

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

Full length articles

The unequal effects of trade and automation across local labor markets[☆]Simon Galle^{a,*}, Linnea Lorentzen^b^a BI Norwegian Business School, Norway^b University of Oslo, Norway

ARTICLE INFO

Dataset link: <https://data.mendeley.com/datasets/jj6kdtcbvy/1>

JEL classification:

F11
F16
J23
J24
E24
E25

Keywords:

International trade
Automation
China shock
Inequality
Local labor markets

ABSTRACT

We quantify the joint impact of the China shock and automation of labor, across US commuting zones (CZs) in the period 2000–2007. To this end, we employ a multi-sector gravity model of trade with Roy-Fréchet worker heterogeneity across sectors, where labor input can be automated. Automation and increased import competition from China are both sector-specific; they lead to contractions in a sector's labor demand and a decline in relative income for CZs more specialized in that sector, amplified by a voluntary reduction in hours worked and an increase in frictional unemployment. The estimated model fits well with the aggregate performance of manufacturing subsectors and with the variation across CZs in changes in average income, the hourly wage, hours worked, the employment rate and employment in manufacturing. By itself, the China shock has stronger distributional effects than automation, but its impact on aggregate gains is less than a third of automation's impact.

1. Introduction

During the 2000–2007 period, employment in US manufacturing abruptly fell by 20.5%, or 3.5 million jobs (US Bureau of Labor Statistics, 2019), resulting in economic stagnation in many regions specialized in manufacturing. In the same period, the Chinese share in US imports roughly doubled, and the associated surge in Chinese import competition has contributed to the relative decline of US manufacturing regions (Autor et al., 2013a). However, while manufacturing employment fell, value added in manufacturing continued to grow. This indicates that the “surprisingly swift” decline of US manufacturing employment is not only due to the “China shock” (Pierce and Schott, 2020), but also a consequence of the rise of labor-saving technology (Acemoglu and Restrepo, 2020).

[☆] We are grateful for insightful feedback from audiences at Antwerp, Berkeley, BI, ECARES, Oslo Macro Group, OsloMet, Paris I, University of Oslo, the Econometric Society Virtual World Congress, the European Trade Study Group, the Latin American Meeting of the Econometric Society, the Mannheim Workshop on Labor Adjustments to New Technologies and Globalization, the Nordic International Trade Seminars, the Nordic Conference on Register Data and Economic Modelling, the Sardinia Empirical Trade Conference, and the Tilburg Growth, Trade, and International Finance Conference. We thank two anonymous referees, the editor, Pamela Bombarda, Fenella Carpena, Malik Çürük, Pawel Gola, Alfonso Irarrazabal, Espen Moen, Andreas Moxnes, Plamen Nenov, and particularly Andrés Rodríguez-Clare for insightful comments, and we are grateful to Andrés Rodríguez-Clare and Moises Yi for sharing data and code. We gratefully acknowledge financial support from the Peder Sather center, United States. All errors are our own.

* Corresponding author.

E-mail address: simon.galle@bi.no (S. Galle).<https://doi.org/10.1016/j.jinteco.2024.103912>

Received 9 March 2021; Received in revised form 26 February 2024; Accepted 27 February 2024

Available online 1 March 2024

0022-1996/© 2024 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license

<http://creativecommons.org/licenses/by/4.0/>.

This paper presents a unifying general equilibrium framework to quantify the impact of trade and automation on the decline of US manufacturing employment and the associated unequal effects across US local labor markets. Many studies examine the impact of either trade or technology on US manufacturing regions,¹ but there is little to no integrated general-equilibrium analysis of their joint impact. Importantly though, the interplay between the trade and automation shocks matters quantitatively. For instance, US productivity shocks due to automation affect world trade patterns, which are central to the calibration of the China shock. Hence, the joint calibration of the China and automation shocks may yield different quantitative results than when the shocks are calibrated separately. Moreover, the interplay of the shocks determines whether their distributional impact will be dampened or amplified – an aspect also missed when the shocks are analyzed in isolation.

Our unified framework starts from the combination of a multi-sector gravity model of trade (Arkolakis et al., 2012), with a Roy-Fréchet setup for labor supply, where workers sort into sectors based on their comparative advantage (Lagakos and Waugh, 2013). Additionally, we introduce equipment as an imperfect substitute for labor, allowing labor to be automated. As explained in Galle et al. (2023), the Roy-Fréchet model generalizes the specific factors intuition to a setting where labor is mobile across sectors. Specifically, the model allows for variation in the pattern of comparative advantage across commuting zones (CZs), and the resulting pattern of sectoral specialization determines how sensitive each CZ is to a particular sector-specific shock. In our model, both increased import competition and automation are sector-specific. Therefore, both can induce a contraction in a sector's labor demand and exert downward pressure on the wages it pays. As a result, CZs more specialized in sectors undergoing increased import competition or automation will tend to experience a relative decline in their average hourly wage. These relative differences are amplified by adjustments in the unemployment rate and in the number of hours worked.

To allow for these changes in unemployment and hours worked, we include a bare-bones search-and-matching framework and an intensive margin adjustment in our model, based on Kim and Vogel (2021). The resulting model has three central elasticities on the labor supply side: (i) the dispersion parameter of the Roy-Fréchet, which governs the cost of reducing sectoral specialization at the CZ-level; (ii) the intensive margin elasticity; and (iii) the elasticity of the employment rate to labor market tightness. We estimate these elasticities by employing a transparent shift-share estimation strategy and obtain parameter values that align with standard values in the literature. Our shift-share instrument captures any national sector-specific shocks, and the model explains how these shocks translate into relative changes in CZs' income.

Armed with our estimates, we jointly calibrate the China and automation shocks and quantify their impact on the US economy. We model the China shock as sector-specific Chinese technological growth and calibrate it by ensuring the model exactly matches the increase in US expenditure shares on Chinese manufacturing goods. In turn, we model automation as a shock to equipment-specific productivity, which lowers the cost share of labor in a sector. This is why we calibrate the automation shock such that the model exactly matches the changes in labor cost shares in US manufacturing sectors.

Our model finds that the combination of the automation and China shocks leads to an increase in aggregate US real income of 3.96%, while aggregate manufacturing employment declines by 0.81 percentage points. Given our parameter estimates, roughly 50% of the aggregate income gain arises from changes in the average hourly wage, while changes in hours worked and the employment rate account for the remaining 20% and 30% of the gain, respectively. Furthermore, we obtain a 0.97% standard deviation for the income gains across CZs, and – perhaps unsurprisingly – there is a strong concentration of low gains around the Rust Belt. More broadly, in line with the patterns described in Austin et al. (2018), large swaths of the Eastern Heartland tend to experience a decline in relative income due to these combined shocks.

When we simulate the individual shocks, we find that the distributional effects of the China shock are larger than the effects of automation. However, the aggregate gain from the rise of China is less than one third the size of automation's impact. Moreover, the impacts across CZs of the individual China and automation shocks are positively correlated. In turn, this positive covariance of the shocks' local effects implies that the variance in income effects of the joint shock is larger than the sum of the variance for the individual shocks. The importance of incorporating both shocks to understand the full distributional effects across CZs helps explain why in previous studies such as GRY, the variance of the model-predicted income effects of an individual shock appears too small compared to the observed variance.

The predictions from our estimated model fit well with the variation across CZs for the different margins of labor market adjustment. In terms of R^2 , our joint China and automation shock explains e.g. 8% of the observed variation in changes in average CZ income and 29% of the variation in changes in the employment rate and the manufacturing employment share. In addition to explaining a substantial share of the variation, the model slightly underpredicts the magnitude of the variance in the observed changes. In part, this is driven by conservative choices for the intensive and extensive margin elasticities.

Turning from the model fit of the cross-CZ variation to the fit of the cross-sector variation, the joint shock continues to perform well. In contrast, for the individual shocks, there is a mismatch between the model and the data in the performance of value added across manufacturing subsectors. Indeed, the individual China shock predicts a decline in all subsectors' value added, while the automation shock predicts an increase for most. In reality, roughly half of the manufacturing subsectors undergo a decline in value added, while the other half experiences an increase – a pattern the combined shock does match well. This finding again underscores the importance of integrating both trade and automation in a single model to properly understand the impact of these correlated shocks both across local labor markets and on US manufacturing in general.

¹ Fort et al. (2018) provide an excellent review of the literature on this topic.

Literature. Our paper is motivated by the large body of reduced-form work on the impact of trade and technology on local labor markets. Following the seminal work of Autor et al. (2013a) – henceforth ADH – many studies have examined the impact of trade on US localities, see for instance Hakobyan and McLaren (2016), Bloom et al. (2019), Greenland et al. (2019), Pierce and Schott (2020) and Besedeš et al. (2021).² In addition, Acemoglu and Restrepo (2020) and Dauth et al. (2021) examine the impact of robotization across local labor markets.³ Autor et al. (2015) is one of the few papers jointly examining trade and technology shocks. We complement this work by providing a unifying general equilibrium framework that can quantify such shocks' impact at the aggregate level.

Like us, Caliendo et al. (2019), Adão et al. (2020), Rodríguez-Clare et al. (2022) and Galle et al. (2023) – henceforth GRY – quantitatively examine the impact of the China shock on US local labor markets employing a gravity model of trade with a Roy-type labor supply side.⁴ We build specifically on GRY since we introduce automation in this type of model in order to examine trade and automation shocks in a unified framework. Moreover, we follow Kim and Vogel (2021) and incorporate additional margins of labor market adjustment, which provide amplification mechanisms for the income effects of any shock. These amplification mechanisms, together with the combined impact of the correlated China and automation shocks, imply that our model matches the observed labor market changes substantially better than GRY.

Atalay et al. (2018) and Burstein et al. (2019) employ Roy models to examine the impact of computerization on inequality. In particular, the latter paper is close to ours since it also has a role for international trade. However, these papers examine the skill premium or the gender wage gap. The focus of our paper is different, as we examine inequality between US commuting zones. Moreover, their setup with a Cobb–Douglas production function for labor and equipment would not be able to account for the fall in the labor share within a manufacturing subsector.⁵

Outside the quantitative trade literature, there is extensive literature studying the impact of technical change on inequality or other labor market outcomes (Krusell et al., 2000; Acemoglu and Restrepo, 2018; Hémous and Olsen, 2022; Moll et al., 2022). Most closely related to us, Acemoglu and Restrepo (2020) present a model for the local impact of technological shocks, but their model does not feature international trade or worker heterogeneity. A further contribution of our paper is that the generalized specific factors mechanism in the Roy–Fréchet framework can help to tractably analyze the impact of sector-specific automation on different worker groups.

This paper is organized as follows. Section 2 introduces the theory, Section 3 discusses the data, and Section 4 examines the model's shift-share approximation and estimates the key labor supply elasticities. Section 5 presents our quantitative results and discusses their fit with the data. Finally, Section 6 concludes.

2. Theory

2.1. Setup

The model consists of three blocks. First, for labor supply we have a discrete choice model of the labor market where workers sort across the S sectors. Specifically, we assume a standard Roy (1951) model, where workers' comparative advantage across sectors determines the sorting pattern. Second, demand for goods in each sector for each of the N countries is governed by a gravity model with an input–output loop as in Costinot and Rodríguez-Clare (2014) and Caliendo and Parro (2015). Finally, we introduce a CES production function where labor and equipment are imperfect substitutes. We model automation as an increase in the productivity of equipment, which leads to a decrease of the labor share in production. We keep track of all the Roman and Greek notation in the model in Appendix Tables A.1 and A.2.

Labor supply. We start from a Roy–Fréchet model of labor supply as in Lagakos and Waugh (2013), and combine it with frictional unemployment and an intensive margin of labor supply as in Kim and Vogel (2021).⁶ To start, wages w_{os} vary by country o and sector s ; and workers differ in their productivity z_s across sectors. Given wages, workers apply to a specific sector to maximize their

² Notable contributions on the local impact of trade shocks outside the US are Topalova (2010), Kovak (2013), Dauth et al. (2014), and Dix-Carneiro and Kovak (2017).

³ The reduced-form literature on the impact of trade and technology on local labor markets typically employs shift-share estimation, and the prominence of this estimation technique has motivated fundamental econometric work on identification and inference with shift-share instruments (Adão et al., 2019; Borusyak et al., 2020; Goldsmith-Pinkham et al., 2020). The resulting deeper understanding of the potential limitations of shift-share estimation indicates the usefulness of a quantitative model to jointly analyze the implications of trade and technology shocks, both on aggregate and cross-sectional moments. Interestingly then, for the latter our model yields predictions both across sectors and across localities, which we will leverage in our model fit discussion.

⁴ Hsieh and Ossa (2016) also quantify the rise of China but do not focus on its distributional impact in the US. Other quantitative trade papers studying worker reallocation via a Roy model are Artuç et al. (2010), Dix-Carneiro (2014), Adão (2016), Curuk and Vannoorenberghe (2017), and Lorentzen (2022), but they all employ a small open economy framework. More similar to us, Lee (2020) employs a gravity model but does not focus on the China shock. In addition, Lee and Yi (2018) examine the impact of trade with China but focus on the skill premium instead of regional inequality. Finally, gravity models are also employed to study the impact of trade on the skill premium (Burstein and Vogel, 2017), or on unequal price changes for different consumption baskets (Fajgelbaum and Khandelwal, 2016).

⁵ In a recent contribution, Bernon and Magerman (2022) generalize the impact of productivity shocks on income inequality arising in general equilibrium in a Roy–Fréchet setup.

⁶ As in Kim and Vogel (2021), we could also add voluntary non-employment by introducing home production as one of the sectors. We refrain from doing so since it would make the estimation less tractable (see Section 8.2 in GRY). As an aside, note that Kim and Vogel (2020) generalizes the setup from Kim and Vogel (2021) and derives sufficient statistics for welfare analysis of trade shocks.

expected utility, which will be a function of $w_{os}z_s$. After applying, there is random matching between vacancies and workers. We assume that unemployed workers have no income.⁷ Conditional on being hired though, workers unilaterally decide how many hours of labor to supply. Specifically, they have a standard utility function with consumption of the final good (C) and hours worked (H) as elements:

$$U(C, H; og) = \delta_{og}C - \frac{H^{1+\mu}}{1+\mu},$$

where consumption is funded by the value of a worker's earnings, at price P_o . After output is realized, the worker and firm engage in Nash bargaining over the surplus of the match.

We are interested in between-group inequality. Each country o therefore has G_o groups of workers, analogous to the setups in [Burstein et al. \(2019\)](#), [Hsieh et al. \(2019\)](#) and GRY. In our application a group will be defined as a commuting zone (CZ), but the theory can be applied to worker groups defined by any pre-determined characteristic. Groups differ in their productivity distributions, which will lead to differences in sectoral specialization across groups. Specifically, a worker from group g in country o has a number of effective units of labor z_s , drawn from a Fréchet distribution with level parameter A_{ogs} and dispersion parameter κ .⁸ Here, the A_{ogs} parameters govern differences in groups' absolute advantage for each sector, while κ determines the dispersion of productivity within a sector. Due to the properties of the Fréchet, κ will also determine the dispersion of workers' comparative advantage, and thereby the elasticity of their labor supply across sectors. L_{og} denotes the measure of workers for group og .

The Nash bargaining between the worker and the firm results in a share v_{og} of revenue going to the worker.⁹ We backwardly solve the other elements of the labor problem in Appendix Section G.1. There, given the properties of the Fréchet, we find that the share of workers in group og that apply to sector s is

$$\pi_{ogs} = \frac{A_{ogs}w_{os}^\kappa}{\Phi_{og}^\kappa},$$

where $\Phi_{og} \equiv (\sum_k A_{ogk}w_{ok}^\kappa)^{1/\kappa}$ is a group-level index of sectoral wages, where the weights indicate the importance of each sector for group og . Next, average real income per worker in og is

$$\frac{v_{og}I_{ogs}}{\pi_{ogs}L_{og}P_o} = \eta e_{og} \delta_{og}^\mu \left(\frac{v_{og}}{P_o}\right)^{\frac{1+\mu}{\mu}} \Phi_{og}^{\frac{1+\mu}{\mu}},$$

with I_{ogs} nominal revenue in s for group og and $\eta \equiv \Gamma\left(1 - \frac{1+\mu}{\mu\kappa}\right)$. Consequently, the share of income obtained by workers of group og in sector s is also given by the sectoral employment share π_{ogs} . Total nominal revenue in group og is

$$I_{og} \equiv \sum_s I_{ogs} = \eta e_{og} \left(\frac{\delta_{og}v_{og}}{P_o}\right)^{\frac{1}{\mu}} \Phi_{og}^{\frac{1+\mu}{\mu}} L_{og}, \quad (1)$$

of which workers earn $v_{og}I_{og}$.

Relatedly, we can show that average expected utility in group og is

$$U_{og} = \frac{\mu}{1+\mu} \eta e_{og} (\delta_{og}v_{og})^{\frac{1+\mu}{\mu}} \left(\frac{\Phi_{og}}{P_o}\right)^{\frac{1+\mu}{\mu}}, \quad (2)$$

and that both hours worked (h_{og}) and hourly income (i_{og}) are also functions of Φ_{og} :

$$h_{og} = \hat{\eta} (\delta_{og}v_{og})^{\frac{1}{\mu}} \left(\frac{\Phi_{og}}{P_o}\right)^{\frac{1}{\mu}}, \quad (3)$$

$$i_{og} = \frac{\eta}{\hat{\eta}} \Phi_{og}. \quad (4)$$

Hence, Φ_{og} (a group's wage index) determines endogenous differences in average hourly labor income across groups and thereby also differences in average hours worked. Moreover, below we will see that it also determines differences in the employment rate, such that we can solve for I_{og} and U_{og} as a function of Φ_{og} and P_o .

Unemployment. A Cobb–Douglas matching function with vacancies \tilde{V}_{ogs} and applicants $\pi_{ogs}L_{og}$ as inputs entails that the employment rate is a function of labor market tightness $\psi_{ogs} \equiv \tilde{V}_{ogs}/(\pi_{ogs}L_{og})$:

$$e_{ogs} \equiv A_{ogs}^M \psi_{ogs}^\chi,$$

⁷ Since employment status is determined purely at random, it is relatively straightforward to introduce unemployment benefits in our framework, since they would not distort any decisions. However, we omit them for simplicity.

⁸ Three parameters that we introduce in this section (κ , μ and χ) are allowed to be og -specific in the theory. However in the empirics we will only estimate one value for each parameter, as is common in the literature. We therefore leave out the og subscript from the start, to ease the notational burden.

⁹ At the time of bargaining, the outside option of the match has zero value for both worker and firm, since hours worked and the cost of posting a vacancy are both sunk. Hence, at this point, the surplus of the match is exactly $u_{os}z_s$, multiplied by the number of hours worked.

where A_{og}^M measures matching efficiency and the employment rate elasticity is $0 < \chi < 1$. For the employer, the cost of posting a vacancy is $c_{og} P_o$, while the expected benefit is the share of revenue per vacancy accruing to the employer: $(1 - v_{og}) I_{ogs} / \tilde{V}_{ogs}$. The implied zero-profit condition together with the matching function determine the employment rate, which is indeed constant across sectors:

$$e_{og} = \left(A_{og}^M \left(\frac{\eta(1 - v_{og})}{c_{og}} \right)^\chi (\delta_{og} v_{og})^\frac{\chi}{\mu} \left(\frac{\Phi_{og}}{P_o} \right)^\frac{\chi(1+\mu)}{\mu} \right)^\frac{1}{1-\chi}, \tag{5}$$

where we can set parameters such that $0 < e_{og} < 1$. Intuitively, a shock that increases the real value of a typical worker (Φ_{og}/P_o), increases the return to posting a vacancy in group og across all sectors, and thereby pushes the employment rate up in all sectors. Importantly, the combinations of Eqs. (2) and (1) with (5) imply that

$$U_{og} \propto \frac{I_{og}}{P_o} \propto \left(\frac{\Phi_{og}}{P_o} \right)^\frac{1+\mu}{(1-\chi)\mu}.$$

Hence, the final good price (P_o) has a common effect on all groups' utility and real income, while sectoral wages lead to differences in utility and income across groups as measured by the group's index of sectoral wages (Φ_{og}). Also note that $I_{og} = e_{og} h_{og} i_{og}$, where the employment rate (e_{og}), hours worked (h_{og}), and the real hourly wage (i_{og}) are all functions of Φ_{og}/P_o . This is the reason for the amplification exponent in the welfare and income equation above, as an amplification of Φ_{og}/P_o through adjustments in employment (governed by $1/(1 - \chi)$) and hours worked (governed by $(1 + \mu)/\mu$).

Trade. There are iceberg trade costs $\tau_{ods} \geq 1$ to export goods from origin country o to destination country d , with $\tau_{oos} = 1$. We work with the multi-sector version of the Eaton and Kortum (2002) gravity model. Hence, within each sector, there is a continuum of goods of measure one, which have constant returns to scale technologies and good-specific productivities drawn from a Fréchet distribution with shape parameter θ_s and level parameter T_{os} in country o and sector s . Preferences across goods within a sector are CES with elasticity of substitution $\sigma_s < \theta_s$.

All this results in the following import shares for sector s in destination country d originating from country o :

$$\lambda_{ods} = \frac{T_{os} (\tau_{ods} c_{os})^{-\theta_s}}{\sum_i T_{is} (\tau_{ids} c_{is})^{-\theta_s}}, \tag{6}$$

where a good's marginal cost c_{os} is determined below. The price index in sector s in country d is then

$$P_{ds} = \tilde{\eta}_s^{-1} \left(\sum_o T_{os} (\tau_{ods} c_{os})^{-\theta_s} \right)^{-1/\theta_s}, \tag{7}$$

where $\tilde{\eta}_s \equiv \Gamma(1 - \frac{\sigma_s - 1}{\theta_s})^{1/(1 - \sigma_s)}$. The final good is the Cobb–Douglas composite of sectoral goods, with expenditure shares β_{ds} , and its price is $P_d = \prod_s P_{ds}^{\beta_{ds}}$.

Factor demand. Each sector has a two-tiered production function, where the upper-tier is Cobb–Douglas:

$$Y_{os} = F_{os}^{\alpha_{os}} K_{os}^{1 - \alpha_{os} - \gamma_{os}} \prod_k N_{oks}^{\gamma_{oks}}, \tag{8}$$

where N_{oks} are intermediate inputs sourced from sector k with $\gamma_{os} \equiv \sum_k \gamma_{oks}$, and K_{os} are structures. Next, F_{os} is a lower tier CES:

$$F_{os} = \xi_{os}^\nu \left[\xi_{os}^\frac{\rho}{\rho} M_{os}^\frac{\rho-1}{\rho} + Z_{os}^\frac{\rho-1}{\rho} \right]^\frac{\rho}{\rho-1},$$

where $Z_{os} = \sum_{g \in G_o} Z_{og}$ are effective units of labor and M_{os} is equipment (or machines). Both K_{os} and M_{os} are units of the final good, purchased at cost P_o .¹⁰ The parameter ξ_{os} appears twice; within the brackets it drives automation by shifting production toward equipment usage, while ν (outside the brackets) regulates the associated productivity gains. This is a tractable reduced-form approach of capturing the impact automation has in a task-based model.¹¹

¹⁰ Since we focus on hat algebra across static equilibria below, we abstract from investment and its dynamics for M_{os} and K_{os} , analogous to Burstein et al. (2019). This abstraction necessarily also entails that our model is not able to analyze the differential impact of automation on the income of asset holders versus workers – let alone the dynamics for these variables, as in Moll et al. (2022).

¹¹ We thank an anonymous referee for pointing us to this setup where ν regulates the productivity gains, and for pointing out how this is a reduced-form approach to modeling the production as in a task-based production setup, which can map into a CES production function (Acemoglu and Restrepo, 2018). Other papers modeling automation employing a CES production function are Krusell et al. (2000) or Hémous and Olsen (2022). These papers are interested in explaining changes in the skill premium, and this in part motivates their assumption that only low-skill labor can be automated. However, since our paper is primarily interested in inequality across CZs and less in inequality across skill groups, we do not introduce a distinction between skill groups in the baseline model. This is similar to Acemoglu and Restrepo (2020), who also focus on inequality across CZs and likewise do not distinguish between skill groups. In contrast to Acemoglu and Restrepo (2020), our model is tractable without introducing differences in tasks, and restrictions on the set of tasks that can be automated. This results from the combination of the CES production function and the imperfectly elastic labor supply to each sector in the Roy–Fréchet setup.

Cost minimization implies that ω_{os} , the labor share of expenditure on F_{os} , is

$$\omega_{os} = \frac{w_{os}^{1-\rho}}{\left[\xi_{os} P_o^{1-\rho} + w_{os}^{1-\rho} \right]}, \quad (9)$$

while c_{os} , the marginal cost for an extra unit of Y_{os} , has the standard form arising from Cobb Douglas production functions. Given perfect competition, total revenue for a sector in country o is $R_{os} \equiv c_{os} Y_{os}$. Note that the above implies that the cost share of capital in production is $(1 - \alpha_{os} - \gamma_{os}) = P_o K_{os} / R_{os}$.

Equilibrium. Cobb Douglas preferences and technologies imply that expenditure on a sector is $X_{ds} = \beta_{ds}(V_d + D_d) + \sum_{k=1}^S \gamma_{dsk} R_{dk}$. Here, value added is $V_d = \sum_{k=1}^S (1 - \gamma_{dk}) R_{dk}$, and trade deficits (D_d) are such that $\sum_d D_d = 0$. Goods market clearing implies that $R_{os} = \sum_d \lambda_{ods} X_{ds}$.

Total payments to labor in a sector are $\alpha_{os} \omega_{os} R_{os}$, while total labor income is $\sum_g \pi_{ogs} I_{og}$,¹² and in the labor market the two need to equalize in equilibrium. Hence, we can write excess labor demand in sector s of country o as

$$ELD_{os} = \alpha_{os} \omega_{os} R_{os} - \sum_g \pi_{ogs} I_{og}. \quad (10)$$

Note that ω_{os} (the labor share), R_{os} (revenue), π_{ogs} (employment shares), and I_{og} (labor income) are functions of the matrix of wages $\mathbf{w} \equiv \{w_{os}\}$. The system $ELD_{os} = 0$ for all o, s is therefore a system of equations in \mathbf{w} whose solution gives the equilibrium wages and prices given a choice of numeraire.

2.2. Counterfactual equilibrium

When interested in the impact of changes in trade costs τ_{ods} , national technology T_{os} , equipment productivity ξ_{os} , or deficits D_d , we can use exact hat algebra, where $\hat{x} \equiv x'/x$, to solve for the proportional change in the endogenous variables (Dekle et al., 2008). Formally, for shocks $\hat{\tau}_{ods}$ for $o \neq d$, \hat{T}_{os} , $\hat{\xi}_{os}$ or \hat{D}_d we compute the counterfactual equilibrium with $ELD'_{os} = 0$ for all o, s .¹³ To this end, we write $ELD'_{os} = 0$ as

$$\sum_g \hat{\pi}_{ogs} \pi_{ogs} \hat{I}_{og} I_{og} = \alpha_{os} \hat{\omega}_{os} \omega_{os} \sum_d \hat{\lambda}_{ods} \lambda_{ods} \left(\beta_{ds} (\hat{V}_d V_d + \hat{D}_d D_d) + \sum_{k=1}^S \gamma_{dsk} \hat{R}_{dk} R_{dk} \right). \quad (11)$$

This NxS system of equations can be solved for $\{\hat{w}_{os}\}$ given data on income levels I_{og} , trade shares λ_{ods} , expenditure shares β_{os} , revenue R_{os} , labor allocation shares π_{ogs} , cost shares α_{os} , ω_{os} , and γ_{ok} , and values for the exogenous shocks (see Appendix Section G.2).

Importantly, the change in a sector's labor share is

$$\hat{\omega}_{os} = \frac{\hat{w}_{os}^{1-\rho}}{\left[(1 - \omega_{os}) \hat{\xi}_{os} \hat{P}_o^{1-\rho} + \omega_{os} \hat{w}_{os}^{1-\rho} \right]}. \quad (12)$$

So conditional on wage and price changes, an increase in equipment productivity ($\hat{\xi}$), or ‘‘automation’’, lowers the labor share. On the trade side, trade or technology shocks will affect demand for each country's goods:

$$\hat{\lambda}_{ods} = \frac{\hat{T}_{os} (\hat{\tau}_{ods} \hat{c}_{os})^{-\theta_s}}{\sum_i \lambda_{ids} \hat{T}_{is} (\hat{\tau}_{ids} \hat{c}_{is})^{-\theta_s}}. \quad (13)$$

Shifts in labor demand lead to wage changes, which induce workers to move across sectors with elasticity κ :

$$\hat{\pi}_{ogs} = \frac{\hat{A}_{ogs} \hat{w}_{os}^\kappa}{\hat{\Phi}_{og}^\kappa}, \quad (14)$$

where $\hat{\Phi}_{og}$ is the change in the average hourly wage in group og (see Eq. (4)), with

$$\hat{\Phi}_{og} = \left(\sum_k \pi_{ogk} \hat{A}_{ogk} \hat{w}_{ok}^\kappa \right)^{\frac{1}{\kappa}}. \quad (15)$$

Given that the change in the employment rate is

$$\hat{e}_{og} = \left(\hat{A}_{og}^M \right)^{\frac{1}{1-\chi}} \left(\frac{\hat{\Phi}_{og}}{\hat{P}_o} \right)^{\frac{\chi(1+\mu)}{(1-\chi)\mu}}, \quad (16)$$

¹² A group's supply of effective labor units to sector s is $Z_{ogs} = I_{ogs}/w_{os}$, and the Fréchet implies that $\pi_{ogs} I_{og} = w_{os} Z_{ogs}$.

¹³ Throughout the analysis in this and the following sections, we assume for simplicity that $\hat{v}_{og} = 1$. It is straightforward to relax this.

and knowing the expressions for nominal income and utility from (1) and (2), the change in real income per worker is

$$\hat{U}_{og} = \frac{\hat{I}_{og}}{\hat{P}_o} = \left(\hat{\Lambda}_{og}^M \right)^{\frac{1}{(1-\chi)}} \left(\frac{\hat{\Phi}_{og}}{\hat{P}_o} \right)^{\frac{1}{(1-\chi)} \frac{1+\mu}{\mu}}. \quad (17)$$

Importantly, $\hat{\Phi}_{og}$ in Eq. (15) has a generalized specific factors interpretation, with κ governing the degree to which workers are a specific factor. When $\kappa \rightarrow 1$, any sector's wage change translates directly to a change in nominal income, weighted by a group's specialization in that sector (π_{ogs}). At the other extreme, when $\kappa \rightarrow \infty$, labor supply is perfectly elastic across sectors and as a result, all sectors experience an identical wage change and all groups undergo the same income change. Hence, between-group distributional effects are stronger for κ close to 1, and disappear when $\kappa \rightarrow \infty$. In Eq. (17), the exponent on $\hat{\Phi}_{og}/\hat{P}_o$ is increasing in the elasticity of the employment rate to labor market tightness (χ), and in the intensive margin elasticity ($1/\mu$). This term therefore represents an amplification effect driven by changes in the employment rate and hours worked.

3. Data

Estimation data. To estimate the labor supply-side parameters of the model, we use worker-level data from IPUMS (Ruggles et al., 2019). There, we restrict the sample to individuals who are between 25 and 60 years old, and also exclude government or non-profit employees, family workers, and institutionalized individuals. IPUMS provides info on total earned income over the past year, total number of hours worked, and employment status. We measure the hourly wage as average income per hour, and the employment rate as the share of employed individuals in the labor force. The model predicts that sectoral employment shares ($\pi_{U,S,g,s}$) in terms of income or hours worked should be perfectly correlated, while in practice the correlation is “only” 97% (see Appendix Fig. A.1). Since the correlation appears a little less strong for the larger $\pi_{U,S,g,s}$, we employ both the income- and hours-based measure for these shares in the estimation.

Our analysis focuses on changes over time with the year 2000 as the start period, which is the first year where IPUMS and our international trade dataset both have available data, and an end period before the onset of the Great Recession. For the year 2000, IPUMS provides the 5% sample of the Census, but for the end period only the American Community Survey (ACS) is available. Similar to the strategy in ADH, we therefore combine the ACS surveys for the years 2005-2006-2007 to ensure a more precise measurement of sector-level variables for all commuting zones. Throughout, we deflate income to 1999 dollars using the CPI.

Simulation data. For the international trade data, we employ the 2016 release of the World Input-Output Database (WIOD) - see Timmer et al. (2015). In the simulations, we measure the labor compensation share ($\alpha_{os}\omega_{os}$) as the labor share of a sector's value added in the WIOD Socio Economic Accounts for all 43 countries in WIOD.

Given the observed values for $\alpha_{os}\omega_{os}$, we disentangle the values of α_{os} (the cost share of the lower-tier CES F_{os}) and ω_{os} as follows. First, we measure the share of structures in the total value of structures and equipment in each sector as $\zeta_{os} \equiv P_o K_{os} / (P_o K_{os} + P_o M_{os})$. Second, we note that $(P_o K_{os} + P_o M_{os})/Y_{os} = 1 - \gamma_{os} - \alpha_{os}\omega_{os}$ and therefore the output share of K_{os} is $(1 - \alpha_{os} - \gamma_{os}) = \zeta_{os}(1 - \gamma_{os} - \alpha_{os}\omega_{os})$, which allows us to solve for α_{os} and ω_{os} . We measure ζ_{os} based on sector-level data on structures and equipment from EU-KLEMS (The Conference Board, 2023; van Ark and Jäger, 2017) and the OECD (OECD, 2019). As in Krusell et al. (2000), we group structures and transportation equipment under K_{os} and the other asset types under M_{os} . Appendix F provides more background, as well as detailed summary statistics for all the cost share parameters.

We employ sector definitions based on ISIC Rev.4, and to create consistent sectors across all estimation and simulation datasets, we aggregate sectors to 23 industries. Of these industries, two are in the primary sector and 11 are in manufacturing.¹⁴ The full list is in Appendix Table A.3.

In the simulations, we need to ensure that groups' labor income $\sum_g I_{ogs}$ is consistent with the WIOD revenue measures: $\sum_g I_{ogs} = \alpha_{os}\omega_{os}Y_{os}^{WIOD}$, where the superscript denotes the data source. To ensure this, we adjust IPUMS income to WIOD revenue as follows:

$$I_{U,S,g,s} = \left(\frac{\nu_{U,S} I_{U,S,g,s}}{\sum_h \nu_{U,S} I_{U,S,h,s}} \right)^{IPUMS} \alpha_{U,S,s} \omega_{U,S,s} Y_{U,S,s}^{WIOD},$$

where we assume that $\nu_{U,S,g}$ is constant across US groups. We then measure group-level sectoral employment shares as $\pi_{U,S,g,s} = I_{U,S,g,s} / \sum_s I_{U,S,g,s}$.

4. Estimation

4.1. A shift-share approximation

Our model features trade and technology shocks at the national level, which lead to changes in wages per effective unit of labor (\hat{w}_{os}), which in turn result in unequal changes in income and welfare across commuting zones (see Eqs. (15) and (17)). While the \hat{w}_{os} are unobservable, we can derive an observable shift-share measure that closely approximates the model-predicted impact of

¹⁴ The main data limitation here is arising from EU-KLEMS, which has data on just 11 manufacturing sectors. (WIOD has data on 19 manufacturing sectors.) Since we then only have these 11 manufacturing sectors, we are missing the additional variation in Chinese import penetration and changes in the labor share within these sectors. This may lead us to understate the aggregate and distributional effects of the China and automation shock.

Table 1
Estimating the model-implied shift-share approximation.

	(1)	(2)	(3)	(4)
	$\ln \hat{I}_g$	$\ln \hat{I}_g$	$\ln \hat{I}_g$	$\ln \hat{I}_g$
$\ln \sum_s \pi_{gs}^{hours} \hat{r}_s$	1.23*** (0.17)	1.18*** (0.27)		
$\ln \sum_s \pi_{gs}^{income} \hat{r}_s$			1.13*** (0.16)	0.94*** (0.25)
Controls	No	Yes	No	Yes
Observations	722	722	722	722

The regressions in this table estimate Equation (19), where the π_{gs} are measured based on hours worked or income. \hat{I}_g is the change in average income per person in the labor force in a CZ, with the unemployed earning zero income. The even-numbered specifications include the following control variables from ADH: dummies for the nine Census divisions, the average offshorability index of occupations, and percentages of employment in manufacturing, college-educated population, foreign-born population, and employment among women, where these percentages are all measured at the start of the period. Standard errors, clustered at the state level, in parentheses. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

changes in national wages across commuting zones. Specifically, assuming $\hat{A}_{og}^M = 1$ and defining sectoral labor income shares as $r_{os} \equiv \sum_g I_{ogs}/I_o$, we show in Appendix Section G.3 that:¹⁵

$$\frac{\hat{I}_{og}}{\hat{I}_o} \approx \left(\sum_s \pi_{ogs} \hat{A}_{ogs} \hat{r}_{os} \right)^{\frac{1+\mu}{\kappa(1-\chi)\mu}}. \tag{18}$$

We bring this approximation to US data with commuting zones as groups. In doing so, we simplify notation by dropping the country subscript and assume that groups face a uniform productivity shock across all sectors: $\hat{A}_{gs} = \hat{A}_g$. We can then estimate:

$$\ln \hat{I}_g = \alpha_{ss} + \beta_{ss} \ln \left(\sum_s \pi_{gs} \hat{r}_s \right) + \varepsilon_{ss,g}, \tag{19}$$

where $\alpha_{ss} \equiv \ln \hat{I}$, $\beta_{ss} \equiv \frac{(1+\mu)}{\kappa(1-\chi)\mu}$ and $\varepsilon_{ss,g} \equiv \ln \hat{A}_g^{\frac{(1+\mu)}{\kappa(1-\chi)\mu}}$. Here, $\sum_s \pi_{gs} \hat{r}_s$ measures a group's exposure to sectors' national expansion or contraction, weighted by initial employment shares. Our shift-share variable, therefore, incorporates the impact of any sector-specific supply or demand shock at the national level, including any trade or technology shocks. Below, we will use the estimate of β_{ss} as a consistency check for the individual estimates of κ , μ , and χ . Note that β_{ss} contains the same amplification term as in our welfare expression (17), arising from adjustments on the intensive and extensive margins. In addition, the term is now divided by κ , which is the parameter that measures workers' scope for reallocation across sectors.

We estimate the elasticity of our shift-share approximation in Table 1, both with and without controls and for our two measures of the employment shares π_{gs} (based on hours or income). While other shift-share estimations hinge on isolating the impact of one particular shock in their estimation (e.g. only the China shock or only robotization), this is not the case here since our shift-share variable incorporates all national-level shocks. We therefore estimate specification (19) using OLS.

The obtained estimates are strongly statistically significant, with coefficients between 0.94 and 1.23. The 95% confidence intervals overlap across all four specifications. These estimates imply strong distributional effects since a coefficient above unity implies amplification effects on top of the specific factors case ($\kappa \rightarrow 1$).

In Appendix B, we document the absence of pre-trends related to the shift-share estimation (Tables B.3, B.4), and that the variation in this variable is strongly affected by the ADH China and computerization shocks, as well as the robotization shock (Table B.5). Finally, we also calculate the Rotemberg weights for this estimation (Tables B.1, B.2).

4.2. Estimating μ , κ , and χ

We now estimate the parameters that govern the behavior of hours worked, the employment rate, and the degree of sectoral mobility.

¹⁵ This approximation builds on a similar derivation in GRY, and here we extend its application to a setting with automation, unemployment and an intensive margin for labor supply. The approximation is exact when $\kappa(1-\chi)\mu/(1+\mu) = 1$. Moreover, using the simulated model, we show that the approximation also performs extremely well for other parameter values (see Appendix Figure B.1).

4.2.1. Estimation equations

Estimating μ . From Eq. (3), we have that the change in average hours worked is $\hat{h}_{og} = (\hat{\Phi}_{og} \hat{\delta}_{og} / \hat{P}_o)^{1/\mu}$, while from Eq. (4) the change in the hourly wage is $\hat{i}_{og} = \hat{\Phi}_{og}$. By combining these results, we obtain the estimation equation for μ :

$$\ln \hat{h}_{og} = \alpha_\mu + \beta_\mu \ln \hat{i}_{og} + \varepsilon_{\mu,og}, \tag{20}$$

where $\beta_\mu \equiv \frac{1}{\mu}$, $\alpha_\mu \equiv -\ln \hat{P}_o^{1/\mu}$ and $\varepsilon_{\mu,og} \equiv \ln \hat{\delta}_{og}^{1/\mu}$. The estimated elasticity ($1/\mu$) governs how the supplied number of labor hours increases with the average hourly wage.

Estimating κ . By combining $\hat{i}_{og} = \hat{\Phi}_{og}$ with Eq. (14), we find that

$$\ln \hat{i}_{og} = \ln \hat{w}_{os} - \frac{1}{\kappa} \ln \hat{\pi}_{ogs} + \ln \hat{A}_{ogs}^{\frac{1}{\kappa}}. \tag{21}$$

Abstracting from local productivity shocks, this expression implies that $\hat{\pi}_{ogs}$ is a sufficient statistic for the change in the hourly wage relative to other groups. Moreover, this relation holds for any sector. To therefore exploit information from all sectors and thereby make this specification less sensitive to measurement in a single sector, we take an average of the previous specification across sectors:

$$\ln \hat{i}_{og} = \alpha_\kappa + \beta_\kappa \sum_s \omega_{\kappa,s} \ln \hat{\pi}_{ogs} + \varepsilon_{\kappa,og}, \tag{22}$$

where $\beta_\kappa \equiv -\frac{1}{\kappa}$, $\alpha_\kappa \equiv \sum_s \omega_{\kappa,s} \ln \hat{w}_{os}$, $\varepsilon_{\kappa,og} \equiv \sum_s \omega_{\kappa,s} \ln \hat{A}_{ogs}^{\frac{1}{\kappa}}$, and where the weights satisfy $0 \leq \omega_{\kappa,s} \leq 1$ and $\sum_s \omega_{\kappa,s} = 1$.¹⁶ Here, the regressor measures the change in the degree of specialization. Since $\sum_s \pi_{ogs} \hat{\pi}_{ogs} = 1$, a higher average $\ln \hat{\pi}_{ogs}$ implies that smaller sectors are expanding while larger sectors are contracting – a decrease in sectoral specialization. Then, the lower is κ , the more negative the impact of decreasing sectoral specialization on the hourly wage.

Estimating χ . From Eq. (1), we find that the change in average income per worker is

$$\hat{i}_{og} \hat{h}_{og} = \frac{\hat{I}_{og}}{\hat{e}_{og}} = \left(\frac{\hat{\delta}_{og}}{\hat{P}_o} \right)^{\frac{1}{\mu}} \hat{\Phi}_{og}^{\frac{1+\mu}{\mu}}.$$

Combining this expression with Eq. (16) and assuming $\hat{e}_{og} = 1$, we obtain an estimation equation for the employment rate elasticity (χ)

$$\ln \hat{e}_{og} = \alpha_\chi + \beta_\chi \ln(\hat{i}_{og} \hat{h}_{og}) + \varepsilon_\chi, \tag{23}$$

where $\beta_\chi \equiv \frac{\chi}{1-\chi}$, $\alpha_\chi \equiv \ln \hat{P}_o^{\chi/(\chi-1)}$ and $\varepsilon_\chi \equiv \frac{1}{1-\chi} \ln \hat{A}_{og}^M$. In the model, the employment rate e_{og} increases with average income per worker, and a higher χ entails that e_{og} is more responsive to such changes, which is reflected in the estimation equation.

IV strategy. In contrast to our shift-share estimation in Eq. (19), the model now implies that for the κ estimation in specification (22) the error term is necessarily correlated with the regressor. Moreover, OLS estimation of the intensive margin elasticity in (20) suffers from division bias (Borjas, 1980).¹⁷ We therefore cannot estimate these two specifications with OLS. However, employing an IV addresses the endogeneity concern for specification (22) and the division bias in (20). Finally, the IV strategy may also address endogeneity concerns not captured by the model.

As an instrument, we utilize our shift-share variable $\sum_s \pi_{gs} \hat{f}_s$.¹⁸ As explained by Goldsmith-Pinkham et al. (2020), a sufficient condition for this instrument to be valid is that the error terms are uncorrelated with the sector shares (π_{gs}). Importantly, the model also suggests this instrument to be highly relevant. It should therefore provide sufficiently strong first stages, in contrast to the ADH or Acemoglu and Restrepo (2020) shocks (see Appendix Table C.1).

4.2.2. Results

Table 2 presents the results of our parameter estimation. For each specification, we employ two versions of our shift-share instrument: in columns 1 and 2, we measure the sectoral shares π_{gs} based on hours worked, while in columns 3 and 4, we measure them based on sectoral income shares. We perform the estimation without and with controls for each version of the instrument in the odd and even columns, respectively. We cluster standard errors at the state level. For all three parameters, our estimates are always strongly statistically significant. Moreover, the first stage is generally sufficiently strong since all the F-statistics are above 10, with one exception where it is 8.5.

¹⁶ In our estimation, we weigh sectors by their employment shares (measured in terms of hours or income units), because percentage changes in employment for smaller sectors are more sensitive to measurement error, in particular at the CZ-level in the IPUMS data.

¹⁷ This division bias arises from calculating the hourly wage rate as the ratio of earnings over hours, such that the hours variable appears on both sides of the regression.

¹⁸ We focus on the instrument in levels instead of logs, to be able to apply the standard identification framework for shift-share instruments – see e.g. Goldsmith-Pinkham et al. (2020). However, the instruments in levels and logarithms have a correlation of 99%, so all our findings are robust to employing the instrument in logarithms instead.

Table 2
Parameter estimation.

(a) Estimation of $\frac{1}{\mu}$				
	(1)	(2)	(3)	(4)
	$\ln \hat{h}_g$	$\ln \hat{h}_g$	$\ln \hat{h}_g$	$\ln \hat{h}_g$
$\ln \hat{i}_g$	0.94*** (0.24)	0.32*** (0.10)	1.04*** (0.28)	0.40*** (0.14)
Implied μ	1.07	3.12	0.96	2.48
F-First Stage	18.1	15.6	14.1	8.53
Instrument	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$
Controls	No	Yes	No	Yes
Observations	722	722	722	722
(b) Estimation of $-\frac{1}{\kappa}$				
	(1)	(2)	(3)	(4)
	$\ln \hat{i}_g$	$\ln \hat{i}_g$	$\ln \hat{i}_g$	$\ln \hat{i}_g$
$\sum_s \pi_{gs}^{hours} \ln \hat{\pi}_{gs}^{hours}$	-0.44*** (0.11)	-0.48*** (0.096)		
$\sum_s r_s \ln \hat{\pi}_{gs}^{income}$			-0.36*** (0.097)	-0.35*** (0.091)
Implied κ	2.27	2.10	2.81	2.88
F-First Stage	71.4	94.3	53.5	34.2
Instrument	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$
Controls	No	Yes	No	Yes
Observations	722	722	722	722
(c) Estimation of $\frac{\chi}{1-\chi}$				
	(1)	(2)	(3)	(4)
	$\ln \hat{e}_g$	$\ln \hat{e}_g$	$\ln \hat{e}_g$	$\ln \hat{e}_g$
$\ln \hat{i}_g \hat{h}_g$	0.39*** (0.052)	0.20*** (0.053)	0.42*** (0.059)	0.27*** (0.079)
Implied χ	0.28	0.17	0.30	0.21
F-First Stage	39.4	16.9	33.2	11.0
Instrument	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{hours} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$	$\sum_s \pi_{gs}^{income} \hat{p}_s$
Controls	No	Yes	No	Yes
Observations	722	722	722	722

Panel (a) presents estimation results for specification (20), where i_g is the average hourly wage and h_g is the average annual supply of hours. Panel (b) estimates specification (22) and Panel (c) specification (23), where e_g is the employment rate and $i_g h_g$ is the average annual income. The even-numbered specifications include the following control variables from ADH: dummies for the nine Census divisions, the average offshorability index of occupations, and percentages of employment in manufacturing, college-educated population, foreign-born population, and employment among women, where these percentages are all measured at the start of the period. Standard errors, clustered at the state level, in parentheses. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For the elasticity of the intensive margin of labor supply ($1/\mu$), Chetty (2012) provides bounds on the estimates in the literature between 0.28 and 0.54.¹⁹ This interval covers our results for the specifications with controls (columns 2 and 4 in panel b). We set our preferred value to 0.4 – our estimate from column 4, with an implied inverse elasticity of $\mu = 2.5$. Our point estimates for κ fall within the interval of 2.1–2.88 (see Panel b). In somewhat different setups, Burstein et al. (2019), Hsieh et al. (2019) and GRY find estimates between 1.26 and 1.81. We therefore set our preferred value at $\kappa = 2.1$. Finally, for the employment elasticity χ , we obtain estimates between 0.17 and 0.3. These estimates are in line with the results from Shimer (2005), who estimates χ between 0.25 and 0.3, while Barnichon and Figura (2015) estimate $\chi = 0.33$.²⁰ We therefore set $\chi = 0.3$.

How do the distributional implications of our parameter values compare to the estimated shift-share elasticity from Table 1 (which informed us on how local exposure to national sectoral reallocation translates into local income changes)? Interestingly, with $\mu = 2.5$, $\kappa = 2.1$ and $\chi = 0.3$ we obtain a value for $\frac{1+\mu}{\kappa(1-\chi)\mu} = 0.95$. Since our estimates for this latter elasticity range from 0.94 to 1.23, our parameter values match the empirical elasticity almost perfectly.

In Appendix C, we perform a series of robustness checks: (i) estimation with benchmark instruments, which are too weak for our setting (Table C.1), (ii) estimating κ separately for each sector (Table C.2), obtaining a median κ estimate of 2.17, and (iii) documentation on the Rotemberg weights for each parameter estimate (Tables C.3, C.4, C.5).

¹⁹ As explained in Kim and Vogel (2021), our labor-leisure preferences have a zero income effect, such that the Marshallian labor supply elasticity equals the Hicksian one. Chetty (2012) focuses on the latter.

²⁰ Our setup is not identical to Shimer (2005) and Barnichon and Figura (2015), who employ a dynamic setting where only the currently unemployed are considered to be job applicants. In our static model, all workers are considered applicants.

5. Counterfactuals

5.1. Calibration

Parameter values. From the estimation we obtained values for the labor-side parameters, namely $\kappa = 2.1$, $\mu = 2.5$ and $\chi = 0.3$. We take values from the literature for ρ and θ . Specifically, for the elasticity of substitution between capital and labor we follow Karabarbounis and Neiman (2014) and set $\rho = 1.28$.²¹ We also set a common trade elasticity for all sectors: $\theta = 5$, as in Head and Mayer (2014).

We also calibrate ν , the parameter that regulates the change in productivity arising from an automation shock. To do so, we target the elasticity from Moll et al. (2022) for productivity increases associated with a certain decline in the labor share. Specifically, Moll et al. (2022) parametrize their model such that an automation-induced decline in the labor share of 13% is associated with a productivity increase of 2.4%. In our calibration, targeting this elasticity for the manufacturing sector implies a value of $\nu = -1.86$.^{22,23}

Shock calibration. With these parameter values, we jointly calibrate the automation shock and the China shock for the manufacturing subsectors.²⁴ First, we model the China shock as changes in China's sectoral productivity ($\hat{T}_{China,s}$). Leveraging Eq. (13), we calibrate these productivity shocks such that the model exactly matches the strong growth in Chinese manufacturing exports to the US ($\hat{\lambda}_{China,US,s}$).²⁵

Second, for the automation shock, recall that the model predicts that automation leads to a decline in sectoral labor shares ($\hat{\omega}_{US,s}$) – see Eq. (12). We therefore calibrate the changes in equipment productivity for US manufacturing sectors such that the model generates exactly the observed $\hat{\omega}_{US,s}$ in these sectors. Here, we categorize all labor-saving technological change under automation, including robotization, computer-driven automation of routine tasks, or other physical or intangible labor-saving technologies.²⁶

As indicated, the calibration targets for the shocks, namely the changes in the labor shares and the changes in the import shares from China, are also a function of equilibrium changes in factor prices (see Eqs. (12) and (13)). Since both the automation and the China shock affect these factor prices, both shocks can affect each of the calibration targets (labor shares and import shares from China). Hence, it is essential to calibrate these shocks jointly, which is what we do.

Fig. 1 details the two sets of targeted moments. With one exception, the US import shares for Chinese goods ($\hat{\lambda}_{China,US,s}$) more than double over the period 2000–2007. In fact, in our data, the average import share for manufacturing sectors increased from 1.02% in 2000 to 3.86% in 2007, which entails an increase of 278 percent over seven years. The textile sector, in particular, experienced a substantial increase in the import share, going from 3.8% in 2000 to 15.9% in 2007. Over the same period, the labor share declines by 8.3% on average in the manufacturing sectors. This pattern is not unique to the US since there is a correlation of 50% with the changes in the labor share in 11 European countries for which EU-KLEMS has detailed sector-level data. We return to this observation in Section 5.7. Such a high correlation is consistent with an automation shock common among these advanced economies. Note that there is also a correlation of 40% between the changes in the labor share and the changes in the import share from China.²⁷ Although these two moments are correlated, in our model, they have very different origins, as we will see in Section 5.4.1.

5.2. Impact of automation and the rise of China

The impact of automation and the rise of China is strongly positive for the US in the aggregate but widely unequal across US commuting zones (see Table 3). At the aggregate level, the US gains 3.96%, but the standard deviation across CZs is elevated

²¹ This value is in line with the recent estimate of $\rho = 1.35$ in Hubmer (2021).

²² Equation 45 in Appendix H explains how ν governs the relationship between output increases and automation-induced changes in the labor share.

²³ In addition to calibrating ν in the above manner, in Appendix H we also estimate it by employing an indirect inference procedure that exploits the link between output increases and automation-induced declines in the labor share. We obtain a point estimate of $\nu = -2.2$, with a standard error of 0.256. Hence, the 95% confidence interval of our estimate includes the value employed here ($\nu = -1.86$). Because the standard error on the indirect inference estimate is substantial and the quantitative results are sensitive to the precise ν value, we prefer to determine ν based on the benchmark productivity effect of automation in Moll et al. (2022).

²⁴ Throughout our simulation analysis, we employ the Alvarez and Lucas (2007) algorithm to find the counterfactual equilibrium. To ensure this algorithm is well behaved, following Ossa (2014) and GRY, we first purge trade deficits from the data using the original Dekle et al. (2008) exercise and then perform the entire analysis on the resulting data with balanced trade.

²⁵ Eq. (13) shows that an increase in $\hat{T}_{China,s}$ leads to an increase in import shares in all destination countries (including the US), taking into account changes in marginal costs. Our shock calibration matches observed import shares in a counterfactual general equilibrium. By design, it therefore incorporates changes in equilibrium prices and wages across all countries and sectors as well as in their marginal costs.

²⁶ It may not be a conservative modeling choice to map the full changes in the labor share into the automation shock, but we will see below that it provides a very good model fit for the cross-sector and the cross-CZ variation. Moreover, our calibrated value of ν , which governs the productivity increases associated with an automation-induced decline in the labor share, is corroborated by the observed relationship between reductions in the labor share and increases in output (see the ν estimation in Appendix H). Hence, our setup does not overstate the productivity gains associated with the observed declines in the labor share. Finally, in Section 5.7, we examine the robustness of our results to targeting the common trend in the labor shares across the US and 11 European countries instead of only the labor share changes in the US. Importantly, the common trends in the labor shares are arguably driven by common technological developments.

²⁷ Interestingly, due to the joint calibration of the China and automation shock, this pattern results in a correlation of 50% between the $\hat{T}_{China,s}$ and $\hat{\xi}_{US,s}$ vectors. Based on the positive correlation between $\hat{\omega}_{US,s}$ and $\hat{\lambda}_{China,US,s}$, and since there is a negative relation between $\hat{\xi}_{US,s}$ and $\hat{\omega}_{US,s}$, we might have expected a negative correlation between the calibrated shocks $\hat{T}_{China,s}$ and $\hat{\xi}_{US,s}$. However, the strong US automation shock in certain sectors, particularly in the petroleum refining sector, requires stronger Chinese productivity growth to match its $\hat{\lambda}_{China,US,s}$ target. This in turn leads to the positive correlation between $\hat{T}_{China,s}$ and $\hat{\xi}_{US,s}$ across manufacturing sectors.




Fig. 1. Changes in import shares from China and labor shares in US manufacturing. For manufacturing sectors over the period 2000–2007, the figure displays $\hat{\lambda}_{China,US,s}$ and $\hat{\omega}_{US,s}$, which are the targeted moments in our calibration. The correlation between the two variables is 40%.

Table 3
Impact of automation and the rise of China across commuting zones.

	Aggregate	Mean	SD	Min.	Max.
\hat{I}_g/\hat{P}	3.96	4.43	0.97	1.24	8.34
\hat{i}_g/\hat{P}	1.96	2.19	0.48	0.62	4.09
\hat{h}_g	0.78	0.87	0.19	0.25	1.61
\hat{e}_g	1.17	1.31	0.28	0.37	2.43
$\Delta\pi_{gM}$	-0.81	-0.82	0.39	-2.57	-0.02

The table shows the impact of automation and the rise of China across US commuting zones. The first row displays the change in average real income, the second on the average hourly wage, the third row on hours worked per employee and the fourth on the employment rate. The final row shows the change in the share of employment in manufacturing. The first column shows the aggregate effect, and the second the average. The third column shows the standard deviation across commuting zones and the fourth and fifth column respectively show the minimum and maximum effect. All variables are measured in percentage changes, except $\Delta\pi_{gM}$ which is measured in percentage points because $\hat{\pi}_{gM}$ is a very noisy measure in our data, especially for low initial π_{gM} (see Appendix Figure D.2).

at 0.97%. Some CZs gain up to 8.34% while for others this is limited to 1.24%.²⁸ The strongest concentration of low gains is concentrated around the Rust Belt (see Fig. 2). More broadly, the non-coastal regions east of the Mississippi river - or the “Eastern Heartland” in the phrasing of Austin et al. (2018) - mostly experience low gains due to these combined shocks.

Recall that changes in the hourly wage drive changes in a CZ’s earnings, hours worked, and the employment rate. Moreover, if there are only national-level shocks, then changes in a group’s total real income (\hat{I}_g/\hat{P}) are perfectly correlated with changes in the real hourly wage, hours worked, and the employment rate. Specifically, we then have that

$$\frac{\hat{i}_g}{\hat{P}} = \left(\frac{\hat{I}_g}{\hat{P}}\right)^{\frac{(1-\chi)\mu}{1+\mu}}, \quad \hat{h}_g = \left(\frac{\hat{I}_g}{\hat{P}}\right)^{\frac{1-\chi}{1+\mu}}, \quad \hat{e}_g = \left(\frac{\hat{I}_g}{\hat{P}}\right)^{\chi}. \tag{24}$$

Given our parameter values, this implies that the hourly wage accounts for 50% of the log change in a group’s real income, while hours worked and the employment rate respectively account for 20% and 30% of that change. For the US in the aggregate, this results in an increase of the real hourly wage by 1.96%, with a standard deviation of 0.48%, and increases in hours worked and the employment rate by 0.78% and 1.17% respectively (see Table 3). Since the trade and automation shocks occur in manufacturing – a sector that only accounts for 16% of employment in the year 2000, these numbers are substantial.

²⁸ When we employ alternative values for θ_s in the simulation, the aggregate gains increase to 4.45% and the standard deviation to 1.76% (please see Appendix Table E.1, Panel a). There, we follow Bartelme et al. (2019) and employ the median value for θ_s from the estimations in Caliendo and Parro (2015), Shapiro (2016), Bagwell et al. (2018) and Giri et al. (2018).




Fig. 2. Geographical impact of the automation and China shock.

The figure plots the geographic distribution of the percentage changes in groups' average real income for US commuting zones due to the combined effect of automation and the China shock.

The change in commuting zones' manufacturing sector ($\Delta\pi_{gM}$) has a correlation of 74% with the income changes.²⁹ As a result of this high correlation, the manufacturing sector's decline also tends to be concentrated in the Eastern Heartland in general and in the Rust Belt in particular (see Appendix Figure D.3). We find that the aggregate decline in the manufacturing sector is 0.81 percentage points, with a standard deviation of 0.39 and a maximum decline of 2.57 percentage points.

5.3. Examining model fit

For all the different margins of adjustment discussed so far – income,³⁰ hourly wage, hours worked, employment rate and employment share in manufacturing – our model predictions fit well with the observed variation across commuting zones (see Appendix Fig. A.3).³¹ To examine the model fit more formally, we regress the observed changes in the data on the model's predicted changes after the combined China and automation shocks (see Table 4). We run this regression with and without control variables for each of the variables of interest. We find that there is a strongly significant positive relationship for all specifications between the model's predictions and the actual changes. To see how much of the variation in the observed changes the model can explain on its own, we focus first on the specifications without controls. There, we explain 28% or more of the variation for changes in hours worked, the employment rate, or the manufacturing share. For average income, the R^2 is 8%, and 4% for changes in the hourly wage. Part of the difference in these R^2 values arises from a stronger presence of outliers for the income variables, which may be due to larger measurement error (see Appendix Fig. A.3, panels a and b).

In addition to examining the share of the observed variation explained by the model, we ask if the magnitude of the model-implied changes is in line with the observed changes. This is the case if the estimated regression coefficient is around unity.³² If the coefficient is larger, the model underpredicts the observed changes because a given change in the model translates into a larger change in the data. Once we add control variables,³³ we cannot reject a unity coefficient for CZ-level income variables ($\ln \hat{I}_g$). However, the degree of underprediction is stronger and significant for the other variables. This might indicate that we have been too conservative in setting the intensive and extensive margin elasticities and that – if anything – we are underestimating the distributional impact of the shocks. Overall, our model with the combined China and automation shock fits the data quite well, regarding the share of the variation it explains and the magnitude of the predicted values.

²⁹ While we measure the other variables in Table 3 in percentage terms, we measure the change in the manufacturing employment share in percentage points ($\Delta\pi_{gM}$). This is because the variable in percentage terms ($\hat{\pi}_{gM}$) is extremely sensitive to the initial π_{gM} , while $\Delta\pi_{gM}$ is not (see Appendix Figure D.2). The wide variance in $\hat{\pi}_{gM}$ is probably due to measurement error in IPUMS, in particular for low initial values of π_{gM} .

³⁰ Since all CZs share the same price index, we focus on variation in nominal income in our model fit analysis.

³¹ Throughout the model fit exercise, we use our WIOD-adjusted measure for changes in the manufacturing share, since this is the measure we need to employ in the simulations (see end of Section 3). This measure has a correlation of 82% with the raw measure from IPUMS.

³² See also the model-fit discussion in Adão et al. (2020).

³³ These control variables capture shocks associated with regional fixed effects, the demographic composition of a CZ, the potential for offshorability of the local jobs, or the secular decline in manufacturing.

Table 4
Model fit of variation across commuting zones.

	$\ln \hat{I}_g$		$\ln \hat{l}_g$		$\ln \hat{h}_g$		$\ln \hat{e}_g$		$\Delta \pi_{gM}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\ln \hat{I}_g$ - Both shocks	2.81 (0.60)	1.81 (0.79)								
$\ln \hat{l}_g$ - Both shocks			1.98 (0.77)	2.81 (1.24)						
$\ln \hat{h}_g$ - Both shocks					8.69 (0.90)	6.95 (1.08)				
$\ln \hat{e}_g$ - Both shocks							4.78 (0.72)	2.95 (0.82)		
$\Delta \pi_{gM}$ - Both shocks									3.62 (0.31)	2.16 (0.33)
R^2	0.08	0.29	0.04	0.20	0.37	0.44	0.29	0.37	0.28	0.43
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	722	722	722	722	722	722	722	722	722	722

The specifications in this table regress observed changes in the data for the period 2000–2007 on the model's predicted changes after the combined China and automation shock. The first two specifications examine average income, specifications (3) and (4) average hourly wage, specifications (5) and (6) hours worked per employee, specifications (7) and (8) the employment rate, and (9) and (10) the manufacturing employment share. We measure $\Delta \pi_{gM}$ in percentage points because $\hat{\pi}_{gM}$ is a very noisy measure in our data, especially for low initial π_{gM} (see Appendix Figure D.2). The even-numbered specifications include the following controls from ADH: dummies for the nine Census divisions, percentage of employment in manufacturing, percentage of college-educated population, percentage of foreign-born population, percentage of employment among women, and the average offshorability index of occupations, where these percentages are all measured at the start of the period. Standard errors, clustered at the state level, in parentheses.

In addition to the observed variation for labor market outcomes across CZs, the model also explains a large part of the aggregate decline in manufacturing employment. As is well known, during the period 2000–2007, the share of manufacturing employment in the US declined exceptionally fast, falling from 16.4% to 13% in our data.³⁴ In our model, the combined China and automation shock leads to a decline by 0.81 percentage points in $\pi_{U,S,M}$ (see Table 3), which means that our model explains 24% of the total decline in manufacturing employment. Other trends in demand and supply then account for the further erosion of manufacturing employment during this period.³⁵

5.4. Comparing the automation and China shocks

We now examine how the individual China and automation shocks contribute to the combined shock. From the calibration of the joint shock, we have obtained the values for productivity changes in China ($\hat{T}_{China,s}$) and on automation shocks in the US ($\hat{\xi}_{U,S,s}$). To simulate the counterfactual impact of the individual China shock and the individual automation shock, we therefore separately shock our model with these previously calibrated values for ($\hat{T}_{China,s}$) and ($\hat{\xi}_{U,S,s}$) respectively.

5.4.1. Comparison of the sector-specific shocks

Through the lens of our model, increased Chinese import penetration and changes in sectors' labor shares clearly have different origins (see Fig. 3). For the individual China shock, the model completely fails to match the changes in the labor shares (Panel a). Analogously, in response to the individual automation shock, import shares from China tend to fall – sometimes substantially – instead of increase (Panel b).

Of course, it may be highly intuitive that US automation does not lead to increased import shares from China and that Chinese technological progress does not induce drastic changes in US sectoral labor shares. In that case, this only emphasizes the need to account for the occurrence of *both* shocks during the period we study. Moreover, the finding that each shock tends to exert downward pressure on the primarily targeted moment for the other shock highlights the importance of *jointly* calibrating the shocks. Specifically, the China shock tends to increase the labor share, while the automation shock tends to lower imports from China. Hence, ignoring the impact of one shock in the calibration of the other shock leads to a type of omitted variable bias, which we overcome in the joint calibration.

Taking both shocks into account also matters for matching the changes in manufacturing subsectors' value added. Specifically, in a regression of the actual on the predicted changes in value-added, the combined shock yields an R^2 of 74%, which is substantially higher than the values of 34% or 57% for the individual China or automation shocks, respectively (see Fig. 4). To understand why,

³⁴ The measured decline of 3.44 percentage points in the model-based measure is closely in line with the 3.78 percentage point decline in the raw IPUMS data. The difference in magnitude is due to adjusting the data to the WIOD value-added data, which is necessary to make the simulation data internally consistent (see Section 3).

³⁵ For instance, consistent with the literature on structural change (e.g. Ngai and Pissarides (2007) and Kehoe et al. (2018)), there has been a substantial downward trend in the consumption share of manufacturing since US final demand for manufacturing falls by two percentage points in our data. Moreover, the increase in global offshoring also negatively affects US manufacturing employment (see e.g. Ebenstein et al. (2014), Feenstra (2017) and Fort (2017)). Finally, the increase in the US trade deficit also puts downward pressure on manufacturing employment. However, as explained in footnote 24, we remove the impact of trade deficits in our analysis by employing data purged from trade deficits.




Fig. 3. Changes in the labor share and imports from China for the individual shocks. In both panels, the blue circles show the actual values for changes in US import shares from China ($\hat{\lambda}_{China,US,s}$) and the changes in the labor shares ($\hat{\omega}_{US,s}$) for US manufacturing subsectors, while the orange squares depict the simulated values after the China shock (Panel a) or the automation shock (Panel b).

Table 5
Impact of the individual shocks on real income across US commuting zones.

	Aggregate	Mean	SD	Min.	Max.	$\Delta\pi_{US,M}$
Only China Shock	0.86	1.47	0.80	-1.71	3.58	-0.60
Only Automation shock	2.90	2.81	0.51	1.44	5.22	-0.28
China and Automation Shock	3.96	4.43	0.97	1.24	8.34	-0.81

The table shows the impact of the individual China shock in the first row, the individual automation shock in the second row, and the combined China and automation shock in the third row. The first four columns display statistics for the changes in groups' real income, with the first column showing the aggregate change, the second the average change, the third the standard deviation, the fourth the minimum, and the fifth the maximum change. All these changes in real income are reported as percentage changes. The final column lists the change in the aggregate US employment share in manufacturing in percentage points. Appendix Table D.1 provides the full breakdown of the different margins of adjustment.

first observe that in reality, roughly half of the manufacturing subsectors experience a decline in value-added, while the other half experience an increase. The individual shocks are unable to match this mixed pattern; the China shock predicts a general decline in manufacturing value-added, while the automation shock predicts an overall increase. Importantly though, when we combine the opposite impact of the individual shocks in the joint shock, the model matches the mixed pattern well (Fig. 4, Panel c) and yields a substantially higher R^2 .

5.4.2. Aggregate and distributional impact of the individual shocks

Turning to the effect of the shocks on aggregate US real income, we find that the impact of the individual automation shock is more than three times as large as that of the individual China shock, namely 2.90% versus 0.86% (see Table 5).^{36,37} Interestingly, we also find that the impact of the combined shock is 0.20 percentage points larger than the sum of the impact of the individual shocks.

Even though the China shock yields a lower aggregate gain than the automation shock, it leads to stronger distributional effects and a larger decline in manufacturing employment. Specifically, the China shock generates a standard deviation of 0.8 percent in gains across CZs and a drop in manufacturing employment of 0.6 percentage points. This compares to a standard deviation for the gains of 0.51 percent and a manufacturing decline of 0.28 for the automation shock.³⁸

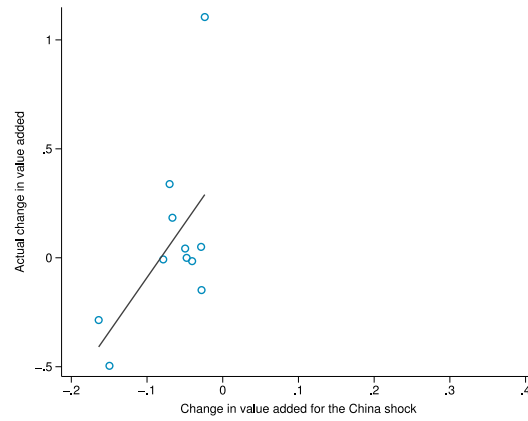
We also find that the distributional impact of the combined shock (measured in variance) is slightly larger than the sum of the effect of the individual shocks. This is due to the income changes for the two shocks being positively correlated (at 8.7% - see

³⁶ Appendix Table D.1 provides the full breakdown for the different margins of adjustment. As explained in Eq. (24), the hourly wage accounts for 50% of the log change in a group's real income, hours worked for 20%, and employment for 30%.


³⁷ When we employ alternative values for θ_s in the simulation, the gains from automation increase to 3.64% and those from the China shock fall to 0.55% (see Appendix Table E.1). We employ the median value for θ_s from some prominent estimates in the literature, as in Bartelme et al. (2019).

³⁸ Interestingly, the model-predicted income changes due to these two shocks are strongly correlated with closely related reduced-form measures. Specifically, the predicted income changes after the China shock have a correlation of -50% with the ADH China shock instrument, while the predicted income changes due to automation have a correlation of -24% with the Acemoglu and Restrepo (2020) robotization shock.

(a) Only the China shock, $R^2 = 34\%$



(b) Only the automation shock, $R^2 = 57\%$



(c) Combined China and automation shock, $R^2 = 74\%$




Fig. 4. Predicted changes in manufacturing value added for the different shocks. The figure shows the percentage change in value added ($\hat{V}_{US,t} - 1$) for US manufacturing sectors for each of the shocks listed. The horizontal axis shows the model's predicted changes, while the vertical axis shows the actual changes..

Appendix Figure D.1), which implies the variance of the combined shock is larger than the sum of the variance of the individual shocks.³⁹

Given the positive but imperfect correlation of the two shocks, the geographic incidence of the China shock and the automation shock is related but far from identical. For instance, the Midwest tends to receive low gains for both shocks. At the same time, the central and southern Appalachians have low gains for the China shock but relatively more positive effects for automation (see Appendix Figures D.5 and D.6). These latter regions have relatively high employment shares in the textile sector – the sector most negatively affected by the China shock – but very little employment in the sectors hit by automation (coke and refined petroleum products). Fig. A.2 has the details on the sector-level impact of each shock.

5.4.3. Model fit across commuting zones for the different shocks

In Table 6, we run the model fit for the same variables as before in Table 4, but now for the individual shocks. Specifically, in the odd columns, we regress the observed values on the predicted changes for both the individual China and the individual automation shock. For comparison, we repeat the analysis from Table 4 for the combined shock in the even columns.

For the specifications with the predictions from both individual shocks, all coefficient values are positive and, with few exceptions, also statistically significant. Hence, just as for the combined shock, also for the individual shocks, there is a positive and usually significant correlation between the observed values in the data and the predictions from the model. This corroborates our model predictions for the individual shocks.

In turn, these significant correlations are reflected in substantial R^2 values for the specifications with the two shock predictions, implying that the model is empirically relevant for explaining the observed variation. More precisely, the R^2 for these specifications tends to be slightly higher than the R^2 for the combined shock. For instance, for CZ income (\hat{I}_g) or the employment rate (\hat{e}_g), the R^2 for the specification with both shocks is 12% and 31% respectively, while it is 8% and 29% for the specification with the combined shock.

Recall that the regression coefficient should be equal to unity for the model to match the magnitude of the variance in the data. There are several cases where we cannot reject this null hypothesis, e.g. for both shock's predictions on the hourly wage (\hat{w}_g), for the China shock's predictions on CZ income, or for automation's predictions on the change in manufacturing employment ($\Delta\pi_{gM}$).

In the other cases, there is some significant underprediction. In part, this could be because the model is understating certain aspects of the labor market impacts of the shock.^{40,41} Alternatively, the model predictions for the different shocks may be correlated with other trends in the economy. When we address this latter issue by introducing controls in our model fit regressions (see Appendix Table D.2), as we also did in Table 4, we indeed notice again that the degree of underprediction falls. In several cases, the coefficient becomes insignificantly different from unity, but not generally so. At the same time, some of the coefficients become insignificant, for instance, the predicted effect of the China shock on CZ income. This is due to the strong correlation of these predictions with CZs' initial manufacturing share. This collinearity problem would raise estimation challenges in a reduced-form setup (see e.g. the estimation results on wage growth in Borusyak et al. (2020)), but does not imply that the China shock has no impact on CZs' income. Indeed, the benefit of our quantitative model is precisely that we can examine the China shock, its impact, and its relation to other shocks without relying only on reduced-form correlations and in isolation from any potential confounds.

Notice also that the predictions for the combined shock always remain statistically significant, in contrast to the predictions for the individual shocks. This is due to the smaller standard errors in case of the combined shock, which arise from the larger variance in the predictions for the joint shock. Therefore, analyzing the impact of the joint shock has the benefit of exploiting all the variation it generates and not being limited to the conditional variation for one shock, holding the impact of other shocks constant.

5.4.4. Comparison to reduced-form analyses of trade and automation

As indicated above, our quantitative analysis is inspired by the seminal work on the China shock by ADH and Pierce and Schott (2020), and on robotization by Acemoglu and Restrepo (2020); i.e. the seminal papers that have put the distributional effects of trade and automation across local labor markets on the agenda. In this section, we first ask what the potential benefits are from analyzing these shocks with our general-equilibrium (GE) framework, and afterwards how our findings compare to those influential reduced-form findings.

The first benefit of our GE approach is that it allows for aggregation. Where the above shift-share analyses can only estimate relative effects across CZs, our approach analyzes both the aggregate and the distributional effects across local labor markets. Second, the estimates in our framework have a clear interpretation and are driven by transparent model-based shocks. This way, we also side-step complex identification issues, such as disentangling Chinese supply from US demand shocks as sources for the observed

³⁹ The correlation in the income changes for the individual China and automation shocks is driven by the size of the manufacturing sector in a CZ. Specifically, when we regress the income change due to one shock on the income change due to the other shock, we obtain a significantly positive coefficient. However, this coefficient turns negative, and significantly so, when we control for the size of a CZ's manufacturing sector.

⁴⁰ For instance, Rodríguez-Clare et al. (2022) argue that downward nominal wage rigidity is important for understanding the labor market impact of the China shock, which is in line with our model underpredicting the impact of the China shock on the employment rate and the decline in manufacturing employment. We view incorporating nominal rigidities into the analysis of the joint equilibrium impact of trade and automation shocks as an excellent topic for further research.

⁴¹ There is a strong association between the explanatory power (R^2) of the individual shock (observed in Appendix Table D.3), and its degree of underprediction. This is because a higher correlation, or more precisely the higher covariance, between the model predictions and the observed changes, increases the value of the coefficient estimate. Interestingly, still in Table D.3, the predictions from the combined shock yield a similar R^2 value as the most predictive individual shock but are associated with a lower degree of underprediction. This is because the predictions for the combined shock exhibit substantially larger variance than those of the individual shock, which lowers the estimated coefficient value.

Table 6
Model fit of the separate shocks.

	$\ln \hat{l}_g$		$\ln \hat{i}_g$		$\ln \hat{h}_g$		$\ln \hat{e}_g$		$\Delta\pi_{gM}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\ln \hat{l}_g$ - China	1.14									
	(0.85)									
$\ln \hat{l}_g$ - Automation	6.08									
	(1.39)									
$\ln \hat{l}_g$ - Both shocks		2.81								
		(0.60)								
$\ln \hat{i}_g$ - China			1.62							
			(0.75)							
$\ln \hat{i}_g$ - Automation			2.58							
			(1.48)							
$\ln \hat{i}_g$ - Both shocks				1.98						
				(0.77)						
$\ln \hat{h}_g$ - China					10.01					
					(1.01)					
$\ln \hat{h}_g$ - Automation					5.25					
					(1.26)					
$\ln \hat{h}_g$ - Both shocks						8.69				
						(0.90)				
$\ln \hat{e}_g$ - China							5.26			
							(0.79)			
$\ln \hat{e}_g$ - Automation							3.41			
							(1.10)			
$\ln \hat{e}_g$ - Both shocks								4.78		
								(0.72)		
$\Delta\pi_{gM}$ - China									4.55	
									(0.44)	
$\Delta\pi_{gM}$ - Automation									1.89	
									(0.51)	
$\Delta\pi_{gM}$ - Both shocks										3.62
										(0.31)
R^2	0.12	0.08	0.04	0.04	0.41	0.37	0.31	0.29	0.31	0.28
Controls	No	No	No	No	No	No	No	No	No	No
Observations	722	722	722	722	722	722	722	722	722	722

The specifications in this table regress observed changes in CZs' labor market outcomes in the data for the period 2000–2007 on the model's predicted changes for the different listed shocks. The odd columns regress the observed values on the predicted changes for both the individual China and the individual automation shock. The even columns repeat the analysis from Table 4 for the combined shock. The first two specifications examine average income, specifications 3 and 4 average hourly wage, specifications (5) and (6) hours worked per employee, specifications (7) and (8) the employment rate, and (9) and (10) the manufacturing employment share. We measure $\Delta\pi_{gM}$ in percentage points because $\hat{\pi}_{gM}$ is a very noisy measure in our data, especially for low initial π_{gM} (see Appendix Figure D.2). Standard errors, clustered at the state level, in parentheses.

growth in Chinese import penetration or distinguishing exogenous changes in automation technology from endogenous technology adoption. While the above-referenced reduced-form studies carefully address these endogeneity concerns, it remains useful to have a quantitative framework available where these identification issues are side-stepped.

An additional advantage is that our framework estimates the full, unconditional effects of trade and technology shocks. In contrast, the reduced-form analyses necessarily focus on the effects of one shock, holding the other shock constant and conditioning on a further set of control variables. As is clear from above, our analysis can estimate the unconditional effects of both the joint and the individual trade and technology shocks, which makes it substantially more versatile than the reduced-form approaches. Relatedly, in some reduced-form analyses, collinearity issues can arise from including all required control variables for proper identification, e.g. when the shock is strongly correlated with the manufacturing employment share. This challenge does not arise in our setting.

Furthermore, our framework incorporates how shocks to one sector affect other sectors through sectoral labor reallocation and input–output linkages. While reduced-form analyses can attempt to approximate these effects, they may not succeed in capturing the full equilibrium cross-sectoral effects. To illustrate this, in our model, we construct an ADH-style exposure term to increased Chinese import penetration in the US after the joint China and automation shock. We then regress the model's predicted income effects for the individual China shock on this ADH-style exposure term and find that the exposure term explains 44% of the variation in the income effects due to the China shock. While this is substantial, it still means that 56% of the variation is not captured by this reduced-form measure. When we plot the model-based ADH exposure term on a map (see Appendix Figure D.7), we indeed notice a strong correlation with the impact of the individual China shock in our model (see Appendix Figure D.5). However, the latter often finds stronger negative effects in the Great Lakes Region (e.g. Ohio and upstate New York). An important part of the explanation for these differences, is that our model takes into account how the China shock to manufacturing differentially affects the various non-manufacturing sectors.

To sum up, conditional on the validity of our model, our quantitative GE analysis is more comprehensive, more precise, and more versatile than a reduced-form approach. Qualitatively though, our findings are closely in line with those in ADH and Acemoglu

Table 7
Counterfactual results on real income across US CZs for $\rho = 0.72$.

	Aggregate	Mean	SD	Min.	Max.	$\Delta\pi_{U,S,M}$
Only China Shock	0.86	1.35	0.78	-1.89	3.18	-0.59
Only Automation shock	3.51	3.76	1.04	-3.13	7.88	-0.77
China and Automation Shock	4.35	5.12	1.28	-1.74	8.07	-1.36

For the model with $\rho = 0.72$, the table shows the impact of the individual China shock in the first row, the individual automation shock in the second row, and the combined China and automation shock in the third row. The first four columns display statistics for the changes in groups' real income, with the first column showing the aggregate change, the second the average change, the third the standard deviation, the fourth the minimum, and the fifth the maximum change. All these changes in real income are reported as percentage changes. The final column lists the change in the aggregate US employment share in manufacturing in percentage points.

and Restrepo (2020), who both find that their respective shocks lead to relative declines in income and employment. Moreover, analogous to Acemoglu and Restrepo (2020), we also find that the distributional effects of the China shock are stronger than those of robotization (see their footnote 27).

Our automation shock is less comparable to the computerization shock in Autor et al. (2015). While this technology shock to routine-task intensive occupations is a type of automation, in manufacturing it primarily occurred in the 1980s and 1990s, i.e. before the 2000–2007 period that we are studying.⁴² In addition, the main impact of the computerization shock is occupational polarization: employment mainly declines in routine-task intensive occupations while it stays stable or increases in other occupations. Our current framework is not set up to study the joint impact of trade and automation on occupational polarization, which is a great topic for further research.

5.5. Alternative elasticity of substitution between labor and equipment

There is debate in the macroeconomic literature on how elastic substitution between labor and equipment is; while Karabarbounis and Neiman (2014) and Hubmer (2021) estimate values for the elasticity of substitution (ρ) around 1.3, Oberfield and Raval (2021) argue that ρ is below unity, with a preferred value of $\rho = 0.72$. In our baseline analysis, we followed the former studies by setting $\rho = 1.28$, but here we examine the robustness of our results by setting $\rho = 0.72$. Due to this substantial fall in ρ , as explained in Appendix Section H.1,⁴³ the effect of automation shocks on productivity would substantially decline – other things equal. However, to ensure that the elasticity of productivity gains associated with automation-induced declines in the labor share remains at the same level as in Moll et al. (2022) – as in our baseline analysis, we now set $\nu = 2.4$.⁴⁴

The effects of the combined shock under this scenario are quite close to our baseline results (see bottom row of Table 7). The aggregate gain in real income is now 4.35% and the standard deviation in the gains is 1.28%, compared to 3.96% and 0.97% in the baseline. In addition, the correlation in the welfare effects across the two versions of the model is substantial, at 41.5% (see Appendix Figure E.1). As a result, the model fit results are also fairly similar, as discussed in Appendix E.1.

5.6. Heterogeneity across education groups

In this section, we examine the heterogeneous effects of trade and automation on workers with different education levels. To this end, we split each CZ into two groups, consisting of workers with or without some college education.

Estimation. Before turning to our counterfactual analysis, we re-estimate the labor-side parameters of the model for each education group. Specifically, we proceed as in the estimation in Table 2, but we estimate each specification separately for workers without or with some college education. Appendix Table E.3 has the results. There, we typically do not find strong patterns of heterogeneity across education groups, and the confidence intervals of the estimates usually overlap. As preferred values for μ and χ , we take the average estimate of the two specifications that include controls but have different instrument construction. This results in the following values for non-college workers: $\mu_{NC} = 2.845$; $\chi_{NC} = 0.185$; and these for college workers: $\mu_{CO} = 3.925$; $\chi_{CO} = 0.172$. Hence, these two worker types exhibit a highly similar employment elasticity (χ), while non-college workers have a more elastic intensive margin labor supply elasticity ($1/\mu$) than college workers. Finally, we obtain that $\kappa_{NC} = 1.445$ and $\kappa_{CO} = 2.105$.

⁴² Autor et al. (2013b) argue that their China and technology shocks are largely uncorrelated, which may create tension with our finding that the impact of the trade and automation shocks has a correlation of 8.7%. Our automation shock being distinct from the computerization shock resolves this tension. Moreover, Autor et al. (2013b) obtain a correlation close to zero after weighting CZs by their population. The unweighted correlation between their trade and technology exposure terms is 31%.

⁴³ Specifically, there we show that:

$$d \ln Y_{s|M_i, Z_i, K_i, N_{ki}} = \left[\nu + \frac{(1 - \omega_s)}{\rho - 1} \right] \alpha_s d \ln \xi_s,$$

where the second term within the brackets switches sign at $\rho = 1$.

⁴⁴ Our value of $\nu = 2.4$, implied by the Moll et al. (2022) relation between productivity gains and automation, is close to the value that we obtain in our indirect inference estimation, namely $\nu = 2.34$, with a standard error of 0.26 (see Appendix Section H.3).

Table 8
Heterogeneity across education groups for the combined shock.

	Aggregate	Mean	SD	Min.	Max.	$\Delta\pi_{U.S.M}$
All groups	3.10	3.56	0.94	-0.43	7.53	-0.78
Non-college workers	3.12	3.58	1.16	-0.43	7.53	-0.77
College workers	3.10	3.53	0.65	1.09	6.58	-0.78

The table shows the impact of the combined China and automation shock for the model with groups defined by commuting zone and education level (some college education or not). The first row shows the effect of the shock on all groups in the top row, on the groups where workers have no college education in the middle row, and on groups with college education in the bottom row. The first four columns display statistics for the changes in groups' real income, with the first column showing the aggregate change, the second the average change, the third the standard deviation, the fourth the minimum, and the fifth the maximum change. All these changes in real income are reported as percentage changes. The final column lists the change in the aggregate US employment share in manufacturing in percentage points.

Welfare effects. For the counterfactual exercise, aside from the estimated parameter values mentioned above, we return to our baseline calibration with $\rho = 1.28$, and $\nu = -1.86$. Table 8 then reports the results for the joint shock. Compared to our baseline results, we now obtain aggregate gains for the combined shock that are 0.86 percentage points lower, at 3.1%. This is mainly due to the lower values for the extensive and intensive margin elasticities (χ and $1/\mu$), which entails lower amplification effects of the shocks.

We do not find substantial heterogeneity in terms of the aggregate effects of the joint shock on non-college versus college workers (see Table 8). However, for college workers, the distributional effects are smaller than for non-college workers, with respective standard deviations in the welfare effects of 0.65% and 1.16%. A first reason is that college workers have a higher reallocation elasticity. Second, since a group's exposure to the joint shock depends mainly on the size of its overall manufacturing sector, this difference in the distributional effects is also due to non-college workers exhibiting a higher standard deviation in their initial share of manufacturing employment across CZs (10.3% versus 6.6%). Specifically, 22.6% of the non-college groups have at least 30% of their workforce employed in manufacturing, compared to 1.5% of college groups for which this is the case.

Turning to the individual shocks, we find that college workers gain slightly more from the automation shock (2.33% versus 2.14% for non-college workers; see Appendix Table E.4). That is because they tend to be less employed in manufacturing, where workers tend to be negatively exposed to automation, and more in services, where employees reap the benefits from automation in manufacturing due to consumer gains. In contrast, college workers gain somewhat less from the China shock (0.60% vs. 0.83% for non-college workers; see Appendix Table E.5). This is because college workers are less specialized in the primary sectors, which gain substantially from the China shock due to increased export demand in these sectors.

One reason our model does not generate strongly different effects for college versus non-college workers, is that here, automation is not biased toward a particular education group. To examine education-biased automation in our setup, one could introduce a CES production function with college and non-college labor as imperfect substitutes (as in Section 7.2 in GRY), where equipment is a stronger substitute for one of the education groups. We leave this exercise for further research.

5.7. Alternative shock calibration

So far, we have calibrated the China shock and the automation shock by respectively matching changes in US import shares from China and changes in US labor shares. Since neither increased imports from China nor automation are unique to the US, in this section, we can examine the robustness of our calibration by focusing on common trends across advanced economies. To this end, we return to our baseline model with a single group per CZ and $\rho = 1.28$, and we perform this analysis in Appendix Section E.3. There, we continue to find that the China shock has the lowest aggregate and the strongest distributional effects.

6. Conclusion

This paper presents a gravity model to examine the aggregate and distributional effects of sector-specific trade and automation shocks on US commuting zones. After transparently estimating the key labor-supply elasticities, the model predictions for the calibrated China and automation shocks match well with both the observed labor market outcomes across CZs and the pattern of value-added growth across manufacturing subsectors. This contrasts with the model fit for the individual China or automation shocks, which generate a mismatch for the pattern of value-added growth in manufacturing. Moreover, since the China and automation shocks are correlated, the individual shocks underpredict the magnitude of the observed variation in income changes across CZs. Taken together, our model provides a more comprehensive understanding of trends in inequality across local labor markets compared to existing quantitative papers that focus on a single shock.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

<https://data.mendeley.com/datasets/jj6kdtcbvy/1>.

Appendix A. Supplementary tables and figures

List of Online Appendices:

- Appendix B. Shift-share approximation (supplementary estimation results)
- Appendix C. Parameter estimation (supplementary results)
- Appendix D. Counterfactual analysis (supplementary results for baseline model)
- Appendix E. Counterfactuals: sensitivity and heterogeneity
 - E.1 Counterfactual results for the model with $\rho = 0.72$
 - E.2 Heterogeneity across education groups
 - E.3 Alternative shock calibration
- Appendix F. Measurement of $\alpha_{os}\omega_{os}, \alpha_{os}^*, \omega_{os}^*$
- Appendix G. Theory
 - G.1 Labor Supply
 - G.2 System of Hat Equations
 - G.3 Derivation of the shift-share approximation
- Appendix H. Indirect Inference for the estimation ν
 - H.1 Estimation equation
 - H.2 Groundwork for the indirect inference
 - H.3 Data and results

Table A.1
List of Roman symbols in the model.

Symbol	Description
A_{ogs}	Level parameter of the Roy-Fréchet distribution
A_{ogs}^M	Efficiency of the matching between workers and employers
c_{os}	Marginal cost of Y_{os}
e_{ogs}	Employment rate
D_d	Trade deficits
F_{os}	Lower-level CES production function with inputs Z_{os}, M_{os}
G_o	Number of groups
h_{og}	Average hours worked per worker
I_{ogs}	Nominal revenue
i_{og}	Average hourly wage rate
K_{os}	Structures, as input in production of Y_{os}
L_{og}	Measure of workers
M_{os}	Equipment, as input in production of F_{os}
N_{oks}	Intermediate inputs in production of Y_{os}
N	Number of sectors
P_o	Price of the final good
R_{os}	Total revenue
S	Number of sectors
T_{os}	Level parameter of the EK-Fréchet distribution
U_{og}	Utility
\tilde{V}_{ogs}	Number of vacancies
V_d	Value added
w_{os}	Wage per effective labor unit
X_{ds}	Expenditure
Y_{os}	Physical output
Z_{ogs}	Supply of effective labor units
z_s	Number of effective units of labor for any worker

Table A.2
List of Greek symbols in the model.

Symbol	Description
α_{os}	Cobb–Douglas share for F_{os} in production of Y_{os}
β_{ds}	Expenditure shares in consumption
Γ	Gamma function
γ_{oks}	Cobb–Douglas share for N_{oks} in production of Y_{os}
δ_{og}	Relative utility weight on consumption versus hours worked
ζ_{os}	Share of structures in the total value of structures and equipment
η	$\Gamma(1 - \frac{1+\mu}{\mu\kappa})$
$\hat{\eta}$	$\Gamma(1 - \frac{1}{\mu\kappa})$
$\tilde{\eta}_s$	$\Gamma(1 - \frac{\sigma_s - 1}{\theta_s})^{1/(1-\sigma_s)}$
θ_s	Dispersion parameter of the EK–Fréchet distribution
κ	Dispersion parameter of the Roy–Fréchet distribution
λ_{ods}	Trade shares
μ	Inverse of the intensive margin labor-supply elasticity
ν_{og}	Nash bargaining share of workers
ξ_{os}	Automation and productivity shifter
π_{og}	Sectoral employment shares
ρ	Elasticity of substitution between labor and equipment
σ_s	Elasticity of substitution across varieties
τ_{ods}	Iceberg trade costs
υ	Elasticity regulating productivity changes driven by automation
Φ_{og}	Index of sectoral wages
χ	Elasticity of the employment rate to labor market tightness
ψ_{ogs}	Labor market tightness
ω_{os}	Cost share of labor in production of F_{os}

Table A.3
Sector classification: overview and summary statistics.

Sector	Sector name	$\beta_{U.S.s}$	$\pi_{U.S.s}$	$\alpha_{U.S.s}$	$\omega_{U.S.s}$	$\gamma_{U.S.s}$
A	Agriculture, forestry, and fishing	0.004	0.011	0.287	0.755	0.593
B	Mining and quarrying	0.003	0.007	0.219	0.772	0.486
10–12	Food, beverages and tobacco	0.032	0.013	0.221	0.598	0.701
13–15	Textiles, apparel, and leatherware	0.014	0.007	0.289	0.838	0.668
16–18	Wood, paper, and printing	0.004	0.019	0.344	0.849	0.629
19	Coke and refined petroleum products	0.008	0.002	0.168	0.289	0.758
20–21	Chemical industry	0.013	0.013	0.355	0.464	0.578
22–23	Rubber, plastics, and other non-metallics	0.003	0.012	0.347	0.708	0.598
24–25	Basic metals and metal products	0.003	0.021	0.354	0.776	0.596
26–27	Electrical and optical equipment	0.032	0.030	0.381	0.735	0.576
28	Machinery and equipment	0.016	0.015	0.357	0.779	0.607
29–30	Transport equipment	0.041	0.023	0.293	0.698	0.668
31–33	Other manufacturing; repair and installation	0.016	0.012	0.445	0.762	0.495
D-E	Electricity, gas and water supply	0.019	0.013	0.238	0.664	0.498
F	Construction	0.074	0.061	0.471	0.897	0.490
G	Wholesale and retail trade, vehicle repair	0.124	0.131	0.503	0.820	0.298
H	Transportation and storage	0.018	0.037	0.394	0.935	0.481
I	Accommodation and food services	0.036	0.029	0.402	0.875	0.441
J	Information and communication	0.054	0.055	0.414	0.732	0.514
K-L	Financial and insurance activities	0.154	0.076	0.182	0.937	0.372
M-N	Professional and administrative services	0.039	0.121	0.525	0.847	0.373
R-S	Arts and entertainment, other services	0.028	0.033	0.517	0.913	0.338
O-P-Q-T-U	Social and personal services	0.262	0.260	0.562	0.959	0.344

This table lists the sector classification used in our estimation and simulations. For each sector, the table provides the value in the US for the listed parameter in the year 2000.

Appendix B. Supplementary analysis

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jinteco.2024.103912>.




Fig. A.1. Comparison of π_{gs} measures. The figure compares different measures of sectoral employment shares (π_{gs}) for US commuting zones used in our estimation. On the horizontal axis, we have employment shares calculated based on labor income in each sector (π_{gs}^{income}), while on the vertical axis the shares are calculated based on the number of hours worked in a sector (π_{gs}^{hours}). The correlation between the two measures is 97%.




Fig. A.2. Pattern of the wage changes for the different shocks. This figure compares the changes in the real wages by sector for each of the shocks.

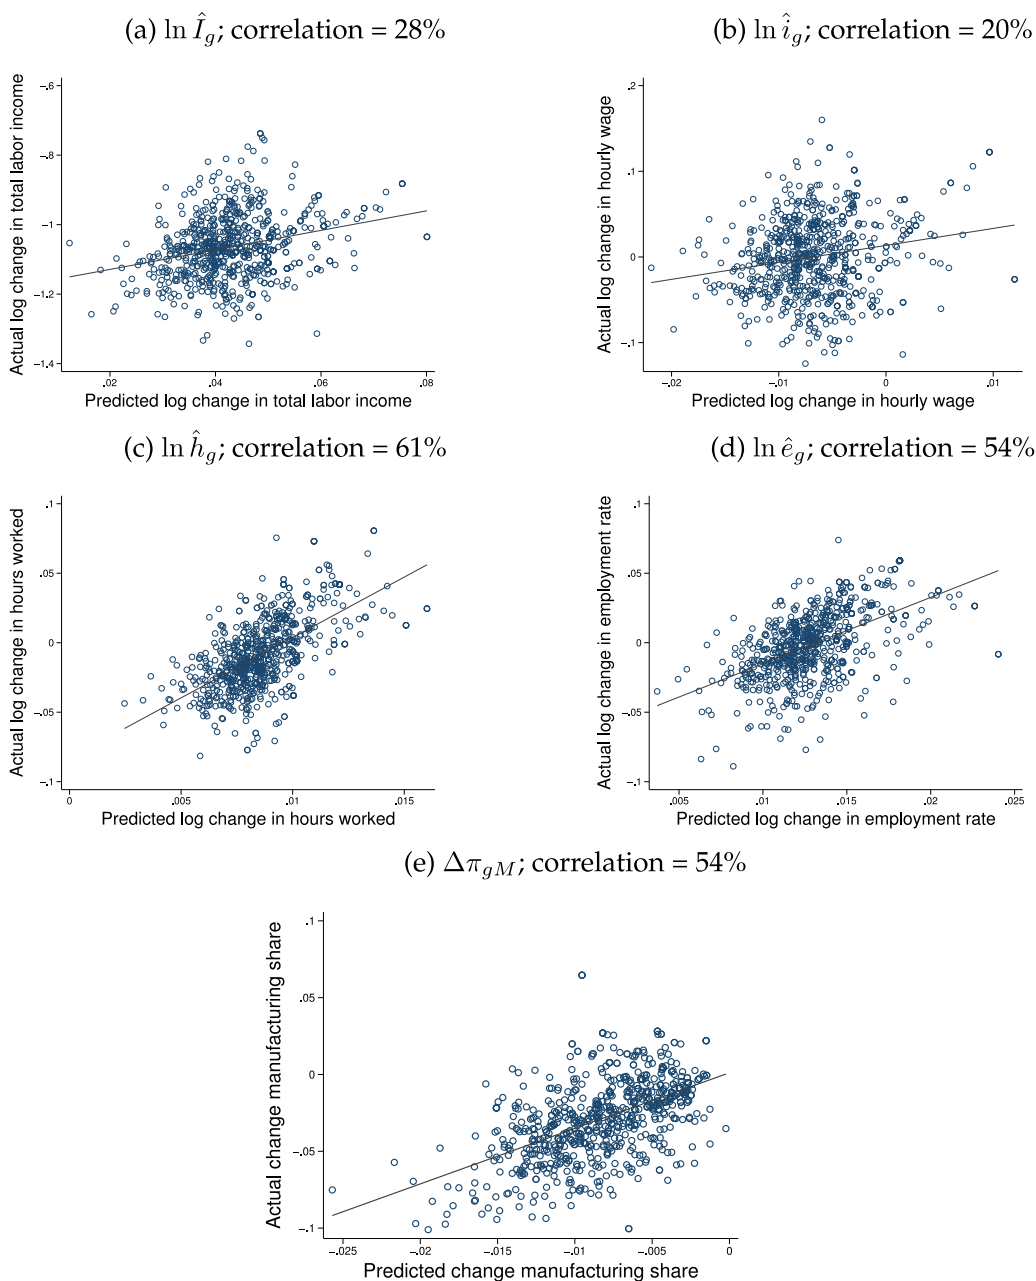


Fig. A.3. Fit of model's predicted changes to non-targeted moments. The table plots the observed values against the model's predicted values for log changes in CZs' income (I_g), average hourly wage (i_g), hours worked (h_g), employment rate (e_g), and manufacturing share (π_{gM}). In contrast to the other variables, the change in the manufacturing share is measured in percentage points.

References

- Acemoglu, D., Restrepo, P., 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *Amer. Econ. Rev.* 108 (6), 1488–1542.
- Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from US labor markets. *J. Polit. Econ.* 128 (6), 2188–2244.
- Adão, R., 2016. Worker heterogeneity, wage inequality, and international trade: Theory and evidence from Brazil. working paper.
- Adão, R., Arkolakis, C., Esposito, F., 2020. General equilibrium effects in space: theory and measurement. NBER working paper.
- Adão, R., Kolesár, M., Morales, E., 2019. Shift-share designs: theory and inference. *Q. J. Econ.* 134 (4), 1949–2010.
- Alvarez, F., Lucas, R.E., 2007. General equilibrium analysis of the Eaton–Kortum model of international trade. *J. Monet. Econ.* 54 (6), 1726–1768.
- van Ark, B., Jäger, K., 2017. Recent trends in Europe's output and productivity growth performance at the sector level, 2002–2015. *Int. Prod. Monit.* (33), 8–23.
- Arkolakis, C., Costinot, A., Rodríguez-Clare, A., 2012. New trade models, same old gains? *Amer. Econ. Rev.* 102 (1), 94–130.
- Artuç, E., Chaudhuri, S., McLaren, J., 2010. Trade shocks and labor adjustment: A structural empirical approach. *Amer. Econ. Rev.* 100 (3), 1008–1045.
- Atalay, E., Phongthientham, P., Sotelo, S., Tannenbaum, D., 2018. New technologies and the labor market. *J. Monetary Econ.* 97, 48–67.
- Austin, B., Glaeser, E., Summers, L.H., 2018. Saving the heartland: Place-based policies in 21st century america. *Brook. Pap. Econ. Act.* 8.
- Autor, D.H., Dorn, D., Hanson, G.H., 2013a. The China syndrome: Local labor market effects of import competition in the United States. *Amer. Econ. Rev.* 103 (6), 2121–2168.
- Autor, D., Dorn, D., Hanson, G., 2013b. The geography of trade and technology shocks in the United States. *Am. Econ. Rev. Pap. Proc.* 103 (3), 220–225.
- Autor, D.H., Dorn, D., Hanson, G.H., 2015. Untangling trade and technology: Evidence from local labour markets. *Econ. J.* 125 (584), 621–646.
- Bagwell, K., Staiger, R.W., Yurukoglu, A., 2018. Quantitative Analysis of Multi-Party Tariff Negotiations. Technical Report, National Bureau of Economic Research.
- Barnichon, R., Figura, A., 2015. Labor market heterogeneity and the aggregate matching function. *Am. Econ. J.: Macroecon.* 7 (4), 222–249.
- Bartelme, D.G., Costinot, A., Donaldson, D., Rodríguez-Clare, A., 2019. The Textbook Case for Industrial Policy: Theory Meets Data. Technical Report, National Bureau of Economic Research.
- Bernon, B., Magerman, G., 2022. Income Inequality in General Equilibrium. National Bank of Belgium.
- Besedes, T., Lee, S.H., Yang, T., 2021. Trade liberalization and gender gaps in local labor market outcomes: dimensions of adjustment in the united states. *J. Econ. Behav. Organ.* 183, 574–588.
- Bloom, N., Handley, K., Kurman, A., Luck, P., 2019. The impact of chinese trade on us employment: The good, the bad, and the debatable. Unpublished draft.
- Borjas, G.J., 1980. The relationship between wages and weekly hours of work: The role of division bias. *J. Hum. Resour.* 15 (3), 409–423.
- Borusyak, K., Hull, P., Jaravel, X., 2020. Quasi-experimental shift-share research designs. working paper.
- Burstein, A., Morales, E., Vogel, J., 2019. Changes in between-group inequality: computers, occupations, and international trade. *Am. Econ. J.: Macroecon.* 11 (2), 348–400.
- Burstein, A., Vogel, J., 2017. International trade, technology, and the skill premium. *J. Polit. Econ.* 125 (5), 1356–1412.
- Caliendo, L., Dvorkin, M., Parro, F., 2019. Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica* 87 (3), 741–835.
- Caliendo, L., Parro, F., 2015. Estimates of the trade and welfare effects of NAFTA. *Rev. Econom. Stud.* 82 (1), 1–44.
- Chetty, R., 2012. Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica* 80 (3), 969–1018.
- Costinot, A., Rodríguez-Clare, A., 2014. Trade theory with numbers: Quantifying the consequences of globalization. *Handb. Int. Econ.* 4, 197.
- Curuk, M., Vannorenbergh, G., 2017. Inter-sectoral labor reallocation in the short run: The role of occupational similarity. *J. Int. Econ.* 108, 20–36.
- Dauth, W., Findeisen, S., Suedekum, J., 2014. The rise of the East and the Far East: German labor markets and trade integration. *J. Eur. Econom. Assoc.* 12 (6), 1643–1675.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N., 2021. The adjustment of labor markets to robots.
- Dekle, R., Eaton, J., Kortum, S., 2008. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Pap.* 55 (3), 511–540.
- Dix-Carneiro, R., 2014. Trade liberalization and labor market dynamics. *Econometrica* 82 (3), 825–885.
- Dix-Carneiro, R., Kovak, B.K., 2017. Trade liberalization and regional dynamics. *Amer. Econ. Rev.* 107 (10), 2908–2946.
- Eaton, J., Kortum, S., 2002. Technology, geography, and trade. *Econometrica* 70 (5), 1741–1779.
- Ebenstein, A., Harrison, A., McMillan, M., Phillips, S., 2014. Estimating the impact of trade and offshoring on american workers using the current population surveys. *Rev. Econ. Stat.* 96 (4), 581–595.
- Fajgelbaum, P.D., Khandelwal, A.K., 2016. Measuring the unequal gains from trade. *Q. J. Econ.* 131 (3), 1113–1180.
- Feenstra, R.C., 2017. Statistics to Measure Offshoring and Its Impact. Technical Report, National Bureau of Economic Research.
- Fort, T.C., 2017. Technology and production fragmentation: Domestic versus foreign sourcing. *Rev. Econom. Stud.* 84 (2), 650–687.
- Fort, T.C., Pierce, J.R., Schott, P.K., 2018. New perspectives on the decline of us manufacturing employment. *J. Econ. Perspect.* 32 (2), 47–72.
- Galle, S., Rodríguez-Clare, A., Yi, M., 2023. Slicing the pie: Quantifying the aggregate and distributional effects of trade. *Rev. Econom. Stud.* 90 (1), 331–375.
- Giri, R., Yi, K.-M., Yilmazkuday, H., 2018. Gains from trade: Does sectoral heterogeneity matter? *J. Int. Econ.* 103429.
- Goldsmith-Pinkham, P., Sorkin, I., Swift, H., 2020. Bartik instruments: What, when, why, and how. *Amer. Econ. Rev.* 110 (8), 2586–2624.
- Greenland, A., Lopresti, J., McHenry, P., 2019. Import competition and internal migration. *Rev. Econ. Stat.* 101 (1), 44–59.
- Hakobyan, S., McLaren, J., 2016. Looking for local labor market effects of NAFTA. *Rev. Econ. Stat.* 98 (4), 728–741.
- Head, K., Mayer, T., 2014. Gravity equations: Workhorse, toolkit, and cookbook. In: *Handbook of International Economics*, vol. 4, pp. 131–195.
- Hémous, D., Olsen, M., 2022. The rise of the machines: automation, horizontal innovation, and income inequality. *Am. Econ. J. Macroecon.* 14 (1), 179–223.
- Hsieh, C.-T., Hurst, E., Jones, C.I., Klenow, P.J., 2019. The allocation of talent and U.S. economic growth. *Econometrica* 87 (5), 1439–1474.
- Hsieh, C.-T., Ossa, R., 2016. A global view of productivity growth in China. *J. Int. Econ.* 102, 209–224.
- Hubmer, J., 2021. The race between preferences and technology. Unpublished Working Paper.
- Karabarbounis, L., Neiman, B., 2014. The global decline of the labor share. *Q. J. Econ.* 129 (1), 61–103.
- Kehoe, T.J., Ruhl, K.J., Steinberg, J.B., 2018. Global imbalances and structural change in the United States. *J. Polit. Econ.* 126 (2), 761–796.
- Kim, R., Vogel, J., 2020. Trade and welfare (across local labor markets). National Bureau of Economic Research working paper.
- Kim, R., Vogel, J., 2021. Trade shocks and labor market adjustment. *Am. Econ. Rev. Insights* 3 (1), 115–130.
- Kovak, B.K., 2013. Regional effects of trade reform: What is the correct measure of liberalization? *Am. Econ. Rev.* 103 (5), 1960–1976.
- Krusell, P., Ohanian, L.E., Ríos-Rull, J.-V., Violante, G.L., 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68 (5), 1029–1053.
- Lagakos, D., Waugh, M.E., 2013. Selection, agriculture, and cross-country productivity differences. *Am. Econ. Rev.* 103 (2), 948–980.
- Lee, E., 2020. Trade, inequality, and the endogenous sorting of heterogeneous workers. *J. Int. Econ.* 125, 103310.
- Lee, E., Yi, K.-M., 2018. Global value chains and inequality with endogenous labor supply. *J. Int. Econ.* 115, 223–241.
- Lorentzen, L., 2022. Domino effects: Understanding sectoral reallocation and its wage implications. BI working Paper.
- Moll, B., Rachel, L., Restrepo, P., 2022. Uneven growth: Automation's impact on income and wealth inequality. working paper.
- Ngai, L.R., Pissarides, C.A., 2007. Structural change in a multisector model of growth. *Am. Econ. Rev.* 97 (1), 429–443.
- Oberfield, E., Raval, D., 2021. Micro data and macro technology. *Econometrica* 89 (2), 703–732.
- OECD, 2019. 9A. Fixed assets by activity and by asset.

- Ossa, R., 2014. Trade wars and trade talks with data. *Amer. Econ. Rev.* 104 (12), 4104–4146.
- Pierce, J.R., Schott, P.K., 2020. Trade liberalization and mortality: evidence from US counties. *Am. Econ. Rev. Insights* 2 (1), 47–64.
- Rodríguez-Clare, A., Ulate, M., Vasquez, J.P., 2022. Trade with nominal rigidities: Understanding the unemployment and welfare effects of the China shock.
- Roy, A.D., 1951. Some thoughts on the distribution of earnings. *Oxf. Econ. Pap.* 3 (2), 135–146.
- Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., Sobek, M., 2019. *Ipums usa: Version 9.0 [dataset]*, Vol. 10. IPUMS, Minneapolis, MN, p. D010.
- Shapiro, J.S., 2016. Trade costs, CO₂, and the environment. *Am. Econ. J. Econ. Policy* 8 (4), 220–254.
- Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *Am. Econ. Rev.* 95 (1), 25–49.
- The Conference Board, 2023. EU KLEMS July 2018 release.
- Timmer, M.P., Dietzenbacher, E., Los, B., Stehrer, R., Vries, G.J., 2015. An illustrated user guide to the world input–output database: the case of global automotive production. *Rev. Int. Econ.* 23 (3), 575–605.
- Topalova, P., 2010. Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India. *Am. Econ. J. Appl. Econ.* 2 (4), 1–41.
- US Bureau of Labor Statistics, 2019. FRED, Federal Reserve Bank of St. Louis, All Employees: Manufacturing [MANEMP]. <https://fred.stlouisfed.org/series/MANEMP>.