares of firms in each market. We denote by $\Omega_{cij} \subset \Omega_{ci}$ the set of firms on island (c, i) that operate in market j . This structure allows us to demonstrate the robustness of our estimation procedure to the

⁸The work by [Gabriele et al. \(2018\)](#) studies the link between the labor share declines with the efficacy of the reallocation of labor across countries.

presence of such shocks.

Demand structure and aggregation

Goods are aggregated by two layers of competitive sectors: the final good sector and the intermediate good sector. On each island, the output of individual firms are combined into the production of an island-specific composite good and the island-specific composites are then again aggregated into a homogenous final good that is sold globally at a normalized (numeraire) price $P = 1$ —with the goods market clearing condition dropped to reflect this assumption.

The final good producers convert a vector of differentiated island composite goods Q_{ci} into Y units of final goods according to the production function

$$Y = \left(\sum_{c=1}^n \sum_{i=1}^m D_{ci}^{\frac{1}{\rho}} Q_{ci}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}, \quad (1)$$

where D_{ci} are fixed weights such that $\sum_{ci} D_{ci} = 1$, and $\rho > 0$ is the elasticity of substitution between goods. Final good producers take prices as given and maximize static profits given by

$$\max_{Y, Q_{ci}} P(t) Y - \sum_{c=1}^n \sum_{i=1}^m P_{ci}(t) Q_{ci},$$

where P_{ci} is the price of the composite good from industry i in city c , and where Y is given by the production function (1). The resulting demand function for the composite good from island i is

$$Q_{ci}(t) = D_{ci} \left(\frac{P_{ci}(t)}{P(t)} \right)^{-\rho} Y,$$

and the normalization of the price of the final good implies

$$P(t) = \left(\sum_{c=1}^n \sum_{i=1}^m D_{ci} P_{ci}(t)^{1-\rho} \right)^{\frac{1}{1-\rho}} \equiv 1.$$

At the island level, a sector of competitive producers aggregates goods of all firms from all markets on the island into a composite bundle Q_{ci} sold to the final good sector at a unit price P_{ci} . The producers solve

$$\max_{q(\omega)} P_{ci}(t) Q_{ci} - \sum_{j=1}^k \int_{\Omega_{cij}} p(t, \omega) q(t, \omega) d\omega,$$

where $p(t, \omega)$ and $q(t, \omega)$ denote the price and quantity of good ω , respectively. That industry uses the production technology

$$Q_{ci} = \prod_{j=1}^k \left(\int_{\Omega_{cij}} d(t, \omega)^{\frac{1}{\theta_j}} q(t, \omega)^{\frac{\theta_j-1}{\theta_j}} d\omega \right)^{\phi_j \frac{\theta_j}{\theta_j-1}}.$$

The terms $d(t, \omega)$ are time-varying demand shocks affecting each individual firm and that follow an arbitrary process. The parameter $\theta_j > 0$ is the elasticity of substitution across goods in market j , and ϕ_j is the intensity of market j . We impose constant returns to scale on the production function and so $\sum_{j=1}^m \phi_j = 1$. By the zero profit condition, equilibrium prices satisfy

$$P_{ci}(t) = \prod_j \left(\frac{P_{cij}(t)}{\phi_j} \right)^{\phi_j},$$

and the price index for market j in city c in industry i at time t is

$$P_{cij}(t) = \left(\int_{\Omega_{cij}} d(t, \omega) p(t, \omega)^{1-\theta_j} d\omega \right)^{\frac{1}{1-\theta_j}}.$$

The demand curve for products produced by firm ω on island (c, i) and operating in market $j(t, \omega)$ is given by

$$q(t, \omega) = d(t, \omega) \left(\frac{p(t, \omega)}{P_{cij(t, \omega)}(t)} \right)^{-\theta_j(t, \omega)} \left(\frac{P_{ci}(t)}{P_{cij(t, \omega)}(t)} \right) \phi_{j(t, \omega)} Q_{ci}(t). \quad (2)$$

Notice that the market $j(t, \omega)$ the firm operates in, and which changes stochastically over time, shifts the overall demand for its goods (through ϕ_j) and the elasticity of that demand is θ_j .

Production technology

The production technology of a firm combines structures k_s , equipment k_e , automation capital m , support labor l_s and production labor l , and it is summarized by the production function⁹

$$F_\omega(k_s, l_s, k_e, l, m) = A(t, \omega) (k_s^{\gamma_\omega} l_s^{1-\gamma_\omega})^{\eta_\omega} \left(a_\omega^{\frac{1}{\sigma_\omega}} (k_e^{\alpha_\omega} l^{1-\alpha_\omega})^{\frac{\sigma_\omega-1}{\sigma_\omega}} + (1 - a_\omega)^{\frac{1}{\sigma_\omega}} m^{\frac{\sigma_\omega-1}{\sigma_\omega}} \right)^{(1-\eta_\omega) \frac{\sigma_\omega}{\sigma_\omega-1}}. \quad (3)$$

The block $k_s^{\gamma_\omega} l_s^{1-\gamma_\omega}$ corresponds to activities (tasks) in support of production activities, which are captured by the last term in parentheses. For instance, support capital k_s might include buildings

⁹The production function pertains to the firm's value added and does not involve intermediate goods.

and structures; support labor l_s could include management, sales force, or any support employees that assist production activities but whose employment is only indirectly affected by automation. These inputs are aggregated in a Cobb-Douglas fashion. The parameter $0 \leq \gamma_\omega \leq 1$ determines the intensity of capital and the parameter $0 \leq \eta_\omega \leq 1$ determines how intensive the firm is in support and productive activities. $A(t, \omega)$ is time-varying total factor productivity.

The second block of the production function (between the large parentheses) describes production activities. These activities can be completed by a mix of equipment k_e and production-level labor l , also combined in a Cobb-Douglas fashion, or by using labor-saving automation capital m . The first option captures traditional techniques for completing production activities and is characterized by a price-invariant labor share determined by the parameter $0 \leq \alpha_\omega \leq 1$. The second option captures automation and involves no labor. The elasticity of substitution $\sigma_\omega > 0$ links the two terms and determines how easily it is to automate activities that would traditionally be produced using an equipment-labor mix. The parameter $0 \leq a_\omega \leq 1$ describes how intensive the production function is in automation capital.

We use the parameter a_ω to distinguish between the intensive and extensive margin of automation investment. The *intensive* margin refers to the accumulation of automation equipment m by firms that are able to use that equipment in production ($a_\omega < 1$). In contrast, the *extensive* margin of adoption, or the spread of automation across firms refers to the distinction between firms that can use automation ($a_\omega < 1$) and those that cannot ($a_\omega = 1$).

The parameters of the production function in (3) are indexed by ω and they may vary across firms. We allow the distributions of these parameters to be arbitrary at this point, although the calibration of the model imposes more structure on how parameters are distributed.

To gain intuition about the notion of automation conveyed by the production function (3), consider a simple task that involves nailing two pieces of material together. This task can be done by a worker (productive labor l) who uses a hammer (equipment k_e) or, alternatively, by an autonomous hammering robot (automation capital m) that performs the task on its own. The elasticity of substitution σ_ω then describes how easy it is for the firm to substitute production labor by automated machinery with any such “automatable” task. Finally, whoever does the hammering (worker or robot) might need a building (support capital k_s) as a place to work and a manager (support labor l_s) to supervise production, which is captured by the term in front of the bracket.

Acemoglu and Restrepo (2019) microfound this interpretation of the above production function using a task-based model. In their theory, the parameter a_ω maps onto the number of tasks for

which automation technology has been developed to date (automatable tasks), and the elasticity of substitution σ_ω maps onto the parameters pertaining to the task-level production function, which determines how much labor and capital the firm chooses per automatable task, given prices.

A notable feature of the above production function is that it is consistent with the [Kaldor \(1961\)](#) fact that the labor share remained constant for several decades. Indeed, if automation capital is impossible to obtain (infinite price for m), (3) collapses to a standard Cobb-Douglas aggregator with a constant labor share equal to

$$LSO(t, \omega) := \frac{\theta_{j(t, \omega)} - 1}{\theta_{j(t, \omega)}} \left(\underbrace{(1 - \eta_\omega)(1 - \alpha_\omega)}_{LSOP} + \underbrace{\eta_\omega(1 - \gamma_\omega)}_{LSON} \right). \quad (4)$$

But once the first automation technologies start to become available, firms might begin using them in production, in which case the labor share $LS(t, \omega)$ will move away from its initial $LSO(t, \omega)$ value. For later use, we refer to LSO as the automation-free labor share, $LSOP$ as the automation-free labor share in production activities and $LSON$ as the automation-free labor share in support (nonproduction) activities. By construction, automation only affects labor share via production activities pertinent to $LSOP$.

Firm problem

We denote the equilibrium wage rate by w . The user costs associated with the three forms of capital available to the firm are assumed to be determined by a fixed global cost of funds r , the price of each type of capital and their corresponding depreciation rates. We denote the user cost of structures by $r^s(t)$, the user cost of equipment by $r^e(t)$, and the user cost of automation by $r^m(t)(1 - s(t, \omega))$, where $s(t, \omega)$ captures any policy-induced time-varying changes in the user cost of automation capital that may vary across firms. Note that $s(t, \omega)$ can vary across firms and across time. Later on, we will map $s(t, \omega)$ to government subsidies for automation investment that we observe in the data and use it to develop our identification strategy.

The user costs of capital are implied by our assumptions of fixed global cost of funds r , the assumption of linear depreciation of capital, and a fixed relative price of each type of capital good in terms of the global final good (which is the numeraire). For instance, if we denote by p^m the price of

one unit of automation capital, then we can write

$$\underbrace{p^m(t)(1-s(t,\omega))(1+r)}_{\text{cost of funds for purchase}} - \underbrace{(1-\delta^m)p^m(t+1)(1-s(t,\omega))}_{\text{value of undepreciated capital}} = r^m(t)(1-s(t,\omega)), \quad (5)$$

where δ^m is the depreciation rate of automation capital and

$$r^m(t) := p^m(t)(1+r) - (1-\delta^m)p^m(t+1). \quad (6)$$

is the user cost of automation capital before subsidies. This formula assumes that the subsidy affects the cost of replacing undepreciated capital $(1-\delta^m)p^m$, which makes sense in the context of our data since the duration of “Made in China 2025” is about a decade and the estimated lifetime of automation equipment and robots is also about a decade. Under this assumption an increase in the subsidy rate s and a decline in the price of capital p^m are equivalent. User costs of other forms of capital are defined analogously to (6).

The problem of the firm boils down to a static profit maximization problem of the form

$$\pi(t,\omega) := \max_q (p(t,\omega) - \lambda(t,\omega))q, \quad (7)$$

where the price $p(t,\omega)$ satisfies the demand equation (2) and $\lambda(t,\omega)$ is the marginal cost of production, which is defined by the cost minimization problem of the form

$$\lambda(t,\omega) := \min_{l,k_e,k_s,m} r^s(t)k_s + r^e(t)k_e + r^m(t)(1-s(t,\omega))m + w(t)(l_s+l), \quad (8)$$

subject to $F_i(k_s, l_s, k_e, l, m) \geq 1$. It is straightforward to show that a standard constant markup pricing rule applies here as long as each firm is atomistic, and so

$$p(t,\omega) = \frac{\theta_{j(t,\omega)}}{\theta_{j(t,\omega)} - 1} \lambda(t,\omega), \quad (9)$$

where $j(t,\omega)$ is the market in which firm ω currently operates.¹⁰ As expected, a change in market j is akin to a markup shock. The lemma below characterizes the marginal cost of production $\lambda(t,\omega)$ as a function of prices and parameters.

¹⁰By assuming a continuum of firms in each sector we assume away the kind of strategic considerations that lead to variable markups as in [Atkeson and Burstein \(2008\)](#).

Lemma 1. *The marginal cost of production $\lambda(t, \omega)$ is given by*

$$\lambda(t, \omega) = \frac{1}{A(t, \omega)} \left(\frac{\lambda^s(t, \omega)}{\eta_\omega} \right)^{\eta_\omega} \left(\frac{(a_\omega \lambda^e(t, \omega))^{1-\sigma_\omega} + (1 - a_\omega) (\lambda^m(t, \omega))^{1-\sigma_\omega}}{1 - \eta_\omega} \right)^{\frac{1}{1-\sigma_\omega} \eta_\omega}, \quad (10)$$

where

$$\lambda^s(t, \omega) = \left(\frac{r^s(t)}{\gamma_\omega} \right)^{\gamma_i} \left(\frac{w(t)}{1 - \gamma_\omega} \right)^{1-\gamma_\omega}, \quad \lambda^e(t, \omega) = \left(\frac{r^e(t)}{\alpha_i} \right)^{\alpha_\omega} \left(\frac{w(t)}{1 - \alpha_\omega} \right)^{1-\alpha_\omega}, \quad \lambda^m(t, \omega) = r^m(t) (1 - s(t, \omega)).$$

Proof. All proofs are in Appendix 7. □

The effect of the price of automation on the marginal cost is captured by the last term in bracket in (10) and it depends on the price of automation capital and the elasticity of substitution between automation capital and labor σ_ω .

Total labor supply in the economy is fixed and the wage rate w is determined by the usual economy-wide market clearing. The assumption of a global market for a homogenous final good and common interest rate r eliminates the need for an explicit statement of the household problem and we omit it. The formal definition of equilibrium is standard and we also omit it.

1.2 Characterization of equilibrium labor share dynamics

We now provide a preliminary characterization of the impact of the price of automation capital on automation and the labor share. These equations are fundamental to the empirical strategy developed in the next section and provide the key intuition for the workings of the model.

We begin with a first result that relates the subsidy $s(t, \omega)$ to the automation to labor ratio $m(t, \omega) / l(t, \omega)$.

Lemma 2. *The automation to labor ratio $m(t, \omega) / l(t, \omega)$ is given by*

$$\log \frac{m(t, \omega)}{l(t, \omega)} = \Theta(t, \omega) - \sigma_\omega \log(1 - s(t, \omega)), \quad (11)$$

where

$$\Theta(t, \omega) := \log \frac{1 - a_\omega}{a_\omega} + \alpha_\omega \log \frac{\alpha_\omega}{1 - \alpha_\omega} + \alpha_i \log \frac{w(t)}{r_\omega^e} + \sigma_\omega \log \frac{\lambda_\omega^e(t)}{r_i^m}. \quad (12)$$

Intuitively, the subsidy s_ω pushes the firm to acquire more automation capital m relative to the number of productive labor l it employs. The magnitude of this effect depends on the elasticity of substitution σ_ω . Since the elasticity is constant, a one percent reduction in the price of automation capital, perhaps from a subsidy, is associated with a σ_ω percent reduction in automation capital per (production) worker, m/l .¹¹

We next move to the labor share, which in the model is given by

$$LS(t, \omega) := \frac{w(t)(l(t, \omega) + l_s(t, \omega))}{y(t, \omega)}, \quad (13)$$

where $y(t, \omega) := p(t, \omega)q(t, \omega)$ denotes the output of firm ω (its value added).

The following lemma completes the characterization of the link between automation and the labor share.

Lemma 3. *The labor share $LS(t, \omega)$ of a firm depends on its automation to labor ratio $m(t, \omega)/l(t, \omega)$ through the expression*

$$\log \frac{LSO(t, \omega) - LS(t, \omega)}{LS(t, \omega) - LSN(t, \omega)} = \Psi(t, \omega) + \frac{\sigma_\omega - 1}{\sigma_\omega} \log \frac{m(t, \omega)}{l(t, \omega)}, \quad (14)$$

where $LSO(t, \omega)$ is given by equation (4),

$$LSN(t, \omega) := \frac{w(t)l_s(t, \omega)}{y(t, \omega)} = \frac{\theta_{j(t, \omega)} - 1}{\theta_{j(t, \omega)}} (1 - \gamma_\omega) \eta_\omega, \quad (15)$$

is the labor share of support workers, and where

$$\Psi(t, \omega) := \frac{1}{\sigma_\omega} \log \frac{1 - a_\omega}{a_\omega} - \alpha_\omega \frac{\sigma_\omega - 1}{\sigma_\omega} \log \left(\frac{\alpha_\omega}{1 - \alpha_\omega} \frac{w(t)}{r_t^e} \right).$$

This lemma implies that automation capital per production-level employee (m/l)—determined by the cost of automation by Lemma 2—affects the labor share. The left-hand side of (14) is a decreasing function of $LS(t, \omega)$, and so any increase in the right-hand side of the equation leads to a lower labor share. For instance, if automation capital m and the equipment-labor bundle $k_e^{\alpha_\omega} l^{1-\alpha_\omega}$ are substitutes ($\sigma_\omega > 1$), the adoption of automation technologies by the firm pushes the labor share down. As expected, the labor share is unaffected by investment in automation in the Cobb-Douglas

¹¹One advantage of normalizing automation capital m by productive labor l (instead of, say, value added) is that the demand shock terms θ_j and ϕ_j do not appear in (11) and (12). The instrument would nonetheless be valid if such shocks were independent of $s(t, \omega)$ but it would not be possible to structurally identify parameter σ in that case.

case ($\sigma_\omega = 1$).

The terms $\Theta(t, \omega)$ and $\Psi(t, \omega)$ in the above lemmas show that the parameter a_ω has an analogous effect on the labor share to the subsidy s_ω , with the elasticity σ_ω determining the sensitivity of the labor share to changes in a_ω .

2 Empirical results

Having laid out the theory, we now turn to the data. We begin by describing our dataset. We then discuss how we use the two key equations derived in Lemmas 2 and 3 to identify the effect of automation on the labor share, and to structurally estimate the elasticity of substitution between labor and automation capital.

2.1 Data

Our data comes from the China Enterprise General Survey (CEGS).¹² The CEGS is a longitudinal large-scale study of manufacturing firms and workers in China conducted in three waves—2015, 2016 and 2018. The 2018 wave covers five provinces across different geographic parts of China: Guangdong in the south, Hubei in the middle, Jiangsu in the east, Sichuan in the West, and Jilin in the North.¹³ Because of its larger coverage, we use the 2018 wave, which retroactively provides data for the years 2015, 2016 and 2017 in a consistent format. Data collection has been meticulously done by a team of economists traveling to site.¹⁴ There are 1618 unique firms in our sample.¹⁵

In our analysis we use information on the wage bill and employment by type of workers and value added. One unique feature of the data is that it distinguishes between various types of capital equipment. We use data on investment in fully-automated industrial robots (Machine-1 in survey), and numerically controlled (CNC) semi-automated machinery (Machine-2), the information on the

¹²The name of the survey changed in 2020 from the China Employer Employee Survey (CEES).

¹³Firms were sampled from the third National Economic Census conducted in 2014. The sampling was conducted in two stages, each using probability proportionate-to-size sampling, with size defined as manufacturing employment. Therefore, the firm-level sample is representative in terms of employment size in China. Employees were surveyed with stratification: 6 to 15 employees were randomly selected in each firm, among which 2 to 3 were middle and senior managers.

¹⁴The survey design team confirmed to us that for large firms it was not uncommon for the reviewers to stay on site for weeks to collect the data. In addition, hard data points, such as those pertaining to a firm's financials, have been taken from accounting records pulled on site.

¹⁵Overall, the response rate of firms was $2019/2417 = 83.5\%$. About 400 firms were dropped in the process of data cleaning. We do not have the information about the exact procedure by which the raw data was cleaned by the data center in Wuhan. We were informed that the data center plan a public release of the cleaning procedure at a later date.

subsidies received from the government for the purchase of each type of machinery, and also the data on all other forms of capital (excluding structures) which we refer to as “ordinary capital”.¹⁶

We define automation investment as the purchase of both Machine-1 and Machine-2 equipment. A Machine-1 piece of equipment is an industrial robot as defined by ISO 8371: “an automatically controlled, reprogrammable multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications.” The Machine-2 category has been specifically designed for this survey to capture advanced labor-saving automation machinery that does not meet the stringent requirements of ISO 8371. It covers production-line numerically controlled equipment that may involve human operators but performs most of the work autonomously.

The key aspect of the survey that enables our analysis is that its timing overlaps with the first phase of “Made in China 2025” (MIC, hereafter)—a vast government-led program that placed high-tech labor-saving automation technologies at the forefront of national industrial policy.¹⁷ MIC introduced sizable subsidies paid to firms as a discount to the purchase price of automation capital. Our dataset provides detailed information on the payments of these subsidies between 2015 and 2017, including the type of equipment they target. Importantly, the implementation of MIC fell largely on local governments, which had some flexibility in the types of policies and subsidy rates to implement. As a result, we see a large amount of variation in subsidies across cities, industries and individual firms.¹⁸ By October 2016, at least 70 provinces, cities and county-level administrations had released local MIC 2025 strategies with specific local priorities.¹⁹ We exploit the variation in this policy to

¹⁶The sample is designed to be representative of the Chinese manufacturing sector, but we do not have the weights needed to construct representative statistics at this point. As a result, our aggregated analysis pertains to the sample.

¹⁷While there are no official figures on the scale of the program, the goals are unprecedented and indicative of a sizable budget. For example, the objective of the program is to increase the adoption rate of automation from 33 percent of firms in 2015 to 64 percent in 2025. Other objectives are equally ambitious. Information about the program can be found at the official website of the State Council of China: <http://www.gov.cn/zhuant/2016/MadeinChina2025-plan/>. The “Notice of the State Council on Made in China 2025” summarizes these goals as follows: “pilot the construction of smart factories/digital workshops in key areas, accelerate the application of technologies and equipment such as human-machine intelligent interaction, industrial robots, smart logistics management, and additive manufacturing in the production process, and promote the simulation and optimization of manufacturing processes, digital control, and status information real-time monitoring and adaptive control.” Translated from official document found at http://www.gov.cn/zhengce/content/2015-05/19/content_9784.htm using Google Translate.

¹⁸One goal of this design was to foster experimentation to identify the most effective policies. In total, there were 30 pilot cities assigned by MIC 2025, fifteen of which are included in the CEES Survey: Seven in Guangdong Province, five in Jiangsu Province, one in Hubei Province, one in Sichuan Province and one in Wuhan Province. The experiments targeted ten industrial sectors: 1) next generation information technology; 2) robotics and advanced automatic machinery; 3) aerospace and aviation equipment; 4) maritime engineering equipment and advanced maritime vessel manufacturing; 5) advanced rail equipment; 6) new energy vehicles; 7) advanced electrical equipment; 8) agricultural machinery and equipment adoption; 9) new material; 10) biomedicine and high-performance medical devices.

¹⁹Appendix 5 provides three examples of MIC implementations, which show broad-based subsidy policy towards

construct an instrument for automation investment. Specifically, as we explain below, we map subsidies onto $s(t, \omega)$ in our model and devise an IV strategy to identify the causal impact of automation on firm-level labor shares. Before we proceed, we provide an overview of the data.

2.2 Data structure and summary statistics

Tables 1 and 2 provide a preliminary characterization of the dataset. These tables group firms into those that report investment in automation in any year between 2015 and 2017 (automating firms) and those that report receiving subsidies for purchases of automation capital in any year during this time period (subsidized firms). We back out the subsidy rate by dividing the subsidy payments received from the government for the purchases of automation equipment by the total amount of automation purchases reported by firms—as opposed to using the actual policy variables for which we do not have a complete information at a firm level. Hence, by definition, all subsidized firms in our data invest in automation capital.

Statistic	All firms	Automating firms	Subsidized firms
Number of cities	60	44	18
Number of industries	31	23	14
Number of city-industry pairs	666	117	31
Number of observations	4602	491	106
Number of (unique) firms	1618	171	37
Share in total employment (in %)	100%	23%	5.6%
Share in total value added (in %)	100%	24%	5.9%

Table 1: Sample Structure (all years 2015-2017)

As Table 1 shows, the data over the sample period 2015-2017 covers 60 cities and 31 industries, altogether 666 city-industry pairs and 1618 firms, amounting to 4602 observations. Of those, there are 171 unique automating firms (with 491 observations in total) during this time period and they are found in 44 cities, 23 industries and 117 city-industry pairs. The sample of firms that report receiving subsidies for automation are spread out over 18 cities, 14 industries and 31 city-industry pairs, and amount to 37 firms (106 observations). In the majority of city-industry pairs in which there is at least one subsidized firm, the majority of automating firms report receiving subsidies (about

purchases of automation equipment. In 2015, MIC wasn't effectively in place because not many localities formulated their programs and MIC was introduced in the middle of the year. We do not have information about the exact timing of the introduction of the subsidies as for 2015, but we know that by the end of 2016 they were already introduced in most places.

80 percent). While there are relatively few producers that automate, they are considerably larger and account for about a quarter of total value added and employment. Firms that report receiving subsidies for automation account for about a quarter of firms that automate.²⁰

The fact that there are fewer subsidized firms than automating firms is a direct consequence of the MIC design. As the examples discussed in Appendix 5 illustrate, not all cities subsidize automation across all industries. The rules vary and they may prevent some firms from taking advantage of the subsidies in place. For instance, some cities require that industrial robots be domestically produced. Other have caps on the total budget of the program such that some firms might be too late to be receive a payment. We will later come back to the different reasons why firms within the same city-industry pairs might not receive the same subsidy when discussing our empirical strategy.

Table 2 provides summary statistics about the three groups of firms in our sample. As we can see from column 2 (N), all variables are well-populated, albeit not perfectly, and there is some variation in coverage across variables. Automating firms tend to be larger and subsidized firms are even larger. Both weighted and unweighted mean labor shares are declining over time, and their decline is more pronounced among subsidized firms.²¹ The mean subsidy rate among subsidized firms is 12 percent, and among automating firms it is below 4 percent. Subsidies and investment in automation are highly skewed, and most automating firms do not receive any subsidies for automation.²²

The decline in the labor share is correlated with investment in automation. Consider a simple OLS specification of the form

$$\Delta LS_{\omega,t,t-1} = \alpha \log(x_{\omega,t-1}) + \delta_{\omega} + \varepsilon_t, \quad (16)$$

²⁰Our definition of automating firms is narrower than the potentially active extensive margin of automation may possibly be in long-run. Many firms that do not report any investment in automation between 2015 and 2017 do report having some stock of automation capital, and hence they too should be considered among firms that could automate. A firm may not be automating not only because it cannot automate but also because the time for automating is not good for idiosyncratic reasons (negative demand shock). Such a narrower definition of automating firms is nonetheless helpful because all information in our study comes from these firms.

²¹Many firms appear to erroneously report their total labor share as the labor share of production workers, since at the same time they report that only a fraction of their employment are production workers, which would be inconsistent. Under the assumption of wages being equalized across production and nonproduction activities, our model implies that the labor share of production employees can be backed out from the share of employment in production. The mean value of this share in the data is .63, which also implies a labor share of production workers of $.496 \times .63 = .31$, a value inconsistent with the raw data number of about .44. To rationalize this number would require higher wages among production workers than nonproduction workers, which we do not find plausible. To address this issue, we have dropped observations in which the labor share in production workers exceeds the total labor share, which gives $\overline{LSP} = 0.32$ and implies $\overline{LSN} = .485 - .33 = .15$.

²²We winsorize the top and bottom 5% of the sample variables. Whenever we construct a ratio we winsorize this ratio similarly. We discount by inflation such that nominal values are measured in 2015 RMB.

Variable ¹	<i>N</i>	Centrality			Dispersion		Skewness
		W.Mean ²	Mean	Median	S.D.	<i>P</i> ₉₀ - <i>P</i> ₁₀	$\frac{P_{90}-2P_{50}+P_{10}}{P_{90}-P_{10}}$
All firms							
Employment	4,602	—	307	100	473	884	.81
- % in production	4,602	.63	.63	.65	.18	.48	-.18
Auto. investment/VA ³	688	.0691	.066	0	.145	.242	1
LS '15 (labor share)	1,330	.496	.559	.535	.295	.791	.03
LS '17	1,491	.473	.537	.516	.276	.776	.06
LS '17–LS '15 (by firm)	1,296	-0.026	-.024	-.012	.1890	.369	-.05
LS of production employees '17	833	.33	.34	.32	.207	.301	.19
Automating firms							
Employment	491	—	656	357	661	1776	.64
- % in production	491	.62	.62	.64	.16	.37	-.09
Auto. investment/VA ³	455	.096	.099	.020	.169	.361	.89
LS '15	151	.50	.558	.534	.294	.814	-.02
LS '17	162	.485	.523	.511	.267	.725	-.05
LS '17–LS '15	147	-.022	-.036	-.019	.185	.343	-.00
LS of production employees '17	101 ⁴	.345	.362	.365	.206	.533	.03
Subsidy rate	320	.036	.037	.000	.11	.12	1
Subsidized firms							
Employment	106	—	791	660	670	1758	.33
- % in production	106	.61	.61	.63	.16	.37	-.17
Auto. investment/VA ³	104	.166	.140	.056	.195	.397	.72
LS '15	34	.502	.533	.534	.318	.855	-.02
LS '17	37	.483	.453	.511	.258	.788	-.11
LS '17–LS '15	34	-.035	-.083	-.02	.211	.382	-.73
LS of production employees '17	22	.35	.32	.343	.217	.58	.03
Subsidy rate	80	.116	.110	.098	.069	.175	.17

Table 2: Selected Summary Statistics

Notes: ¹Unless a specific year is noted, the values pertain to the average for the three-year period under consideration: 2015, 2016 and 2017. ² Weighted mean is obtained by weighting each observation by firm's share in total value added averaged across all years 2015-2017. ³ Output is measured by value added after subtracting all subsidies. ⁴ As described in text, many firms appear to be reporting total wage bill and we drop observations when the two are equal. *N* means the number of observations. We first add nominal values for the three years without discounting and then take the ratios.

where $\Delta LS_{\omega,t,t-1}$ is the change in labor share between two consecutive years, $t \in \{2015, 2016, 2017\}$, and $x_{\omega,t-1}$ denotes investment in machines by firm ω in year $t - 1$, δ_{ω} is a firm fixed effect, and ε_t is an error term. To contrast the effect of automation and that of other capital, we consider both automation investment (x_A) and other machine investment (x_O)—the later including all other capital excluding structures. The measures of investment in capital are normalized by value added (*VA*) to capture their importance compared to the size of the firm.²³

²³We exclude from value added the subsidies paid by government for automation to avoid spurious correlation.

As we can see from Table 3. It is clear from columns (1) and (2) that automation investment exhibits a negative and statistically significant association with the labor share. The size of the effect is also economically large. In contrast, columns (3) and (4) show that the relationship between investment in other types of machine and the labor share is only marginally significant. The magnitude of the effect is also much weaker, with the marginal impact of automation investment larger by a factor of twenty.²⁴

Dependent Variable: $LS_{\omega,t-1} - LS_{\omega,t}$				
Explanatory variables	Full Sample			
	(1)	(2)	(3)	(4)
$\log\left(\frac{x_A}{VA}\right)_{t-1}$	-.0320* (.0186)	-.0120** (.00565)		
$\log\left(\frac{x_Q}{VA}\right)_{t-1}$			-.0073* (0.027)	-.00142 (0.00561)
Firm fixed effect	Yes	No	Yes	No
Year fixed effect	Yes	Yes	Yes	Yes
R-squared	0.81	0.03	0.59	.00
Observations	186	186	233	233

Table 3: Change in Labor Share and Automation Investment (OLS).

Notes: Least squares estimation. *** indicates significance at 1 %, ** at 5% and * at 10% level of confidence. Standard errors in parenthesis. Sample is restricted to observations with positive investment to ensure log values are well-defined. Value added calculation excludes government subsidies for investment that could bias the results.

While these results show that the decline in labor share exhibits a higher partial correlation with investment in automation compared to that with investment in other machines, interpreting the automation investment coefficient in terms of structural parameters is fraught with problems. First, OLS does not address the issue of causality, i.e. changes in labor shares and investment in automation are endogenous responses of firms to changes in goods and factor prices, and so we need exogenous variation to provide evidence of a causal effect of automation on the labor share. Second, even with an instrument to uncover causality, we need a specification informed by a structural model

Normalizing by production labor instead of value does not have a meaningful impact on the results.

²⁴We test the robustness of the above relationship by using alternative specifications where we replace firm fixed effects with industry, city and year fixed effects but include commonly used time-varying firm specific attributes as control variables, namely, i) value added per worker ($\frac{VA}{L}$) which controls for productivity; ii) capital per worker ($\frac{K}{L}$) which controls for the capital intensity of a firm; iii) the ratio of total debt to total assets which controls for the extent to which a firm is leveraged (LR); and iv) the export to total sales ratio ($\frac{X}{TS}$) which controls for export intensity (higher export intensity tends to be associated with higher productivity, higher quality of products, more skilled workforce, and a greater product scope). All control variables are lagged by one period. The results are qualitatively the same and we report them in Table ?? of the Technical Appendix. Export intensity, value added per worker, and leverage exhibit statistical significance.

in order to estimate key parameters of the model. The next section discusses how we address these challenges using subsidies for automation as an instrument in a 2SLS setup.

2.3 Identifying the effect of automation: theory

The equations derived in lemmas 2 and 3 readily suggest a path to estimate the causal impact of automation on the labor share and the average elasticity $\mathbb{E}[\sigma_\omega]$ by using the observed variation in subsidy rates s_ω as an instrument.²⁵ Indeed, from (11) and (14) we know that s_ω affects the labor share only through its impact on the m/l ratio. Accordingly, a subsidy rate that is appropriately orthogonal to the firm’s parameters and shocks can be used as a valid instrument to identify the effect of automation on the labor share. However, using these lemmas for the estimation is not possible in practice. First, we do not observe the terms $LSO(\omega, t)$ and $LSN(\omega, t)$ in the right-hand side of (14). As a result, we cannot compute the dependent variable in the second stage regression. Second, we do not have reliable information on the *stock* of automation capital. We now show how this theory can be extended to estimate the effect of automation using the data that we have.

In what follows, let τ denote the three-year period 2015-2017 covered by the survey, and $\tau - 1$ denoted the previous three-year period (which, in principle, we do not observe). Unless otherwise noted, we compute flows in period τ as the discounted sum of the flows for the years 2015, 2016 and 2017 (defined precisely in next section). For stocks, we take the average of the three years. To calculate the change in the labor share across periods—lacking data for the preceding period—we take the difference between endpoints; that is, years 2017 and 2015. Such an approach, if anything, likely understates the change in the labor share and would work against finding a strong effect of automation.

Assumptions

To obtain identification, we impose two assumptions on the relationship between the subsidy rates and other elements of the model. We will discuss the empirical relevance of these assumptions in the next section.

²⁵Throughout, we work with an expectation operator $\mathbb{E}[x]$, that is to be understood as the mean value of any variable or parameter x across all firms:

$$\bar{x} := \mathbb{E}[x] = \frac{1}{n} \sum_i \sum_c \int_{\omega \in \Omega_{ci}} x(t, \omega) d\omega,$$

and analogously for any conditional expectation operator.

Our first assumption is on the probabilistic structure of the subsidies.

Assumption 1. *The subsidy $s(\tau, \omega)$ follows*

$$s(\tau, \omega) = s_i(\tau) + s_c(\tau) + \varepsilon^s(\tau, \omega), \quad (17)$$

where s_i and s_c are industry and city-specific random variation at time τ and $\varepsilon^s(\tau, \omega)$ is a mean-zero i.i.d. firm specific subsidy residual.

Our second and key identifying assumption requires that the subsidy residual $\varepsilon^s(\tau, \omega)$ be orthogonal to other exogenous variables or parameters.

Assumption 2. *The random variable $\varepsilon^s(\tau, \omega)$ is orthogonal to any parameter, shock or factor price (or their combination) $z(t, \omega)$ in the sense that*

$$\mathbb{E}[z(t, \omega) | \varepsilon^s(\tau, \omega)] = \mathbb{E}[z(t, \omega)]$$

for all t .

The above assumptions play a key role in our analysis. They ensure that the residual of the subsidy rate, after being projected onto city and industry fixed effects, satisfies the exclusion restriction for instrumental variable (IV) estimation. Note that the second assumption applies not only to the exogenous variables but also to the endogenous wage rate w , to industry-level price indices, as well as to the subcomponents of the subsidy: s_i and s_c in equation (17). This requires the subsidy program to be relatively small compared to the size of the relevant markets, including the labor market, so that the subsidy residual ε^s does not move these equilibrium objects “too much.” Since this residual has mean zero, and only a fraction of firms are subsidized, we do not consider this to be a serious limitation. In the quantitative section we will nonetheless test the validity of this assumption using our calibrated model.

Furthermore, Assumption 2 implies that previous subsidy rates, which act as parameters, are similarly orthogonal to the residual $\varepsilon^s(\tau, \omega)$. While we view this assumption as reasonable, we show in Appendix 8 that the estimates presented below are conservative, in the sense that they provide a lower bound on the mean elasticity of substitution $\mathbb{E}[\sigma_\omega]$, so long as the residuals $\varepsilon^s(\tau - 1, \omega)$ and $\varepsilon^s(\tau, \omega)$ are positively correlated.

Identification results

We next work out the link between a firm's investment in automation capital to the investment subsidy it faces. The result hold up to a first-order approximation (denoted by the symbol \approx) of the left-hand side of 14 and of the firm's policy functions with respect to the subsidy rate s . We also validate these approximations using numerical simulations of our model in the next section.

To be consistent with the institutional framework described in Section 2.2, we focus on investment decisions between two periods $\tau - 1$ and τ , defined as

$$x(\tau, \omega) := p_m(\tau) m(\tau, \omega) - (1 - \delta_i^m) p_m(\tau - 1) m(\tau - 1, \omega). \quad (18)$$

The following lemma relates investment $x(\tau, \omega)$ to the subsidy residual.

Lemma 4. *The automation investment intensity is related to the subsidy residual through the equation*

$$\mathbb{E} \left[\frac{x(\tau, \omega)}{l(\tau, \omega)} | \varepsilon_\omega^s(\tau, \omega) \right] \approx \mathcal{L} \varepsilon_\omega^s(\tau, \omega) + cte, \quad (19)$$

where \mathcal{L} and cte are some constants.

This lemma is analogous to Lemma 2 and the same intuition applies, with the crucial difference that here it involves the investment rate x normalized by labor employed in production instead of the stock of automation capital m .

We now turn to our second result that links the conditional expectation of the change in a firm's labor share to the conditional expectation of its investment in automation.

Lemma 5. *The change in the labor share is related to automation investment intensity through the equation*

$$\mathbb{E} [LS(\tau, \omega) - LS(\tau - 1, \omega) | \varepsilon_\omega^s(\tau, \omega)] \approx \mathcal{B} \mathbb{E} \left[\frac{x(\tau, \omega)}{l(\tau, \omega)} | \varepsilon_\omega^s(\tau, \omega) \right] + cte \quad (20)$$

where cte is a constant and where \mathcal{B} is another constant given by

$$\mathcal{B} = -\frac{1}{1 - \bar{s}} \frac{1}{\mathcal{L}} \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right)^{-1} \mathbb{E} [\sigma_\omega - 1], \quad (21)$$

where \bar{s} is the mean subsidy rate across firms, \overline{LS} is the average labor share, and \overline{LSO} and \overline{LSN} are the averages of (4) and (15), respectively.

The constant terms \bar{s} , \overline{LS} , \overline{LSN} and \overline{LSO} are the points around which first-order approximations are taken. We have some flexibility in choosing their values, but to minimize the approximation error they should be as close as possible to the actual observations, and the mean values are a natural choice.

The following proposition puts the two lemmas together to form the basis for our IV estimation in the next section.

Proposition 1. *The coefficients \mathcal{B} and \mathcal{L} can be consistently estimated from a two-stage least squares regression of the form:*

$$LS(\tau, \omega) - LS(\tau - 1, \omega) \approx cte + \mathcal{B} \frac{x(\tau, \omega)}{l(\tau, \omega)} + FE_i + FE_c + e(\omega), \quad (22)$$

$$\frac{x(\tau, \omega)}{l(\tau, \omega)} \approx cte + \mathcal{L}s_\omega + FE_i + FE_c + u(\omega), \quad (23)$$

where e, u are error terms, \mathcal{B} is given by (21) and FE_i, FE_c are a set of industry and city fixed effects, respectively.

This proposition shows that we can consistently estimate the causal impact of a change in automation investment on the labor share using the subsidy s_ω as an instrument. Furthermore, it implies that we can back out the average elasticity estimate from the estimated regression coefficients using the formula

$$\bar{\sigma} := \mathbb{E}[\sigma_\omega] = \mathcal{B}\mathcal{L}(1 - \bar{s}) \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right), \quad (24)$$

where the coefficients \mathcal{B} and \mathcal{L} come from estimating the system (22)–(23). The calculation requires values for \bar{s} , \overline{LS} , \overline{LSN} and \overline{LSO} . We will describe in the next section how we come up with these numbers.

Measurement error

Before moving to the estimation results, we briefly discuss how robust our identification strategy is to measurement errors in the subsidy rate and the investment intensity (x/l). This discussion will be important in the implementation of our approach.

First, measurement errors in the second stage explanatory variable x/l can be present without affecting our estimate of $\mathbb{E}[\sigma_\omega]$, even if they are of a multiplicative form rather than classical (orthogonal additive noise). In other words, measuring $\tilde{c} \times (x/l)$ instead of the true investment intensity

x/l is sufficient for identification, as long as \tilde{c} is an orthogonal random variable with a positive mean that satisfies Assumption 2. The reason for this is that the identification of the average elasticity $\mathbb{E}[\sigma_\omega]$ involves the product $\mathcal{B}\mathcal{L}$ and, by (19) and (20), the same measurement error appears on the left- and right-hand side of the two equations for a null effect on the product itself. In particular, we can readily see that estimating the first stage with the mismeasured investment rate $\tilde{c} \times (x/l)$ will yield the coefficient $\mathcal{L}' = \mathcal{L}/\mathbb{E}[\tilde{c}]$ while the second stage will yield coefficient $\mathcal{B}' = \mathcal{B}\mathbb{E}[\tilde{c}]$, implying $\mathcal{L}'\mathcal{B}' = \mathcal{L}\mathcal{B}$. Adding classical measurement error to this relation, that is, $\tilde{c} \times (x/l) + \tilde{d}$, would not change this conclusion because the IV removes this kind of measurement error.²⁶

Second, measurement errors in the subsidy rates s_ω do not invalidate our identification of the second stage coefficient but may invalidate the estimate of the elasticity, to the extent that measurement errors bias the first-stage coefficient \mathcal{L} without an offsetting effect on the second-stage coefficient \mathcal{B} . For example, suppose that the econometrician observes $s' = \tilde{b} \times s$, where s is the true subsidy rate that enters the decision process of the firm and where \tilde{b} is a constant or some i.i.d. random variable. Then, since the left-hand side of (23) is the outcome variable, the first-stage coefficient \mathcal{L}' will be equal $\mathcal{L}' = \mathbb{E}[\mathcal{L}\tilde{b}]$. However, the predicted values of x/l will be unchanged and so the second stage regression of 2SLS will not change, implying a biased point-estimate of $\mathbb{E}[\sigma_\omega]$ by (24). Classical measurement errors (orthogonal additive noise) are of a lesser concern here. Such an error will result in an attenuation bias in the first stage coefficient, but the same attenuation bias will tend to have an offsetting effect on the second stage coefficient by reducing the variation of the predicted value of the explanatory variable in the second stage (x/l).

2.4 Estimation results

In this section, we employ Proposition 1 to estimate the causal effect of automation investment using two-stage least squares regressions. The dependent variable is the change $\Delta LS(\omega)$ in the labor share of a firm ω between 2015 and 2017, and we use as instrument subsidies for automation which, for robustness and to increase statistical power, we construct in several different ways.

²⁶If \tilde{c} is a random variable that is correlated with x/l , however, we would only recover the coefficient in the second stage as long as the error term of the linear projection of $\tilde{c} \times (x/l)$ onto x/l has desirable properties, otherwise our estimation may be biased simply because the linearity assumption is not valid.

Specification using firm-level subsidies

We begin with the most basic specification that uses the average subsidy rate calculated for each firm individually as an instrument for automation investment per production worker, x/l . We construct this subsidy rate by averaging the subsidy rates reported by each firm for the years 2015, 2016 and 2017. The subsidy rate for a firm in a given year is calculated by dividing the subsidy payments self-reported by the firm for purchases of automation capital by its total automation investment in that year.²⁷

As for the variable we instrument for, we consider two specifications. The first adheres closely to the theory and uses time average of investment in automation capital over the entire sample period for each firm individually, deflated by the CPI and depreciation, and divides it by the corresponding average employment in production between the years 2015 and 2017. Formally, let δ_m be the depreciation rate of automation capital and let π_t be the CPI inflation in each year.²⁸ Let then the cumulative investment in 2015 renminbi (RMB) between years \underline{T} and \bar{T} be

$$x_{\underline{T}-\bar{T}}(\omega) := \frac{\sum_{t=\underline{T}, \dots, \bar{T}} x(\omega, t) \left(\frac{1-\delta_m}{1+\pi_t} \right)^{\bar{T}-t}}{1 + \bar{T} - \underline{T}},$$

where x_t is the surveyed firm's nominal spending on automation purchases in renminbi (RMB). Our first specification uses $x_{15-17}(\omega) / l_{15-17}(\omega)$ as instrument, where l_{15-17} is the average in production labor over those years.

Our second specification relies on the ratio $x_{17}(\omega) / l_{17}(\omega)$ instead, that effectively treats 2015 and 2016 as predetermined years. The reasoning here is that since MIC was announced in mid 2015 and effectively started to come online throughout 2016, investment rates prior to 2017 might not reflect the impact of the subsidies and, as a result, add noise to our estimation of the first stage at the expense of introducing an error to the second stage by not adhering to the theory's implied relation.²⁹ Depending on how important these errors are this may increase statistical power. As

²⁷Alternatively, one can take the sum of subsidies and investment for all years and then take the ratio. This specification leads to qualitative similar coefficients but to smaller standard errors. See the technical appendix for these results.

²⁸We set $\delta_m = 10\%$ following Table C.1-5 "List of depreciation rates under the new asset code classification — Industrial machinery" from the Bureau of Economic Analysis. We use the Consumer Price Index (All Items for China, Index 2015=100, Annual, Not Seasonally Adjusted Values): CPI(2015) = 1, CPI(2016) = 1.02, CPI(2017) = 1.03625. Source: Economic Research Division Federal Reserve Bank of St. Louis.

²⁹Investment decisions are planned ahead of time, and while some firms might have received a subsidy in 2016, it need not imply that the subsidy drove them to do so. Decisions in 2015 are unlikely to have been influenced at all by MIC and if anything firms might have withheld investment. If that is the case, the first stage under the second

discussed, the IV deals quite well with measurement errors implied by not exactly adhering to the theory-implied law of motion.

The results of the 2SLS estimation for both specifications are shown in Table 4. We focus on the intensive margin of automation investment and therefore restrict the sample to firms with positive investment. In line with the model and Assumption (2), we include city and industry fixed effects, and so our identification relies on the variation in subsidies that are not captured by these fixed effects. We also include value added as a regressor to control for firm size.

The results show that the first specification suffers from low statistical power. This may indicate that the early years in our sample firms were not affected by the subsidies yet. The coefficients are significant but only at the ten percent level. The F statistic is quite low and indicates a weak instrument. In contrast, the specification that drops the years 2015 and 2016 from the investment variable fares significantly better in terms of coefficient estimates, although the instrument is still somewhat weak. We will consider an alternative specification with stronger instruments in the next section to address this shortcoming and other concerns.

The results of Table 4 are robust to various changes in the specifications. In the Technical Appendix we show that dropping the size of the firm as a control has no meaningful impact on the results. The point estimates are stable with and without fixed effects and across the two specifications. Finally, we check for heteroscedasticity in the error terms using the Pagan-Hall test. We find p -values close to one so that we cannot reject the null hypothesis that the error terms are homoscedastic.

In all the specifications included in Table 4, the signs of the coefficients are consistent with an elasticity of substitution between labor and automation capital that is larger than 1. Our estimation therefore suggest that higher automation investment has a negative and sizable effect on the labor share. We can see this through a back-of-the-envelope calculation. Using the most conservative coefficients in Table 4, the estimates imply that a one percent subsidy, or decline in the

specification will have more power, while the second stage will involve mismeasurement because it no longer adheres to the laws of motion implied by theory. For this measurement error to not become a problem it must be the case that the linear projection of one variable onto the other,

$$\frac{x_{17}}{l_{17}} = a_{\omega} \left(\frac{x_{16-17}}{l_{15-17}} \right) + a_i + a_c + e_{\omega},$$

satisfies Assumption 2 in the sense that the error term e and the coefficient of the projection (e_{ω}, a_{ω}) are orthogonal to the subsidy residual. This follows from the second part of the proof of Proposition 1 (version of Frisch–Waugh–Lovell theorem), the discussion of measurement error in that section, and the fact that IV estimation removes the attenuation bias normally implied by the presence of classical measurement error e . We omit the formal proof.

	(1)	(2)	(3)	(4)	(5)	(6)
Second stage dependent variable $\Delta LS_\omega = LS_\omega(2017) - LS_\omega(2015)$						
Automation investment (x_{15-17}/l_{15-17})	-0.0659*	-0.0725*	-0.070**			
	(0.034)	(0.041)	(0.045)			
Automation investment (x_{17}/l_{17})				-0.053**	-0.054**	-0.043**
				(0.023)	(0.029)	(0.020)
First stage dependent variable Automation investment (x/l)						
Subsidy rate (s)	4.867*	4.455*	3.853*	6.022***	6.372**	6.086***
	(2.6)	(2.36)	(2.36)	(2.15)	(2.65)	(2.76)
Industry / city fixed effects	No/Yes	Yes/No	Yes/Yes	No/Yes	Yes/No	Yes/Yes
Size (log VA)	Yes	Yes	Yes	Yes	Yes	Yes
# observations	143	143	143	143	143	143
F -statistic	3.5	3.5	2.9	7.9	5.8	4.9
$\bar{\sigma}$ ($\overline{LSO} = 0.6$)	4.27	4.30	3.75	4.25	4.51	3.67
$\bar{\sigma}$ ($\overline{LSO} = .66$)	3.43	3.44	3.04	3.41	3.61	2.98

Table 4: Impact of automation on labor share using firm-level subsidies as an instrument.

Notes: Two-stage least squares estimation (2SLS). *** indicates significance at 1 %, ** at 5% and * at 10% level of confidence. Robust standard errors in parenthesis. Sample is restricted to observations with positive automation investment. The instrument is the subsidy rate on firm-level calculated by taking the ratio of subsidy received for purchases of automation capital (as defined in text) and total investment during the sample period.

price of automation equipment, leads to a decline in the labor share of $0.043 \times 6.086 \times 1\% = 0.26\%$ —a large number. We can put this number in perspective through the following back-of-the-envelope calculation. Suppose that the price of automation capital declines by 80%, which is well within the range of industry quality-adjusted estimates since the 1990s (see discussion in footnote (6)). Then, our results would imply a decline in the labor share across automating firms in manufacturing by 20.8 percentage points. To give an order of magnitude, the aggregate labor share declined by about 12 percentage points from 1990 in the United States.

That being said, the above back-of-the-envelope calculation needs to be taken with a grain of salt. First, the coefficients of Table 4 pertain to a locally identified partial equilibrium effect that does not take into account offsetting forces, in particular the adjustment of prices and wages, that may

be dampening the overall effect. Second, our estimation only covers manufacturing firms that (actively) invest in automation. Since these firms have witnessed a larger decline in their labor share, our estimates would predict a much smaller decline when considering all manufacturing firms. Third, we used linear approximations to derive our structural equations and these approximation might lose their validity for large changes occurring over decades. We will come back to the aggregate impact of automation prices when we explore the calibrated version of the model in the next section.

Structural interpretation of the estimated coefficients

We now combine the estimates from Table 4 with equation (21) to obtain a structural estimate of the average elasticity of substitution between labor and automation capital, given by the equation

$$\bar{\sigma} := \mathbb{E}[\sigma_\omega] = -\mathcal{B}\mathcal{L}(1 - \bar{s}) \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right) + 1, \quad (25)$$

where, recall, \mathcal{B} is the second-stage coefficient on automation investment, and \mathcal{L} is the first-stage coefficient on the subsidy rate. Since our identification comes from automating firms, we approximate the expressions around $\bar{s} = 0.12$, which is the average subsidy across firms receiving a subsidy, and $\overline{LS} = 0.485$, which is the average labor share across automating firms in 2017. As for the labor share of nonproduction workers, we use .33, which is the value reported in Table 2 for automating firms.

Computing the average elasticity $\bar{\sigma}$ from (25) requires a value for \overline{LSO} which, recall, would be the average labor share if the firms were not using any automation capital in production ($m = 0$). Given our estimated coefficients, the expression (25) is decreasing in \overline{LSO} , such that higher values of \overline{LSO} are pushing for lower elasticities of substitution. Since \overline{LSO} is a counterfactual quantity, we have no hope of directly observing it in the data. In the next section, we use the calibrated model to find a consistent value with the data, and find that $\overline{LSO} = 60\%$ provides the best fit. Here, we also report estimates of $\mathbb{E}[\sigma_\omega]$ using the textbook value of the labor share of $\overline{LSO} = 66\%$. We view this number as quite conservative for two reasons. First, in the presence of profits, the value of \overline{LSO} is bounded from above not by one but by a lower number that accommodates markups. Second, the manufacturing labor share in China was 57% in 1997—when very few, if any, modern industrial robots were at work.³⁰ On the other hand, the US labor share back in the 1960 was closer to the textbook value. To the extent that US history is a better indicator, the value of $\overline{LSO} = 66\%$ is

³⁰Data from the China Statistical Yearbook, 1997.

nonetheless appropriate.³¹

We report the estimates of $\mathbb{E}[\sigma]$ in the last rows of Table 4. For the textbook value of the labor share, $\overline{LSO} = 66\%$, the average elasticity ranges from 3.0 to 3.6. For our baseline value of $\overline{LSO} = 60\%$ the analogous range is instead from 3.7 to 4.5. Importantly, all values are above one, confirming that a decline in the price of automation equipment has a negative effect on the labor share. They are also relatively large, but in drawing such a conclusion one should keep in mind the magnitude of the decline in the US manufacturing labor share.³² We return to the discussion of the estimated elasticity of substitution in the quantitative section.

Regression setup using city-industry average subsidies

Investment decisions are affected by the subsidy rate that the firm *expects* to receive for these investments and not the subsidy that the firm actually received. We do not observe these expectations and proxy for them using information available in the data. In the previous section, we relied on the observed subsidy payments received by the firms to compute the subsidy rate and assumed that what firms expected to receive is equal to what the firm actually received. But if firms which received zero subsidies ex post did not expect such an outcome ex ante, or firms which received full subsidy considered the risk of receiving zero subsidy ex ante, the measurement of subsidy rate could induce a bias that is multiplicative, and as explained in the previous section, such a bias could adversely affect the estimated elasticity.³³ For this reason, here we consider an alternative instrument computed by averaging the subsidy rates reported by all the firms in a given industry-city grouping, and then applying it to all firms for that grouping. This partially addresses the problem since the average realization of any random variable should bring us closer to its ex ante expectation, and hence if the coefficient do not change, this is an indication that our estimates are immune to this problem. An additional advantage of this setup is that it increases the statistical power since more firms are “effectively” subsidized.

³¹ Kehrig and Vincent (2018) report that the labor share in U.S. manufacturing declined from 61 percent in 1967 to 41 percent in 2012, which would suggest that historic evidence from the U.S. points to a higher value. By this measure, then, the textbook number of $\overline{LSO} = 66\%$ would be a conservative upper bound.

³²Since the 1990s the US labor share index in manufacturing declined by 19% according to U.S. Bureau of Labor Statistics “Productivity and Costs” (1990-2018). See also discussion in footnote 31.

³³It is possible that firms are simply misreporting the subsidy payments that they receive for their investments. It is also possible that, when making investment decisions, firms were facing some uncertainty about the size of the subsidy payments that they would receive. In either case, averaging the firm-level rates across all firms in a city-industry pair might get us closer to the true policy subsidy that was expected by the firms. Finally, it is possible that some firms might have “lobbied” for more generous subsidies for themselves, perhaps introducing an undesirable correlation between subsidies and other firm characteristics. This second instrument would also alleviate these concerns.

Formally, for any given city-industry pair (c, i) and year $t \in \{2015, 2016, 2017\}$, we construct the instrument for a firm ω as follows:

$$s_{\omega}^{\Omega_{ci}} = \frac{\sum_{t=2015}^{2017} \sum_{\omega \in \Omega'_{ci}} S_{ci}(t, \omega)}{\sum_{t=2015}^{2017} \sum_{\omega \in \Omega'_{ci}} x_{ci}(t, \omega)},$$

where $S_{ci}(t, \omega)$ is the subsidy transfer (in RMB) deflated by the CPI inflation and reported by firm ω from city c and industry i in period t , $x_{ci}(t, \omega)$ is automation investment similarly deflated by CPI; the set $\Omega'_{ci} \subseteq \Omega_{ci}$ includes all automating firms from city-industry pair (c, i) . To back out the average elasticity in this case, we use $\bar{s} = 0.109$ and otherwise the analysis is unchanged.

The regression results are presented in Table 5. The estimated coefficients are similar across different variants in terms of their product and also similar to those reported in Table 4. The standard errors are however smaller with several coefficient significantly different than zero at the one percent threshold. Statistical power also increases, and the F-statistics are indicative of a strong instrument for two specifications, including for our preferred specification that uses city and industry fixed effects in column 6. The estimated elasticities $\mathbb{E}[\sigma_{\omega}]$ are similar to those of Table 4.³⁴

Differential impact of ordinary capital

For completeness, we ran an analogous exercise for ordinary capital (Machine-3), for which we also have capital-type specific subsidy information. We found coefficients to be insignificant.³⁵ Together with our previous OLS results, this indicates an important difference between ordinary and automation capital: while we find that automation capital has a clear negative impact on the labor share, there is no evidence that ordinary capital has a similar effect.³⁶

2.5 Limitations and robustness

We conclude this section by discussing further limitations of our instrument and discuss how we addressed them.

First, it is possible that subsidies during the sample period may be correlated with some other subsidies that were previously in place. Our theory, recall, assumes that preexisting subsidies are uncorrelated with the residual of the current set of subsidies after controlling for city and industry

³⁴Controlling for size of the firms using value added does not have a meaningful impact on the results.

³⁵See Technical Appendix.

³⁶The lack of significance in this case may indicate that the instrument is not relevant and it should not necessarily be interpreted as evidence against any relationship between price of capital and the labor share.

	(1)	(2)	(3)	(4)	(5)	(6)
Second stage dependent variable $\Delta LS_\omega = LS_\omega(2017) - LS_\omega(2015)$						
Automation investment (x_{15-17}/l_{15-17})	-0.0928*	-0.0815*	-0.0734*			
	(0.051)	(0.045)	(0.045)			
Automation investment (x_{17}/l_{17})				-.0382***	-.0458*	-.0377***
				(0.014)	(0.026)	(0.012)
First stage dependent variable Automation investment (x/l)						
Subsidy rate ($s_\omega^{\Omega_{ci}}$)	3.584	4.319**	3.838*	7.772***	6.401**	7.092***
	(2.26)	(2.13)	(2.36)	(2.15)	(2.65)	(2.76)
Industry / city fixed effects	No/Yes	Yes/No	Yes/Yes	No/Yes	Yes/No	Yes/Yes
Size (log VA)	Yes	Yes	Yes	Yes	Yes	Yes
# observations	143	143	143	143	143	143
F-statistic	3.5	3.5	2.9	20.8	5.6	14.9
$\bar{\sigma}$ ($\overline{LSO} = 0.6$)	4.44	4.64	3.91	4.06	4.03	3.76
$\bar{\sigma}$ ($\overline{LSO} = .66$)	3.55	3.70	3.16	3.28	3.25	3.05

Table 5: Impact of automation on labor shares using city-industry average subsidies as an instrument.

Notes: Two-stage least squares estimation (2SLS). *** indicates significance at 1 %, ** at 5% and * at 10% level of confidence. Robust standard errors in parenthesis. Sample is restricted to observations with positive automation investment. The instrument is the subsidy rate on firm-level calculated by taking the ratio of subsidy received for purchases of automation capital (as defined in text) and total investment during the sample period.

fixed effect. While we are unaware of any previous major national initiative that specifically targeted automation investment on that scale, other economic policies were in place in China before MIC 2025 and it is possible that some small subsidies for automation were also in place. We investigate how this would affect our results in Appendix 8 and find that, in the likely case of a positive correlation between the residuals of the pre-existing subsidies and subsidies under MIC 2025, our estimation of the average elasticity $\bar{\sigma}$ would be biased downward, providing a lower bound for the true elasticity. Since we find high values to begin with, this does not change the conclusions from our findings. Intuitively, a positive correlation reduces the link between current subsidies and labor share changes as the firm would have already responded to the subsidies.

Second, subsidy programs may have targeted specific firms or industries whose labor shares have

been changing for reasons unrelated to their investment in automation. In this case, our instrument would mistakenly identify a causal link between automation and the labor share where none exists. We have read through official documentation and found no evidence of this type of policy being used to target specific firms, albeit there exists a possibility that this is the case because subsidized firms tend to be larger. But as is clear from the regression, controlling for size does not make a difference. We also show in the Technical Appendix that including controls for party connections and “lobbying power” do not invalidate our result.

3 Quantitative results

In this section we calibrate our model using a simulated method of moments and conduct two exercises using the calibrated economy. First, we analyze the consistency of our estimates of the microeconomic elasticity of substitution σ with the macroeconomic changes observed in the sample. These aggregate changes involve not only the firm-level elasticity but also patterns of substitution across firms so that the full model is needed to capture the aggregate impact of a change in automation prices. Second, we validate our identification procedure on model-generated data. We find that despite the assumptions that our identification involves, our reduced-form empirical strategy is quite robust and can accurately back out the true elasticity of substitution even in the presence of large parameter heterogeneity. In what follows, all mean values calculated across firms are weighted using the value-added, unless otherwise noted.

3.1 Parameterization

The setup of the model is as described in Section 1. While our exercise focuses on automating firms, we include firms that cannot automate ($a_\omega = 1$) to study the model’s implications for the whole sample. The presence of these firms only matters for the general equilibrium feedback between automation and prices.

Sample structure

The model’s time horizon is infinite but, as before, we focus on two subsequent periods: a pre-subsidy period “ $\tau - 1$ ” and a post-subsidy period “ τ .” As before, we rely on the differences between values in year 2015 and 2017 and relate those to model implied differences between τ and $\tau - 1$. The

calibrated model is designed to approximate the city-industry structure of our dataset. As in the data, there are 44 cities, 23 industries, and 117 city-industry groupings. To avoid random variation induced by the small numbers of firms in the sample, we roughly double the number of firms in the model economy. Specifically, the model includes 351 automating firms—three per grouping—and 3192 firms in total—14 per grouping. While the additional firms are useful to remove variation coming from a particular draw of random numbers, it only has a limited impact on the properties of the model, and when we run regressions on model-generated data to test our empirical strategy we make sure to sample the same number of firms as in the data so as to back out the correct standard errors.

Setting the share of automation activities $a < 1$ is insufficient to ensure that a firm automates (invests in automation capital) in the model. For instance such a firm facing a negative demand shock might not want to invest to increase its size. To ensure a consistent number of automating and nonautomating firms, we therefore pick firms at random from a larger pool to ensure that we populate each grouping as specified (three automating firms per grouping). This selection procedure is a function of equilibrium prices and it is part of the fixed point to solve the model. By construction, then, automating firms are the ones which do not experience large negative shocks in period τ .

Technology parameters

We fix markups at 33% for each firm, implying a constant θ that solves $\frac{\theta}{\theta-1} = 1.33$. This level of markups is consistent with the aggregate operating (pre-tax) income of automating firms relative to their value added that we see in our data.³⁷

To model firm-level heterogeneity in α, γ, η , we assume that these technology parameters follow independent beta distributions with shape parameters $\zeta_1(h)$ and $\zeta_2(h)$ for $h \in \{\alpha, \gamma, \eta\}$.³⁸ To restrict the amount of parameters to calibrate, we further assume that α and γ follow the same distribution.³⁹

³⁷We calculate the Lerner index as the ratio of operating income to value added. Operating income corresponds to item D.3. in the enterprise cost and profit statement questions in the questionnaire. The questionnaire can be found in the supporting files.

³⁸We normalized the distributions to the interval $[0.05, 0.95]$ to avoid issues with parameter values that are too close to the extremes. To do so, we use the transformed random variable $Y = X(b-l) + l$, where $b = .95$ and $l = .05$, and X is a beta distributions on $[0, 1]$.

³⁹This assumption is motivated by the data. In the presence of heterogeneity in η , any large differences in the mean values of α and γ would imply a nonzero correlation between the labor share and the importance of production activities in firms' overall operations. For example, if we assumed $\alpha = 1/3$ and set γ close to zero, we would see that firms which do not have many production activities have a much lower labor share, as in their case α would be of greater importance for the division of income between labor and capital. Similarly, a strong correlation between α and γ would induce a high labor share among firms with a large share of production activities. Since we are not aware of such patterns in historic data, we choose independent distributions with equal shape parameters.

Our calibration strategy takes advantage of the fact that, in the model, the mean of various quantities is related to the average pre-automation labor share \overline{LSO} . We write these means in terms of \overline{LSO} and then pick that last quantity to best fit the data, which corresponds to the value $\overline{LSO} = .6$ used in the previous section. For instance, we know from (4) that

$$\bar{\alpha} = \bar{\gamma} = 1 - \frac{\theta}{\theta - 1} \overline{LSO} = 1 - 1.33 \overline{LSO},$$

where the overline $\bar{\cdot}$ denotes the value added-weighted average. For a given value of \overline{LSO} , we can use that equation to discipline the shape parameters of the distributions of α and γ .

We proceed similarly to discipline the distribution of η . To calculate the mean value $\bar{\eta}$, we note that automation in the model reduces the labor share from its automation-free level \overline{LSO} by lowering the labor share of production workers according to the relation $\overline{LS}_\tau - \overline{LSO} = \overline{LSP}_\tau - \overline{LSOP}$, where, as defined in (4), $\overline{LSOP} := \frac{\theta-1}{\theta} (1 - \bar{\eta})(1 - \bar{\alpha})$. Given the calculated value added-weighted data averages for $\overline{LS}_\tau = .485$ and $\overline{LSP}_\tau = .345$ across automating firms in 2017 (see Table 2), and together with (4), we obtain

$$\bar{\eta} = 1 - \frac{\overline{LSOP}}{\overline{LSO}} = 1 - \frac{.345 + \overline{LSO} - .485}{\overline{LSO}} = \frac{.155}{\overline{LSO}}.$$

Finally, we set the value of the elasticity of substitution σ as a function of \overline{LSO} using the estimated coefficients from our city-industry regression that includes all fixed effects (our preferred specification). We assume that this elasticity is identical across firms, sectors and cities. Equation (24) states the relationship between σ and \overline{LSO} .

Our estimation procedure varies, among other thing, \overline{LSO} to match a set of moments. While the parameters are jointly estimated, \overline{LSO} has a particular importance for one key targeted moment: the average growth rate of value added produced by automating firms. Below, we describe the mechanism through which \overline{LSO} affects that quantity and why that parameter is properly identified through the estimation procedure.

The choice of \overline{LSO} pins down the mean values of α , γ and η but more data is needed to pin down the variance of these parameters and hence their distribution. To do so, we choose two data targets that are particularly relevant: 1) the coefficient of variation of the share of production activities as measured by the ratio of production employment to total employment (SPE17) of .25 in 2017 (which we associate with period τ); 2) the coefficient of variation of the ratio of the labor share of

production employees to the labor share of nonproduction employees of 1.0 in 2017. Together with the mean values listed above, these targets identify $\zeta_1(h)$ and $\zeta_2(h)$ for $h \in \{\alpha, \gamma, \eta\}$ (recall that the distributions of γ and α are assumed identical).

Shocks and trend growth rates

We target moments related to the growth of value added and the initial employment distribution across firms to pin down d and A . Since d and A are equivalent in the model, we use d to model the dispersion of employment across firms and shocks across periods. We use A to model the relative size of automating firms relative to non-automating firms and the growth rate of wages (and productivity) in the economy as a whole.⁴⁰

Specifically, the initial value of $\log(d(\omega, \tau - 1)) \sim \text{iid } \mathcal{N}(0, \sigma_{d(\tau-1)})$ is log-normally distributed with a standard deviation $\sigma_{d(\tau-1)}$ that is set to match the coefficient of variation of total employment across automating firms in 2015, which we find to be close to one. Since automating firms account for 23.5 percent of value added produced by all firms in 2017, we normalize $A(\omega, \tau - 1) = 1$ for nonautomating firms and set the TFP of automation firms to a constant $A(\omega, \tau - 1) = \bar{A} > 1$ that is large enough to hit this target.

To capture shocks that affect idiosyncratic differences in firm growth rates, we assume that the demand shifter $\log(d(\omega, \tau)) \sim \text{iid } \mathcal{N}(\log(d(\omega, \tau - 1)), \sigma_{d(\tau)})$ is log-normally distributed around the previous value $d(\omega, \tau - 1)$, implying a random walk processes for the evolution of firm sizes. We choose the standard deviation $\sigma_{d(\tau)}$ that matches the standard deviation of the firm-level growth in value added between 2015 and 2017, which in our data is .55. The period τ productivity A is assumed to grow at a constant rate across all firms so that the model exactly matches the average growth rate of manufacturing real wages in China of 12.5 percent between 2015 and 2017.⁴¹

User cost of capital

We use (5) to compute the user cost of each type of capital. We assume that the real cost of funds is 6 percent on an annual basis, which is based on the estimated global weighted average cost

⁴⁰The distinction between A and d would have mattered had d be assumed to be “quality” reflected in the price of the good, which we did not assume by treating it as a “hidden” preference parameter. “Quality” shocks in our model are captured by a markup shock.

⁴¹Wage growth in manufacturing is between 2015 and 2017 for consistency with other variables that are measured this way and it comes from the National Bureau of Statistics of China. It is deflated by the all items CPI.

of capital (WACC) by Aswath Damodaran.⁴² We assume that the annual depreciation rates are 5 percent for and 10 percent for equipment and automation capital—which is within the range of estimates of the useful life of equipment.⁴³ These numbers are not particularly important for the results. We assume that they do not change over time and are common across firms and sectors. We compound annual rates to arrive at three-year equivalents for the model.

To get the user cost of automation capital from (5) also requires information about the price of capital $p_m(\tau)$ to other machinery and the subsidy rate $s(\tau, \omega)$. Since $p_m(\tau)$ in the model and the mean value of the technology parameter \bar{a} have an analogous effect, we set $\bar{a} = .5$ and assume no dispersion in this parameter for simplicity. We then set the initial price $p_m(\tau - 1)$ to account for the 2015 value of the labor share of automating firms in the data, and which is 0.5. This target implicitly determines how much automation capital ($m_{\tau-1}$) automating firms had as of 2015 to bring down the value of their labor share $LS(\omega, \tau - 1)$ from the automation-free reference value $LSO(\omega)$.

Finally, the subsidy rate $s(\tau, \omega)$ is assumed to vary at the city-industry level with a positive subsidy received by only those city-industry groupings that receive positive subsidies in the data. The subsidy rate is drawn from the beta distribution with a mean across subsidized firms within subsidized groupings of 10.9% and the coefficient of variation of 44%.⁴⁴

Computing the user cost of capital using equation (5) requires information on $p_m(\tau)$. Here, we target the ratio $\mathbb{E}_w m(\omega, \tau) / \mathbb{E}_w m(\omega, \tau - 1)$, where \mathbb{E}_w denote the value added-weighted average, and require it to be consistent with the average industrial robot density in Chinese manufacturing industry between 2015 and 2017 which, according to the International Federation of Robotics (IFR), increased by 98 percent between 2015 and 2017.^{45,46} This moment pins down the rate of decline in the price of automation capital, which in the baseline parameterization is $\Delta_{\%} p_m := (p_m(\tau) -$

⁴²We use a rounded average emerging market dollar estimate of WACC by Aswath Damodaran (New York University Stern) for 2015 after subtracting US inflation of 2 percent ($8.96 - 2 = 5.96$). The dataset for 2015 can found at <http://people.stern.nyu.edu/adamodar/...>

⁴³The useful life of industrial robots and digitally controlled machinery is around 10 years. See, for example, <https://www150.statcan.gc.ca/...> The International Federation of Robotics uses a slightly higher value of 12 years. As for other categories, we take the values from <https://assets.ey.com/...> (page 37), which report the useful life for plant, property and equipment in China of 10 years (k_e) and 20 years for buildings (k_s) (for tax purposes).

⁴⁴We limit the domain of subsidy rates so that it does not exceed 50 percent and so the beta distribution is defined on the interval from 0 to .5.

⁴⁵Robot density in Chinese manufacturing reported by the IFR was: 25 units per 10,000 workers in 2013, 36 units in 2014, 49 units in 2015, 68 units in 2016, and 97 in 2017. A similar rate of growth applies to earlier periods. See 2015, 2016, 2017 and 2018 World Robotics Survey, International Federation of Robotics.

⁴⁶Without knowing the productivity of automation capital, it is not possible to take advantage of data on investment in automation to pin down $p_m(\tau)$. Note that assuming a higher initial value of $p_m(\tau - 1)$ and a lower value of \bar{a} yields an equivalent model, yet the measurements of investment in automation are different because the assumed price of a unit of capital is different.

$p_m(\tau - 1)/p_m(\tau - 1) \times 100\% = -3.9\%$. While this may seem like a small number, the price of automation equipment is determined in world markets and the Chinese exchange rate depreciated during this period, implying that in local currency foreign goods became more expensive in China. The evidence on the transacted price of industrial robots suggests that the price of industrial robots was flat during this time period in dollar terms.⁴⁷ The real decline of 3.9 percent thus only captures quality improvements that might have positively affected the productivity of automation capital.

We calibrate the model using a simulated method of moment that jointly targets all the moments listed in the previous paragraphs. Table 6 summarizes these moments and the obtained parameter values, together with an approximate mapping between parameters and the moments that each parameter influences the most at the calibrated value.

Moment	Data	Model	Parameter	Value
w.m. $LSP(\omega, \tau)$.35	.35	$\zeta_1(\eta)$	1.10
			$\zeta_1(\alpha) = \zeta_1(\gamma)$	0.70
c.v. $l(\omega, \tau) / L(\omega, \tau)$.25	.24	$\zeta_2(\eta)$	3.61
c.v. $LSN(\omega, \tau) / LSP(\omega, \tau)$	1.0	1.0	$\zeta_2(\alpha) = \zeta_2(\gamma)$	2.77
c.v. $L(\omega, \tau)$	1.0	1.0	$\sigma_{d(\tau-1)}$	1.20
markup	33%	33%	$\bar{\theta}$	3.95
$\sum_{\omega:x>0} y(\omega, \tau) / \sum_{\omega} y(\omega, \tau)$	23.5%	23.5%	\bar{A}	1.179
$\sum_{\omega:x>0} y(\omega, \tau) / \sum_{\omega:x>0} y(\omega, \tau - 1)$	25%	25%	\overline{LSO}	0.60
w.m. $LS(\omega, \tau - 1)$	0.5	0.5	$p_m(\tau)$	5.24
$\Delta\% \sum_{\omega} m(\omega, \tau)$	98%	98%	$p_m(\tau) / p_m(\tau - 1)$	0.961
$w_{\tau} / w_{\tau-1}$	12.5%	12.5%	$A_{\tau} / A_{\tau-1}$	1.066
s.d. $\Delta\% y(\omega, \tau)$.55	.55	$\sigma_{d(\tau)}$	0.48

Table 6: Targeted moments and model Fit

Notes: The table lists the key moment conditions we use to parameterize the model. $\Delta\%$ denotes percentage growth rate. "w.m." denotes value-added weighted sample mean. The weights pertain to an average value added of a firm over the time period 2015-2017. "s.d." denotes standard deviation and "c.v." denotes coefficient of variation. Standard deviations and coefficients of variation are not weighted. Wage growth in manufacturing is between 2015 and 2017 and it comes from the National Bureau of Statistics of China. We deflate it by all items CPI as referenced in text. The calibration of the parameters is joint and the mapping of parameters to moments is approximate.

3.2 Mechanism and model-based identification of \overline{LSO}

In our reduced-form regressions of Section 2.4, we needed to pick a value for \overline{LSO} to back out an estimate of the average elasticity of substitution $\bar{\sigma}$ from the regression coefficients. We looked at how various plausible values of \overline{LSO} affected $\bar{\sigma}$. Among them was $\overline{LSO} = 0.60$, the number that is selected by the calibration as providing the best fit to the data and that is visible in Table 6.

⁴⁷See <https://www.statista.com/...>

Here, we describe what features of the data push the calibration toward that number. Since the intuition does not rely on variations across firms, we simplify the exposition and consider a version of the model without heterogeneity and zero subsidies. We also focus on two parameterizations: the baseline parameterization that emerges from the calibration and a “lower σ ” parameterization that imposes $\overline{LSO} = .66$ (the textbook labor share value) , resulting in $\sigma = 3.05$.⁴⁸ This elasticity is calculated from (25) by using that value of \overline{LSO} together with regression coefficients from the last column of Table 5.

Table 7 shows the difference across the two parameterizations in terms of the general equilibrium drivers of automation: i) the trend increase in total factor productivity $\Delta\%A := (A(\tau) - A(\tau - 1))/A(\tau) \times 100\%$, ii) the equilibrium increase in the real wage rate $\Delta\%w$ and iii) the decline in the price of automation capital $\Delta\%p_m$. These values, recall, have been calibrated to match the assumed wage growth of 12.5 percent and ensure that the stock of automation capital doubles in period τ , that is $(\mathbb{E}_w m(\omega, \tau)/\mathbb{E}_w m(\omega, \tau - 1) \approx 2)$. Note that technological progress in the model is not labor augmenting so that, since the amount labor is fixed, the increase in productivity spurs automation by increasing wages by more than the rise in total factor productivity. We see that the main difference is in terms of $\Delta\%p_m$. Higher value of \overline{LSO} (and lower elasticity σ) requires a steeper drop in automation prices to explain the observed increase in automation capital.

Parameterization	$\Delta\%A$	$\Delta\%w$	$\Delta\%p_m$
Baseline ($\overline{LSO} = .6, \sigma = 3.76$)	8.9%	12.5%	-8.6%
“Lower σ ” ($\overline{LSO} = .66, \sigma = 3.05$)	9.4%	12.5%	-10.8%

Table 7: No Heterogeneity Parameterization: Aggregate Trends.

Parameterization	$LS_{\tau-1}$	ΔLS_{τ}	$\Delta\%L_{\tau}$	$y_{\tau-1}$	$\Delta\%y_{\tau}$	$m_{\tau-1}$	$\Delta\%m_{\tau}$
<i>A. Baseline ($\overline{LSO} = .6, \sigma = 3.76$)</i>							
Automating firm	.50	-.04	1%	1.33	25%	.19	98%
Nonautomating firm	.60	0.	0%	.62	13%	0.	-
<i>B. “Lower σ” ($\overline{LSO} = .66, \sigma = 3.05$)</i>							
Automating firm	.50	-.05	5%	1.28	32%	.28	97%
Nonautomating firm	.61	0.	0%	.63	.28	0.	-

Table 8: No Heterogeneity Parameterization: Firm Dynamics

Table 8 compares the two parameterizations along various outcomes for automating and nonau-

⁴⁸We recalibrate the full model without \overline{LSO} , and its associated moment (growth rate of value added) as a parameter. We also drop this moment to better show what goes wrong when we deviate from $\overline{LSO} = .6$.

tomating firms. Recall that since there is no heterogeneity in this case, there are only two types of firms: automating firm and nonautomating firms and all city-industry groupings are identical.

The table reveals two key differences between the parameterizations and which inform us about the mechanisms that are behind the identification of \overline{LSO} —as indicated in Table 6. First, the overall impact on the labor share of the automating firms, ΔLS_τ , is smaller in the baseline parameterization despite the micro-elasticity σ being higher in that case. To understand this somewhat counterintuitive outcome, note that the “lower σ ” parameterization assumes $\overline{LSO} = .66$, so as to get a lower σ from the same regression coefficient estimates. This implies that, for the model to be consistent with an initial level of the labor share of $LS_{\tau-1} = .5$, the initial stock of automation capital ($m_{\tau-1} = .28$) must be larger in the “lower” σ parameterization to reduce the labor share from a higher automation-free value of $\overline{LSO} = .66$ and despite a lower sensitivity of the labor share to automation capital due to the lower σ . Now, since our calibration targets that the stock of automation capital doubles between 2015 and 2017, the larger initial value of $m_{\tau-1}$ implies a larger increase in the stock of automation capital in absolute terms in the “lower σ ” case. As a result, the labor share declines more but also the value added of firms rises much more.⁴⁹

As a result, the “lower σ ” parameterization implies that the model grossly overshoots the targeted value of 25 percent increase in total value added produced by automating firms $\Delta\%y_\tau$ —which in the baseline case is targeted to obtain the value of $\overline{LSO} = .6$ and in the “lower σ ” parameterization is dropped from calibration targets to accommodate the higher assumed value of $\overline{LSO} = .66$ (and lower σ). To summarize, the lower micro-elasticity σ leads to the inference that firms must have been much more automated back in 2015, which in combination with the fact that the stock of automation capital doubled between 2015 and 2017 implies a bigger increase in the value added produced by automating firms and a larger decline of labor share in these firms.

It is important to note that the data strongly rejects values that would be anywhere near the unit elasticity that leads to labor shares that are invariant to the price of automation capital. Such parameter values would imply an enormous growth of value added produced by automating firms. Note also that \overline{LSO} cannot be higher than .7 when average markups are at 33 percent and so the .66 value that we use in the previous section is close to the upper-bound of what is possible given the assumed level of markups.

⁴⁹Note that the difference in elasticity is only reflected here in the assumed change in the price of automation capital. In the “lower σ ” parameterization, it is calibrated to decline by $\Delta\%p_m = -10.8\%$ while in the baseline case by $\Delta\%p_m = -8.6\%$.

3.3 Results

We are now ready to discuss our model’s aggregate implications and to analyze the robustness of our procedure to estimate the average elasticity σ using model-generated data.

Equilibrium effect of automation

Our reduced-form exercises focused on identifying the microeconomic elasticity of substitution σ . That number predicts how changes in economic conditions affect the labor share at the firm-level, but movements in the aggregate labor share depend also on substitution patterns between firms. Here, we take advantage of our calibrated model to study these patterns and to figure out if our estimated micro elasticity is consistent with the aggregated changes that we observe in the sample.

Table (10) compares the aggregate predictions of the baseline calibration for each type of firms. As we can see, the model does a reasonable good job at matching the labor share of automating and subsidize firms in period $\tau - 1$, while it somewhat missed the labor share of nonautomating firms. Overall, we see that the model is able to roughly replicate the change in the aggregate labor share between τ and $\tau - 1$ —a key quantity of interest—although its performance for automating and subsidized firms is better when looking at the unweighted average. When we take all firms together, the unweighted mean performs slightly worse. What this implies is that larger firms in the data are the ones in which the labor share is changing more sluggishly than in the model. The last row of Table 9 reports aggregate automation investment in the model and in the data. As we can see, the model performs well in this dimension, as it roughly matches the difference between automating and subsidized firms.

One key lesson from Table (10) is that even with a relatively large micro-elasticity of substitution ($\sigma = 3.76$) and large movements in automation prices and productivity, the overall impact on the labor share is muted. While the labor share of automating firms does decline sizably (by .041), these firms account for only 23.5 percent of value added so that their impact on the aggregate labor share’s decline is close to $.235 \times .041 = 0.0096$. We interpret these results as showing that the seemingly large micro-elasticity of σ is roughly consistent with the aggregate patterns in our sample and does not indicate excessive sensitivity vis-a-vis the data.

Statistic		Firm type					
		Automating		Subsidized		All	
		Data	Model	Data	Model	Data	Model
Agg. labor share in $\tau - 1$	$LS(\omega, \tau - 1)$ w.m.	.500	.500	.502	.505	.496	.573
Agg. labor share in τ	$LS(\omega, \tau)$ w.m.	.484	.460	.483	.435	.473	.563
Change in labor share	ΔLS_ω w.m.	-.022	-.041	-.035	-.070	-.026	-.010
	ΔLS_ω m.	-.036	-.039	-.083	-.065	-.024	-0.004
Agg. auto. investment ¹	$\sum_\omega p_m x_\omega / \sum_\omega y_\omega$	15.2%	15.2%	22.6%	24.6%	3.6%	3.6%

Table 9: Quantitative Results from the Calibrated Model

Notes: The table compares the baseline calibrated model to the data. $\Delta\%$ denotes percentage growth rate. "w.m." denotes value-added weighted sample mean. The weights pertain to an average value added of a firm over the time period 2015-2017. "m." denotes (unweighted) mean. "s.d." denotes standard deviation and "c.v." denotes coefficient of variation (unweighted across firms). Standard deviations or coefficients of variation are not weighted. ¹ As explained, there is an indeterminacy between the price of automation capital and productivity A . We pick the units of automation capital so that the model matches perfectly the aggregate investment by automating firms.

Validation of elasticity estimation procedure

Here we use the calibrated model to evaluate the performance of the empirical strategy that we used in Section 2. That strategy was grounded in Proposition 1, which formally established its validity under two assumptions. First, it relies on linear approximations of the firms' policy functions. Second, it assumes that the subsidy residual $\varepsilon^s(\tau, \omega)$ from (17) is uncorrelated with various quantities, among them prices such as the wage rate and the the price indices of goods at the city and industry-level. In the calibrated model, these assumptions may only hold approximately. For instance, the linear approximations are only valid locally and large subsidies might make higher-order terms relevant. In addition, some prices might be affected by the subsidies, even after controlling for city and industry fixed effects. To test whether these deviations threaten our estimation procedure, we test their importance one by one. To simplify the exposition of the main forces at work, we first consider a version of the calibrated model without parameter heterogeneity. We reintroduce that heterogeneity later on.

We begin by examining the importance of the linear approximation by considering whether the first and second stage regressions involve data that are related in a linear relationship. That information is provided by the first two panels of 1. As we can see from the first panel, the first-stage relationship is roughly linear although there is a slight nonlinearity that might introduce a bias in the estimation results when the variability of the subsidies across firms is substantial. The second panel illustrates the second stage of our instrumental variable estimation. Here, the model suggests that the approximation is more accurate, despite the nonlinearity in the first stage, with the predicted

values from the first-stage regression varying roughly linearly with the dependent variable. The second panel also displays the slope estimated by our IV approaches as well as the slope that would be estimated using a standard ordinary least squares estimator. In this case, without firm heterogeneity, the OLS and the 2SLS estimates are very close to each other, which is not surprising.

The third panel of Figure 1 focuses directly on the Taylor approximation that underlies the proof of Proposition 1 (and Lemma 5); that is, the replacement of

$$\log \frac{LSO(\tau, \omega) - LS(\tau, \omega)}{LS(\tau, \omega) - LSN(\tau, \omega)} - \frac{LSO(\tau - 1, \omega) - LS(\tau - 1, \omega)}{LS(\tau - 1, \omega) - LSN(\tau - 1, \omega)} \quad (26)$$

by a linear approximation that uses $LS_\tau - LS_{\tau-1}$ and a scaling factor involving mean values of the involved variables (see (36) in the Appendix). Again, we see that the relationship between these two quantities is close to linear, and most importantly the slope is approximately one (compare to the 45 degrees line indicated in the figure).

To confirm that the slight departure from linearity observed in the first-stage regression does not introduce a bias in the estimation of σ , we replicate our estimation procedure on model-generated data and compares the estimated σ to its true value. The results are presented in Table 10. As we can see, the bias is not very large with an estimated elasticity of $\sigma = 3.8$ versus a true value of 3.76.⁵⁰

We next examine the potential impact of the endogenous response of prices to changes in subsidies. The fourth row of Table 10 reports the elasticity as estimated on model-generated data when we counterfactually take the period- τ prices from an economy with zero subsidies, which is labelled “estimated σ w/ no GE feedback”. As we can see, the estimated elasticity is little changed by using counterfactual prices and so the general equilibrium feedback effect through prices is quite small.

We repeat the same exercise on data generated from the baseline model with parameter heterogeneity. The main findings, also reported in Table 10, do not change. The fit is similar, with an estimated elasticity σ of 3.72 versus a true value of 3.76. The bootstrapped standard error is informative in this case and its value of .27 implies a reasonably tight confidence interval.⁵¹

Figure 2 shows the first-stage and second-stage regressions in panels (a) and (b), and the quality of the linear approximation in panel (c). These panels illustrates the enormous level of heterogeneity generated by the calibrated model. This heterogeneity does not, however, have a strong impact on

⁵⁰For larger variations in the subsidy rates this error tends to increase. The estimated elasticity converges to the true value as we decrease the variance of the subsidy residual.

⁵¹We compute that number by averaging the standard error of the 2SLS estimator from 1000 random draws of 143 firms (out of 351), the same as in the sample.

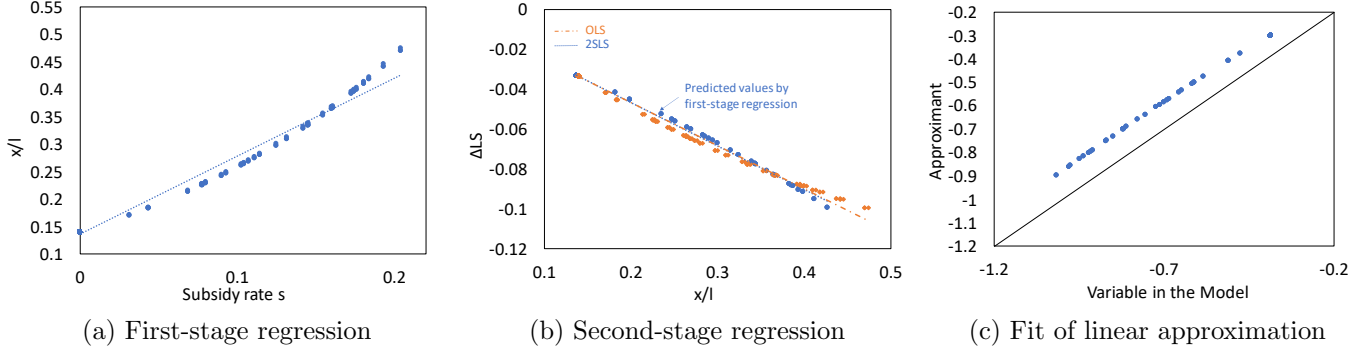


Figure 1: 2SLS Elasticity Estimates from Data Generated by Baseline Model with No Parameter Heterogeneity

Statistic	Model	
	No heterogeneity	Baseline
Assumed σ	3.76	3.76
Estimated σ	3.80	3.72
S.E.	—	(0.27)
Estimated σ w/ no GE feedback	3.80	3.74
S.E.	—	(0.24)
Observations (out of 342)	143	143

Table 10: 2SLS Elasticity Estimated from Model Data

Notes: We followed the estimation procedure implied by Proposition 1. Bootstrapped standard errors in parenthesis. To calculate them, we run regressions 1000 times by drawing a random sample of automating firms from each city-industry grouping for a total of 143 random observations across all groupings. We assumed a fixed number of draws per grouping to ensure that the fraction of subsidized firms is fixed across all samples in the bootstrap procedure.

our estimated σ , as we can see from Table 10. Although with a limited sample size heterogeneity might have an impact on the statistical power of our identification. As the table shows, identification is indeed possible in this case and the estimated (bootstrapped) error is reasonable.⁵² Overall, these figures suggest that the procedure we propose in Proposition 1 is consistent with the model, as required, and that the approximations that were used to derive our reduced-form estimation procedure do not meaningfully affect the estimated coefficients.

Finally, Figure 2 highlights the importance of using an instrument to tease out the impact of automation on the labor share in the presence of massive firm heterogeneity. We plot in panel (b) the slopes of the OLS and 2SLS estimators. While these lines overlapped when firms were homogeneous, they differ markedly in the full model and the OLS estimator underestimates the

⁵²This might be surprising in view of the somewhat poor fit of the linear approximation visible in panel (c). But these errors do not have a meaningful impact on the estimates since they are not systematically related with the firm characteristics that enter the first and the second stage regressions. Indeed notice that the best linear fit as a slope very close to one.

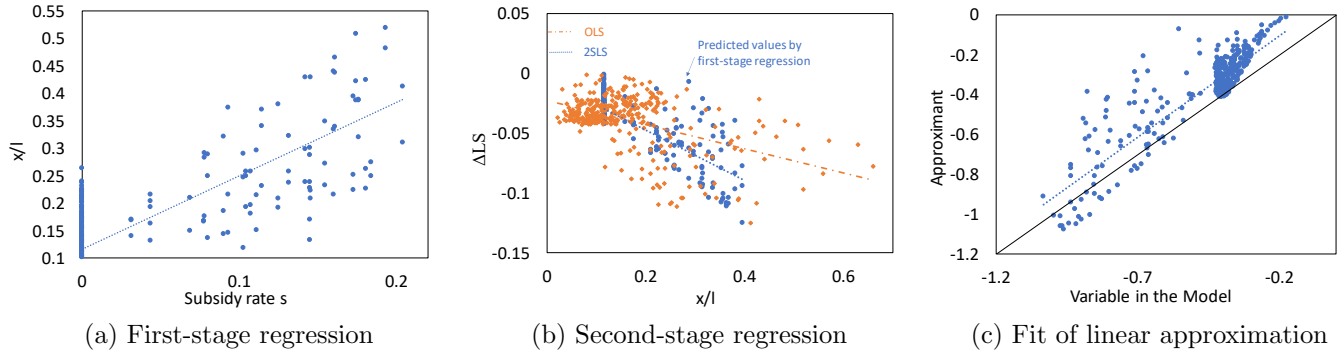


Figure 2: 2SLS Elasticity Estimate from Data Generated by Baseline Model

impact of automation on the labor share.

4 Conclusion

We have developed a new methodology to estimate the impact of automation on the labor share and have used it on micro data from China. We found a negative and large causal impact of automation on the labor share of automating firms. We also estimated an elasticity of substitution between labor and automation capital of about 3.8—the key structural parameter that describes how the labor share is affected by the price of automation. Our findings suggest that further declines in the price of industrial robots may lead to a sizable redistribution of a firm’s income away from workers and toward capital owners among automating firms. One has to be cautious in translating our results to aggregate effects. Throughout the paper, we have focused on the intensive margin of automation investment, i.e. how much firms that are already automating react to changes in the price of robots. Future research could also explore the role of the extensive margin to better understand how automation technologies are initially introduced in a workplace. Another important question we leave unanswered is how our findings extend to other sectors of the economy, outside manufacturing.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, tasks and technologies: Implications for employment and earnings,” in *Handbook of labor economics*, Elsevier, vol. 4, 1043–1171.
- ACEMOGLU, D. AND P. RESTREPO (2018): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108, 1488–1542.

- (2019): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*.
- ARNTZ, M., T. GREGORY, AND U. ZIERAHN (2016): “The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis,” .
- ATKESON, A. AND A. BURSTEIN (2008): “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 98, 1998–2031.
- AUTOR, D. AND A. SALOMONS (2018): “Is automation labor-displacing? Productivity growth, employment, and the labor share,” *Brookings Papers on Economic Activity*, BPEA Conference Drafts.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The skill content of recent technological change: An empirical exploration,” *The Quarterly journal of economics*, 118, 1279–1333.
- BRYNJOLFSSON, E. AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company.
- DAO, M. C., M. M. DAS, Z. KOCZAN, AND W. LIAN (2017): *Why is labor receiving a smaller share of global income? Theory and empirical evidence*, International Monetary Fund.
- DAVID, H. AND D. DORN (2013): “The growth of low-skill service jobs and the polarization of the US labor market,” *American Economic Review*, 103, 1553–97.
- DVORKIN, M. A. AND A. MONGE-NARANJO (2019): “Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations,” *FRB St. Louis Working Paper*.
- FORD, M. (2009): *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future*, CreateSpace Independent Publishing Platform.
- FREY, C. B. (2020): *The Technology Trap: Capital, Labor, and Power in the Age of Automation*, Princeton University Press.
- FREY, C. B. AND M. A. OSBORNE (2017): “The Future of Employment: How Susceptible are Jobs to Computerisation?” *Technological Forecasting and Social Change*, 114, 254 – 280.
- GABRIELE, C., R. A. DUVAL, AND D. FURCERI (2018): “Employment Protection Deregulation and Labor Shares in Advanced Economies,” *IMF Working Papers*.
- GOOS, M. AND A. MANNING (2007): “Lousy and lovely jobs: The rising polarization of work in Britain,” *The review of economics and statistics*, 89, 118–133.
- (2018): “Robots at Work,” *Review of Economics and Statistics*, C.
- HUBMER, J. (2020): “The Race Between Preferences and Technology,” *unpublished manuscript*.
- HUMLUM, A. (2019): “Robot Adoption and Labor Market Dynamics,” *Princeton University Working Paper*.
- KALDOR, N. (1961): “Capital Accumulation and Economic Growth,” in *The Theory of Capital*, ed. by F. Lutz and D. Hague, St. Martins Press, 177–222.
- KARABARBOUNIS, L. AND B. NEIMAN (2013): “The Global Decline of the Labor Share*,” *The Quarterly Journal of Economics*, 129, 61–103.
- KEHRIG, M. AND N. VINCENT (2018): “The Micro-Level Anatomy of the Aggregate Labor Share Decline,” *NBER Working Paper No. 25275*.

- MANYIKA, J., J. MISCHKE, J. BUGHIN, J. WOETZEL, M. KRISHNAN, AND S. CUDRE (2019): “A new look at the declining labor share of income in the United States,” *McKinsey Global Institute Discussion paper*.
- MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): “Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years,” *Review of Economics and Statistics*, 96, 60–77.
- MURO, M., R. MAXIM, AND J. WHITON (2019): “Automation and artificial intelligence: How machines are affecting people and places,” .
- NATIONAL BUREAU OF STATISTICS OF CHINA (1997): “China Statistical Yearbook,” .

Appendices

5 Implementation of MIC: Selected examples

Here we provide three examples of how automation subsidies under MIC have been implemented by different cities. The information is based on official program descriptions.

Foshan (Guangdong Province): “assign 130 million RMB per year to support automation and robotic machinery [...] the government provides subsidies on robots purchases — 12% of machines value if the robots are made in Foshan (the maximum subsidy cannot be larger than 3 million RMB per year); 8% of machines value if the robots are made elsewhere (the maximum subsidy cannot be larger than 2 million per year) [...] every year, the government awards 8 million RMB per firm to 10 selected firms as the automation demonstration based on the following criterions...”^{53,54}

Huzhou (Zhejiang Province): “to encourage automation, the government provides subsidies based on the following three categories: a) 6% of machines value (one time claim cannot be larger than 10 million RMB) for four industries — metal material, furniture, modern textile, and fashion products; b) 8% of machines value (one time claim cannot be larger than 15 million RMB) for three industries — information technology, advanced machinery, and biomedicine. c) 10% of machines value (one time claim cannot be larger than 20 million RMB) for the industrial areas of integrated circuit, new energy (including adoption to electronic vehicles, battery, and machinery), logistics equipment, aerospace and aviation equipment, new medical technology.”⁵⁵

Wuhan (Hubei Province): “the government provides subsidies at the rate 12% of machines value (one time claim cannot be larger than 3 million).”

As we can see from these examples, some cities offer blanket subsidies for automation while others target specific sectors of activity.⁵⁶

6 Variable definitions

Value-added (VA) : Sum of four components: 1) employees’ wages and compensations; 2) depreciation of fixed assets; 3) net taxes on production; 4) business earning surplus. This ap-

⁵³The explicit criterions are “a) the firm’s business registration is in the city of Foshan; b) the plant is in the city of Foshan; c) the annual sales of main business is at least 50 million RMB; d) total investment on automation is at least 30 million RMB; e) at least 60% of machines investment is for automation; e) the firm development strategy needs to be consistent with national industrial policy.

⁵⁴Text from the Bureau of Industry and Information Technology of Foshan City available at <http://fsiit.foshan.gov.cn/>.

⁵⁵Text from the Bureau of Economy and Information of Huzhou City available at <http://hzjx.huzhou.gov.cn/>.

⁵⁶Text from the official government website of Wuhan City at <http://www.wuhan.gov.cn/>.

proach follows the standard value-added calculation from the National Bureau of Statistics of China (unit: 10,000 RMB) but excludes subsidies that enter our analysis. Importantly, government subsidies are not part of value-added to avoid polluting the construction of the labor share variable. **Robots investment** (I_R) Purchase value of robots (unit: 10,000 RMB). **Automatic machine investment** (I_{AM}): The purchase value of automatic machines robots and semi-automatic machines (unit: 10,000 RMB). **Other machine investment** (I_{OM}) The purchase value of other machines (unit: 10,000 RMB). **Robots subsidy** (S_R) The amount of subsidy for purchasing robots (unit: 10,000 RMB) **Automatic machine subsidy** (S_{AM}) The amount of subsidy for purchasing automatic and semi-automatic machines (unit: 10,000 RMB). **Other machine subsidy** (S_{OM}) The amount of subsidy for purchasing other machines (unit: 10,000 RMB). **Employment** (L) **Total employees** (unit: person). **Lerner index** (LI) Operating income to value added. **Labor share** (ls) The ratio of labor wages and compensation to value-added. **Labor share of production employees** (also production employment share) (lsp) The ratio of labor wages and compensation of production workers to value-added. **Export intensity** (X/VA) The ratio of export sales to value-added.

7 Omitted proofs and derivations

This appendix includes derivations and additional results related to the model of Section 1.

Proof of Lemma 1:

The first-order conditions corresponding to the cost minimization problem (8) are

$$\begin{aligned}
k_s : r^s(t) k_s &= \lambda \gamma_\omega \eta_\omega \\
l_s : w l_s &= \lambda (1 - \gamma_\omega) \eta_\omega \\
k_e : r^e(t) k_e &= \lambda a_\omega^{\frac{1}{\sigma_\omega}} \alpha_\omega (1 - \eta_\omega) (k_e^{\alpha_\omega} l^{1-\alpha_\omega})^{\frac{\sigma_\omega-1}{\sigma_\omega}} \Gamma^{-1} \\
m : r^m(t) (1 - s_\omega) m &= \lambda (1 - a_i)^{\frac{1}{\sigma_\omega}} (1 - \eta_i) m^{\frac{\sigma_\omega-1}{\sigma_\omega}} \Gamma^{-1} \\
l : w l &= \lambda a_\omega^{\frac{1}{\sigma_\omega}} (1 - \alpha_\omega) (1 - \eta_\omega) (k_e^{\alpha_\omega} l^{1-\alpha_\omega})^{\frac{\sigma_\omega-1}{\sigma_\omega}} \Gamma^{-1}
\end{aligned} \tag{27}$$

where we define

$$\Gamma_\omega \equiv a_\omega^{\frac{1}{\sigma_\omega}} (k_e^{\alpha_\omega} l^{1-\alpha_\omega})^{\frac{\sigma_\omega-1}{\sigma_\omega}} + (1 - a_\omega)^{\frac{1}{\sigma_\omega}} m^{\frac{\sigma_\omega-1}{\sigma_\omega}}. \tag{28}$$

and where λ is the Lagrange multiplier on the constraint $F_\omega(k_s, l_s, k_e, l, m) \geq 1$. Because of the constant returns to scale, the Lagrange multiplier λ equals the firm's marginal cost. Combining the first-order conditions, it is straightforward to show that (10) holds and we omit an explicit derivation. Q.E.D.

Proof of Lemma 2:

We use the first-order conditions derived in the proof of Lemma 1. Combining the first-order conditions with respect to k_e and l , we have

$$(k_e^{\alpha_\omega} l^{1-\alpha_\omega})^{\frac{1}{\sigma_\omega}} = \lambda a_\omega^{\frac{1}{\sigma_\omega}} (1 - \eta_\omega) \left(\frac{\alpha_\omega}{r_i^e} \right)^{\alpha_\omega} \left(\frac{1 - \alpha_\omega}{w} \right)^{1-\alpha_\omega} \Gamma_\omega^{-1}, \quad (29)$$

which, together with the first-order condition with respect to m , gives

$$\left(\frac{m}{k_e^{\alpha_\omega} l^{1-\alpha_\omega}} \right)^{\frac{1}{\sigma_\omega}} = \left(\frac{1 - a_\omega}{a_\omega} \right)^{\frac{1}{\sigma_\omega}} \left(\left(\frac{\alpha_\omega}{r_i^e} \right)^{\alpha_\omega} \left(\frac{1 - \alpha_\omega}{w} \right)^{1-\alpha_\omega} \right)^{-1} (r_i^m (1 - s_\omega))^{-1}.$$

From the first-order conditions with respect to k_e and l we also have $\alpha_\omega w l = (1 - \alpha_\omega) r_e k_e$, and so the previous expression becomes

$$\frac{m}{l} = \left(\frac{1 - a_\omega}{a_\omega} \right) \left(\frac{\alpha_\omega w}{1 - \alpha_\omega r_i^e} \right)^{\alpha_\omega} \left(\left(\frac{\alpha_\omega}{r_i^e} \right)^{\alpha_\omega} \left(\frac{1 - \alpha_\omega}{w} \right)^{1-\alpha_\omega} \right)^{-\sigma_\omega} (r_m (1 - s_\omega))^{-\sigma_\omega}$$

which corresponds to (11) after simplifications. Q.E.D.

Proof of Lemma 3:

The labor share of support (nonproduction) workers is

$$LSN(t, \omega) := \frac{w(t) l_s(t, \omega)}{y(t, \omega)} = \frac{w(t) l_s(t, \omega)}{\frac{\theta_{j(t, \omega)}}{\theta_{j(t, \omega)} - 1} \lambda(t, \omega) q(t, \omega)} = \frac{\theta_{j(t, \omega)} - 1}{\theta_{j(t, \omega)}} (1 - \gamma_\omega) \eta_\omega,$$

where the last equality follows from the first-order condition for labor (scaled up to an arbitrary output q) from Lemma 1. Similarly for productive labor we can compute the labor share $LSP(t, \omega) := \frac{w(t) l(t, \omega)}{y(t, \omega)}$ so that $LS(t, \omega) = LSN(t, \omega) + LSP(t, \omega)$ by the definition of the labor share (13).

By combining the (scaled) first-order condition with respect to l with the definition of Γ (see Lemma 1) and the fact that $\alpha_\omega w l = (1 - \alpha_\omega) r_i^e k$ and $p = \frac{\theta_j}{\theta_j - 1} \lambda$, we have that

$$\begin{aligned} LSP(t, \omega) &:= \frac{w(t) l(t, \omega)}{y(t, \omega)} = \frac{\theta_{j(t, \omega)} - 1}{\theta_{j(t, \omega)}} (1 - \eta_\omega) (1 - \alpha_\omega) \\ &\quad \times \left(1 + \left(\frac{1 - a_\omega}{a_\omega} \right)^{\frac{1}{\sigma_\omega}} \left(\frac{1 - \alpha_\omega r_\omega^e}{\alpha_\omega w} \right)^{\alpha_\omega \frac{\sigma_\omega - 1}{\sigma_\omega}} \left(\frac{m(t, \omega)}{l(t, \omega)} \right)^{\frac{\sigma_\omega - 1}{\sigma_\omega}} \right)^{-1} \end{aligned}$$

or, equivalently,

$$\frac{1}{LSP(t, \omega)} = \left(\frac{\theta_j(t, \omega) - 1}{\theta_j(t, \omega)} (1 - \eta_\omega) (1 - \alpha_\omega) \right)^{-1} \left(1 + \left(\frac{1 - a_\omega}{a_\omega} \right)^{\frac{1}{\sigma_\omega}} \left(\frac{1 - \alpha_\omega r_i^e}{\alpha_\omega w} \right)^{\alpha_\omega \frac{\sigma_\omega - 1}{\sigma_\omega}} \left(\frac{m(t, \omega)}{l(t, \omega)} \right)^{\frac{\sigma_\omega - 1}{\sigma_\omega}} \right).$$

Combining with $LS(t, \omega) = LSN(t, \omega) + LSP(t, \omega)$ and using the definition of $LSO(t, \omega)$ in equation (4) yields the result. Q.E.D.

Proof of Lemma 4:

From (18), we can write

$$\frac{x(\tau, \omega)}{l(\tau, \omega)} = p_m(\tau) \frac{m(\tau, \omega)}{l(\tau, \omega)} - (1 - \delta_i^m) p_m(\tau - 1) \frac{m(\tau - 1, \omega)}{l(\tau - 1, \omega)} \frac{l(\tau - 1, \omega)}{l(\tau, \omega)}.$$

Using (11), the last expression becomes

$$\frac{x(\tau, \omega)}{l(\tau, \omega)} = p_m(\tau) \exp(\Theta(\tau, \omega)) (1 - s_\omega)^{-\sigma_\omega} \quad (30)$$

$$- p_m(\tau - 1) (1 - \delta_i^m) \exp(\Theta(\tau - 1, \omega)) (1 - s_{\omega, -1})^{-\sigma_\omega} \frac{l(\tau - 1, \omega)}{l(\tau, \omega)}, \quad (31)$$

where $\Theta(\tau, \omega)$ is given by (12) and where $s_{\omega, -1} := s(\tau - 1, \omega)$.

We handle the two terms on the right-hand side of that equation separately. Taking a first-order approximation of the first term with respect to s_ω around some level \bar{s} we find

$$\exp(\Theta(\tau, \omega)) (1 - s_\omega)^{-\sigma_\omega} \approx \exp(\Theta(\tau, \omega)) (1 - \bar{s})^{-\sigma_\omega} + \sigma_\omega \exp(\Theta(\tau, \omega)) (1 - \bar{s})^{-\sigma_\omega - 1} (s_\omega - \bar{s}).$$

As for the second term, we linearize it with respect s_ω to obtain

$$p_m(\tau - 1) (1 - \delta_i^m) \exp(\Theta(\tau - 1, \omega)) (1 - s_{\omega, -1})^{-\sigma_\omega} \frac{l(\tau - 1, \omega)}{l(\tau, \omega)} \approx \mathcal{N}(\tau - 1, \tau, \omega) + \mathcal{H}(\tau - 1, \tau, \omega) (s_\omega - \bar{s}),$$

where \mathcal{N} and \mathcal{H} are terms that depend on the linearization point but that are not functions of s_ω itself (just \bar{s}); that is \mathcal{N} and \mathcal{H} only depend on exogenous variables such as parameters and current level of shocks (other than s_ω). Here note that the linearization error is introduced only by linearization with respect to s_ω and not the value of the shocks; in other words, we do not linearize with respect to shock values and only take a linear approximation around their current value with respect to s_ω .

Putting these expressions back into (30) we find

$$\begin{aligned} \frac{x(\tau, \omega)}{l(\tau, \omega)} &= p_m(\tau) \exp(\Theta(\tau, \omega)) (1 - \bar{s})^{-\sigma_\omega} \\ &+ (p_m(\tau) \sigma_\omega \exp(\Theta(\tau, \omega)) (1 - \bar{s})^{-\sigma_\omega - 1} - \mathcal{H}(\tau - 1, \tau, \omega)) (s_\omega - \bar{s}) - \mathcal{N}(\tau - 1, \tau, \omega). \end{aligned}$$

Plugging in $s_\omega = s_i + s_c + \varepsilon_\omega^s$ from (17), taking expectation of both sides conditional on ε_ω^s , and using the assumed probability structure of subsidies together with Assumption 2, gives

$$\mathbb{E} \left[\frac{x(\tau, \omega)}{l(\tau, \omega)} \middle| \varepsilon_\omega^s \right] = cte + \mathbb{E} \left[p_m(\tau) \sigma_\omega \exp(\Theta(\tau - 1, \tau, \omega)) (1 - \bar{s})^{-\sigma_\omega - 1} - \mathcal{H}(\tau - 1, \tau, \omega) \right] \varepsilon_\omega^s,$$

where cte is a constant. Q.E.D.

Proof of Lemma 5:

We simplify the notation by writing $\varepsilon_\omega^s := \varepsilon^s(\tau, \omega)$, which is without loss under Assumption 2. A first-order approximation of (11) around a point \bar{s} yields

$$\log \left(\frac{m(\tau, \omega)}{l(\tau, \omega)} \right) \approx c_i(\tau, \omega) + \frac{\sigma_\omega}{1 - \bar{s}} s(\tau, \omega), \quad (32)$$

where $c(\tau, \omega) := \Theta(\tau, \omega) - \sigma_\omega \log(1 - \bar{s}) + \frac{\sigma_\omega}{1 - \bar{s}} \bar{s}$ includes all the constants. Combining (14) with (32) yields.

$$\log \frac{LSO(\tau, \omega) - LS(\tau, \omega)}{LS(\tau, \omega) - LSN(\tau, \omega)} = \Psi(\tau, \omega) + \frac{\sigma_\omega - 1}{\sigma_\omega} c_i(\tau, \omega) + \frac{\sigma_\omega - 1}{1 - \bar{s}} s(\tau, \omega). \quad (33)$$

Now computing the difference of that equation between the periods τ and $\tau - 1$ and taking the conditional expectation, we get,

$$\mathbb{E} \left[\log \frac{LSO_\tau - LS_\tau}{LS_\tau - LSN_\tau} - \log \frac{LSO_{\tau-1} - LS_{\tau-1}}{LS_{\tau-1} - LSN_{\tau-1}} \middle| \varepsilon_\omega^s \right] = cte + \mathbb{E} \left[\frac{\sigma_\omega - 1}{1 - \bar{s}} \right] \varepsilon_\omega^s, \quad (34)$$

where we have again used Assumption 2 and the probability structure (17) of the subsidy policy, and where we have abbreviated $LSO(\omega, \tau) \equiv LSO_\tau$ and so on. Now, note that the left-hand side of that expression can be written as

$$\begin{aligned} \log \frac{LSO_\tau - LS_\tau}{LS_\tau - LSN_\tau} - \log \frac{LSO_{\tau-1} - LS_{\tau-1}}{LS_{\tau-1} - LSN_{\tau-1}} &= \log(LS_{\tau-1} - LSN_{\tau-1}) \\ &\quad - \log(LS_\tau - LSN_\tau) \\ &\quad + \log(LSO_\tau - LS_\tau) \\ &\quad - \log(LSO_{\tau-1} - LS_{\tau-1}), \end{aligned} \quad (35)$$

and that we can approximate each term using the Taylor's expansion

$$\log(x - y) \approx \log(\bar{x} - \bar{y}) + \frac{1}{\bar{x} - \bar{y}} (x - \bar{x}) - \frac{1}{\bar{x} - \bar{y}} (y - \bar{y}).$$

So, for instance, we obtain

$$\begin{aligned} \log(LS_{\tau-1} - LSN_{\tau-1}) &\approx \log(\overline{LS} - \overline{LSN}) + \frac{1}{\overline{LS} - \overline{LSN}} (LS_{\tau-1} - \overline{LS}) \\ &\quad - \frac{1}{\overline{LS} - \overline{LSN}} (LSN_{\tau-1} - \overline{LSN}). \end{aligned}$$

Repeating this argument on all the terms on the right-hand side of (35), we find that it can be written as

$$\begin{aligned} \log \frac{LSO_\tau - LS_\tau}{LS_\tau - LSN_\tau} - \log \frac{LSO_{\tau-1} - LS_{\tau-1}}{LS_{\tau-1} - LSN_{\tau-1}} & \quad (36) \\ &= \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right) [LS_{\tau-1} - LS_\tau] \\ &\quad - \left(\frac{1}{\overline{LSO} - \overline{LS}} \right) [LSO_\tau - LSO_{\tau-1}] \\ &\quad + \frac{1}{\overline{LS} - \overline{LSN}} [LSN_\tau - LSN_{\tau-1}] + cte, \end{aligned}$$

where cte is a constant. We are interested in the conditional expectation of that quantity, and so

$$\begin{aligned} &\mathbb{E} \left[\log \frac{LSO_\tau - LS_\tau}{LS_\tau - LSN_\tau} - \log \frac{LSO_{\tau-1} - LS_{\tau-1}}{LS_{\tau-1} - LSN_{\tau-1}} \middle| \varepsilon_\omega^s \right] \\ &= cte + \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right) \mathbb{E} [LS_{\tau-1} - LS_\tau | \varepsilon_\omega^s], \end{aligned}$$

where we also use Assumption 2 to include terms in the constant. We combine this last expression with (19) and (34) to obtain

$$\left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right) \mathbb{E} [LS_{\tau-1} - LS_\tau | \varepsilon_\omega^s] = cte + \mathbb{E} \left[\frac{\sigma_\omega - 1}{1 - \bar{s}} \right] \frac{1}{\mathcal{L}} \mathbb{E} \left[\frac{x(\tau, \omega)}{l(\tau, \omega)} \middle| \varepsilon_\omega^s \right].$$

Q.E.D.

Proof of Proposition 1

Before we begin, we establish a preliminary technical result on the equivalence of two related statistical models.

Suppose that the data generating process follows

$$\begin{aligned} Z &= X\beta_A + v_A \\ X &= M_Y V \alpha_A + u_A, \end{aligned}$$

where α_A and β_A are coefficient vectors, v_A and u_A are error term vectors, X and V are predetermined

matrices, and where

$$M_Y = I - P_Y = I - Y(Y'Y)^{-1}Y'$$

is the orthogonal (linear) projection matrix (also known as annihilator matrix) that returns residual from projecting any conformable matrix on the subspace generated by Y (a predetermined variable). P_Y is the (linear) projection matrix that returns predicted values (also known as prediction matrix). We will refer to this system of equations as model A.

Consider an alternative model (model B) that follows the equations

$$\begin{aligned} Z &= X\beta_B + Y\gamma_B + v_B \\ X &= V\alpha_B + Y\delta_B + u_B \end{aligned}$$

where α_B , β_B , δ_B and γ_B are coefficient vectors and v_B and u_B are error term vectors that are orthogonal to Y in the sense that $M_Y v_B = M_Y u_B = 0$. Note that the above definition implies that the error terms are related as follows

$$X\beta_A + v_A = X\beta_B + Y\gamma_B + v_B$$

and

$$V\alpha_B + Y\delta_B + u_B = M_Y V\alpha_A + u_A,$$

which after multiplying both sides by the orthogonal projection matrix M_Y implies

$$M_Y V\alpha_B + M_Y u_B = M_Y V\alpha_A + M_Y u_A. \tag{37}$$

The parameters for model A and model B are identical and their estimates have identical properties, as stated in the lemma below.

Lemma 6. *The two-stage least squares estimators of model A and model B yield the same point estimates for β_A and β_B and for α_A and α_B , and these estimators have the same variance.*

The proof of this fact is a straightforward corollary of a standard result in econometrics: the Frisch–Waugh–Lovell theorem. Since it is standard and technical, we relegate it to the Technical Appendix.

The proof of the Proposition 1 now trivially follows from Lemma (4) and Lemma (5) and the basic properties of a conditional expectations operator, i.e., the fact that for any random variable X and Y (with bounded expectation) we have $Y = \mathbb{E}[Y|X] + \varepsilon$, we know 1) $\mathbb{E}[\varepsilon|X] = 0$, 2) $\mathbb{E}[\varepsilon] = 0$, 3) for any function $h(X)$ such that $\mathbb{E}[|h(X)\varepsilon|] < \infty$, $\mathbb{E}[h(X)\varepsilon] = 0$. It is clear that the first stage boils down to model A. Under the independence assumption the terms s_i, s_c are random variables that can be estimated by a fixed effect regression of the form $s = FE_c + FE_i + \varepsilon$ (see Woolridge

(2002), section 10.2.1). This is equivalent to estimating model B as stated in the proposition by the above lemma. Detailed proof can be found in Technical Appendix. Q.E.D.

8 Robustness to correlated pre-existing subsidies

In this section, we generalize the theoretical results of Section 1 to allow for a pre-existing subsidy for automation investment in the period $\tau - 1$ —which would have violated Assumption 2. More specifically, we assume that the part of subsidies $s(\tau, \omega)$ and $s(\tau - 1, \omega)$ that are orthogonal to the industry and city variables are related as follows:

$$\varepsilon^s(\tau - 1, \omega) = \rho_\omega^s \varepsilon^s(\tau, \omega) + v_\omega^s, \quad (38)$$

where ρ_ω^s and v_ω^s are i.i.d. shocks with mean 0. The results that are affected by this change are Lemma 4, Lemma 5 and Proposition 1. Since modifying the proofs is fairly straightforward we relegate the proofs to the Technical Appendix. The resulting statement of the proposition is identical. The only difference is that the constant \mathcal{B} that is carried around is given by

$$\mathcal{B} = -\frac{1}{1 - \bar{s}} \frac{1}{\mathcal{L}} \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right)^{-1} \mathbb{E}[(\sigma_\omega - 1)(1 - \rho_\omega^s)], \quad (39)$$

instead of by (21).

Implications for estimating the average elasticity $\mathbb{E}[\sigma_\omega]$

We now show that the above modification implies that, when $0 < \mathbb{E}[\rho_\omega^s]$, our estimated elasticity is biased toward a smaller $\mathbb{E}[\sigma_\omega]$ and, as a result, against finding a large effect of automation on the labor share. To see this, note that we can find $\mathbb{E}[\sigma_\omega]$ by using the estimated coefficients $\hat{\mathcal{B}}$ and $\hat{\mathcal{L}}$. Namely, from (39) we can write

$$\mathbb{E}[\sigma_\omega - 1] = -\hat{\mathcal{L}}\hat{\mathcal{B}} \frac{1 - \bar{s}}{\mathbb{E}[1 - \rho_\omega^s]} \left(\frac{1}{\overline{LS} - \overline{LSN}} + \frac{1}{\overline{LSO} - \overline{LS}} \right) \quad (40)$$

where we have used the fact that ρ_ω^s is independent of other random variables. It is clear that the right-hand side of (40) is a increasing function of $\mathbb{E}[\rho_\omega^s]$ (the product $\hat{\mathcal{L}}\hat{\mathcal{B}}$ is always negative in our estimates and $1 - \bar{s}$, $\overline{LS} - \overline{LSN}$ and $\overline{LSO} - \overline{LS}$ are positive by the restrictions of the model). As a result, by assuming that $\mathbb{E}[\rho_\omega^s] = 0$ as in our benchmark exercises, we would be biasing our estimate toward a smaller $\mathbb{E}[\sigma_\omega]$.