The Dynamics of Agricultural Productivity Gaps: An Open-Economy Perspective*

Douglas Gollin
Tufts University and University of Oxford

David Lagakos
Boston University and NBER

Xiao Ma
Peking University

Shraddha Mandi
Boston University

June 4, 2026

Abstract

This paper draws on cross-country census data to study how agricultural productivity gaps have evolved over the last four decades. We find little tendency for gaps to decline on average despite global decreases in agricultural employment shares. We analyze the dynamics of agricultural productivity gaps through the lens of an open-economy model of structural change. We calibrate the model using international trade data, which help discipline changes in comparative advantage over time. Quantitatively, the model predicts that relatively faster physical productivity growth in the non-agricultural sector has, in many countries, offset the movement of labor out of agriculture, leading to persistently lower value added per worker in agriculture than in the non-agricultural sector.

*Email: douglas.gollin@tufts.edu, lagakos@bu.edu, xiaoma@phbs.pku.edu.cn, mandis@bu.edu.

For helpful feedback we thank Tasso Adamopoulos, Chaoran Chen, Raahil Madhok, Michael Peters, Xincheng Qiu, Natalia Ramondo, Mike Waugh, Shengxing Zhang, Xiaodong Zhu, and seminar/conference participants at Edinburgh, Oxford, the Minneapolis Fed Growth and Development Conference 2025, the Northeast Universities Development Conference 2025 (NEUDC), Peking University HSBC Business School, Peking University National School of Development, Shanghai Jiao Tong University, Jinan University, Zhejiang University of Finance and Economics, Xiamen University, Shanghai University of Finance and Economics, and the World Bank Development Economics Research Group. We thank Katia Del Frari for valuable inputs on a previous version of this paper. We thank Isaac Otoo for excellent research assistance. Ma acknowledges support from the National Natural Science Foundation of China (Grant No. 72522018, 72203012) and the Shenzhen Science and Technology Program (Grant No. JCYJ20250604175918024). The authors are listed in alphabetical order. All potential errors are our own.

Introduction

In nearly every country in the world, value added per worker is substantially higher in non-agricultural activities than in agriculture. The ratio of the two, known as the *agricultural productivity gap*, has been the focus of an active literature over the last decade. A central issue is whether these gaps represent labor misallocation; if so, there may be gains in efficiency as workers transition from agriculture to more productive sectors (Gollin, Lagakos, and Waugh, 2014). A number of studies embrace this view and assign agricultural productivity gaps a prominent role in theories of structural transformation and development (e.g. Bryan and Morten, 2019; Diao, McMillan, and Rodrik, 2019; Storesletten, Zhao, and Zilibotti, 2019; Gai, Guo, Li, Shi, and Zhu, 2025). Others are more skeptical, pointing to smaller gaps in wages than in average productivity (e.g. Vollrath, 2014; Herrendorf and Schoellman, 2015), modest income gains for workers who switch sectors (Hamory, Kleemans, Li, and Miguel, 2021; Alvarez, 2020), and evidence of selection on unobservable characteristics (Young, 2013; Herrendorf and Schoellman, 2018).

The goal of this paper is to make progress in understanding agricultural productivity gaps by studying their dynamics. To do so, we first construct a new panel data set that draws on national account data and population censuses from 68 countries that span four decades. Analyzing these data, we find that, on average, agricultural productivity gaps have declined little, if at all, since the 1980s. Yet individual countries have experienced markedly different patterns over this time period. Data from China and India, for example, show significantly widening gaps over time; for others, such as Brazil, the gaps have closed steadily.

To understand the dynamics of agricultural productivity gaps in our panel of countries, we turn to a dynamic model of structural change with many countries. The model focuses on two broad explanations. The first is the tendency for workers to move from agriculture to non-agricultural activities in the course of economic development (see Herrendorf, Rogerson, and Valentinyi, 2013). Sectoral reallocation helps close agricultural productivity gaps, as workers leave agriculture for manufacturing and service activities where their labor is more valuable. The second is differences in real productivity growth rates across sectors. To the extent that real labor productivity growth is faster in manufacturing and services than in agriculture, agricultural productivity gaps will widen over time, all else equal.

A central premise of our analysis is that an open-economy perspective is useful in understanding the dynamics of agricultural productivity gaps. International trade in agricultural and manufactured goods is significant in practice and has grown substantially as a share of total value added over this period. Our analysis draws on international trade data to help make inferences about the real patterns of comparative advantage that underlie trade by sector. This allows us to calibrate our model to data that are independent of our measurements of the agricultural productivity gaps that we are trying to explain.

We incorporate and quantify these features in a model that builds on the foundational studies of structural change in open economies (e.g. Stokey, 2001; Matsuyama, 2009; Uy, Yi, and Zhang, 2013; Tombe, 2015; Levchenko and Zhang, 2016; Teignier, 2018). Households have Armington preferences over varieties of agricultural and non-agricultural goods, and international trade is subject to asymmetric sector-specific iceberg trade costs. We adopt PIGL preferences to capture income effects in which agriculture's expenditure share falls as income rises; this allows the model to match expenditure shares by sector over a wide range of income levels in a flexible way (Boppart, 2014; Eckert and Peters, 2022). The model allows the physical productivity of labor to grow at different rates in each country, sector, and year.

We parameterize the model and use it to quantify the importance of sectoral reallocation forces and differential real productivity growth in explaining the dynamics of agricultural productivity gaps. We calibrate trade costs to match estimates from gravity equations by sector, and we allow costs to depend on exporters' development levels following Waugh (2010). We target each country's physical productivity levels by sector and year to match fixed effect estimates in gravity equations similar to Eaton and Kortum (2002). We assume frictional movement of labor but abstract from the specific sources of migration costs. In particular, we assume that a randomly selected 2 percent of the population can change sectors each period, while the remainder cannot. This parsimonious choice is consistent with the average decline in employment shares in agriculture that we observe across all of our countries.

Despite its simplicity, the quantitative model matches various non-targeted country-level moments in the data reasonably well. The initial levels of the agricultural productivity gaps are highly correlated with their counterparts in the data, with a correlation coefficient of 0.64. The main test of the model is whether its predicted *changes* in agricultural productivity gaps—which are not targeted in the calibration—are comparable to those in the data. Comparing the first and last years for each country, we find that the correlation between log changes in gaps in the model and data is 0.80. The correlations are also strong when conditioning on the productivity gap increasing or decreasing over the period, signaling that the model gets the symmetry in the data about right. In both the model and the data, the average country has a negligible change in its productivity gap despite substantial changes in sectoral employment.

To better understand the underlying drivers of the model's predicted agricultural productivity gap dynamics, we conduct a series of counterfactual experiments. The first holds the allocation of labor across sectors constant, so as to focus on the effects of physical productivity growth by sector. This counterfactual produces a reasonably high correlation with the observed changes in productivity gaps. But the productivity gaps are shifted upward substantially: the average change is now far too high in almost every country (27 log points in the model vs. -7 log points in the data). The lesson is that heterogeneous real productivity growth rates across sectors are crucial for understanding the variation in agricultural productivity gap dynamics. However, all else equal, real sectoral productivity growth by itself would have raised gaps substantially rather than keeping them constant or reducing them.

In a second counterfactual, we assess the role of worker sectoral reallocation by relaxing the fixed employment assumption of the first experiment while also holding sectoral productivity levels constant. In this scenario, the average country shows a substantial decline in its agricultural productivity gap, matching our initial intuition that movements out of agriculture should close sectoral gaps. However, this contradicts the data, as gaps show no tendency to decline in our set of countries. Furthermore, this counterfactual explains little of the variation in gap dynamics between countries, with only a 0.08 correlation between model and data. The reason is that workers move out of agriculture at similar rates across countries, generating declines of comparable magnitude. This uniform pattern cannot match the observed heterogeneity in the data, in which some countries experience rising agricultural productivity gaps while others see them falling.

A natural question is how the model's conclusions would have differed by using direct measures of real productivity by sector for calibration purposes. The main challenge here is the absence of time-series data on sectoral prices for many countries. The GGDC Productivity Level Database reports sectoral price data for only 2005, 2011, and 2017 for example. Herrendorf, Rogerson, and Valentinyi (2026) estimate real productivity levels by sector for a smaller set of countries and years. We find that when calibrating our model directly to these data, the model's predicted changes in gaps for this set of countries are still largely in line with their empirical counterparts. Though the *levels* of the gaps match the data significantly worse. Brazil, China, and India have gaps that are around three times as large as the data in the 1990s, for example.

To help validate the mechanisms at work in our model, we conclude by examining the raw correlations between observed APG dynamics and country characteristics suggested by our theory. Consistent with the model's mechanisms, growth in agricultural exports is associated with a decline in the gap, whereas non-agricultural export growth is linked to an expansion of the gap. Increases in GDP and non-agricultural expenditure shares are associated with larger APGs; increases in non-agricultural employment are associated with smaller ones. For a smaller set of countries, direct measures of non-agricultural (agricultural) sector output-per-worker growth are associated with rising (falling) gaps.

Overall, our findings support the view that agricultural productivity gaps reflect differences in physical productivity levels and growth rates across countries, combined with frictions that limit labor mobility across sectors, rather than faulty national-account data, as emphasized by e,g, Jerven (2013). An implication of our study is that, because countries experience differential growth rates of physical productivity across sectors, agricultural productivity gaps need not disappear over time. So long as worker reallocation across sectors responds slowly and imperfectly to sectoral income gaps, these gaps may remain as permanent fixtures rather than petering out as countries grow.

Our study contributes to a large literature in macroeconomics studying the role of agriculture in structural change and development (e.g. Gollin, Parente, and Rogerson, 2002; Caselli, 2005; Restuccia, Yang, and Zhu, 2008). Our work is particularly related to the more recent studies emphasizing the importance of open-economy issues for structural change (Alessandria, Johnson, and Yi, 2023; Tombe, 2015; Betts, Giri, and Verma, 2017; Lewis, Monarch, Sposi, and Zhang, 2022; Farrokhi and Pellegrina, 2023). By employing the gravity equation to estimate productivity levels, we follow a longstanding literature that infers productivity from trade patterns (e.g., MacDougall, 1951; Balassa, 1963; Golub and Hsieh, 2000; Eaton and Kortum, 2002; Waugh, 2010; Malmberg, 2025). Our paper takes real productivity levels and their growth rates as exogenous, though the literature has emphasized variation in farm size, institutions, technology diffusion, intermediate input use, and capital inputs as key proximate causes of low agricultural productivity (e.g. Adamopoulos and Restuccia, 2022, 2014; Donovan, 2021; Moscona and Sastry, 2025; Boppart, Kiernan, Krusell, and Malmberg, 2023).

As noted above, we also abstract from any detailed modeling of the causes and consequences of barriers to rural-urban migration in developing economies. This topic has given rise to a great deal of recent literature (Bryan and Morten, 2019; Lagakos, Marshall, Mobarak, Vernot, and Waugh, 2020; Donovan and Schoellman, 2023). Relatively few countries have explicit policy barriers to internal migration, with China being the principal exception through its Hukou policy (see e.g. Lu and Xia, 2016; Fan, 2019; Ngai, Pissarides, and Wang, 2019; Tombe and Zhu, 2019; Liu and Ma, 2023). Yet migration costs appear to be high, and even the costs of switching sectors within locations may be high (Baysan, Dar, Emerick, Li, and Sadoulet, 2024). Limited information among potential migrants can act as a barrier to rural-urban migration (Baseler, 2023), as can frictions in land or labor markets (de Janvry, Emerick, Gonzalez-Navarro, and Sadoulet, 2015; Chen, 2017; Gottlieb and Grobovšek, 2019; Adamopoulos, Brandt, Leight, and Restuccia, 2022; Adamopoulos, Brandt, Chen, Restuccia, and Wei, 2024; Marshall, 2025). Migration costs may be highly heterogeneous, and they may vary systematically by age and gender; Cao, Chen, Xi, and Zuo (2024) show that in China, women face particularly large migration costs. See Lagakos (2020), Lucas (2021), and McKenzie (2024) for recent summaries of the literature on rural-urban migration.

I Data and Motivating Facts

Our paper makes use of a new data set that we constructed, drawing on available data sources that allowed us to produce consistent and high-quality estimates of agricultural productivity gaps for a large set of countries. We measure these gaps following previous studies, but in contrast to earlier work, we focus on time-series patterns rather than cross-sectional analysis. The facts we document here motivate our model and the quantitative analysis that follows.

A Data

Following Gollin et al. (2014), we define the agricultural productivity gap (APG) as the ratio of value added per worker in non-agriculture to value added per worker in agriculture. Letting v_a be the share of value added in agriculture, and l_a be the share of employment in agriculture, the APG can be expressed as:

$$APG = \frac{(1-v_a)/(1-l_a)}{v_a/l_a}. \tag{1}$$

We calculate sectoral employment shares using data from population censuses and nationally representative labor force surveys, collected and made available by IPUMS International. These contain individual-level information on employment status and industry (Minnesota Population Center, 2020). We only use country-years for which we can compute sectoral employment shares ourselves from the underlying individual-level micro data, since different governments and international organizations follow inconsistent approaches in constructing aggregate variables. By working with the underlying micro data, we can produce measures of employment shares that are consistently calculated over time and that are (largely) comparable across countries, subject to some modest differences.

Sectoral value added comes from the United Nations National Accounts Main Aggregates Database: Basic Data Selection (United Nations Statistics Division, 2020). This source supplies data for many countries from 1970 onward. We use sectoral value added shares measured in current prices and national currency units. Since our measures of the productivity gaps are across sectors within countries at a moment in time, it makes sense to calculate these in current prices and national currency. GDP per capita comes from the Penn World Tables, and we focus on output-side real GDP at chained PPPs (Feenstra, Inklaar, and Timmer, 2015). We restrict our analysis to countries for which we can calculate the APG in at least two years, where the first and last observations are at least ten years apart. This allows us to capture medium- to long-run dynamics.

To construct a dataset of employment and industry that minimizes measurement error and enables cross-country and cross-time comparisons, we undertook extensive efforts to harmonize the data. While we restricted our scope to data available from IPUMS, each survey's documentation was carefully reviewed to ensure its inclusion supported consistency and comparability. For example, where the choice was available, we prioritized censuses over labor force surveys, as the latter often operate on a rolling basis, complicating comparability. Censuses, in contrast, are nationally representative and capture the working population at a specific point in time, providing a more reliable foundation for analysis. We also assessed the geographic representativeness of each census. Surveys that excluded economically significant areas or large portions of the population were deemed not nationally representative and excluded.

Surveys also differed in their definitions of the labor force and employment. To address these discrepancies, we limited our analysis to individuals aged 15 to 99 with known employment status and industry of employment. We included, as employed, those who supplied labor in paid, unpaid, and in-kind work arrangements. We excluded those who worked only in home services such as cooking and childcare. Although most surveys adhered to this definition, variations remained in their treatment of minimum work hours, subsistence work, and short-term unemployment. We included only surveys in which economic activity could be constructed consistently and comparably based on clearly defined employment and industry measures. The IPUMS data offer a harmonized variable for individuals' industry of employment. Although this is convenient to use, values were missing for some country-years. To address this, we imputed missing industry data using information on occupation and other employment variables, and by going back to the raw (unharmonized) census data.[1]

The final dataset encompasses 68 countries and 241 country-year observations and spans the period 1970 to 2017. See Appendix Table A.1 for a complete list of countries and years. Of the countries included, 19 percent have only two data points, while 57 percent have three. Eleven

[1] Industry was imputed only if more than 4 percent of the working population's industry data was missing and if the process could maintain consistency across surveys within a country. When harmonized industry variables were unavailable or contained substantial missing values, we followed a systematic approach to address these gaps. Original, non-harmonized industry variables were used whenever available. If these were unavailable, occupation variables were used to impute industry data. Surveys were excluded if neither industry nor occupation data could be used for imputation. Surveys with significant missing values—defined as exceeding 15 percent after imputation or having unusually high levels of missing data in specific years—were excluded from the analysis.

countries, including Brazil, India, Ireland, Mexico, and the United States, feature panels with more than five data points. On average, the dataset provides a panel spanning 27 years per country. Ireland, Mexico, and the United States contribute the longest series, with panels spanning 45 years. This selection of countries and years enables the analysis of both medium-term and long-run effects across developed and developing countries.

B Dynamics of Gaps: Empirical Facts

Figure 1 plots the change in the APG from the earliest to the latest available years in each country (*y*-axis) against the APG level in the earliest year (*x*-axis).[2] The size of the dot represents the population of the country. Many countries cluster near the 45-degree line, meaning similar gaps in the first and last years. Out of the 68 countries, 33 showed gaps changing by less than one unit - for instance, Turkey's gap moved marginally from 6.5 to 7. The remaining countries, including China, India, and Brazil, experienced larger shifts. Tanzania and Togo, for example, saw dramatic reductions: Tanzania's gap fell from nearly 12 to around 5, while Togo's decreased from around 9 to just under 2.

The fact that so many countries experienced minimal changes in their APGs seems puzzling given the global reallocation of labor out of agriculture in recent decades. Figure 2 plots APG changes against changes in agricultural employment share. All but two countries appear in the two left quadrants, meaning that they experienced declines in agriculture's share of employment (the exceptions are Lesotho, pictured in Q1, and Zambia, with a change in APG of 9.4). Although more countries fall into the bottom-left quadrant (Q3), where APGs declined, than into the top-left quadrant (Q2), where APGs increased, roughly one-third of countries still lie in Q2 with rising APGs despite declining agricultural employment shares. Notably, Q2 includes several highly populous countries—such as China and India—which together account for just over 80 percent of the global population.

The stark heterogeneity in APG dynamics present in some of the largest economies is reported in more detail in Figure 3. Brazil has the most expected patterns. Agriculture's share of value added and employment have both fallen since 1970, with the latter falling at a faster rate, reducing the APG from 6 to a little over 2. China and India are more surprising. Like Brazil, both countries have experienced declines in agriculture's share of employment and value added. However, the

[2]For expositional reasons, Figure 1 excludes the seven countries with the highest APG values: Burkina Faso (earliest: 28.3, latest: 19.2) , Cameroon (earliest: 14.5, latest: 8.9), Guinea (earliest: 21.4, latest: 4.7), Liberia (earliest: 15.2, latest: 0.5), Rwanda (earliest: 14.6, latest: 8.2), Thailand (earliest: 11.8, latest: 14.2) , and Zambia (earliest: 6.6, latest: 15.9).

Figure 1: Change in APGs: Earliest vs Latest Years

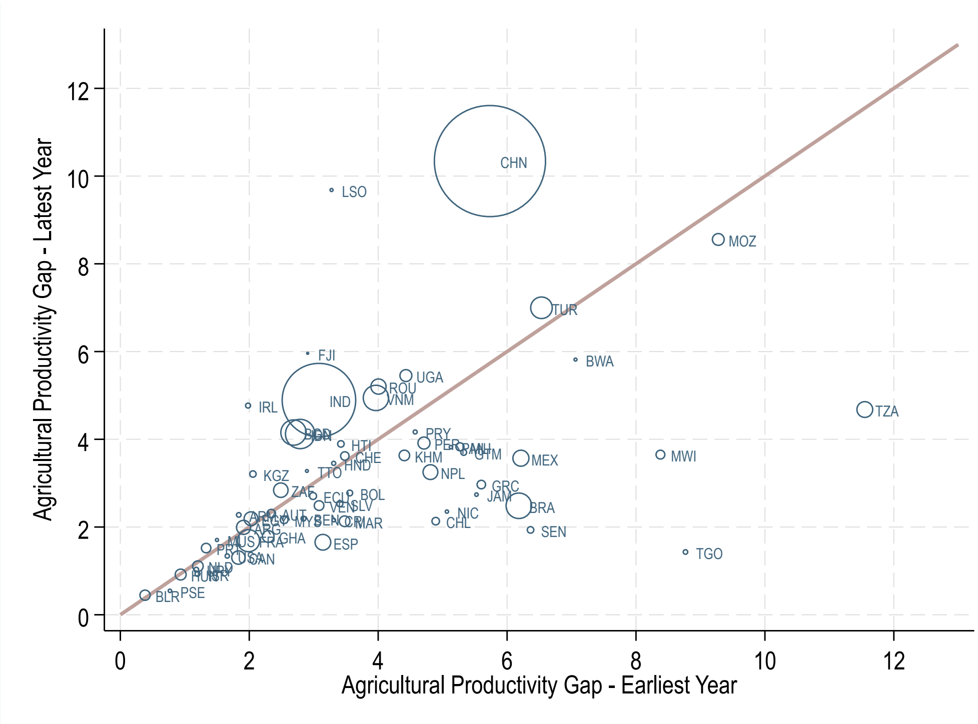

share of value added has fallen faster, leading to an increasing APG. China's gap rose from a little less than 6 to a little over 10 from 1980 to 2000; India's rose from 3 to around 5 from 1983 to 2009. The United States saw little change in agriculture's share of either employment or value added, which were quite low to begin with. Its APG changed little as a result, hovering at just under 2, with a modest dip around 2005.

In most countries, school completion rates at all levels increased steadily during the 1980s, 1990s, and 2000s (see e.g. Barro and Lee, 2013). If higher schooling levels result mostly in better educated non-agricultural workers, this could serve to raise APGs and explain part of the puzzle in Figures 1 and 2. Ma (2024) for example, shows that China's tertiary education expansion provided a boost for their export-oriented manufacturing sectors, and Porzio, Rossi, and Santangelo (2022) and Hobijn, Schoellman, and Vindas Q. (2018) use cross-country data to show that education is a key driver of structural change. To address this possibility, we imputed human capital per sector using educational attainment levels from our census data as in Gollin et al. (2014). These data, covering 32 countries, present a qualitatively similar story. Figure A.1 shows that we continue to have a wide variation in dynamics for the APGs even after human-capital adjustments. For example, Bangladesh still sees an increase in their gap while Brazil still has a decrease. Other countries are relatively stagnant over time. Figure A.2 shows that the differences between the

Figure 2: Changes in APGs and Agricultural Shares of Employment

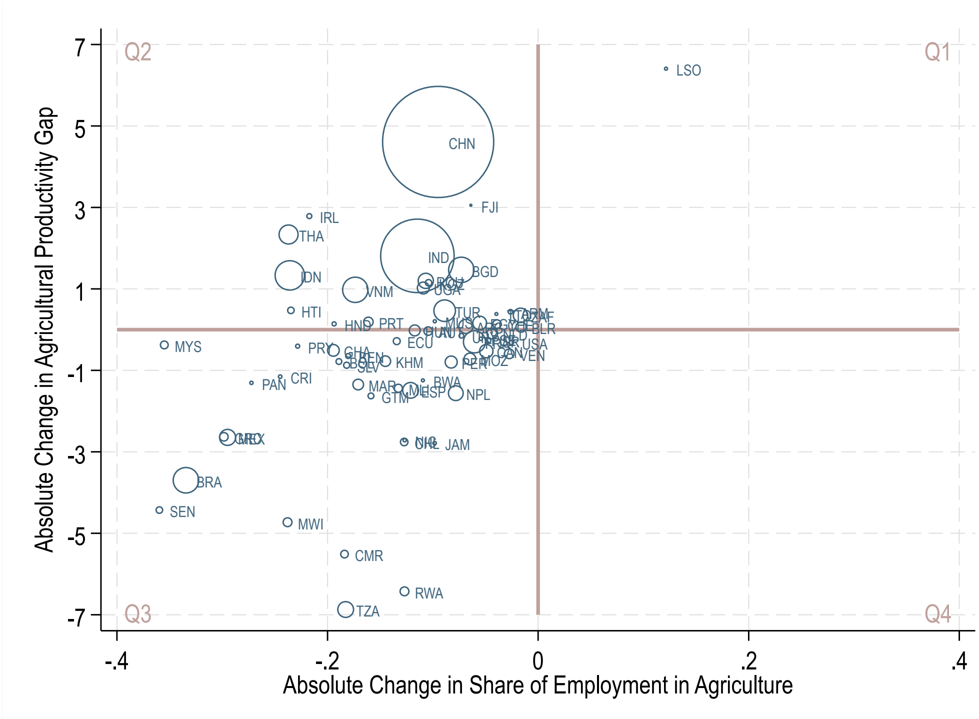

APGs with and without adjustments for human capital are quite modest.[3]

One takeaway from this analysis is that APGs showed no aggregate trend over the last few decades. In spite of a secular movement of labor out of agriculture, the average APG does not appear to have changed very much. The second point, however, is that there is substantial heterogeneity in experiences across countries. Some countries saw large and clear increases in their APGs. Others saw significant declines. The goal of the remainder of our paper is to understand these patterns.

II Quantitative Model

Motivated by the empirical patterns above, we now develop a multi-country model to understand the drivers of APG dynamics across countries. The model features real sectoral productivity growth to capture sector-specific trends at the country level, and non-homothetic preferences to capture

[3]For a smaller set of countries we also considered whether excluding the public sector or mining sector changed our findings, but found that it did not. Labor shares in production could also differ by sector in ways that could help explain variation in APGs. Gollin et al. (2014) concluded that, at least on average, the data did not support a significantly higher labor share in non-agricultural activities than in agriculture as would be need to explain the large observed APGs. For this reason we abstract from factors of production other than labor in the model that follows.

Figure 3: Agriculture Productivity Gap Trends for Select Countries

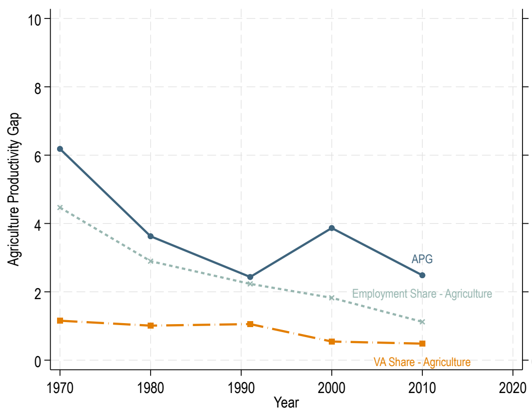

(a) Brazil

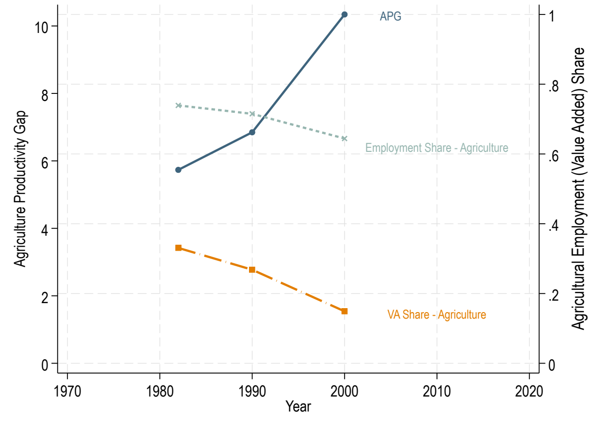

(b) China

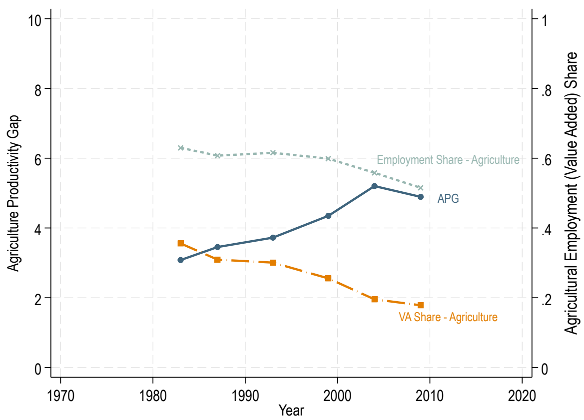

(c) India

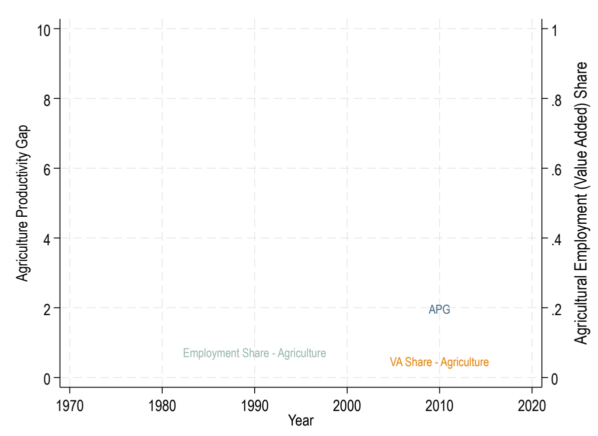

(d) United States

income effects in household expenditures. We also assume frictional labor mobility to capture patterns of sectoral employment shares, and in particular declining shares of agricultural employment.

Specifically, we assume that the world has many countries, $i = 1,...,I$. Each country i hosts both agricultural and non-agricultural sectors, $j \in \{a,n\}$, and each sector produces a homogeneous good. Agricultural and non-agricultural goods can be traded across countries subject to iceberg trade costs, and consumers view goods from different countries as imperfect substitutes (e.g., Armington, 1969). Each country has many households who each supply one unit of labor, and households are imperfectly mobile across sectors.

A Production and Trade

In each country i, many perfectly competitive producers supply a homogeneous non-agricultural good with a constant-returns-to-scale production function:

$$q_i^n = z_i^n l_i^n, \tag{2}$$

where z_i^n is the non-agricultural physical productivity in country i, capturing the impact of many factors such as technology or institutions. We denote the wage rate in country i's non-agriculture sector as w_i^n. Similarly, competitive producers in agriculture use labor (l_i^a) to produce a homogeneous agricultural good according to the production function:

$$q_i^a = z_i^a l_i^a, \tag{3}$$

where z_i^a is the fundamental productivity in agriculture. We denote w_i^a as the wage rate in country i's agriculture sector.

Goods can be traded across countries. Shipping one unit of a good from country i to country k in sector j incurs an iceberg trade cost $d_{ik}^j \geq 1 \ \forall \ i, j, k$, with domestic trade costs normalized to 1, $d_{ii}^j = 1$. Thus, under perfect competition, the unit price of goods produced in country i and sold in country k is:

$$p_{ik}^j = \frac{w_i^j d_{ik}^j}{z_i^j}. \tag{4}$$

Following Armington (1969), consumers view goods from different countries as distinct varieties with imperfect substitutability. Each country i has a composite sectoral consumption good produced from varieties sourced across countries:

$$y_i^j = \left(\sum_k (y_{ki}^j)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}. \tag{5}$$

Parameter $\eta > 1$ governs the elasticity of substitution between goods y_{ki}^j from different countries of origin k and thus determines the trade elasticity. Combining prices in equation (4) and the cost minimization of consumption baskets, we can solve the share of expenditures sourced from country

k within consumers' overall budget in sector j and country i:

$$\pi_{ki}^{j} = \frac{(p_{ki}^{j})^{1-\eta}}{\sum_{k'}(p_{k'i}^{j})^{1-\eta}} = \frac{(w_{k}^{j}d_{ki}^{j}/z_{k}^{j})^{1-\eta}}{\sum_{k'}(w_{k'}^{j}d_{k'i}^{j}/z_{k'}^{j})^{1-\eta}}. \tag{6}$$

Clearly, consumers in country i would source more from country k if country k had cheaper production costs (lower wages, w_k^j, or higher productivity, z_k^j) or lower trade costs to ship goods to country i (d_{ki}^j). The trade flow from country k to i in sector j is

$$X_{ki}^{j} = \pi_{ki}^{j}E_{i}^{j}, \tag{7}$$

where E_i^j is the total expenditures on sector j in country i.

The aggregate price per unit of the composite sectoral good is given by:

$$p_{i}^{j} = \left(\sum_{k}(p_{ki}^{j})^{1-\eta}\right)^{\frac{1}{1-\eta}}. \tag{8}$$

B Household Preferences

In country i, there is a measure l_i of households. Building on the existing literature (e.g. Boppart, 2014), we assume that household preferences belong to the non-homothetic Price-Independent Generalized Linear (PIGL) preferences class. More specifically, we parameterize the indirect utility on agricultural and non-agricultural goods following the functional form in Eckert and Peters (2022):

$$V_{i}(w) = \frac{1}{\zeta}\left(\frac{w}{(p_{i}^{a})^{\phi}(p_{i}^{n})^{1-\phi}}\right)^{\zeta} - v\log\left(\frac{p_{i}^{a}}{p_{i}^{n}}\right) \tag{9}$$

Here, $\zeta > 0$, $\phi > 0$, and $v > 0$ are structural parameters, and w is the wage level of the worker. As we show below, these preferences are important in matching expenditure shares and relative prices of agricultural goods in the data. Applying Roy's identity to this indirect utility function, we can obtain the share of expenditures on agricultural goods:

$$\vartheta_{i}(w) = \phi + v\left(\frac{w}{(p_{i}^{a})^{\phi}(p_{i}^{n})^{1-\phi}}\right)^{-\zeta} \tag{10}$$

Hence, ϕ represents the expenditure share on agricultural goods as income approaches its limit. The expression $w/((p_i^a)^\phi(p_i^n)^{1-\phi})$ can be viewed as real income, and therefore, ζ determines how agricultural expenditure shares decline with real income levels, reflecting the relative Engel curves. The parameter v determines agriculture's expenditure share for an "average" income level. Given this indirect utility function, the consumption levels of agricultural and nonagricultural goods for a worker with wage w are given by $c_i^a(w) = \vartheta_i(w)w/p_i^a$, and $c_i^n(w) = (1-\vartheta_i(w))w/p_i^a$.

As shown by Eckert and Peters (2022), the elasticity of substitution between consumption levels in the two sectors is given by $\rho = 1 + \zeta(\vartheta_i - \phi)^2/(\vartheta_i(1-\vartheta_i))$. This elasticity is close to 1 in our calibration, suggesting a minor role for substitution between the two sectoral goods in driving relative sectoral price changes. Therefore, the major drivers of relative sectoral price changes are the interactions between differential productivity growth and income effects. The relative prices between agricultural and non-agricultural goods impact consumption shares only modestly. This aligns with the findings by Comin, Lashkari, and Mestieri (2021), who demonstrate that the majority of structural change from agriculture originates from income effects rather than substitution effects. In Appendix C.5, we follow Boppart (2014) to allow for a nonunitary elasticity of substitution between two sectoral goods; our quantitative findings are robust to this model extension.

C Households' Sectoral Choices

Each household in country i supplies one unit of labor. At the beginning of the period, a portion $\lambda_i^{a,0}$ of workers stay in agriculture. Households face a cost to move between sectors. For simplicity, we assume that only a proportion κ of workers have the opportunity to switch sectors in each period. The compensation from working in sector j includes wage income w_i^j. A potential mover will choose the sector with higher utility, by comparing the wages of working in non-agriculture (w_i^n) and in agriculture (w_i^a). The resulting share of workers in agriculture is λ_i^a. Thus, the measure of workers in agriculture is $l_i^a = l_i\lambda_i^a$ and the measure of workers in non-agriculture is $l_i^n = l_i(1-\lambda_i^a)$.

D Equilibrium

In equilibrium, the goods market will clear for non-agriculture and agriculture in each country. This implies:

$$w_i^j l_i^j = \sum_k \pi_{ik}^j E_k^j. \tag{11}$$

The left-hand side is the overall production value of each sector in country i. Given perfect competition, this equals wages multiplied by labor supply. On the right-hand side, the total demand for goods (or labor) in country i and sector j is a summation of demand across all destinations, where $E_k^j = p_k^j \sum_{j'} l_k^{j'} c_k^j \left(w_k^{j'}\right)$ is country k's aggregate expenditure in sector j (aggregated over consumption of households working in different sectors j'), and π_{ik}^j is the share of expenses spent on goods sourced from country i. The labor supplies l_i^n and l_i^a are determined by households' sectoral choices with $l_i^n + l_i^a = l_i$.

In this model, the APG in country i is given by the ratio of value added per worker in non-agriculture relative to agriculture. Due to the linearity of the production function and perfect competition, the APG is simply the ratio of sectoral wages:

$$APG = \frac{(\sum_k X_{ik}^n)/l_i^n}{(\sum_k X_{ik}^a)/l_i^a} = \frac{w_i^n}{w_i^a}. \tag{12}$$

Note that sectoral wages are in general different from one another due to the migration frictions that prevent workers from moving freely across sectors. Thus, APGs in the model will not be equal to one. According to equation (11), the APGs are affected by trade shares, π_{ik}^j (shaped by physical productivity levels, $\{z_i^j\}$ and trade costs, $\{d_{ik}^j\}$); expenditures, E_k^j (governed by demand parameters $\{\phi, \nu, \zeta\}$); and sectoral labor supply, l_i^a (governed by initial employment share $\lambda_i^{a,0}$, the portion of workers that can switch sectors, κ, and total population, l_i). We will back out these parameters from the data in Section III to understand the drivers of APGs.

Given a set of parameter values, $\{z_i^j, d_{ik}^j, \eta, \phi, \nu, \zeta, \lambda_i^{a,0}, \kappa, l_i\}$, we can combine equations (6), (8), (10), (11), and households' utility maximization to solve for endogenous variables $\{\pi_{ik}^j, p_i^j, \vartheta_i, \lambda_i^a, w_i^j\}$. We solve the model numerically by applying the algorithm of Alvarez and Lucas (2007).

III Calibration

We now use our model to quantitatively understand the drivers of APGs. Based on data availability, we choose to calibrate a version of our model with 77 countries and a 'Rest of World' for the period 1980–2015. In what follows, we first describe the data sources and then discuss how we calibrate the model parameters.

A Data

We choose the countries in our quantitative analysis based on the availability of the data on trade flows, value added, and employment during the 1980–2015 period. The countries used in the calibration are listed in Appendix C.1. These closely, but not completely, overlap with the countries presented earlier. The reason we expand the number of countries for the quantitative work is to get better coverage of the world trade network. We add Japan and Germany, for example, even though we do not have APG data for either one.

Trade Data. For international trade data between 1980–2015, we utilize the United Nations Comtrade Database, which offers detailed information on bilateral trade flows (Feenstra and Romalis, 2014). We follow the procedure in Feenstra and Romalis (2014) to clean this database to obtain bilateral trade flows between countries based on 4-digit SITC products. We then map SITC products to agriculture and non-agricultural categories based on the WTO classification.The bilateral trade costs and contiguity measures are obtained from Conte, Cotterlaz, and Mayer (2022).

Sectoral Value Added and Employment. For each country involved, the hand-collected APG data has information on the agricultural shares of GDP and employment for several years, and we use linear interpolation to fill in and extrapolate the logarithm of these shares to other years with missing values. Combining this with the GDP and employment data from the Penn World Table, we obtain the yearly nominal value added (and value added per capita) of the agricultural and non-agriculture sectors for each country (Feenstra et al., 2015).

Because many countries' data are not available in the APG dataset described in the previous section, we supplement the APG data with observations from the OECD Structural Analysis Database (STAN), which reports information on nominal value added and employment by sector for 38 countries since 1980 (OECD, 2022). Appendix Figure B.1 confirms that our main empirical finding—declining agricultural employment shares alongside divergent changes in APGs—persists when using these OECD data. Finally, Appendix C.6 shows that our key quantitative results remain robust even when the calibration is limited to the set of countries in the APG data.

B Calibration Procedure

We allow trade costs and productivity levels to change over time to capture the fact that many developing countries have experienced dramatic shifts in productivity and trade openness (e.g., China, India). We take each year's employment l_{it} directly from the data. All other parameters are time-invariant. We describe each step of the calibration procedure in turn.

Calibrating Productivity Levels and Trade Costs We estimate sectoral fundamental productivity levels and trade costs using bilateral trade data. Starting from the trade value in equation (7), we can obtain the relative trade values (Eaton and Kortum, 2002):

$$\frac{X_{kit}^{j}}{X_{iit}^{j}} = \frac{(w_{kt}^{j} d_{kit}^{j} / z_{kt}^{j})^{1-\eta}}{(w_{it}^{j} / z_{it}^{j})^{1-\eta}}. \tag{13}$$

By using relative values, aggregate prices and quantities cancel out in the formula. For tractability of estimation, we also need to impose functional form restrictions on trade costs d_{kit}^{j}. We follow the literature (e.g., Head and Mayer, 2014) and assume that trade costs take the following functional form:

$$d_{kit}^{j} = \exp(\beta_{1t}^{j} \log dist_{kit} + \beta_{2t}^{j} contig_{kit} + \beta_{3t}^{j} \log GDPPC_{kt}) \tag{14}$$

where $dist_{kit}$ is the distance between the capitals of country k and country i; $contig_{kit}$ is a dummy variable that indicates whether two countries are contiguous. Finally, we allow trade costs to depend on GDP per capita of the exporting countries, reflecting that poorer countries may have higher export iceberg costs as shown by Waugh (2010). In Waugh (2010), a group of country-specific dummy variables is introduced to estimate export iceberg costs associated with development levels. However, due to the increase in estimation imprecision caused by introducing an additional full set of country-specific dummy variables (especially considering the estimation performed for each year and sector), we opt to directly model trade costs as a function of GDP per capita.

It is also worth noting that for some origin-destination pairs, X_{kit}^{j} could be zero, and thus the logarithm of these trade flows does not exist, biasing the estimates if we perform the regression using log values. Thus, we follow Silva and Tenreyro (2006) to estimate equation (13) using the Poisson pseudo-maximum-likelihood estimator, allowing for zero bilateral trade flows:

$$\frac{X_{kit}^{j}}{X_{iit}^{j}} = \exp\left(S_{kt}^{j} - S_{it}^{j} + (1-\eta)\left(\beta_{1t}^{j} \log dist_{kit} + \beta_{2t}^{j} contig_{kit} + \beta_{3t}^{j} \log GDPPC_{kt}\right)\right) + \varepsilon_{kit}, \tag{15}$$

where fixed effects $S_{kt}^{j} = (1-\eta)(\log w_{kt}^{j} - \log z_{kt}^{j})$ capture marginal costs, measuring the "competitiveness" of country k in sector j. It is worth noting that in equation (15), we can only estimate S_{kt}^{j} relative to a reference country (note that changing all S_{kt}^{j} by the same amount does not affect the regression fit).[4] We estimate this equation for each sector and year.

Recovering trade costs. Note that η cannot be identified separately from β^j_{1t}, β^j_{2t}, and β^j_{3t}. We use the trade elasticity $\eta - 1 = 4$ according to common estimates in the literature (Head and Mayer, 2014; Simonovska and Waugh, 2014). We recover yearly and sectoral trade costs from the estimates:

$$\hat{d}^j_{kit} = \exp\left(\hat{\beta}^j_{1t}\log dist_{kit} + \hat{\beta}^j_{2t} contig_{kit} + \hat{\beta}^j_{3t}\log GDPPC_{kt}\right).$$

Recovering productivity levels. To recover estimates of fundamental productivity levels z^j_{it} from $\hat{S}^j_{it}$, we need information on wage rates w^j_{it}. It is worth noting that wage rates are endogenous variables in our model: if we apply a set of country-sector-year productivity levels $\{z^j_{kt}\}$ in the model, we can endogenously solve wage levels $\{w^j_{kt}\}$ from market clearing conditions in equation (11) and obtain the model-predicted level of $\{S^{j,model}_{kt}\}$. Given this observation, we search for the set of productivity levels $\{z^j_{kt}\}$ in the model that minimize the summed absolute difference between the model-predicted "competitiveness" level $\{S^{j,model}_{it}\}$ and the model estimates $\{\hat{S}^j_{it}\}$, $\sum_t \sum_i \sum_j |\hat{S}^j_{it} - S^{j,model}_{it}|$. Furthermore, since we estimate S^j_{it} solely in relation to a reference country as depicted in equation (13), we can only derive relative productivity levels through this process. To determine absolute productivity levels for all countries, we must also know the productivity levels of the reference country. We opt for the United States as our reference country, normalize its sectoral productivity to one in the initial year (1980), and employ sectoral growth data of output per worker from FRED and USDA to measure the U.S. sectoral productivity levels. Subsequently, we obtain the productivity levels of all other countries. The estimates of productivity growth obtained from this process are broadly consistent with direct evidence on sectoral output per worker for a smaller set of countries; we return to this issue below.

Calibrating Parameters in Utility Function We proceed to calibrate the parameters $\{\zeta, \phi, v\}$ within the indirect utility function. The choice of ζ is critical as it determines the elasticity of agricultural expenditure shares with respect to real income. Following the estimate conducted by Eckert and Peters (2022) using historical U.S. data, we select a value of 0.93 for ζ. As for ϕ, representing the asymptotic expenditure share in agriculture, we base our selection on the U.S. expenditure share in 2015, which leads us to choose a value of 0.01. The parameter v governs

[4] An alternative approach to addressing this issue is to introduce an additional constraint. For example, Eaton and Kortum (2002) and Waugh (2010) assume that the sum of fixed effects is zero. The distinction between our method and this alternative approach is that all S^j_{kt} can differ by a constant, which will not affect the quantitative results.

Figure 4: Levels of APGs: Model vs Data

the magnitude of the variable portion of agricultural expenditure shares. We choose v to match the APG for the United States in the initial year of our sample period, resulting in a value of $v = 0.008$.

Calibrating Sectoral Labor Adjustments Finally, we choose the proportion of workers who have the opportunity to switch sectors, $\kappa = 0.02$, to match the average decline in agricultural employment shares across our sample countries and years. We set initial (1980) agricultural employment shares to match the data directly.

IV Quantitative Evaluation

In this section, we begin by comparing the fit of our calibrated model with the actual data. Next, we conduct several counterfactual experiments to identify the factors driving the dynamics of APGs. Finally, we show that our model-estimated productivity growth aligns well with the empirical data.

We begin by comparing APG levels between the model and the data, in logs, for all country-years in our sample, none of which were targeted in our calibration procedure. As shown in Figure 4, the APG predicted by our model exhibits a strong positive correlation with the observed gaps, with a correlation coefficient of 0.71. The correlation is 0.64 if we focus on the first observation of

Figure 5: Model Fit in Four Select Economies

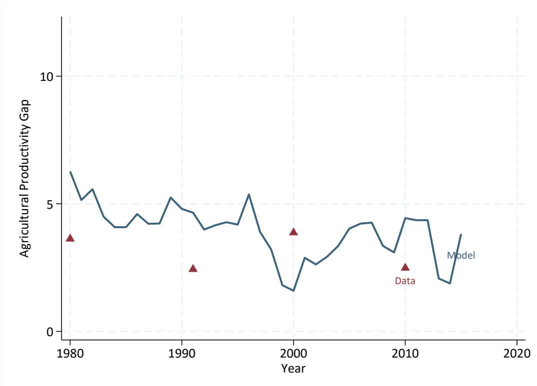

(a) Brazil

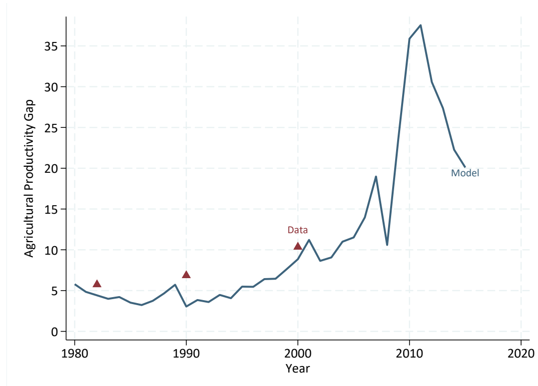

(b) China

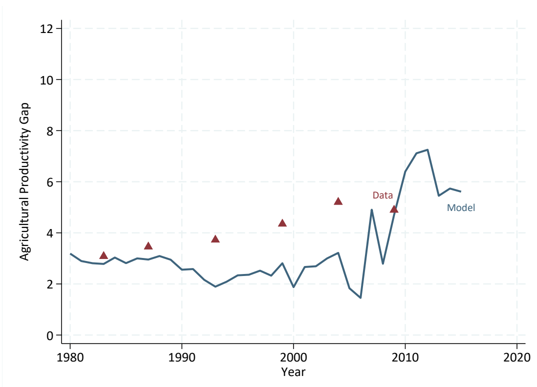

(c) India

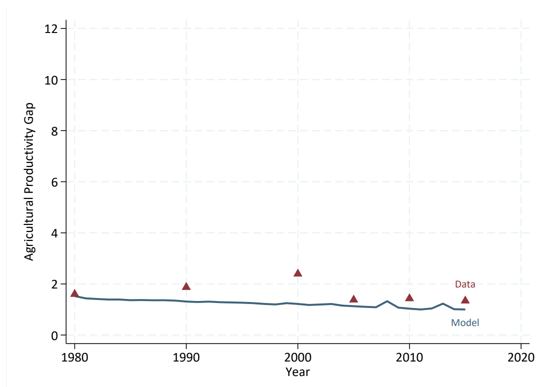

(d) United States

the APGs for each country.

Our primary focus is on the dynamics of gaps. Figure 5 illustrates the evolution of APGs in the model and data for Brazil, China, India, and the United States, which were not targeted in the calibration. The model captures the temporal patterns observed in these countries fairly well. Figure 6 shows the actual dynamics alongside the model-generated changes for all of our countries. In the data, we take the difference between the first observation and the last observation for each country with at least two years of data observed during the 1980–2015 period. In the model, for each country, we compute the changes in APGs using the same time window as in the data. The correlation between the model's projected changes and the observed changes is 0.80. We find that the model's predictions are highly correlated with the data when conditioning on whether the APG increased or decreased over the period. The model fit is also similar for more open and less open countries; see Figure B.5 and Table B.1 for the full set of model predictions.[5]

[5]Figure B.4 shows that the model-generated changes in agricultural employment shares align reasonably well

Figure 6: Dynamics of APGs: Model vs Data

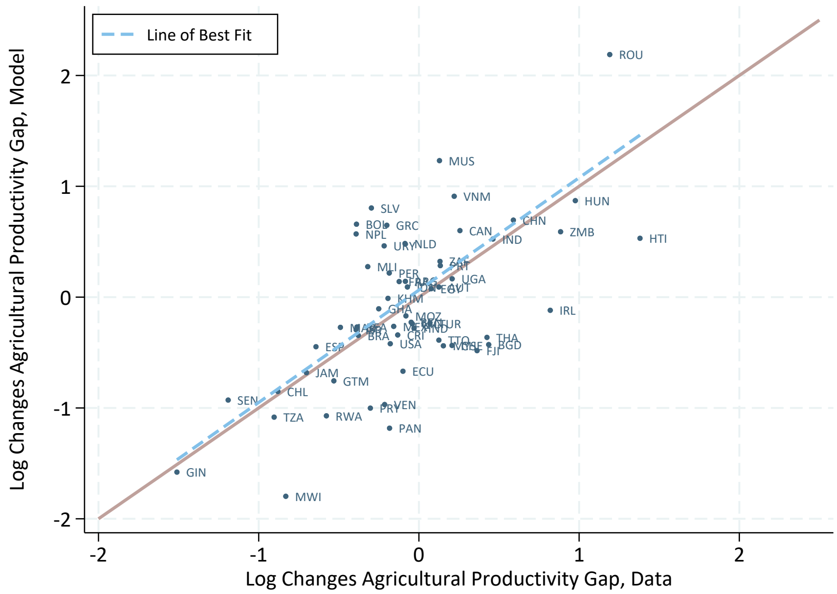

A Drivers of Dynamics in APGs

Our model encapsulates time-varying trade costs, time-varying productivity levels, and sectoral labor reallocation, all of which contribute to the evolution of APGs. Without these ingredients, the gaps in agricultural productivity would have remained static. To assess the importance of these dynamic channels, we perform several counterfactual exercises. Since different channels may interact with each other, the marginal contribution of any one channel depends on what is assumed about the others. Therefore, we follow the Shapley decomposition of Hao, Sun, Tombe, and Zhu (2020) to assess the marginal effect of each channel on APGs by averaging its marginal contributions across all possible orderings of the other two channels.

Role of Real Sectoral Productivity Growth. First, we examine the influence of sectoral productivity growth by comparing the calibrated model to a version with productivity fixed at 1980 levels. We do so by averaging over all combinations of time-varying versus fixed trade costs and

with the data. Our model tends to overestimate changes in agricultural employment (weighted by employment) in China and India, consistent with the large migration barriers found in those countries (e.g, Ngai et al., 2019; Tombe and Zhu, 2019). In Appendix C.3, we utilize a country-specific proportion of workers with opportunities to switch sectors, denoted κ_i, based on the average proportional change in agricultural employment shares over the years for each country. This extended model matches the observed changes in agricultural employment somewhat more closely, though leaves the predictions for the dynamics of APGs largely unaffected.

Table 1: Effect of Real Sectoral Productivity Growth on APG Changes

	Δ Log APG		
	(1)	(2)	(3)
Average Productivity Growth	0.603***		0.427***
	(0.188)		(0.0772)
Relative Productivity Growth		0.662***	0.626***
		(0.0446)	(0.0259)
Observations	59	59	59
Adjusted R-squared	0.199	0.833	0.938

Notes: The dependent variable is the change in model-generated APG in the baseline model measured over the same time window as in the data for each country. The independent variables are the average sectoral productivity growth and the productivity growth differential between non-agriculture and agriculture over that same period. Regressions are weighted by employment in each country in the initial year of observation. Robust standard errors are in parentheses. Significant levels: * 10%, ** 5%, *** 1%.

with or without sectoral labor mobility.[6] This counterfactual predicts a strong correlation between the model's APG changes and those in the data. See panel (b) of Figure 7. Yet the predicted productivity gaps are shifted upward substantially (and counterfactually), with an average change that is 27 log points in the model versus -7 in the data.

The average increase in APGs can be attributed both to a decline in the relative demand for agricultural goods as overall productivity growth accelerates, and to relative productivity growth itself. Table 1 provides evidence for this. In Column (1), we regress the model-generated APG changes in the baseline specification on average sectoral productivity growth and find that higher average productivity growth is associated with larger APGs. Column (2) instead incorporates the productivity growth differential between non-agriculture and agriculture. The results indicate that a wider differential leads to larger APGs and accounts for a greater share of the variation in the model-generated APG dynamics than average productivity growth alone. Consistent with this pattern, Column (3) shows that adding the productivity growth differential to average productivity growth explains a greater share of the variation in model-generated APG dynamics than average productivity growth alone.

To further examine the drivers of APG dynamics, we conduct a simple variance decomposition where we consider the R-squared from a regression of APG changes on various independent variables. Table 2 shows that real sectoral productivity growth alone accounts for 58.3 percent

[6]There are four possible combinations in this case: both trade costs and sectoral employment shares fixed at their 1980 levels; trade costs varying over time while sectoral employment shares remain fixed at 1980 levels; trade costs fixed at 1980 levels while sectoral employment shares change over time; and both trade costs and sectoral employment shares varying over time. We isolate the contribution of productivity growth by comparing APG dynamics with 1980–2015 productivity growth to a version fixed at 1980 levels for each combination and then averaging the results.

Figure 7: Dynamics of APGs in Counterfactual Scenarios

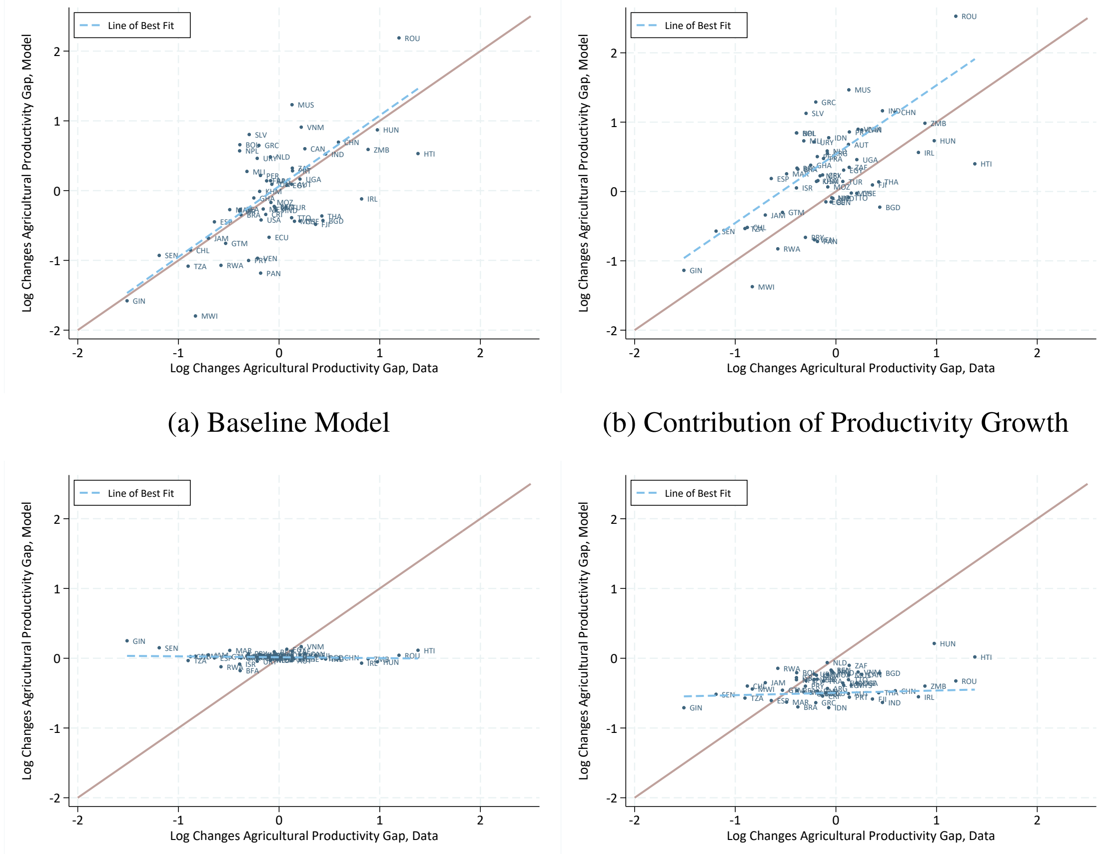

(a) Baseline Model

(b) Contribution of Productivity Growth

(c) Contribution of Changes in Trade Costs

(d) Contribution of Sectoral Labor Movements

of the cross-country variation in observed gap dynamics, again pointing to its important role in explaining the patterns at hand, even if it makes inaccurate predictions for the level changes.[7]

Role of Labor Sectoral Reallocation. We next assess the role of labor mobility by comparing the calibrated model, which allows workers to move across sectors (subject to frictions), to a counterfactual in which sectoral employment shares are fixed at their 1980 levels. As workers move from agriculture to the non-agricultural sector, driven by higher wages in the latter, the result is a reduction in APGs in the majority of countries. This is shown in panel (d) of Figure 7.

This counterfactual exercise matches the intuition that workers moving out of agriculture tend to help equate value added per worker in the agriculture and non-agriculture sectors. Yet it does not

[7]In Appendix C.2, we present an alternative decomposition that breaks down gap dynamics into movements in the relative price and real relative productivity between the two sectors. This decomposition similarly highlights that real productivity changes are the primary driver of country variation in APG changes.

Table 2: Variance Decomposition of APG Changes

Share of Variance Explained	Data Change
1. Productivity Growth	0.583
2. Trade Costs	0.018
3. Worker Sectoral Reallocation	0.008
4. Interaction Effects	0.025
5. Overall	0.634

Note: Rows 1–3 report R-squared values from regressing data changes on APG changes attributable to productivity growth (Row 1, Panel b in Figure 7), trade costs (Row 2, Panel c), and worker sectoral movements (Row 3, Panel d). Row 5 reports the R-squared for the baseline model combining all three channels, and Row 4 shows the difference between Row 5 and the sum of Rows 1–3, capturing interaction effects among the channels.

help explain much of the country variation in APG changes. In the data, there has been a global tendency for agricultural employment shares to decline, but marked diversity in APG dynamics across nations. As such, worker movements across sectors explain a negligible share of the cross-country patterns of gap dynamics, at just 0.8 percent.

Declining employment shares in agriculture are nonetheless an important part of the explanation for the observed patterns of APG changes. Changes in real sectoral productivity by themselves tend to overstate the magnitude of changes in gaps relative to the data, and worker reallocation out of agriculture brings down the average APG change, offsetting the upward bias from productivity growth alone.

Role of Trade Costs. The literature documents a decrease in trade barriers and a simultaneous increase in globalization levels in recent decades (e.g., Caliendo, Feenstra, Romalis, and Taylor, 2015; Caliendo and Parro, 2015), which can catalyze structural transformation (Cravino and Sotelo, 2019). Consequently, we examine the role of changing trade costs by comparing the calibrated model—which incorporates post-1980 changes in trade costs d^j_{kit}—with a counterfactual that holds trade costs fixed at their 1980 levels. As highlighted in panel (c) of Figure 7, we observe that the impact of trade cost changes on APGs tends to be relatively insignificant compared to the role played by productivity growth, explaining only 1.8 percent of cross-country variation in gap dynamics as shown by Table 2.

B Real Sectoral Productivity in the Model and Data

The model's predictions for real productivity by sector play a central role in explaining the dynamics of APGs. For a subset of our countries and a shorter time period, we can compare our model's predictions to direct measures of real sectoral output per worker. We make these comparisons in two data sets. The first is the EU KLEMS database, which estimates the labor productivity (based on quantities of value added per hour worked) for many European countries and the United States; these data are available for the 1975–2015 period (The Conference Board, 2023). The second data set is from Herrendorf et al. (2026), who estimate labor productivity (based again on quantities of value added per worker) for a larger set of countries, including many developing countries; these data are available for the 1990–2018 period. For each country-sector pair, in the data, we compute the productivity growth between the first observation and the last observation in the period 1980–2015; in the model, we compute productivity growth using the same time window as in the data.

Figure 8 compares the model-estimated sectoral productivity growth against the data, with an auxiliary 45-degree line. We find that our model-estimated sectoral productivity growth is positively correlated with both data sets. The correlation coefficient between model and data is 0.44 using the EU KLEMS estimates and 0.50 using those of Herrendorf et al. (2026). This suggests that our model estimates of sectoral productivity growth and theirs reflect some common underlying patterns of real changes in output per worker at the country and sector level.

As an alternative to the main parameterization, we calibrate the model to the productivity estimates from Herrendorf et al. (2026). Due to data availability, this analysis covers 53 countries over the shorter 1990–2015 period (see Appendix C.7 for more details). The APG dynamics implied by this alternative calibration remain closely aligned with the data, similar to the baseline model. This consistency reflects the reasonably strong correspondence between our estimated productivity growth rates and those computed by Herrendorf et al. (2026).

However, using productivity estimates from Herrendorf et al. (2026) leads to a significantly weaker fit for the *levels* of APGs. Figure 9 plots the levels of APGs in Brazil, China, India, and the United States in our main calibration and when using the sectoral output per worker estimates of Herrendorf et al. (2026). For Brazil, China, and India, the alternative calibration predicts APG levels 2.5–4.4 times higher than the data in the 1990s. This pattern holds more generally, and not just for these three countries (see Figure C.8). We find that this discrepancy arises mainly because Herrendorf et al. (2026) report relatively higher levels of non-agricultural productivity than those implied by our model, particularly for poorer countries. We plot their productivity levels by sector against ours in Figure C.9.

Figure 8: Comparison of Productivity Growth between Data and Model Estimates

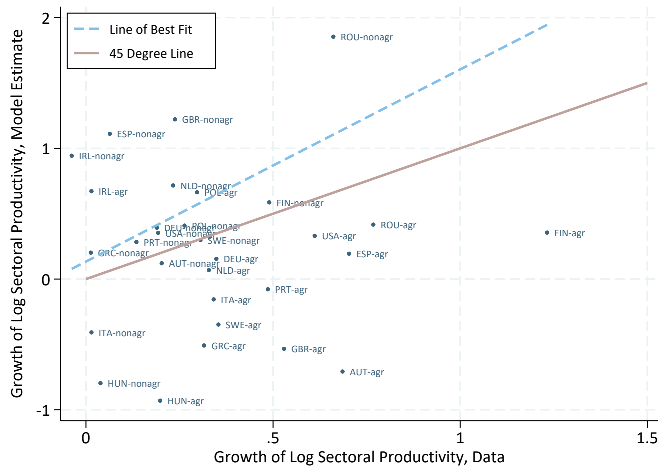

(a) EU KLEMS

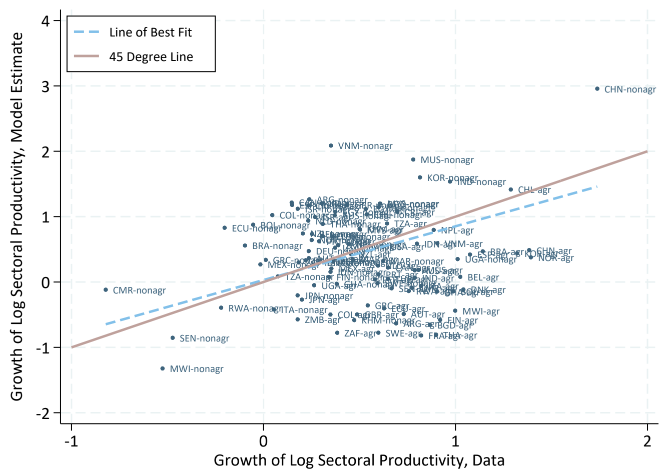

(b) Data from Herrendorf et al. (2026)

C The Role of Non-Homothetic Preferences

In our baseline model, we considered PIGL preferences for households. This type of preferences will generally lead to an increase in relative demand for non-agriculture as the economy grows, mainly driven by income effects (Comin et al., 2021), which would intensify the APG. To evaluate this shift in relative sectoral demand, we now assume instead that the utility function is homothetic and Cobb-Douglas:

$$U_i = (c_i^a)^{\alpha_i}(c_i^n)^{1-\alpha_i}. \tag{16}$$

Figure 9: APG Levels using Sectoral Productivity Data from Herrendorf et al. (2026)

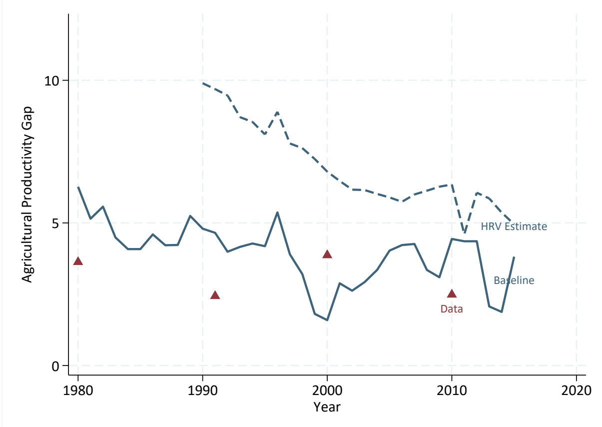

(a) Brazil

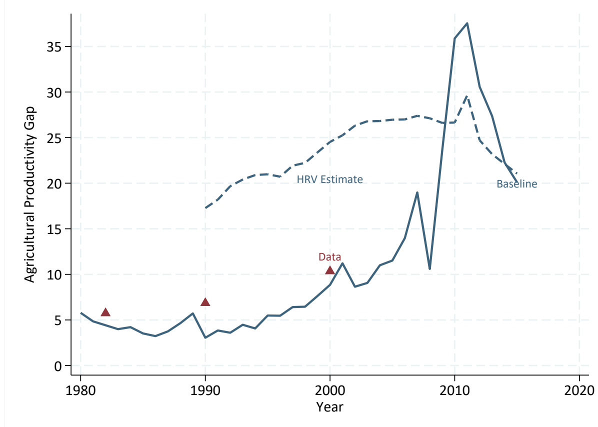

(b) China

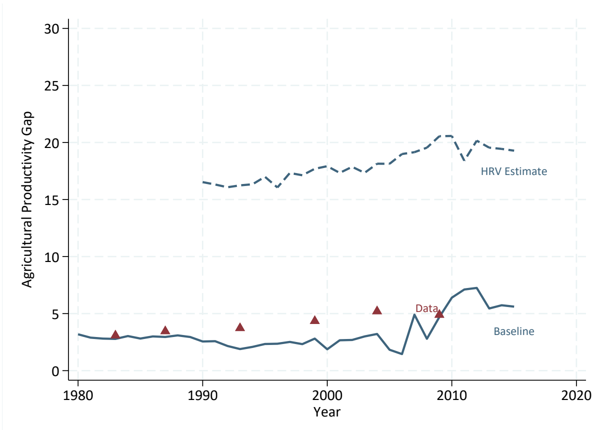

(c) India

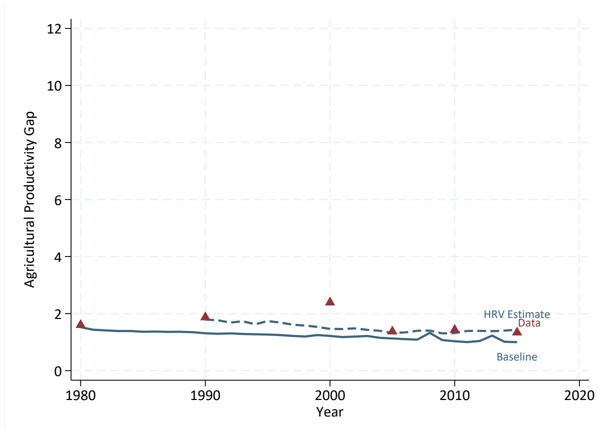

(d) United States

We calibrate α_i for each country i such that our model generates the same sectoral expenditure shares in the initial year as in the baseline calibration. We keep all other parameters unchanged.

In Figure 10a, we illustrate the model-generated changes in APGs (now calculated using Cobb-Douglas preferences) compared with the baseline results. We find that for the majority of points, the model-projected changes in APGs are smaller under homothetic preferences than under non-homothetic preferences. For many countries, the model even predicts declines in APGs where positive changes are observed in the data (Figure B.8). We find that the share of variance in observed APG dynamics explained by the three channels decreases from 63.4 percent in the baseline model (Table 2) to 49.8 percent, suggesting that income effects account for roughly 13.6 percent of the variance explained by the three channels. Appendix Figure B.9 shows that the contribution of productivity growth to APG dynamics is smaller with homothetic Cobb-Douglas preferences

Figure 10: Comparison of Dynamics of APGs in Different Scenarios

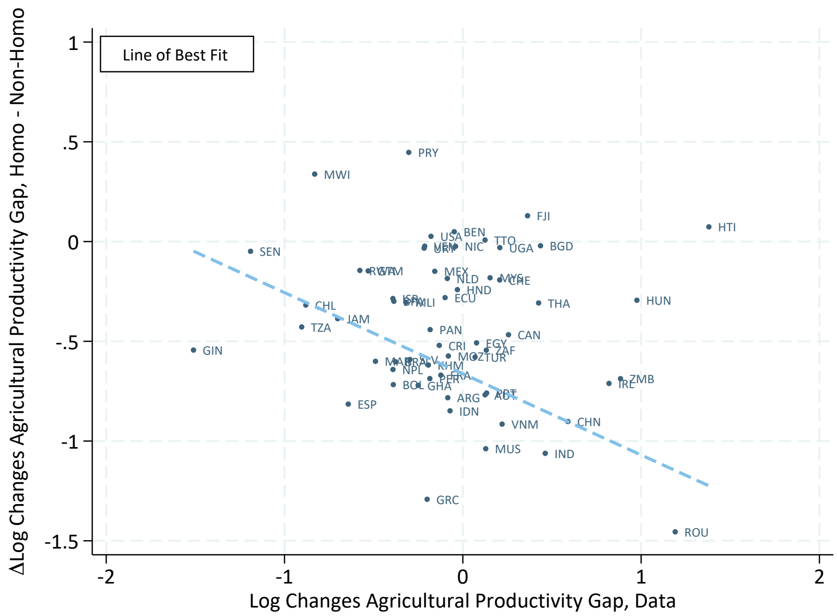

(a) Homothetic versus Non-Homothetic Preferences

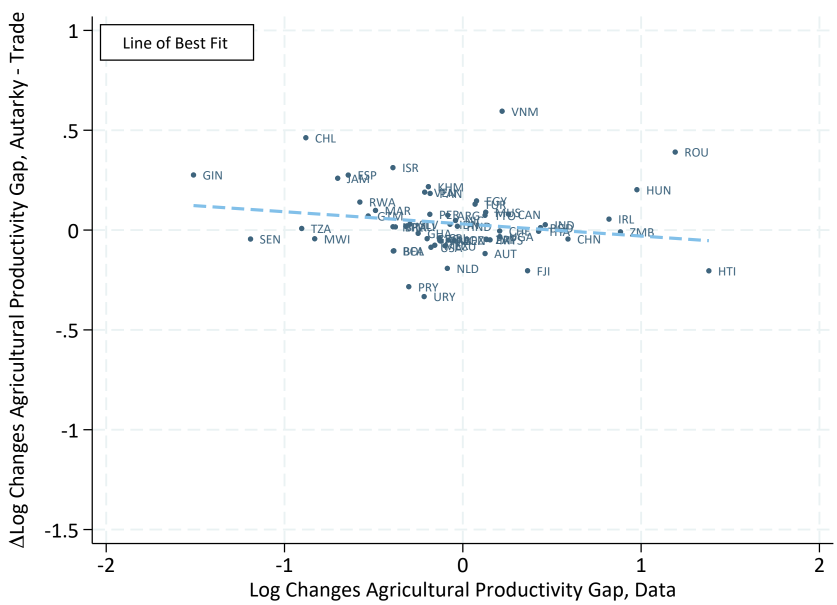

(b) Autarky versus Trade Openness

Note: In Panel (a), we plot the difference of combined effects of productivity growth, trade growth, and labor movements between the model with Cobb-Douglas preferences and the baseline model (with non-homothetic preferences). In Panel (b), we plot the difference of combined effects of productivity growth and labor movements between autarky and the baseline scenario (with trade openness).

than in the baseline scenario (Figure 7), while the contributions of trade costs and labor movements remain similar. This highlights the importance of non-homothetic preferences—particularly the income effects of productivity growth—in amplifying APGs by shifting demand away from agriculture.

D The Role of International Trade

Although Figure 7 demonstrated that changes in trade costs do not significantly influence the dynamics of APGs, keeping trade open could still affect the impact of other factors. In Figure 10b, we examine the combined effects of productivity growth and sectoral shifts in labor on the dynamics of APGs, comparing the baseline model with trade openness to an autarkic scenario (where trade costs between countries are assumed to be prohibitively high and other parameter values remain unchanged at the baseline levels). While the average differences are near zero, we observe considerable variation across countries, with a standard deviation of 0.17 on a logarithmic scale. This suggests that by considering a closed economy rather than an open one, our predictions for changes in APGs could err by tens of percentage points for a typical country.

While many factors may jointly cause such different predictions in our general equilibrium framework, one key driver is that trade openness enables countries to leverage their comparative advantages. Specifically, in Figure B.10, we illustrate that the degree of underprediction is more pronounced in countries with higher relative productivity growth in agriculture. For instance, Chile experienced rapid relative productivity growth in agriculture, but under autarky, it would be unable to fully take advantage of its agricultural comparative advantages. As a result, this would constrain agricultural wage income and widen the APG in the country.

Another key motivation for using an open-economy framework is that it allows us to leverage the gravity equation to estimate productivity growth. In Appendix C.8, we explore an alternative approach that estimates each country's sectoral productivity growth using non-agricultural expenditure shares and real GDP per capita in a closed-economy setting. We find that this method tends to underestimate sectoral productivity growth for large exporting countries (e.g., non-agriculture in China and India), resulting in a poorer fit of the model's APG dynamics to the observed data compared with the baseline model.

Table 3: Agricultural Productivity Dynamics and Country Characteristics

	Change in Log(Observed APG)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Agricultural export growth	-0.208						
	(0.142)						
Nonagricultural export growth	0.321^{***}						
	(0.0985)						
Log (GDP per capita)		-0.208^{**}					
		(0.0836)					
Growth of GDP per capita			0.479^{***}				
			(0.148)				
Change in nonagr expenditure share				4.819^{***}			
				(0.546)			
Change in nonagr employment share					-3.911^{***}		
					(1.332)		
Human capital growth						-0.112	
						(0.884)	
Agricultural productivity growth							-10.96^{***}
							(3.978)
Nonagricultural productivity growth							7.380^{***}
							(1.136)
Observations	59	59	59	59	59	58	39
Adjusted R-squared	0.254	0.129	0.103	0.630	0.166	-0.017	0.743

Note: The dependent variable is the change in the observed APG. For each country with at least two years of data between 1980 and 2015, we measure this as the difference between the first and last available observations. The key independent variables are the log growth of sectoral exports, the change in nonagricultural employment share, the change in nonagricultural expenditure shares, the log growth and the level of GDP per capita, and the log growth of human capital index for the same time window, with the latter two drawn from the Penn World Table. We measure exports and domestic expenditures by combining sectoral value added and trade data. The productivity measure follows Herrendorf et al. (2026), which is available only after 1990; accordingly, we compute annualized productivity growth up to each country's last observation year. We weight the regression by the observed employment. Robust standard errors are in parentheses; the significance levels are: * 10%, ** 5%, *** 1%.

V Model Validation: Variation in Gap Dynamics

To help validate the model's predictions in a way that relies less on the structure of the model, we examine the correlations between observed APG dynamics and country characteristics that closely relate to our theory.

Table 3 reports country-level regressions of log APG changes over the 1980–2015 period on a set of country characteristics. Column (1) shows that APG growth is positively associated with growth in nonagricultural exports and negatively associated with growth in agricultural exports, highlighting a close empirical link between trade patterns and APG dynamics, as we posit. Column

(2) indicates that high-income countries tend to experience smaller changes in APGs. Columns (3)–(5) show that changes in APGs are positively correlated with income growth and increases in the non-agricultural expenditure share, but negatively correlated with increases in the nonagricultural employment share. In contrast, APG changes exhibit little correlation with human capital growth, consistent with our adjustments described above. Finally, Column (7) incorporates sectoral productivity growth from Herrendorf et al. (2026) for the 39 countries with available data and shows that APG growth is positively related to non-agricultural productivity growth and negatively related to agricultural productivity growth in these countries.

These observed empirical patterns cannot prove causal relationships between the independent variables and the observed APG changes that are the focus of this paper. Yet they help underscore the role of the real economic forces that underlie our theory, and help reinforce the relationships we posited in our quantitative model.

VI Conclusion

Development economists have long been interested in the determinants of sectoral differences in output per worker (see e.g. Lewis, 1954; Harris and Todaro, 1970; Fields, 2005). Recent research has focused on the substantial gaps in value added per worker between the non-agricultural and agricultural sectors (Vollrath, 2009; Gollin et al., 2014; Herrendorf and Schoellman, 2015). This paper draws on new evidence to show that these gaps have shown little tendency to decline over the past three decades. Instead, they have significantly widened in some major economies, such as China and India, and fallen markedly in others, with roughly similar sized gaps observable now, on average, as in the 1980s and 1990s.

Our quantitative open-economy model of structural change does a reasonable job of explaining the gap dynamics overall, from China's widening gap amid rapid non-agricultural export growth to Brazil's narrowing gap as a major agricultural exporter. The main forces in explaining rising gaps, according to our model, are faster real productivity growth in the non-agriculture sector, and frictions in re-allocating workers out of agriculture. As in the data, countries with declining agricultural productivity gaps tend to have relatively fast real labor productivity growth in the agriculture sector. The global pattern of declining employment shares in agriculture helps to reduce agricultural productivity gaps, all else equal, which helps to offset the relatively faster growth of real non-agricultural output per worker experienced in most countries.

Our findings provide little support for the possibility that variation in APGs over time largely reflects measurement error in the national income accounts, as emphasized by many observers (e.g.

Jerven, 2013). While questions will always persist about the reliability of national statistics in low-income countries, our study points to the role of standard economic forces in shaping measured gaps in sectoral value added per worker. The sluggish pace of labor migration out of agriculture, which helps sustain the sectoral gaps, remains an important topic for future research.

References

ADAMOPOULOS, T., L. BRANDT, C. CHEN, D. RESTUCCIA, AND X. WEI (2024): "Land Security and Mobility Frictions," *Quarterly Journal of Economics*, 139, 1941–1987.

ADAMOPOULOS, T., L. BRANDT, J. LEIGHT, AND D. RESTUCCIA (2022): "Misallocation, Selection, and Productivity: A Quantitative Analysis With Panel Data From China," *Econometrica*, 90, 1261–1282.

ADAMOPOULOS, T. AND D. RESTUCCIA (2014): "The size distribution of farms and international productivity differences," *American Economic Review*, 104, 1667–1697.

——— (2022): "Geography and agricultural productivity: Cross-country evidence from micro plot-level data," *The Review of Economic Studies*, 89, 1629–1653.

ALESSANDRIA, G., R. C. JOHNSON, AND K.-M. YI (2023): "Perspectives on Trade and Structural Transformation," *Oxford Development Studies*, 51, 455–475.

ALVAREZ, F. AND J. R. E. LUCAS (2007): "General Equilibrium Analysis of the Eaton–Kortum Model of International Trade," *Journal of Monetary Economics*, 54, 1726–1768.

ALVAREZ, J. A. (2020): "The Agricultural Wage Gap: Evidence From Brazilian Micro-Data," *American Economic Journal: Macroeconomics*, 12, 153–173.

ARMINGTON, P. S. (1969): "A Theory of Demand for Products Distinguished by Place of Production," *Staff Papers (International Monetary Fund)*, 16, 159–178.

BALASSA, B. (1963): "An Empirical Demonstration of Classical Comparative Cost Theory," *The Review of Economics and Statistics*, 45, 231–238.

BARRO, R. J. AND J. W. LEE (2013): "A New Data Set of Educational Attainment in the World, 1950–2010," *Journal of Development Economics*, 104, 184–198.

BASELER, T. (2023): "Hidden Income and the Perceived Returns to Migration," *American Economic Journal: Applied Economics*, 15, 321–352.

BAYSAN, C., M. H. DAR, K. EMERICK, Z. LI, AND E. SADOULET (2024): "The agricultural wage gap within rural villages," *Journal of Development Economics*, 168, 103270.

BETTS, C., R. GIRI, AND R. VERMA (2017): "Trade, Reform, and Structural Transformation in South Korea," *IMF Economic Review*, 65, 745–791.

BOPPART, T. (2014): "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82, 2167–2196.

BOPPART, T., P. KIERNAN, P. KRUSELL, AND H. MALMBERG (2023): "The Macroeconomics of Intensive Agriculture," Working Paper 31101, NBER.

BRYAN, G. AND M. MORTEN (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127, 2229–2268.

CALIENDO, L., R. C. FEENSTRA, J. ROMALIS, AND A. M. TAYLOR (2015): "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades," Working Paper 21768, NBER.

CALIENDO, L. AND F. PARRO (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 82, 1–44.

CAO, H., C. CHEN, X. XI, AND S. X. ZUO (2024): "Family Migration and Structural Transformation," *Canadian Journal of Economics*, 57, 753–776.

CASELLI, F. (2005): "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion and S. Durlauf., Elsevier, vol. 1, 679–741.

CHEN, C. (2017): "Untitled Land, Occupational Choice, and Agricultural Productivity," *American Economic Journal: Macroeconomics*, 9, 91–121.

COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural Change with Long-run Income and Price Effects," *Econometrica*, 89, 311 – 374.

CONTE, M., P. COTTERLAZ, AND T. MAYER (2022): "The CEPII Gravity Database," CEPII Working Paper 2022-05, CEPII, data accessed January 2023; 202211 version used. Downloaded from `https://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=8`.

CRAVINO, J. AND S. SOTELO (2019): "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, 11, 289–326.

de Janvry, A., K. Emerick, M. Gonzalez-Navarro, and E. Sadoulet (2015): "Delinking Land Rights from Land Use: Certification and Migration in Mexico," *American Economic Review*, 105, 3125–3149.

Diao, X., M. McMillan, and D. Rodrik (2019): "The Recent Growth Boom in Developing Economies: A Structural-Change Perspective," in *The Palgrave Handbook of Development Economics*, ed. by M. Nissanke and J. A. Ocampo, Palgrave Macmillan, 281–334.

Donovan, K. (2021): "The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity," *Review of Economic Studies*, 88, 2275–2307.

Donovan, K. and T. Schoellman (2023): "The role of labor market frictions in structural transformation," *Oxford Development Studies*, 51, 362–374.

Eaton, J. and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–79.

Eckert, F. and M. Peters (2022): "Spatial Structural Change," Working Paper 30489, NBER.

Fan, J. (2019): "Internal Geography, Labor Mobility, and the Distributional Impacts of Trade," *American Economic Journal: Macroeconomics*, 11, 252–288.

Farrokhi, F. and H. S. Pellegrina (2023): "Trade, Technology, and Agricultural Productivity," *Journal of Political Economy*, 131, 2509–2555.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015): "The Next Generation of the Penn World Table," *American Economic Review*, 105, 3150–3182, data accessed May 2022. Downloaded from `https://www.rug.nl/ggdc/productivity/pwt/pwt-releases/pwt100?lang=en`.

Feenstra, R. C. and J. Romalis (2014): "International Prices and Endogenous Quality," *The Quarterly Journal of Economics*, 129, 477–527, data accessed June 2022. Downloaded from `https://cid.ucdavis.edu/world-trade-flows` and `https://cid.ucdavis.edu/nber-united-nations-trade-data-1962-2000`.

Fields, G. S. (2005): "A Welfare Economic Analysis of Labor Market Policies in the Harris-Todaro Model," *Journal of Development Economics*, 76, 127–146.

GAI, Q., N. GUO, B. LI, Q. SHI, AND X. ZHU (2025): “Rural Pensions, Labor Reallocation, and Aggregate Income: An Empirical and Quantitative Analysis of China,” *Econometrica*, 93, 1663–1696.

GOLLIN, D., D. LAGAKOS, AND M. E. WAUGH (2014): “The Agricultural Productivity Gap,” *Quarterly Journal of Economics*, 129, 939–993.

GOLLIN, D., S. L. PARENTE, AND R. ROGERSON (2002): “The Role of Agriculture in Development,” *American Economic Review*, 92, 160–164.

GOLUB, S. S. AND C.-T. HSIEH (2000): “Classical Ricardian Theory of Comparative Advantage Revisited,” *Review of International Economics*, 8, 221–234.

GOTTLIEB, C. AND J. GROBOVŠEK (2019): “Communal Land and Agricultural Productivity,” *Journal of Development Economics*, 138, 135–152.

HAMORY, J., M. KLEEMANS, N. Y. LI, AND E. MIGUEL (2021): “Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata,” *Journal of the European Economic Association*, 19, 1522–1555.

HAO, T., R. SUN, T. TOMBE, AND X. ZHU (2020): “The Effect of Migration Policy on Growth, Structural Change, and Regional Inequality in China,” *Journal of Monetary Economics*, 113, 112–134.

HARRIS, J. R. AND M. P. TODARO (1970): “Migration, Unemployment and Development: A Two-Sector Analysis,” *American Economic Review*, 60, 126 – 142.

HEAD, K. AND T. MAYER (2014): “Gravity equations: Workhorse, Toolkit, and Cookbook,” in *Handbook of international economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 4, 131–195.

HERRENDORF, B., C. HERRINGTON, AND A. VALENTINYI (2015): “Sectoral Technology and Structural Transformation,” *American Economic Journal: Macroeconomics*, 7, 104–133.

HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103, 2752–2789.

HERRENDORF, B., R. ROGERSON, AND Á. VALENTINYI (2026): “Unconditional Convergence in Manufacturing: A Reassessment,” *American Economic Review: Insights*, forthcoming. Data provided directly by the authors.

HERRENDORF, B. AND T. SCHOELLMAN (2015): "Why is Measured Productivity so Low in Agriculture?" *Review of Economic Dynamics*, 18, 1003–1022.

——— (2018): "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 10, 1–23.

HOBIJN, B., T. SCHOELLMAN, AND A. VINDAS Q. (2018): "Structural Transformation by Cohort," Unpublished Working Paper.

HSIEH, C.-T. AND P. J. KLENOW (2007): "Relative Prices and Relative Prosperity," *American Economic Review*, 97, 562–585.

INKLAAR, R., R. MARAPIN, AND K. GRÄLER (2023): "Tradability and Sectoral Productivity Differences across Countries," GGDC Research Memorandum 195, Groningen Growth and Development Centre, data accessed May 2024. Downloaded from `https://www.rug.nl/ggdc/productivity/pld/releases/pld-2023?lang=en`.

JERVEN, M. (2013): *Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It*, Cornell University Press.

LAGAKOS, D. (2020): "Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities?" *Journal of Economic Perspectives*, 34, 174–192.

LAGAKOS, D., S. MARSHALL, A. M. MOBARAK, C. VERNOT, AND M. E. WAUGH (2020): "Migration Costs and Observational Returns to Migration in the Developing World," *Journal of Monetary Economics*, 113, 138–154.

LEVCHENKO, A. A. AND J. ZHANG (2016): "The Evolution of Comparative Advantage: Measurement and Welfare Implications," *Journal of Monetary Economics*, 78, 96–111.

LEWIS, A. W. (1954): "Economic Development with Unlimited Supplies of Labor," *The Manchester School*, 22, 139–191.

LEWIS, L. T., R. MONARCH, M. SPOSI, AND J. ZHANG (2022): "Structural Change and Global Trade," *Journal of the European Economic Association*, 20, 476–512.

LIU, C. AND X. MA (2023): "Migration, Tariffs, and China's Export Surge," *Journal of International Economics*, 140, 103696.

Lu, M. and Y. Xia (2016): "Migration in the People's Republic of China," Working Paper 593, Asian Development Bank Institute.

Lucas, R. E. (2021): *Crossing the Divide: Rural to Urban Migration in Developing Countries*, Oxford University Press.

Ma, X. (2024): "College Expansion, Trade, and Innovation: Evidence from China," *International Economic Review*, 65, 315–351.

MacDougall, G. D. A. (1951): "British and American Exports: A Study Suggested by the Theory of Comparative Costs. Part I," *The Economic Journal*, 61, 697–724.

Malmberg, H. (2025): "Skill Bias and Development: Evidence from Trade and Industry Data," Unpublished Working Paper.

Marshall, S. (2025): "Spatial Labor Market Power in Sub-Saharan Africa: The Roles of Self-Employment and Migration," Unpublished Working Paper.

Matsuyama, K. (2009): "Structural Change in an Interdependent World: A Global View of Manufacturing Decline," *Journal of the European Economic Association*, 7, 478–486.

McKenzie, D. (2024): "Fears and Tears: Should More People Be Moving within and from Developing Countries, and What Stops this Movement?" *The World Bank Research Observer*, 39, 75–96.

Minnesota Population Center (2020): "Integrated Public Use Microdata Series, International: Version 7.3," Dataset. Data accessed May 2022. Downloaded from https://international.ipums.org/international/.

Moscona, J. and K. Sastry (2025): "Inappropriate technology: Evidence from global agriculture," Working Paper 33500, NBER.

Ngai, L. R., C. A. Pissarides, and J. Wang (2019): "China's Mobility Barriers and Employment Allocations," *Journal of the European Economic Association*, 17, 1617–1653.

OECD (2022): "STAN Database for Structural Analysis," Dataset. Data accessed June 2022. Downloaded from https://data-explorer.oecd.org/vis?tm=DF_STAN_2025&snb=1&df[ds]=dsDisseminateFinalDMZ&df[id]=DSD_STAN%40DF_STAN_2025&df[ag]=OECD.STI.PIE&df[vs]=1.0&dq=A.FRA._T...&pd=2015%2C&to[TIME_PERIOD]=false&vw=tb. OECD now provides an updated 2025 version.

——— (2023): "Inter-Country Input-Output Tables, 2023 Release, Regular ICIO," Dataset. Data accessed April 2024. Downloaded from `https://www.oecd.org/en/data/datasets/inter-country-input-output-tables.html`.

PORZIO, T., F. ROSSI, AND G. SANTANGELO (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112, 2774–2814.

RESTUCCIA, D., D. T. YANG, AND X. ZHU (2008): "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis," *Journal of Monetary Economics*, 55, 234–250.

SILVA, J. M. C. S. AND S. TENREYRO (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

SIMONOVSKA, I. AND M. E. WAUGH (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92, 34–50.

STOKEY, N. L. (2001): "A Quantitative Model of the British Industrial Revolution, 1780–1850," *Carnegie-Rochester Conference Series on Public Policy*, 55, 55–109.

STORESLETTEN, K., B. ZHAO, AND F. ZILIBOTTI (2019): "Business Cycle During Structural Change: Arthur Lewis' Theory from a Neoclassical Perspective," Working Paper 26181, NBER.

TEIGNIER, M. (2018): "The Role of Trade in Structural Transformation," *Journal of Development Economics*, 130, 45–65.

THE CONFERENCE BOARD (2023): "EU KLEMS July 2018 Release," Dataset, Version 1.1. Data accessed September 2022. Downloaded from `https://doi.org/10.34894/6GDD7Q`.

TOMBE, T. (2015): "The Missing Food Problem: Trade, Agriculture, and International Productivity Differences," *American Economic Journal: Macroeconomics*, 7, 226–258.

TOMBE, T. AND X. ZHU (2019): "Trade, Migration, and Productivity: A Quantitative Analysis of China," *American Economic Review*, 109, 1843–1872.

UNITED NATIONS STATISTICS DIVISION (2020): "National Accounts Main Aggregates Database: Basic Data Selection," Dataset. Data accessed October 2020. Downloaded from `https://unstats.un.org/unsd/snaama/basic`.

UY, T., K.-M. YI, AND J. ZHANG (2013): "Structural Change in an Open Economy," *Journal of Monetary Economics*, 60, 667–682.

VOLLRATH, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity?" *Journal of Development Economics*, 88, 325–334.

——— (2014): "The Efficiency of Human Capital Allocations in Developing Countries," *Journal of Development Economics*, 108, 106–118.

WAUGH, M. E. (2010): "International Trade and Income Differences," *American Economic Review*, 100, 2093–2124.

WORLD TRADE ORGANIZATION (2020): "Composition, Definitions & Methodology," in *World Trade Statistical Review 2020*, World Trade Organization, chap. V, 60–75, data accessed June 2022 for sector classification on page 66. Downloaded from `https://www.wto.org/english/res_e/statis_e/wts2020_e/wts2020chapter05_e.pdf`.

YOUNG, A. (2013): "Inequality, the Urban-Rural Gap, and Migration," *Quarterly Journal of Economics*, 128, 1727–1785.