

UC Berkeley

CEGA Working Papers

Title

Premature Deindustrialization and the Agricultural Productivity Gap

Permalink

<https://escholarship.org/uc/item/2tt8g8h3>

Authors

Cejka, Tim

O'Neil, Flint

Publication Date

2026-10-15

Series Name: WPS
Paper No.: 268
Issue Date: October 15, 2026

Premature Deindustrialization and the Agricultural Productivity Gap

Tim Cejka and Flint O'Neil



CEGA

Center for Effective Global Action

Working Paper Series

Center for Effective Global Action
University of California



This paper is posted at the eScholarship Repository, University of California. http://escholarship.org/uc/cega_wps Copyright © 2023 by the author(s).

The CEGA Working Paper Series showcases ongoing and completed research by faculty affiliates of the Center. CEGA Working Papers employ rigorous evaluation techniques to measure the impact of large-scale social and economic development programs, and are intended to encourage discussion and feedback from the global development community.

Recommended Citation:

Cejka, Tim and Flint O'Neil (2026): Premature Deindustrialization And The Agricultural Productivity Gap. CEGA Working Paper Series No. WPS-268. Center for Effective Global Action. University of California, Berkeley. Text. <https://doi.org/10.26085/C30882>

Premature Deindustrialization and the Agricultural Productivity Gap

Tim Cejka and Flint O’Neil *

October 2026

Abstract

Why does manufacturing employment in emerging markets peak at half the levels advanced economies reached? We investigate this question through a quantitative trade model with structural change and rural-to-urban migration. Emerging markets have a comparative advantage in manufacturing, so trade alone should make them industrialize more, not less. Instead, we show that what holds them back are large rural-urban labor market wedges, which may reflect migration frictions or skill barriers. We estimate the wedges from wage gaps and migration flows. Without them, the median emerging economy does not prematurely deindustrialize, and instead reaches a peak manufacturing employment share roughly 10 percentage points higher than the data. By contrast, the rise of China contributes little to premature deindustrialization. Consistent with the model, we provide quasi-experimental evidence that a labor-saving agricultural technology increased industrialization in Brazil only in municipalities with young, mobile populations.

Keywords: international trade, structural transformation, premature deindustrialization, rural-urban mobility frictions, dynamic spatial equilibrium

JEL Classification: F11, F16, F43, O14, R23

*We thank Adrien Auclert, Adrien Bilal, Lorenzo Caliendo, Gerard Martín Escofet, Kyra Frye, Maria Elena Guerrero, Patrick Kehoe, Benny Kleinman, Emi Nakamura, Pascal Noel, Elena Pastorino, Stephen Redding, Venance Riblier, Andrés Rodríguez-Clare, Nico Rotundo, Yash Singh, Tomáš Štěpák, and Filip Tokarski for helpful comments and discussions. We especially thank Michael Sposi for sharing useful replication files. Cejka: UC Berkeley. E-mail: tim_cejka@berkeley.edu. O’Neil: Stanford University. E-mail: foneil@stanford.edu.

1 Introduction

As countries develop, the share of employment allocated to agriculture declines, services increases and manufacturing is hump-shaped with a peak. This process of structural transformation was emphasized by Kuznets’s (1973) Nobel Prize address as one of the central features of development. The peak manufacturing employment share of today’s emerging economies is roughly half what it was historically for advanced economies, and it is reached at a substantially lower level of income. Rodrik (2016) calls this pattern *premature deindustrialization*. This pattern, illustrated in Figure 1, is one of the most prominent stylized facts of post-war structural change and the central puzzle this paper addresses.




FIGURE 1: INDUSTRIAL PROFILES FOR SELECT ECONOMIES

Notes: Smoothed manufacturing employment shares for nine early, middle, and late developers, plotted against calendar year in Panel A and against the non-agricultural employment share in Panel B. Manufacturing is broad manufacturing, the sum of manufacturing, mining, utilities, and construction employment; the non-agricultural employment share is one minus the agricultural employment share. Panels A and B are smoothed with a fourth-order polynomial over time. Dashed paths in warm colors are early and middle developers, solid paths in cool colors late developers. Source: GGDC 10-sector database.

Benchmark models of structural change attribute premature deindustrialization to uneven productivity growth across sectors and countries. The core argument is that faster relative manufacturing productivity growth today relative to history means that fewer workers are needed to produce manufactured goods. Two features of the data sit in tension with this explanation.

First, the productivity difference between rich countries and poor countries in agriculture is more than double their productivity difference in the overall economy. Herrendorf, Rogerson, and Valentinyi (2026), Gollin, Lagakos, and Waugh (2014a), Caselli (2005), and Restuccia, Yang, and Zhu (2008) all document that poor countries lag behind far more in agriculture than any other sector of the economy. For instance, Herrendorf, Rogerson, and Valentinyi (2026) document that the richest 10% of countries are around 100 times more

productive than the poorest 10% of countries in agriculture, but only 40 times more productive in manufacturing. Given that services are less tradable, this means that developing economies have comparative advantage in manufacturing over agriculture. Thus, in a world with trade we might expect emerging markets to be very industrial – exporting manufactures and importing agriculture. But in the data we observe roughly balanced trade between rich and poor countries in both agriculture and manufacturing. This points to some force that is making emerging markets produce “too much” agriculture and “too little” manufacturing in a world with trade.

Second, a worker’s revenue productivity is roughly 3–5 times lower in agriculture than in manufacturing and services in the typical emerging market (Gollin, Lagakos, and Waugh 2014b). In a frictionless labor market, workers would leave agriculture and enter manufacturing until revenue productivity is equal between sectors. These two facts suggest that some force keeps workers in agriculture despite higher pay and productivity in manufacturing. This paper models this force as a fixed-cost spatial wedge that workers must pay in order to transition between rural and urban labor markets. This wedge could be either a migration cost or a skill differential between urban and rural workers.

This paper shows that once the spatial wedge is estimated to match migration and rural-urban wage gaps, the standard spatial trade model of Caliendo, Dvorkin, and Parro (2019) extended to incorporate structural change can generate premature deindustrialization. We estimate the model fundamentals at a region and sector level. Our estimation finds that, if not for the spatial wedge, emerging markets would have a substantial comparative advantage in manufacturing over agriculture. We show that this latent comparative advantage is strong enough that emerging markets would specialize in manufacturing over agriculture, and prematurely *industrialize* rather than deindustrialize. That is, trade is strong enough to *flip* the sign of the deindustrialization result in benchmark models without a rural-to-urban labor market wedge.

Two further empirical patterns point to the same wedge. First, across late developers, the rural-urban labor revenue productivity gap covaries negatively with the peak manufacturing share. Countries with high agricultural wage gaps have low peak manufacturing shares. Second, while emerging markets have exhibited strong premature deindustrialization in employment shares, premature deindustrialization in value added shares is weaker. Both patterns are consistent with the claim that barriers between rural and urban labor markets are needed to explain premature deindustrialization.

To analyze these patterns, we build a model of structural change that matches global trade flows, internal rural-urban migration, and the production network of the economy. Before quantitatively solving, we exploit the model’s budget constraints to perform a de-

scriptive decomposition exercise. We extend the approach of the open-economy accounting decomposition of Smitkova (2024) by incorporating within-country regional structure. The exercise decomposes the observed change in manufacturing employment shares from 1970 to 2014 into six channels. The channels are: trade specialization ($\tilde{\Pi}$), secular trends (preferences ($\tilde{\Lambda}$), IO weights ($\tilde{\Gamma}$) and wages ($\tilde{w}$)), migration (μ), trade imbalances ($\tilde{D}$), natural population dynamics and a Baumol term, the employment-share change attributable to manufacturing’s relative value added per worker rather than to its value-added share. We apply the decomposition to the six emerging markets that are both present in the WIOD and GGDC datasets. For these six economies, the rural to urban migration channel contributes about +7.1 percentage points to the total change in manufacturing employment shares. In fact, this channel is the single largest positive channel in four of the six emerging economies (China +11.9, Brazil +10.2, Korea +8.0, Mexico +5.2 pp). In developed economies this channel is broadly negligible. We then turn to a quantitative model that can account for these patterns.

Our quantitative model is based on the multi-sector trade model of Caliendo, Dvorkin, and Parro (2019) (CDP). We extend CDP by incorporating structural change through non-homothetic preferences, and allowing for time variation in the labor market wedges and input-output weights. Each country splits into a rural and an urban region. Workers solve a discrete-choice costly migration problem. Sectors are tradable up to bilateral iceberg trade costs estimated from observed flows. Consumers have non-homothetic Stone-Geary CES preferences over agriculture, manufacturing, and services. Although uneven technological growth between sectors is the primary driver of structural change in our model, non-homothetic preferences are necessary in order to match the speed of the global drop in agriculture and rise in services. The model jointly matches trade flows, the broad patterns of sectoral employment, and rural-urban wages from 1970–2013. The model is consistent with both observed agricultural wage gaps and observed trade patterns by construction. Because rural and urban labor markets clear at different wages, the model also reproduces the fact that premature deindustrialization appears in employment shares but is weaker in value added shares.

In the general equilibrium counterfactual that holds productivities and trade frictions at their estimated values but shuts off the rural-urban migration wedges, late developers do not deindustrialize prematurely. Instead, comparative advantage pulls agricultural labor directly into manufacturing. Late developers then reach manufacturing employment peaks well above those in the data. With the migration wedge, this comparative-advantage channel is shut down: rural labor is locked in agriculture and cannot supply the urban manufacturing sector. The Baumol pull from manufacturing into non-tradable services then dominates, and

the early, low manufacturing peak that we observe in the data emerges.

We recover large and persistent rural-urban wedges in developing economies. These wedges have larger consequences for today’s emerging markets than they would have had for the advanced economies because development is now compressed into a shorter window. Advanced economies historically developed at 1-2 percent annual rate of growth, while today’s emerging markets grow at rates closer to 5 percent. If it takes multiple generations to move workers off the farm, then this friction bites less if the development window is long at 1 percent annual GDP growth. But if the development window is shorter, then these movement frictions will bite far more and drive up the rural-urban wage gap. This large wage gap leads us to estimate large, biting wedges that generate premature deindustrialization in our model.

Since these wedges are recovered from the model rather than observed, we test the mechanism they imply outside of the model. In the model, the wedge blocks the push of labor out of agriculture. A labor-saving agricultural productivity shock should therefore raise manufacturing employment shares more in places with lower rural-urban wedges, i.e., in places that rural workers can leave more easily.

We test this prediction of our model using Brazilian census microdata. We borrow a research design analyzing a labor-saving agricultural productivity shock from the literature, specifically, the roll-out of genetically modified soy in Brazil (Bustos, Caprettini, and Ponticelli 2016). Bustos, Caprettini, and Ponticelli (2016) exploits the introduction of the new, labor-saving agricultural technology and municipalities’ cross-sectional exposure to it, finding that it raised employment shares in manufacturing but lowered them in agriculture. We augment their design with predetermined proxies of municipal migration propensity, including the youth share and past realized urban-directed out-migration, to examine whether migration propensity moderated the shock’s effect. The youth share captures mobility through the well-documented concentration of migration among young adults, while past out-migration records actual moves. We find that, across a range of proxies, functional forms, and a rich set of controls, the shock exhibits a strong gradient with a municipality’s mobility. More mobile municipalities experienced larger effects, with the shock more strongly increasing their manufacturing shares and decreasing their agricultural shares. In contrast, the shock had no detectable effect on employment shares in the least mobile municipalities. We provide evidence consistent with our identifying assumptions by showing that differentially exposed and mobile municipalities were not on differential manufacturing employment trends in the pre-period.

Related Literature. This paper relates to three strands of literature. The first piece of literature we relate to are the quantitative theories of premature deindustrialization. Huneus

and Rogerson (2024) and Fujiwara and Matsuyama (2024) generate the pattern in efficient closed-economy environments through differential productivity growth and technology adoption lags. Sposi, Yi, and Zhang (2026) (henceforth, SYZ) is perhaps the closest paper to ours in the literature. They construct a multi-country Eaton-Kortum trade model with capital accumulation and nonhomothetic preferences and production.

Our model departs from these benchmarks in two ways. First, Huneus and Rogerson (2024) and Fujiwara and Matsuyama (2024) are closed economy models. Closed-economy models do not survive opening to trade. In a closed economy, faster manufacturing productivity growth leads to less employment in manufacturing. In an open economy however, fast manufacturing productivity growth increases manufacturing employment because of specialization through trade. In our 43-country panel, countries that experience relatively faster manufacturing productivity growth do not reach lower manufacturing peaks, as Proposition 2 in Huneus and Rogerson (2024), and Fujiwara and Matsuyama (2024) predicts. A dominant open economy channel is, however, consistent with the trade model of Sposi, Yi, and Zhang (2026). This point that trade can flip the sign of industrial dynamics is not novel. Matsuyama (1992) shows that more productive agricultural technology pushes workers into manufacturing in a closed economy but into agriculture in a pure small open economy. Likewise, Glaeser (2014) shows this reversal due to trade for urbanization. A set of complementary papers (Peters, Zhang, and Zilibotti 2026; Fan, Peters, and Zilibotti 2023) documents that labor leaving agriculture in India and Sub-Saharan Africa has moved largely into non-tradable services rather than manufacturing. We take this channel as given.

Second, each of the benchmark models makes the simplifying assumption that labor reallocates frictionlessly across sectors at a single national wage. At a single wage, value added shares are equal to employment shares, so any model that fits manufacturing value added shares will tend to overshoot employment shares. Interpreting the single wage as per efficiency unit avoids this overshoot, but at the cost of remaining silent on employment headcounts. In either case, while a single-wage model can be useful for studying premature deindustrialization in value added shares, it cannot directly address employment shares. These observations motivate us to write a model that features both global trade and internal labor market frictions.

If emerging markets seem to have a comparative advantage and higher wages in manufacturing, why does trade deepen premature deindustrialization in Sposi, Yi, and Zhang (2026)? SYZ show that trade interacts with sector-biased technical change to transmit falling manufacturing prices across countries. Since SYZ has frictionless labor markets, SYZ shows that, holding migration and worker sorting fixed, trade would thus lower manufacturing's value-added share through relative prices. By contrast, we show that if rural-urban labor

market barriers were different, then trade would instead raise manufacturing’s employment share. This difference in counterfactual framing is what distinguishes our different headline interpretations. If the agricultural wage gap is due entirely to worker composition then the SYZ interpretation in efficiency units is correct. If the agricultural wage gap is due to other barriers, such as migration frictions, then a single wage inversion does not capture them which motivates a dynamic spatial model. Our model is a dynamic spatial model, and so it allows for either interpretation. We discuss this difference in more depth in Section 5.2.

The second strand of literature this paper relates to regards the agricultural productivity gap. Gollin, Lagakos, and Waugh (2014b) documents large gaps in value added per worker between agriculture and the rest of the economy across developing countries, and Bryan and Morten (2019) estimates that removing movement costs in Indonesia would raise aggregate output by 7.5 percent, and by 22 percent together with amenity differences. Whether the gap is due to barriers or sorting of workers is contested. Wage and consumption microdata in Herrendorf and Schoellman (2018) and Young (2013) attribute most of the gap to selection on worker productivity. Similarly, Hamory *et al.* (2021) finds that individual fixed effects absorb 67–92 percent of the gap. Other evidence shows the gap is not entirely selection. Imbert and Papp (2020) shows that the same rural Indian worker earns 35 percent more on seasonal trips to urban construction than on public works at home, and still stays. Bryan, Chowdhury, and Mobarak (2014) shows that purchasing an \$8.50 bus fare for workers raises Bangladeshi seasonal migration by 22 percentage points and consumption among induced migrants by 30–35 percent. In Kenya, information interventions to rural families about urban incomes raise family income by 24 percent (Baseler 2023). In Indonesia, Bryan and Morten (2019) finds that lowering movement costs raises output primarily by allowing workers to sort more effectively.

Our model and empirical evidence are agnostic about the source of the rural-urban wedges. In our model, the rural-urban wedge can be interpreted as migration frictions or skill barriers; crucially, however, both of these interpretations have the observable implication of lower migration propensity. Our quasi-experimental evidence from Brazil tests this implication, using the youth share of a municipality as a proxy for migration propensity. The proxy is consistent with migration frictions based on the well-documented age gradient in mobility. Across countries and time periods, migration is concentrated among the young, with migration propensity peaking in the twenties and then declining (Rogers and Castro 1981; Bernard, Bell, and Charles-Edwards 2014). The proxy is also consistent with skill barriers based on the empirical cohort pattern of labor reallocation out of agriculture. In particular, young cohorts hold the human capital that is valued outside of agriculture, with new cohorts entering the labor market driving nearly half of the global decline in agricul-

tural employment (Porzio, Rossi, and Santangelo 2022). Distinguishing between the two interpretations is beyond the scope of this paper and is left for future work.

Finally, it is worth noting that evidence on whether agricultural productivity shocks have an industrializing effect remains mixed. In Brazil, the labor-saving agricultural productivity shock of genetically modified soy raised manufacturing employment shares (Bustos, Caprettini, and Ponticelli 2016), with the resulting savings in agriculture channeled into manufacturing and services in other regions (Bustos, Garber, and Ponticelli 2020), and the released workers reallocated within local labor markets rather than across them (Bustos *et al.* 2023). Also in Brazil, Albert *et al.* (2026) finds that workers who stay in drought-hit regions reallocate toward manufacturing while climate migrants end up in small firms outside manufacturing at their destinations, an asymmetry it attributes to spatial labor market frictions. In contrast, agricultural productivity shocks have not led to local industrialization in India after the Green Revolution (Foster and Rosenzweig 2004; Moscona 2019), over the Ogallala aquifer (Hornbeck and Keskin 2015), and along India’s canals, where structural change took the form of the emergence and growth of nearby towns rather than non-farm growth in the irrigated villages (Asher *et al.* 2024). Our evidence relates to this debate only by showing that, within Brazil, the industrializing effect of the soy shock depends on spatial mobility costs.

Roadmap.— Section 2 describes the data, and documents motivating cross-country patterns in the data. Section 3 develops the dynamic spatial trade model as well as the accounting decomposition, while Section 4 presents the main theoretical propositions representing the extended dynamic hat algebra solution methods and analytical results. Section 5 presents quantitative results and counterfactuals, including the frictionless-labor counterfactuals. Section 6 reports the Brazilian quasi-experimental test. Section 7 concludes.

2 Data

2.1 Sources

We use data from the following sources for our empirical analysis.

(i.) **Macro database.** We create our database from the following sources: GGDC 10 (de Vries *et al.* 2015; Timmer, de Vries, and de Vries 2015), (Feenstra *et al.* 2005), CEPII BACI (Gaulier and Zignago 2010), UNIDO (United Nations Statistics Division 2015), EU KLEMS (O’Mahony and Timmer 2009), World Input Output Database (WIOD) (Timmer *et al.* 2015), and Penn World Tables (PWT) (Feenstra, Inklaar, and Timmer 2015).

The GGDC 10 Sector database provides long-term data on sectoral developments at the

10 sector level. We use GGDC, UNIDO MEI and EU KLEMS to construct the main data series on nominal and real value-added, as well as employment shares per worker for 43 countries. The idea to increase coverage of GGDC by additionally using the EU KLEMS extension file follows SYZ. We thank them for sharing their replication package which was helpful for us to learn from.

As is standard in the literature, we define agriculture as agriculture, hunting, and forestry and fishing. We define manufacturing, mining, construction and utilities as manufacturing. Services are the aggregate of the remaining sectors: trade, restaurants and hotels, transport and communication, finance, insurance and real estate, government services and community social and personal services. We use the Penn World Tables price index series exchange rate xr to convert the value-added of these sectors back to nominal USD values.

We then stitch together the Feenstra and BACI series with WIOD (and WIOD LR (Woltjer, Gouma, and Timmer 2021)) to construct a trade database that includes self-expenditure shares. As is well known in trade, self-expenditure shares are critical for estimating trade barriers and welfare (see, for example, Arkolakis, Costinot, and Rodríguez-Clare 2012).

Since our trade coverage is comprehensive, our coverage is primarily restricted by the smaller country sample present in GGDC10, EU KLEMS, and UNIDO. We rely on either GGDC10, UNIDO, or EU KLEMS for sector-level employment shares and value added. We attempt to cover most of world GDP as well as sampling major emerging markets. The final panel includes 43 countries from 1970 through 2013.

Input-output structure and gross-output prices. In addition to value added and employment sector-level data, we need internationally comparable gross-output prices for each sector and the input-output (IO) structure. We build IO matrices from WIOD intermediate use tables where WIOD covers the country. We stitch the 2000-2014 WIOD series with the WIOD long run series which covers 1965-2000. Countries not in WIOD are imputed through a k -NN imputation on log GDP per capita, region, and sectoral value-added shares which imputes their IO structure by looking at similar economies. More details are provided in the appendix. Sectoral value-added deflators come from UNIDO MEI and GGDC. These are grossed up to a gross-output price index through the country’s IO structure. We anchor the price indices to the 2005 cross-country price level recorded in the GGDC Productivity Level Database (Inklaar and Timmer 2014) where it covers the country. If the country is not covered we resort to a scaled up version of World Bank ICP 2005 (World Bank 2008). If a country is in neither we use a Balassa-Samuelson imputation. Sector level wages are built primarily from UNIDO MEI and GGDC sectoral value added and employment series, with the level anchored to PWT 10.01 nominal aggregate GDP. Self-expenditure shares come straight from WIOD where available. The rest are imputed by either wage bill imputation

or a logit-Lasso. Full details are in Appendix A.

Rural-urban panel. While the above covers the trade part of the model, we also need to cover the dynamic spatial dynamics of the model separately. This dynamic block needs rural and urban employment stocks and wages within each country. The first challenge is to figure out how much of a sector’s activity is rural vs. urban. Denote the fraction of sector- j workers in country n at time t that are rural by $\alpha_{n,t}^j \in [0, 1]$. To obtain this share we combine IPUMS occupational classifications at the micro level with World Bank national urbanization series (World Bank 2026). We assume all agriculture is rural, so $\alpha^{\text{ag}} = 1$. Rural and urban stocks $R_{n,t}, U_{n,t}$ are then employment-weighted aggregates of sectoral counts. Rural and urban wages $w_{n,t}^R, w_{n,t}^U$ come from a spatial market-clearing assumption where we assume that workers can move frictionlessly within regions between sectors. Net rural-to-urban flows fall out of an identity that we describe in the appendix. This panel covers the same 43 countries from 1970 to 2013. We detail this as well as the gross migration flow inversion in Appendix A.

(ii.) **Administrative data from Brazil.** We construct the dataset for our quasi-experimental analysis in Brazil from raw census microdata and FAO-GAEZ grids (IIASA/FAO 2012), closely following each step of data construction based on Bustos, Caprettini, and Ponticelli (2016).¹ We use the population censuses from Brazil’s national statistical institute (IBGE) from years 1980, 1991, 2000, and 2010, which we aggregate to 4,260 minimum comparable areas (AMCs), which maintain municipal boundaries constant over time. This dataset contains variables on sectoral employment, age structure, and other demographic characteristics. We use the FAO-GAEZ v3.0 potential yields under low and high input levels, which we also aggregate to AMCs and merge with the population census. We use this dataset for our exposure measure.

2.2 Patterns

Before moving directly to the quantification of our model, we describe some historical features of the data. Figure 1, Panel A, shows the smoothed manufacturing employment shares over time for select economies.² We plot profiles for Great Britain, France, Spain, Japan, South Korea, Brazil, Peru, Philippines, and Indonesia. The choice of countries for this plot was essentially arbitrary other than trying to capture major economies that were early, middle, and late developers. Figure 2, Panel A, plots the peak manufacturing share over time for all

¹Details on the institutional setting and data construction are in Section 6 and Appendix F. For more details of the data construction, please see Bustos, Caprettini, and Ponticelli (2016).

²To construct these smoothed paths we take the raw data on employment shares from GGDC 10 and fit a fourth-order polynomial over time for each country.

countries in the GGDC database. For countries with no rapid-growth episode in the data window, we use an extremely conservative imputation setting the year rapid growth began at 1900.




FIGURE 2: INDUSTRIAL PROFILES

Notes: Both panels cover all countries in the GGDC 10 dataset (de Vries *et al.* 2015; Timmer, de Vries, and de Vries 2015). Panel A plots each country's peak manufacturing share against the year in which its manufacturing employment share began growing rapidly. For countries with no rapid growth in the data, we take an extremely conservative approach and impute their rapid-growth start to the year 1900. The two lines are lines of best fit, one for the full period 1900 onward and one for 1948 onward. Panel B plots each country's manufacturing share against its non-agricultural share. Both panels smooth with a fourth-order polynomial.

Figure 1, Panel A, and Figure 2, Panel A, highlight that there is substantial heterogeneity in the employment shares of countries. Even in these selected time paths there is a tight relationship between *when* a country began industrializing and the industrial employment share that it eventually reaches. For instance, Great Britain was the world's first industrializer and reached an extraordinary manufacturing employment share peak of almost 50 percent. Middle developers like Spain, Japan, and South Korea obtain more modest shares and the most recent developers like Peru and Philippines seem to have attained very low peaks. These patterns are discussed in more detail in the original paper Rodrik (2016).

Another way to view a similar relationship is to plot the manufacturing employment share of these countries against the non-agricultural employment share. We do this in Figure 1, Panel B, and Figure 2, Panel B. Plotting against the non-agricultural employment share is common in the literature due to the negative monotone relationship between agricultural employment and economic development, pointed out by Kuznets (1973). Thus, using the non-agricultural employment share is a convenient way to visualize economic development on the x-axis. It is also convenient because it highlights another dimension of heterogeneity in the industrialization profiles of countries. Figure 1, Panel B, shows that the manufacturing employment share has peaked at lower levels of non-agricultural employment. That

demonstrates that there is also heterogeneity in the size of the service sector. That is, even at similar levels of non-agricultural employment the service sector of late developers is now larger than it was for early developers.

Another feature of the data we would like to draw attention to is the relationship between economic openness to trade and the peak manufacturing share. We proxy economic openness via the self-expenditure share of a given economy. Figure 3, Panel A, plots the relationship between self-expenditure share in 2000 and the peak manufacturing share attained for all developing economies for which we have data. We observe that more open economies tend to have higher peak manufacturing shares. This is consistent with the observation that developing economies have a comparative advantage in manufacturing.³ If an economy is more open it is more able to take advantage of its comparative advantage, producing manufactured goods and exchanging them for agricultural and service goods.




FIGURE 3: INDUSTRIAL PROFILES OF DEVELOPING ECONOMIES

Notes: Each point is a developing economy. Panel A plots the peak manufacturing employment share against the self-expenditure share in 2000. Panel B plots the peak against the ratio of manufacturing to agricultural labor productivity in 2000. The peak is smoothed over a five-year moving average. Dashed lines are 95 percent confidence bands. The slope of the OLS is -0.407 (s.e. 0.219) in Panel A and -0.0052 (s.e. 0.0021) in Panel B.

How much does trade shape industrialization? As discussed in the introduction, if countries are autarkic, then faster manufacturing productivity growth should lead to a lower peak manufacturing share. If countries are open, then faster manufacturing productivity growth should lead to a higher peak share (due to specialization). Huneus and Rogerson's (2024) Proposition 2 illustrates this relationship in the closed economy case. In their closed-economy benchmark, they prove that the peak manufacturing share is *decreasing* in $g \equiv g_m - g_s$ and *increasing* in g_a . Sposi, Yi, and Zhang (2026) studies the open-economy counterpart to this point. Their trade-versus-autarky counterfactuals predict that countries


³Fujiwara and Matsuyama (2024) for instance discusses this point at length.

with faster manufacturing productivity growth reach higher peaks in Chinese and Indian examples.


We empirically test this in Figure 4. For our 43-country panel, we compute country-level annualized labor productivity growth rates (g_a, g_m, g_s) , measured as real value added per worker. We then regress the peak manufacturing employment share of each country on three relative average manufacturing growth rates for each country: $g_m - \bar{g}$, $g_m - g_a$ and $g_m - g_s$. Our sample is the 29 countries whose peak falls inside their sample window. We overlay the data with the peaks predicted by the model of Huneus and Rogerson (2024) (HR). We start the model from each country’s observed initial sectoral employment shares and feed own growth rates in the same manner as HR. The model predicts negative slopes, as in their Proposition 2. In the data the slopes are flat.⁴ These regressions suggest that in the data, countries that get relatively better at manufacturing do not reach lower manufacturing peaks, as the closed economy predicts.⁵ This is consistent with open-economy forces offsetting the closed-economy mechanism, as in Sposi, Yi, and Zhang (2026).

⁴Appendix Figure D.1 repeats this exercise over a full sample and replaces labor productivity with inverted TFP. With inverted TFP results are similar. All four data slopes are positive with marginally better fits.


⁵Appendix Figures D.4 and D.5 show the full 43 countries which have positive data slopes.



A: Relative to overall productivity growth



B: Relative to agricultural productivity growth



C: Relative to services productivity growth

—◇— Model: HR (2024) □ Model: HR (2024), peak at window edge • Data

FIGURE 4: PEAK MANUFACTURING SHARES AND PRODUCTIVITY GROWTH

Notes: Peak manufacturing employment share against manufacturing productivity growth relative to that country's overall productivity growth (Panel A), agricultural productivity growth (Panel B) and services productivity growth (Panel C). The sample consists of the 29 countries whose peak lies inside their sample window of their first and last recorded data in GGDC. Productivity growth is labor productivity, measured as changes in sector level real GDP per person. This is the same definition and production function as used in Huneus and Rogerson (2024) (HR). The peak is measured over a five-year moving average. Each panel reports the OLS slope, standard error and the correlation, for the data (black) and the HR model (orange). Model peaks are computed from HR, started from each country's observed employment shares at the start of the sample.

One of the most striking features of the data is also the relationship between the peak manufacturing employment share and the manufacturing-agricultural LRP productivity gap. We plot this in Figure 3, Panel B, for all the developing economies for which we have data. Figure 3, Panel B, shows that there is a negative relationship between the LRP productivity gap and attained peak manufacturing shares. We interpret this as highly suggestive evidence of a wedge between agricultural and non-agricultural labor (either a migration barrier or a difference in worker composition) that constrains the manufacturing sector. For instance, countries like Tanzania exhibit huge LRP productivity differentials of over 10-fold between agricultural and manufacturing workers. At the same time, these countries exhibit peak manufacturing employment shares of below 10 percent - very low. This is consistent with the idea that where laborers are inefficiently locked into agriculture the “push” of labor from agriculture into manufacturing is insufficient to counterbalance the “pull” from services.

3 Theoretical Framework

3.1 Static Equilibrium

We follow CDP but generalize the model to feature structural change. Households are indexed by region n only. Within a region, workers reallocate frictionlessly across sectors, so a single regional wage w_t^n clears the local labor market. Per-capita supernumerary income is $\tilde{Y}_t^n \equiv w_t^n - P_t^{n,a}\bar{c}_a - P_t^{n,s}\bar{c}_s = (1 - S_t^n)w_t^n$, where we define the per-sector *Stone–Geary share* for $j \in J_{SG} \equiv \{a, s\}$ and the net share as

$$s_t^{n,a} \equiv \frac{P_t^{n,a}\bar{c}_a}{w_t^n} \in (0, 1), \quad s_t^{n,s} \equiv \frac{P_t^{n,s}\bar{c}_s}{w_t^n} \in (-\infty, 0], \quad S_t^n \equiv s_t^{n,a} + s_t^{n,s}. \quad (1)$$

By construction, the agricultural shift is a *subsistence floor* ($s_t^{n,a} > 0$) and the services shift is a *luxury premium* ($s_t^{n,s} < 0$). Manufacturing has $\bar{c}_m = 0$.

The household’s real consumption is the supernumerary CES dual,

$$C_t^n = \frac{w_t^n - P_t^{n,a}\bar{c}_a - P_t^{n,s}\bar{c}_s}{\tilde{P}_t^n} = \frac{(1 - S_t^n)w_t^n}{\tilde{P}_t^n}.$$

The associated sequential value function and instantaneous utility are

$$V_t^n(\{C_t^m\}) = \mathbb{E}_t \left[\sum_{s=0}^{\infty} \beta^s U(C_{t+s}^m) \right], \quad U(C_t^m) = \log C_t^m,$$

with a Stone–Geary CES aggregator that introduces a per-capita subsistence floor $\bar{c}_a > 0$ on

agriculture and a luxury shift $\bar{c}_s < 0$ on services:

$$C_t^n = \left(\alpha_a^{1/\sigma} (c_t^{n,a} - \bar{c}_a)^{(\sigma-1)/\sigma} + \alpha_m^{1/\sigma} (c_t^{n,m})^{(\sigma-1)/\sigma} + \alpha_s^{1/\sigma} (c_t^{n,s} - \bar{c}_s)^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)},$$

with the subsistence constraints $c_t^{n,a} \geq \bar{c}_a$ and $c_t^{n,s} \geq 0$. We restrict attention to allocations where the services expenditure share, $\Lambda_t^{n,s}$, is positive. Defining the supernumerary quantities as $\tilde{c}_t^{n,j} \equiv c_t^{n,j} - \bar{c}_j$ for all j (where $\bar{c}_m = 0$), then the bundle is a standard homothetic CES in $\tilde{c}$. This supernumerary CES dual price index has the usual closed form,

$$\tilde{P}_t^n = \left(\sum_{k=1}^J \alpha_k (P_t^{nk})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \quad (2)$$

First-order conditions yield supernumerary CES expenditure shares reminiscent of the usual homothetic CES forms, and expenditure shares $\Lambda_t^{n,j} \equiv P_t^{n,j} \tilde{c}_t^{n,j} / w_t^n$ that split into a Stone–Geary piece (only for $j \in J_{SG}$) and a supernumerary piece:

$$\tilde{\lambda}_t^{n,j} \equiv \frac{\alpha_j (P_t^{n,j})^{1-\sigma}}{\sum_k \alpha_k (P_t^{n,k})^{1-\sigma}}, \quad \Lambda_t^{n,j} = \mathbf{1}\{j \in J_{SG}\} s_t^{n,j} + (1 - S_t^n) \tilde{\lambda}_t^{n,j}, \quad \sum_j \Lambda_t^{n,j} = 1. \quad (3)$$

Setting $\bar{c}_a = \bar{c}_s = 0$ recovers the homothetic CES in the appendix of CDP with $\Lambda^{n,j} = \tilde{\lambda}^{n,j} = \lambda^{n,j}$ and $\tilde{P}^n = P^n$.

Production follows the multisector form of CDP. The output for a producer of a variety is given by

$$q_t^{nj} = z^{nj} (A_t^{nj} l_t^{nj})^{\gamma_t^{nj}} \left(\prod_{k=1}^J (M_t^{nj,nk})^{\gamma_t^{nj,nk}} \right).$$

Note that $M_t^{nj,nk}$ represents producer materials sourced for sector j from sector k in region n . These materials themselves come from local output of sector k . The total local sectoral output in sector k divides into that which goes to total consumption locally, and that which goes into this material bundle. Consumption is also not traded. The final good that supplies both consumption and materials is the CES aggregator,

$$Q_t^{nj} = \left(\int (\tilde{q}_t^{nj}(z^j))^{1-1/\eta^{nj}} d\phi^j(z^j) \right)^{\frac{\eta^{nj}}{\eta^{nj}-1}}$$

We also omit capital, modeling sectoral gross output as labor value added combined through an input–output network. Herrendorf, Herrington, and Valentinyi (2015) show that the main technological driver of structural change is differential (labor-augmenting)

productivity growth between sectors. They argue that cross-sector differences in capital intensity are second order and a common capital share fits the data about as well as flexible CES estimates. SYZ follow this route by keeping capital, but fixing a common capital share of 0.33 in every sector, so that sectoral differences enter through intermediate intensity.

The unit price of the overall input bundle is:

$$x_t^{nj} = B^{nj} (w_t^n)^{\gamma^{nj}} \prod_{k=1}^J (P_t^{nk})^{\gamma^{nj,nk}}. \quad (4)$$

Crucially, by the Eaton-Kortum structure of trade in this model the unit cost of a variety is $\frac{x_t^{nj}}{z^{nj} (A_t^{nj})^{\gamma^{nj}}}$. If iceberg trade costs are denoted by κ then

$$p_t^{nj}(z^j) = \min_i \left\{ \frac{\kappa_t^{nj,ij} x_t^{ij}}{z^{ij} (A_t^{ij})^{\gamma^{ij}}} \right\}.$$

Now, let Q_t^{nj} be the quantity produced of aggregate sectoral goods j in region n and $\tilde{q}_t^{nj}(z^j)$ be the quantity demanded of an intermediate good of a variety from the lowest-cost supplier. Given the Fréchet structure we obtain the price of the sectoral aggregate good as:

$$P_t^{nj} = \Gamma^{nj} \left(\sum_{i=1}^N (x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma^{ij}} \right)^{\frac{-1}{\theta^j}} \quad (5)$$

where $\Gamma^{nj} = \Gamma(1 + (1 - \eta^{nj})/\theta^j)^{1/(1-\eta^{nj})}$ is the Fréchet-CES constant (Γ_{EK} in Appendix A.5, where the price index and trade shares are derived from the Fréchet primitives).

$$\pi_t^{nj,ij} = \frac{(x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma^{ij}}}{\sum_{m=1}^N (x_t^{mj} \kappa_t^{nj,mj})^{-\theta^j} (A_t^{mj})^{\theta^j \gamma^{mj}}}. \quad (6)$$

With these objects defined we can move to the main budget constraint equation of this model. Region n expenditure on sector j is given by final demand plus materials demand. Final demand consists of household consumption out of income, $w_t^n L_t^n$, and a period-by-period transfer D_t^n that tracks trade imbalances. Writing these together gives the expenditure market clearing condition:

$$X_t^{nj} = \sum_{k=1}^J \gamma^{nk,nj} \sum_{i=1}^N \pi^{ik,nk} X_t^{ik} + \Lambda_t^{n,j} (w_t^n L_t^n + D_t^n), \quad (7)$$

where $\Lambda_t^{n,j} = \mathbf{1}\{j \in J_{SG}\} s_t^{n,j} + (1 - S_t^n) \tilde{\lambda}_t^{n,j}$ from above, and the regional wage bill $w_t^n L_t^n =$

$\sum_k w_t^{nk} L_t^{nk}$ replaces the sum-over-sectors form because $w_t^{nk} = w_t^n$ for all k . We allow for exogenous trade imbalances via D_t^n . For simplicity, households consume out of wage income only, while D_t^n then finances additional expenditure that is spent across sectors with the same expenditure shares $\Lambda_t^{n,j}$ but does not enter household utility. Each country's trade imbalance is computed as $D_t^c = (d_t^c - \bar{d}_t) \sum_{m \in c} w_t^m L_t^m$, where d_t^c is the smoothed-current account deficit (Müller *et al.* 2025) as a share of GDP and $\bar{d}_t$ is the income-adjusted world mean, $\bar{d}_t \equiv \sum_c d_t^c \sum_{m \in c} w_t^m L_t^m / \sum_m w_t^m L_t^m$, so that we force imbalances to sum to zero each period. We allocate D_t^c across the country's two regions in proportion to the regional wage bill, $D_t^n = (w_t^n L_t^n / \sum_{m \in c} w_t^m L_t^m) \cdot D_t^c$. When $\bar{c}_a = \bar{c}_s = 0$, $s_t^{n,a} = s_t^{n,s} = 0$, and $\Lambda_t^{n,j} = \tilde{\lambda}_t^{n,j} = \lambda_t^{n,j}$.

Finally, static labor market clearing aggregates the per-sector clearing of CDP to the region level:

$$w_t^n L_t^n = \sum_{j=1}^J \gamma^{nj} \sum_{i=1}^N \pi_t^{ij,nj} X_t^{ij}. \quad (8)$$

Sectoral employment L_t^{nj} is no longer a state variable; it is a static residual pinned by labor demand at the regional wage,

$$L_t^{nj} = \frac{\gamma^{nj} \sum_{i=1}^N \pi_t^{ij,nj} X_t^{ij}}{w_t^n}, \quad \sum_{j=1}^J L_t^{nj} = L_t^n. \quad (9)$$

3.2 Dynamic Equilibrium

Workers compare *regions* so that their idiosyncratic taste shock is drawn over destination locations $i \in \{1, \dots, N\} (= \{R, U\})$, and migration wedges $\tau_t^{n,i}$ depend only on origin and destination region. Then the sequential problem after the taste shock is realized is

$$\begin{aligned} v_t^n &= U(C_t^n) + \max_i \left\{ \beta \mathbb{E} [v_{t+1}^i] - \tau_t^{n,i} + \nu \epsilon_t^i \right\} \\ &= \log \left(\frac{(1 - S_t^n) w_t^n}{\tilde{P}_t^n} \right) + \max_i \left\{ \beta \mathbb{E} [v_{t+1}^i] - \tau_t^{n,i} + \nu \epsilon_t^i \right\}. \end{aligned}$$

Before learning the taste shock, integrating the Gumbel max yields

$$\begin{aligned} V_t^n &= U(C_t^n) + \mathbb{E}_\nu \left[\max_i \left\{ \beta V_{t+1}^i - \tau_t^{n,i} + \nu \epsilon_t^i \right\} \right] \\ &\equiv \log \left(\frac{(1 - S_t^n) w_t^n}{\tilde{P}_t^n} \right) + \nu \log \left(\sum_{i=1}^N \exp \left(\frac{\beta V_{t+1}^i - \tau_t^{n,i}}{\nu} \right) \right). \end{aligned} \quad (10)$$

The migration mass then takes the standard logit form, now over destination regions

only:

$$\begin{aligned}\mu_t^{n,i} &= \mathbb{P}(i = \arg \max \{ \beta V_{t+1}^i - \tau_t^{n,i} + \nu \epsilon_t^i \}) \\ &= \frac{\exp\left(\frac{\beta V_{t+1}^i - \tau_t^{n,i}}{\nu}\right)}{\sum_{m=1}^N \exp\left(\frac{\beta V_{t+1}^m - \tau_t^{n,m}}{\nu}\right)}.\end{aligned}\tag{11}$$

Combined with regional labor market accounting

$$L_{t+1}^n = \sum_{i=1}^N \mu_t^{i,n} L_t^i.\tag{12}$$

Within-region sectoral employment L_{t+1}^{nj} is then determined in the static equilibrium.

The state variable of the model at time t is the distribution of workers in each location $L_t = \{L_t^n\}_{n=1}^N$. The fundamentals of the model are denoted by Θ_t and are $\Theta_t \equiv (A_t, \kappa_t, \gamma_t, \tau_t, d_t)$. That is, productivities, trade costs, IO and value added weights, migration costs, and trade imbalances. We now state the equilibrium of the model in a similar fashion to Caliendo, Dvorkin, and Parro (2019). We choose world income as the numeraire, $\sum_{n=1}^N w_t^n L_t^n = 1$.

Definition 3.1 (Static equilibrium). Given (L_t, Θ_t) , a *static equilibrium* is a vector of wages, $w(L_t, \Theta_t)$, prices $P_t^{nj}(L_t, \Theta_t)$, trade shares $\pi_t^{nj,ij}(L_t, \Theta_t)$ and expenditures $X_t^{nj}(L_t, \Theta_t)$, that solves equations (1)-(9).

Definition 3.2 (Dynamic equilibrium). Given $(L_0, \{\Theta_t\}_{t=0}^\infty)$, a *dynamic equilibrium* is a sequence $\{L_t, \mu_t, V_t, w(L_t, \Theta_t)\}_{t=0}^\infty$, that solves equilibrium conditions (10)-(12) and the static equilibrium at each t .

Note that the migration wedge τ_t is agnostic to whether the cause of rural-urban wage gaps is a skill gap, migration frictions, or some other wedge. Suppose for instance that most of the gap were due to worker composition. Our migration wedge is then simply the marginal worker's skill discount. If rural workers face no costs in migrating but would earn only some fraction h of the prevailing urban wage if they migrated, then a one-off wedge of $\frac{\beta}{1-\beta} \log(1/h)$ would reproduce the decisions, flows and wage gap exactly.⁶ This point is discussed in Alvarez (2020) and Caliendo, Dvorkin, and Parro (2019, Appendix G), who note that wages and flows cannot distinguish the two economies.⁷ What changes is the interpretation of the cause of premature deindustrialization and what policies would reduce

⁶This exact formula holds under log utility, homothetic CES and no return migration.

⁷Caliendo, Dvorkin, and Parro (2019), footnote 53, offers this reading of their own sectoral costs too. They note that the wedges “capture, in part, the skill composition in each industry, and, therefore, how costly it is for certain skill groups to switch across industries”.

the wedge. Our paper shows that late developers have low manufacturing peaks because labor does not flow to manufacturing when manufacturing productivity rises. Whether a skill gap or a migration cost is responsible is outside the scope of this paper and is a contentious topic in the development literature.

3.3 Accounting Decomposition of Manufacturing Share Changes

Before solving the model, we show how to decompose the budget constraint and employment stocks to see how migration has shaped the manufacturing employment shares across countries. This section draws closely on the insight of Smitskova (2024). To disentangle the effect of migration on manufacturing employment we simply add a migration channel to the decomposition. We present the overall idea here. The value-added steps follow Smitskova (2024). Figure 5 presents the headline results.

Region n in country c has budget identity

$$X_t^{nj} = \sum_{k=1}^J \gamma_t^{nk,nj} \sum_{i=1}^N \pi_t^{ik,nk} X_t^{ik} + \Lambda_t^{n,j} w_t^n L_t^n, \quad (13)$$

This identity suppresses the trade imbalances D_t^n for simplicity, which we add below. Let Y^{nj} be the gross output of sector j in region n , $Y^{nj} = \sum_i \pi^{ij,nj} X^{ij}$. Substituting into equation (13), we obtain

$$Y^{nj} = \sum_i \pi^{ij,nj} \Lambda^{i,j} w^i L^i + \sum_{i,k} \pi^{ij,nj} \gamma^{ik,ij} Y^{ik}. \quad (14)$$

Using the product rule on (14) and grouping by shock,

$$\begin{aligned} \tilde{Y}^{nj} = & \underbrace{\sum_i \left(\phi_{nj,i}^F + \sum_k \phi_{nj,ik}^M \right) \tilde{\pi}^{ij,nj}}_{T_{nj}^\Pi : \text{specialization}} + \underbrace{\sum_i \phi_{nj,i}^F \tilde{\Lambda}^{i,j}}_{T_{nj}^\Lambda : \text{preferences}} + \underbrace{\sum_{i,k} \phi_{nj,ik}^M \tilde{\gamma}^{ik,ij}}_{T_{nj}^\Gamma : \text{IO}} \\ & + \underbrace{\sum_i \phi_{nj,i}^F \widetilde{(wL)}^i}_{T_{nj}^{wL} : \text{labor income}} + \underbrace{\sum_{i,k} \phi_{nj,ik}^M \tilde{Y}^{ik}}_{\text{IO-feedback}}. \end{aligned} \quad (15)$$

where

$$\phi_{nj,i}^F \equiv \frac{\pi^{ij,nj} \Lambda^{i,j} w^i L^i}{Y^{nj}}, \quad \phi_{nj,ik}^M \equiv \frac{\pi^{ij,nj} \gamma^{ik,ij} Y^{ik}}{Y^{nj}}. \quad (16)$$

Since value added shares are given by $v^{nj} = \gamma^{nj} Y^{nj} / (w^n L^n)$, we can add trade imbalances

and analogously decompose value added into components

$$\Delta v_m^c = \underbrace{\Delta v_m^c(\tilde{\Pi})}_{\text{specialization}} + \underbrace{\Delta v_m^c(\tilde{\text{sec}})}_{\text{secular}} + \underbrace{\Delta v_m^c(\tilde{D})}_{\text{borrowing}} + \varepsilon_m^c. \quad (17)$$

To map value added shares into employment shares it is convenient to define an industry's value added per worker, relative to economy average as

$$\psi_m^c \equiv \frac{V_m^c/L_m^c}{V^c/E^c}, \quad V_m^c \equiv \sum_{n \in R(c)} V^{nm}, \quad V^c \equiv \sum_{n \in R(c)} \sum_k V^{nk}.$$

By construction we have

$$v_m^c = S_m^c \cdot \psi_m^c,$$

where $S_m^c \equiv \frac{L_m^c}{E^c}$ is the fraction of labor in sector m in the entire economy. One can then show that

$$\Delta S_m^c = \underbrace{\frac{\Delta S_m^c / \Delta \ln S_m^c}{\Delta v_m^c / \Delta \ln v_m^c}}_{\kappa^c} \Delta v_m^c - \underbrace{\frac{\Delta S_m^c}{\Delta \ln S_m^c} \Delta \ln \psi_m^c}_{\text{Baumol wedge}}. \quad (18)$$

In this paper, we divide each country into two regions, R and U , with employment shares $\omega_R + \omega_U = 1$, so that $S_m^c = \omega_R s_R + \omega_U s_U$. Writing in changes, the product rule gives the identity

$$\begin{aligned} \Delta S_m^c &= \underbrace{(\omega_R^0 \Delta s_R + \omega_U^0 \Delta s_U)}_{\text{within}} + \underbrace{(s_R^0 \Delta \omega_R + s_U^0 \Delta \omega_U)}_{\text{between}} + \underbrace{(\Delta \omega_R \Delta s_R + \Delta \omega_U \Delta s_U)}_{\text{interaction}} \\ &= (\omega_R^0 \Delta s_R + \omega_U^0 \Delta s_U) + (s_U^0 - s_R^0) \Delta \omega_U + \Delta \omega_U (\Delta s_U - \Delta s_R) \end{aligned}$$

The between term splits into a net rural-to-urban flow part, Mig^c , and a natural growth remainder, Nat^c that is the gap between observed and model-implied net flows:

$$\text{Mig}^c \equiv (s_U^0 - s_R^0) \Delta \omega_U^{\text{mig}}, \quad \text{Nat}^c \equiv (s_U^0 - s_R^0) (\Delta \omega_U - \Delta \omega_U^{\text{mig}}) \quad (19)$$

In this way, we can divide changes in manufacturing employment into a migration, natural migration-fit remainder, within, and interaction component. Then

$$\Delta S_m^c = \text{Mig}^c + \text{Nat}^c + \text{Int}^c + \text{within}^c.$$

Adding and subtracting $\kappa^c \Delta v_m^c$ we get, up to the term $\kappa^c \varepsilon_m^c$, which we fold into R^c :

$$\begin{aligned}\Delta S_m^c &= \kappa^c \Delta v_m^c + (\text{within} - \kappa^c \Delta v_m^c) + \text{Mig}^c + \text{Nat}^c + \text{Int}^c \\ &= \underbrace{\kappa^c \Delta v_m^c(\tilde{\Pi})}_{\text{spec.}} + \underbrace{\kappa^c \Delta v_m^c(\widetilde{\text{sec}})}_{\text{secular}} + \underbrace{\kappa^c \Delta v_m^c(\tilde{D})}_{\text{borrow}} + \underbrace{(\text{within} - \kappa^c \Delta v_m^c)}_{\text{Baumol}} \\ &\quad + \text{Mig}^c + \text{Nat}^c + \text{Int}^c\end{aligned}$$

In our specific case we use WIOD for the value added decomposition and GGDC for the employment changes and Baumol wedge. This means we have a residual term that incorporates WIOD–GGDC data discrepancies and second-order terms $\kappa^c \varepsilon_m^c$. This residual is,

$$R^c \equiv \kappa^c \left(\Delta v_{m,\text{ggdc}}^c - [\Delta v_m^c(\tilde{\Pi}) + \Delta v_m^c(\widetilde{\text{sec}}) + \Delta v_m^c(\tilde{D})] \right)$$

and so our final decomposition is

$$\begin{aligned}\Delta S_m^c &= \underbrace{\kappa^c \Delta v_m^c(\tilde{\Pi})}_{\text{spec.}} + \underbrace{\kappa^c \Delta v_m^c(\widetilde{\text{sec}})}_{\text{secular}} + \underbrace{\kappa^c \Delta v_m^c(\tilde{D})}_{\text{borrow}} + \underbrace{(\text{within} - \kappa^c \Delta v_{m,\text{ggdc}}^c)}_{\text{Baumol}} \\ &\quad + \text{Mig}^c + \text{Nat}^c + \text{Int}^c + R^c\end{aligned}$$

We map this decomposition to the intersection of countries that are present in both WIOD and GGDC data. Results are pictured below in Figure 5, and Table 1 reports the migration channel by country. The reallocation of labor from rural to urban areas has been a large component of rising manufacturing shares in the emerging markets we present. The migration channel averages 7.1 percentage points in the emerging markets in our panel. We find that the channel accounts for +11.9 points in China, +10.2 in Brazil, +8.0 in Korea, +5.2 in Mexico, +3.8 in Taiwan, and +3.4 in India. In advanced economies the migration channel contributes -0.1 on average. Migration accounts for about 53% of the total change in manufacturing employment share in China, 42% in Taiwan and 37% in India, 113% in Korea and 298% in Mexico. Brazil’s net change is roughly zero, so the +10.2 pp migration push is cancelled by the other channels rather than appearing in the headline number. Our decomposition is partial equilibrium and mechanical. Nevertheless, we believe the large contribution of migration to manufacturing in emerging markets is suggestive evidence that spatial immobility could be first order to understanding labor market dynamics in global manufacturing. Section 6 complements this mechanical accounting with quasi-experimental evidence on the same push in Brazil.




FIGURE 5: DECOMPOSITION OF CHANGES IN MANUFACTURING EMPLOYMENT SHARES

Notes: Smitkova (2024) style decomposition of manufacturing employment shares into constituent components for countries present in both WIOD and GGDC. Productivity-bridge decomposition of the 1970 to 2014 change in each country's manufacturing (ISIC C–F) employment share, in percentage points, for the countries present in both WIOD and GGDC, ordered from the largest increase to the largest decline. Bars stack the channels of Section 3.3: trade specialization, secular trends (preferences, input-output weights and wages), borrowing (trade imbalances), the productivity wedge (net of migration, natural growth and interaction), rural-to-urban migration, other reallocation, and the WIOD versus GGDC residual. The channels sum exactly to the observed change, marked by a cross.

TABLE 1: CONTRIBUTION OF RURAL-TO-URBAN MIGRATION TO MANUFACTURING
EMPLOYMENT SHARE CHANGES BY COUNTRY

Country	Mig ^c (pp)	Observed ΔS_m^c (pp)	Share (%)
<i>Emerging markets</i>			
China	+11.9	+22.6	53
Brazil	+10.2	−0.3	−3400
Korea	+8.0	+7.0	113
Mexico	+5.2	+1.7	298
Taiwan	+3.8	+8.9	42
India	+3.4	+9.0	37
<i>Average</i>	+7.1		
<i>Advanced economies</i>			
Spain	+2.2	−14.7	−15
Denmark	+0.6	−18.0	−3
Italy	+0.4	−11.9	−4
France	+0.3	−18.5	−2
Sweden	+0.0	−16.3	−0.3
Germany	−0.1	−30.2	+0.3
USA	−0.2	−15.2	+1
UK	−0.2	−25.6	+1
Australia	−0.3	−21.2	+1
Finland	−0.7	−17.6	+4
Japan	−1.2	−10.8	+11
Portugal	−1.8	−12.3	+15
<i>Average</i>	−0.1		

Notes: Mig^c is the rural-to-urban migration term from the employment-share decomposition in Figure 5, in percentage points of the 1970–2014 change in the manufacturing (ISIC C–F) employment share. Observed ΔS_m^c is the total observed change. Share is defined to mean $\equiv \text{Mig}^c / \Delta S_m^c$. Averages use uniform weights within each cohort.

4 Model Solution

The previous section presents a mechanical decomposition of manufacturing employment share changes into secular, trade, and migration channels. However, in order to study the full general equilibrium role of migration and trade in shaping structural change we must solve the full model we have laid out. We take the approach of Caliendo, Dvorkin, and Parro (2019) by solving the model using dynamic hat algebra. Since our model features Stone-Geary CES preferences and time varying IO weights and migration wedges we cannot take the CDP dynamic hat algebra off the shelf. Unfortunately, time varying IO weights in

our context mean we need to track the level of TFP as an additional initial condition. The solution to the full static problem is presented in the following proposition.

Proposition 1. *Given the static allocation $\{L_t^n, L_t^{nj}, \pi_t, X_t, w_t^n, P_t, A_t, s_t^a, s_t^s, \tilde{\lambda}_t\}$, the solution to the temporary equilibrium at $t+1$ for given changes $\dot{L}_{t+1}^n, \dot{\Theta}_{t+1}$ and imbalances d_{t+1} solves*

$$\dot{x}_{t+1}^{nj} = \dot{B}_{t+1}^{nj} (\dot{w}_{t+1}^n)^{\gamma_{t+1}^{nj}} \prod_{k=1}^J (\dot{P}_{t+1}^{nk})^{\gamma_{t+1}^{nj,nk}} \prod_{k=1}^J \left(\frac{P_t^{nk}}{w_t^n} \right)^{\Delta \gamma_{t+1}^{nj,nk}} \quad (\text{P1.1})$$

$$\dot{P}_{t+1}^{nj} = \left(\sum_{i=1}^N \pi_t^{nj,ij} (\dot{x}_{t+1}^{ij} \dot{\kappa}_{t+1}^{nj,ij})^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (A_t^{ij})^{\theta^j \Delta \gamma_{t+1}^{ij}} \right)^{-1/\theta^j} \quad (\text{P1.2})$$

$$\pi_{t+1}^{nj,ij} = \pi_t^{nj,ij} \left(\frac{\dot{x}_{t+1}^{ij} \dot{\kappa}_{t+1}^{nj,ij}}{\dot{P}_{t+1}^{nj}} \right)^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (A_t^{ij})^{\theta^j \Delta \gamma_{t+1}^{ij}} \quad (\text{P1.3})$$

$$s_{t+1}^{n,a} = s_t^{n,a} \frac{\dot{P}_{t+1}^{n,a}}{\dot{w}_{t+1}^n}, \quad s_{t+1}^{n,s} = s_t^{n,s} \frac{\dot{P}_{t+1}^{n,s}}{\dot{w}_{t+1}^n}, \quad S_{t+1}^n = s_{t+1}^{n,a} + s_{t+1}^{n,s} \quad (\text{P1.4})$$

$$\tilde{\lambda}_{t+1}^{n,j} = \tilde{\lambda}_t^{n,j} \frac{(\dot{P}_{t+1}^{n,j})^{1-\sigma}}{\sum_k \tilde{\lambda}_t^{n,k} (\dot{P}_{t+1}^{n,k})^{1-\sigma}}, \quad \Lambda_{t+1}^{n,j} = \mathbf{1}\{j \in J_{SG}\} s_{t+1}^{n,j} + (1 - S_{t+1}^n) \tilde{\lambda}_{t+1}^{n,j} \quad (\text{P1.5})$$

$$X_{t+1}^{nj} = \sum_{k=1}^J \gamma_{t+1}^{nk,nj} \sum_{i=1}^N \pi_{t+1}^{ik,nk} X_{t+1}^{ik} + \Lambda_{t+1}^{n,j} \left(\dot{w}_{t+1}^n \dot{L}_{t+1}^n w_t^n L_t^n + D_{t+1}^n \right) \quad (\text{P1.6})$$

$$\dot{w}_{t+1}^n \dot{L}_{t+1}^n w_t^n L_t^n = \sum_{j=1}^J \gamma_{t+1}^{nj} \sum_{i=1}^N \pi_{t+1}^{ij,nj} X_{t+1}^{ij} \quad (\text{P1.7})$$

$$L_{t+1}^{nj} = \frac{\gamma_{t+1}^{nj} \sum_{i=1}^N \pi_{t+1}^{ij,nj} X_{t+1}^{ij}}{\dot{w}_{t+1}^n w_t^n} \quad (\text{P1.8})$$

$$\dot{B}_{t+1}^{nj} = \frac{(\gamma_{t+1}^{nj})^{-\gamma_{t+1}^{nj}} \prod_k (\gamma_{t+1}^{nj,nk})^{-\gamma_{t+1}^{nj,nk}}}{(\gamma_t^{nj})^{-\gamma_t^{nj}} \prod_k (\gamma_t^{nj,nk})^{-\gamma_t^{nj,nk}}} \quad \tilde{\lambda}_{t+1}^{nj} = \dot{\tilde{\lambda}}_{t+1}^{nj} \tilde{\lambda}_t^{nj}, \quad \dot{\tilde{\lambda}}_{t+1}^{nj} = \frac{(\dot{P}_{t+1}^{nj})^{1-\sigma}}{\sum_{k=1}^J \tilde{\lambda}_t^{nk} (\dot{P}_{t+1}^{nk})^{1-\sigma}}.$$

where $\Delta \gamma_{t+1}^{nj,nk} \equiv \gamma_{t+1}^{nj,nk} - \gamma_t^{nj,nk}$ and

$$\tilde{\lambda}_0^{nj} = \frac{\alpha_j (P_0^{nj})^{1-\sigma}}{\sum_{k=1}^J \alpha_k (P_0^{nk})^{1-\sigma}}, \quad s_0^{n,a} = \frac{P_0^{n,a} \bar{c}_a}{w_0^n}, \quad s_0^{n,s} = \frac{P_0^{n,s} \bar{c}_s}{w_0^n}.$$

The supernumerary CES dual price index $\tilde{P}_t^n$ has the same closed form as the homothetic

P_t^n in (α, P^{nk}) , so it satisfies $(\dot{P}_{t+1}^n)^{1-\sigma} = \sum_k \tilde{\lambda}_t^{nk} (\dot{P}_{t+1}^{nk})^{1-\sigma}$.

Proposition 1 mimics the dynamic hat proposition 1 of Caliendo, Dvorkin, and Parro (2019) with some changes. Three changes are particularly noteworthy. First, our setting has no structures. Instead, we match trade imbalances through the exogenous wedge D_t^n . Value added in our model is only labor. This contrasts with CDP, which employs structures in value added with share ξ , giving $(\dot{L})^{\xi\gamma}$ in the cost bundle. Second, given the more secular nature of our paper we have time-varying IO weights. With time-varying IO weights, $\dot{x}$ picks up a correction term $\prod_k (P_t^{nk}/w_t^n)^{\Delta\gamma_{t+1}^{nj,nk}}$. Price indices and trade shares also pick up an extra factor $(A_t^{ij})^{\theta^j \Delta\gamma_{t+1}^{ij}}$, which involves the *level* of productivity. Finally, since our framework also features non-homothetic Stone-Geary preferences, we pick up extra per-region scalar state variables, $s_t^{n,a} > 0$ and $s_t^{n,s} < 0$.

Although the above proposition seems messier than Proposition 1 of CDP, the extra terms are small corrections and the core idea of the solver is the same. Below we present the solution to the dynamic problem as a second proposition. This follows CDP as well. It also features small correction terms due to our inclusion of time varying migration wedges and Stone-Geary CES preferences.

Proposition 2. *Conditional on an initial allocation of the economy, $(\{L_0^n\}, \pi_0, X_0, \{\mu_{-1}^{n,i}\}, \tilde{\lambda}_0, s_0^a, s_0^s, w_0, P_0, A_0, \gamma_0)$, and given an anticipated path of fundamentals $\{\dot{\Theta}_t\}_{t=1}^\infty$ and imbalances $\{d_t\}_{t=1}^\infty$, the dynamic equilibrium in time differences solves*

$$\mu_{t+1}^{n,i} = \frac{\mu_t^{n,i} \exp((V_{t+2}^i - V_{t+1}^i)\beta/\nu) \exp(-\Delta\tau_{t+1}^{n,i}/\nu)}{\sum_{m=1}^N \mu_t^{n,m} \exp((V_{t+2}^m - V_{t+1}^m)\beta/\nu) \exp(-\Delta\tau_{t+1}^{n,m}/\nu)} \quad (\text{P2.1})$$

$$\exp(V_{t+1}^n - V_t^n) = \frac{\dot{w}_{t+1}^n}{\dot{P}_{t+1}^n} \left(\sum_{i=1}^N \mu_t^{n,i} \exp((V_{t+2}^i - V_{t+1}^i)\beta/\nu) \exp(-\Delta\tau_{t+1}^{n,i}/\nu) \right)^\nu \quad (\text{P2.2})$$

$$\dot{w}_{t+1}^n = \frac{\dot{w}_{t+1}^n - s_t^{n,a} \dot{P}_{t+1}^{n,a} - s_t^{n,s} \dot{P}_{t+1}^{n,s}}{1 - S_t^n} = \frac{1 - S_{t+1}^n}{1 - S_t^n} \dot{w}_{t+1}^n$$

$$L_{t+1}^n = \sum_{i=1}^N \mu_t^{i,n} L_t^i \quad (\text{P2.3})$$

where $\Delta\tau_{t+1}^{n,i} \equiv \tau_{t+1}^{n,i} - \tau_t^{n,i}$.

When rural-urban wedges are constant over time we have $\Delta\tau = 0$ and the $\exp(-\Delta\tau/\nu)$ factors collapse to 1. This recovers the standard CDP equations from proposition 2 with a Stone-Geary adjustment.

4.1 A Closed-Form Special Case

While we solve the full model above numerically, here we provide a closed-form solution to a special case to highlight the model's central mechanisms. We restrict attention to a special case of the quantitative model in order to obtain analytical expression. Consider a small open economy in agriculture and manufacturing. In this economy, the rural region produces only agriculture and the urban region produces manufacturing and non-traded services. There are no intermediate inputs, no cross-border migration (restrictions (i)–(vi) in Appendix C), and trade is balanced. We use homothetic CES preferences ($\bar{c}_a = \bar{c}_s = 0$). As in Section 3.3, the employment share of sector j is given by $S_{j,t} \equiv S_{j,t}^c = L_{j,t}^c/E^c$. Steady state value omit t such that the initial steady state value is S_j , and the rural share is S_a . Recall that we denote trade elasticities by θ^a, θ^m . At steady state define the objects:

$$\zeta \equiv 1 - \lambda^s \left(1 - \frac{\theta^a}{1 + \theta^a} \frac{w^R}{w^U} \right), \quad B \equiv \lambda^s S_m - \sigma(1 - \lambda^s) S_s, \quad M \equiv \frac{\zeta S_a(1 - S_a)}{\bar{\nu} + (1 - S_a)/(1 + \theta^a)}.$$

We call B the *services pull*. Holding farm employment fixed, a higher urban wage lowers the manufacturing share in proportion to B because it makes non-traded services more expensive. When goods and services are gross complements, i.e., $\sigma < 1$, households spend a larger share of income on services, and services then hire workers away from manufacturing. Agricultural worker wages don't change upon impact, so they buy fewer services. Consequently, the services pull, B is positive if and only if the agricultural employment share is below a threshold $\bar{S}$ (Lemma 3). Analogously, we call M the *farm push*. In the long run, a higher urban wage also draws workers out of agriculture, raising the manufacturing share in proportion to M . At a given urban wage, a share ζ of these workers goes to manufacturing. Part of the inflow goes to services because their higher incomes raise the demand for services. The long-run response of the manufacturing share therefore has the sign of $M - B$. Migration costs affect only $\bar{\nu}$ in M . The object $\bar{\nu}$ measures how strongly workers respond in the long run to the rural-urban wage gap. We discuss this further in Appendix C, equation (61). Finally, we define a “round-trip” migration cost, $\tau \equiv \tau^{RU} + \tau^{UR}$.

We start from a steady state with $S_a \in (0, 1)$ and $S_m, S_s > 0$. At $t = 0$, urban manufacturing productivity rises unexpectedly and permanently by an MIT shock $d \log A^{Um} > 0$. Agents have perfect foresight after the shock. As in CDP, employment in each region is fixed within a period, so workers can only respond to the shock by migrating in later periods. Throughout, we denote by d deviations from the initial steady state. Section 6 examines this migration margin in Brazil, using a labor-saving agricultural shock rather than a manufacturing productivity shock.

Proposition 3. *Consider the special case with the steady state and shock described above and*

restrictions (i)–(vi) stated formally in Appendix C. Let $0 \leq \sigma < 1$, and let $dS_{m,t}/d\log A^{Um}$ be the first-order response of the manufacturing share t periods after the shock. Consider the family of economies that share ν , all other fundamentals, and the initial allocation, and differ only in their round-trip cost, $\tau^{RU} + \tau^{UR} > 0$. Then there exist thresholds $0 < \underline{S} < \bar{S} < 1$, which depend on σ , θ^a , ν , β , and the initial λ^s and w^R/w^U but not on θ^m or on the round-trip cost (Lemma 3 of Appendix C), such that:

(a) If $S_a < \bar{S}$:

- $dS_{m,0}/d\log A^{Um} = -\theta^m S_m B / [(1 + \theta^m)S_m - B] < 0$.

(b) If $\underline{S} < S_a$:

- $dS_{m,t}/d\log A^{Um}$ is strictly increasing in t , with

$$\lim_{t \rightarrow \infty} \frac{dS_{m,t}}{d\log A^{Um}} = \theta^m S_m \frac{M - B}{M - B + (1 + \theta^m)S_m} > 0.$$

- There is a unique date $t^*(\tau) \geq 0$ such that $dS_{m,t}/d\log A^{Um} < 0$ if and only if $t < t^*(\tau)$. We call $t^*(\tau)$ the time to industrialization.
- Time to industrialization, $t^*(\tau)$, is weakly increasing in τ . If also $S_a < \bar{S}$, then $t^*(\tau) \geq 1$ and $t^*(\tau) \rightarrow \infty$ as $\tau \rightarrow \infty$.




FIGURE 6: THE MANUFACTURING SHARE AFTER A MANUFACTURING PRODUCTIVITY SHOCK

Notes: The figure illustrates Proposition 3. The bars depict regions by initial farm employment share S_a where each part of Proposition 3 applies. The curves plot illustrative paths of $dS_{m,t}/d\log A^{Um}$ over time t in each region. The dashed lines depict zero, and “cost” is the round-trip migration cost $\tau^{RU} + \tau^{UR}$.

Condition a) says that upon impact, as long as the agricultural share of the economy is not too large, then services pull immediately acts to deindustrialize the economy. The logic here is standard: higher manufacturing productivity raises urban wages. Since manufacturing prices are set globally, manufacturing prices do not rise but services prices do rise. Because of gross complementarity between goods this increases relative expenditure on services. Because services are non-tradable they cannot be imported so workers move from manufacturing to services. Resources also cannot move from agriculture to manufacturing on impact because migration takes time.

The interesting condition is when $S_a > \bar{S}$ such that the economy actually industrializes on impact. This sheds light on another force that acts in the background. When manufacturing productivity rises, this pulls up urban wages in both manufacturing and services. Prices of agriculture and manufacturing stay constant and are set by the rest of the world, but services prices rise. Agricultural workers are thus worse off in real terms due to more expensive services and consume less services. If agriculture is a sufficiently large portion of the economy then the reduction in agricultural real wages can swamp the increase in urban real wages and decrease overall demand for services. This creates industrialization on impact.

Condition b) implies three separate claims. Each relying on there being a sufficient amount of agricultural labor for the farm push to bite. First, if agricultural labor is abundant enough then in infinite time the migration push eventually dominates the services pull for any migration cost. Second, for any round trip migration cost there is a unique threshold date where before this date services pull dominates and deindustrializes the economy, but after this date, agricultural push eventually supplies enough labor to dominate. If $S_a > \bar{S}$, this date occurs at time zero. Thirdly, this threshold date occurs further out the higher the migration costs. When condition b) does not hold, agricultural labor that can be pushed is already exhausted and services pull can dominate even in the long run.

In the next section, we present the results from this model and implications for premature deindustrialization.

5 Quantitative Results and Estimation

This section explores the quantitative results of the model in the 43-country panel that spans 1970-2013. We first show the fit of the baseline model, paying particular attention to the rural-urban labor market wedges and the rural-urban wage gap. We then analyze the frictionless benchmark comparison by taking $\tau_t^{n,i}, \nu \rightarrow 0$. We then finally compare the peak manufacturing shares of our model and the benchmark compared to the data.

5.1 Baseline Fit: Urbanization Rates and the Rural-Urban Wage Gap

Figures A.2 and A.3 show the paths of the rural-urban population shares and wage gaps in the model vs. the data for each country between 1970-2013. As we detail in depth in the appendix, the model smooths various parts of the data series, including IO weights and productivity. The result is that the wage gaps generated by the model are smoother than the data. The data is quite volatile because nominal value added per worker moves with sectoral

prices. The rural-urban shares in the data are smooth, so the data is not meaningfully more volatile than the model in those plots.

Figures A.2 and A.3 are direct tests of the performance of our rural-urban labor market wedge inversion. The idea of the inversion is that in the data for each country and region n, r we observe the path of supernumerary real incomes $\{(\tilde{w}_t^{n,r})/\tilde{P}_t^n\}_{t=0}^T$, where $\tilde{w}_t^{n,r} = w_t^{n,r} - K_t^n$ with $K_t^n \equiv P_t^{n,a}\bar{c}_a + P_t^{n,s}\bar{c}_s$. We assume agents have perfect foresight over this wage path, and thus know the value of each region for any given moving costs $\tau_t^{n,ru}$

$$V_t^{n,r} = \ln \frac{\tilde{w}_t^r}{\tilde{P}_t^n} + \mathbb{E}_\epsilon \left[\max_u \left(\beta V_{t+1}^{n,u} - \tau_t^{n,ru} + \nu \epsilon_t^u \right) \right]$$

Since we know their values, we also know their moving decisions for any given costs

$$\mu_t^{n,ru} = \mathbb{P} \left(u = \arg \max \left\{ \beta V_{t+1}^{n,u} - \tau_t^{n,ru} + \nu \epsilon_t^u \right\} \right)$$

We thus choose the path of costs $\{\tau_t^{n,ru}\}_{t=0}^T$ that make the observed moving decision look like the data. Critically, this only works if the model generates the same relative wages as the data.

The model tracks urbanization closely. Across the 41 economies with rural regions (Hong Kong and Singapore have no defined rural region in our model), the rural population share path has a median root mean squared error of 1.6 percentage points. The interquartile range is 1.1 - 2.8 points. The median within-country correlation between model and data paths is 0.99. In 2013, the median urban share is 72 percent against 71 percent in the data. Botswana has poor fit.

The model also reproduces the paths of the rural-urban wage gaps. In the data, the median urban-to-rural wage ratio among developing countries falls from 4.1 in 1970 to 3.5 in 2013. Among advanced economies, it rises from 0.9 to 1.4. The median root mean squared error of the log wage ratio is 0.30. The root mean squared error is not a direct test though, since the model is by design smoothed whereas the raw data features heavy relative sector price swings.

The model broadly matches global patterns of structural change. Agricultural employment falls from 60 to 35 percent in the model against a 56 to 28 percent fall in the data. Services rise from 22 to 44 percent in the model, while in the data they move from 26 percent to 48 percent. Manufacturing employment rises from 18 to 23 percent, while the model starts at 18 percent and rises to 21 percent by 2013. The model matches the level of global manufacturing employment at the start of the sample and replicates the data pattern of no global deindustrialization.

5.2 The Frictionless Benchmark

We now compare our model to a frictionless benchmark trade economy by taking $\tau_t^{n,i}, \nu \rightarrow 0$. In this limit the option value $\nu \log \sum_i \exp((\beta V_{t+1}^i - \tau_t^{n,i})/\nu)$ collapses to $\beta \max_i V_{t+1}^i$ and wages are equalized in equilibrium between regions. In this frictionless economy the migration block drops out and we thus solve a standard trade model. That is, we solve for a sequence of temporary equilibria. We hold productivities, trade costs, input-output weights, preferences, and trade balances at their baseline values. Appendix E recalibrates the preference parameters to give the frictionless economy its best fit. The over-industrialization of late developers survives in that calibration too. The most relevant point to make here is that we do not re-invert productivity at a single national wage. Instead, the frictionless economy inherits the region- and sector-level productivities of Appendix A. This choice (as well as capital) differentiates our frictionless economy from Sposi, Yi, and Zhang (2026).

Our productivity estimates differ from SYZ in proportion to the wage gap between agriculture and the rest of the economy. Sposi, Yi, and Zhang (2026) (SYZ) recover sectoral productivity from prices, self-trade shares and one national wage. By contrast, we invert productivity from sector and region level wages. Let $A_n^j(t)$, $w_n^j(t)$, and $\pi_{nn}^j(t)$ denote country n 's sector j productivity, wage, and trade openness at time t respectively. Model inversion for productivity sets

$$A_n^j(t) \propto \frac{w_n^j(t)}{P_n^j(t)} \cdot \pi_{nn}^j(t)^{1/\theta}.$$

So productivity moves proportional to which wage is fed in.

To see how the difference works intuitively, consider a typical developing economy where agricultural workers earn a third of what manufacturing workers earn. SYZ price labor in each sector at the same national wage. A single wage inversion observes that agricultural output is cheap but also sees a high national wage so assumes agriculture must be productive. Equivalently, in terms of efficiency units, the single wage inversion sees that agriculture uses few efficiency units of labor and is thus productive. Under this inversion, poor countries have a comparative advantage in agriculture. By contrast, when we feed in sector-level wages, we input a much lower agricultural wage. In our inversion, agricultural output is cheap because labor is cheap, rather than because labor is productive. Thus, we find that developing economies have a comparative advantage in manufacturing.

Peters, Zhang, and Zilibotti (2026) likewise recover sectoral productivity from region-level earnings rather than a national wage, and find a large rural–urban productivity gap in tradables. Our relative productivity estimates are broadly consistent with Herrendorf, Rogerson, and Valentinyi (2026), Gollin, Lagakos, and Waugh (2014a), Caselli (2005), and Restuccia, Yang, and Zhu (2008), who all document that poor countries lag far more in

agriculture than any other sector of the economy.

The interpretation of the frictionless benchmark depends on the exact nature of the spatial wedge. If it is a migration barrier then removing the wedge simply reflects removing spatial frictions in the allocation of economic activity. If the wedge is due to a skill gap then the frictionless benchmark reflects a rise in rural workers' skills until they are as skilled as the average urban worker. The higher manufacturing peaks then represent what happens under a rural human capital shock.

One important point is that even if the agricultural productivity gap reflects pure selection then our benchmark does not cleanly collapse to SYZ. If rural workers are less skilled in general then a low rural wage reflects few efficiency units. In that case, our model coincides with an SYZ interpretation where comparative advantage is in agriculture. But if the skill gaps are urban-specific then it signals low farm productivity. In this case, there is a latent comparative advantage in manufacturing until rural workers acquire urban skills.

The evidence suggests that much of the gap is due to such a lack of urban-specific skills. In Hamory *et al.* (2021), and Alvarez (2020) the selection gap is primarily absorbed by individual fixed effects. These fixed effects include things such as language differences between rural and urban labor markets which represent a disadvantage that an agricultural worker carries into urban labor markets but not at home. For example, language differences measurably reduce migration between Indian districts (Kone *et al.* 2018). Similarly, Zambia has over eighty dialects and around thirty languages. Seven of these are official national languages. English is used mainly in urban areas, while rural areas rely on national languages (Marten and Kula 2008).

Figure 7 shows the industrial paths of the baseline spatial model against our frictionless benchmark model for the three biggest developing economies. Each of these economies peak above the data in the frictionless benchmark, and premature deindustrialization does not appear. The spatial model pulls down industrial shares in the three emerging economies, and does generate premature deindustrialization. In the appendix, we report alternative calibrations of the frictionless model in E.3 that try to get these economies to prematurely deindustrialize in the frictionless benchmark.




FIGURE 7: MAJOR EMERGING MARKET MANUFACTURING EMPLOYMENT PATHS

Notes: Manufacturing employment shares for the three biggest emerging markets from 1971–2013: Brazil, India and China. Data in black, baseline spatial model in blue dashed, and frictionless model in red dotted. Table 2 summarizes the fit.

5.3 Peak Manufacturing Employment Shares

Figure 8 is the central exhibit of this paper. For each country we record the manufacturing employment share over 1970–2013 from the data. We smooth the data series with a centered five-year moving average, and record the value and year of the maximum. For the model, we record the manufacturing employment share by aggregating over urban and rural manufacturing. Panel A plots the baseline model peaks on the y-axis against the smoothed data peaks on the x-axis. Panel B does the same for the frictionless benchmark. Table 2 records summary statistics.




FIGURE 8: FIT OF THE PEAK MANUFACTURING EMPLOYMENT SHARE, 1970–2013

Notes: Each point is a country. The left panel plots the peak of the dynamic spatial model with rural-urban labor market wedges against the peak in the data. The right panel plots the peak of the frictionless benchmark economy with $\tau, \nu \rightarrow 0$. Red markers are developing countries, blue markers are advanced economies. The 45-degree line denotes perfect fit. Each panel also reports the correlation, the OLS slope of the model peak on the data peak, the root mean squared error and the mean model minus data gap for developing countries. Table 2 summarizes fit moments.

Figure 8 shows that the frictionless economy clearly over-industrializes late developers. Peak manufacturing employment shares exceed the data by 14.1 percentage points on average across the 26 developing countries. The median peak overshoots by 10.3 points. This is exactly what the comparative advantage argument would suggest. Once rural labor frictionlessly can access urban labor markets, the open-economy sends a large number of workers into manufacturing. The core feature of the data that drives this reallocation is the high wages (and hence high inferred productivity) of manufacturing in the data. This high inferred productivity shows up as strong comparative advantage. Developing economies in the model specialize in manufacturing and import agriculture. The correlation between frictionless and observed peaks is 0.04 and the slope is 0.05. Recalibrating preferences to the frictionless economy alone moves the median overshoot to 8.9 percentage points and leaves the correlation near zero. See Appendix E, for details, and Table E.1.

The dynamic spatial model with costly migration wedges removes roughly nine tenths of over-industrialization bias. The spatial model over-industrializes developing economies by 1.9 percentage points on average. However, most of this is driven by outliers. The model

slightly under-industrializes the median developing economy by 0.8 percentage points. The correlation with the data rises to 0.59 and the slope to 0.66. South Korea, Botswana, and Costa Rica are the largest model outliers. These countries' baseline peaks exceed the data by 36, 31, and 19 percentage points respectively. Excluding these countries flips the mean over-prediction into an under-prediction of 1.6 points. Botswana's over-industrialization is driven by its mining sector, which the model inversion infers as high manufacturing productivity.

As might be expected, the migration wedge does not materially move advanced economies. Both the frictionless and frictional models under-predict the advanced economy mean peaks by about 8 percentage points. The shortfall is inherited from the calibration of our trade block.

TABLE 2: FIT OF PEAK MANUFACTURING EMPLOYMENT SHARES, 1970–2013

	Baseline	Frictionless
Correlation with data peaks	0.59	0.04
OLS slope of model peak on data peak	0.66	0.05
Root mean squared error (pp)	10.7	17.7
Model – data, developing, median (pp)	–0.8	+10.3
Mean peak, developing (data 0.24)	0.26	0.38
Mean peak, advanced (data 0.39)	0.31	0.31
Median peak year (data 1986)	1989	1973

Notes: 41 countries. Hong Kong and Singapore excluded. Developing (26) is every country except the 15 advanced economies AUS, AUT, DEU, DNK, ESP, FIN, FRA, GBR, IRL, ITA, JPN, NLD, PRT, SWE, and USA, so South Korea and Taiwan count as developing. Peaks are maxima of five-year moving averages of the manufacturing employment share over 1970–2013.

5.4 Peak Levels, Peak Timing, and Income

Figure 9 plots the peak manufacturing share against log GDP per worker in 1990 in the baseline model, in the frictionless economy and in the data. In the data the relationship is steeper and starker than our frictional model. A country that is one log point richer in 1990 reaches a peak manufacturing employment share 5.7 percentage points higher, and peaks 12 years earlier. This is the premature deindustrialization fact of Figure 2, Panel A, restated for our sample data.

The frictionless model gets the sign of this relationship wrong. The slope is -0.021 (with standard error 0.017) because poorer countries over-industrialize. The baseline removes this overshoot and the slope is $+0.026$ (standard error 0.015). Neither version produces a relationship that is as stark as the data. Both models under-industrialize rich economies by about 8 percentage points. The migration wedge successfully rationalizes why late developers peak far below their comparative advantage (where they peak in the frictionless benchmark).

Richer countries also peak earlier in the data than in the model. The median peak year moves from 1973 in the frictionless economy, to 1989 in the baseline, against 1986 in the data.




FIGURE 9: PEAK MANUFACTURING SHARE AND INCOME PER WORKER

Notes: Each panel plots a country's peak manufacturing employment share over 1970 to 2013 against its log GDP per worker in 1990 for the baseline model (left), the frictionless model, $\tau, \nu \rightarrow 0$ (middle), and the data (right). Peaks are maxima of five-year moving averages of the manufacturing employment share. Colors give the peak year. Circles mark interior peaks, and diamonds mark peaks at the first or last year of a country's window. Hong Kong and Singapore are excluded since they have no rural region in the model.

5.5 China Shock

In the previous sections we argue that a rural-urban labor market wedge is blocking the industrialization paths of today's emerging markets. One common argument is that premature deindustrialization could be caused by China. This section adjudicates the role of China in premature deindustrialization. We run a counterfactual that freezes Chinese fundamentals at their 1970 values. Our baseline follows Caliendo, Dvorkin, and Parro (2019) by explicitly modeling the Chinese manufacturing export boom as a target of the model. Because running counterfactuals relies on a baseline that matches the data, we also run a version of the model that incorporates sector switching costs into the model. Changing rural-urban switching costs for direct sector switching costs follows the approach of Caliendo, Dvorkin, and Parro (2019) and makes sure that the baseline follows the data as closely as possible. We also run the rural-urban model (with frictionless labor movement within regions) and report results for both.

In this version of the model each country now has three labor markets: agriculture,

manufacturing, and services. Each has a switching cost inverted similarly to the rural vs. urban discrete choice model. We run the counterfactual in this version of the model in addition to the frictionless within region version. The sector-switching model tracks sector employment closely. The correlation between model and data peak manufacturing shares in this version of the model is 0.96. This version also imposes a time varying drift in the manufacturing preference weight so that it drifts down at a common rate of 1.5 percent a year over 1970–2013.

Table 3 reports the resulting difference in manufacturing employment shares between the baseline and the counterfactual. Figure 10 plots the baseline and counterfactual employment paths against the data for the United States and nine countries which the baseline fits well.

In the sector-switching model, we find that the rise of China decreased manufacturing employment by 0.67 percentage points in advanced economies and 0.29 in developing countries. The rise of China decreased U.S. manufacturing employment by 0.23 in our model which is below the 0.36 found by Caliendo, Dvorkin, and Parro (2019). Naturally, we find larger effects in 2000 and 2013 in the region switching model, where the less frictional reallocation of labor leads to a 0.52 percentage point decrease in manufacturing employment in the United States. Our smaller estimate in the sector-switching model is consistent with our broad manufacturing industry classification which includes construction, mining, and utilities. These are either non-tradable sectors or sectors that likely boomed via input demand to Chinese manufacturing. Our model captures these forces inside of the broad manufacturing definition, which likely mutes the result relative to the definition in CDP.

This aggregate effect masks substantial cross-country heterogeneity. The rise of China had the biggest deindustrialization effects by 2013 in Australia (−2.0), Ireland (−1.9), Ethiopia (−1.8) and France (−1.0). By contrast, Hong Kong (+4.0) and Malaysia (+0.8) found their manufacturing shares rise with China. Effects on peak manufacturing shares are smaller than the 2013 effects. Advanced economies’ peaks are unchanged, since they occur before China’s rise. Developing economies are 0.22 points lower in the sector-switching model. In the region-switching model, advanced and developing peaks are 0.38 and 0.09 points lower respectively.

Two points are noteworthy in the interpretation of our estimates. First, we aggregate manufacturing across different types of goods. So Chinese manufacturing competition falls equally over all types of manufactures in this model’s counterfactual. In reality, Chinese manufacturing competition was concentrated in lower value added, labor-intensive goods. This overlap is stronger with other emerging economies than for advanced nations. Thus, in practice emerging economies were probably *more* exposed while advanced economies were probably *less* exposed than our counterfactual suggests. Our counterfactual averages over

these exposures.

Second, consistent with Huneus and Rogerson (2024), we use the broad sector manufacturing classification. This includes mining, utilities, and construction into the category of manufacturing. The global commodity supercycle of the 2000s was largely driven by China, and so shows up in our estimates. In this sense, commodity intensive exporters like Australia, Brazil, and Chile would have deindustrialized more than the counterfactual suggests if we were to define manufacturing more narrowly.

Figure 11 shows manufacturing real wages for the select countries. By 2013 the rise of China lowers manufacturing real wages in seven of the ten pictured economies, most strongly in France, where it falls by 6.9 percent. Since switching between sectors is costly, Chinese competition lowers manufacturing real wages despite increasing real wages in the economy overall. The switching costs mean that real wages move more dramatically than employment shares. In the global economy as a whole, the rise of China raises the average real wage in 2013 by 3.1 percent in advanced and 4.3 percent in developing economies, where real wages are deflated by the model’s own price index with the drifting preference weights.

TABLE 3: MANUFACTURING EMPLOYMENT SHARES: BASELINE AND CHINA FROZEN AT 1970 COUNTERFACTUAL

	Effects of the Rise of China	
	Sector switching	Region switching
Advanced, in 2000	−0.38	−0.52
Developing, in 2000	−0.13	−0.41
Advanced, in 2013	−0.67	−1.23
Developing, in 2013	−0.29	−0.77
Advanced, peak share	0.00	−0.38
Developing, peak share	−0.22	−0.09
United States, in 2013	−0.23	−0.52

Notes: The table reports the manufacturing employment share in the baseline minus the counterfactual (pp). This counterfactual holds Chinese productivities and trade costs at 1970 levels. The labor market wedges, Chinese value-added and input-output shares and China’s deficit-to-income ratio follow baseline paths. The column labeled “Sector switching” reports results for the model with sector switching costs described in this section. The column labeled “Region switching” reports results for the rural-urban migration model used in the rest of the paper. Manufacturing employment share sums manufacturing employment over a country’s rural and urban regions. The rows for advanced and developing economies report simple means over the 15 advanced economies and the 25 developing economies excluding China, Hong Kong and Singapore.




FIGURE 10: MANUFACTURING EMPLOYMENT SHARES: BASELINE AND COUNTERFACTUAL

Notes: The figure reports the manufacturing employment share, in percent, in the baseline, counterfactual and data. The counterfactual is defined as in Table 3. Results are for the model with sector switching costs described. The top row reports advanced economies and the bottom row developing economies.




FIGURE 11: MANUFACTURING REAL WAGES IN BASELINE AND COUNTERFACTUAL

Notes: The figure reports the manufacturing real wage in the baseline, counterfactual and data. The counterfactual is defined as in Table 3. Results are for the model with sector switching costs described in this section. The top row reports advanced economies and the bottom row developing economies. The countries are those of Figure 10. Model real wages are the manufacturing wage deflated by the CES price index of the manufacturing region.

The finding that China is not a primary driver of premature deindustrialization should not be surprising given the stylized facts we observe in the data. If China were the cause of broad global deindustrialization then we should expect to see low and falling wages in manufacturing. Yet, in the data manufacturing wages remain stubbornly higher than overall

wages in emerging markets. Secondly, premature deindustrialization was already a stark feature of the data before China’s accession to the WTO in 2001, which is widely considered to be the start of the China shock. If China were responsible for premature deindustrialization then the timing of premature deindustrialization would sit awkwardly against Chinese global import penetration numbers before 2000. Finally, CDP already study this question in the context of the United States. Although CDP document substantial negative impacts to U.S. manufacturing from the China shock, their numbers are an order of magnitude lower than the numbers that show up in the premature deindustrialization literature.

5.6 Discussion

We draw three conclusions from the results presented in this section. First, the baseline dynamic spatial model is internally consistent with the data used for inversion.

Second, Section 2.2 documented suggestive evidence as to why premature deindustrialization should take account of global trade. The figures in this section document this comprehensively. Our frictionless benchmark of structural change that incorporates trade with frictionless labor markets (and productivities measured per worker) tends to over-industrialize late developers by about 10 percentage points. Trade-induced over-industrialization eliminates, and slightly inverts, premature deindustrialization. The open-economy is thus quantitatively active, and swamps the closed economy Baumol channel.

Third, estimated migration wedges close almost nine tenths of the trade-induced over-industrialization and bring emerging economies much closer to the data. Our quantitative model still under-industrializes advanced economies. The next section tests the mechanism behind these results outside the model, using a labor-saving agricultural shock in Brazil.

6 Evidence from Brazil

An implication of our model is that whether a labor-saving agricultural productivity shock moves employment from agriculture into manufacturing depends on rural-urban mobility costs. We examine this implication using the setting, data construction, and specification of Bustos, Caprettini, and Ponticelli (2016) (henceforth BCP), which studies the 2003 legalization of genetically modified soy in Brazil, a labor-saving shock that raised manufacturing employment shares in more exposed municipalities. We augment their specification by allowing the effect of the shock to vary with a pre-period proxy for municipalities’ mobility costs.

In 2003, Brazil legalized the planting of genetically modified soy, which decreased the

necessary labor per hectare by allowing farmers to plant without tillage and with fewer herbicide applications while maintaining soy yield. The following decade simultaneously saw the continued spread of a second harvesting season for maize due to modern cultivation techniques, which increased output per hectare by allowing farmers to grow maize as a second crop within the same agricultural year. We exploit the fact that the soy technology was labor-saving rather than a pure productivity shock. In municipalities that sell crops in open markets, a productivity shock that does not save labor would draw workers into agriculture, as BCP finds for the land-augmenting second maize season. The labor-saving soy shock, however, releases labor out of agriculture, which allows us to test our model’s prediction about how mobility costs determine whether they find jobs in manufacturing.

Our exposure measure is the difference in FAO-GAEZ potential yields under high and low-input technologies separately for soy and maize, which measures how much a region stands to gain from adopting the new technology for that crop. We use the population census microdata waves 2000 and 2010 to measure changes in sectoral employment shares, and the 1991 wave for our baseline control variables. Table F.1 shows the summary statistics of all main variables used.

We proxy for mobility costs with migration propensity, the observable implication of low mobility costs under either of their interpretations, migration frictions or skill barriers. Specifically, workers leave agriculture more easily where migration frictions are low and find urban jobs more easily where skill barriers are low. Our baseline measure is the youth share, i.e., the share of a municipality’s 1991 population aged 18 to 25, measured more than a decade before the shock. Migration is strongly concentrated among younger adults in every country and period for which age-specific rates exist. This is the well-documented age gradient in migration propensity (Rogers and Castro 1981; Bernard, Bell, and Charles-Edwards 2014). The mean age of migration between Brazilian municipalities lies in young adulthood in the 1991, 2000 and 2010 censuses (dos Santos, Barbieri, and Amaral 2023), and those leaving rural Brazil have become increasingly young (Camarano and Abramovay 1998). The same gradient appears in international migration: countries with a larger share of their population aged 15 to 29 have higher emigration rates (e.g., Mayda 2010). Young cohorts also drive structural change by entering non-agricultural work, carrying the human capital valued outside it (Porzio, Rossi, and Santangelo 2022). A young population therefore signals low mobility costs under both interpretations. In Appendix F, we show that our results are similar under alternative age-based proxies and under realized urban-directed out-migration from 1995 to 2000, which measures migration propensity directly. Past out-migration also matters for later moves, as migrants from the same origin facilitate access to destination labor markets (Munshi 2003).

We keep the youth share as our baseline because it is largely determined by fertility decisions in the 1960s and 1970s, three decades before the shock, and thus cannot respond to local economic conditions in the years just before the shock, unlike realized out-migration. The two proxies reflect migration frictions and skill barriers differently and thus need not rank municipalities in the same way. While the youth share proxy measures potential movers, i.e., residents at the most mobile ages, the out-migration proxy measures actual moves between 1995 and 2000.

Our research design augments the continuous difference-in-differences specification of BCP by interacting the technological exposure measures with quintiles of migration propensity:

$$\begin{aligned} \Delta y_j = & \Delta\delta + \beta^{\text{soy}} \Delta A_j^{\text{soy}} + \beta^{\text{maize}} \Delta A_j^{\text{maize}} + \sum_{k=2}^5 \beta_{Q_k}^{\text{soy}} \Delta A_j^{\text{soy}} \cdot \mathbf{1}\{Q(j) = k\} \\ & + \sum_{k=2}^5 \beta_{Q_k}^{\text{maize}} \Delta A_j^{\text{maize}} \cdot \mathbf{1}\{Q(j) = k\} + \sum_{k=2}^5 \theta_{Q_k} \mathbf{1}\{Q(j) = k\} + \mathbf{X}'_{j,1991} \boldsymbol{\omega} + \Delta e_j \end{aligned} \quad (20)$$

where Δy_j is the 2000-2010 change in an AMC j 's sectoral employment share, ΔA_j^{soy} and $\Delta A_j^{\text{maize}}$ are the crop exposure measures, and $\mathbf{X}_{j,1991}$ collects the baseline controls, specifically, the rural share, log population density, literacy rate, and log income per capita from 1991, as in BCP. The interactions with $Q(j)$, the quintiles of an AMC j 's migration propensity, identify how the effect on sectoral employment shares is moderated by mobility costs. The total effect of the soy shock in a given quintile is given by the sum of the coefficients of interest, $\beta^{\text{soy}} + \beta_{Q_k}^{\text{soy}}$, which our model predicts to increase with mobility. We include the maize exposure term and its interactions with migration propensity in each specification as controls but do not use this shock for our analysis as its average effect size is only a fifth of the size of the soy effect, too small to detect heterogeneity with precision. We cluster standard errors at the microregion level, a grouping of municipalities above the AMC level, except in columns that replicate BCP, where we use robust standard errors.

Table 4 presents our main results. The first column replicates BCP, showing that a one tonne per hectare increase in soy exposure increases the manufacturing employment share by 2.1 percentage points from 2000 to 2010. The second column shows our baseline specification.⁸ Our key finding is that the effect on manufacturing employment shares exhibits

⁸Our estimation sample includes slightly more AMCs than BCP (our 4,260 against their 4,149) because their estimation sample requires filtering on non-missing values on variables that we do not study. Throughout our analysis, we explicitly test if our results are robust to limiting our analysis to the AMCs in the original BCP sample. We label these columns with a ‘‘BCP sample’’ row indicator, such as column 4 of this table.

a strong gradient across migration propensity. In particular, municipalities with the lowest migration propensity experience an effect indistinguishable from zero; this effect then rises monotonically with migration propensity reaching its maximum total size of 2.8 percentage points in the largest quintile. Figure 12, Panel A, depicts this gradient in the total effects on manufacturing employment shares, while Panel B shows the mirror image of this gradient for the agricultural sector. Column 3 shows that the gradient is robust to a rich set of development controls, specifically, state fixed effects and the 1991 measures of urbanization, log population density, literacy rate, and log income per capita, each demeaned and interacted with our two exposure measures, and column 4 replicates our results in the sub-sample of AMCs used in BCP.

TABLE 4: MIGRATION PROPENSITY GRADIENT IN THE SOY EFFECT ON
MANUFACTURING

	Δ Manufacturing employment share			
	(1)	(2)	(3)	(4)
ΔA^{soy}	0.021*** (0.002)	0.003 (0.006)	0.003 (0.006)	0.003 (0.006)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 2}$		0.014** (0.007)	0.011* (0.007)	0.014** (0.007)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 3}$		0.014* (0.008)	0.012* (0.007)	0.014* (0.008)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 4}$		0.017** (0.007)	0.016** (0.006)	0.016** (0.007)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 5}$		0.025*** (0.008)	0.026*** (0.007)	0.027*** (0.008)
Joint test, soy coefficients on quintiles 2-5 equal 0:				
<i>p</i> -value		0.028	0.011	0.016
Observations	4,149	4,260	4,260	4,149
Fixed Effects and Controls:				
Baseline controls	✓	✓	✓	✓
Development controls	-	-	✓	-
BCP sample	✓	-	-	✓

Notes: The table reports estimates from equation (20). Column (1) restricts all migration propensity terms to zero, which replicates the BCP specification (Table 9, column 2). Columns (2)–(4) estimate the unrestricted equation. The outcome variable for each column is the 2000-2010 change in the manufacturing employment share. Migration propensity is measured as the 1991 youth share (ages 18 to 25), and its lowest quintile is the omitted category. Baseline controls are the 1991 rural share, log population density, literacy rate, and log income per capita. Column (3) adds development controls, specifically state fixed effects and the 1991 urbanization rate, log population density, literacy rate, and log income per capita, demeaned and interacted with our soy and maize exposure measures. Column (4) restricts the sample to the AMCs used in BCP. The row labeled “p-value” reports the p-value from an F-test of the null hypothesis that the soy exposure $\times$ quintile coefficients are jointly zero. Standard errors are in parentheses, clustered at the microregion level in columns (2)–(4) and robust in column (1), as in BCP. ***, **, and * denote significance at the 1, 5, and 10 percent levels.




FIGURE 12: MIGRATION PROPENSITY GRADIENT IN THE TOTAL SOY EFFECTS ON MANUFACTURING AND AGRICULTURE

Notes: The figure explores the relationship between the soy shock's effect on sectoral employment shares and migration propensity. We divide AMCs into five groups depending on their quintile of migration propensity, measured as the 1991 youth share (ages 18 to 25). We estimate the unrestricted equation (20) and plot the total effects $\hat{\beta}^{\text{soy}} + \hat{\beta}_{Q_k}^{\text{soy}}$ for each quintile along with the 95 percent confidence interval. Each panel includes the baseline controls, specifically the 1991 rural share, log population density, literacy rate, and log income per capita. Panel A corresponds to column (2) of Table 4 and Panel B to Column (2) of Table F.2. Confidence intervals are based on standard errors clustered at the microregion level.

Table F.2 estimates our baseline specification across the four sectors, showing that the positive mobility gradient of the effect on manufacturing employment shares is mirrored by a negative mobility gradient of the effect on agricultural employment shares. We interpret this evidence as consistent with a reallocation across sectors, from agriculture into manufacturing, for municipalities with high enough migration propensity. Indeed, in each of the four upper quintiles, the fall in the agricultural share is accompanied by a comparable rise in the manufacturing share, while the services share barely moves.

Our results are remarkably robust across different migration propensity definitions and functional forms. Table F.3 shows that we obtain similar results when measuring migration propensity with the young-adult share (ages 18-35) or the prime-age share (ages 15-55) as with our baseline measure, the youth share (ages 18-25). We also obtain similar results using pre-shock realized out-migration into urban destinations, from 1995 to 2000. Table F.4 presents similar results after replacing our migration propensity quintiles with a continuous version, normalized to the unit interval in our sample. Our baseline specification, in column (2), shows that the effect of a one tonne per hectare increase in soy exposure rises by as much as 4.9 percentage points at the highest migration propensity region relative to the lowest, where it is indistinguishable from zero.

Identification in this design rests on two main assumptions. First, as in BCP, exposure must be exogenous conditional on the baseline controls. That is, regions with differential

potential yield gains would have followed similar trends in sectoral employment absent the new technologies. Second, migration propensity must be quasi-exogenous, meaning that it was fixed before the shock and that it does not stand in for other municipality characteristics that change the shock’s effect. It was fixed before the shock, since we measure it in 1991, over a decade before the shock and before genetically modified soy was commercially available anywhere. It also does not appear to stand in for development or region, since our results are robust to development controls, which include state fixed effects and 1991 measures of development interacted with exposure, and to alternative proxy measures. The out-migration proxy provides a further check because it moves with development in the opposite direction from the youth share. While the youth share is higher in richer municipalities, out-migration is lower there. Were the gradient driven by development, it would imply opposite patterns for the two proxies. Instead, both proxies produce the same pattern, with no detectable effect in the least mobile quintile and larger effects in all four upper quintiles.

We additionally test both assumptions with a pre-trend test, asking whether municipalities with different soy exposure were already on different sectoral employment trends and whether these differences varied with migration propensity. This would be the case, for example, if municipalities with high migration propensity and high soy exposure were already industrializing faster, independent of the shock.

Table F.5 reports the pre-trend tests. Following BCP, we regress the 1991-2000 and 2000-2010 changes in log manufacturing employment on exposure and exposure interacted with the post-shock period (2000-2010) in a stacked specification with lagged baseline controls. Column (1) replicates BCP’s pre-trend test, testing whether soy shock exposure predicted changes in manufacturing employment in the years preceding the shock, between years 1991 and 2000.⁹ As BCP, we find a large effect in the post-period and a null effect in the pre-period, consistent with the quasi-exogeneity of exposure. Columns (2) to (4) add our migration propensity quintiles and their interactions with exposure and with the post-period indicator. Consistent with Table 4, we find evidence of substantial effects in the post-period, moderated by migration propensity, with no detectable effect in the least mobile quintile of municipalities. In contrast, across all of the specifications, we can never reject the joint null hypothesis of no effects across the five quintiles in the pre-period. While our individual total effects show a clear, robust pattern across migration propensity in the post-period, they are noisy, around zero, and overwhelmingly not significant in the pre-period.

⁹While our and BCP’s exposure measures, built using the same FAO-GAEZ v3.0 grids, are correlated at 0.997, they are not numerically the same. This explains why our replicated coefficients are essentially, but not numerically, identical. Our pre-period coefficient in column (1) is < 0.001 (s.e. 0.020) while theirs is 0.007 (s.e. 0.020), and our post-period differential, not shown in the table, is 0.238 (s.e. 0.025) while theirs is 0.234 (s.e. 0.025).

While our evidence is consistent with mobility costs acting as a moderator, two limitations remain. First, we measure migration propensity using proxies, which rely on the age-gradient of migration or the realized pre-period out-migration directed to urban places. These proxies may capture municipality characteristics other than mobility costs, such as its level of development or regional characteristics. While the robustness of our results speaks against these main alternative mechanisms, we cannot rule out every alternative reading of the gradient. Second, we observe sectoral shares rather than movements of workers. While we show that in mobile municipalities, employment shifted from agriculture to manufacturing, we cannot link workers who left agriculture to those who entered manufacturing at the worker level as the censuses do not track workers over time.

Overall, our quasi-experimental evidence is consistent with the implication of our model in Section 3 that a labor-saving agricultural productivity shock moves employment from agriculture into manufacturing only where mobility costs are low. In Brazil, comparing municipalities with the same exposure to the shock but different migration propensities, those with high migration propensity industrialized in response to the shock while those with low migration propensity did not.

7 Conclusion

Today’s late developers reach manufacturing employment peaks about half as high as the early industrializers did, and they reach them at lower incomes. This paper argues that through the lens of a spatial model inversion, international trade should push these economies to industrialize more, not less. We rationalize the fact that emerging markets are not industrializing through a rural-urban labor market wedge, which can be interpreted as either an internal migration friction, or a worker skill barrier. This wedge stops manufacturing from pulling labor out of agriculture, and keeps workers in agriculture.

We recover this wedge from the joint path of rural-urban wage gaps and population stocks in 43 countries. We find it to be large in late developers, where comparative advantage would otherwise pull labor into manufacturing, and falling over time in most countries. Because today’s late developers grow much faster than the early industrializers did, this wedge binds within a far shorter development window. We show that without this wedge a trade model over-industrializes the median developing economy by over 10 percentage points and eliminates premature deindustrialization. Because rural and urban labor markets clear at different wages, the model also delivers the fact that premature deindustrialization appears in employment shares but is weaker in value added shares.

We augment this result with an accounting decomposition of the standard trade model

budget constraint. The accounting decomposition suggests the same story: internal migration is the largest contributor to manufacturing employment growth in four of the six emerging markets in WIOD. By contrast, migration appears to be a small channel in advanced economies.

We use microdata and a quasi-experiment from the literature to examine the observable implication of our model that agricultural productivity shocks should have a stronger industrializing effect on places where spatial mobility costs are lower. Following BCP, we use the 2003 legalization of genetically modified soy in Brazil as a labor-saving agricultural productivity shock. We show that the productivity shock’s positive effect on local manufacturing employment shares is larger in municipalities with higher pre-period migration propensity and absent in those with the lowest pre-period migration propensity, consistent with our model. Together, our accounting decomposition, structural counterfactual, and quasi-experiment all point to the same bottleneck, the barriers slowing the flow of labor out of agriculture. We leave for future work whether differences in the rural-urban wedge, in addition to the factor bias of technical change emphasized by Bustos, Caprettini, and Ponticelli (2016), can explain the mixed empirical evidence on the industrializing effects of agricultural productivity shocks across settings.

The paper suggests multiple avenues for future work. First, the model under-industrializes advanced economies by about eight percentage points. The median advanced economy manufacturing peak also comes about nine years too late. The level shortfall is independent of the migration wedge and could require more sophisticated non-homothetic preferences, such as those in Comin, Lashkari, and Mestieri (2021).

Second, wage and migration flows cannot distinguish moving costs from skill discounts for the marginal rural worker. Our results suggest that late developers do not industrialize because labor does not flow from rural to urban areas, but we do not take a stance on whether this is due to skill or barrier. Whether the agricultural productivity gap reflects selection or barriers is a long-standing debate. Incorporating a dynamic Roy model of worker selection to our framework could separate the two forces.

Our results suggest that rural-urban migration is a critical channel through which opening to trade works to deliver welfare gains. When this channel is blocked, comparative advantage cannot be efficiently exploited and industrial dynamics end up coincidentally resembling those of closed economies in Huneus and Rogerson (2024) and Fujiwara and Matsuyama (2024).

References

- ALBERT, CHRISTOPH, BUSTOS, PAULA, MESTIERI, MARTÍ, & PONTICELLI, JACOPO. 2026. *The Effects of Climate Change on Labor Reallocation*. Working paper, April 2026; earlier version NBER Working Paper 28995 (2021).
- ALVAREZ, JORGE A. 2020. The Agricultural Wage Gap: Evidence from Brazilian Micro-data. *American Economic Journal: Macroeconomics*, **12**(1), 153–173.
- ARKOLAKIS, COSTAS, COSTINOT, ARNAUD, & RODRÍGUEZ-CLARE, ANDRÉS. 2012. New Trade Models, Same Old Gains? *American Economic Review*, **102**(1), 94–130.
- ASHER, SAM, CAMPION, ALISON, GOLLIN, DOUGLAS, & NOVOSAD, PAUL. 2024. *The Long-Run Development Impacts of Agricultural Productivity Gains: Evidence from Irrigation Canals in India*. Working paper, May 2024 version.
- BASELER, TRAVIS. 2023. Hidden Income and the Perceived Returns to Migration. *American Economic Journal: Applied Economics*, **15**(4), 321–352.
- BERNARD, AUDE, BELL, MARTIN, & CHARLES-EDWARDS, ELIN. 2014. Life-Course Transitions and the Age Profile of Internal Migration. *Population and Development Review*, **40**(2), 213–239.
- BRYAN, GHARAD, & MORTEN, MELANIE. 2019. The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy*, **127**(5), 2229–2268.
- BRYAN, GHARAD, CHOWDHURY, SHYAMAL, & MOBARAK, AHMED MUSHFIQ. 2014. Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh. *Econometrica*, **82**(5), 1671–1748.
- BUSTOS, PAULA, CAPRETTINI, BRUNO, & PONTICELLI, JACOPO. 2016. Agricultural Productivity and Structural Transformation: Evidence from Brazil. *American Economic Review*, **106**(6), 1320–1365.
- BUSTOS, PAULA, GARBER, GABRIEL, & PONTICELLI, JACOPO. 2020. Capital Accumulation and Structural Transformation. *Quarterly Journal of Economics*, **135**(2), 1037–1094.
- BUSTOS, PAULA, CASTRO-VINCENZI, JUANMA, MONRAS, JOAN, & PONTICELLI, JACOPO. 2023. *Industrialization without Innovation*. Working Paper 1444. Barcelona School of Economics. September 2023 version, released April 2024; earlier version NBER Working Paper 25871.
- CALIENDO, LORENZO, DVORKIN, MAXIMILIANO, & PARRO, FERNANDO. 2019. Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, **87**(3), 741–835.
- CAMARANO, ANA AMÉLIA, & ABRAMOVAY, RICARDO. 1998. Êxodo rural, envelhecimento e masculinização no Brasil: panorama dos últimos cinquenta anos. *Revista Brasileira de Estudos de População*, **15**(2), 45–65.
- CASELLI, FRANCESCO. 2005. Accounting for Cross-Country Income Differences. *Chap. 9, pages 679–741 of: AGHION, PHILIPPE, & DURLAUF, STEVEN N. (eds), Handbook of Economic Growth*, vol. 1A. Amsterdam: Elsevier.
- COMIN, DIEGO, LASHKARI, DANIAL, & MESTIERI, MARTÍ. 2021. Structural Change with Long-

- Run Income and Price Effects. *Econometrica*, **89**(1), 311–374.
- DE VRIES, GAAITZEN, DE VRIES, KLAAS, GOUMA, REITZE, PAHL, STEFAN, & TIMMER, MARCEL. 2015. *GGDC 10-Sector Database: Contents, Sources and Methods*. Database documentation. Groningen Growth and Development Centre, University of Groningen. Version January 2015.
- DOS SANTOS, REINALDO ONOFRE, BARBIERI, ALISSON FLÁVIO, & AMARAL, ERNESTO FRIEDRICH DE LIMA. 2023. Transiciones del curso de vida y migración interna en el Brasil: un análisis basado en datos de múltiples períodos. *Notas de Población*, 105–135.
- FAN, TIANYU, PETERS, MICHAEL, & ZILIBOTTI, FABRIZIO. 2023. Growing Like India—The Unequal Effects of Service-Led Growth. *Econometrica*, **91**(4), 1457–1494.
- FEENSTRA, ROBERT C., LIPSEY, ROBERT E., DENG, HAIYAN, MA, ALYSON C., & MO, HENGYONG. 2005 (Jan.). *World Trade Flows: 1962–2000*. NBER Working Paper 11040. National Bureau of Economic Research.
- FEENSTRA, ROBERT C., INKLAAR, ROBERT, & TIMMER, MARCEL P. 2015. The Next Generation of the Penn World Table. *American Economic Review*, **105**(10), 3150–3182.
- FOSTER, ANDREW D., & ROSENZWEIG, MARK R. 2004. Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000. *Economic Development and Cultural Change*, **52**(3), 509–542.
- FUJIWARA, IPPEI, & MATSUYAMA, KIMINORI. 2024. A Technology-Gap Model of ‘Premature’ Deindustrialization. *American Economic Review*, **114**(11), 3714–3745.
- GAULIER, GUILLAUME, & ZIGNAGO, SOLEDAD. 2010 (Oct.). *BACI: International Trade Database at the Product-Level. The 1994–2007 Version*. CEPII Working Paper 2010-23. CEPII.
- GLAESER, EDWARD L. 2014. A World of Cities: The Causes and Consequences of Urbanization in Poorer Countries. *Journal of the European Economic Association*, **12**(5), 1154–1199.
- GOLLIN, DOUGLAS, LAGAKOS, DAVID, & WAUGH, MICHAEL E. 2014a. Agricultural Productivity Differences across Countries. *American Economic Review*, **104**(5), 165–170.
- GOLLIN, DOUGLAS, LAGAKOS, DAVID, & WAUGH, MICHAEL E. 2014b. The Agricultural Productivity Gap. *The Quarterly Journal of Economics*, **129**(2), 939–993.
- HAMORY, JOAN, KLEEMANS, MARIEKE, LI, NICHOLAS Y., & MIGUEL, EDWARD. 2021. Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata. *Journal of the European Economic Association*, **19**(3), 1522–1555.
- HERRENDORF, BERTHOLD, & SCHOELLMAN, TODD. 2018. Wages, Human Capital, and Barriers to Structural Transformation. *American Economic Journal: Macroeconomics*, **10**(2), 1–23.
- HERRENDORF, BERTHOLD, HERRINGTON, CHRISTOPHER, & VALENTINYI, ÁKOS. 2015. Sectoral Technology and Structural Transformation. *American Economic Journal: Macroeconomics*, **7**(4), 104–133.
- HERRENDORF, BERTHOLD, ROGERSON, RICHARD, & VALENTINYI, ÁKOS. 2026. Unconditional Convergence in Manufacturing: A Reassessment. *American Economic Review: Insights*, **8**(3), 303–319.

- HORNBECK, RICHARD, & KESKIN, PINAR. 2015. Does Agriculture Generate Local Economic Spillovers? Short-Run and Long-Run Evidence from the Ogallala Aquifer. *American Economic Journal: Economic Policy*, **7**(2), 192–213.
- HUNEEUS, FEDERICO, & ROGERSON, RICHARD. 2024. Heterogeneous Paths of Industrialization. *The Review of Economic Studies*, **91**(3), 1746–1774.
- IIASA/FAO. 2012. *Global Agro-ecological Zones (GAEZ v3.0)*. IIASA, Laxenburg, Austria and FAO, Rome, Italy.
- IMBERT, CLÉMENT, & PAPP, JOHN. 2020. Costs and Benefits of Rural-Urban Migration: Evidence from India. *Journal of Development Economics*, **146**, 102473.
- INKLAAR, ROBERT, & TIMMER, MARCEL P. 2014. The Relative Price of Services. *Review of Income and Wealth*, **60**(4), 727–746.
- KONE, ZOVANGA L., LIU, MAGGIE Y., MATTOO, AADITYA, ÖZDEN, ÇAĞLAR, & SHARMA, SIDDHARTH. 2018. Internal Borders and Migration in India. *Journal of Economic Geography*, **18**(4), 729–759.
- KUZNETS, SIMON. 1973. Modern Economic Growth: Findings and Reflections. *American Economic Review*, **63**(3), 247–258.
- MARTEN, LUTZ, & KULA, NANCY C. 2008. Zambia: ‘One Zambia, One Nation, Many Languages’. *Pages 291–313 of*: SIMPSON, ANDREW (ed), *Language and National Identity in Africa*. Oxford: Oxford University Press.
- MATSUYAMA, KIMINORI. 1992. Agricultural Productivity, Comparative Advantage, and Economic Growth. *Journal of Economic Theory*, **58**(2), 317–334.
- MAYDA, ANNA MARIA. 2010. International Migration: A Panel Data Analysis of the Determinants of Bilateral Flows. *Journal of Population Economics*, **23**(4), 1249–1274.
- MOSCONA, JACOB. 2019. *Agricultural Development and Structural Change, Within and Across Countries*. Working paper, Massachusetts Institute of Technology, May 2019.
- MÜLLER, KARSTEN, XU, CHENZI, LEHBIB, MOHAMED, & CHEN, ZILIAN. 2025. *The Global Macro Database: A New International Macroeconomic Dataset*. NBER Working Paper 33714. National Bureau of Economic Research.
- MUNSHI, KAIVAN. 2003. Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market. *Quarterly Journal of Economics*, **118**(2), 549–599.
- O’MAHONY, MARY, & TIMMER, MARCEL P. 2009. Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database. *Economic Journal*, **119**(538), F374–F403.
- PETERS, MICHAEL, ZHANG, YODAN, & ZILIBOTTI, FABRIZIO. 2026. Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World. *Chap. 8 of*: DUPAS, PASCALINE, GOLDBERG, PINELOPI KOUJIANOU, & PANDE, ROHINI (eds), *Handbook of Development Economics*, vol. 6A. Amsterdam: Elsevier. Forthcoming. Earlier version NBER Working Paper 34692, January 2026.
- PORZIO, TOMMASO, ROSSI, FEDERICO, & SANTANGELO, GABRIELLA. 2022. The Human Side of

- Structural Transformation. *American Economic Review*, **112**(8), 2774–2814.
- RESTUCCIA, DIEGO, YANG, DENNIS TAO, & ZHU, XIAODONG. 2008. Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis. *Journal of Monetary Economics*, **55**(2), 234–250.
- RODRIK, DANI. 2016. Premature Deindustrialization. *Journal of Economic Growth*, **21**(1), 1–33.
- ROGERS, ANDREI, & CASTRO, LUIS J. 1981 (Nov.). *Model Migration Schedules*. Research Report RR-81-030. International Institute for Applied Systems Analysis, Laxenburg, Austria.
- SIMONOVSKA, INA, & WAUGH, MICHAEL E. 2014. The Elasticity of Trade: Estimates and Evidence. *Journal of International Economics*, **92**(1), 34–50.
- SMITKOVA, LIDIA. 2024 (Oct.). *Dissecting Structural Change in an Open Economy*. Working paper, October 2024.
- SPOSI, MICHAEL, YI, KEI-MU, & ZHANG, JING. 2026. The Role of Trade in Premature Deindustrialization and Industry Polarization. *Econometrica*, **94**(5), 1779–1815.
- TIMMER, MARCEL P., DIETZENBACHER, ERIK, LOS, BART, STEHRER, ROBERT, & DE VRIES, GAAITZEN J. 2015. An Illustrated User Guide to the World Input–Output Database: The Case of Global Automotive Production. *Review of International Economics*, **23**(3), 575–605.
- TIMMER, MARCEL P., DE VRIES, GAAITZEN J., & DE VRIES, KLAAS. 2015. Patterns of Structural Change in Developing Countries. *Pages 65–83 of: WEISS, JOHN, & TRIBE, MICHAEL (eds), Routledge Handbook of Industry and Development*. London: Routledge.
- UNITED NATIONS STATISTICS DIVISION. 2015. *National Accounts Main Aggregates Database: GDP and Its Breakdown at Current Prices and at Constant 2005 Prices in US Dollars*. United Nations, New York.
- WOLTJER, PIETER, GOUMA, REITZE, & TIMMER, MARCEL P. 2021 (Dec.). *Long-run World Input-Output Database: Version 1.1 Sources and Methods*. GGDC Research Memorandum 190. Groningen Growth and Development Centre, University of Groningen.
- WORLD BANK. 2008. *Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Program*. Washington, DC: World Bank.
- WORLD BANK. 2026. *World Development Indicators*. World Bank, Washington, DC. Series SP.URB.TOTL.IN.ZS, last updated April 8, 2026.
- YOUNG, ALWYN. 2013. Inequality, the Urban-Rural Gap, and Migration. *Quarterly Journal of Economics*, **128**(4), 1727–1785.

Online Appendix

Appendix Contents

A Construction of Model Inputs and Inversion of Fundamentals	54
A.1 Internationally Comparable Gross-Output Prices	55
A.2 Sectoral Wages	57
A.3 Bilateral Trade Flows and Self-Trade Shares	60
A.4 Rural-Urban Decomposition	61
A.5 Static Inversion: Sectoral and Regional TFP	63
A.6 An Alternative Expression for TFP Based on Labor Productivity Chaining .	67
A.7 Static Inversion: Bilateral Trade Costs	70
A.8 Dynamic Inversion: Spatial Labor Market Wedges	71
B Proofs of Propositions 1 and 2	76
B.1 Proof of Proposition 1	76
B.2 Proof of Proposition 2	80
C Proof of Proposition 3	83
C.1 The Static Equilibrium of the Special Case	83
C.2 Dynamics after a Manufacturing Productivity Shock	84
C.3 Thresholds on the Agricultural Employment Share	87
C.4 Proof of the Closed-Form Proposition	87
D Manufacturing Peaks and Sectoral TFP Growth	88
E Frictionless Benchmark with Recalibrated Preferences	95
F Empirical Appendix	98
F.1 Definition of Variables	98
F.2 Additional Empirical Results and Robustness	100

A Construction of Model Inputs and Inversion of Fundamentals

This appendix describes how we recover the structural fundamentals of the model from data. The main fundamentals include sectoral productivities $\{A_n^j(t)\}$, bilateral trade costs $\{\kappa_{ni}^j(t)\}$, IO weights $\Gamma^{n,j}(t)$ and rural-urban labor wedges $\{\tau_n^{RU}(t), \tau_n^{UR}(t)\}$. We set $\theta^a = 4$ in agriculture, $\theta^m = 4$ in manufacturing and $\theta^s = 4$ in services, in line with the estimates from Simonovska and Waugh (2014). The sector-switching model of Section 5.5 instead uses $\theta^a = 3.92$, $\theta^m = 1.61$ and $\theta^s = 1.91$, and its productivities and trade costs are inverted at these internally calibrated trade elasticities. The across-variety CES elasticity is set to $\eta = 2$, following Caliendo, Dvorkin, and Parro (2019).

Sections A.1–A.4 cover the underlying panels. Sections A.5–A.8 cover the model inversion. Table A.1 summarizes calibrated parameters and targets. The model covers 43 countries from 1970 to 2013.

TABLE A.1: CALIBRATED PARAMETERS

Parameter	Symbol	Value	Source
<i>Preferences</i>			
Discount factor	β	0.96	annual
CES preference elasticity	σ	0.56	internal
Ag. subsistence	$\bar{c}_a$	calibrated	target median $s_0^{n,a} = 0.27$
Svc. luxury shift	$\bar{c}_s$	calibrated	target median $ s_0^{n,s} = 0.14$
Sectoral weights	α_j	calibrated	match labor shares in 1970
Migration Gumbel scale	ν_n	country-spec.	NLS on $\{3, \dots, 100\}$
<i>Production</i>			
Fréchet elasticity	θ^j	4, 4, 4	Simonovska and Waugh (2014)
Across-variety elasticity	η	2	Caliendo, Dvorkin, and Parro (2019)
IO weights	$\gamma^{nj}, \gamma^{nj,nk}$	data	WIOD, k -NN imputation, RAS
<i>Geography & numerics</i>			
Sectors	J	3	ag, manu, svc
Countries	C	43	ROW excluded
Regions per country	R_c	2	rural / urban
Total regions	N	84	$C \times R_c$
Model base year	–	1970	
Price benchmark year	–	2005	GGDC, ICP
Time horizon	T	100	yrs

Notes: 1970 Base Year, 43 Countries, 84 Regions

A.1 Internationally Comparable Gross-Output Prices

The static inversion below requires, for each country n , sector j , and year t , an internationally comparable gross-output price $P_n^j(t)$ in common units.

We start from a sectoral value-added deflator

$$\hat{P}_n^{j,\text{va}}(t) = \frac{V_n^j(t)}{Q_n^j(t)},$$

where V and Q are nominal and constant-price value added in sector j . For every country except Taiwan, Hong Kong and Bolivia we use UNIDO MEI (United Nations Statistics Division 2015). They report sectoral value added in both current U.S. dollars and constant 2005 U.S. dollars.¹⁰ UNIDO coverage begins later than GGDC coverage for most countries. For the years before the first year in which UNIDO reports all three sectors, we use the GGDC 10 sector database, scaled to match the UNIDO deflator. GGDC is in local currency units so we multiply the price series by the PWT market exchange-rate ratio $\text{xr}_n(2005)/\text{xr}_n(t)$. Taiwan, Hong Kong and Bolivia are not in UNIDO. For them we use GGDC throughout. We smooth these series before applying them so that the model smooths through exchange-rate volatility.

All value-added series are smoothed before taking the deflator. Different to most papers, we do not use the HP filter. The reason is that, due to the famous exchange rate disconnect, nominal exchange rate adjustments can have very long half lives that the HP filter is not designed to smooth through. Even with a large λ the HP filter can leave decade-long humps that pick up exchange rate factors. These humps flowed through to fundamentals and led to non-convergence in our dynamic model solver. Instead, for each country and sector, we regress the log of the series on a time trend and its square by a robust Huber regression. We replace the series with these fitted values. If the fitted quadratic changes direction inside the sample, we use the linear fit instead. This smoother is used at every smoothing step in the model inversion, except for rural employment shares, which use a five-year centered moving average. This is more aggressive smoothing than the HP filter, which Sposi, Yi, and Zhang (2026) uses. It is less aggressive smoothing than Huneus and Rogerson (2024), which picks end points and does simple average growth rates.

Our model needs the *gross-output* price, which has to absorb intermediate inputs. Let Γ_n be the country's 3×3 matrix of IO coefficients γ_n^{jk} . Let ν_n^j be the share of value added in gross output. We measure ν_n^j directly as value added over gross output and smooth it over time when we have both value added and gross output data at sector level. We then scale each sector's input coefficients so that $\sum_k \gamma_n^{jk} = 1 - \nu_n^j$. The Cobb-Douglas constant

¹⁰So exchange-rate movements are built in.

is $B_n^j = (\nu_n^j)^{-\nu_n^j} \prod_k (\gamma_n^{jk})^{-\gamma_n^{jk}}$. Equilibrium gross-output prices then must satisfy

$$\log P_n^j = \log B_n^j + \nu_n^j \log P_n^{j,\text{va}} + \sum_k \gamma_n^{jk} \log P_n^k,$$

and solving the Leontief system gives the expression for gross output sector prices, $\tilde{P}_n^j(t)$,

$$\tilde{P}_n^j(t) \equiv \exp \left[\left((I - \Gamma_n(t))^{-1} (\nu_n(t) \odot \log \hat{P}_n^{\text{va}}(t) + \log B_n(t)) \right)_j \right].$$

The GGDC Productivity Level Database of Inklaar and Timmer (2014) provides cross country comparable gross output series in 2005. The remaining countries use the World Bank ICP 2005 price indices (World Bank 2008). We convert those levels to dollars at the 2005 market exchange rate. We then scale them by the United States benchmark from the Productivity Level Database. Costa Rica appears in neither source so we impute it with a Balassa-Samuelson regression of log relative prices on log GDP per capita. This gives us a complete set of gross output prices for each sector that are internationally comparable at 2005. We denote this series, $\bar{P}_n^j(2005)$. Then, to obtain a series of prices for each year we chain differences off of this anchored 2005 data. That is, we set

$$P_n^j(t) = \frac{\tilde{P}_n^j(t)}{\tilde{P}_n^j(2005)} \cdot \bar{P}_n^j(2005).$$

This approach is similar to SYZ, arrived to independently. We learned to impute Costa Rica using the Balassa-Samuelson imputation from SYZ (Section 4.1, Appendix A, eq. E.1).

Input-output structure. The $\Gamma_n(t)$ input-output weights come from WIOD. We splice the 2000-2014 WIOD series with the new WIOD Long Run dataset (Woltjer, Gouma, and Timmer 2021) which covers 1965-2000. We scale the long-run dataset so that the levels match at the 2000 boundary.

For countries not in WIOD we impute their Input-Output structure. We prioritize getting value-added shares right, so we identify the value-added share $\nu_n^j(t)$ first, and impute the value added shares of missing countries first. We do this via k -NN imputation that fits value added based off the average of the three closest economies. “Closeness” of economies means economies that are the closest in terms of log GDP per capita, sector value added shares of total GDP, and a region dummy. Once the value added share is imputed, we then impute the weights on each of the inputs via the same process but restricting that the weights must sum up to $1 - \nu_n^j(t)$. This prioritizes getting the level of value added shares right, and the relative input weights right rather than the level of all the inputs right. As one final step to

make sure the matrices satisfy national accounts we run RAS on the original WIOD weights of sample countries to impose two identities: 1) intermediates plus final demand must equal sector j 's gross output, 2) intermediate purchases by a sector must equal non-value-added payments. This final RAS step follows SYZ and was learned from that paper.

Throughout this process we smooth IO weights three times. We smooth once before the nearest-neighbor imputation, once after it, and once again after the balancing.

A.2 Sectoral Wages

GGDC reports sectoral value added and employment in local currency. The first step we do is fill in countries with short gaps of missing years with simple extrapolation from average growth anchored to the nearest units where we have data. GGDC sector value added data stops at 2011 or 2012 for most countries, so we fill in the gaps to 2014 from either WIOD or simple extrapolation. The simple extrapolation is just a log-linear trend fitted to the last 10 available years. Countries without internationally comparable gross-output price data constructed earlier are dropped. We drop any years where more than half of the remaining countries lack these gross-output prices which is why our balanced panel starts at 1970, despite GGDC having data for many countries as far back as 1950. Employment shares and sector wage differentials are then built on this balanced panel from 1970-2014. We now describe how the employment shares and wage differentials are constructed. The first piece is a sectoral wage premium, that we define as

$$\text{prem}_n^j(t) = \frac{\text{VA}_n^j(t)/\text{EMP}_n^j(t)}{\sum_k \text{VA}_n^k(t)/\sum_k \text{EMP}_n^k(t)},$$

where $\text{VA}_n^j(t)$ and $\text{EMP}_n^j(t)$ are GGDC value added and employment in sector j , country n , year t respectively. We do this so that we have a unitless measure of wage differentials between sectors. Missing data points are imputed using k -NN on log GDP per capita and sectoral employment shares.

For levels, we then anchor aggregate GDP to PWT 10.01 nominal aggregate GDP $G_n(t) = \text{cgdp}_n(t) \cdot \text{pl_gdp}_n(t)$. We choose the scalar $\bar{w}_n(t)$ so that the sectoral wage bill matches $G_n(t)$:

$$\bar{w}_n(t) = \frac{G_n(t)}{\sum_j \text{prem}_n^j(t) \cdot L_n^j(t)}$$

where $L_n^j(t)$ denotes PWT population times the GGDC sectoral employment share. We call $L_n^j(t)$ “EMP POP” since it represents the share of the *population* that would be associated with a given sector rather than a specific employee headcount. This definition makes the

model consistent with national accounts, since in the model every household works and represents a worker. The level of wages in each sector is then obtained using this PWT anchor as:

$$w_n^j(t) = \text{prem}_n^j(t) \cdot \bar{w}_n(t).$$

In short, PWT gives us national aggregates to anchor, while GGDC gives us sector-level differences.

Robustness: stripping capital income. Because our sectoral wage is value added per worker, part of the agricultural wage gap could in principle be due to differences in capital intensity. Since our model absorbs payments to capital into payments to labor, if payments to capital are vastly higher in manufacturing then this could be inflating the wage. The agricultural wage gap could in that case be an illusion.

To check this, we rebuild the wage panel under two capital-stripped variants. The first is labor compensation per worker, $w_n^{1,j} = \text{LS}_n^j w_n^j$, with LS_n^j denoting a self-employment-adjusted sectoral labor share (WIOD Socio-Economic Accounts relative structure, k -NN imputed for the 21 countries not in them, anchored to the PWT aggregate labor share). The second is value added net of investment, $w_n^{2,j} = w_n^j(1 - \lambda_n(1 - \text{LS}_n^j))$. Here, λ_n denotes the country's investment-to-capital-income ratio. Figure A.1 plots the manufacturing-to-agricultural wage ratio under the three conventions. We find that stripping capital only compresses the premium modestly. At 2010 the cross-country median falls from 2.51 to 2.30. The rank correlation with the raw ratio is 0.97. The large premiums in poor countries remain. For example, Tanzania falls from 10.7 to 7.0, Indonesia from 6.3 to 3.6 and Brazil from 4.3 to 3.2. Differences in capital intensity between sectors cannot account for the wage gap between agriculture and manufacturing.




FIGURE A.1: AGRICULTURAL WAGE GAP WITH AND WITHOUT CAPITAL INCOME

Notes: Each panel plots the ratio of the manufacturing wage to the agricultural wage on a log scale, by country and year, under three wage concepts: value added per worker (raw), labor compensation per worker, and value added net of investment per worker. The gray line marks parity. Rest of world is excluded.

A.3 Bilateral Trade Flows and Self-Trade Shares

Bilateral goods trade for 1962-2000 comes from Feenstra World Trade Flows (Feenstra *et al.* 2005), and for 1995-2023 from BACI HS92 (Gaulier and Zignago 2010). We stitch these datasets at the 1995 overlap by scaling the whole Feenstra series so that aggregate importer data matches in 1995. This adjusts aggregate levels without disturbing the bilateral share structure of Feenstra. For agriculture, we then replace part of the stitched matrix with WIOD. For every year and every importer that WIOD covers we replace the trade data with WIOD trade data. This is because trade data is relatively poorer on agricultural goods than other goods. As such, WIOD is better for these goods because it is imputed to be consistent with national accounts rather than relying directly on cleaned customs data, which is the approach that Feenstra and BACI use. WIOD members outside our 43 countries are pooled into one WIOD ROW lump. Within this lump we revert to trade shares from BACI and Feenstra. Manufacturing keeps the trade data from Feenstra and BACI rather than WIOD.

Is this switch from Feenstra and BACI to WIOD in agriculture worth it? Simple accounting says that sector's wage bill must equal its sales times the value-added share ν_n^j . In agriculture, moving from trade data to WIOD moves the discrepancy from these two sources from over 100% to just 10%.

For services trade we again use WIOD data as the backbone. We impute services trade in three steps. First, for pairs covered by WIOD we use the WIOD trade levels but record down trade in shares, not levels. Pairs that WIOD does not cover use a logit-lasso regression of bilateral services trade shares on log services exports, log services imports, log GDP per capita, and sectoral employment shares of both partners. A third step then converts these trade shares into levels using WIOD's implied levels of services trade relative to goods trade.

Next, the inversion needs domestic absorption (self trade shares) π_{nn}^j , which trade data does not record, but is again recorded via input-output tables in WIOD. This self trade object is filled via one of four routes. Where we can find self-trade shares from WIOD we use those. Where WIOD has no observation, we close the diagonal from the wage bill, as set out below. Where this closure is infeasible we fall back on a logit-Lasso self-share imputation. The Lasso is estimated on sector dummies, year fixed effects, employment shares, log GDP per capita and trade openness. We clip the re-export hubs, Singapore and Hong Kong. Where the predictors run out before 2014 we carry the last observed self-share forward as simple extrapolation. Here, flow_{ni} denotes n 's purchases from i . Where we observe WIOD and where we use the Lasso fallback we fill the self-trade levels as

$$\text{flow}_{nn} = \left(\sum_{i \neq n} \text{flow}_{ni} \right) \cdot \frac{S_n^j}{1 - S_n^j},$$

where S_n^j is the self trade share in country n , sector j obtained from either WIOD or Lasso.

For entries that we close via the wage bill we proceed as follows. Denote sector gross output R_n^j as wage bill divided by value-added share,

$$R_n^j(t) = \frac{w_n^j(t) L_n^j(t)}{\nu_n^j(t)}.$$

Then accounting means that we must have

$$\text{flow}_{nn} = R_n^j(t) - \sum_{i \neq n} \text{flow}_{in}.$$

The self-share is then the diagonal over total absorption. We do not touch observed off-diagonal exports, so bilateral patterns are preserved. This approach fails for re-export (entrepot) hubs such as Singapore and Hong Kong, which motivates our Lasso fallback. Of the 6450 country-sector-years in the panel, 2240 use the WIOD self-share, and 1120 are agriculture columns replaced by WIOD. 2908 are closed from the wage bill. The remaining 182 use the Lasso imputation.

Two more steps stand between this rescaling, and the matrices used in the model. The services diagonal is adjusted by imperfect coverage:

$$S_n^{\text{adj},j} = \frac{0.95 S_n^j}{1 - 0.05 S_n^j}.$$

This corrects for the roughly five percent of countries with no imputed services trade flows. We then rebuild trade levels from these shares.

A.4 Rural-Urban Decomposition

The dynamic block needs location-specific (rural/urban) employment stocks and wages within each country.

Rural shares. Let $\alpha_{n,t}^j \in [0, 1]$ be the rural share of sector j employment in country n at time t . We assume that agriculture is conducted entirely rurally, while services and manufacturing can be done in any region. Thus we set $\alpha_{n,t}^{\text{ag}} = 1$. We use IPUMS international census micro-data, from which we compute rural-rural (RR) and urban-urban (UU) stocks at the OCCISCO occupation level which we then map to three sectors. This sets

$$\hat{\alpha}_{n,t}^{j,\text{IP}} = \frac{\sum_{c \in S_j} RR_c}{\sum_{c \in S_j} (RR_c + UU_c)}.$$

IPUMS labels are noisy and only available at census frequency, so we project them onto a Lasso in country-year variables to get smooth predictions $\hat{\alpha}_{n,t}^j$ on the full panel. The lasso variables are log GDP per capita, sectoral employment shares, year as well as the World Bank urbanization share (World Bank 2026). We fit a cross-validated Lasso for manufacturing and services separately. The training sample is 49 country-year observations across 20 countries. We exclude China from the training sample. Its hukou-based urban definition is incompatible with both the World Bank urbanization and the IPUMS measures. We then use the fitted values from these predicted weights for all countries, sectors, and year - including where we have the raw IPUMS data because the World Bank urbanization rate is very predictive. In order to make the final shares consistent with the overall urban share of that country, we then replace the five-year centered moving average of the Lasso predictions, clipped to $[0, 1]$, with the closest pair of weights $(\alpha_{n,t}^{\text{manu}}, \alpha_{n,t}^{\text{svc}})$, that satisfies

$$s_{n,t}^{\text{manu}} \alpha_{n,t}^{\text{manu}} + s_{n,t}^{\text{svc}} \alpha_{n,t}^{\text{svc}} = (1 - \text{urban_share}_{n,t}^{\text{WB}}) - s_{n,t}^{\text{ag}},$$

where $s_{n,t}^j$ is sector employment shares, and $0 \leq \alpha_{n,t}^j \leq 1$. when $s_{n,t}^{\text{ag}} + \text{urban_share}_{n,t}^{\text{WB}} > 1$ then the adding up constraint is infeasible at $\alpha^{\text{ag}} = 1$. In words, this is just saying that agricultural employment is greater than the World Bank urban share implied rural population. In that case, we relax the assumption that all agriculture is rural and set $\alpha^{\text{manu}} = \alpha^{\text{svc}} = 0$ and back out $\alpha^{\text{ag}} < 1$ from the World Bank identity. This regime covers 14.3 percent of the country-year observations. Hong Kong and Singapore are defined as fully urban in every sector and do not have rural regions. Taiwan has the year-by-year average of the Korean and Japanese shares. This is because the World Bank reports no urban share for Taiwan.

Stocks. Aggregate rural and urban employment within a country are

$$R_{n,t} = \sum_j \alpha_{n,t}^j L_{n,t}^j, \quad U_{n,t} = N_{n,t} - R_{n,t},$$

We normalize for population growth using the factor, $g_t = N_{n,t+1}/N_{n,t}$. The net movement of workers from rural to urban areas (or vice versa) is given by:

$$\text{net_RU}_{n,t} = R_{n,t} \cdot \frac{N_{n,t+1}}{N_{n,t}} - R_{n,t+1}.$$

Rural and urban wages. We assume that wages are pinned down at the region level and not the sector level, and that workers can move frictionlessly within regions. A rural worker

earns $w_{n,t}^R$ and an urban worker $w_{n,t}^U$. The sectoral wage we observe is then the employment weighted of the two regional wages:

$$w_{n,t}^j = \alpha_{n,t}^j w_{n,t}^R + (1 - \alpha_{n,t}^j) w_{n,t}^U, \quad j \in \{\text{ag}, \text{manu}, \text{svc}\}.$$

Three sector-level wages cannot be matched by only two regional wages. We thus impose this identity for agriculture and for the combined manufacturing and services wage bill. Since agriculture is entirely rural we are able to pin down the rural wage as the agricultural wage, $w_{n,t}^R = w_{n,t}^{\text{ag}}$. In some countries, commodity booms and high capital intensity send agricultural wages higher than the rest of the economy. Usually this reflects commodity boom measurement problems. We thus cap any potential rural wage premia at the Australian level. Australia represents the most efficient and capital intensive agricultural sector in the world. So when other countries, such as Argentina, temporarily have implied rural wages that boom higher than Australia, we cap them at Australian values. To solve for the urban wage we then simply solve for the wage bill identity, given the known rural wage:

$$w_{n,t}^U = \frac{w^{\text{manu}} L^{\text{manu}} + w^{\text{svc}} L^{\text{svc}} - (\alpha^{\text{manu}} L^{\text{manu}} + \alpha^{\text{svc}} L^{\text{svc}}) w^{\text{ag}}}{(1 - \alpha^{\text{manu}}) L^{\text{manu}} + (1 - \alpha^{\text{svc}}) L^{\text{svc}}}.$$

In the capped observations, and observations where $\alpha_{n,t}^{\text{ag}} < 1$, where urban farm workers earn $w_{n,t}^U$, regional wage bills match national wage bills only approximately.

A.5 Static Inversion: Sectoral and Regional TFP

Firms in country n and sector j draw productivities from a Fréchet distribution with location parameter T_n^j and shape θ . We define the value-added-level TFP A_n^j implicitly through $T_n^j = (A_n^j)^{\nu_n^j \theta}$. This normalization interprets A as a productivity at the value-added level. Given the wage and price panels of Sections A.1–A.2, the unit cost of producing gross output in sector j is

$$x_n^j(t) = B_n^j(t) \cdot (w_n^j(t))^{\nu_n^j(t)} \cdot \prod_{k=1}^3 (P_n^{k,\text{int}}(t))^{\gamma_n^{jk}(t)}.$$

The model's domestic self-trade share satisfies

$$\pi_{nn}^j = \Gamma_{EK}^{-\theta} \cdot T_n^j \cdot (x_n^j)^{-\theta} \cdot (P_n^j)^\theta,$$

where $\Gamma_{EK} = \Gamma(1 + (1 - \eta)/\theta)^{1/(1-\eta)}$. Note that Γ_{EK} is common across countries but is not neutral across countries. It enters A_n^j as Γ_{EK}^{1/ν_n^j} , so it moves with the IO structure. Inverting

for TFP gives the classic expression

$$A_n^j(t) = \left[\Gamma_{EK} \cdot \frac{x_n^j(t)}{P_n^j(t)} \cdot \pi_{nn}^j(t)^{1/\theta} \right]^{1/\nu_n^j(t)}. \quad (21)$$

The self-trade share term $\pi_{nn}^j(t)^{1/\theta}$ is the Ricardian-selection correction. Finicelli, Pagano, and Sbracia (2013) highlights that international competition holds average productivity of home firms above the fundamental $T_n^j(t)$ by a factor of $\pi_{nn}^j(t)^{-1/\theta}$. The Cobb-Douglas constant B_n^j shows up in x_n^j in the numerator and in $P_n^{j,\text{int}}$ in the denominator (via the Leontief inversion of Section A.1). Because the same RAS-adjusted IO structure feeds both, B_n^j cancels in the ratio. Recomputing $P^{\text{model}} = \Gamma_{EK} \cdot x/A^\nu \cdot \pi_{nn}^{1/\theta}$ from the recovered A reproduces the input price to machine precision. This equation is where a substantive trade inversion difference from SYZ arises. Our x_n^j here uses the sectoral wage w_n^j , while theirs uses a national wage.

Derivation of (21). We obtain (21) directly in the following steps.

Step 1 – Fréchet/CES primitives at time t . Producers in country n , sector j draw productivities from a Fréchet distribution with location $T_n^j(t) = (A_n^j(t))^{\nu_n^j(t)\theta}$ and shape θ . Unit costs follow from the Cobb Douglas bundle,

$$x_n^j(t) = B_n^j(t) \cdot (w_n^j(t))^{\nu_n^j(t)} \cdot \prod_k (P_n^{k,\text{int}}(t))^{\gamma_n^{jk}(t)}, \quad \nu_n^j(t) + \sum_k \gamma_n^{jk}(t) = 1. \quad (22)$$

Fix t and suppress the notation for simplicity. Due to perfect competition, the origin i delivers variety with draw z to market n at price $\kappa_{ni}^j x_i^j / z$. The price distribution from i to n is given by the standard formula

$$G_{ni}^j(p) = \Pr \left[\frac{\kappa_{ni}^j x_i^j}{z} \leq p \right] = 1 - F_i^j \left(\frac{\kappa_{ni}^j x_i^j}{p} \right) = 1 - e^{-b_{ni}^j p^\theta}, \quad b_{ni}^j \equiv T_i^j (\kappa_{ni}^j x_i^j)^{-\theta}.$$

The above distribution, $F_i^j(z) = \exp(-T_i^j z^{-\theta})$ is the Fréchet cdf. Each variety is bought from the cheapest origin which yields

$$G_n^j(p) = 1 - \prod_i (1 - G_{ni}^j(p)) = 1 - e^{-\Phi_n^j p^\theta},$$

where we define $\Phi_n^j \equiv \sum_i b_{ni}^j$. Let $dG_n^j(p) = \theta \Phi_n^j p^{\theta-1} e^{-\Phi_n^j p^\theta} dp$. Substituting $u = \Phi_n^j p^\theta$

implies

$$(P_n^j)^{1-\eta} = \int_0^\infty p^{1-\eta} dG_n^j(p) = \int_0^\infty \left(\frac{u}{\Phi_n^j}\right)^{\frac{1-\eta}{\theta}} e^{-u} du = (\Phi_n^j)^{-\frac{1-\eta}{\theta}} \Gamma\left(1 + \frac{1-\eta}{\theta}\right).$$

This expression is finite iff $\eta < 1 + \theta$. Raising to the power $\frac{1}{1-\eta}$ generates the sector level gross-output price index

$$P_n^j(t) = \Gamma_{EK} \cdot [\Phi_n^j(t)]^{-1/\theta}, \quad \Phi_n^j(t) \equiv \sum_i T_i^j(t) (\kappa_{ni}^j(t) x_i^j(t))^{-\theta}. \quad (23)$$

Note that $\Phi_n^j(t)$ is unobserved, but the price-index relation (23) substitutes it out of the algebra.

Step 2 – gravity expressions Origin i undercuts all rivals on a given variety with probability

$$\begin{aligned} \pi_{ni}^j &= \int_0^\infty \prod_{m \neq i} (1 - G_{nm}^j(p)) dG_{ni}^j(p) = \int_0^\infty e^{-(\Phi_n^j - b_{ni}^j)p^\theta} \theta b_{ni}^j p^{\theta-1} e^{-b_{ni}^j p^\theta} dp \\ &\stackrel{s=p^\theta}{=} b_{ni}^j \int_0^\infty e^{-\Phi_n^j s} ds = \frac{b_{ni}^j}{\Phi_n^j}. \end{aligned}$$

Computing the integral up to p rather than ∞ gives $\Pr[p_{ni}^j \leq p, i \text{ cheapest}] = \frac{b_{ni}^j}{\Phi_n^j} (1 - e^{-\Phi_n^j p^\theta}) = \pi_{ni}^j G_n^j(p)$. Thus, conditional on supplying, every origin's transaction price is distributed G_n^j , so expenditure per variety is equalized across origins and the expenditure share equals the sourcing probability π_{ni}^j . Setting $i = n$ with $\kappa_{nn}^j = 1$ and restoring t we get,

$$\pi_{nn}^j(t) = \frac{T_n^j(t) (x_n^j(t))^{-\theta}}{\Phi_n^j(t)}.$$

Solving (23) for the multilateral aggregate, dividing by Γ_{EK} , and then raising both sides to the power $-\theta$ so that the exponent on Φ_n^j becomes $(-\frac{1}{\theta}) \cdot (-\theta) = 1$ we get

$$\frac{P_n^j(t)}{\Gamma_{EK}} = [\Phi_n^j(t)]^{-1/\theta} \implies \Phi_n^j(t) = \left(\frac{P_n^j(t)}{\Gamma_{EK}}\right)^{-\theta} = \Gamma_{EK}^\theta (P_n^j(t))^{-\theta}.$$

Substituting this expression into the domestic share eliminates Φ_n^j and yields,

$$\pi_{nn}^j(t) = \Gamma_{EK}^{-\theta} \cdot T_n^j(t) \cdot (x_n^j(t))^{-\theta} \cdot (P_n^j(t))^\theta. \quad (24)$$

Step 3 – solve for T_n^j . Rearranging (24) and using $T_n^j(t) = (A_n^j(t))^{\nu_n^j(t)\theta}$ yields,

$$(A_n^j(t))^{\nu_n^j(t)\theta} = \Gamma_{EK}^\theta \cdot \frac{(x_n^j(t))^\theta}{(P_n^j(t))^\theta} \cdot \pi_{nn}^j(t). \quad (25)$$

Both the gravity equation and $T = A^{\nu\theta}$ normalization hold for all t .

Step 4 – value added $1/(\nu_n^j(t)\theta)$. Taking (25) to the power $1/(\nu_n^j(t)\theta)$ we get,

$$A_n^j(t) = \left[\Gamma_{EK}^\theta \cdot \frac{(x_n^j(t))^\theta}{(P_n^j(t))^\theta} \cdot \pi_{nn}^j(t) \right]^{1/(\nu_n^j(t)\theta)}. \quad (26)$$

Pulling the common $1/\nu_n^j(t)$ gives,

$$A_n^j(t) = \left[\Gamma_{EK} \cdot \frac{x_n^j(t)}{P_n^j(t)} \cdot \pi_{nn}^j(t)^{1/\theta} \right]^{1/\nu_n^j(t)}. \quad (27)$$

Step 5 – model and code gross-output price indices Equation (27) still references the model object $P_n^j(t)$, which is defined off the unobserved $\Phi_n^j(t)$. The pipeline's $P_n^{j,\text{int}}(t)$ (Section A.1) is built as the Leontief inversion of the value-added deflators $P_n^{V A,j}(t)$ using the *same* weights $\{\nu_n^j(t), \gamma_n^{jk}(t)\}$ and constant $B_n^j(t)$ that enter (22). In this case,

$$\log P_n^{j,\text{int}}(t) = \nu_n^j(t) \log P_n^{V A,j}(t) + \sum_k \gamma_n^{jk}(t) \log P_n^{k,\text{int}}(t) + \log B_n^j(t).$$

The calibration imposes $P_n^j(t) \equiv P_n^{j,\text{int}}(t)$ to absorb the productivity residual into $A_n^j(t)$. As a corollary, $B_n^j(t)$ appears in $\log x_n^j(t)$ and in $\log P_n^{j,\text{int}}(t)$ with the same weight and so cancels in the ratio $x_n^j(t)/P_n^{j,\text{int}}(t)$, as already noted above.

Step 6 – country-level inversion. Substituting $P_n^j(t) \rightarrow P_n^{j,\text{int}}(t)$ in (27) completes the steps to yield the inversion of Section A.5. The final expression is thus:

$$A_n^j(t) = \left[\Gamma_{EK} \cdot \frac{x_n^j(t)}{P_n^{j,\text{int}}(t)} \cdot \pi_{nn}^j(t)^{1/\theta} \right]^{1/\nu_n^j(t)}, \quad (28)$$

which is (21). Note that Γ_{EK} is fixed by (θ, η) , $x_n^j(t)$ is built in Section A.2 wages and Section A.1 materials prices, $P_n^{j,\text{int}}(t)$ is the Leontief-grossed-up gross-output price, $\pi_{nn}^j(t)$ is the domestic self-share from Section A.3, and $\nu_n^j(t)$ is determined from the time- t IO weights.

Step 7 – the location-conditional inversion. The above working gives us an inversion for a sector's TFP, but now we need to see how that splits across regions. Our model splits the labor market into rural and urban with wages: $w_n^R(t), w_n^U(t)$. So, wages differ across

regions and the region's trade openness (self-trade share) also differs. We assume that the IO structure of sectors within the same country is the same across regions. We also assume free-trade between regions within a country so that price indices are also the same. The region- Z unit cost is (22) with $w_n^j \rightarrow w_n^Z$:

$$x_{Z,n}^j(t) = B_n^j(t) \cdot (w_n^Z(t))^{\nu_n^j(t)} \cdot \prod_k (P_n^{k,\text{int}}(t))^{\gamma_n^{jk}(t)}.$$

This gives region-sector level unit costs. But now we need a measure of a region's trade openness in order to complete the Ricardian selection term in the trade inversion. We know that a country's domestic share is the sum of the two regional shares, $\pi_{nn}^j(t) = \pi_{nR}^j(t) + \pi_{nU}^j(t)$. The assumption of free trade within countries means we take the simplest approach that a region's openness corresponds to its wage bill share within the country in that sector. That is,

$$\pi_{nZ}^j(t) = \varphi_n^{Zj}(t) \pi_{nn}^j(t), \quad \text{where } \varphi_n^{Zj}(t) \equiv \frac{w_n^Z(t) L_n^{Zj}(t)}{w_n^R(t) L_n^{Rj}(t) + w_n^U(t) L_n^{Uj}(t)}, \quad (29)$$

where $L_n^{Rj}(t) = \alpha_{n,t}^j L_{n,t}^j$ and $L_n^{Uj}(t) = (1 - \alpha_{n,t}^j) L_{n,t}^j$. Note that this split is not a separate assumption, but is implied by the assumption of free trade within the country, as well as identical international trade costs and value-added shares across regions. The initial period's allocation proceeds similarly. It uses each country-sector's 1970 sales, and so value added is split between the rural and urban regions in proportion to $\varphi_n^{Zj}(1970)$. The resulting region-conditional inversion is

$$A_{Z,n}^j(t) = \left[\Gamma_{EK} \cdot \frac{x_{Z,n}^j(t)}{P_n^{j,\text{int}}(t)} \cdot (\varphi_n^{Zj}(t) \pi_{nn}^j(t))^{1/\theta} \right]^{1/\nu_n^j(t)}, \quad (30)$$

A.6 An Alternative Expression for TFP Based on Labor Productivity Chaining

The above section derived location-conditional TFP through inversion. However, it relies very heavily on the quality of our gross output price index series. This section outlines a procedure to reduce dependence on this series.

We refer to the inversion formula from the previous section as Option A, and now we present Option B.

Option B: 2005-anchored LP propagation. Option A loads the time path of $A_{Z,n}^j$ with two wedges that are not productivity. One is a country-level drift between PWT-anchored GDP and GGDC sectoral value added. This creeps in through the period-by-period wage anchor. The other issue is a sector-level wedge between the UNIDO and GGDC value-added vintages used to build $P^{j,\text{int}}$. Option B anchors at $t_0 = 2005$ from the unsmoothed Option A series. From there we propagate away from this 2005 anchor via simple chaining arithmetic of the model. Define GGDC labor productivity $\text{LP}_n^j(t) = \text{VA_Q05}_n^j(t)/\text{EMP}_n^j(t)$. Also, recall the wage-bill share is denoted by $\varphi_n^{Zj}(t)$, and defined in (29). We can express $A_{Z,n,t}^j$ as:

$$A_{Z,n,t}^j = A_{Z,n,2005}^j \cdot \frac{\text{LP}_{n,t}^j}{\text{LP}_{n,2005}^j} \cdot \frac{w_{n,t}^Z/w_{n,t}^j}{w_{n,2005}^Z/w_{n,2005}^j} \cdot \left(\varphi_{n,t}^{Zj} \pi_{nn,t}^j\right)^{1/(\nu_n^j(t)\theta)} \cdot \left(\varphi_{n,2005}^{Zj} \pi_{nn,2005}^j\right)^{-1/(\nu_n^j(2005)\theta)} \cdot \Gamma_{EK}^{1/\nu_n^j(t)-1/\nu_n^j(2005)}.$$

Here, $w_{n,t}^Z$ is the regional wage of Section A.4 and $w_{n,t}^j$ the sectoral wage of Section A.2. This expression means that we only need to carry the price indices in 2005 inside the $A_{Z,n,2005}^j$ term.

Derivation of the Option B formula. We now re-derive the Option-B level formula from the static EK inversion.

Step 1 – static wage-plug identity. Take

$$A_{Z,n,t}^j = \left[\Gamma_{EK} \cdot \frac{x_{Z,n,t}^j}{P_n^{j,\text{int}}(t)} \cdot \left(\varphi_{n,t}^{Zj} \pi_{nn,t}^j\right)^{1/\theta} \right]^{1/\nu_n^j(t)}, \quad (31)$$

and note

$$x_{Z,n,t}^j = B_n^j(t) \cdot \left(w_n^Z(t)\right)^{\nu_n^j(t)} \cdot \prod_k \left(P_n^{k,\text{int}}(t)\right)^{\gamma_n^{jk}(t)}, \quad (32)$$

with $\nu_n^j(t) + \sum_k \gamma_n^{jk}(t) = 1$.

Step 2 – logs and the two channels of ν_t dependence. Taking logs of (31) and multiplying both sides by $\nu_n^j(t)$ gives

$$\begin{aligned} \nu_n^j(t) \log A_{Z,n,t}^j &= \log \Gamma_{EK} + \log B_n^j(t) + \nu_n^j(t) \log w_n^Z(t) \\ &\quad + \sum_k \gamma_n^{jk}(t) \log P_n^{k,\text{int}}(t) - \log P_n^{j,\text{int}}(t) + \frac{1}{\theta} \log \varphi_{n,t}^{Zj} + \frac{1}{\theta} \log \pi_{nn,t}^j. \end{aligned} \quad (33)$$

Step 3 – cancel materials via the Leontief identity. Recall that the gross-output (inter-

mediate) price can be written in logs as:

$$\log P_n^{j,\text{int}}(t) = \nu_n^j(t) \log P_n^{VA,j}(t) + \sum_k \gamma_n^{jk}(t) \log P_n^{k,\text{int}}(t) + \log B_n^j(t). \quad (34)$$

Substituting (34) into (33), the $\log B_n^j(t)$ terms cancel and the materials bundle $\sum_k \gamma_n^{jk}(t) \log P_n^{k,\text{int}}(t)$ cancels, leaving

$$\nu_n^j(t) \log A_{Z,n,t}^j = \log \Gamma_{EK} + \nu_n^j(t) [\log w_n^Z(t) - \log P_n^{VA,j}(t)] + \frac{1}{\theta} [\log \varphi_{n,t}^{Zj} + \log \pi_{nn,t}^j]. \quad (35)$$

Step 4 – collapse w/P^{VA} to labor productivity. Note that we can define labor productivity, $\text{LP}_{n,t}^j$ via:

$$\frac{V_n^j(t)/L_n^j(t)}{V_n^j(t)/V_n^{j,Q05}(t)} = \frac{V_n^{j,Q05}(t)}{L_n^j(t)} \equiv \text{LP}_{n,t}^j.$$

Value added per worker, $V_n^j(t)/L_n^j(t)$, measures the sectoral wage $w_n^j(t)$ rather than the regional wage $w_n^Z(t)$. Note that option B assumes that the sectoral real wage is proportional to GGDC real value added per worker, $w_n^j(t)/P_n^{VA,j}(t) = c_n^j \text{LP}_{n,t}^j$ with c_n^j constant, so

$$\frac{w_n^Z(t)}{P_n^{VA,j}(t)} = c_n^j \text{LP}_{n,t}^j \frac{w_n^Z(t)}{w_n^j(t)}.$$

Step 5 – the pointwise Option B identity. Combining Steps 3 and 4 and dividing by $\nu_n^j(t)$ gives

$$\log A_{Z,n,t}^j = \frac{\log \Gamma_{EK}}{\nu_n^j(t)} + \log \text{LP}_{n,t}^j + \log c_n^j + \log \frac{w_n^Z(t)}{w_n^j(t)} + \frac{1}{\nu_n^j(t) \theta} [\log \varphi_{n,t}^{Zj} + \log \pi_{nn,t}^j]. \quad (36)$$

Step 6 – difference between t and $t_0 = 2005$. Applying (36) at t and t_0 yields:

$$\begin{aligned} \log A_{Z,n,t}^j - \log A_{Z,n,t_0}^j &= [\log \text{LP}_{n,t}^j - \log \text{LP}_{n,t_0}^j] + \log \frac{w_n^Z(t)/w_n^j(t)}{w_n^Z(t_0)/w_n^j(t_0)} \\ &\quad + \frac{1}{\nu_n^j(t) \theta} \log(\varphi_{n,t}^{Zj} \pi_{nn,t}^j) - \frac{1}{\nu_n^j(t_0) \theta} \log(\varphi_{n,t_0}^{Zj} \pi_{nn,t_0}^j) \\ &\quad + \log \Gamma_{EK} \left[\frac{1}{\nu_n^j(t)} - \frac{1}{\nu_n^j(t_0)} \right]. \end{aligned} \quad (37)$$

Step 7 – comparison without time varying IO weights ratio formula. If IO weights are

constant over time we get

$$[\log A_{Z,n,t}^j - \log A_{Z,n,t_0}^j]_{\text{naive}} = \log \left(\frac{\text{LP}_{n,t}^j}{\text{LP}_{n,t_0}^j} \right) + \log \frac{w_n^Z(t)/w_n^j(t)}{w_n^Z(t_0)/w_n^j(t_0)} + \frac{1}{\nu_n^j(t) \theta} \log \left(\frac{\varphi_{n,t}^{Zj} \pi_{nn,t}^j}{\varphi_{n,t_0}^{Zj} \pi_{nn,t_0}^j} \right). \quad (38)$$

With time-varying IO weights, $\nu_n^j(t_0) \rightarrow \nu_n^j(t)$, in (37) we pick up the extra term

$$\varepsilon = \left[\frac{1}{\nu_n^j(t) \theta} - \frac{1}{\nu_n^j(t_0) \theta} \right] \log(\varphi_{n,t_0}^{Zj} \pi_{nn,t_0}^j) + \left[\frac{1}{\nu_n^j(t) \theta} - \frac{1}{\nu_n^j(t_0) \theta} \right] \log \Gamma_{EK}. \quad (39)$$

Step 8 – exponentiate to the level formula. Rearranging (37) and exponentiating yields:

$$\begin{aligned} A_{Z,n,t}^j &= A_{Z,n,t_0}^j \cdot \frac{\text{LP}_{n,t}^j}{\text{LP}_{n,t_0}^j} \cdot \frac{w_n^Z(t)/w_n^j(t)}{w_n^Z(t_0)/w_n^j(t_0)} \\ &\quad \cdot (\varphi_{n,t}^{Zj} \pi_{nn,t}^j)^{1/(\nu_n^j(t) \theta)} \cdot (\varphi_{n,t_0}^{Zj} \pi_{nn,t_0}^j)^{-1/(\nu_n^j(t_0) \theta)} \cdot \Gamma_{EK}^{1/\nu_n^j(t) - 1/\nu_n^j(t_0)}. \end{aligned} \quad (40)$$

All countries use this TFP inversion. Since China is an important economy where wages are low relative to other productivity measures, we also multiply inverted manufacturing productivity in China by an export-import penetration wedge. This is a common factor $\exp(\zeta_t)$ on rural and urban TFP, with $\zeta_{1970} = 0$ and ζ_t piecewise linear between dates 1975, 1980, ..., 2010 and 2013 and constant thereafter. We choose this wedge so that the model best matches China's share of global manufacturing absorption which rises from 0.05 percent in 1975 to 3.8 percent in 2013.

A.7 Static Inversion: Bilateral Trade Costs

Given the bilateral expenditure-share matrices π_{ni}^j from Section A.3, we invert for bilateral iceberg trade costs κ_{ni}^j through the symmetric Head-Ries inversion. This process uses the gravity equation to obtain

$$\frac{\pi_{ni}^j \cdot \pi_{in}^j}{\pi_{nn}^j \cdot \pi_{ii}^j} = (\kappa_{ni}^j \kappa_{in}^j)^{-\theta},$$

since country-specific terms (A^j, x^j, P^j) cancel. Imposing symmetry $\kappa_{ni}^j = \kappa_{in}^j$ means that we get

$$\kappa_{ni}^j = \left(\frac{\pi_{ni}^j \cdot \pi_{in}^j}{\pi_{nn}^j \cdot \pi_{ii}^j} \right)^{-1/(2\theta)}. \quad (41)$$

By construction, $\kappa_{nn}^j = 1$. The raw Head-Ries ratio is noisy and in many zero trade pairs undefined. For each country pair and sector we therefore fit a linear or monotone quadratic trend to $\log \kappa_{ni}^j$ over the years in which it is finite, and use the fitted values. Outside the

window of years in which a pair is observed we hold κ_{ni}^j constant at its fitted value in the last observed year. We also cap the absolute value of the year-on-year change in inverted $\log \kappa_{ni}^j$ at 0.10 and floor κ_{ni}^j at 1.

A.8 Dynamic Inversion: Spatial Labor Market Wedges

Setup

Note that the within-country value functions are

$$V_t^R = \log\left(\frac{\tilde{w}_t^R}{\tilde{P}_t^n}\right) + \nu \log\left[\exp\left(\frac{\beta V_{t+1}^R}{\nu}\right) + \exp\left(\frac{\beta V_{t+1}^U - \tau_t^{RU}}{\nu}\right)\right],$$

$$V_t^U = \log\left(\frac{\tilde{w}_t^U}{\tilde{P}_t^n}\right) + \nu \log\left[\exp\left(\frac{\beta V_{t+1}^R - \tau_t^{UR}}{\nu}\right) + \exp\left(\frac{\beta V_{t+1}^U}{\nu}\right)\right],$$

with migration shares given by

$$\mu_t^{RU} = \frac{1}{1 + \exp((\beta \Delta V_{t+1} + \tau_t^{RU})/\nu)}, \quad \mu_t^{UR} = \frac{1}{1 + \exp((- \beta \Delta V_{t+1} + \tau_t^{UR})/\nu)},$$

where $\Delta V_{t+1} \equiv V_{t+1}^R - V_{t+1}^U$.

Backward recursion for ΔV

The price index $\tilde{P}_t^n$ is common to the two regions of country n . The index, $\tilde{P}_t^n$, thus cancels when expressing between region differences, but the subsistence term, K_t^n , does not. Let $\text{softplus}(x) \equiv \log(1 + e^x)$:

$$\begin{aligned} \Delta V_t = & \log\left(\frac{\tilde{w}_t^R}{\tilde{w}_t^U}\right) + \beta \Delta V_{t+1} \\ & + \nu \text{softplus}\left(\frac{-\beta \Delta V_{t+1} - \tau_t^{RU}}{\nu}\right) - \nu \text{softplus}\left(\frac{\beta \Delta V_{t+1} - \tau_t^{UR}}{\nu}\right). \end{aligned} \quad (42)$$

We set the terminal condition $\Delta V_{\bar{t}+1} = \Delta V^*$, where $\bar{t}$ is the last year of the inversion sample and ΔV^* is the fixed point of (42) which holds incomes and fundamentals constant for years after the sample. This fixed point exists and is unique because the slope of (42) in ΔV_{t+1} , $\beta(1 - \mu_t^{RU} - \mu_t^{UR})$, lies in $(-\beta, \beta)$. The path $\{\Delta V_t\}_{t=\bar{t}}$ is computed through backward iteration path from the observed path of $(\tilde{w}_t^R)/(\tilde{w}_t^U)$, and a guessed spatial labor market wedge path $(\nu, \tau_t^{RU}, \tau_t^{UR})$.

Estimation objective

Given that we know the path of values ΔV , we also know the path of migration shares that follow in closed form, μ_t^{RU}, μ_t^{UR} . These are gross flows. In the data we observe *net* flows. Net flows must satisfy the basic constraint that

$$\widehat{\text{net_RU}}_t = g_t \cdot (\mu_t^{RU} \cdot R_t - \mu_t^{UR} \cdot U_t),$$

where R_t and U_t are known rural and urban population stocks from the data and $g_t = N_{t+1}/N_t$ controls for the population growth rate. The task is thus to estimate $(\nu, \tau_t^{RU}, \tau_t^{UR})$ so that we back out gross migration flows μ_t^{RU}, μ_t^{UR} that reconstruct the observed net flows, $\text{net_RU}_t^{\text{data}}$. In order to preserve smoothness and avoid under-identification, we parametrize spatial labor market wedges at country level as follows:

$$\tau_{n,t}^{RU} = \max(0, \bar{\tau}_n^{RU} + \delta_n^{RU} s_t), \quad \tau_{n,t}^{UR} = \max(0, \bar{\tau}_n^{UR} + \delta_n^{UR} s_t),$$

where $s_t = (\text{year} - 1992)/21.5$. For each country we choose $(\bar{\tau}^{RU}, \delta^{RU}, \bar{\tau}^{UR}, \delta^{UR})$ by nonlinear least squares to solve

$$\min_{\bar{\tau}^{RU}, \delta^{RU}, \bar{\tau}^{UR}, \delta^{UR}} \sum_{t=1}^{\bar{t}} (\widehat{\text{net_RU}}_t - \text{net_RU}_t^{\text{data}})^2. \quad (43)$$

We pick migration elasticity ν_n separately and then choose the elasticity that gives the best fit. We run the NLS on a 21-point grid of ν_n . The grid is the integers 3 to 15, then 17, 20, 25, 30, 40, 50, 75, and 100. We pick the ν with the highest cumulative R^2 against the observed cumulative net-migration trajectory. This whole dynamic inversion procedure runs in a matter of seconds in Julia. There is a median cumulative R^2 of 0.99. The 25th and 75th percentiles are 0.97 and 0.998.




FIGURE A.2: FIT OF THE RURAL POPULATION SHARE, 1970–2013

Notes: Rural population share by country, 1970 to 2013: model (blue, solid) and data (red, dashed). Hong Kong and Singapore have no rural region and are omitted.




FIGURE A.3: FIT OF THE RURAL-URBAN WAGE GAP, 1970–2013

Notes: Ratio of the urban to the rural wage, w_U/w_R , by country, 1970 to 2013: model (blue, solid) and data (red, dashed). Hong Kong and Singapore have no rural region and are omitted.




FIGURE A.4: EMPLOYMENT SHARES IN DATA VS. MODEL, 1970–2013

Notes: Global rural population share and employment share by sector from 1970 to 2013. Model (solid) and data (dashed).

B Proofs of Propositions 1 and 2

This appendix provides derivations of the dynamic-hat-algebra system in Propositions 1 and 2. There are two technical generalizations relative to Caliendo, Dvorkin, and Parro (2019): (i) Stone–Geary shifts in agriculture ($\bar{c}_a > 0$) and services ($\bar{c}_s < 0$), (ii) time-varying input-output weights $\gamma_t^{nj}, \gamma_t^{nj,nk}$.

Notation. For any positive quantity y_t denote $\dot{y}_{t+1} \equiv y_{t+1}/y_t$. For any object z_t denote $\Delta z_{t+1} \equiv z_{t+1} - z_t$. We also make repeated use of the identity

$$a_{t+1} \log b_{t+1} - a_t \log b_t = a_{t+1} \log \dot{b}_{t+1} + \Delta a_{t+1} \log b_t, \quad (44)$$

which follows by adding and subtracting $a_{t+1} \log b_t$.

B.1 Proof of Proposition 1

Cost bundle (P1.1)

The unit cost of the input bundle at any date is

$$x_t^{nj} = B_t^{nj} (w_t^n)^{\gamma_t^{nj}} \prod_{k=1}^J (P_t^{nk})^{\gamma_t^{nj,nk}}, \quad B_t^{nj} \equiv (\gamma_t^{nj})^{-\gamma_t^{nj}} \prod_k (\gamma_t^{nj,nk})^{-\gamma_t^{nj,nk}}, \quad (45)$$

where B_t^{nj} is the standard Cobb–Douglas constant. Take logs of (45) at $t+1$ and t and difference:

$$\log \dot{x}_{t+1}^{nj} = \log \dot{B}_{t+1}^{nj} + [\gamma_{t+1}^{nj} \log w_{t+1}^n - \gamma_t^{nj} \log w_t^n] + \sum_k [\gamma_{t+1}^{nj,nk} \log P_{t+1}^{nk} - \gamma_t^{nj,nk} \log P_t^{nk}]. \quad (46)$$

Apply identity (44) to each bracket:

$$\begin{aligned} \gamma_{t+1}^{nj} \log w_{t+1}^n - \gamma_t^{nj} \log w_t^n &= \gamma_{t+1}^{nj} \log \dot{w}_{t+1}^n + \Delta \gamma_{t+1}^{nj} \log w_t^n, \\ \gamma_{t+1}^{nj,nk} \log P_{t+1}^{nk} - \gamma_t^{nj,nk} \log P_t^{nk} &= \gamma_{t+1}^{nj,nk} \log \dot{P}_{t+1}^{nk} + \Delta \gamma_{t+1}^{nj,nk} \log P_t^{nk}. \end{aligned}$$

Substituting back into (46) collects the level terms into

$$\Delta \gamma_{t+1}^{nj} \log w_t^n + \sum_k \Delta \gamma_{t+1}^{nj,nk} \log P_t^{nk}.$$

The CRS condition at t and $t + 1$ gives $\Delta\gamma_{t+1}^{nj} + \sum_k \Delta\gamma_{t+1}^{nj,nk} = 0$, so $\Delta\gamma_{t+1}^{nj} = -\sum_k \Delta\gamma_{t+1}^{nj,nk}$. The level collection above therefore consolidates into $\sum_k \Delta\gamma_{t+1}^{nj,nk} \log(P_t^{nk}/w_t^n)$. Exponentiating:

$$\dot{x}_{t+1}^{nj} = \dot{P}_{t+1}^{nj} \cdot (\dot{w}_{t+1}^n)^{\gamma_{t+1}^{nj}} \prod_k (\dot{P}_{t+1}^{nk})^{\gamma_{t+1}^{nj,nk}} \prod_k \left(\frac{P_t^{nk}}{w_t^n} \right)^{\Delta\gamma_{t+1}^{nj,nk}}. \quad (47)$$

This matches (P1.1) of Proposition 1.

Sectoral price index (P1.2)

Raise both sides of the price-index equation to $-\theta^j$:

$$(P_{t+1}^{nj})^{-\theta^j} = (\Gamma^{nj})^{-\theta^j} \sum_{i=1}^N (x_{t+1}^{ij} \kappa_{t+1}^{nj,ij})^{-\theta^j} (A_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}}. \quad (48)$$

In each summand, multiply and divide by the base-period factor $(x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma_t^{ij}}$. The trade-share term and trade-cost term factor as $(x_{t+1}^{ij} \kappa_{t+1}^{nj,ij})^{-\theta^j}$. Apply identity (44) (in exponent θ^j) to the productivity ratio:

$$\frac{(A_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}}}{(A_t^{ij})^{\theta^j \gamma_t^{ij}}} = (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} \cdot (A_t^{ij})^{\theta^j \Delta\gamma_{t+1}^{ij}}. \quad (49)$$

Substituting:

$$(P_{t+1}^{nj})^{-\theta^j} = (\Gamma^{nj})^{-\theta^j} \sum_i (x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma_t^{ij}} \cdot (x_{t+1}^{ij} \kappa_{t+1}^{nj,ij})^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (A_t^{ij})^{\theta^j \Delta\gamma_{t+1}^{ij}}.$$

The base-period factor inside the sum, by the level definitions of P_t^{nj} and $\pi_t^{nj,ij}$, satisfies

$$(\Gamma^{nj})^{-\theta^j} (x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma_t^{ij}} = (P_t^{nj})^{-\theta^j} \cdot \pi_t^{nj,ij}. \quad (50)$$

Substituting and rearranging:

$$\left(\dot{P}_{t+1}^{nj} \right)^{-\theta^j} = \sum_i \pi_t^{nj,ij} (x_{t+1}^{ij} \kappa_{t+1}^{nj,ij})^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (A_t^{ij})^{\theta^j \Delta\gamma_{t+1}^{ij}}.$$

Raising to $-1/\theta^j$ yields (P1.2). □

Bilateral trade shares (P1.3)

The level definition of $\pi_{t+1}^{nj,ij}$ is the ratio of i 's contribution to the price denominator. By (49), the numerator factors as $(x_t^{ij} \kappa_t^{nj,ij})^{-\theta^j} (A_t^{ij})^{\theta^j \gamma_t^{ij}} \cdot (\dot{x}_{t+1}^{ij} \dot{\kappa}_{t+1}^{nj,ij})^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (\dot{A}_t^{ij})^{\theta^j \Delta \gamma_{t+1}^{ij}}$. By (50) the first factor equals $(\Gamma^{nj})^{-\theta^j} (P_t^{nj})^{-\theta^j} \pi_t^{nj,ij}$. The denominator equals $(\Gamma^{nj})^{-\theta^j} (\dot{P}_{t+1}^{nj})^{-\theta^j} (P_t^{nj})^{-\theta^j}$. Canceling the common $(\Gamma^{nj})^{-\theta^j} (P_t^{nj})^{-\theta^j}$:

$$\pi_{t+1}^{nj,ij} = \pi_t^{nj,ij} \cdot \left(\frac{\dot{x}_{t+1}^{ij} \dot{\kappa}_{t+1}^{nj,ij}}{\dot{P}_{t+1}^{nj}} \right)^{-\theta^j} (\dot{A}_{t+1}^{ij})^{\theta^j \gamma_{t+1}^{ij}} (\dot{A}_t^{ij})^{\theta^j \Delta \gamma_{t+1}^{ij}}.$$

This is (P1.3). □

Stone–Geary share recursion (P1.4)

By definition $s_t^{n,j} \equiv P_t^{n,j} \bar{c}_j / w_t^n$ for $j \in J_{SG} = \{a, s\}$. Moreover,

$$\frac{s_{t+1}^{n,j}}{s_t^{n,j}} = \frac{P_{t+1}^{n,j} \bar{c}_j / w_{t+1}^n}{P_t^{n,j} \bar{c}_j / w_t^n} = \frac{\dot{P}_{t+1}^{n,j}}{\dot{w}_{t+1}^n},$$

delivering (P1.4). □

Supernumerary CES expenditure share and total-income share (P1.5)

Step 1: level shares. The household maximizes $\log C_t^n$ subject to $\sum_k P_t^{n,k} c_t^{n,k} = w_t^n$. Substituting $c^{n,k} = \tilde{c}^{n,k} + \bar{c}_k$:

$$\sum_k P_t^{n,k} \tilde{c}_t^{n,k} = w_t^n - \sum_{j \in J_{SG}} P_t^{n,j} \bar{c}_j \equiv \tilde{Y}_t^n = (1 - S_t^n) w_t^n.$$

Since C_t^n is a homothetic CES in $\tilde{c}$, the standard FOC gives

$$P_t^{n,j} \tilde{c}_t^{n,j} = \tilde{\lambda}_t^{n,j} \tilde{Y}_t^n, \quad \tilde{\lambda}_t^{n,j} \equiv \frac{\alpha_j (P_t^{n,j})^{1-\sigma}}{\sum_k \alpha_k (P_t^{n,k})^{1-\sigma}}, \quad \sum_j \tilde{\lambda}_t^{n,j} = 1.$$

Step 2: total-income share. Total per-capita expenditure on j is $P_t^{n,j} c_t^{n,j} = \mathbf{1}\{j \in J_{SG}\} P_t^{n,j} \bar{c}_j + P_t^{n,j} \tilde{c}_t^{n,j}$. Dividing by w_t^n and using $\tilde{Y}_t^n / w_t^n = 1 - S_t^n$:

$$\Lambda_t^{n,j} = \mathbf{1}\{j \in J_{SG}\} s_t^{n,j} + (1 - S_t^n) \tilde{\lambda}_t^{n,j}, \quad \sum_j \Lambda_t^{n,j} = S_t^n + (1 - S_t^n) = 1.$$

Step 3: hat recursion for $\tilde{\lambda}$. From the level definition,

$$\frac{\tilde{\lambda}_{t+1}^{n,j}}{\tilde{\lambda}_t^{n,j}} = (\dot{P}_{t+1}^{n,j})^{1-\sigma} \cdot \frac{\sum_l \alpha_l (P_t^{n,l})^{1-\sigma}}{\sum_k \alpha_k (P_{t+1}^{n,k})^{1-\sigma}}.$$

For the denominator ratio, multiply and divide each numerator term by $(P_t^{n,k})^{1-\sigma}$:

$$\frac{\sum_k \alpha_k (P_{t+1}^{n,k})^{1-\sigma}}{\sum_l \alpha_l (P_t^{n,l})^{1-\sigma}} = \sum_k \tilde{\lambda}_t^{n,k} (\dot{P}_{t+1}^{n,k})^{1-\sigma}. \quad (51)$$

Therefore

$$\tilde{\lambda}_{t+1}^{n,j} = \tilde{\lambda}_t^{n,j} \cdot \frac{(\dot{P}_{t+1}^{n,j})^{1-\sigma}}{\sum_k \tilde{\lambda}_t^{n,k} (\dot{P}_{t+1}^{n,k})^{1-\sigma}},$$

giving (P1.5). Apply the same step to $\tilde{P}_t^n = (\sum_k \alpha_k (P_t^{n,k})^{1-\sigma})^{1/(1-\sigma)}$ to obtain the supernumerary price-index hat $(\dot{\tilde{P}}_{t+1}^n)^{1-\sigma} = \sum_k \tilde{\lambda}_t^{n,k} (\dot{P}_{t+1}^{n,k})^{1-\sigma}$, used in Proposition 2. $\square$

Goods market clearing (P1.6)

The level identity at $t + 1$ is

$$X_{t+1}^{nj} = \sum_{k=1}^J \gamma_{t+1}^{nk,nj} \sum_i \pi_{t+1}^{ik,nk} X_{t+1}^{ik} + \Lambda_{t+1}^{n,j} (w_{t+1}^n L_{t+1}^n + D_{t+1}^n).$$

The materials term is in change form already. For the consumption term, multiply and divide $w_{t+1}^n L_{t+1}^n$ by $w_t^n L_t^n$:

$$X_{t+1}^{nj} = \sum_k \gamma_{t+1}^{nk,nj} \sum_i \pi_{t+1}^{ik,nk} X_{t+1}^{ik} + \Lambda_{t+1}^{n,j} (\dot{w}_{t+1}^n \dot{L}_{t+1}^n w_t^n L_t^n + D_{t+1}^n),$$

which is (P1.6). Substituting from (P1.5) and using $s_{t+1}^{n,j} \dot{w}_{t+1}^n = s_t^{n,j} \dot{P}_{t+1}^{n,j}$ (from P1.4) and $(1 - S_{t+1}^n) \dot{w}_{t+1}^n = (1 - S_t^n) \dot{w}_t^n$ (proved below) gives the equivalent split form in which the consumption term decomposes into a Stone–Geary outlay $\mathbf{1}\{j \in J_{SG}\} s_t^{n,j} \dot{P}_{t+1}^{n,j}$ and a supernumerary CES expenditure $\tilde{\lambda}_{t+1}^{n,j} (1 - S_t^n) \dot{w}_{t+1}^n$, each multiplied by $(1 + d_{t+1}^c - \bar{d}_{t+1}) \dot{L}_{t+1}^n w_t^n L_t^n$. $\square$

Labor-market clearing and sectoral residual (P1.7–P1.8)

The level labor-market clearing $w_{t+1}^n L_{t+1}^n = \sum_j \gamma_{t+1}^{nj} \sum_i \pi_{t+1}^{ij,nj} X_{t+1}^{ij}$ becomes, after multiplying and dividing the LHS by $w_t^n L_t^n$,

$$\dot{w}_{t+1}^n \dot{L}_{t+1}^n w_t^n L_t^n = \sum_j \gamma_{t+1}^{nj} \sum_i \pi_{t+1}^{ij,nj} X_{t+1}^{ij},$$

which is (P1.7). Sectoral employment is then the static residual $L_{t+1}^{nj} = \gamma_{t+1}^{nj} \sum_i \pi_{t+1}^{ij,nj} X_{t+1}^{ij} / (\dot{w}_{t+1}^n w_t^n)$ — equation (P1.8). Note the absence of L_t^{nj} on the RHS, confirming that within-region sectoral employment carries no dynamic state, in contrast to Caliendo, Dvorkin, and Parro’s (2019) model with persistence through migration over (n, j) pairs. $\square$

Θ Level Dependence

Inspecting (P1.1)–(P1.8), the only level objects entering the date- $(t + 1)$ system are the date- t allocation $\{L_t^n, L_t^{nj}, \pi_t^{nj,ij}, X_t^{nj}, w_t^n, P_t^{nj}, A_t^{nj}, s_t^{n,a}, s_t^{n,s}, \tilde{\lambda}_t^{nj}\}$ and the level A_t^{ij} , as well as the IO weights γ_{t+1} and the imbalances $d_{t+1}^c - \bar{d}_{t+1}$. The Stone–Geary parameters $\bar{c}_a, \bar{c}_s$ enter only through the initial states $s_0^{n,a}, s_0^{n,s}$. From $t = 1$ onwards the recursions only use $\dot{\Theta}_{t+1}$, $\Delta\gamma_{t+1}$, the level d_{t+1} , and the state variables. $\square$

B.2 Proof of Proposition 2

Migration logit recursion (P2.1)

The Gumbel-max calculation of Section 3 gives the level migration logit, integrated over destination regions i :

$$\mu_t^{n,i} = \frac{\exp((\beta V_{t+1}^i - \tau_t^{n,i})/\nu)}{\sum_{m=1}^N \exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu)}. \quad (52)$$

Take the ratio at $t + 1/t$ and group exponentials:

$$\frac{\mu_{t+1}^{n,i}}{\mu_t^{n,i}} = \exp\left(\frac{\beta(V_{t+2}^i - V_{t+1}^i)}{\nu}\right) \exp(-\Delta\tau_{t+1}^{n,i}/\nu) \cdot \frac{\sum_m \exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu)}{\sum_m \exp((\beta V_{t+2}^m - \tau_{t+1}^{n,m})/\nu)}. \quad (53)$$

For the ratio of summations, multiply and divide each numerator term by $\exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu)$:

$$\sum_m \exp((\beta V_{t+2}^m - \tau_{t+1}^{n,m})/\nu) = \sum_m \exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu) \exp\left(\frac{\beta(V_{t+2}^m - V_{t+1}^m)}{\nu}\right) \exp(-\Delta\tau_{t+1}^{n,m}/\nu).$$

Dividing by $\sum_m \exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu)$ and using (52) at date t yields

$$\frac{\sum_m \exp((\beta V_{t+2}^m - \tau_{t+1}^{n,m})/\nu)}{\sum_m \exp((\beta V_{t+1}^m - \tau_t^{n,m})/\nu)} = \sum_m \mu_t^{n,m} \exp\left(\frac{\beta(V_{t+2}^m - V_{t+1}^m)}{\nu}\right) \exp(-\Delta\tau_{t+1}^{n,m}/\nu). \quad (54)$$

Substituting (54) into (53):

$$\mu_{t+1}^{n,i} = \frac{\mu_t^{n,i} \exp(\beta(V_{t+2}^i - V_{t+1}^i)/\nu) \exp(-\Delta\tau_{t+1}^{n,i}/\nu)}{\sum_{m=1}^N \mu_t^{n,m} \exp(\beta(V_{t+2}^m - V_{t+1}^m)/\nu) \exp(-\Delta\tau_{t+1}^{n,m}/\nu)},$$

which is (P2.1). The level $\tau_t^{n,i}$ has dropped out — only the additive change $\Delta\tau_{t+1}^{n,i}$ enters — while the level migration matrix $\mu_t^{n,m}$ enters as a state. $\square$

Supernumerary real-consumption growth

Indirect utility under Stone–Geary CES is $U(C_t^n) = \log C_t^n = \log(\tilde{Y}_t^n/\tilde{P}_t^n) = \log((1 - S_t^n)w_t^n/\tilde{P}_t^n)$. Time-differencing:

$$U(C_{t+1}^n) - U(C_t^n) = \log(\dot{\tilde{Y}}_{t+1}^n/\dot{\tilde{P}}_{t+1}^n) = \log \dot{C}_{t+1}^n. \quad (55)$$

Compute $\dot{\tilde{Y}}_{t+1}^n$ from the level definition $\tilde{Y}_t^n = w_t^n - \sum_{j \in J_{SG}} P_t^{n,j} \bar{c}_j$. In the numerator at $t+1$, write $w_{t+1}^n = \dot{w}_{t+1}^n w_t^n$ and $P_{t+1}^{n,j} \bar{c}_j = \dot{P}_{t+1}^{n,j} \cdot s_t^{n,j} w_t^n$ (using $P_t^{n,j} \bar{c}_j = s_t^{n,j} w_t^n$). The denominator equals $(1 - S_t^n)w_t^n$. Therefore

$$\dot{w}_{t+1}^n \equiv \dot{\tilde{Y}}_{t+1}^n = \frac{\dot{w}_{t+1}^n - s_t^{n,a} \dot{P}_{t+1}^{n,a} - s_t^{n,s} \dot{P}_{t+1}^{n,s}}{1 - S_t^n}. \quad (56)$$

Equivalently, $\tilde{Y}_{t+1}^n/\tilde{Y}_t^n = (1 - S_{t+1}^n)w_{t+1}^n/((1 - S_t^n)w_t^n)$ gives

$$\dot{w}_{t+1}^n = \frac{1 - S_{t+1}^n}{1 - S_t^n} \dot{w}_{t+1}^n, \quad (57)$$

which is the form used in (P2.2). The two forms are algebraically identical (use $s_t^{n,j} \dot{P}_{t+1}^{n,j} = s_{t+1}^{n,j} \dot{w}_{t+1}^n$ from P1.4).

For the supernumerary price-index hat, applying (51) to $\tilde{P}_t^n$ yields $(\dot{\tilde{P}}_{t+1}^n)^{1-\sigma} = \sum_k \tilde{\lambda}_t^{n,k} (\dot{P}_{t+1}^{n,k})^{1-\sigma}$. Combining,

$$\dot{C}_{t+1}^n = \frac{\dot{w}_{t+1}^n}{\dot{\tilde{P}}_{t+1}^n}. \quad (58)$$

In the homothetic limit $\bar{c}_a = \bar{c}_s = 0$, $S_t^n = 0$ and $\dot{w}_{t+1}^n = \dot{w}_{t+1}^n$, $\dot{\tilde{P}}_{t+1}^n = \dot{P}_{t+1}^n$, and (58) reduces to Caliendo, Dvorkin, and Parro's (2019) $\dot{C}_{t+1}^n = \dot{w}_{t+1}^n/\dot{P}_{t+1}^n$. $\square$

Bellman in time differences (P2.2)

The Bellman in levels, with $U(C_t^n) = \log C_t^n$ and integrated taste shock, is

$$V_t^n = \log C_t^n + \nu \log \sum_{i=1}^N \exp((\beta V_{t+1}^i - \tau_t^{n,i})/\nu).$$

Time-differencing:

$$V_{t+1}^n - V_t^n = \log \dot{C}_{t+1}^n + \nu \log \frac{\sum_i \exp((\beta V_{t+2}^i - \tau_{t+1}^{n,i})/\nu)}{\sum_i \exp((\beta V_{t+1}^i - \tau_t^{n,i})/\nu)}.$$

The ratio of summations was already evaluated in (54). Substituting and exponentiating:

$$\exp(V_{t+1}^n - V_t^n) = \dot{C}_{t+1}^n \left(\sum_{i=1}^N \mu_t^{n,i} \exp(\beta(V_{t+2}^i - V_{t+1}^i)/\nu) \exp(-\Delta \tau_{t+1}^{n,i}/\nu) \right)^\nu.$$

Substituting (58) for $\dot{C}_{t+1}^n$ delivers (P2.2). The Stone–Geary modification enters *only* through the supernumerary real-consumption growth $\dot{w}_{t+1}^n/\dot{P}_{t+1}^n$, replacing CDP's $\dot{w}_{t+1}^n/\dot{P}_{t+1}^n$. $\square$

Labor flow accounting (P2.3)

The labor stock in region n at $t + 1$ is the sum of inflows from all origin regions i , $L_{t+1}^n = \sum_i \mu_t^{i,n} L_t^i$ — a level identity that requires no time-differencing. No Stone–Geary modification arises because $\bar{c}_j$ does not enter the migration logit. $\square$

Θ level dependence

Equations (P2.1)–(P2.3) depend only on (i) the changes $\dot{w}_{t+1}^n, \dot{P}_{t+1}^n, \Delta \tau_{t+1}^{n,i}$ and (ii) the contemporaneous state $\{\mu_t^{n,i}, L_t^i, V_{\geq t+1}^i\}$. From the proof of Proposition 1, $\dot{w}_{t+1}^n$ and $\dot{P}_{t+1}^n$ require some levels of Θ . Specifically, the A_t^{ij} correction in (P1.2)–(P1.3), the IO-weight levels γ_{t+1} , and the imbalances d_t^c . Stone–Geary parameters $\bar{c}_a, \bar{c}_s$ enter via $s_0^{n,a}, s_0^{n,s}$, spatial labor market wedges $\tau_t^{n,i}$ enter through $\mu_{-1}^{n,i}$. From $t = 1$ onwards we need $\Delta \tau_{t+1}^{n,i}$ the level d_{t+1} , and $\Delta \gamma_{t+1}^{nj,nk}$. $\square$

C Proof of Proposition 3

This appendix proves Proposition 3 in the special case of Section 4.1.

C.1 The Static Equilibrium of the Special Case

We make the following restrictions on the primitives of country c .

- (i) No intermediate input output network. So $\gamma^{nj} = 1$ for $n \in \{R, U\}$, and $x^{nj} = w^n$.
- (ii) R produces only agriculture and U only manufacturing and services. This is the special case limit $A^{Rm}, A^{Rs}, A^{Ua} \rightarrow 0$.
- (iii) Services are completely non-tradable across countries but freely tradable within countries. This is in contrast to quantitative model where services are internationally tradable with Head-Ries inverted trade costs. This means that for services we have $\kappa^{ns, is} = \kappa^{is, ns} = \infty$ for $n \in \{R, U\}$ and $i \notin c$ but $\kappa^{Rj, Uj} = \kappa^{Uj, Rj} = 1$.
- (iv) The economy is a completely small open economy in agriculture and manufacturing. No trade costs in goods, and international market shares $\pi^{ij, nj} \rightarrow 0$.
- (v) Homothetic preferences and balanced trade each period such that $\bar{c}_a = \bar{c}_s = 0$ and $D^c = D^n = 0$.
- (vi) No international migration, $\tau^{n, i} = \tau^{i, n} = \infty$ for $n \in c, i \notin c$.

Write $L^R = S_a E^c$ and $L^U = (1 - S_a) E^c$. Because the country is small and there is free trade within countries, every price except services is set by the international market and is fixed. Thus, $P^s = \Gamma^{Us} w^U / A^{Us}$ moves one-for-one with the urban wage. With homothetic preferences, every household has the same services expenditure share

$$\lambda^s = \frac{\alpha_s (P^s)^{1-\sigma}}{\sum_k \alpha_k (P^k)^{1-\sigma}}.$$

Note that this expenditure share depends on w^U , and that the consumption gap between the regions equals the wage gap, so that we have $\log C^U - \log C^R = \log(w^U / w^R)$.

Each region faces a constant-elasticity foreign demand for its traded output. Region n sales in sector $j \in \{a, m\}$ are $K^{nj} (w^n / A^{nj})^{-\theta^j}$, where K^{nj} summarizes foreign demand.¹¹ Labor market clearing in agriculture, manufacturing, and services is consequently

$$w^R L^R = K^{Ra} \left(\frac{w^R}{A^{Ra}} \right)^{-\theta^a}, \quad w^U L^{Um} = K^{Um} \left(\frac{w^U}{A^{Um}} \right)^{-\theta^m}, \quad w^U L^{Us} = \lambda^s (w^U L^U + w^R L^R), \quad (59)$$

with $L^{Um} + L^{Us} = L^U$.

¹¹The foreign demand constant is $K^{nj} = \sum_{i \notin c} (\Gamma^{ij} \kappa^{ij, nj} / P^{ij})^{-\theta^j} X^{ij}$.

C.2 Dynamics after a Manufacturing Productivity Shock

Consider an MIT shock at $t = 0$. Suppose at $t = -1$ the economy is in a steady state with $S_a \in (0, 1)$ and $S_m, S_s > 0$. The shock raises A^{Um} unexpectedly and permanently by $\hat{A} \equiv d \log A^{Um} > 0$.

Migration. Let $\Delta V_t \equiv V_t^R - V_t^U$ be as in Appendix A.8. The logit and the Bellman equation from Section 3 yield

$$\mu_t^{RU} = [1 + e^{(\tau^{RU} + \beta \Delta V_{t+1})/\nu}]^{-1}, \quad \mu_t^{UR} = [1 + e^{(\tau^{UR} - \beta \Delta V_{t+1})/\nu}]^{-1},$$

$$\Delta V_t = -\log \frac{w_t^U}{w_t^R} + \beta \Delta V_{t+1} + \nu \log \frac{\mu_t^{UU}}{\mu_t^{RR}},$$

and $S_{a,t+1} = (1 - \mu_t^{RU})S_{a,t} + \mu_t^{UR}(1 - S_{a,t})$. Let $c \equiv \mu^{RU} + \mu^{UR}$ be the steady-state churn rate and $\chi \equiv 1 - c$ be the fraction of stayers. In a steady state, flows must balance, so $S_a \mu^{RU} = (1 - S_a) \mu^{UR}$. Thus, $\mu^{RU} = (1 - S_a)c$ and $\mu^{UR} = S_a c$. Multiplying together the fractions of movers in each direction gives,

$$\mu^{RU} \mu^{UR} = e^{-(\tau^{RU} + \tau^{UR})/\nu} \mu^{RR} \mu^{UU},$$

which yields

$$ac^2 + c = 1 \tag{60}$$

where

$$a \equiv S_a(1 - S_a)(e^{(\tau^{RU} + \tau^{UR})/\nu} - 1).$$

The left hand side increases in c and in a , and exceeds 1 at $c = 1$. So churn $c \in (0, 1)$ falls strictly as the cost $\tau \equiv \tau^{RU} + \tau^{UR}$ rises. Since $ac^2 \leq 1$, it follows that $c \rightarrow 0$ as $\tau \rightarrow \infty$.

To see how $\tau \equiv \tau^{RU} + \tau^{UR}$ affects an economy, we compare economies with the same parameters, fundamentals, and initial allocations and that differ only in τ . For each cost, exactly one split between τ^{RU} and τ^{UR} supports the allocation as a steady state. By the definition of B , $H = (1 + \theta^m - \lambda^s)S_m + \sigma(1 - \lambda^s)S_s > 0$. Since migration costs don't enter into (59), ζ , B , and H are constant in τ . Only χ is an increasing function of τ .

Outline. The proof now analyzes the MIT shock in three steps.

- Within a period, agricultural employment is fixed, so the shock only moves urban workers between manufacturing and services. Services pull B decides which way they move. This is lemma 1 parts (a)–(b)).
- Over time, a higher urban wage draws workers into urban labor markets. At a given urban wage, a share ζ of them end up in manufacturing. This is the farm push M and is Lemma 1 part (c).

- Migration costs determine the speed and size of this worker movement given the new wages (Lemma 2). This is equation (61) below).

How strongly migration responds. Define

$$\bar{\nu} \equiv \frac{\nu(1 - \beta\chi)}{\beta(1 + \chi)}. \quad (61)$$

In the long run, when the urban wage premium w^U/w^R rises by 1%, workers leave agriculture until $(1 - S_a)/S_a$ has risen by $1/\bar{\nu}$ percent (Lemma 1). A smaller $\bar{\nu}$ means a larger contraction and farm push, M .

Write M_0 for M at $\tau = 0$. That is, at $\bar{\nu} = \nu/\beta$. This is the weakest the farm push will ever be, since $M > M_0$ for every $\tau > 0$.

Lemma 1. *To first order in $\hat{A}$:*

(a) *Given $S_{a,t}$, the static equilibrium of period t satisfies*

$$\begin{aligned} dS_{m,t} &= -\frac{\theta^m S_m B}{H} \hat{A} - \frac{\zeta(1 + \theta^m) S_m}{H} dS_{a,t}, \\ d \log \frac{w_t^U}{w_t^R} &= \frac{\theta^m S_m}{H} \hat{A} + W dS_{a,t}, \quad W \equiv \frac{\zeta}{H} + \frac{1}{(1 + \theta^a) S_a} > 0. \end{aligned}$$

(b) *On impact, the move in manufacturing employment satisfies $dS_{m,0} = -\theta^m S_m B \hat{A}/H$.*

(c) *If $\tau > 0$, then $dS_{a,t} = (1 - \rho^t) dS_{a,\infty}$ with $dS_{a,\infty} < 0$, where $\rho \in (0, \chi)$ is the smaller root of*

$$f(r) \equiv \beta\chi r^2 - (1 + \beta\chi^2 + \Omega W)r + \chi, \quad \Omega \equiv \beta S_a(1 - S_a)(1 - \chi^2)/\nu. \quad (62)$$

So $dS_{m,t}$ rises strictly in t to $dS_{m,\infty} = \theta^m S_m \hat{A}(M - B)/(M + H)$.

Proof. (a) and (b). Write $\hat{w} \equiv d \log w^U$, and take farm employment as given. Formally, (59), $S_s = \lambda^s(1 - S_a + S_a w^R/w^U)$, and $d \log \lambda^s = (1 - \sigma)(1 - \lambda^s)\hat{w}$ give

$$d \log w^R = -\frac{dS_a}{(1 + \theta^a) S_a}, \quad dS_m = S_m[\theta^m \hat{A} - (1 + \theta^m)\hat{w}], \quad dS_s = B\hat{w} - (1 - \zeta) dS_a.$$

Since $-dS_a = dS_m + dS_s$, we get $H\hat{w} = \theta^m S_m \hat{A} + \zeta dS_a$. Substituting $\hat{w}$ into dS_m and into $\hat{w} - d \log w^R$, and using $H + B = (1 + \theta^m) S_m$, gives (a). Since nobody moves on impact, ($dS_{a,0} = 0$), from which (b) follows immediately.

(c). Linearizing the migration block gives

$$dS_{a,t+1} = \chi dS_{a,t} + \Omega d\Delta V_{t+1}, \quad d\Delta V_t = -d \log \frac{w_t^U}{w_t^R} + \beta\chi d\Delta V_{t+1}.$$

In the second equation, we can see that the value gap is forward looking in all future gaps. These future gaps are discounted at $\beta\chi$, because starting in a location raises the chance of being there next period by $\chi = \mu^{RR} - \mu^{UR}$. Higher costs thus raise χ and make migration respond more to a more persistent wage gap.

(a) can be used to eliminate ΔV , which then gives a second-order difference equation in $dS_{a,t}$. Define the characteristic polynomial of this difference equation as f . Since $f(0) = \chi > 0 > f(\chi) = -\Omega W\chi$ and roots multiply to $1/\beta$, f has one root $\rho \in (0, \chi)$ and one other root above 1. The value gap can jump at $t = 1$, thus the only non-explosive path is $dS_{a,t} = (1 - \rho^t) dS_{a,\infty}$. Therefore, each period, a fraction $1 - \rho$ of the remaining adjustment takes place.

As $t \rightarrow \infty$, the two equations give $dS_a = -[S_a(1 - S_a)/\bar{\nu}] d \log(w^U/w^R)$. Combining this with (a), and using $W + \bar{\nu}/[S_a(1 - S_a)] = \zeta/H + \zeta/M$, gives

$$dS_{a,\infty} = -\frac{\theta^m S_m M \hat{A}}{\zeta(M + H)} < 0.$$

From (a), every worker who exits agriculture raises S_m to first order by $dS_{m,t} = dS_{m,0} - \zeta(1 + \theta^m)S_m(1 - \rho^t) dS_{a,\infty}/H$. So $dS_{m,t}$ rises with $1 - \rho^t$, to $dS_{m,\infty} = \theta^m S_m \hat{A} (M - B)/(M + H)$. $\square$

Lemma 2. *Migration slows down in τ . This is because ρ and $\rho(1 + \chi)/\chi$ rise strictly in τ , and $\rho \rightarrow 1$ as $\tau \rightarrow \infty$.*

Proof. χ is an increasing function of τ that increases from 0 toward 1. Let $\omega \equiv \beta S_a(1 - S_a)W/\nu$. This term does not depend on τ . Let $z \equiv \chi/\rho > 1$. Then $f(\rho) = 0$ implies that z is the larger root of

$$g(z, \chi) \equiv (z - 1)(z - \beta\chi^2) - (1 - \chi^2)\omega z.$$

At this root,

$$z \frac{\partial g}{\partial z} = z^2 - \beta\chi^2 > 0 \quad \text{and} \quad (1 + \chi) \frac{\partial g}{\partial \chi} = \frac{2\chi(z - 1)(z - \beta)}{1 - \chi} > 0.$$

By the implicit function theorem, $\frac{d}{d\chi} \log \frac{z}{1 + \chi} = -\frac{z \partial g / \partial z + (1 + \chi) \partial g / \partial \chi}{(1 + \chi) z \partial g / \partial z} < 0$. So $\rho(1 + \chi)/\chi = (1 + \chi)/z$ increases in χ . The root $\rho = \frac{\chi}{1 + \chi} \cdot \frac{\rho(1 + \chi)}{\chi}$ also increases in χ . Finally, $g = 0$ and $z - \beta\chi^2 \geq (1 - \beta)z$ imply $(1 - \beta)(z - 1) \leq (1 - \chi^2)\omega \rightarrow 0$, so $\rho = \chi/z \rightarrow 1$. $\square$

C.3 Thresholds on the Agricultural Employment Share

Note that for a fixed λ^s and w^R/w^U , the farm share pins down $S_s = \lambda^s(1 - S_a + S_a w^R/w^U)$ and $S_m = 1 - S_a - S_s$, and hence B and M_0 . For $0 \leq \sigma < 1$ define

$$\bar{S} \equiv \frac{(1 - \sigma)(1 - \lambda^s)}{(1 - \sigma)(1 - \lambda^s) + [\lambda^s + \sigma(1 - \lambda^s)]w^R/w^U} \in (0, 1).$$

Lemma 3. Fix $0 \leq \sigma < 1$ and the initial λ^s and w^R/w^U .

(a) $B > 0$ if and only if $S_a < \bar{S}$.

(b) There is a unique $\underline{S} \in (0, \bar{S})$ such that $M_0 > B$ if and only if $S_a > \underline{S}$.

Neither threshold depends on θ^m or on τ .

Proof. (a). Substituting S_s and S_m into B gives $B = \lambda^s(1 - \sigma)(1 - \lambda^s)(1 - S_a/\bar{S})$.

(b). By (a) and the definition of M_0 ,

$$\frac{M_0 - B}{1 - S_a} = \frac{\zeta S_a}{\nu/\beta + (1 - S_a)/(1 + \theta^a)} - \lambda^s \left[(1 - \sigma)(1 - \lambda^s) - [\lambda^s + \sigma(1 - \lambda^s)] \frac{w^R}{w^U} \frac{S_a}{1 - S_a} \right].$$

This expression is strictly increasing in S_a . It is negative as $S_a \rightarrow 0$ and positive at $S_a = \bar{S}$, at which point we have $B = 0$. Thus, the expression changes sign exactly once, at some $\underline{S} \in (0, \bar{S})$. Note that B and M_0 do not involve θ^m or τ . $\square$

C.4 Proof of the Closed-Form Proposition

Proof of Proposition 3. Use the thresholds in Lemma 3. Across comparable economies that differ in τ , note that only χ changes. Define $H = (1 + \theta^m)S_m - B$.

(a). If $S_a < \bar{S}$, then $B > 0$ and $dS_{m,0} < 0$ by Lemma 1 part (b).

(b). If $S_a > \underline{S}$, then $B < M_0 < M$. That is, the agricultural push beats the services pull at every value of τ . By Lemma 1 part (c), note that $dS_{m,t}$ rises strictly to a positive limit. Thus, $t^* \equiv \min\{t \geq 0 : dS_{m,t} \geq 0\}$ is finite and $dS_{m,t} < 0$ if and only if $t < t^*$.

If $S_a \geq \bar{S}$ then it follows that $B \leq 0$, and $t^* = 0$ at every τ . Now, let $S_a < \bar{S}$, so that $t^* \geq 1$ by (a). Higher τ implies that there are two opposing effects on t^* :

- migration is slower (ρ rises), which delays industrialization.
- the eventual push is larger (M rises), which speeds it up.

We now show that the first force always dominates. Since $dS_{m,t} = dS_{m,\infty} - \rho^t(dS_{m,\infty} - dS_{m,0})$, we have $t^* > k$ if and only if $\rho^k > \Psi$, where we define

$$\Psi \equiv \frac{dS_{m,\infty}}{dS_{m,\infty} - dS_{m,0}} = \frac{H}{H + B} \left(1 - \frac{B}{M} \right) < \frac{H}{H + B} < 1$$

as the share of the gap still left to close when $S_m = S_{m,0}$. By (61), we have that $1/M = 1/M_0 - q\chi/(1 + \chi)$ where $q \equiv (1 + \beta)\nu/[\beta\zeta S_a(1 - S_a)] > 0$. Thus,

$$\frac{\rho}{\Psi} = \frac{H + B}{H} \cdot \frac{\rho(1 + \chi)/\chi}{(1 - B/M_0)(1 + \chi)/\chi + qB}.$$

By the argument in Lemma 2, we know that the numerator is increasing in τ . Since $0 < B < M_0$, the denominator is decreasing. So for each $k \geq 1$, $\rho^k/\Psi = \rho^{k-1} \cdot \rho/\Psi$ rises in τ , and t^* is weakly increasing. Since $\rho \rightarrow 1$ while $\Psi < H/(H + B) < 1$ for all τ , for each k it follows that $\rho^k > \Psi$ for τ large enough, so $t^* \rightarrow \infty$. $\square$

D Manufacturing Peaks and Sectoral TFP Growth

Figure 4 of the main text measures productivity growth in terms of labor productivity growth as measured by GGDC's measure of real value added growth per worker. This appendix repeats that exercise but replaces labor productivity with inverted TFP for all 43 countries over 1965–2014. This adjusts labor productivity by the domestic expenditure share, such that:

$$A_t^{nj} = A_{t_0}^{nj} \cdot \frac{LP_t^{nj}}{LP_{t_0}^{nj}} \cdot (\pi_{nn,t}^{nj})^{1/(\gamma_t^{nj}\theta^j)} \cdot (\pi_{nn,t_0}^{nj})^{-1/(\gamma_{t_0}^{nj}\theta^j)} \cdot \Gamma_{EK}^{1/\gamma_t^{nj}-1/\gamma_{t_0}^{nj}}.$$

All four slopes in the data are positive. Slopes are steeper than labor productivity counterparts (Figure D.2). The slope β moves from +3.10 to +3.12 in (A), +7.43 to +7.66 in (B), +2.98 to +3.11 in (C) and +3.78 to +4.07 in (D), while ρ rises from +0.53 to +0.55 in (C) and rises from +0.45 to +0.50 in (D).




FIGURE D.1: PEAK MANUFACTURING SHARES AND MANUFACTURING TFP GROWTH

Notes: Four Moments. Panels, left to right and top to bottom: (A) Raw g_m ; (B) $g_m - \bar{g}$; (C) $g_m - g_a$ (Tradable–Tradable Margin); (D) $g_m - g_s$ (= HR's g). All Four Slopes Are Positive. Sample: 43 Countries, 1965–2014. As Figure 4, with sectoral TFP in place of labor productivity. TFP is the unsmoothed series anchored at 2005. Growth rates are annualized log changes over each country's first and last observed year within 1965–2014, and the sample is all 43 countries.




FIGURE D.2: PEAK MANUFACTURING SHARES AND MANUFACTURING LABOR PRODUCTIVITY GROWTH

Notes: As Figure D.1, with labor productivity in place of TFP. Panels, left to right and top to bottom: (A) raw g_m ; (B) $g_m - \bar{g}$; (C) $g_m - g_a$; (D) $g_m - g_s$. Sample: 43 countries, 1965–2014.

Figure D.3 is a 43-country, 1965-2014 version of Figure 4 but measures the peak as the unsmoothed annual maximum of the manufacturing employment share, rather than the smoothed five-year moving average. The slope β is +3.15 in (A), +7.55 in (B), +3.03 in (C), and +3.83 in (D). All four are positive and statistically significant.





FIGURE D.3: UNSMOOTHED PEAK MANUFACTURING SHARES AND MANUFACTURING PRODUCTIVITY GROWTH


Notes: A 43 country version of Figure 4, with the peak now measured as the unsmoothed maximum of manufacturing employment share. Panels, left to right and top to bottom: (A) raw g_m ; (B) $g_m - \bar{g}$; (C) $g_m - g_a$; (D) $g_m - g_s$. Sample: 43 countries, 1965-2014.

Figures D.4 and D.5 show our 43 countries over the same time period as Figure 4. In Figure D.4 the model of Huneus and Rogerson (2024) is started from each country's observed initial shares. Model generated slopes are positive but flatter than in the data. In both model and data, the positive slopes mostly reflect where countries start. A country's peak correlates at 0.90 with its initial manufacturing employment share in the data and at 0.87 in the model. This means that countries that were more industrial at the start subsequently


had faster relative manufacturing productivity growth. In Figure D.5 every country starts from the HR model's common benchmark initial state. This is as HR Proposition 2. The model is run until each country's manufacturing share peaks. In this case the model slopes are negative.



A: Relative to overall productivity growth



B: Relative to agricultural productivity growth




C: Relative to services productivity growth


—◇— Model: HR (2024) ● Data

FIGURE D.4: PEAK MANUFACTURING SHARES AND MANUFACTURING PRODUCTIVITY GROWTH, ALL COUNTRIES


Notes: Same exercise as Figure 4 for all 43 countries.



A: Relative to overall productivity growth



B: Relative to agricultural productivity growth



C: Relative to services productivity growth

—◇— Model: HR (2024) • Data

FIGURE D.5: MANUFACTURING PEAK SHARES AND PRODUCTIVITY GROWTH, BENCHMARK MODEL INITIAL CONDITION

Notes: Data as in Figure D.4. Orange diamonds are the peaks of the Huneus and Rogerson (2024) model started from its benchmark initial state in that paper, instead of observed starting employment shares. The 13 countries with $g_m < g_s$ and $g_a > 0$ have no model peak and are omitted from the model fit.

E Frictionless Benchmark with Recalibrated Preferences

The frictionless benchmark of Section 5.2 uses the preference parameters of the spatial model. These parameters do not give the frictionless economy its best chance. This appendix recalibrates preferences for the frictionless economy alone to improve its fit to the data. We keep fundamentals, trade elasticities and trade balances at baseline values. The household elasticity of substitution between sectors and Stone-Geary terms are re-chosen. Common preference weights α_j are re-solved, as in the baseline, so that the model matches average 1970 labor shares. Base-year expenditure shares, wage bills and trade flows are recomputed. We choose the parameters to minimize the sum over 43 countries and 1970–2014 of squared deviations of model from data manufacturing employment shares plus 0.3 times the same sum for services. This is the rough criterion used for the baseline calibration. This search is over a grid of $(\sigma, \bar{c}_a, \bar{c}_s)$.

We report results in Table E.1. The best frictionless fit pushes σ to the lower bound of the search, 0.05. Sectors become almost Leontief complements in this range. Preference weights on manufacturing fall from 0.33 to 0.20. Frictionless model mean square loss falls from 89.1 to 82.9 against 36.1 for the spatial baseline.

Figure E.1 plots the recalibrated peaks against the data. The median developing economy still over-industrializes by 8.9 percentage points. This is relative to the 10.3 percentage points in the original calibration. The newly calibrated preferences lower manufacturing almost everywhere and not just developing economies. The mean advanced-economy peak falls from 0.31 to 0.27 relative to 0.39 in the data. The correlation between model and data peaks also falls from 0.04 to 0.03. The slope of the peak share on log GDP per worker in 1990 remains oppositely signed to the data, and appears in Figure E.2. Figure E.3 shows how the time paths change for Brazil, India, and China under the new calibration. Brazil’s frictionless performance is much better but China still over-industrializes substantially relative to data and the spatial model. India over-industrializes only until the mid-1990s and then falls below the data. Overall, we conclude that the over-industrialization of emerging markets in the frictionless benchmark does not appear to be due to the spatial model’s calibration.

TABLE E.1: FRICTIONLESS BENCHMARK WITH RECALIBRATED PREFERENCES

	Baseline	Frictionless	Recalibrated frictionless
<i>Preferences</i>			
Elasticity of substitution, σ	0.56	0.56	0.05
Preference weights ($\alpha_a, \alpha_m, \alpha_s$)	(0.07, 0.33, 0.60)	(0.07, 0.33, 0.60)	(0.18, 0.20, 0.62)
Median subsistence exp. share s_a , 1970	0.28	0.28	0.10
Median luxury exp. share $ s_s $, 1970	0.14	0.14	0.37
<i>Quantitative Fit</i>			
Sector share target	36.1	89.1	82.9
Correlation with data peaks	0.59	0.04	0.03
Model peak vs. data peak OLS slope	0.66	0.05	0.03
RMSE (pp)	10.7	17.7	17.4
Model – data, developing median (pp)	–0.8	+10.3	+8.9
Mean developing peak (data 0.24)	0.26	0.38	0.35
Mean advanced peak (data 0.39)	0.31	0.31	0.27
Median peak year (data 1986)	1989	1973	1973
Peak vs. 1990 log GDP per worker slope (data +0.057)	+0.026	–0.021	–0.028

Notes: Baseline and Frictionless calibrations as in Table 2 and Figure 9. The recalibrated frictionless economy re-chooses σ , $\bar{c}_a$, and $\bar{c}_s$, with common α_j re-solved to match global 1970 sectoral labor shares. Peak statistics cover the 43 countries minus HK and SGP from 1970-2013.




FIGURE E.1: PEAK MANUFACTURING EMPLOYMENT SHARE, FRICTIONLESS BENCHMARK WITH RECALIBRATED PREFERENCES

Notes: As Figure 8, with the right panel (No frictions) showing the frictionless economy under the recalibrated preferences of Table E.1. The left panel (Baseline) is unchanged. Recalibration seeks to maximize fit of frictionless model.




FIGURE E.2: PEAK MANUFACTURING SHARE VS. GDP PER WORKER IN 1990, FRICTIONLESS BENCHMARK WITH RECALIBRATED PREFERENCES

Notes: As Figure 9, but with the middle panel demonstrating the recalibrated frictionless economy from Table E.1.




FIGURE E.3: MANUFACTURING EMPLOYMENT SHARE PATHS IN BRAZIL, INDIA AND CHINA FOR RECALIBRATED FRICTIONLESS BENCHMARK

Notes: Manufacturing employment share paths, 1971–2013 in data (black), baseline spatial model (blue) and frictionless economy with the recalibrated preferences of Table E.1 (red).

F Empirical Appendix

F.1 Definition of Variables

- (i) **Sectoral employment shares.** The share of an AMC's employed persons working in agriculture, manufacturing, services, or other sectors, from the 2000 and 2010 population censuses. Outcomes are the 2000-2010 changes in these shares.
- (ii) **Manufacturing employment growth.** The change in the log number of an AMC's employed people working in manufacturing, from the 1991, 2000, and 2010 censuses. This is the outcome of the pre-trend test, as in BCP, since employment shares cannot be compared across the 1991 and 2000 waves. The 1991 and 2000 censuses define employment differently, which affects total employment but barely affects manufacturing employment.
- (iii) **Soy and maize exposure.** The difference in potential yields of the crops under high versus low input farming, measured in tonnes per hectare. We take this variable from FAO-GAEZ v3.0 and average over the AMC's area.
- (iv) **Baseline controls.** The 1991 share of the population living in rural areas, log income per capita, log population density, and the literacy rate. The pre-trend test, which additionally uses the 1991-2000 changes, uses the 1980 values of these control variables as well.
- (v) **Development controls.** State fixed effects and the 1991 urbanization rate, log population density, literacy rate, and log income per capita, which we demean across AMCs and interact with our soy and maize exposure measures.
- (vi) **Youth share.** The share of an AMC's 1991 population aged 18-25. The young-adult share (ages 18-35) and the prime-age share (ages 15-55) are defined analogously.
- (vii) **Urban-directed out-migration.** The share of an AMC's 1995 residents who by 2000 lived in a different, urban AMC.

TABLE F.1: SUMMARY STATISTICS

Variable	Mean	SD	p10	p50	p90	N
<i>Panel A: Outcomes (2000 to 2010 changes in employment shares)</i>						
Δ Manufacturing share	0.013	0.057	-0.036	0.004	0.078	4,260
Δ Agriculture share	-0.063	0.074	-0.158	-0.057	0.015	4,260
Δ Services share	0.032	0.057	-0.035	0.031	0.103	4,260
Δ Other-sector share	0.018	0.039	-0.024	0.015	0.064	4,260
<i>Panel B: Agricultural exposures</i>						
ΔA^{soy}	1.806	0.847	0.628	1.837	2.882	4,260
ΔA^{maize}	3.055	1.808	0.778	2.727	5.753	4,260
<i>Panel C: 1991 baseline controls</i>						
Rural share	0.458	0.231	0.133	0.470	0.761	4,260
Log pop. density	3.202	1.358	1.588	3.232	4.672	4,260
Literacy rate	0.717	0.164	0.476	0.773	0.899	4,260
Log income per capita	4.525	0.588	3.746	4.556	5.281	4,260
<i>Panel D: Moderators</i>						
Youth share, 1991	0.144	0.015	0.125	0.144	0.162	4,260
Young-adult share, 1991	0.286	0.033	0.240	0.290	0.327	4,260
Prime-age share, 1991	0.530	0.047	0.465	0.539	0.586	4,260
Urban outmigration, 1995 to 2000	0.231	0.092	0.122	0.223	0.351	4,260

Notes: This table presents summary statistics of our data. Panel A reports the 2000-2010 changes in sectoral employment shares, the outcome variables of equation (20). Panel B reports our soy and maize exposure measures, constructed as the difference between FAO-GAEZ potential yields under high and low inputs, in tonnes per hectare. Panel C reports the baseline controls, measured in the 1991 population census. Panel D reports the four measures of migration propensity. “Youth share” is the share of an AMC’s 1991 population aged 18 to 25; the prime-age (ages 15 to 55) and the young-adult (ages 18 to 35) shares are defined analogously. “Urban outmigration” is the share of an AMC’s 1995 residents who by 2000 lived in a different, urban AMC. The unit of observation is the AMC.

F.2 Additional Empirical Results and Robustness

TABLE F.2: MIGRATION PROPENSITY GRADIENT IN THE SOY EFFECT: ALL SECTORS

	Δ Employment share							
	Agriculture		Manufacturing		Services		Other Sectors	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔA^{soy}	-0.020*** (0.002)	0.006 (0.008)	0.021*** (0.002)	0.003 (0.006)	-0.002 (0.002)	-0.012* (0.006)	0.001 (0.001)	0.002 (0.004)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 2}$		-0.020** (0.009)		0.014** (0.007)		0.007 (0.007)		> -0.001 (0.005)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 3}$		-0.028*** (0.010)		0.014* (0.008)		0.010 (0.007)		0.004 (0.004)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 4}$		-0.025*** (0.009)		0.017** (0.007)		0.009 (0.007)		-0.001 (0.005)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 5}$		-0.035*** (0.010)		0.025*** (0.008)		0.016** (0.007)		-0.006 (0.005)
Joint test, soy coefficients on quintiles 2-5 equal 0:								
<i>p</i> -value		0.009		0.028		0.163		0.050
Observations	4,149	4,260	4,149	4,260	4,149	4,260	4,149	4,260

Notes: The table reports estimates from equation (20) for each sector. The odd-numbered columns restrict all migration propensity terms to zero, which replicates the BCP specification (Table 10). The even-numbered columns estimate the unrestricted equation. The outcome variable for each column is the 2000-2010 change in the employment share of the respective sector. Migration propensity is measured as the 1991 youth share (ages 18 to 25), and its lowest quintile is the omitted category. Baseline controls are the 1991 rural share, log population density, literacy rate, and log income per capita. The odd-numbered columns restrict the sample to the AMCs used in BCP. Column (4) corresponds to column (2) of Table 4. The row labeled “p-value” reports the p-value from an F-test of the null hypothesis that the soy exposure $\times$ quintile coefficients are jointly zero. Standard errors are in parentheses, clustered at the microregion level in the even-numbered columns and robust in the odd-numbered columns, as in BCP. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

TABLE F.3: MIGRATION PROPENSITY GRADIENT IN THE SOY EFFECT ON
MANUFACTURING: ALTERNATIVE PROXIES

	Δ Manufacturing employment share			
	Youth share (1)	Young- adult share (2)	Prime- age share (3)	Urban outmig. (4)
ΔA^{soy}	0.003 (0.006)	-0.011 (0.008)	-0.005 (0.007)	< 0.001 (0.005)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 2}$	0.014** (0.007)	0.026*** (0.009)	0.021*** (0.008)	0.019*** (0.007)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 3}$	0.014* (0.008)	0.039*** (0.010)	0.028*** (0.008)	0.031*** (0.007)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 4}$	0.017** (0.007)	0.033*** (0.010)	0.033*** (0.009)	0.023*** (0.007)
$\Delta A^{\text{soy}} \times \text{Mig. Prop. Quintile 5}$	0.025*** (0.008)	0.035*** (0.010)	0.024*** (0.009)	0.020*** (0.006)
Joint test, soy coefficients on quintiles 2-5 equal 0:				
p -value	0.028	0.001	0.006	< 0.001
Observations	4,260	4,260	4,260	4,260

Notes: This table explores the robustness of our results to alternative definitions of the migration propensity moderator. Each column estimates the unrestricted equation (20) with migration propensity measured as indicated in the heading of the column. The outcome variable for each column is the 2000-2010 change in the manufacturing employment share. “Youth share” is the share of an AMC’s 1991 population aged 18 to 25; the young-adult share (ages 18 to 35) and the prime-age share (ages 15 to 55) are defined analogously. “Urban outmig.” is urban-directed out-migration, the share of an AMC’s 1995 residents who by 2000 lived in a different, urban AMC. In each column, the lowest quintile of migration propensity is the omitted category. Baseline controls are the 1991 rural share, log population density, literacy rate, and log income per capita. Column (1) corresponds to column (2) of Table 4. The row labeled “p-value” reports the p-value from an F-test of the null hypothesis that the soy exposure $\times$ quintile coefficients are jointly zero. Standard errors are in parentheses and clustered at the microregion level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

TABLE F.4: MIGRATION PROPENSITY GRADIENT IN THE SOY EFFECT ON
MANUFACTURING: CONTINUOUS MIGRATION PROPENSITY

	Δ Manufacturing employment share			
	(1)	(2)	(3)	(4)
ΔA^{soy}	0.021*** (0.002)	-0.003 (0.009)	-0.008 (0.010)	-0.005 (0.009)
$\Delta A^{\text{soy}} \times \text{Mig. Prop.}$		0.049** (0.021)	0.059*** (0.021)	0.055** (0.022)
Test, soy coefficient on Mig. Prop. equal 0:				
p -value		0.022	0.006	0.011
Observations	4,149	4,260	4,260	4,149
Fixed Effects and Controls:				
Baseline controls	✓	✓	✓	✓
Development controls	-	-	✓	-
BCP sample	✓	-	-	✓

Notes: This table explores the robustness of our results to a continuous definition of the migration propensity moderator. The table reports estimates from equation (20), where we replace the migration propensity quintiles with a continuous version, normalized to the unit interval in our sample. Column (1) restricts the migration propensity term to zero, which replicates the BCP specification (Table 9, column 2). Columns (2)–(4) estimate the unrestricted equation. The outcome variable for each column is the 2000-2010 change in the manufacturing employment share. Migration propensity is measured as the 1991 youth share (ages 18 to 25). Baseline controls are the 1991 rural share, log population density, literacy rate, and log income per capita. Column (3) adds development controls, specifically state fixed effects and the 1991 urbanization rate, log population density, literacy rate, and log income per capita, demeaned and interacted with our soy and maize exposure measures. Column (4) restricts the sample to the AMCs used in BCP. The row labeled “p-value” reports the p-value from a test of the null hypothesis that the soy exposure $\times$ migration propensity coefficient is zero. Standard errors are in parentheses, clustered at the microregion level in columns (2)–(4) and robust in column (1), as in BCP. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

TABLE F.5: MIGRATION PROPENSITY GRADIENT IN MANUFACTURING EMPLOYMENT GROWTH: BEFORE AND AFTER THE SOY SHOCK

	Δ Log manufacturing employment			
	(1)	(2)	(3)	(4)
Soy exposure, 1991 to 2000				
All AMCs	< 0.001 (0.020)			
Total effect on:				
Mig. Prop. Quintile 1		0.146* (0.087)	0.102 (0.084)	0.197** (0.086)
Mig. Prop. Quintile 2		0.053 (0.053)	0.064 (0.054)	0.055 (0.053)
Mig. Prop. Quintile 3		-0.041 (0.045)	-0.022 (0.050)	-0.023 (0.042)
Mig. Prop. Quintile 4		-0.019 (0.045)	-0.008 (0.046)	-0.007 (0.044)
Mig. Prop. Quintile 5		0.003 (0.045)	-0.004 (0.052)	0.007 (0.043)
Joint test, pre-period soy slopes equal 0:				
<i>p</i> -value		0.403	0.637	0.222
Soy exposure, 2000 to 2010				
All AMCs	0.238*** (0.018)			
Total effect on:				
Mig. Prop. Quintile 1		0.073 (0.078)	0.017 (0.075)	0.052 (0.071)
Mig. Prop. Quintile 2		0.227*** (0.052)	0.155*** (0.052)	0.240*** (0.052)
Mig. Prop. Quintile 3		0.181*** (0.041)	0.123*** (0.042)	0.174*** (0.040)
Mig. Prop. Quintile 4		0.232*** (0.038)	0.190*** (0.043)	0.215*** (0.037)
Mig. Prop. Quintile 5		0.233*** (0.040)	0.228*** (0.043)	0.233*** (0.040)
Joint test, post-period soy slopes equal 0:				
<i>p</i> -value		< 0.001	< 0.001	< 0.001
Observations	7,984	8,440	8,440	7,984
Fixed Effects and Controls:				
Baseline controls	✓	✓	✓	✓
Development controls	-	-	✓	-
BCP pre sample	✓	-	-	✓

Notes: This table explores whether the migration propensity gradient was present before the soy shock. The table reports estimates from equation (20) stacked over the periods 1991-2000 and 2000-2010, where all exposure and migration propensity terms are interacted with an indicator for the post-period, following BCP (Online Appendix Table A6). Column (1) restricts all migration propensity terms to zero, which replicates the BCP specification (column 1). Columns (2)–(4) estimate the unrestricted equation. The outcome variable for each column is the change in log manufacturing employment over the period. The table reports the total effect of soy exposure in each period, for all AMCs in column (1) and for each quintile of

migration propensity in columns (2)-(4). Migration propensity is measured as the 1991 youth share (ages 18 to 25). Baseline controls are the rural share, log population density, literacy rate, and log income per capita, measured in 1980 for the first period and in 1991 for the second. Column (3) adds development controls, as in Table 4, interacted with the indicator for the second period. Column (4) restricts the sample to the AMCs used in BCP. The rows labeled “p-value” report the p-value from an F-test of the null hypothesis that the five total effects of the period are jointly zero. Standard errors are in parentheses, clustered at the microregion level in columns (2)-(4) and robust in column (1), as in BCP. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Appendix References

- FINICELLI, ANDREA, PAGANO, PATRIZIO, & SBRACIA, MASSIMO. 2013. Ricardian Selection. *Journal of International Economics*, **89**(1), 96–109.