PHBS WORKING PAPER SERIES

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

Douglas Gollin
Tufts University and CEPR

David Lagakos
Boston University and NBER

Xiao Ma
Peking University

Shraddha Mandi
Boston University

May 2025

Working Paper 20250501

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 are measured independently from sectoral value added and employment data. 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. Consistent with the model's predictions, previous exports by sector are strong predictors of agricultural productivity gaps in the current cross-section of countries.

JEL Classification: E01, F11, O11, O41

Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055, China

The Dynamics of Agricultural Productivity Gaps: An Open-Economy Perspective
Douglas Gollin, David Lagakos, Xiao Ma, and Shraddha Mandi
NBER Working Paper No. 33529
February 2025
JEL No. E01, F11, O11, O41

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 are measured independently from sectoral value added and employment data. 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. Consistent with the model's predictions, previous exports by sector are strong predictors of agricultural productivity gaps in the current cross-section of countries.

Douglas Gollin
Tufts University
School of Arts and Sciences, Economics,
Ballou Hall,
Oxford OX1 3TB
Medford, MA 02155
and CEPR
Douglas.Gollin@tufts.edu

David Lagakos
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
lagakos@bu.edu

Xiao Ma
Peking University
HSBC Business School
xiaomaecon@gmail.com

Shraddha Mandi
Boston University
mandis@bu.edu

1 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, 2024). 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).

Another ubiquitous pattern in the data is the tendency for workers to move from agriculture to non-agricultural activities in the course of economic growth (see Herrendorf, Rogerson, and Valentinyi, 2013). One might expect that this sectoral reallocation would help close agricultural productivity gaps, as workers leave agriculture for manufacturing and service activities, in which their labor is more valuable. With diminishing marginal products of labor, reallocation should tend to drive an increase in the marginal (and hence average) product of labor in agriculture and a corresponding decline in non-agriculture.

In this paper, we examine changes in the agricultural productivity gaps across a large set of countries over the last four decades. To do this, we first construct a new panel data set that draws on national accounts data and population censuses from 68 countries of all income levels. Analyzing these data, we show that in the average country, the agricultural productivity gap shows no meaningful decline over since the 1980s, in spite of movements – often large – of workers out of agriculture. In fact, several prominent economies, including China and India, have experienced significantly widening gaps over time.

The remainder of the paper seeks to understand and explain the dynamics of agricultural productivity gaps. Our main premise is that an open-economy perspective is useful in understanding the dynamics of the gaps. In particular, our analysis draws on international trade data that are collected separately from our data on sectoral value added and employment. This allows us to calibrate our model to data that are effectively independent of our measurements of agricultural productivity gaps. We thus avoid directly targeting the productivity gaps that we seek to explain.

Our model is based largely 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 symmetric sector-specific iceberg trade costs. To capture income effects in the model, we posit PIGL preferences, with agriculture being an inferior good; this structure 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). We also allow the physical productivity of labor to grow at different rates in each country, sector, and year.

We calibrate the model using data on exports by sector, which helps discipline the model's physical productivity growth rates. 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). These imply decreasing trade costs over our time period in general. We target each country's productivity levels by sector and year to match fixed-effects estimates in gravity equations similar to Eaton and Kortum (2002). We assume frictional movement of labor and intentionally take an abstract view on migration costs. In particular, we assume that a randomly selected 2 percent of the population can change sectors each period, if they wish to, while the rest of the population cannot move. This parsimonious choice is transparent and also 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 0.77. The main test of the model is whether its predicted *changes* in agricultural productivity gaps – which are not targeted in the calibration in any way – are comparable to those in the data. In general, we find that they are. Comparing the first and last years for each country, the correlation between log changes in gaps in the model and data is 0.86. 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 in spite of substantial changes in sectoral employment.

To understand how the model matches the data so well, we conduct a series of counterfactual experiments. The first allows for physical productivity growth by sector but holds fixed the allocation of workers across sectors and also holds constant trade costs over time. Interestingly, 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 (30 log points in the model vs. -7 log points in the data). The lesson is that heterogeneous real productivity growth rates by sectors are crucial for understanding the variation in agricultural productivity gap dynamics, but *ceteris paribus*, sectoral productivity growth by itself would have tended to raise gaps rather than to keep them the same or lower them.

In a second counterfactual, we shut off real productivity growth by sector but allow worker movements and non-constant trade costs. Now, 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 poorly explains the variation in gap dynamics between countries, with only a 0.19 correlation between model and data. The reason is that workers move out of agriculture in a much more uniform way across countries, with the majority of countries experiencing declines of similar magnitude; this movement can thus not match the observed patterns, in which some countries see rising agricultural productivity gaps and other countries see the gaps falling.

In a third counterfactual we replace the model's non-homothetic preferences with simple Cobb-Douglas preferences, thereby shutting off income effects. The model fit becomes substantially worse in this case, with gaps declining in most countries in the model. In a fourth and final counterfactual we assume that each country is in autarky over the entire period, though with the same productivity trends as before. We show that this model leads to significant gaps between the data and model in many countries as well. The takeaway from these two counterfactuals is that non-homothetic preferences and frictional trade are important to match the data – but real productivity growth patterns by sector are still the most important force in explaining the cross-country heterogeneity.

Our model predicts that *past* real productivity trends are a key driver of agricultural productivity gaps. To test this, we examined how past productivity trends correlate with current agricultural productivity gaps across countries. In particular, we conduct simple descriptive regressions of log agricultural productivity gaps on a host of observable country characteristics, including cumulative growth of agricultural and non-agricultural exports. The results align with our model: past agricultural export growth predicts lower current gaps, while past non-agricultural growth predicts higher ones. These relationships persist across different controls and significantly improve the model fit beyond GDP per capita controls.

Our findings support the view that agricultural productivity gaps are created by differences in physical productivity levels and growth rates across countries, in conjunction with frictions that limit sectoral movements of labor, rather than faulty measurement. An implication of our study is that, because countries experience differential growth rates of sectoral productivity, there is no reason to assume that agricultural productivity gaps will necessarily disappear as labor reallocates across sectors.

This paper contributes to a long line of work 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 (Tombe, 2015; Betts, Giri, and Verma, 2017; Lewis, Monarch, Sposi, and Zhang, 2022; Farrokhi and Pellegrina, 2023). See Alessandria, Johnson, and Yi (2023) for a recent overview of this literature. 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, 2022; Boppart, Kiernan, Krusell, and Malmberg, 2023). Our model also abstracts from the role of sectoral linkages and multinationals in structural change (Alviarez, Chen, Pandalai-Nayar, Varela, Yi, and Zhang, 2022; Sposi, 2019).

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; Nakamura, Sigurdsson, and Steinsson, 2022; Donovan and Schoellman, 2023). Relatively few countries have explicit policy barriers to internal migration, with China 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 very high, and even the costs of switching sectors within locations may be high, as in recent work in India by Baysan, Dar, Emerick, Li, and Sadoulet (2024). Limited information among potential migrations can act as a barrier to rural-urban migration (Baseler, 2023), as can frictions in land markets, especially those that make it harder for potential migrants to sell or secure rural land (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 (2023) argues that barriers to migration can interact with labor market power, further lowering wages in rural areas where firms are scarce. 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 (2023) for recent summaries of the literature on rural-urban migration.

2 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, particularly Gollin et al. (2014), 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 follow.

2.1 Data

The agricultural productivity gap (APG) is defined 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 agricultural productivity gap 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 (Ruggles, Cleveland, Lovatón Dávila, Sarkar, Sobek, Burk, Ehrlich, Heimann, and Lee, 2024). We only use country-years for which we can compute sectoral labor 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 small differences described in the appendix.

Sectoral value added comes from the United Nations National Accounts Main Aggregates Database. 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. 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. Economic activity was defined broadly to include paid, unpaid, and in-kind productive labor, excluding housework. 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 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 contributed 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.

[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, provided that this approach did not compromise the internal consistency of the data set. 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.

Figure 1: Change in Agricultural Productivity Gaps: Earliest vs Latest Years

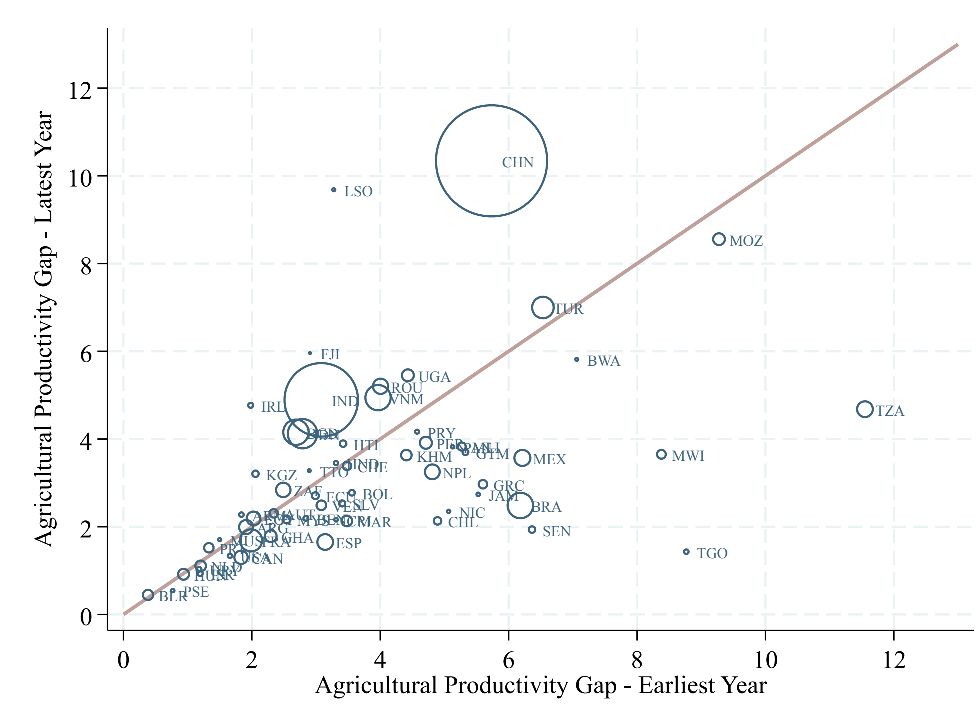

2.2 Dynamics of Gaps: Empirical Facts

Figure 1 plots APGs from earliest to latest available years per country, where the size of the dot represents the population of the country[2]. Many countries cluster near the 45-degree line, meaning similar gaps in the first and last years. Of 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 is puzzling given the widespread reallocation of labor outside of agriculture. 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.[3] Yet the countries

[2]Figure 1 excludes 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). Their agricultural productivity gaps are too large, so they are excluded for visual clarity.

[3]These are Lesotho, pictured in Q1, and Zambia, with a change in agricultural productivity gap of 9.4.

Figure 2: Changes in APGs and Agricultural Shares of Employment

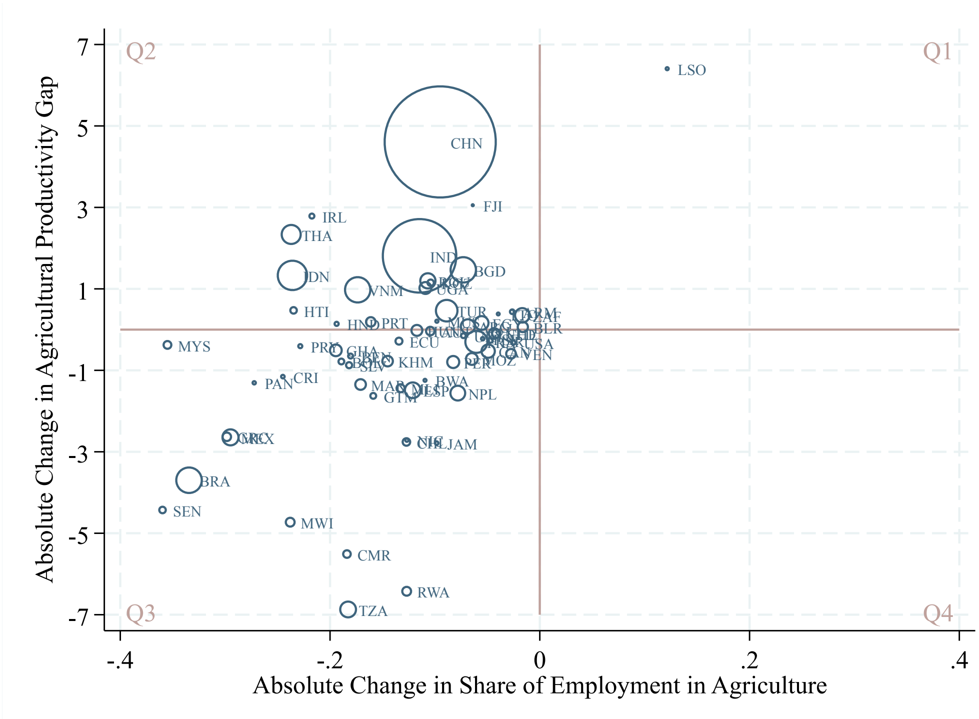

are not more likely to be in the bottom left quadrant (Q3), in which the APG also declined, than in the top left one (Q2), where the APG rose. In fact, around one-third of the countries, and just over 80 percent of the population, are in Q2, with increasing APGs despite declining shares of agricultural employment.

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 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. (2024) 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 agricultural productivity gaps. Bangladesh sees an increase in their gap while Brazil saw a decrease. Other countries are relatively stagnant over time. A.2 emphasizes that the difference between the agricultural productivity gap with and without adjustments for human capital are quite similar. Human capital adjustments make a minimal difference to the agricultural productivity gap for most countries. See Appendix A.[4]

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 declines. The goal of our quantitative analysis is to understand these patterns. One explanation, always lurking when we look at data of this kind, is that the observed variation across countries and over time simply reflects messy data. Simple noise or measurement error could plausibly generate patterns like this. Our conjecture, however, is that the patterns we observe are reflective of meaningful variation over time and across countries. It is striking that the APGs for China and India – both of which experienced growth in non-agricultural exports during the period – rose sharply. In what follows, we explore this conjecture further by introducing a model in which APGs and sectoral exports are determined by patterns of real sectoral productivity growth, in conjunction with frictions that limit the movement of workers across space or sector.

[4]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. Details are available on request. We did not systematically analyze changes in hours worked per sector since historical data on hours worked by sector are not available for enough countries.

Figure 3: Agriculture Productivity Gap Trends for Select Countries

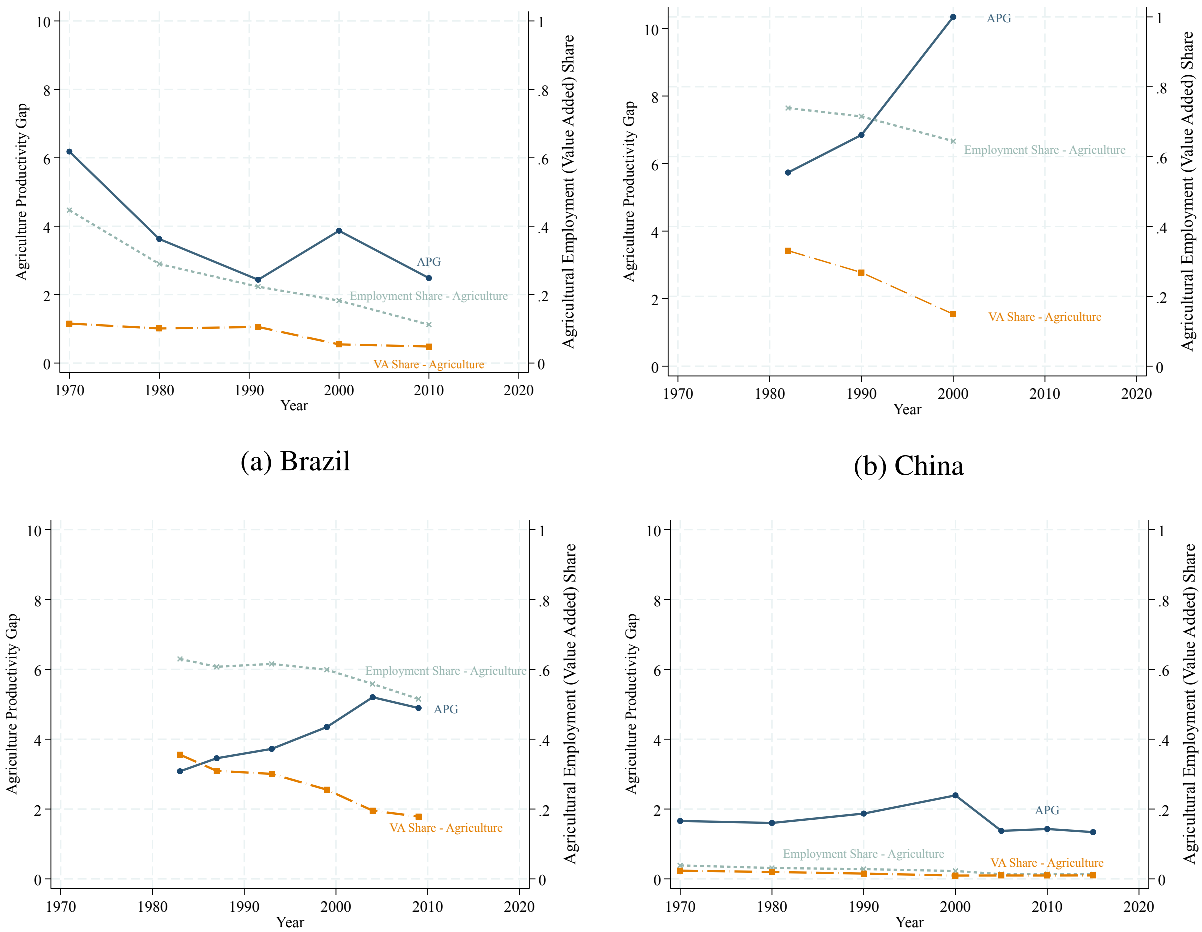

(a) Brazil

(b) China

(c) India

(d) United States

3 Quantitative Model

We now develop a multi-country model to understand the drivers of agricultural productivity gaps over time and across countries. The world has many countries, $i = 1,...,I$. Each country i hosts both agricultural and nonagricultural 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.

3.1 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 fundamental 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 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 enjoyed cheaper production costs (lower wage 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:

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

3.2 Household Preferences

In country i, there is a measure l_i of households. Building on the existing literature (e.g. Boppart, 2014), we adopt the assumption 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 nonagricultural 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 how the expenditure share that varies with real income levels. 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 $\rho = 1+\zeta\frac{(\vartheta_i-\phi)^2}{\vartheta_i(1-\vartheta_i)}$. 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.2, 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.

3.3 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 nonagriculture (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 $l_i^n = l_i(1-\lambda_i^a)$.

3.4 Equilibrium

In the 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 agricultural productivity gap in country i is given by the ratio of value added per worker in nonagriculture 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 in general not be equal to one. According to equation (11), the agricultural productivity gaps 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 $\{\phi, v, \zeta\}$, 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 4 to understand the drivers of agricultural productivity gaps.

Finally, given a set of parameter values $\{z_i^j, d_{ik}^j, \eta, \phi, v, \zeta, \lambda_i^{a,0}, \kappa, l_i\}$, 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 by applying the algorithm of Alvarez and Lucas (2007).

4 Calibration

We now use our model to quantitatively understand the drivers of agricultural productivity gaps. Given the 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 this section, we first describe the data sources and then discuss how we calibrate the model parameters.

4.1 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 Table A.1. These closely, but not completely, overlap with the countries presented earlier due to data availability issues.

Trade Data. For international trade data between 1980–2015, we utilize the United Nations Comtrade Database, which offers detailed information on bilateral trade flows. 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.

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 interpolate 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 agriculture and non-agriculture sectors for each country.

Because many countries' data are not available in the APG data, 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.[5]

[5] For some countries, there are available data for fewer years, and thus we interpolate the data using agricultural shares in GDP similarly as we did for the APG list.

4.2 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 levels and trade openness levels (e.g., China, India). We take each year's employment l_{it} directly from the data. All other parameters are time-invariant. We now describe our procedure to calibrate all the parameters.

4.2.1 Calibrating Productivity Levels and Trade Costs

We estimate sectoral fundamental productivity levels and trade costs using bilateral trade data. To begin with, from the trade value in equation (7), we can obtain the relative trade values (Eaton and Kortum, 2002):

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

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 the 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 capitals of country k and country i. $contig_{kit}$ is a dummy variable that indicates whether two countries are contiguous. Finally, we also 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:

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

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).[6] We estimate this equation for each sector and year.

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

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

Recovering productivity levels. To recover estimates of fundamental productivity levels z_{it}^{j} from $\hat{S}_{it}^{j}$, we need information on wage rates w_{it}^{j}. It is worth noting that wage rates are endogenous variables in our model: if we apply a set of country-sector-year productivity levels $\{z_{kt}^{j}\}$ in the model, we can endogenously solve wage levels $\{w_{kt}^{j}\}$ from market clearing conditions in equation (11) and obtain the model-predicted level of $\{S_{kt}^{j,model}\}$. Given this observation, we therefore search for the set of productivity levels $\{z_{kt}^{j}\}$ in the model that minimize the summed absolute difference between the model-predicted "competitiveness" level $\{S_{it}^{j,model}\}$ and the model estimates $\{\hat{S}_{it}^{j}\}$, $\sum_t \sum_i \sum_j |\hat{S}_{it}^{j} - S_{it}^{j,model}|$. Furthermore, since we estimate S_{it}^{j} 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 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 appear to be reasonable. First, in line with existing research that observes faster productivity growth in agriculture (e.g., Huneeus and Rogerson, 2024), we find that agricultural productivity growth tends to surpass non-

[6] 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_{kt}^{j} can differ by a constant, which will not affect the quantitative results.

Figure 4: Productivity Growth (Relative to Year 1980)

(a) Brazil

(b) China

(c) India

(d) United States

agricultural productivity growth on average. Across 77 countries and the Rest of World, the average annual productivity growth is 1.7 percent in agriculture and 0.9 percent in non-agriculture during the 1980–2015 period. Additionally, as shown by Figure 4, Bangladesh, China and India, with higher non-agricultural productivity growth than agricultural productivity growth, exhibit significant increases in agricultural productivity gaps based on our evidence in Section 2. Finally, as we recovered productivity growth mainly from trade data, in Appendix Figure B.1, we show that Bangladesh, China, and India enjoyed faster export growth in non-agriculture than agriculture during the 1980–2015 period.

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

4.2.3 Calibrating Sectoral Labor Adjustments

Finally, we choose the proportion of workers who have the opportunity to switch sectors, $\kappa = 0.02$, according to the average change in the share of agricultural employment in our countries and years. We set the agricultural employment shares in the initial year, 1980, to directly match the data.

5 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 agricultural productivity gaps. Finally, we show that our model-estimated productivity growth aligns well with the empirical data.

To evaluate how well our calibrated model aligns with the data, we start by comparing the levels of agricultural productivity gaps between the model and the actual data, in logs, for all country years in our sample. These were not targeted in our calibration procedure. As shown in Figure 5, the agricultural productivity gap predicted by our model exhibits a strong positive correlation with the observed gaps, with a correlation coefficient of 0.73. The correlation is 0.77 if we focus on the initial level (the first observation) of the agricultural productivity gaps for each country.

As our focus is on the dynamics of agricultural productivity gaps, Figure 6 shows the actual dynamics alongside the model-generated dynamics of these productivity gaps, which were not also not targeted in our calibration.[7] Detailed numbers can be found in Appendix Table B.1. In the

[7]Figure B.2 shows that the model-generated changes in agricultural employment shares align well with the data. Our model tends to overestimate changes in agricultural employment (weighted by employment), as there were slow sectoral movements in some populous countries such as China and India, consistent with the large migration barriers

Figure 5: Levels of Agricultural Productivity Gaps: Model vs Data

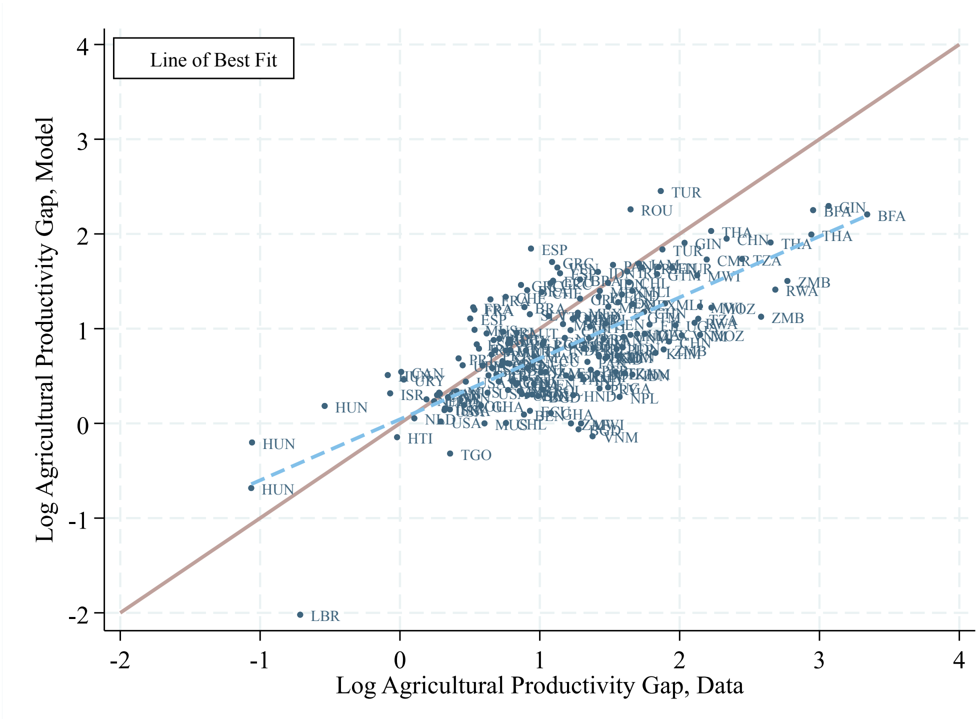

data, we take the difference between the first observation and the last observation for each country with at least two years of data observed during the 1980–2015 period. In the model, for each country, we compute the changes in agricultural productivity gaps using the same time window as in the data. Based on this graph, it appears that our model can partially explain the fluctuations in agricultural productivity gaps over time, as demonstrated by the positive correlation between the model's projected changes and the observed data, with a slope coefficient of 1.14. Figure 7 illustrates the evolution of agricultural productivity gaps in both the model and the actual data for the four countries analyzed in Figure 3. The graph continues to suggest that our model partially captures the temporal patterns observed in these countries. Finally, Figure B.4 shows that our model fit of the dynamics of agricultural productivity gaps is good for both more open and less open countries.

found in those countries (e.g, Ngai et al., 2019; Tombe and Zhu, 2019). In Appendix Figure B.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. Our analysis reveals that this extended model closely matches the observed changes in agricultural employment. However, predictions on the dynamics of agricultural productivity gaps remain largely unaffected.

Figure 6: Dynamics of Agricultural Productivity Gaps: Model vs Data

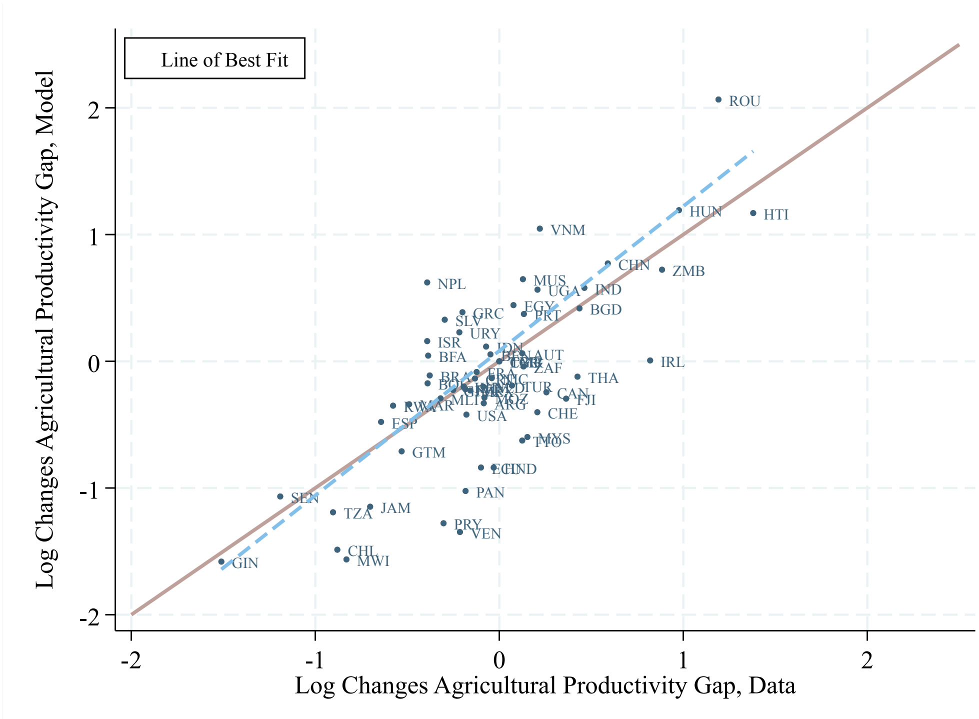

5.1 Drivers of Dynamics in Agricultural Productivity Gaps

Our model encapsulates varying trade costs, productivity levels, and labor market frictions (i.e., barriers to sectoral mobility), all of which contribute to the evolution of agricultural productivity gaps. Without these ingredients (and without time variation in at least some of these elements), the gaps in agricultural productivity would have remained static. In order 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 a specific channel depends on the introduction of changes in other channels within the model. Therefore, we take the approach of Hao, Sun, Tombe, and Zhu (2020) to assess the resulting changes in agricultural productivity gaps for a channel by averaging its effects conditional on all possible combinations of the other two channels.

Role of Productivity Growth. First, we examine the influence of sectoral productivity growth by comparing the calibrated model with and without variations in productivity levels z_{it}^{j} over time. We observe that, after 1980, the agricultural productivity gaps rose for most countries, as evidenced by panel (b) in Figure 8. This increase in APGs can be attributed to a decrease in the relative demand for agricultural goods as income grows. Thus, the evidence indicates that productivity

Figure 7: Model Fit in Four Select Economies

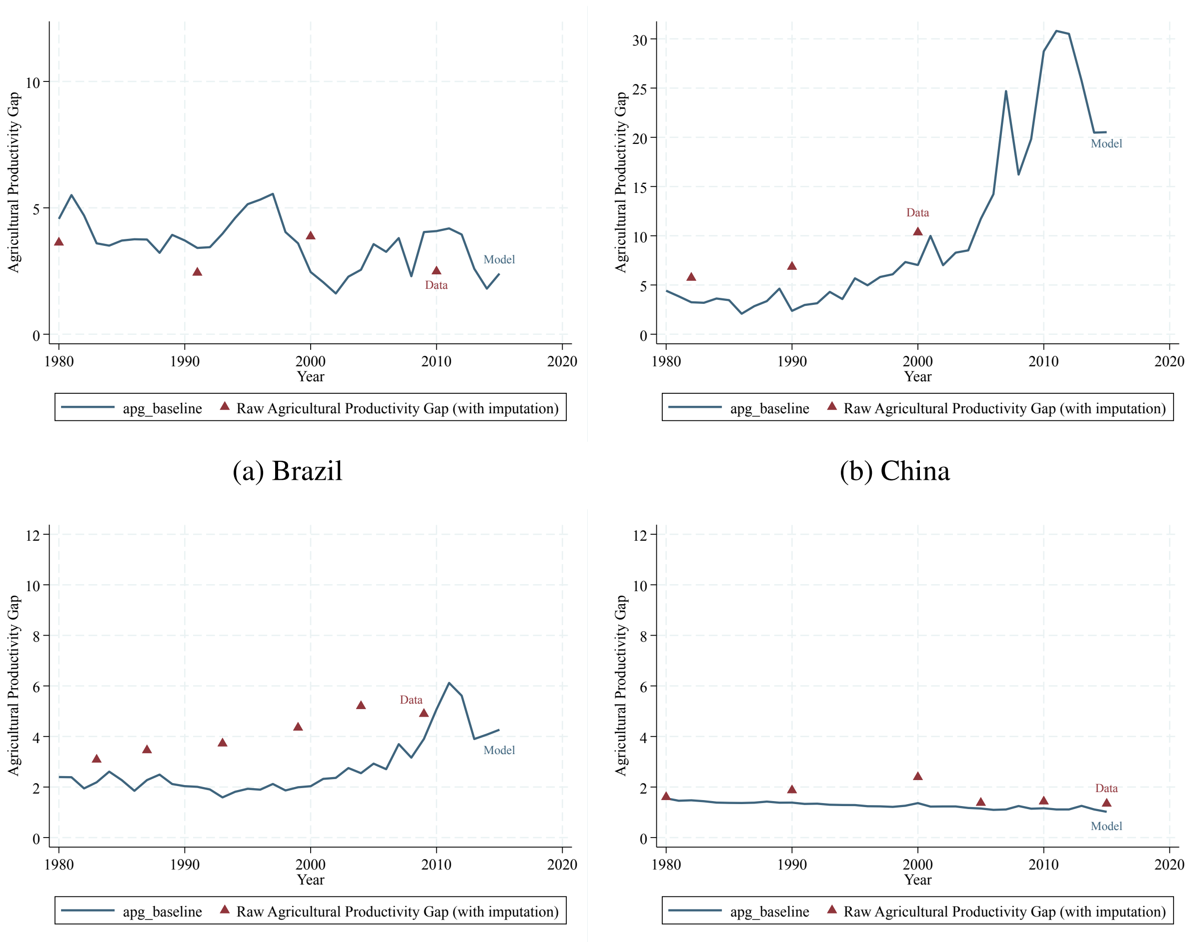

(a) Brazil

(b) China

(c) India

(d) United States

growth has served as the main catalyst for the escalation of agricultural productivity gaps in these countries. However, by itself, productivity growth in the model tends to over-predict the increase in agricultural productivity gaps.

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 explore the influence of changing trade costs by comparing the calibrated model with and without alterations in trade costs d_{kit}^{j} after 1980. As highlighted in panel (c) of Figure 8, we observe that the impact of trade cost changes on agricultural productivity gaps tends to be relatively insignificant compared to the role played by productivity growth. Further-

Figure 8: Dynamics of Agricultural Productivity Gaps in Counterfactual Scenarios

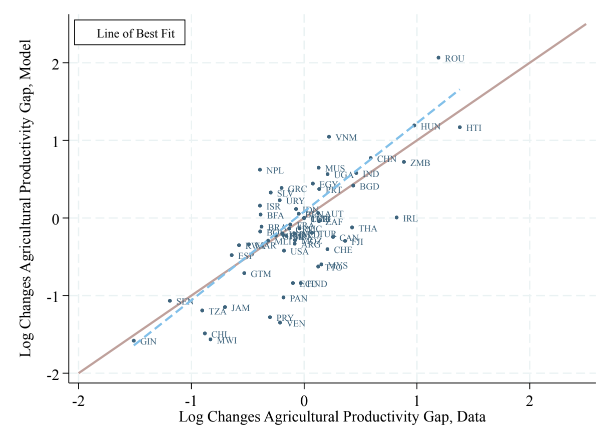

(a) Baseline Model

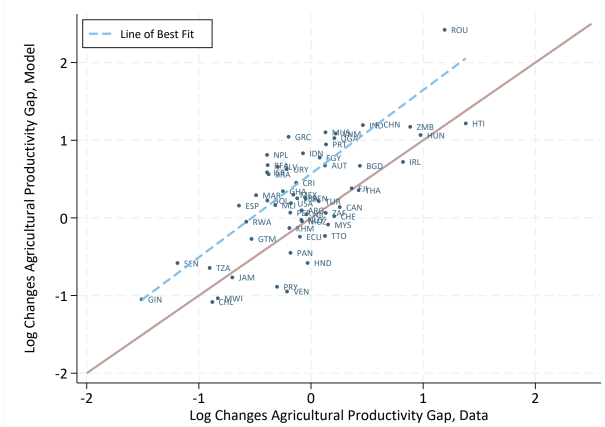

(b) Only Productivity Growth

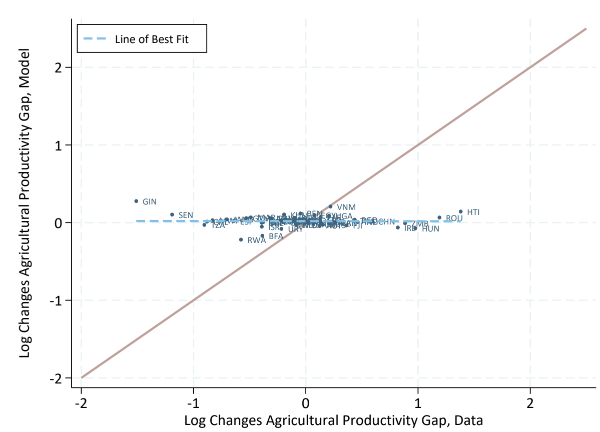

(c) Only Changes in Trade Costs

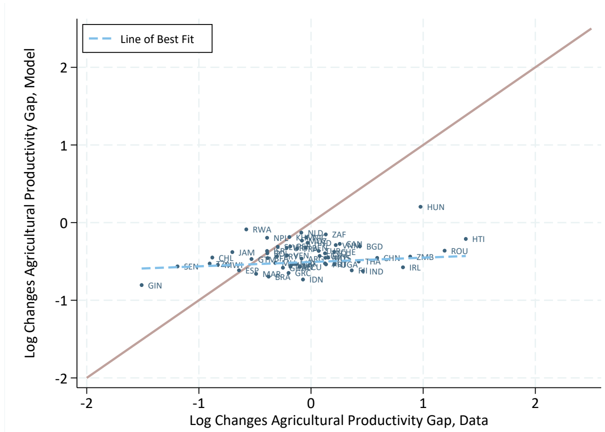

(d) Only Sectoral Labor Movements

more, this impact varies across countries, reflecting the heterogeneity of trade liberalization across sectors and nations.

Role of Labor Sectoral Adjustments. Lastly, we demonstrate the effects of barriers to the movement of labor across sectors by comparing the calibrated model with and without allowing workers to transition between sectors. As workers move from agriculture to non-agriculture, driven by relatively higher wages in the non-agricultural sector, these sectoral movements of labor tend to mitigate the agricultural productivity gaps in the majority countries, as shown in panel (D) of Figure 8.

Decomposition of Dynamics into Relative Price and Productivity Effects. According to equation (12), our model has a simple decomposition of agricultural productivity gaps:

$$APG = \frac{w_i^n}{w_i^a} = \underbrace{\frac{w_i^n/z_i^n}{w_i^a/z_i^a}}_{\text{relative price of nonagriculture}} \times \underbrace{\frac{z_i^n}{z_i^a}}_{\text{relative productivity of nonagriculture}} .$$

As we consider perfect competition, marginal costs equal prices for each sector. Thus, the first component in this decomposition, $\frac{w_i^n/z_i^n}{w_i^a/z_i^a}$, captures relative prices of nonagricultural versus agricultural goods. In our model, the relative prices are determined by the demand and supply of goods in each sector, which are affected by all of our three channels (productivity growth, trade, and barriers to labor movements). The second component, $\frac{z_i^n}{z_i^a}$, captures the relative productivity of non-agricultural versus agricultural goods, which directly relates to the productivity channel we studied.

To understand the relative contributions of price changes and productivity growth to the dynamics of agricultural productivity gaps, we perform the following regression using the model-generated data:

$$y_i = \beta_0 + \beta_1 \Delta \log APG_i + \varepsilon_i,$$

where $\Delta \log APG_i$ is the model-generated changes in agricultural productivity gaps, using the same time window for each country as in Figure 6. We separately use changes in log relative prices and changes in log relative productivity as the dependent variable. In this specification, the coefficients from these two regressions shall add up to 1.

Table 1: Relative Contributions of Price Changes and Productivity Growth

Dep var	log rel prices (1) unweighted	log rel productivity (2) unweighted	log rel prices (3) weighted	log rel productivity (4) weighted
log APG	-0.135** (0.058)	1.135*** (0.058)	0.027 (0.069)	0.973*** (0.069)
Obs	59	59	59	59
R-squared	0.098	0.885	0.005	0.857

Notes: Regressions in Columns (1)–(2) are unweighted, whereas regressions in Columns (3)–(4) are weighted by employment in each country in the initial year of observation. Robust standard errors are in parentheses. Significant levels: * 10%, ** 5%, *** 1%.

Table 1 shows that relative productivity growth between non-agriculture and agriculture is the primary driver of the dynamics of agricultural productivity gaps. It is worth noting that this decomposition is based on model estimates. Due to the lack of long time series on both prices and productivity, we are unable to directly perform the decomposition with the observed data. However, relying on the GGDC Productivity Level Database, which provides data on relative prices based on the 2011 International Comparisons Program (ICP), we observe that the model-generated relative price of agriculture aligns reasonably well with the data, as shown in Figure B.5. We will also show that the model's estimates for sectoral productivity growth align well with the data in Section 5.3.

5.2 Model Validation: Cross-Sectional Variation in Gaps

In order to validate the model's predictions, we examine whether the agricultural productivity gaps in the most recent cross section of countries are correlated with past sector export growth in a way consistent with the model.

Table 2 presents a regression analysis of the observed agricultural productivity gap for the most recent year of observation in each country versus past export growth since 1980. The results in column (1) show that higher agricultural export growth is linked to smaller observed agricultural productivity gaps, while higher non-agricultural export growth is linked to larger gaps. This is just as the model predicts. In column (2), we regress the observed gap on GDP per capita. Although GDP per capita accounts for some of the variation in gaps between countries, its explanatory power, as measured by R-squared, is less than that of export growth in column (1). Column (3) includes both export growth and GDP per capita in the regression, showing that the coefficients for export growth remain significant and only modest changed from column (1). Columns (4) and (5) integrate climate and institutional variables into the regression, following Gollin et al. (2014). These additional covariates do not significantly alter the coefficients for export growth. Finally, Column (6) demonstrates that the explanatory power of the regression for gaps decreases when export growth is excluded.

5.3 Sectoral Productivity Growth in Model and Data

The model's predictions for real productivity growth by sector play a central role in explaining the dynamics of APGs. For a subset of our countries, we can compare our model's predictions to direct measures of real sectoral productivity growth. We make these comparisons in two data sets. The first is the EU KLEMS database, which estimates the labor productivity (based on quantities

Table 2: Correlations of Observed Agricultural Productivity Gaps and Country Characteristics

	Log (Observed Agricultural Productivity Gap)					
	(1)	(2)	(3)	(4)	(5)	(6)
Past agricultural export growth	-8.26**		-7.24**	-7.04**	-7.05*	
	(4.15)		(3.18)	(3.41)	(3.82)	
Past nonagricultural export growth	15.84***		11.43**	9.25***	10.39***	
	(2.46)		(2.54)	(2.98)	(3.32)	
Log (GDP per capita)		-0.59***	-0.31***	-0.34***	-0.23***	-0.44***
		(0.11)	(0.07)	(0.09)	(0.08)	(0.11)
Ruggedness				1.09	0.65	1.26
				(0.84)	(0.87)	(1.16)
Fertile soil				-0.54**	-0.24	-0.45
				(0.25)	(0.47)	(0.69)
Tropical climate				0.09	-0.19	-0.02
				(0.28)	(0.37)	(0.55)
Rule of law index					-0.91	-0.20
					(0.73)	(1.15)
Restrictions on domestic movement					-0.02	0.33*
					(0.22)	(0.20)
Ethnic fractionalization					0.35	-0.18
					(0.36)	(0.34)
Observations	62	62	62	56	56	56
Adjusted R-squared	0.70	0.59	0.80	0.82	0.82	0.75

Note: The dependent variable is the observed agricultural productivity gap in the most recent year of observation (during 1980–2015) for each country. The independent variables include the observed annualized export growth from 1980 to the year of observation, as well as climate- and institution-related variables. We weight the regression by the observed employment. Robust standard errors are in parentheses. Significant levels: * 10%, ** 5%, *** 1%.

of value added per hour worked) for many European countries and the United States; these data are available for the 1975–2015 period. The second data set is from Herrendorf, Rogerson, and Ákos Valentinyi (2022) (HRV2022 hereafter), who estimate the labor productivity (based 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 during 1980–2015; in the model, we compute the productivity growth using the same time window as in the data.

Figure 9 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 the data estimate, for both data sets. The correlation coefficient between

Figure 9: Comparison of Productivity Growth between Data and Model Estimates

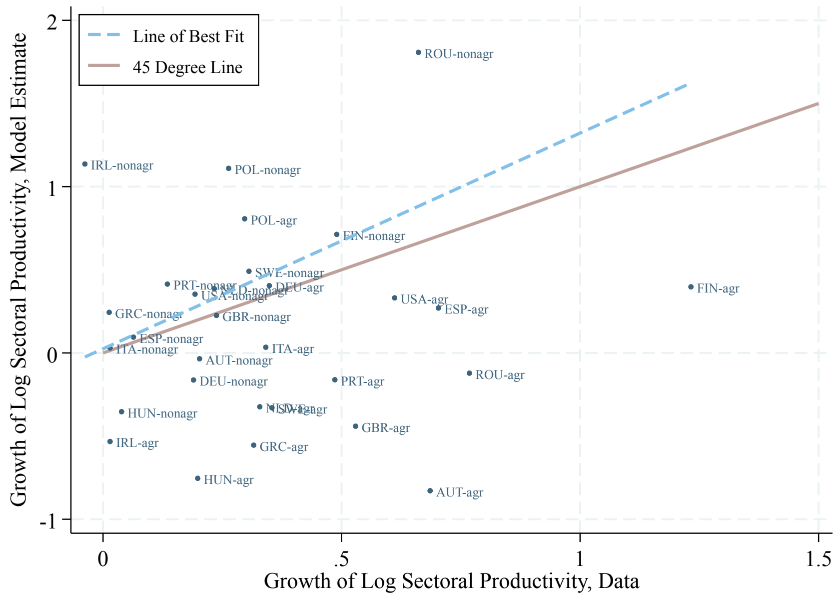

(a) EU KLEMS

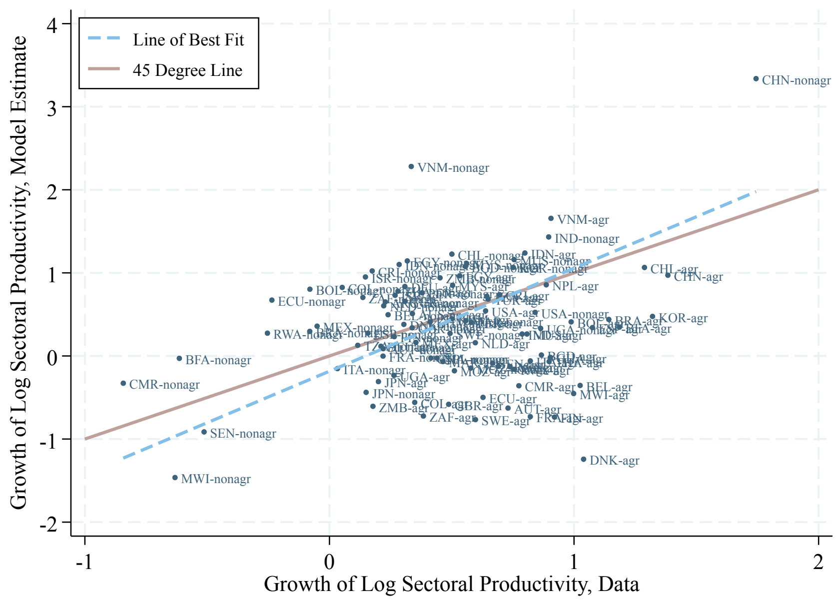

(b) Data from HRV2022

model and data is 0.47 using the data estimate from EU KLEMS, and it is 0.59 using the data estimate from Herrendorf et al. (2022). This indicates that our model estimates of sectoral productivity growth capture at least a portion of the actual sectoral productivity growth. In Figure B.6, we further plot the comparison of growth in relative productivity (agricultural productivity divided by nonagricultural productivity) between the data and the model. We find that there is always a positive correlation between the model estimates and the data estimates of relative productivity growth.

5.4 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 agricultural productivity gap. To evaluate this shift in relative sectoral demand, we now assume instead that the utility function is homothetic and Cobb-Douglas:

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

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 agricultural productivity gaps (now calculated using homogeneous preferences) compared with the baseline results. We find that for the majority of points, the model-projected changes in agricultural productivity gaps are smaller under homogeneous preferences than under non-homothetic preferences. For many countries, the model even predicts declines in agricultural productivity gaps where positive changes are observed in the data (Figure B.7). This underscores the role of non-homothetic preferences (particularly income effects) in exacerbating agricultural productivity gaps by diverting demand away from agriculture.

5.5 The Role of International Trade

Although Figure 8 demonstrated that changes in trade costs do not significantly influence the dynamics of agricultural productivity gaps, 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 agricultural productivity gaps, 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 stan-

Figure 10: Comparison of Dynamics of Agricultural Productivity Gaps 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).

dard deviation of 0.16 on a logarithmic scale. This suggests that by considering a closed economy rather than an open one, our predictions for changes in agricultural productivity gaps 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.8, we illustrate that the degree of underprediction is more pronounced in countries with higher relative productivity growth in agriculture. For instance, Tanzania 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 agricultural productivity gap in the country.

6 Conclusion

Development economists have long been interested in the determinants of sectoral gaps in output per worker (see e.g. Lewis, 1954; Harris and Todaro, 1970; Fields, 2005). Recent research has highlighted the substantial differences in value added per worker between non-agriculture and agriculture (Vollrath, 2009; Gollin et al., 2014; Herrendorf and Schoellman, 2015). This paper draws on new evidence to show that these gaps have remained remarkably persistent on average over the past decades, despite massive shifts in the share of workers in agriculture. The gaps have even widened in several major economies, including some that have seen a dramatic exodus of labor from agriculture.

We argue that an open-economy perspective helps to understand the dynamics of agricultural productivity gaps. Our quantitative open-economy model explains the dynamics of productivity gaps quite well, from China's widening gap amid rapid non-agricultural export growth to Brazil's narrowing gap as a major agricultural exporter. More generally, our model predicts that countries experiencing faster growth of non-agricultural (agricultural) exports will end up with larger (smaller) productivity gap. These patterns are borne out in our data for the period 1980 to 2015.

Our findings strongly suggest that the observed patterns of APGs across countries and over time are not simply the product of noisy data. Instead, standard economic forces seem to be shaping productivity gaps, in particular differential growth rates of physical productivity by sector, barriers to labor mobility, and changing trade costs. While questions will always persist about the reliability of national statistics in low-income countries (Jerven, 2013), our study has the advantage that we quantify physical sectoral productivity growth using trade data – which are collected independently from the data used for measuring APGs, and which are generally thought to reflect with some accuracy the country-level patterns of comparative advantage. Our analysis supports the view that agricultural productivity gaps, both in levels and changes over time, reflect real forces. The sluggish pace of labor migration out of agriculture, which sustains the productivity 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.

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, 153–73.

ALVIAREZ, V. I., C. CHEN, N. PANDALAI-NAYAR, L. VARELA, K.-M. YI, AND H. ZHANG (2022): "Multinationals and Structural Transformation," .

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

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," Unpublished Working Paper, University of Minnesota.

BRYAN, G. AND M. MORTEN (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127, 2229–68.

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," *NBER Working Paper No. 21768*.

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.

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.

CONTE, M., P. COTTERLAZ, AND T. MAYER (2022): "The CEPII Gravity Database," *CEPII Working Paper N°2022-05*.

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

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*, Palgrave Macmillan.

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*.

FAN, J. (2019): "Internal Geography, Labor Mobility, and the Distributional Impacts of Trade," *American Economic Journal: Macroeconomics*, 11, 252–88.

FARROKHI, F. AND H. S. PELLEGRINA (2023): "Trade, Technology, and Agricultural Productivity," *Journal of Political Economy*, 131.

FEENSTRA, R. AND J. ROMALIS (2014): "International Prices and Endogenous Quality," *The Quarterly Journal of Economics*, 129.

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 (2024): "Rural Pensions, Labor Reallocation, and Aggregate Income: An Empirical and Quantitative Analysis of China," Unpublished Working Paper, University of Hong Kong.

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 Papers and Proceedings*, 92.

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.

HEAD, K. AND T. MAYER (2014): "Gravity Equations: Workhorse, Toolkit, and Cookbook," *Handbook of International Economics*, 4.

HERRENDORF, B., C. HERRINGTON, AND VALENTINYI (2015): "Sectoral Technology and Structural Transformation," *American Economic Journal: Macroeconomics*, 7, 104–33.

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

HERRENDORF, B., R. ROGERSON, AND ÁKOS VALENTINYI (2022): "New Evidence on Sectoral Labor Productivity: Implications for Industrialization and Development," *NBER Working Paper 29834*.

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. (2024): "Structural Transformation by Cohort," Unpublished Working Paper, Federal Reserve Bank of Minneapolis.

Huneeus, F. and R. Rogerson (2024): "Heterogeneous Paths of Industrialization," *Review of Economic Studies*, 91, 1746–1774.

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, 139–91.

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.

Lu, M. and Y. Xia (2016): "Migration in the People's Republic of China," Tech. Rep. 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.

Marshall, S. (2023): "Spatial Labor Market Power in Sub-Saharan Africa: The Roles of Self-Employment and Migration," Unpublished Working Paper, University of Warwick.

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. (2023): "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.

Moscona, J. and K. Sastry (2022): "Inappropriate technology: Evidence from global agriculture," *Available at SSRN 3886019*.

Nakamura, E., J. Sigurdsson, and J. Steinsson (2022): "The Gift of Moving: Intergenerational Consequences of a Mobility Shock," *Review of Economic Studies*, 89, 1557–1592.

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.

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.

Ruggles, S., L. L. Cleveland, R. Lovatón Dávila, S. Sarkar, M. Sobek, D. Burk, D. E. Ehrlich, Q. Heimann, and J. Lee (2024): "IPUMS International: Version 7.5," .

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.

Sposi, M. (2019): "Evolving Comparative Advantage, Sectoral Linkages, and Structural Change," *Journal of Monetary Economics*, 103, 75–87.

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," NBER Working Paper 26181.

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

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.

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*, 2093–2124.

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

A Data Appendix

Table A.1: Countries and years with APG data

Country	Years of Data		
	APG	Model	HCAPG
Argentina	1970 1980 2001	1980 2001	1970 1980 2001
Armenia	2001 2011		
Austria	1971 1981 1991 2001 2011	1981 1991 2001 2011	
Bangladesh	1991 2001 2011	1991 2001 2011	1991 2001 2011
Belarus	1999 2009		
Benin	1979 1992 2002 2013	1992 2002 2013	
Bolivia	1976 1992 2001 2012	1992 2001 2012	
Botswana	1991 2001 2011		
Brazil	1970 1980 1991 2000 2010	1980 1991 2000 2010	1970 1980 1991 2000
Burkina Faso	1985 1996	1985 1996	
Cambodia	1998 2004 2008 2013	1998 2004 2008 2013	1998 2004 2008 2013
Cameroon	1976 2005	2005	1976 2005
Canada	1971 1981 1991 2001 2011	1981 1991 2001 2011	
Chile	1970 1982 1992 2002	1982 1992 2002	1970 1982 1992 2002
China	1982 1990 2000	1982 1990 2000	
Costa Rica	1973 1984 2000 2011	1984 2000 2011	1973 1984 2000 2011
Ecuador	1982 1990 2001 2010	1982 1990 2001 2010	1982 1990 2001 2010
Egypt	1986 1996 2006	1986 1996 2006	
El Salvador	1992 2007	1992 2007	1992 2007
Fiji	1976 1986 2014	1986 2014	
France	1975 1982 1990 1999	1982 1990 1999	
Ghana	1984 2000 2010	1984 2000 2010	1984 2000 2010
Greece	1971 1981 1991 2001 2011	1981 1991 2001 2011	
Guatemala	1973 1981 1994 2002	1981 1994 2002	1973 1981 1994 2002
Guinea	1983 1996 2014	1983 1996 2014	1983 1996 2014

Table A.1 Continued: Countries and years with APG data

Haiti	1971 1982 2003	1982 2003	
Honduras	1974 1988 2001	1988 2001	1974 1988 2001
Hungary	1970 1980 1990 2001 2011	1980 1990 2001 2011	
India	1983 1987 1993 1999 2004 2009	1983 1987 1993 1999 2004 2009	
Indonesia	1971 1976 1980 1985 1990 2000 2010	1980 1985 1990 2000 2010	
Ireland	1971 1981 1986 1991 1996 2002 2006 2011 2016	1981 1986 1991 1996 2002 2006 2011	
Israel	1972 1983 1995 2008	1983 1995 2008	
Jamaica	1982 1991 2001	1982 1991 2001	
Kyrgyz Republic	1999 2009		
Lesotho	1996 2006		
Liberia	1974 2008	2008	1974 2008
Malawi	1987 1998 2008	1987 1998 2008	1987 1998 2008
Malaysia	1970 1980 1991 2000	1980 1991 2000	1980 1991
Mali	1987 1998 2009	1987 1998 2009	
Mauritius	1990 2000 2011	1990 2000 2011	1990 2000 2011
Mexico	1970 1990 1995 2000 2010 2015	1990 1995 2000 2010 2015	1970 1990 1995 2000 2010 2015
Morocco	1982 1994 2004 2014	1982 1994 2004 2014	
Mozambique	1997 2007	1997 2007	
Nepal	2001 2011	2001 2011	2001 2011
Netherlands	1971 2001 2011	2001 2011	
Nicaragua	1971 1995 2005	1995 2005	
Palestine	1997 2007 2017		1997 2017
Panama	1970 1980 1990 2000 2010	1980 1990 2000 2010	1970 1980 1990 2000 2010
Paraguay	1972 1982 1992 2002	1982 1992 2002	1972 1982 1992 2002
Peru	1993 2007	1993 2007	1993 2007
Portugal	1981 1991 2001 2011	1981 1991 2001 2011	
Romania	1977 1992 2002 2011	1992 2002 2011	
Rwanda	2002 2012	2002 2012	2002 2012

Table A.1 Continued: Countries and years with APG data

Senegal	1988 2002 2013	1988 2002 2013	1988 2002 2013
South Africa	1996 2001 2007	1996 2001 2007	1996 2001 2007
Spain	1981 1991 2001 2011	1981 1991 2001 2011	
Switzerland	1970 1980 1990 2000	1980 1990 2000	
Tanzania	1988 2002 2012	1988 2002 2012	1988 2002 2012
Thailand	1970 1980 1990 2000	1980 1990 2000	1970 1980 1990 2000
Togo	1970 2010	2010	
Trinidad And Tobago	1980 1990 2000	1980 1990 2000	1980 1990 2000
Turkey	1985 1990 2000	1985 1990 2000	
Uganda	1991 2002 2014	1991 2002 2014	1991 2002 2014
United States	1970 1980 1990 2000 2005 2010 2015	1980 1990 2000 2005 2010 2015	
Uruguay	1975 1985 1996 2006	1985 1996 2006	
Venezuela	1981 1990 2001	1981 1990 2001	1981 1990 2001
Vietnam	1989 1999 2009	1989 1999 2009	
Zambia	1990 2000 2010	1990 2000 2010	1990 2000 2010

Note: The "Model" column refers to data points where we have both empirical and the estimated APG. The model was calibrated for countries that have the full set of calibration information and all years from 1980 to 2015 for countries where the full set of calibration data is available.

A.1 Construction of Variables using IPUMS

To determine whether an individual is engaged in agricultural economic activity, we use primarily the harmonized employment and industry variables provided by IPUMS. In cases where IPUMS does not provide a harmonized employment variable for a particular survey, we refer to the original survey variables to establish the employment status. If such variables are absent or inadequate, we consider an individual to be employed if their industry is known.

However, there are instances where the harmonized variable for industry is either unavailable or has significant missing values for employed individuals. Given the requirement that the industry must be known for inclusion in the analysis, we make several efforts to impute or derive the industry information using alternative variables. The process is as follows:

1. **Attempt to use original, non-harmonized industry variables**: If these are available in the

original survey but not harmonized by IPUMS, we use them to assign an industry classification.

2. **Use occupation variables for imputation**: If the original industry variable is not available, we use occupation variables to impute industry data.

3. **Exclusion of surveys**: If neither industry nor occupation data is available or usable for imputation, the survey is excluded from the analysis.

Additionally, in cases where the industry is unknown for more than 4 percent of employed individuals, we proceed with imputing the missing values using occupation variables, provided they are available. We do not impute industry values using occupation if such imputation would introduce inconsistencies within a country's dataset, or if using occupation would not sufficiently resolve the issue of missing industry data.

Imputing years of schooling can be challenging, as inconsistencies in the underlying data can arise. We rely on the harmonized years of schooling variable available in IPUMS for consistency, provided that we are confident that the imputation of schooling years is consistent across time and countries. This means we only use the variable if we believe that IPUMS' imputation methods do not overestimate or underestimate years of schooling for a biased subsample of the population. Of the 45 countries for which years of schooling is made available by IPUMS, we calculate average years of schooling by agriculture and non-agriculture for 32 countries.

A.2 Human Capital Adjustment

The analysis presented in equation 12 does not account for differences in human capital per worker, which is a critical factor influencing productivity across sectors. Sectoral employment shares alone do not capture this dimension of productivity variation. To address this gap, we calculate a *Human Capital Adjusted Agricultural Productivity Gap* (HCAPG). Human capital is measured using years of schooling, with the assumption that the return to a year of schooling is consistent across countries, set at 10 percent based on standard values for the Mincerian returns to schooling.

The Human Capital Gap (HCG) is defined as:

$$HCG \equiv \frac{exp(0.1 * \frac{S_n}{L_n})}{exp(0.1 * \frac{S_a}{L_a})} \tag{17}$$

where S_n and S_a represent total years of schooling in non-agriculture and agriculture sectors, respectively, and L_n and L_a are the respective labor forces in these sectors.

Figure A.1: Change in HCAPG: Earliest vs Latest Years

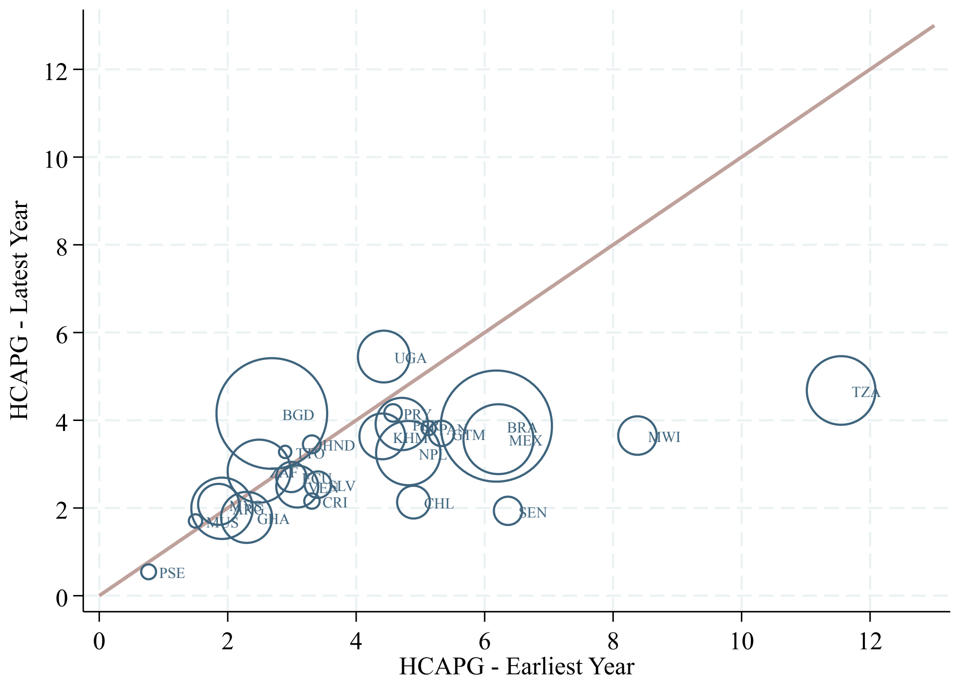

The Human Capital Adjusted Agricultural Productivity Gap (HCAPG) is then calculated by dividing the original APG by the human capital gap:

$$\text{HCAPG} = \frac{APG}{HCG} \tag{18}$$

We calculate the HCAPG for 32 countries, where information on education levels is available. As we can see in Figure A.2, the HCAPG makes a minimal difference to the APG for most country-year data points. Brazil and Cambodia are the only countries where the HCAPG is more than one unit lower than the APG. Figure A.2 shows that we still have an interesting mix of APG dynamics with some countries, like Brazil, seeing a decline in their APG while others, like Bangladesh, See an increase in their APG. We also continue to see a clustering of countries with stagnant Agricultural Productivity Gaps.

Figure A.2: Agricultural Productivity Gap vs HCAPG

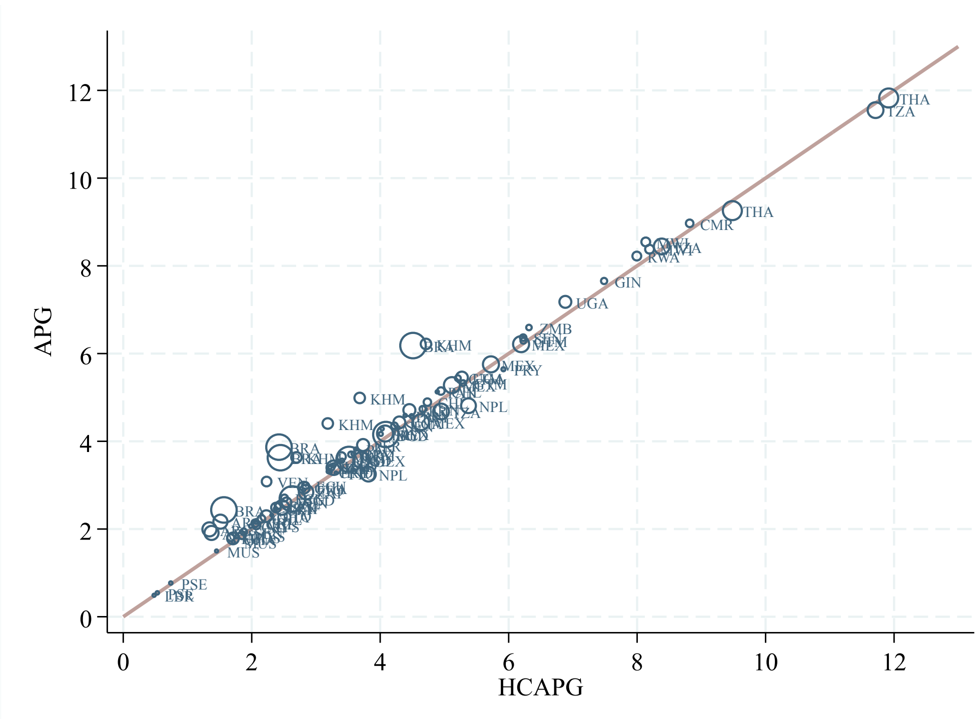

B Additional Tables and Graphs

Table B.1: Log Changes in Agricultural Productivity Gaps, Model versus Data

		Model-generated Changes			
Country	Data Change	Baseline	Productivity Growth	Trade Costs	Labor Sectoral Move
ARG	-0.08	-0.33	0.10	0.03	-0.46
AUT	0.12	0.06	0.67	-0.04	-0.53
BGD	0.44	0.42	0.67	0.04	-0.31
BEN	-0.05	0.05	0.25	0.12	-0.32
BOL	-0.39	-0.17	0.22	0.00	-0.39
BRA	-0.38	-0.11	0.56	0.01	-0.70
BFA	-0.39	0.04	0.68	-0.17	-0.46
KHM	-0.19	-0.20	-0.13	0.10	-0.19
CAN	0.26	-0.24	0.14	0.00	-0.27
CHL	-0.88	-1.49	-1.08	0.01	-0.45
CHN	0.59	0.77	1.21	0.01	-0.45
CRI	-0.13	-0.14	0.45	0.02	-0.53
ECU	-0.10	-0.84	-0.24	0.06	-0.57
EGY	0.08	0.44	0.78	0.09	-0.43
SLV	-0.30	0.33	0.66	0.02	-0.31
FJI	0.36	-0.29	0.38	-0.03	-0.61
FRA	-0.12	-0.08	0.25	0.02	-0.33
GHA	-0.25	-0.23	0.34	0.02	-0.58
GRC	-0.20	0.39	1.04	0.01	-0.65
GTM	-0.53	-0.71	-0.27	0.06	-0.47
GIN	-1.51	-1.58	-1.05	0.28	-0.80
HTI	1.38	1.17	1.22	0.14	-0.21
HND	-0.03	-0.84	-0.58	0.03	-0.26
HUN	0.98	1.19	1.07	-0.07	0.20
IND	0.46	0.58	1.20	0.01	-0.62
IDN	-0.07	0.12	0.83	0.01	-0.73
IRL	0.82	0.01	0.72	-0.06	-0.58
ISR	-0.39	0.16	0.59	-0.05	-0.37
JAM	-0.70	-1.15	-0.77	0.04	-0.38
MWI	-0.83	-1.56	-1.04	0.03	-0.54
MYS	0.15	-0.60	-0.09	-0.03	-0.45
MLI	-0.32	-0.29	0.17	0.06	-0.52
MUS	0.13	0.65	1.10	0.01	-0.45
MEX	-0.16	-0.23	0.30	0.02	-0.52

Table B.1 Continued: Changes in Agricultural Productivity Gaps

MAR	-0.49	-0.34	0.29	0.07	-0.66
MOZ	-0.08	-0.29	-0.04	-0.03	-0.23
NPL	-0.39	0.62	0.81	0.00	-0.20
NLD	-0.09	-0.20	-0.02	-0.03	-0.13
NIC	-0.04	-0.13	0.05	0.04	-0.19
PAN	-0.18	-1.02	-0.45	0.04	-0.57
PRY	-0.30	-1.28	-0.89	0.06	-0.43
PER	-0.18	-0.22	0.07	0.02	-0.31
PRT	0.13	0.37	0.95	0.00	-0.54
ROU	1.19	2.07	2.42	0.07	-0.36
RWA	-0.58	-0.35	-0.05	-0.22	-0.09
SEN	-1.19	-1.07	-0.58	0.10	-0.56
ZAF	0.13	-0.04	0.06	0.07	-0.15
ESP	-0.64	-0.48	0.16	0.02	-0.61
CHE	0.21	-0.40	0.02	-0.01	-0.38
TZA	-0.90	-1.19	-0.64	-0.03	-0.53
THA	0.43	-0.12	0.36	0.02	-0.50
TTO	0.12	-0.63	-0.23	0.04	-0.40
TUR	0.07	-0.19	0.22	-0.02	-0.36
UGA	0.21	0.57	1.03	0.09	-0.54
USA	-0.18	-0.42	0.19	0.03	-0.55
URY	-0.22	0.23	0.63	-0.08	-0.32
VEN	-0.21	-1.35	-0.95	0.03	-0.42
VNM	0.22	1.05	1.09	0.21	-0.29
ZMB	0.88	0.72	1.17	-0.01	-0.44
Mean	-0.07	-0.15	0.27	0.02	-0.42
Slope		1.14	1.08	0.00	0.06
Slope if data change≥0		1.42	1.44	-0.06	0.03
Slope if data change<0		1.08	1.05	-0.01	0.04
Correlation		0.86	0.83	-0.01	0.16
Correlation if data change≥0		0.72	0.73	-0.26	0.04
Correlation if data change<0		0.63	0.55	-0.05	0.06

Notes: The slope and correlation are calculated using each country's initial employment as weights.

Figure B.1: Export Patterns across Countries

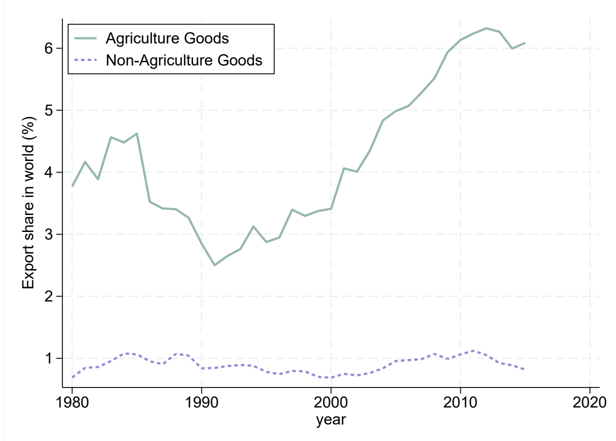

(a) Brazil

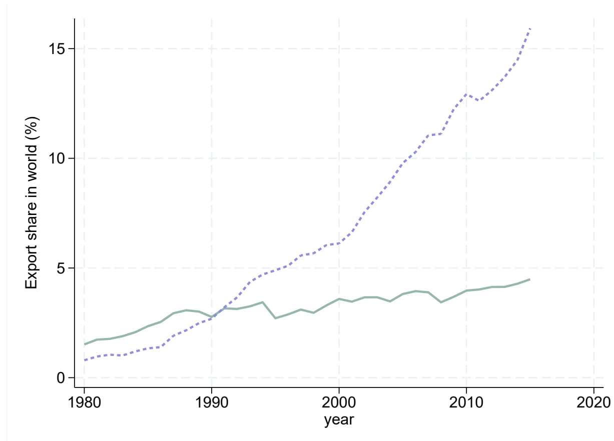

(b) China

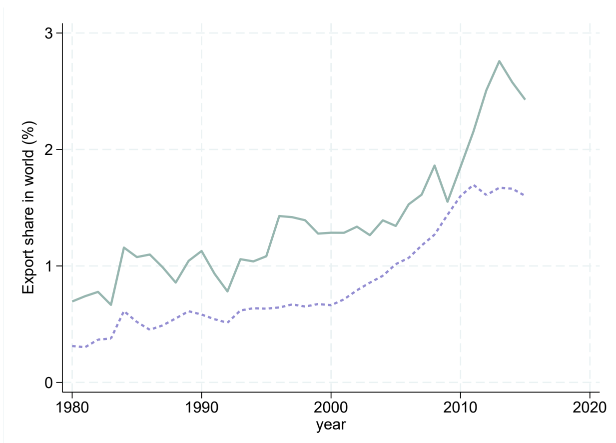

(c) India

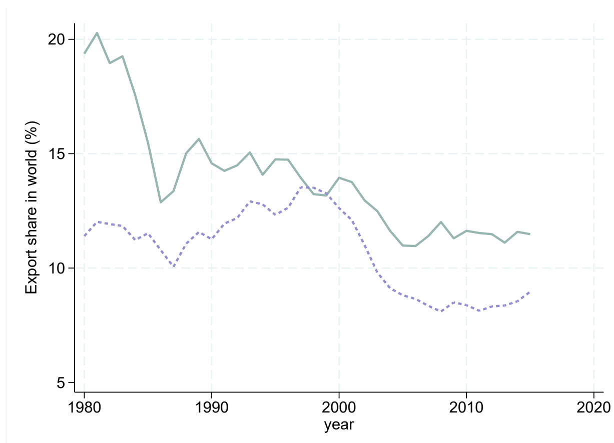

(d) United States

Note: The graphs show the country's exports as a percentage of the global total for each respective year and sector.

Figure B.2: Changes in Agricultural Employment Shares

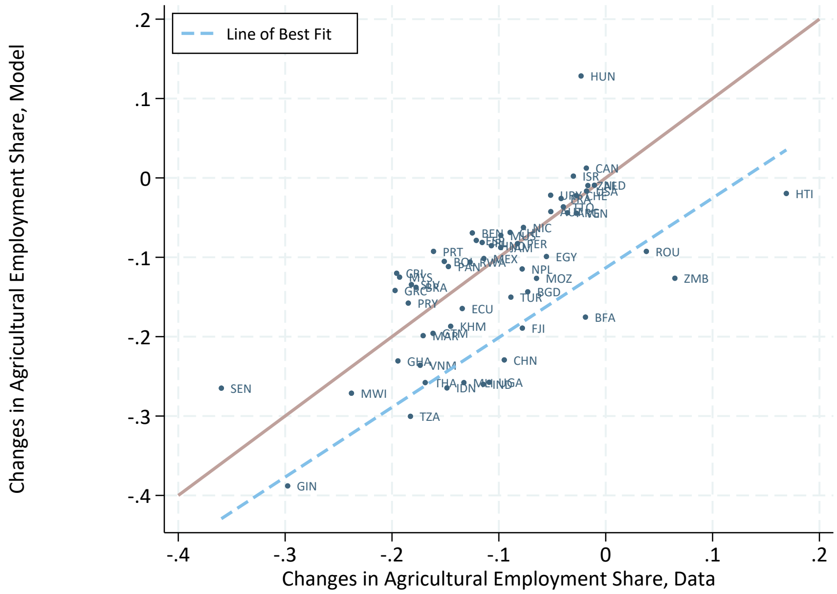

Figure B.3: Quantitative Results with Country-Specific Opportunities for Sector Switching

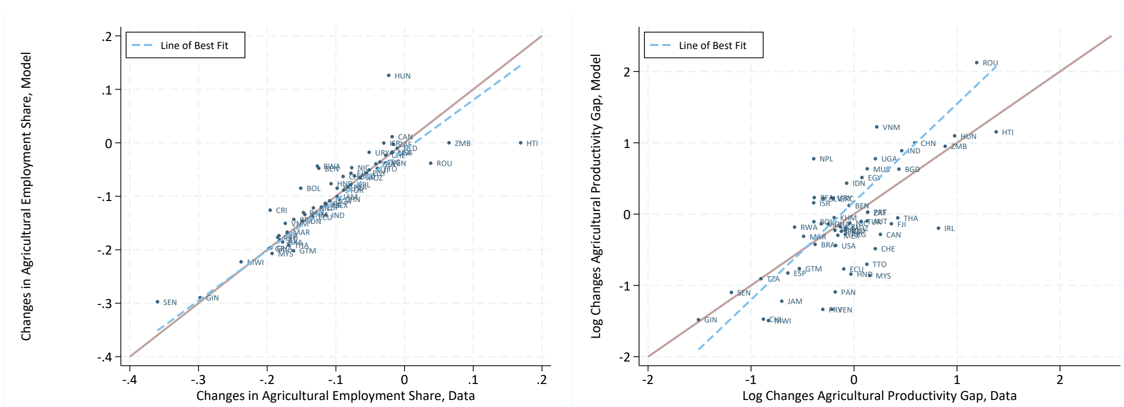

(a) Changes in Agricultural Employment Shares (b) Dynamics of Agricultural Productivity Gaps

Notes: We utilize a country-specific proportion of workers with opportunities to switch sectors, based on the average proportional change in agricultural employment shares over the years for each country. All other parameters remain unchanged from the baseline model. The regression is weighted by each country's initial employment.

Figure B.4: Dynamics of Agricultural Productivity Gaps: Model vs Data

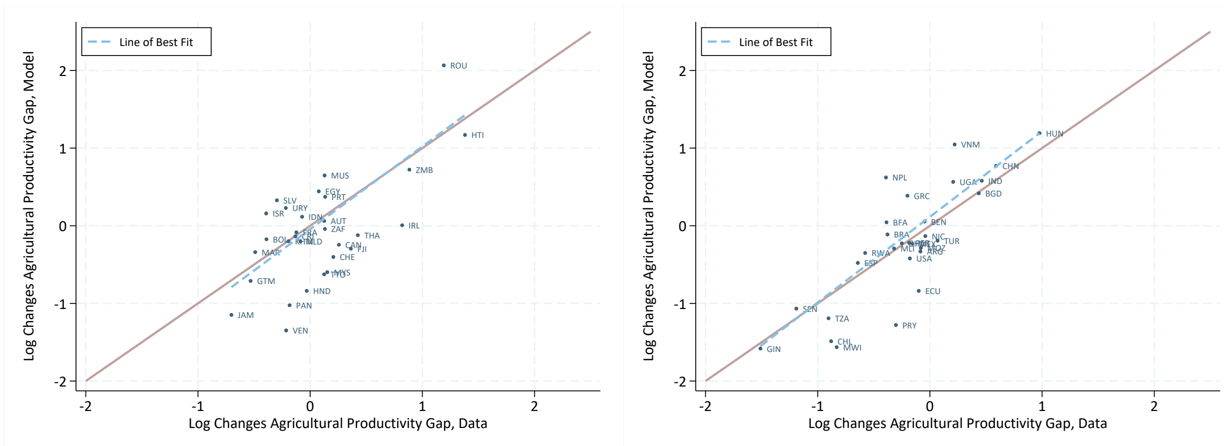

(a) More Open Countries

(b) Less Open Countries

Note: We measure the openness for each country by $\frac{\text{Imports+Exports}}{\text{GDP}}$ for the first year of observation in the 1980–2015 period. Countries are categorized as either more open or less open based on whether their openness measure surpasses the median.

Figure B.5: Relative Prices of Agriculture across Countries

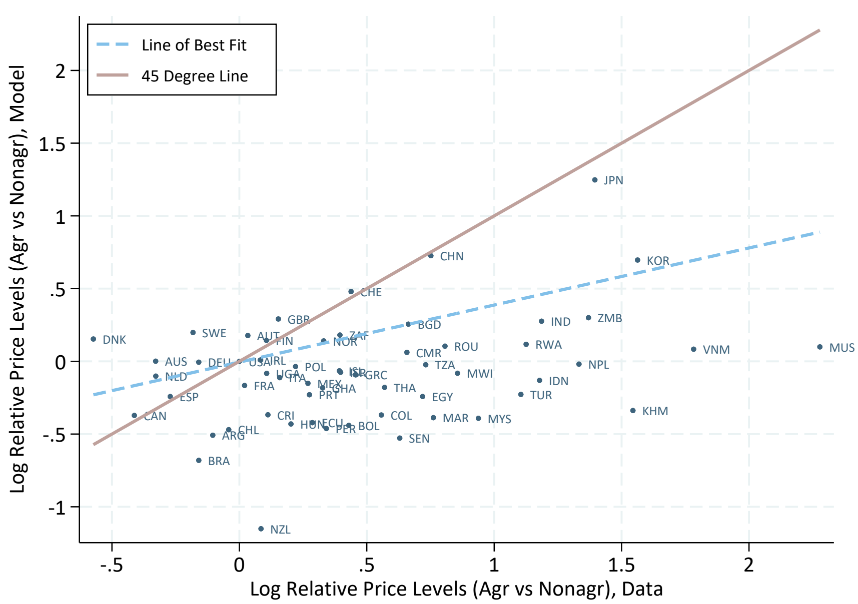

Note: We rely on the GGDC Productivity Level Database, which presents data on relative prices and labor productivity across 84 countries and 12 sectors, for 2011. These data are largely based on the results of the 2011 International Comparisons Program (ICP). In the data, the price index of the US is normalized to 1 in each year and sector. For ease of comparison, we perform the same normalization in each sector and year for the model-generated data.

Figure B.6: Comparison of Relative Productivity Growth between Data and Model Estimates

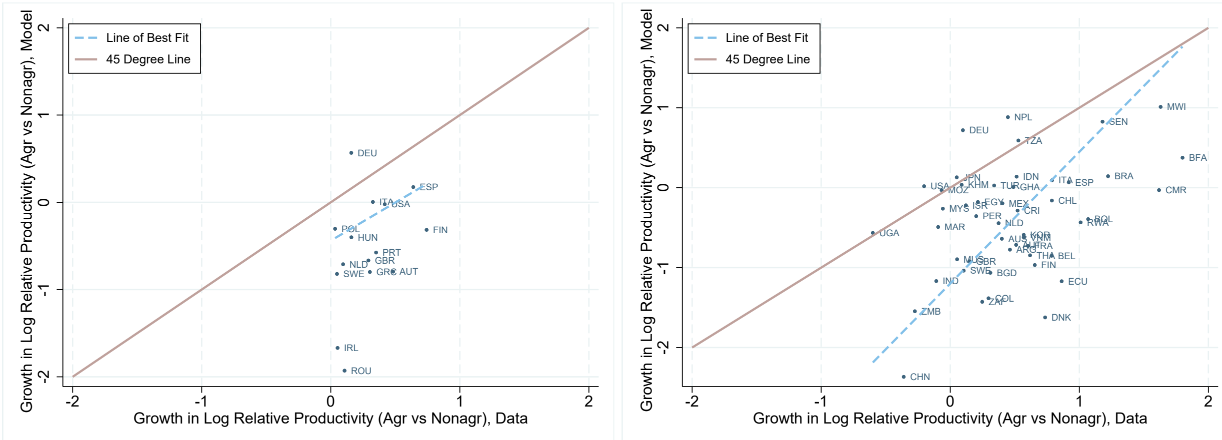

(a) EU KLEMS

(b) Data from HRV2022

Note: For each country, in the data, we compute the relative productivity growth between the first observation and the last observation during 1980–2015; in the model, we compute the productivity growth using the same time window as in the data.

Figure B.7: Changes in Agricultural Productivity Gaps, with Cobb-Douglas Preferences

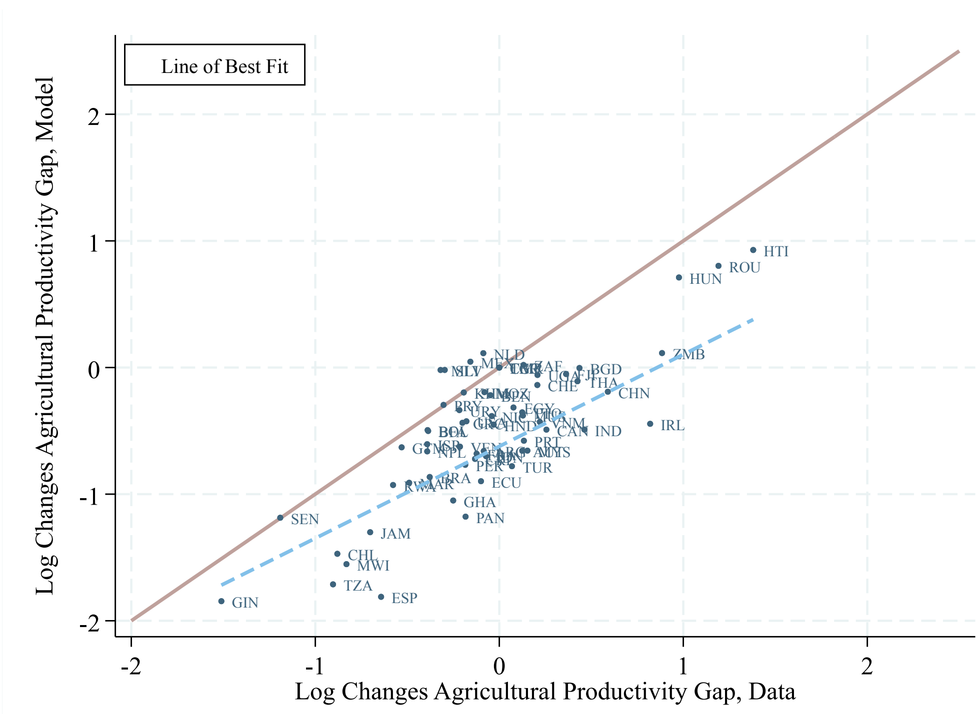

Figure B.8: Dynamics of Agricultural Productivity Gaps, Autarky versus Baseline Model

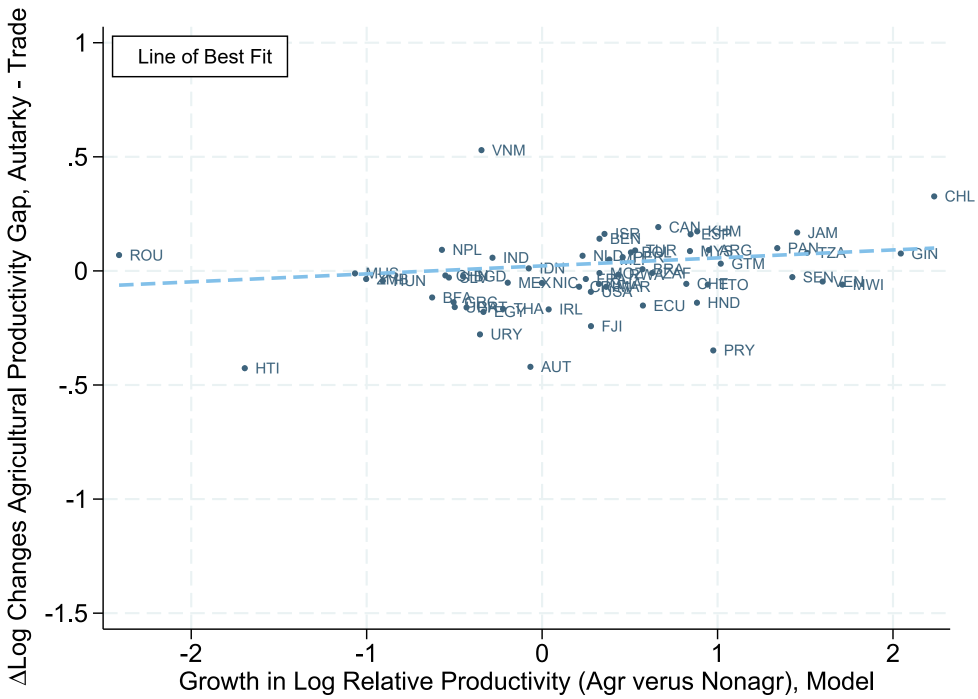

Note: We plot the difference of combined effects of productivity growth and labor movements between the baseline and the autarkic scenarios against relative productivity growth between agriculture and nonagriculture during the same time window.

C Additional Quantitative Results

C.1 Using Price and Wage Data to Recover Productivity Levels

In equation (13), only relative S_{it}^{j} matters, and therefore we can only estimate S_{it}^{j} relative to a reference country (and thus lack one degree of freedom for each year and sector). In the baseline calibration, we resolved this issue by choosing the US as the reference country and directly employing sectoral growth data of output per worker from FRED and USDA to measure the U.S. sectoral productivity levels. We notice that a large literature also uses price data in the estimation process to recover productivity levels for the US (e.g., Herrendorf, Herrington, and Valentinyi, 2015). To apply price data in our estimation process, after estimating equation (13) and obtaining the estimates relative to the US level ($\hat{S}_{it} - \hat{S}_{USt}$), we note from equation (8) that:

$$\hat{p}_{USt}^{j} = \left[\sum_{k} \exp(\hat{S}_{kt}^{j}) (\hat{d}_{kUSt}^{j})^{1-\eta} \right]^{1/(1-\eta)} = \left[\sum_{k} \exp(\hat{S}_{USt}^{j}) \exp(\hat{S}_{kt}^{j} - \hat{S}_{USt}^{j}) (\hat{d}_{kUSt}^{j})^{1-\eta} \right]^{1/(1-\eta)}. \tag{19}$$

Thus, we use the US price data $\hat{p}_{USt}^{j}$ for each year and sector to recover $\hat{S}_{USt}^{j}$ and other countries' estimates $\hat{S}_{it}^{j}$. We extract data on the US price indices from the EU KLEMS database.

We then proceed with two approaches to calibrate sectoral productivity levels z_{kt}^{j} for the US and other countries.

1. **Method of Moment.** We employ a similar approach as in our baseline calibration. Specifically, we search for the set of productivity levels $\{z_{kt}^{j}\}$ in the model that minimize the summed absolute difference between the model-predicted "competitiveness" level $\{S_{it}^{j,model}\}$ and the model estimates $\{\hat{S}_{it}^{j}\}$, $\sum_t \sum_i \sum_j |\hat{S}_{it}^{j} - S_{it}^{j,model}|$. Since we used price data from the US to determine the absolute level of S_{it}^{j}, there is no longer a need for a reference country. This allows us to estimate productivity levels for each country and sector, aligning them with the levels of S_{it}^{j} through this estimation process.

   Figure C.1 shows that with the productivity levels estimated using price data, the model predictions align well with the observed changes in agricultural productivity gaps.

Figure C.1: Dynamics of Agricultural Productivity Gaps in Counterfactual Scenarios (alternative calibration based on price data)

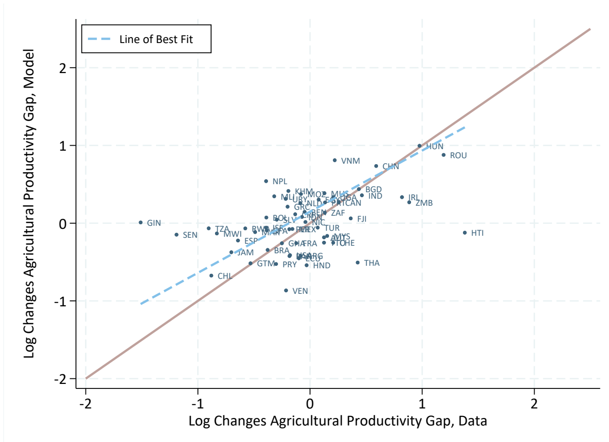

(a) Baseline

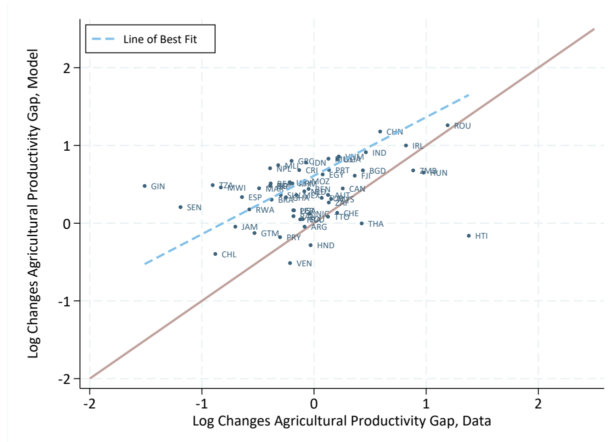

(b) Only Productivity Growth

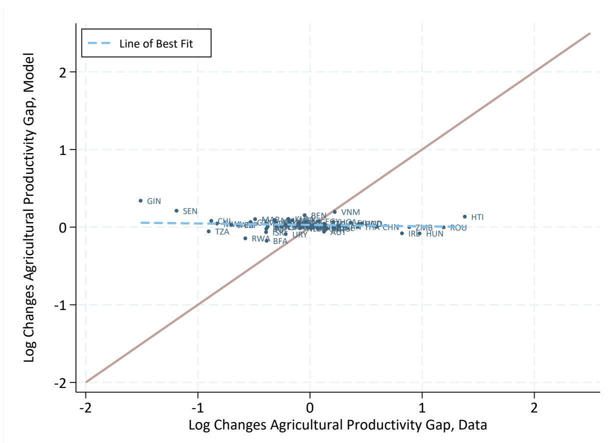

(c) Only Changes in Trade Costs

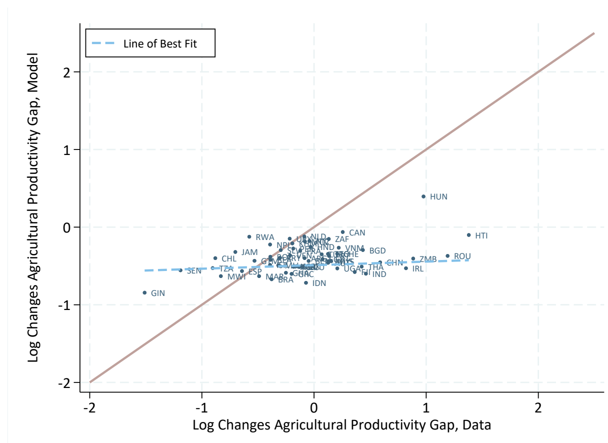

(d) Only Sectoral Labor Movements

2. **Using Wage Data.** Recall that $S_{kt}^j = (1-\eta)(\log w_{kt}^j - \log z_{kt}^j)$. We note from the expression of marginal costs that:

$$S_{it}^a = (1-\eta)\left[\log w_{it}^a - \log z_{it}^a\right]; \quad S_{it}^n = (1-\eta)\left[\log w_{it}^n - \log z_{it}^n\right]$$

To recover productivity estimates $\hat{z}_{it}^j$ from the estimates $\hat{S}_{it}^j$, we require yearly sectoral wage rates w_{kt}^j. To avoid directly incorporating the dynamics of agricultural productivity gaps in the calibration, for each country, we first use sectoral value added per worker to compute sectoral wage rates in the initial year (1980). We then combine these initial values with the yearly growth of GDP per capita (common to two sectors) to construct sectoral wage rates in later years.

Figure C.2: Model Fit of Dynamics of Agricultural Productivity Gaps (alternative calibration based on price and wage data)

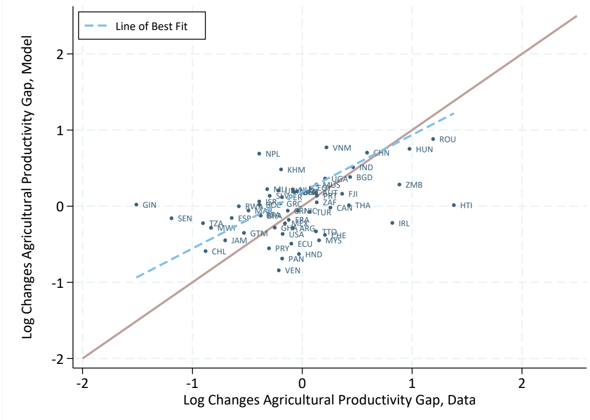

(a) Baseline Model

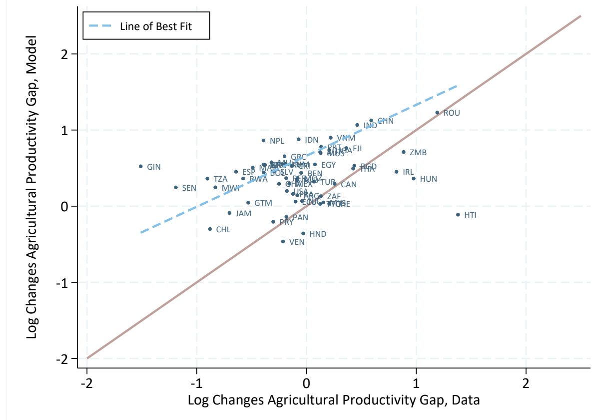

(b) Only Productivity Growth

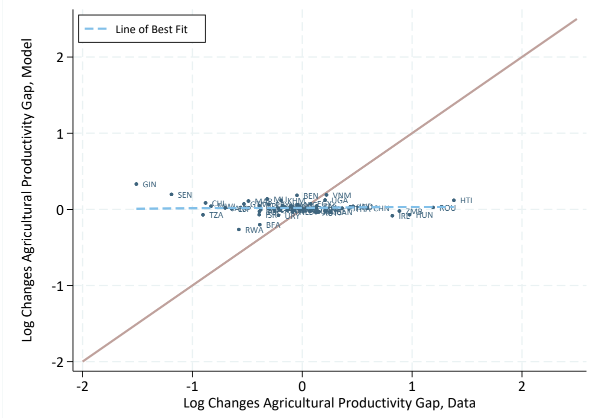

(c) Only Changes in Trade Costs

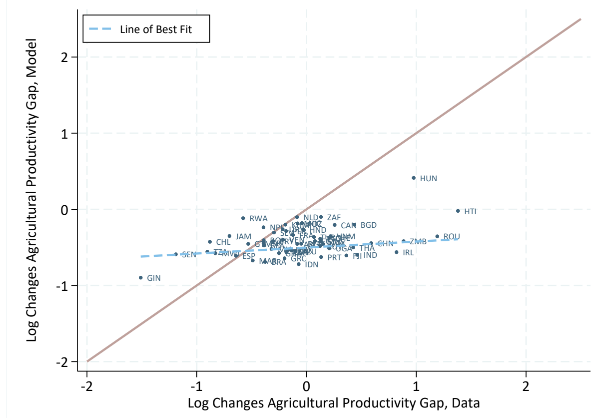

(d) Only Sectoral Labor Movements

Figure C.2 shows that with the productivity levels estimated using price and wage data, the model predictions still align well with the observed changes in agricultural productivity gaps.

C.2 Alternative Functional Form of Indirect Utility

Alternatively, we can use the indirect utility from Boppart (2014) in the case of two goods:

$$V_i(w) = \frac{1}{\varepsilon}\left(\frac{w}{p_i^n}\right)^{\varepsilon} - \frac{\nu}{\gamma}\left(\frac{p_i^a}{p_i^n}\right)^{\gamma} - \frac{1}{\varepsilon} + \frac{\nu}{\gamma},$$

with $1 > \gamma > 0$, $1 > \varepsilon > 0$, and $\nu > 0$. This implies a share of expenditures on agricultural goods:

$$\vartheta_i(w) = \nu \left(\frac{p_i^n}{w} \right)^{\varepsilon} \left(\frac{p_i^a}{p_i^n} \right)^{\gamma}.$$

The elasticity of substitution between two goods are given by:

$$\rho = 1 - \gamma - \frac{\nu \left(\frac{p_i^a}{p_i^n} \right)^{\gamma}}{\left(\frac{w}{p_i^n} \right)^{\varepsilon} - \nu \left(\frac{p_i^a}{p_i^n} \right)^{\gamma}} (\gamma - \varepsilon),$$

which, depending on the values of γ, ε, and ν, does not necessarily exceed 1. The elasticity of substitution converges to $1-\gamma$ when income w goes to infinity.

To estimate ε and γ, we perform the following regression:

$$\log \vartheta_{it} = \beta_0 + \varepsilon \log(p_{it}^n / w_{it}) + \gamma \log(p_{it}^a / p_{it}^n) + \varepsilon_{it} \tag{20}$$

We calculate the share of expenditures on agricultural goods for each country and year using our sectoral value-added and trade data. We use GDP per employee for w_{it}. We estimate country-year-level price index by $\hat{p}_{it}^j = \left[\sum_k \exp(\hat{S}_{USt}^j) \exp(\hat{S}_{kt}^j - \hat{S}_{USt}^j)(\hat{d}_{kit}^j)^{1-\eta} \right]^{1/(1-\eta)}$, where $\hat{S}_{USt}^j$ represents the fixed effects for the US, obtained from matching the US price level according to equation (19).

We conduct the regression analysis and find $\hat{\varepsilon} = 0.66$ and $\hat{\gamma} = 0.44$.[8] The estimate $\hat{\varepsilon} = 0.66$ corresponds to an income elasticity of agricultural consumption of 0.34, which is close to the 0.35 estimate reported in Boppart et al. (2023). This suggests that Engel curves for agricultural goods are flatter compared to those for non-agricultural goods. Additionally, $\hat{\gamma} > 0$ indicates that the two sectoral goods are gross complements. Finally, aligning with our baseline calibration, we still select the value of ν to align with the agricultural productivity gap for the US at the beginning of our sample period, resulting in $\nu = 0.023$.

Figure C.3 shows that with the alternative version of indirect utility function (the calibration for the parameters of other model elements remain unchanged), the model predictions still align well with the observed changes in agricultural productivity gaps.

[8] To avoid bias from interpolation, we conduct the regression analysis using only the country-years with actual APG data on agricultural shares of GDP and employment. The regression coefficients obtained from the full sample with interpolated data are quite similar.

Figure C.3: Model Fit of Dynamics of Agricultural Productivity Gaps (alternative functional form of indirect utility)

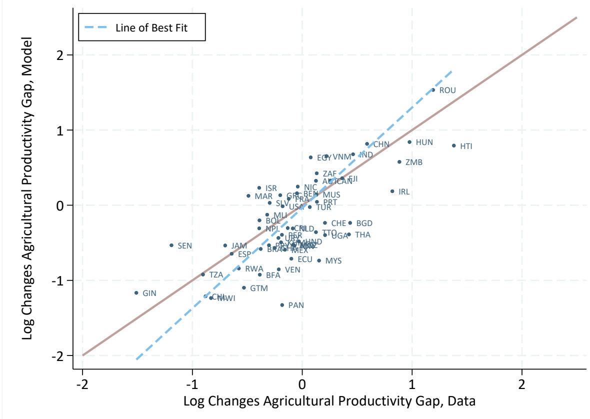

(a) Baseline Model

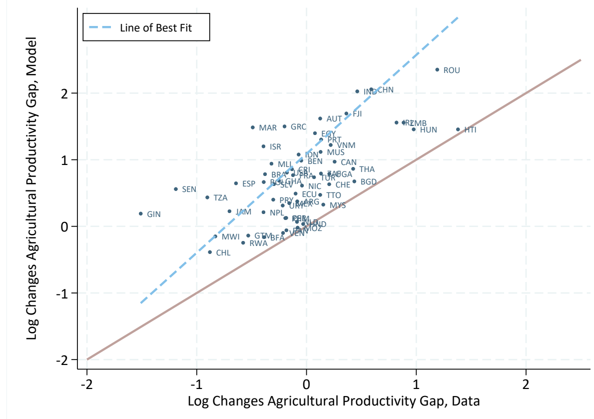

(b) Only Productivity Growth

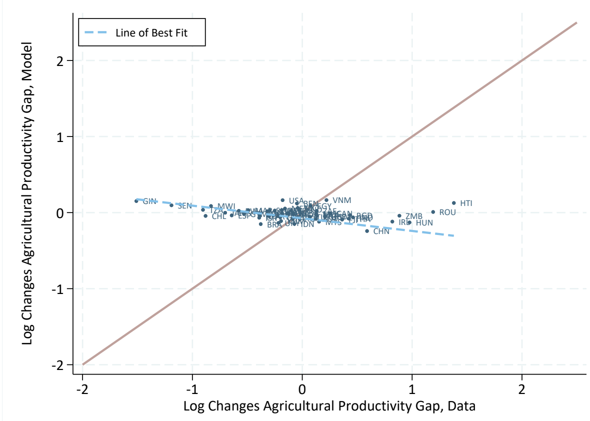

(c) Only Changes in Trade Costs

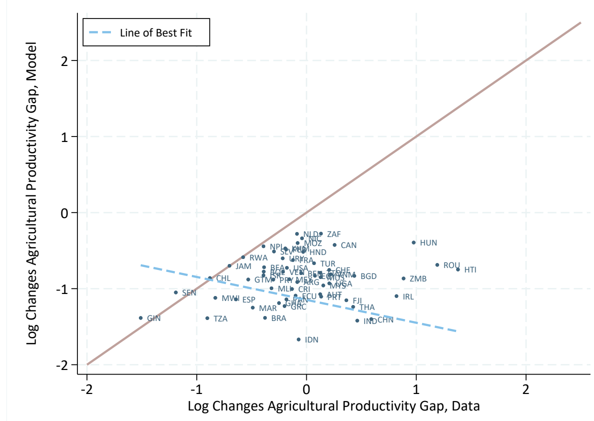

(d) Only Sectoral Labor Movements