The Macroeconomics of Poverty across Countries*

Nezih Guner
CEMFI and Banco de España

Duc Nguyen
Amherst College

Diego Restuccia
University of Toronto and NBER

Juan Carlos Ruiz-Garcia
European Commission

Guillaume Vandenbroucke
FRB St Louis

March 2026

Abstract

We assemble a harmonized cross-country dataset of household income distributions from censuses and surveys covering 91 countries and 192 country-years. Poverty declines sharply with development, but it is not evenly distributed across sectors: it is disproportionately concentrated in agriculture, especially in economies where agriculture absorbs a large share of employment. Across countries, agricultural workers earn less than non-agricultural workers, the agricultural income gap is larger in poorer economies, and incomes near the bottom of the distribution rise more than proportionally with labor productivity. To interpret these facts, we develop a parsimonious two-sector model of structural transformation with heterogeneous households, sectoral self-selection, and non-homothetic demand. The model is calibrated to aggregate, sectoral, and household-level moments for the United States in 1950 and then used for cross-country quantitative experiments. It accounts well for the joint relationships between development, poverty, agricultural employment, and the agricultural income gap. Quantitatively, economy-wide productivity is the main force behind poverty reduction across countries, while agricultural productivity further lowers poverty by raising incomes within agriculture. Sectoral distortions matter less at very low levels of development. Cross-country differences in poverty are therefore primarily differences in productivity, amplified by the concentration of poor households in agriculture.

Keywords: Poverty, inequality, income, households, structural transformation, productivity.
JEL classification: O11, O14, O4.

*Preliminary. Comments welcome. All remaining errors are our own. Restuccia gratefully acknowledges the support from the Canada Research Chairs program and the Bank of Canada Fellowship program. The views expressed herein are those of the authors and should not be attributed to Banco de España or the Eurosystem, the Bank of Canada or its Governing Council, nor the Federal Reserve Bank of St. Louis or the Federal Reserve System. Contact: nezih.guner@cemfi.es, dnguyen@amherst.edu, diego.restuccia@utoronto.ca, ruizgarcia.juancar@gmail.com, and guillaumevdb@gmail.com.

1 Introduction

Large differences in income per capita across countries remain one of the central facts of economic development. But countries do not differ only in average income. They also differ markedly in the distribution of living standards. Growing inequality in high-income countries, coupled with the availability of better data, has generated a large literature in macroeconomics on placing household heterogeneity at its center. Poorer countries exhibit much higher poverty rates, much larger income gaps between rich and poor households, and much greater concentration of deprivation in agriculture.

Why is there so much poverty in less developed countries? A natural answer emphasizes the distribution of resources within countries: poverty reflects unequal opportunities, missing markets, and barriers that keep some households trapped at the bottom (e.g., Balboni et al., 2022). But a more basic fact is that less-developed countries are poor across most of their income distribution. In the poorest economies, even households near the top of the local distribution often have living standards that would place them near the bottom in rich countries. This suggests that poverty across countries cannot be understood only as a problem of who gets what within a given economy. It is also a macroeconomic problem. When average productivity is low, the entire income distribution is compressed toward subsistence, and poverty becomes the natural outcome for a large share of the population. The central question is then not only why some households are poor, but why entire countries remain poor.

This perspective is consistent with two broad lessons from the literature. First, where people live is a first-order determinant of their economic fate. Milanovic (2015) shows that country-level circumstances account for a very large share of the variation in individual incomes worldwide, implying that global inequality is to a large extent a matter of location. Second, growth in average income is typically growth in the income of the poor. Dollar et al. (2016) show that incomes in the bottom 20 and bottom 40 percent tend to rise in proportion to mean incomes, because changes in their income shares are generally small and only weakly related to growth.

A key feature of those forces is structural transformation. Poverty in poor countries is not evenly spread across the economy. It is heavily concentrated in agriculture, where workers earn less, income is more dispersed, and a far larger fraction of households live below standard poverty lines. At the same time, in poor countries, much more labor is allocated to agriculture, and within countries, agricultural work is concentrated among poorer households. Development, therefore, involves more than a uniform increase in incomes. It also involves a reallocation of workers away from low-income agricultural activities and toward more productive nonagricultural sectors. Understanding cross-country poverty thus requires understanding the joint determination of productivity, sectoral allocation, and the distribution of household incomes. This is the focus of this paper.

We begin by constructing a harmonized household-level dataset spanning 91 countries and 192 country-years, using census and survey microdata from IPUMS, EU-SILC, LIS, and LSMS-ISA. These data allow us to compare household income per capita, poverty, inequality, and sector of employment across countries at very different stages of development.

We document several facts. First, poverty rates decline sharply with development. Second, poverty is disproportionately concentrated in agriculture, and in poor countries, the vast majority of the poor work in that sector. Third, agricultural employment declines strongly with development, and within countries, poorer households are much more likely to work in agriculture. Fourth, agricultural workers also earn substantially less than nonagricultural workers in average, with the agricultural income gap larger in poorer countries. Finally, incomes at the bottom of the distribution rise more steeply with labor productivity than incomes near the median or the top. These facts suggest that poverty across countries is tightly linked to the same forces that govern aggregate productivity and structural transformation.

To interpret these patterns, we develop a parsimonious two-sector general equilibrium Roy (1951) model with agriculture and nonagriculture, heterogeneous households, and nonhomothetic preferences (Lagakos and Waugh, 2013; Adamopoulos et al., 2022). Households differ in their sector-specific earning abilities and sort across sectors according to compara-

tive advantage, subject to a sectoral mobility wedge that lowers the effective return to nonagricultural work. Non-homothetic preferences generate Engel curves that make structural transformation an equilibrium response to development. On the production side, economy-wide productivity affects both sectors, while agricultural productivity affects agriculture specifically. This simple structure allows us to study how changes in economy-wide productivity, agricultural productivity, and sectoral distortions jointly shape occupational choice, the sectoral allocation of labor, and the cross-sectional distribution of household income.

We calibrate a benchmark economy to aggregate, sectoral, and micro moments from the United States in 1950. This benchmark is useful for two reasons. First, it is rich enough to discipline both sectors while still featuring nontrivial poverty and meaningful agricultural employment. Second, it predates the major expansion of modern antipoverty programs, allowing the model to map more directly from earnings and sectoral allocation into household income.

We then use the calibrated model to conduct cross-country quantitative experiments in which, relative to the benchmark economy, we vary only three parameters: economy-wide productivity, agricultural-specific productivity, and a sectoral mobility wedge. The quantitative exercise delivers three main results. First, the model does well quantitatively: although the cross-country experiments are disciplined only by productivity and sectoral-allocation forces, the model reproduces the strong cross-country relationships between development and aggregate poverty, sectoral poverty rates, the share of poverty in agriculture, agricultural employment, the agricultural income gap, and the inverted-U pattern of income dispersion. Second, economy-wide productivity is the dominant force behind poverty reduction and structural transformation. Starting from a poor economy, raising aggregate productivity alone accounts for most of the decline in both poverty and agricultural employment, because it shifts the entire income distribution upward and, through non-homothetic demand, pulls workers out of agriculture. Third, agricultural productivity also matters importantly, especially by raising incomes within agriculture and further reducing poverty, whereas lowering

mobility frictions on its own has only modest effects at very low levels of development and can even raise poverty through general-equilibrium increases in food prices when agriculture remains highly unproductive. The broader implication is that cross-country differences in poverty are best understood primarily as a productivity phenomenon, amplified by the concentration of poor households in agriculture.

Our paper contributes to three strands of work. First, it relates to the macro-development literature on cross-country income differences (Chari et al., 1996; Klenow and Rodriguez-Clare, 1997; Prescott, 1998; McGrattan and Schmitz Jr, 1999; Caselli, 2005; Jones, 2016). We contribute to this literature by shifting attention from mean income to the entire household income distribution. The question is not only why some countries are poorer on average, but why low productivity translates into much greater poverty and much larger lower-tail compression. The relationship between economic development and inequality was first highlighted by Kuznetz (1955). Another strand of literature focuses on how inequality can affect development in the presence of credit market imperfections and non-convexities; Banerjee and Newman (1993) and Galor and Zeira (1993) are classic references. Recent papers that study how inequality changes with economic development include, among others, Erosa et al. (2010), Lagakos et al. (2018), and Guner and Ruggieri (2022).

Second, it contributes to the literature on growth and poverty by providing a quantitative framework in which the strong empirical link between average income growth and the incomes of the poor emerges naturally from sectoral allocation and non-homothetic demand. Recent empirical contributions on economic growth and poverty include Milanovic (2015) and Dollar et al. (2016). The paper is also related to recent work that organizes and presents facts on income inequality and dynamics, such as Krueger et al. (2010) and Guvenen et al. (2022). And more broadly, on papers that study how inequality and macroeconomics interact; Ahn et al. (2018) and Gomez (2025) provide recent reviews.

Third, it contributes to the literature on structural transformation (Kongsamut et al., 2001; Gollin et al., 2002; Ngai and Pissarides, 2007; Restuccia et al., 2008; Acemoglu and

Guerrieri, 2008; Boppart, 2014; Duarte and Restuccia, 2010; Comin et al., 2021; Herrendorf et al., 2014). We contribute to this literature by showing that the same forces that govern labor reallocation across sectors also account for major differences in poverty and inequality across countries.

The rest of the paper proceeds as follows. In the next section, we describe the household-level data and present the main empirical observations across countries. Section 4 describes the model and characterizes the equilibrium solution. In Section 5, we calibrate the model to aggregate, sectoral, and disaggregate data for the United States in 1950, and report the results from the main quantitative experiments. We conclude in Section 6.

2 Data

We construct an ambitious cross-country dataset based on household surveys for 91 countries at different points in time and different levels of development (192 country-year surveys), totaling millions of individual-level observations. The poorest country in our sample is Niger in 2011 with a GDP per capita of $800 (2011 international dollars, adjusted for purchasing power parity), while the richest country in our sample is Luxembourg in 2011 with a GDP per capita of $92,000. The dataset is constructed from four sources:

1. IPUMS International, which covers 15 countries from 1960 to 2010 and provides 38 country-year observations.

2. European Union Statistics on Income and Living Conditions (EU-SILC), which covers 31 countries from 2003 to 2011 and provides 61 country-year observations.

3. Luxembourg Income Study (LIS), which covers 38 countries from 1974 to 2016 and provides 86 country-year observations.

These three sources primarily cover middle and high-income countries. To be able to include low-income countries in our sample, we turn to the World Bank's Living Standard

Measurement Study, Integrated Survey of Agriculture (LSMS-ISA):

4. While the LSMS provides household surveys for many countries, the data is not provided in a homogeneous, readily comparable format. Furthermore, not all surveys provide income data. Most importantly, in very poor countries, many households operate near subsistence agriculture and therefore do not report any earnings or labor market income. To obtain a comprehensive measure of income, we construct a measure of agricultural income from detailed information on agricultural outputs and inputs. We also compute income from business activities outside the primary sector. The details are provided in Appendix X. From the LSMS, we cleaned and processed data for 7 countries from 2004 to 2013, which adds 7 country-year observations.

2.1 Data Harmonization

Because the raw surveys differ substantially across sources, all datasets are first harmonized into a common individual-level structure. The harmonization stage constructs comparable variables for demographics, education, labor market status, sector of employment, hours worked, and income. Education is grouped into primary, secondary, and tertiary levels, while industries are collapsed into agriculture, industry, and services.

Individuals are classified as employed if they report positive hours worked, positive labor income, or report being employed in the survey data. To ensure consistency, labor market variables are reconciled with hours and income information. For example, individuals reporting positive hours worked, positive wage income, or positive labor earnings are reclassified as employed and as participating in the labor force even if the original employment indicator differs.

2.2 Income Construction

The central variable used in the analysis is total household income per capita. Income includes labor earnings, business income (wages plus income from farm and non-farm businesses), capital income and transfers where available in the survey.

For most surveys (IPUMS, LIS, and EU-SILC), income variables are reported directly in the microdata. These are harmonized into three main measures: wage income, labor earnings, and total income. For LSMS surveys that do not directly report income components, we construct income from detailed information on agricultural production, non-farm businesses, and transfers.

Household income is constructed by aggregating individual income across all household members. To obtain household income per capita, we divide total household income by the number of household members. In some robustness checks, we also compute equivalised household income using the OECD equivalence scale to account for household composition.

Treatment of zero incomes. Households with zero total income are excluded from the baseline sample when the zero value appears inconsistent with reported labor supply or employment status. In particular, individuals who report working positive hours but have zero labor income are treated as having missing income rather than true zero earnings.

However, households that legitimately report zero income, particularly in subsistence agricultural settings, are retained when agricultural production information is available and used to construct income.

Price adjustments and international comparability. To ensure comparability across countries and time, all income variables are converted into 2011 international dollars. The conversion proceeds in two steps.

First, incomes reported in local currency are converted into constant 2011 local currency units using country-specific consumer price indices (CPI). Second, these values are converted

into international dollars using 2011 purchasing power parity (PPP) exchange rates.

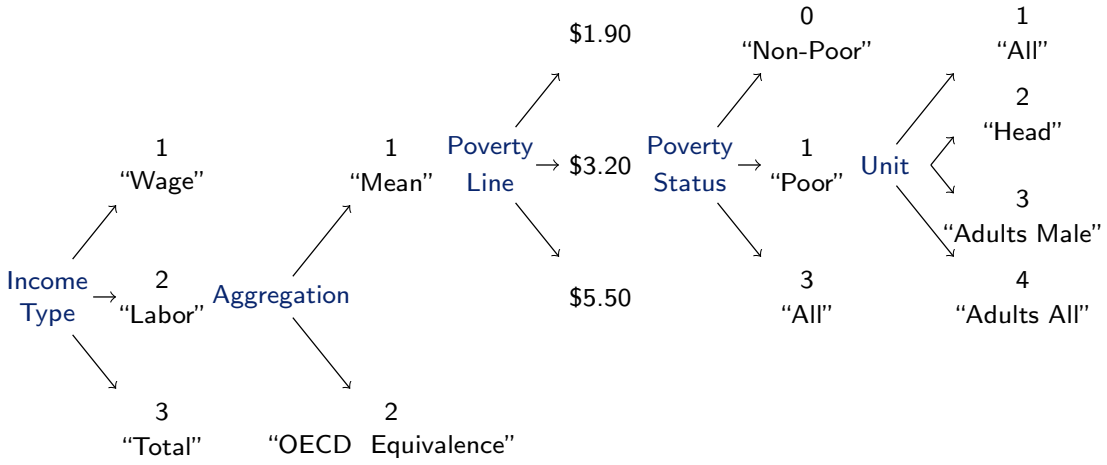
Finally, mean household income in each country–year is scaled to match GDP per capita from national accounts, preserving the within-country income distribution observed in the microdata.

2.3 Sample restrictions.

The analysis focuses on the working-age population aged 18 to 64. Observations outside this age range are excluded to maintain comparability across countries.

Additional consistency checks are applied to ensure valid household structure. For example, households must contain exactly one household head and heads must be adults. Observations with inconsistent household identifiers or missing key variables are dropped.

Figure 1: Sample Restrictions for Poverty Statistics



Poverty measurement. Poverty is defined at the household level. We compute household income per capita by dividing total household income by household size. A household is classified as poor if its per-capita income falls below a given poverty threshold, and all members of that household are considered poor. Poverty rates are then calculated as the share of individuals living in poor households using survey weights. Throughout the analysis we use the international poverty lines of \$1.90, \$3.20, and \$5.50 per person per day (2011

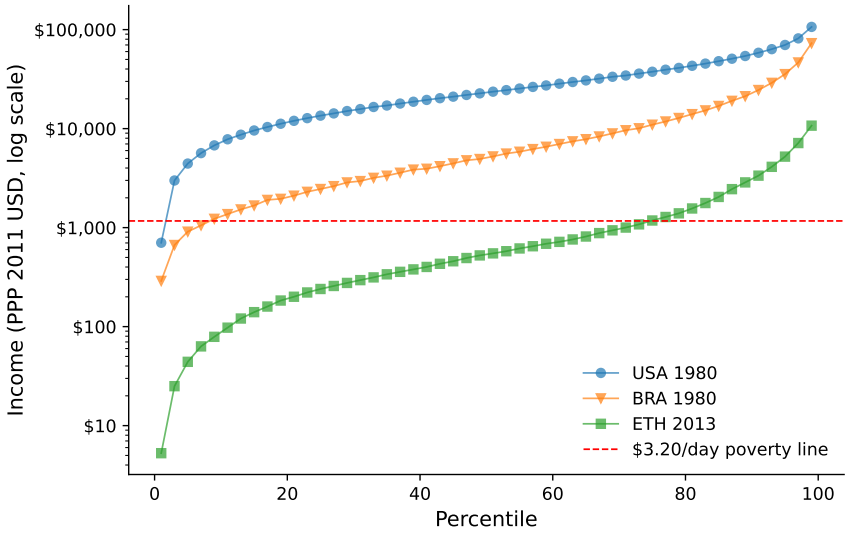
PPP dollars). Figure 1 illustrates the decision tree for our classification of individuals as poor and non-poor and the different units for which statistics can be reported: all household members, head of household, male adults, and all adults in the household.

3 Facts

In this section, we present the key patterns that emerge from the data.

Illustration. While it is well known the great disparities in average income per capita across nations (Klenow and Rodriguez-Clare, 1997; Caselli, 2005; Jones, 2016), less know are the distributional differences and, in particular, the determinants of levels of poverty across countries. To illustrate some of the properties we document in detail below across countries, Figure 2 reports the percentile distribution of household income per capita (fifty percentile points from p1 to p99) across three countries, USA 1980, Brazil 1980 and Ethiopia 2013.

Figure 2: Household Income per Capita by Percentiles



Notes: Household income per capita by percentile for the USA in 1980, Brazil in 1980 and Ethiopia in 2013. Income is household income adjusted for household size (OECD equivalence scale). Dashed red line is the World Bank $3.20/day poverty line (2011 PPP, $1,168/year). Data are from household microdata.

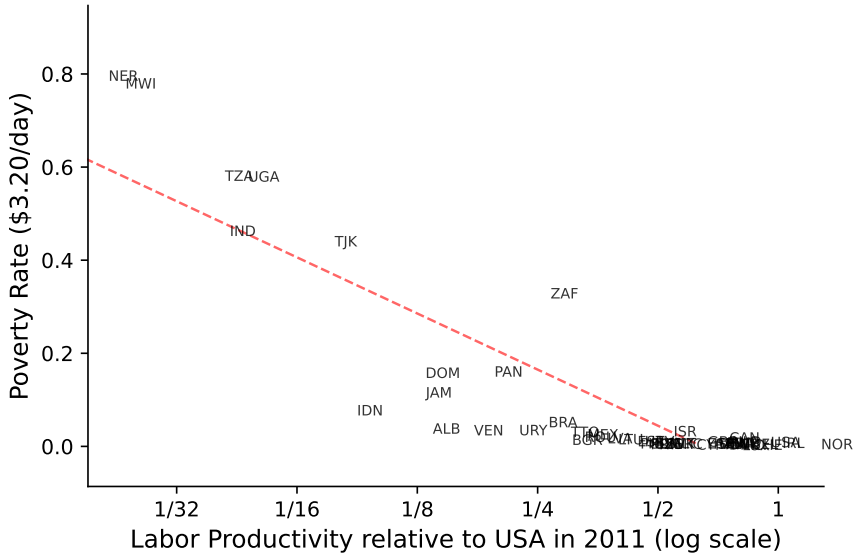
Relative to USA 1980, average income per capita in Brazil 1980 is 37% and Ethiopia 2013

is 4.5%, conforming to the large differences in average income typically emphasized in the literature. The figure illustrates important differences in the distribution of income, with more dispersion in Brazil and Ethiopia than the United States. This matters for poverty for example, the poverty rate (according to a poverty line of 3.20/day (PPP 2011 $) poverty line) is 0.8% in USA 1980 and 78% in Ethiopia. It also illustrates how pure redistributive policies would be limited in alleviating poverty in Ethiopia, a redistribution that leaves average income the same may actually increase poverty. In the observations that follow, we generalize this characterization across all countries in our sample and connect them with the sectoral composition of the economy, productivity and with factors that may keep people in agriculture beyond productivity.

3.1 Poverty Declines Sharply with Development

The poverty rate (at the $3.20/day, 2011 PPP, threshold) is virtually zero in rich countries and exceeds 75% in the poorest countries (see Figure 3).

Figure 3: Poverty Rate across Countries ($3.20/day)

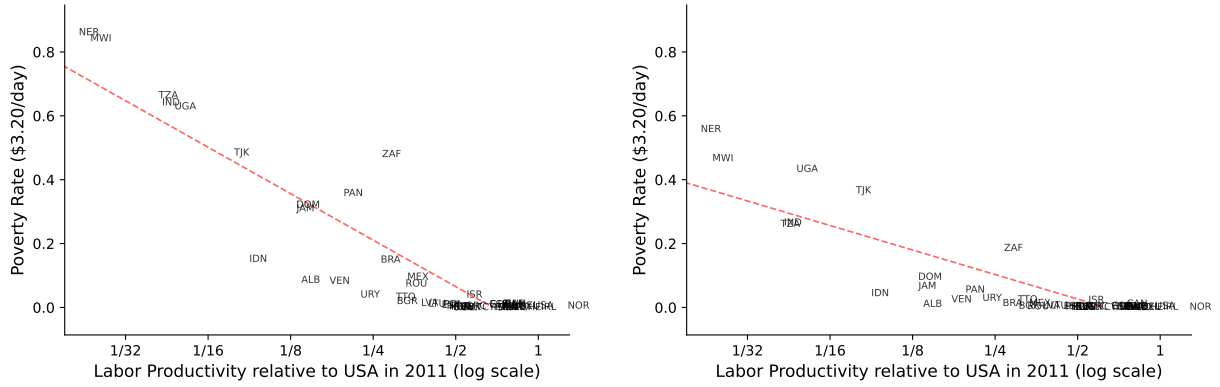


Notes: Poverty rate at the World Bank $3.20/day line (2011 PPP). Data from household microdata; aggregate labor productivity from PWT 11.0.

The relationship between the poverty rate and aggregate labor productivity holds in both

sectors, though individuals engaged in agricultural work face substantially higher poverty rates at every level of development (see Figure 4).

Figure 4: Poverty Rate in each Sector across Countries ($3.20/day)



(a) Agriculture (b) Non-Agriculture

Notes: Poverty rate at the World Bank $3.20/day line (2011 PPP). Panels share the same y-axis scale. Sector assigned by household head's industry. Data from household microdata; aggregate labor productivity from PWT 11.0.

In poor countries, the vast majority of the poor are engaged in agriculture. The share of all poor individuals engaged in agriculture rises sharply at lower development levels (Figure 5). Individuals in agriculture face substantially higher poverty rates everywhere, and the gap between agricultural and non-agricultural poverty widens in poorer economies. This composition effect means that understanding poverty requires understanding the agricultural sector.

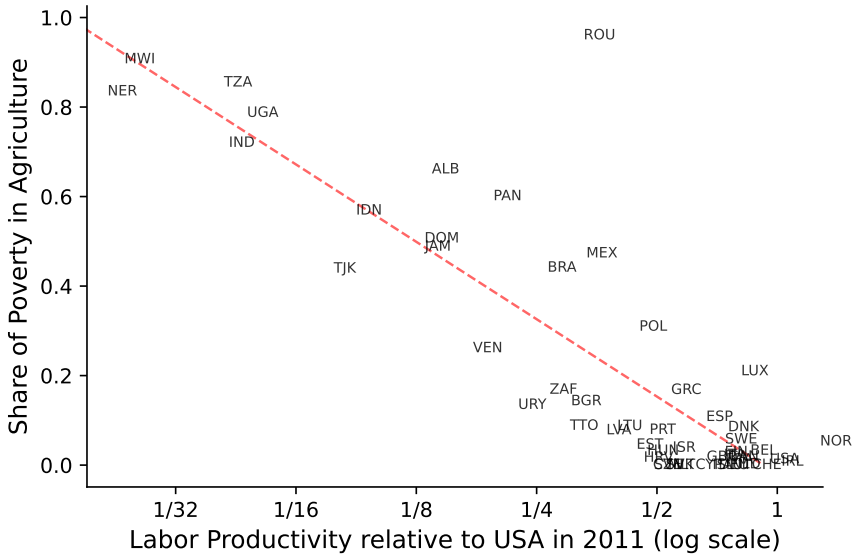
3.2 Agricultural Employment Declines with Development

The share of employment in agriculture declines steeply with labor productivity across countries (Figure 6), from over 80% in the poorest countries to below 5% in the richest.

3.3 Agricultural Employment is Concentrated among the Poor

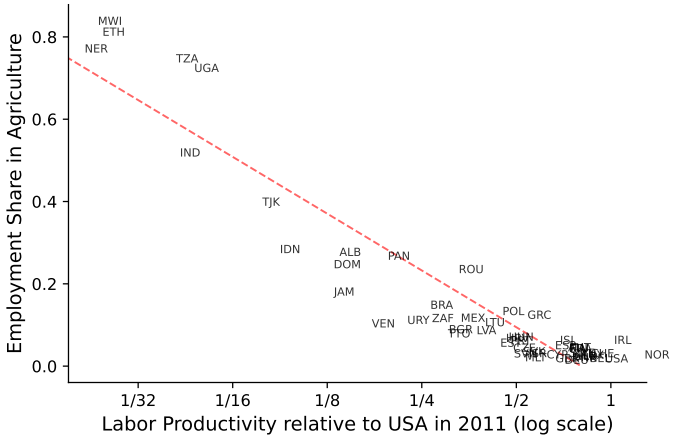
Within countries, poorer households are more likely to be engaged in agriculture at every level of development (Figure 7). The gradient is steep in middle-income economies (Brazil in

Figure 5: Share of Poverty in Agriculture across Countries



Notes: Share of poverty in agriculture is the fraction of all individuals below the $3.20/day poverty line (2011 PPP) who are employed in agriculture. Each label is a country (ISO3 code). X-axis is labor productivity (GDP per worker from PWT 11.0) relative to USA in 2011 on a log scale. Dashed red line is a log-linear fit. Data from household microdata.

Figure 6: Agricultural Employment Share across Countries

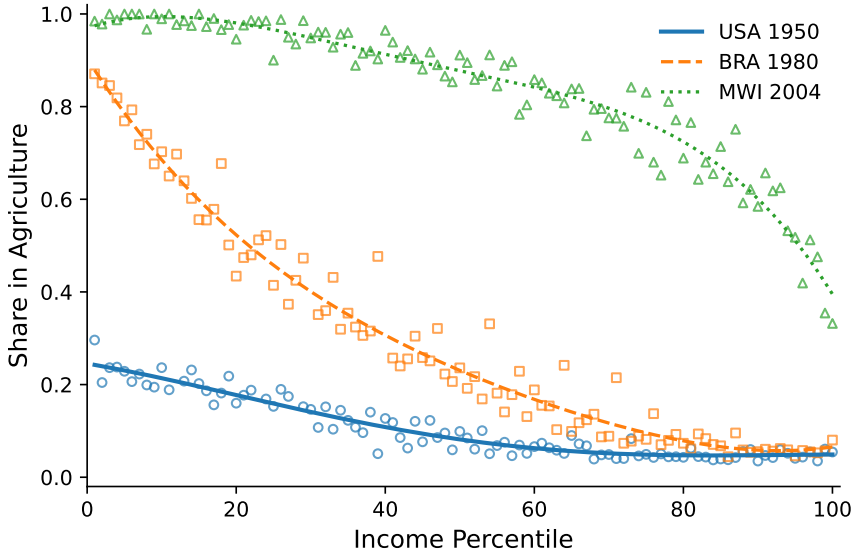


Notes: Each label is a country (ISO3 code). X-axis is labor productivity (GDP per worker from PWT 11.0) relative to USA in 2011 on a log scale. Dashed red line is a log-linear fit. Data from household microdata surveys across 54 countries.

1980) and flatter but high in the poorest economies (Malawi in 2010). In the United States in 1950, the poorest deciles have the highest agricultural employment, but levels are already low overall.

The same pattern holds over time. As countries grow, workers move out of agriculture, especially among the poor. This is documented for the USA (1950 to 2015) and Brazil (1980 to 2010) in Figure 8.

Figure 7: Agricultural Employment Share by Income Percentile: Cross-Country



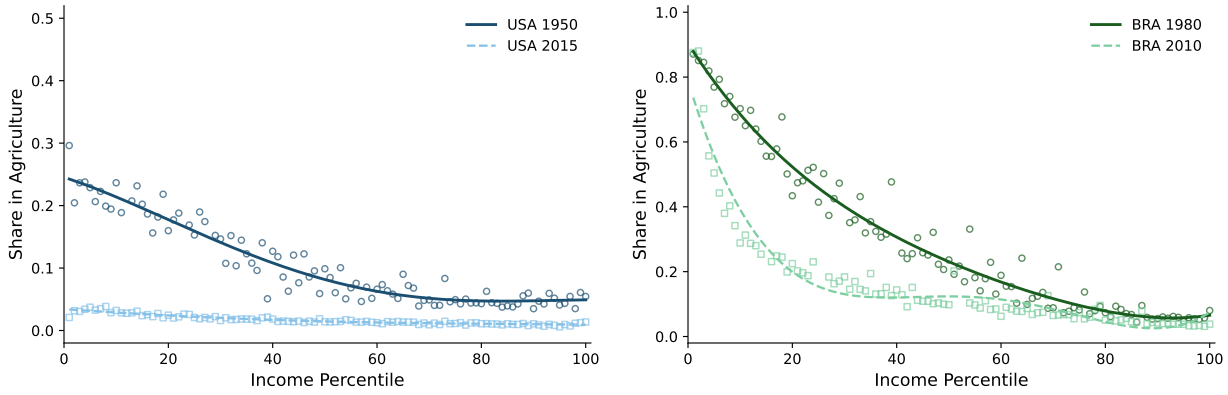
Notes: Scatter points show data from household microdata (IPUMS for USA and Brazil; LSMS-ISA for Malawi). Lines are degree-4 polynomial fits. Income percentiles based on household income adjusted for household size (OECD equivalence scale).

3.4 Income in Agriculture Lower than in Non-Agriculture

There is a persistent average agricultural income gap (AIG): average income in agriculture is below average income in non-agriculture across all development levels (Figure 9). The AIG is larger in poorer countries.

At a more disaggregated level, we plot income by percentile separately for agriculture and non-agriculture. In the USA in 1950 (Figure 10), agricultural income is systematically below non-agricultural income at every within-sector percentile. The overall poverty rate is 7.7%,

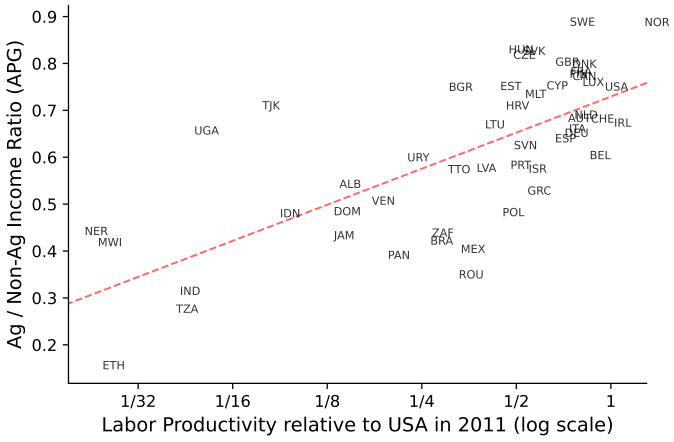
Figure 8: Agricultural Employment Share by Income Percentile: Over Time



(a) USA: 1950 vs. 2015 (b) Brazil: 1980 vs. 2010

Notes: Scatter points show data from household microdata (IPUMS). Lines are degree-4 polynomial fits. Darker shades denote earlier years. Income percentiles based on household income adjusted for household size.

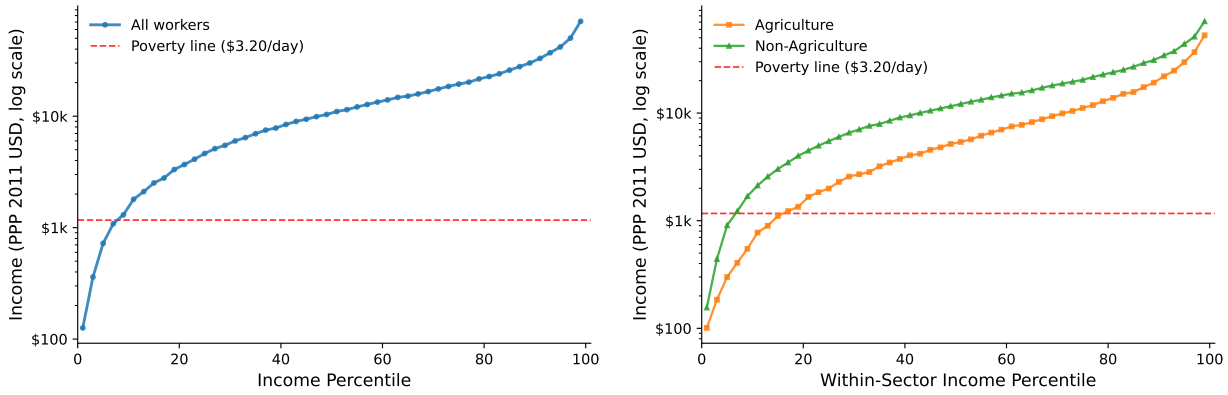
Figure 9: Agricultural Income Gap across Countries



Notes: Average agricultural income gap is average income in agriculture divided by average income in non-agriculture. Each label is a country (ISO3 code). X-axis is labor productivity (GDP per worker from PWT 11.0) relative to USA in 2011 on a log scale. Dashed red line is a log-linear fit. Data from household microdata surveys across 54 countries.

but the rate among those in agriculture is 16% compared with only 6.4% among those in non-agriculture. Although agriculture accounts for just 10.7% of employment, it accounts for about 22% of all poverty. This directly links the preceding facts: poverty is concentrated in agriculture because agriculture income is lower.

Figure 10: Income by Percentile: USA 1950



(a) All workers

(b) Agriculture vs. Non-Agriculture

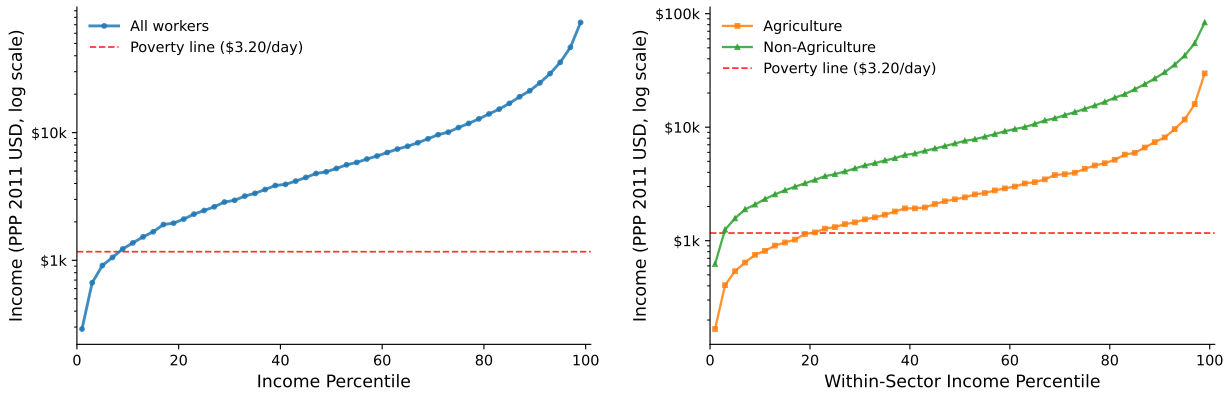
Notes: Scatter points show data from US 1950 Census microdata (IPUMS). Lines are degree-4 polynomial fits. Panel (b) uses within-sector percentiles. Dashed red line is the World Bank $3.20/day poverty line (2011 PPP, $1,168/year). Income is household income adjusted for household size (OECD equivalence scale).

The same pattern holds in poorer economies but with dramatically different levels. In Brazil 1980 (Figure 11), the overall poverty rate is 8.0%, but the poverty rate among those in agriculture is 20.6% compared with only 2.2% among those in non-agriculture. Because of the size of the agricultural sector that makes up up to 30% of employment, over 77% of all poverty is accounted for by those in agriculture. Agricultural income is again below non-agricultural income at every percentile of the income distribution, and the poverty line cuts through a much larger portion of the agricultural income distribution.

In Malawi 2004, a much poorer country (Figure 12), the overall poverty rate is 78%, with 84% of those in agriculture and 47% of those in non-agriculture below the poverty line. About 91% of all poverty from agriculture which comprises 84% of employment. At this low development level, nearly the entire income distribution falls below the poverty line and agricultural and non-agricultural incomes are compressed together at very low levels, though

the ag–non-ag gap persists.

Figure 11: Income by Percentile: Brazil 1980



(a) All workers

(b) Agriculture vs. Non-Agriculture

Notes: Scatter points show data from Brazil 1980 Census microdata (IPUMS International). Panel (b) uses within-sector percentiles. Dashed red line is the World Bank $3.20/day poverty line (2011 PPP, $1,168/year). Income is household income adjusted for household size (OECD equivalence scale), converted to PPP 2011 USD.

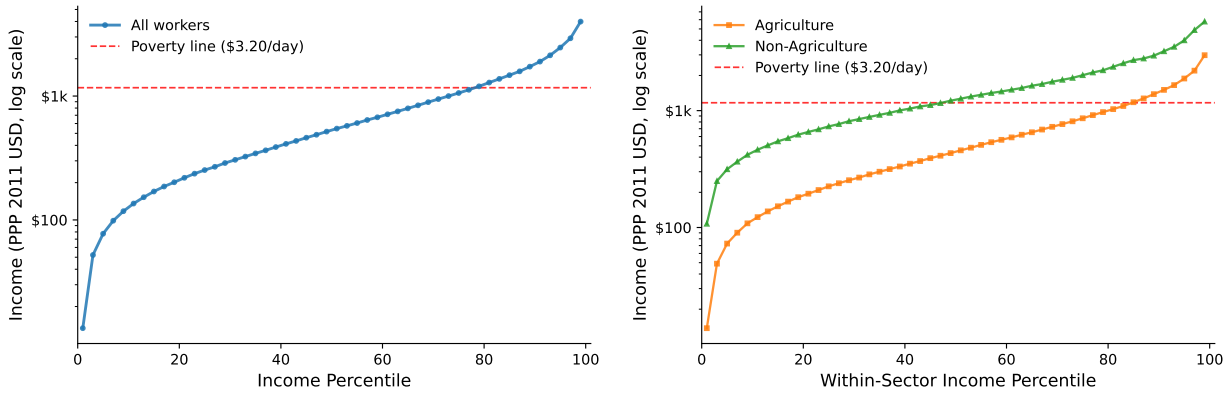
3.5 The Bottom of Income Distribution Increases Faster with Development than the Median and the Top

Income at the bottom of the distribution is more elastic to aggregate productivity than the median. Figure 13 plots income at the 5th, 50th and 95th percentiles against labor productivity across countries. The fitted log-log slopes show that the elasticity of bottom incomes to aggregate productivity exceeds one: a doubling of labor productivity is associated with more than a doubling of income at the 5th percentile. By contrast, median and the top income grows approximately one-for-one with productivity. This implies that development disproportionately lifts the poorest, compressing the lower tail of the income distribution.

4 Model

We develop a general equilibrium Roy (1951) model of self-selection with two sectors (agriculture and non-agriculture) following Lagakos and Waugh (2013) and Adamopoulos et al.

Figure 12: Income by Percentile: Malawi 2004

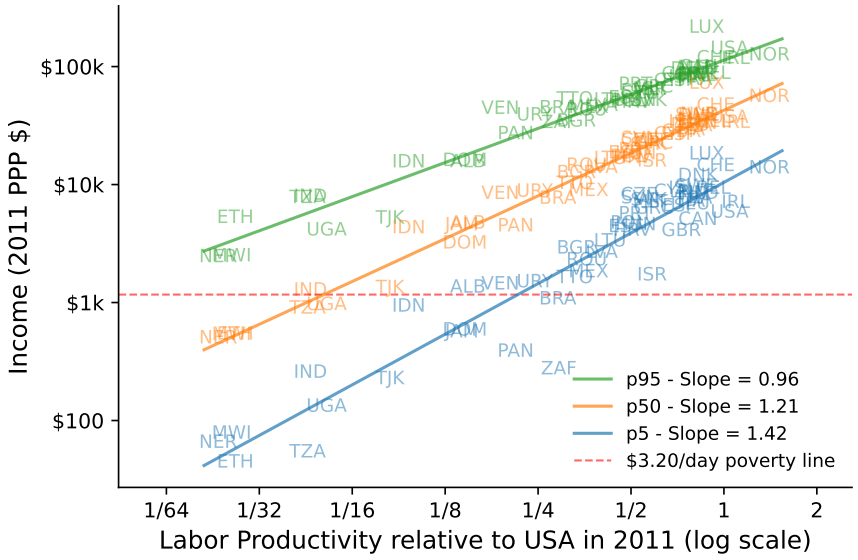


(a) All workers

(b) Agriculture vs. Non-Agriculture

Notes: Scatter points show data from Malawi 2004 household survey (LSMS-ISA). Panel (b) uses within-sector percentiles. Dashed red line is the World Bank $3.20/day poverty line (2011 PPP, $1,168/year). Income is household income adjusted for household size (OECD equivalence scale), converted to PPP 2011 USD.

Figure 13: Income at Selected Percentiles vs. Labor Productivity



Notes: Each label is a country (ISO3 code) at the latest available year. Lines are log-log OLS fits for each percentile. Income is household income adjusted for household size (OECD equivalence scale), converted to PPP 2011 USD. X-axis is labor productivity (GDP per worker from PWT 11.0) relative to USA in 2011. Dashed red line is the World Bank $3.20/day poverty line.

(2022). Individuals are heterogeneous in sectoral abilities and face frictions in their occupational choice. Production technologies are linear in effective labor, and, following Boppart (2014), households have Price-Independent Generalized Linear (PIGL) preferences over agricultural and non-agricultural goods. The model is static and delivers sectoral income differences and agricultural income gaps that depend on productivity, agricultural technology, and labor mobility frictions.

4.1 Economic Environment

Technologies. On each date, two goods (agricultural and non-agricultural) are produced in two sectors. Agricultural output Y_a is produced according to

$$Y_a \;=\; A\kappa\, H_a, \tag{1}$$

where $A > 0$ is an economy-wide productivity parameter, $\kappa > 0$ is an agriculture-specific productivity parameter, and H_a is the aggregate effective labor in agriculture.

Non-agricultural output Y_n is produced according to

$$Y_n \;=\; A\, H_n, \tag{2}$$

where H_n is aggregate effective labor in non-agriculture.

Preferences and endowments. There is a unit mass of individuals endowed with one unit of time each period supplied inelastically to the market. Individuals have Price-Independent Generalized Linear (PIGL) preferences over agricultural and non-agricultural goods as specified below. Individuals are endowed with a pair of sector-specific abilities

$$(z_a, z_n) \in \mathbb{R}^2_+,$$

drawn from a joint distribution $F(z_a, z_n)$ with a continuous density. The distribution F governs ex-ante heterogeneity in comparative advantage across sectors.

An individual with abilities (z_a, z_n) engaged in agriculture supplies effective labor $h_a = z_a v$, where z_a is the agricultural ability and v is an idiosyncratic agricultural productivity shock that is realized for those engaged in agriculture after occupation decisions are made. The shock v is i.i.d. across workers and drawn from a distribution $G(v)$ with $\mathbb{E}[v] = 1$, capturing shocks that affect agricultural output, such as weather. Aggregate effective labor in agriculture is

$$H_a \;=\; \int_{(z_a, z_n)} o(z_a, z_n)\, z_a \, dF(z_a, z_n), \tag{3}$$

noting that the ex-post shock v disappears from this equation since $\mathbb{E}[v] = 1$.

An individual with abilities (z_a, z_n) engaged in non-agriculture supplies effective labor $h_n = z_n$, and aggregate effective labor in non-agriculture is

$$H_n \;=\; \int_{(z_a, z_n)} (1 - o(z_a, z_n))\, z_n \, dF(z_a, z_n). \tag{4}$$

Market structure. Suppose the price of nonagricultural goods is the numeraire, $p_n = 1$, and let p_a be the relative price of agricultural goods. We denote sectoral wages per effective unit of labor in sector i by w_i. Under competitive factor markets, wages equal marginal products:

$$w_a \;=\; p_a A \kappa, \qquad w_n \;=\; A. \tag{5}$$

Occupational choice. Occupational choice takes place before the realization of the agricultural shock v. The key friction is a common labor mobility barrier θ that potentially reduces the effective income to non-agricultural work.

Given prices (p_a, w_a, w_n), pre-shock sectoral incomes based on abilities are

$$\tilde{I}_a(z_a) = w_a z_a, \tag{6}$$

and

$$I_n(z_n) = w_n z_n. \tag{7}$$

The effective payoff from working in non-agriculture is affected by the labor mobility barrier θ, hence the non-agricultural *effective* income is $\theta I_n(z_n)$.

Because the agricultural shock v is realized *after* occupation (sector) choice, individuals choose the sector that maximizes *expected indirect utility*. Under PIGL preferences, the indirect utility is given by

$$\mathcal{V}(p_a, I) = \frac{1}{1-\varphi} I^{1-\varphi} - \frac{\nu}{1-\eta} p_a^{1-\eta} - \frac{1}{1-\varphi} + \frac{\nu}{1-\eta}, \tag{8}$$

where I is household income, p_a is the relative price of agriculture, $\varphi \in (0,1)$ governs the income elasticity of demand for agricultural goods and $\eta > 0$ governs the price elasticity of demand for agricultural goods. The parameter $\nu > 0$ controls the share of agricultural expenditure in total income. Hence, the occupational choice rule is to work in non-agriculture if the following condition holds:

$$\mathcal{V}(p_a, \theta I_n(z_n)) \;\geq\; \mathbb{E}_v\Big[\mathcal{V}\Big(p_a, \tilde{I}_a(z_a)v\Big)\Big]. \tag{9}$$

A key simplification that arises with PIGL preferences is that all price-only terms in $\mathcal{V}$ cancel in the comparison of indirect utility across sectors. Using equation (8), the occupational choice condition in (9) is equivalent to

$$(\theta I_n(z_n))^{1-\varphi} \;\geq\; \mathbb{E}_v\left[\left(\tilde{I}_a(z_a)v\right)^{1-\varphi}\right]. \tag{10}$$

Therefore, occupational choice depends on a comparison of (power-transformed) incomes.

As v is independent of (z_a, z_n), then

$$\mathbb{E}_v\left[\left(\tilde{I}_a(z_a)v\right)^{1-\varphi}\right] = \tilde{I}_a(z_a)^{1-\varphi} \cdot \mathbb{E}\left[v^{1-\varphi}\right],$$

and condition (10) reduces to a simple cutoff rule

$$\theta I_n(z_n) \;\geq\; \tilde{I}_a(z_a)\left(\mathbb{E}[v^{1-\varphi}]\right)^{\frac{1}{1-\varphi}}. \tag{11}$$

Equivalently, in wages and abilities,

$$\theta\, w_n z_n \;\geq\; w_a z_a\left(\mathbb{E}[v^{1-\varphi}]\right)^{\frac{1}{1-\varphi}}. \tag{12}$$

The occupational choice function $o(z_a, z_n) \in \{0, 1\}$, where $o = 1$ denotes agriculture and $o = 0$ denotes non-agriculture, is

$$o(z_a, z_n) = \begin{cases} 0, & \text{if } \theta\, w_n z_n \;\geq\; w_a z_a\left(\mathbb{E}[v^{1-\varphi}]\right)^{\frac{1}{1-\varphi}}, \\ 1, & \text{otherwise.} \end{cases} \tag{13}$$

Consumption allocations under PIGL preferences. Given realized income I, the budget constraint is

$$p_a c_a + c_n = I, \qquad c_a \geq 0,\; c_n \geq 0. \tag{14}$$

PIGL preferences imply Hicksian demand of the form

$$c_a(I, p_a) = \begin{cases} \nu\, I^{\varphi}\, p_a^{-\eta} & \text{if } p_a c_a < I, \\ \frac{I}{p_a} & \text{otherwise.} \end{cases} \tag{15}$$

Given $c_a(I, p_a)$, $c_n(I, p_a)$ is solved from the budget constraint in equation (14).

4.2 Equilibrium

A *competitive equilibrium* comprises of a set of prices (p_a, w_a, w_n), occupational choice $o(z_a, z_n)$, individual income $I(z_a, z_n, v)$, consumption allocations $(c_a(p_a, I), c_n(p_a, I))$ and sectoral allocations (H_a, H_n, N_a, N_n) such that:

(i) Given prices, profit maximization in each sector implies that wages equal the marginal products of effective labor, given by equation (5).

(ii) Given prices and ability, occupational choice $o(z_a, z_n)$ maximizes individual's expected utility, and satisfies equation (13), and yields individual's income

$$I(z_a, z_n, v) = o(z_a, z_n) w_a z_a v + (1 - o(z_a, z_n)) \theta w_n z_n. \tag{16}$$

(iii) Given prices and realized income, consumption allocations solve the household's problem in equations (15) and (14).

(iv) Markets clear:

$$H_a = \int_{(z_a, z_n)} z_a o(z_a, z_n) dF(z_a, z_n), \quad H_n = \int_{(z_a, z_n)} z_n (1 - o(z_a, z_n)) dF(z_a, z_n).$$

$$N_a = \int_{(z_a, z_n)} o(z_a, z_n) dF(z_a, z_n), \quad N_n = 1 - N_a.$$

$$Y_a = A\kappa H_a = \int_{(z_a, z_n)} c_a(z_a, z_n)\, dF(z_a, z_n), \quad Y_n = A H_n = \int_{(z_a, z_n)} c_n(z_a, z_n)\, dF(z_a, z_n).$$

Agricultural income gap (AIG). Given an equilibrium allocation, average income in agriculture and non-agriculture are given by

$$\bar{I}_a = \frac{1}{N_a} \mathbb{E}[v] \int_{(z_a, z_n)} o(z_a, z_n)\, \tilde{I}_a(z_a, z_n)\, d\mu(z_a, z_n), \qquad \tilde{I}_a(z_a, z_n) = p_a A \kappa\, z_a, \tag{17}$$

$$\bar{I}_n = \frac{1}{N_n} \int_{(z_a, z_n)} (1 - o(z_a, z_n))\, I_n(z_a, z_n)\, d\mu(z_a, z_n), \qquad I_n(z_a, z_n) = A\, z_n. \tag{18}$$

We define the agricultural income gap (AIG) as

$$\text{AIG} = \frac{\bar{I}_a}{\bar{I}_n}. \tag{19}$$

In the quantitative analysis, we study how AIG, sectoral income distributions, and poverty rates vary with (A, κ, θ), holding fixed the joint ability distribution $F(z_a, z_n)$ and the agricultural shock distribution $G(v)$.

5 Quantitative Analysis

We calibrate a benchmark economy to match data moments for the United States in 1950. We chose the USA in 1950 as this predates the substantial poverty-alleviation policies implemented in the 1960s (Bailey and Duquette, 2014). We then conduct a series of cross-country experiments to assess the roles of economy-wide productivity, agricultural productivity, and sectoral distortions in shaping poverty and macroeconomic outcomes across countries.

5.1 Calibration

We assume that sectoral abilities z_a and z_n are drawn from a joint lognormal distribution with normalized means and standard deviations σ_a, σ_n, and correlation ρ_{an}. The ex-post agricultural productivity shock v is also log-normal with mean $\mathbb{E}[v] = 1$ and standard deviation σ_v.

The model features ten parameters to calibrate: the dispersion of agricultural and non-agricultural abilities (σ_a, σ_n), the correlation of sectoral abilities ρ_{an}, the dispersion of ex-post agricultural productivity shock σ_v, the PIGL preference parameters (φ, η, ν), productivity parameters (A, κ), and labor mobility barrier θ.

A set of six parameters is either normalized or disciplined by outside evidence. We set the income-elasticity parameter to $\varphi = 0.70$ and the price-elasticity parameter to $\eta = 0.40$ to match Engel curve patterns for agricultural expenditure shares (Boppart, 2014). We set

$\rho_{an} = 0.36$ based on estimates in literature (Lagakos and Waugh, 2013). Economy-wide and agricultural productivity parameters are normalized to one: $A = \kappa = 1$. The mobility friction θ is also normalized to one in the benchmark economy.

We calibrate the remaining four parameters $(\sigma_n, \sigma_a, \sigma_v, \nu)$ to match four moments from the U.S. 1950 Census microdata: (1) the share of employment in agriculture, (2) the agricultural income gap (average income in agriculture divided by average income in non-agriculture), (3) the standard deviation of log income, and (4) the standard deviation of log income in agriculture. Calibration targets are computed directly from the U.S. 1950 Census microdata using household income adjusted for household size (OECD equivalence scale), with sector assigned by the household head's industry. See Table 1 for a summary of parameter values and targeted moments.

Table 1: Calibration of Benchmark Economy to USA 1950

Parameter	Description	Value	Target moment	Model	Data
σ_n	Non-ag ability dispersion	1.236	SD log non-ag income	1.185	1.184
σ_a	Ag ability dispersion	1.318	Ag income gap	0.596	0.575
σ_v	Ex-post ag productivity	0.757	SD log ag income	1.350	1.350
ν	Ag preference weight	0.264	Ag emp share (%)	10.7	10.7

Poverty line. The poverty line is set at the 6th percentile of the overall income distribution in the benchmark economy, corresponding to the World Bank \$3.20/day line in 2011 PPP (the US 1950 poverty rate at this threshold is 7.7% in the data). The model generates a 6% poverty rate at the benchmark, closely matching the data.

To compare poverty across economies with different relative agricultural prices p_a, we use the PIGL indirect utility function to compute *equivalent incomes*. Specifically, we find the income level in each counterfactual economy that delivers the same utility as the poverty line income in the benchmark economy:

$$V(y^*_{\text{ctf}}, p_a^{\text{ctf}}) = V(y^*_{\text{BE}}, p_a^{\text{BE}}),$$

where ctf and BE refer to the counterfactual and the benchmark, respectively. This is a welfare-consistent equivalent variation approach that generates an *income-dependent* cost-of-living index. Unlike standard PPP conversions, which apply a uniform price adjustment regardless of income level, the PIGL-based conversion accounts for the fact that poor households spend a larger share of their income on food (Engel's law). When agricultural prices are high—as in poor economies—the welfare cost falls disproportionately on the poor, and our approach captures this non-homotheticity.

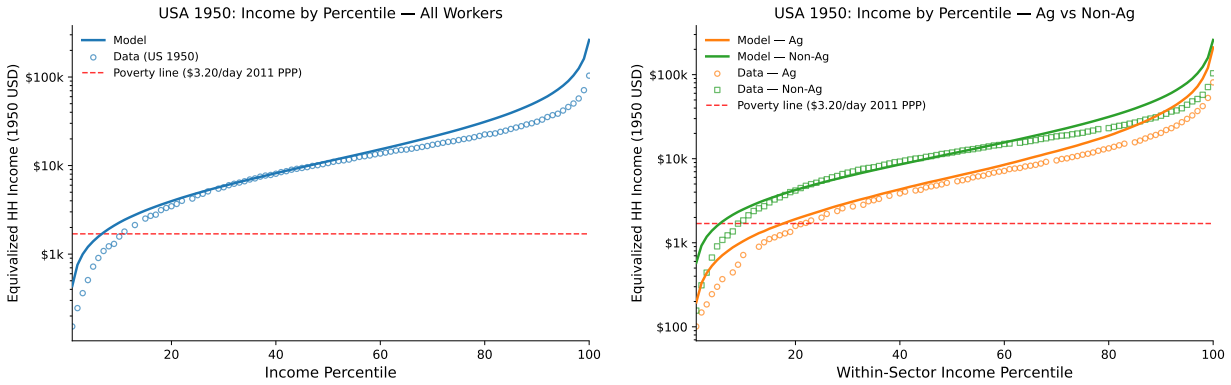
Model fit. Figure 14 documents income by percentile in the benchmark economy, comparing the calibrated model (lines) against USA 1950 microdata (scatter points). Panel (a) shows the overall income distribution. The model replicates the shape of the income distribution well. Panel (b) shows agricultural and non-agricultural incomes separately, where percentiles are computed within each sector. The model captures the key pattern that income in the agricultural sector is systematically lower than in the non-agricultural sector across all percentiles and that there is a large overlap in the sectoral income distributions. The income gap between sectors is a central feature of the data and implies that poverty is concentrated in agriculture.

Model implications of income by percentiles across development levels. We use the model to compare income distributions at development levels of the USA in 1950 (benchmark) with those of Brazil in 1980, and Malawi in 2004. To this end, we select A, κ, and θ values, which are all normalized to 1 for the benchmark, so that the counterfactual economy has the GDP per worker observed in the data.[1] Income is converted to PPP 2011 USD using the model's poverty-line anchoring.

The model reproduces the key pattern from the data: poverty is mainly driven by high poverty rates in agriculture. In the benchmark economy (USA 1950, Figure 14), the model

[1]The values for$A$, κ, and θ are 0.4531, 0.3689, and 0.2643 for Brazil in 1980, and 0.2000, 0.1500, and 0.1000 for Malawi in 2004

Figure 14: Income by Percentiles in Benchmark Economy: Model vs. Data (USA 1950)

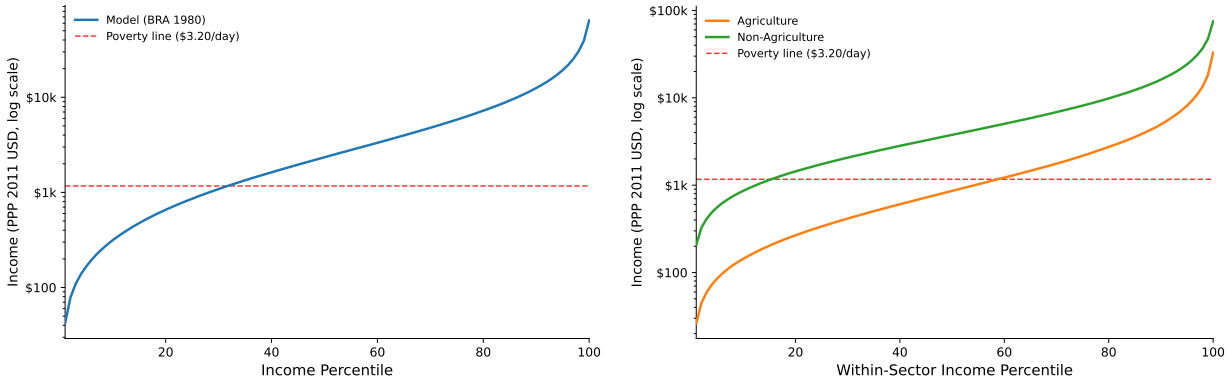


(a) All workers (b) Agriculture vs. Non-Agriculture

Notes: Lines show model implications; hollow markers show USA 1950 Census microdata (IPUMS). Income is household income adjusted for household size, with sector assigned by the household head's industry. Panel (b) uses within-sector percentiles. Model income is normalized to 1950 USD by matching median income. Dashed red line indicates the World Bank $3.20/day poverty line (2011 PPP).

generates a poverty rate of 6% with an agricultural employment share of 10.8%. Agricultural income is systematically below non-agricultural income at every percentile, consistent with the data pattern that poverty is concentrated in agriculture.

Figure 15: Model: Income by Percentile — Brazil 1980



(a) All workers (b) Agriculture vs. Non-Agriculture

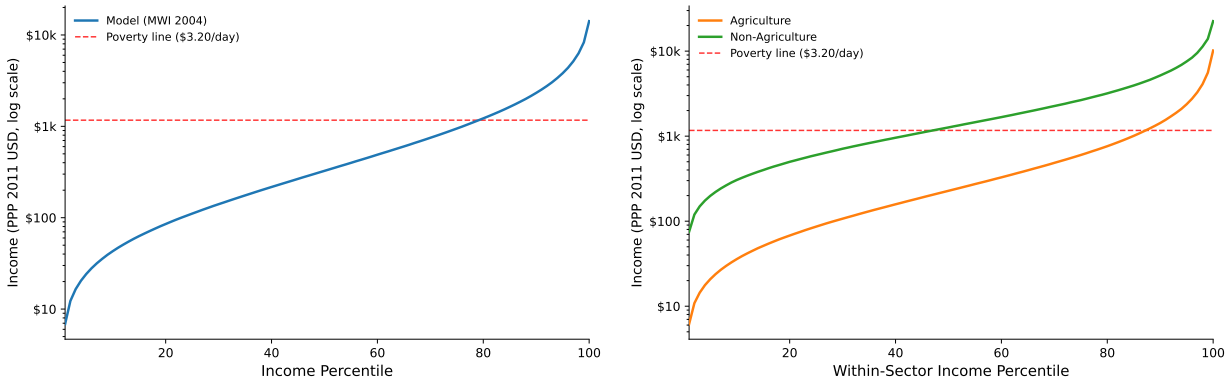
Notes: Model counterfactual matched to Brazil 1980 by GDP per worker (PWT 11.0). Model poverty rate 31%, ag share 38%. Income converted to PPP 2011 USD via poverty-line anchoring. Panel (b) uses within-sector percentiles. Dashed red line is the World Bank $3.20/day poverty line.

At the Brazil 1980 development level (Figure 15), the model generates a poverty rate of 31% with an agricultural employment share of 38%. Agricultural income is again below

non-agricultural income at every percentile, and the poverty line cuts through the bulk of the agricultural income distribution.

At the Malawi 2004 development level (Figure 16), the model generates a poverty rate of 79% with an agricultural employment share of 80%. Both agricultural and non-agricultural incomes are very low, but agriculture is still systematically lower. Even at very low development levels, the sectoral income gap remains a major driver of poverty.

Figure 16: Model: Income by Percentile — Malawi 2004



(a) All workers

(b) Agriculture vs. Non-Agriculture

Notes: Model counterfactual matched to Malawi 2004 by GDP per worker (PWT 11.0). Model poverty rate 79%, ag share 80%. Income converted to PPP 2011 USD via poverty-line anchoring. Panel (b) uses within-sector percentiles. Dashed red line is the World Bank $3.20/day poverty line.

5.2 Cross-Country Experiments

We compute cross-country economies by jointly varying three parameters across a grid of counterfactual economies:

- Economy-wide productivity A: varied from 0.20 to 1 in the benchmark economy to 3.5 to generate a range of non-agricultural labor productivity of approximately 1/5 between the benchmark economy and poorest economy (Restuccia et al., 2008).

- Agricultural productivity κ: varied from 0.15 to 3.5 to match the range of agricultural employment shares (N_a) observed between rich and poor countries. In this context κ

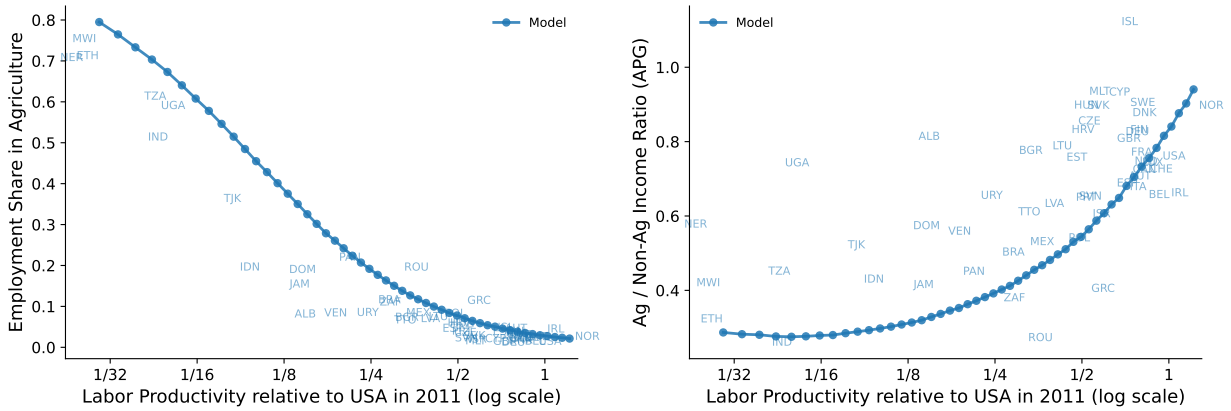
is a catch-all for policies and institutions that affect agricultural productivity such as misallocation within agriculture (Restuccia et al., 2008; Adamopoulos and Restuccia, 2014; Chen et al., 2023).

- Sectoral distortion θ: varied from 0.10 to 3.0 to match the range of the agricultural productivity gap (APG) between rich and poor countries (Gollin et al., 2014).

For each counterfactual economy, we solve for the equilibrium and compute poverty rates, income dispersion, agricultural employment share, and other moments. We compare the model-implied moments against cross-country data from 54 countries, using labor productivity (GDP per worker) from PWT 11.0 normalized to USA in 2011 as the common x-axis.

Model cross-country fit. We overlay model-implied economies on the cross-country data. The model replicates the key cross-country patterns well. The model captures the strong negative relationship between labor productivity and the share of employment in agriculture (Figure 17a), the agricultural productivity gap (Figure 17b).

Figure 17: Model vs. Data: Agriculture Share and Agricultural Income Gap

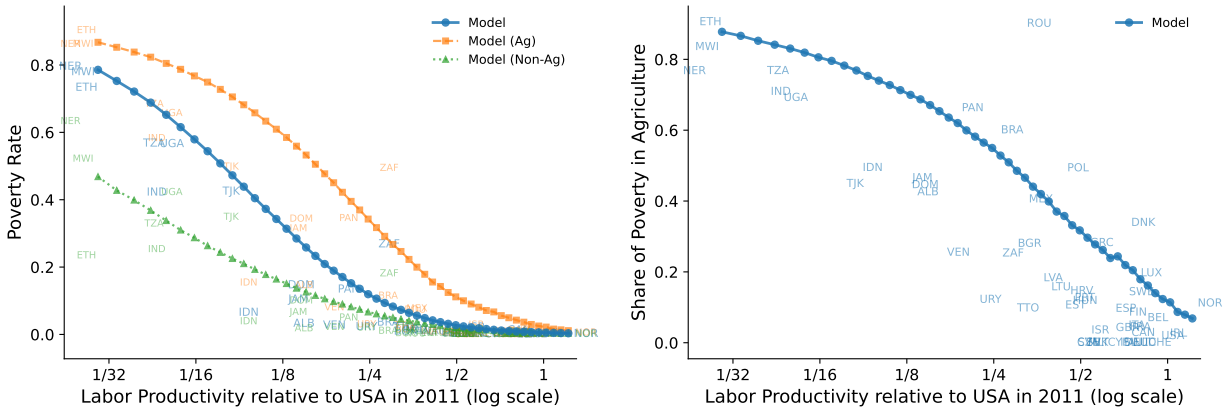


(a) Employment share in agriculture (b) Ag / Non-Ag income ratio

Notes: Lines with markers represent model economies. ISO3 labels represent cross-country data. Share of poverty in agriculture is the fraction of all poor individuals who are employed in the agricultural sector. Data from household microdata; aggregate labor productivity from PWT 11.0.

Even though the variation in (A, κ, θ) does not target poverty rates, the model economies replicates the cross-country patterns in poverty rates in aggregate and by sector (Figure 18a).

Figure 18: Model vs. Data: Poverty Rate and Share of Poverty in Agriculture



(a) Poverty rate

(b) Share of Poverty in Agriculture

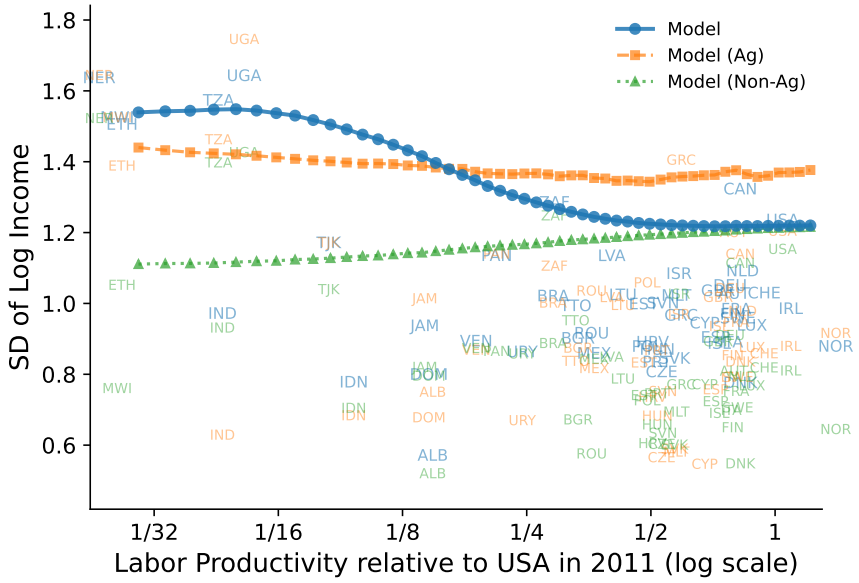
Notes: Lines with markers represent model counterfactuals. ISO3 labels represent cross-country data. Share of poverty in agriculture is the fraction of all poor individuals who are employed in the agricultural sector. Data from household microdata; labor productivity from PWT 11.0.

Poverty rates for all workers, and separately for agricultural and non-agricultural workers, decline sharply with labor productivity, consistent with the data. The model captures the level and slope of poverty rates across the development spectrum. The model economies also match well the concentration of poverty in agriculture, measured by the share of poverty accounted for by the agricultural sector. In particular, the model generates the steep rise in poverty composition toward agriculture at lower development levels (Figure 18b).

For income dispersion, the model matches the relatively flat pattern in the standard deviation of log income within each sector and the larger dispersion in agriculture than in non-agriculture (Figure 19). Larger income inequality in less developed economies arises from the reallocation of employment to agriculture where income dispersion is larger.

Agricultural share of employment by income percentile. A key non-targeted moment is the distribution of agricultural employment across the income distribution. Using household-level microdata for the USA (1950), Brazil (1980), and Malawi (2004), we compute the share of household heads employed in agriculture at each income percentile. For each country, we match the closest model counterfactual by labor productivity relative to

Figure 19: Model vs. Data: SD of log Income



Notes: Standard deviation of log household income per capita. Lines with markers represent model experiments; ISO3 labels represent cross-country data. Sector assigned by the household head's industry.

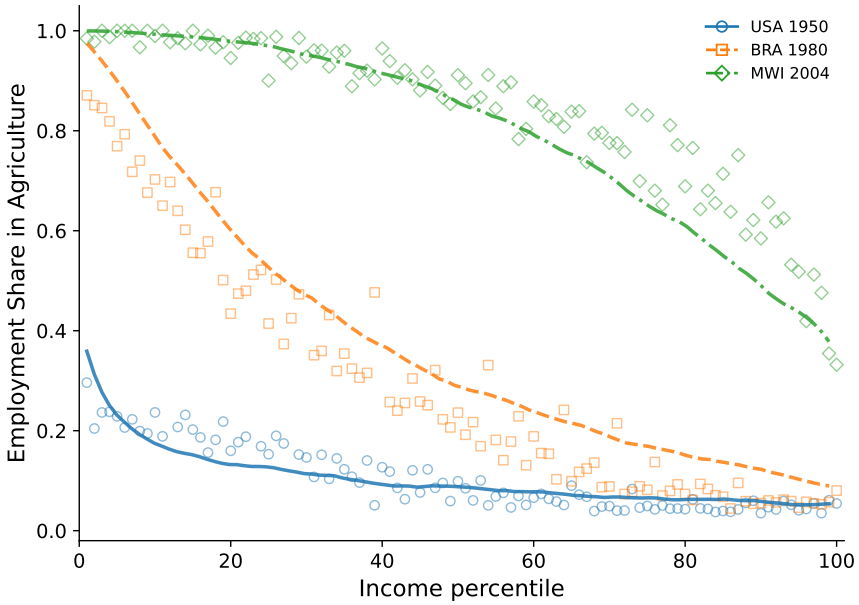
the United States in 2011. For USA 1950, the model counterpart is the benchmark economy itself.

Figure 20 plots the results. The model (lines) captures the key qualitative patterns in the data (scatter points) remarkably well: (i) within each economy, the share of employment in agriculture declines monotonically with income, (ii) the profile is flatter and high in the poorest economy (Malawi), steeply declining in the middle-income economy (Brazil), and low and relatively flat in the richest economy (USA), and (iii) the overall level of agricultural employment shifts dramatically across development levels.

Within-country variation over time. We further examine the model's ability to match within-country changes over time. For the United States, we compare model predictions against microdata from 1950 and 2015. For Brazil, we use 1980 and 2010. In each case, we match the model counterfactual whose output per worker level (relative to the benchmark) is closest to the country's labor productivity in that year.

Figure 21a shows the USA over time. As the economy grows from 1950 to 2015, both data

Figure 20: Model vs. Data: Agricultural Share of Employment by Income Percentile



Notes: Hollow markers show data from household microdata (IPUMS for USA and Brazil; LSMS-ISA for Malawi). Lines show model counterparts: benchmark economy for USA 1950, closest counterfactual by labor productivity for Brazil 1980 and Malawi 2004. Income percentiles based on household income adjusted for household size. Model curves smoothed with a Savitzky-Golay filter.

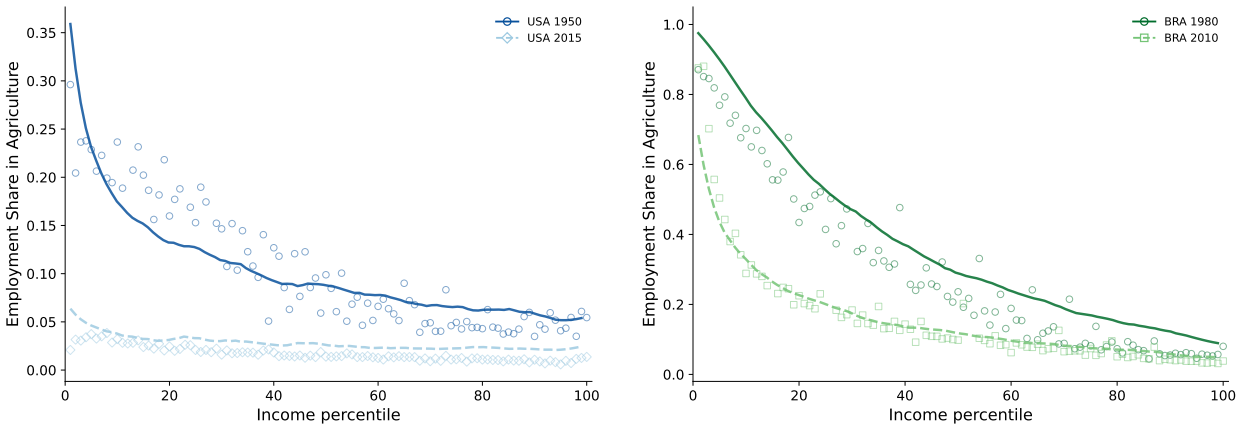
and model show a dramatic decline in agricultural employment across all income percentiles. Figure 21b shows the same pattern for Brazil between 1980 and 2010.

5.3 Counterfactual Decomposition

To isolate the contribution of economy-wide productivity A, agriculture-specific productivity κ, and sectoral mobility wedge θ to poverty reduction and structural transformation, we perform counterfactual decompositions starting from the poorest economy on the joint parameter set ($A = 0.20$, $\kappa = 0.15$, $\theta = 0.10$). We sweep parameters to the richest economy using log-linear interpolation in 40 steps, considering three scenarios: varying A only, varying A and κ, and varying θ only. We contrast these decompositions with the baseline where all parameters are changed.

Figure 22 reports the results of these decompositions. Panel (a) shows the share of employment in agriculture and panel (b) the poverty rate in each case. For both structural

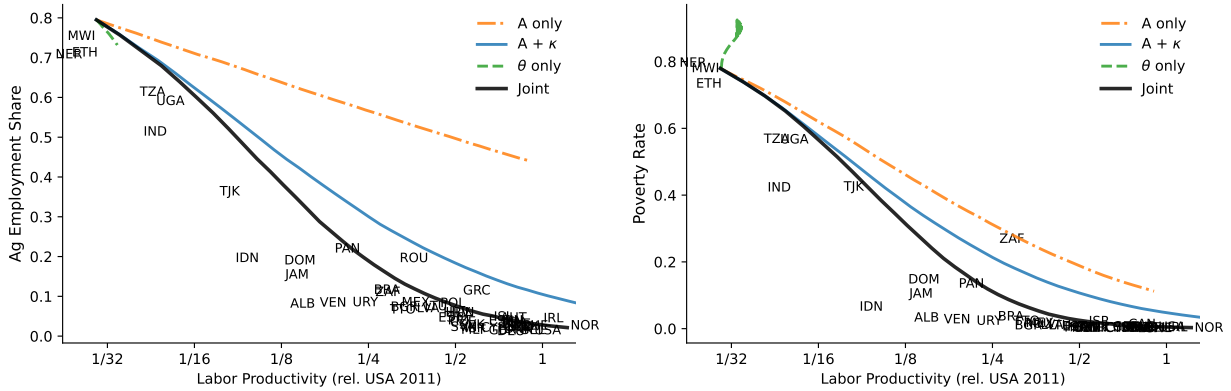
Figure 21: Agricultural Share of Employment by Income Percentile over Time



Notes: Shades of blue denote USA for different years (dark = 1950, light = 2015). Shades of green denote BRA for different years (dark = 1980, light = 2010). Data from IPUMS USA Census microdata. Model economies matched by labor productivity (PWT 11.0). Model curves smoothed with a Savitzky-Golay filter.

transformation and poverty reduction, economy-wide productivity A is the dominant driver.

Figure 22: Counterfactual Decomposition



(a) Agricultural employment share (b) Poverty rate

Notes: Each line shows the agricultural employment share and the poverty rate as a function of labor productivity (relative to USA 2011) when the indicated parameter(s) changes from the poorest to richest values while all others are held at their poorest-economy levels. "Joint" counterfactual changes all three parameters simultaneously. ISO3 labels show cross-country data.

Increasing A alone from its lowest to highest value reduces agricultural employment and the poverty rate substantially, accounting for the majority of the joint path's total decline. Note also that even though A increases from 0.2 in the poorest economy to one in the benchmark, aggregate labor productivity on this factor alone changes from 0.09 to 0.66,

covering an important portion of the range observed across countries.

Agricultural-specific productivity κ also contributes meaningfully to structural transformation and poverty reduction through its effect on income and on agricultural wages. The main take away is that productivity differences through A and κ alone capture the bulk of the cross-country variation in aggregate labor productivity, structural transformation, and poverty.

The sectoral mobility wedge θ, by contrast, has minimal impact on both structural transformation, aggregate productivity, and poverty. Reducing the mobility friction alone improves the selection of workers across sectors, raising average productivity and promoting structural transformation, but the quantitative effect is relatively small when productivity is low. That is, in the poorest economy—at very low A and κ, reshuffling workers between two low-productivity sectors does not materially change living standards.

Regarding the effect of θ on the poverty rate, on its own, reducing the mobility friction actually increases poverty in our model as illustrated in Figure 22b. Lower mobility barriers makes non-agriculture relatively more attractive for workers, but because productivity remains low, in equilibrium the relative price of agriculture must increase substantially in order to support a relatively high agricultural employment. It is possible that with price support policies, the poverty rate may not increase, but as with the share of employment in agriculture, the quantitative effect of reducing mobility friction on its own is limited.

Table 2 reports the counterfactual decomposition results. Starting from the low-income economy with aggregate labor productivity of 9% relative the benchmark (column 1), increasing A to its benchmark value alone (column 2) reduces the agricultural share of employment by 21.7 percentage points (from 79.5% to 57.8%) and cuts the poverty rate by 44.3 percentage points (from 77.9% to 33.6%), while aggregate labor productivity rises sevenfold to 0.66 relative to benchmark—substantially reducing the aggregate labor productivity gap with the benchmark economy (column 5). Increasing also κ to its benchmark level (column 3) induces further structural transformation, reducing agricultural employment by an additional 34.2

Table 2: Counterfactual Decomposition from Poorest Economy

	(1) Low-income economy	(2) Increase A only	(3) Increase (A, κ) only	(4) Increase θ only	(5) Benchmark all
Aggregate labor prod.	0.09	0.66	0.98	0.10	1.00
Ag. labor prod.	0.01	0.07	0.74	0.01	1.00
Non-ag. labor prod.	0.37	1.45	1.10	0.35	1.00
Ag. emp. share (%)	79.5	57.8	23.6	75.1	10.6
Poverty rate (%)	77.9	33.6	16.0	87.0	5.6
Relative price ag. p_a	2.89	1.12	0.04	23.15	0.19

Notes: Column (1) "Low-income economy" corresponds to the economy with parameters ($A = 0.20$, $\kappa = 0.15$, $\theta = 0.10$). Column (2) changes only A to one while holding the other parameters unchanged at the low-income economy. Column (3) changes (A, κ) to one while holding θ at the low-income economy level. Column (4) changes only θ to one while holding (A, κ) at the low-income economy level. (5) Benchmark all sets all parameters to their benchmark values ($A = \kappa = \theta = 1$). Aggregate and sectoral labor productivity are reported as relative to the benchmark economy.

percentage points (to 23.6%) and reducing the poverty rate by an additional 17.6 percentage points (to 16.0%), with aggregate productivity reaching 0.98, reducing almost entirely the aggregate gap with the benchmark economy. The reduction in the aggregate productivity gap is driven largely by a 73-fold rise in agricultural labor productivity from 0.01 to 0.74. The small remaining productivity gap is due to removing sectoral mobility frictions θ.

Column (4) isolates the role of θ: reducing mobility friction alone leaves aggregate productivity nearly unchanged at 0.10 and barely moves agricultural employment, declining by only 4.4 percentage points (to 75.1%), while the poverty rate rises by 9.1 percentage points to 87.0%—even above the low-income economy. This increase in poverty rate results from the general equilibrium effect on the relative price of agriculture p_a, which surges eightfold from 2.89 to 23.15, as discussed earlier.

We evaluate the effect on poverty of reducing mobility barriers in the poorest economy under different scenarios for the change in the relative price of agriculture. Table 3 reports the

Table 3: Poverty Rate and Mobility Barriers θ under Alternative p_a

	(1)	Increase $\theta = 1$ only		
	Low-income economy	(2) GE p_a (8×)	(3) PE p_a (4×)	(4) PE p_a (1×)
Relative price ag. p_a	2.89	17.31	11.54	2.89
Poverty rate (%)	77.9	81.2	71.4	38.3

Notes: Column (1) reports the low-income economy ($A = 0.20$, $\kappa = 0.15$, $\theta = 0.10$). Columns (2)–(4) all set $\theta = 1$ while holding (A, κ) and all allocations fixed at the Increase-θ general equilibrium. They differ only in the price used to compute household income: column (2) uses the general equilibrium (GE) price $p_a = 23.15$ (approximately 8× the low-income level); columns (3)–(4) impose counterfactual partial equilibrium (PE) prices of 4× and 1× the low-income relative price p_a and recomputes utility-equivalent income for households under these different price scenarios.

results when the price of agriculture increases less than in the general equilibrium economy, including the case where the relative price of agriculture remains the same as in the economy with high mobility barriers. Reducing mobility barriers makes the more productive non-agricultural sector more attractive, raising average income measured in units of the non-agricultural good. However, this reallocation also drives up the relative price of agricultural goods in general equilibrium—and it is this price increase that accounts for the adverse poverty effect. Under the GE price response of 8× (column 2), the poverty rate rises by 9.1 percentage points despite the income gains from reallocation. Even halving the price response to 4× (column 3) yields only a modest 6.5 percentage point decline in the poverty rate relative to the low-income baseline. Only when the general equilibrium price effect is shut down entirely (column 4) does reducing the mobility friction produce a large decline in poverty of 39.6 percentage points. The limited—and even negative—effect of θ on poverty therefore reflects a fundamental tension: reducing mobility friction without accompanying improvements in agricultural productivity results in rising food prices, affecting lower income households the most.

6 Conclusions

This paper brings a macroeconomic perspective to poverty by combining harmonized household income microdata from 91 countries and 192 country-years with a parsimonious model of structural transformation. The data reveal a set of tightly connected facts. Poverty declines sharply with development, but it is also highly uneven across sectors: it is disproportionately concentrated in agriculture, where workers earn less and where poor households are overrepresented. In poorer economies, agriculture accounts for both a much larger share of employment and a much larger share of total poverty. These patterns suggest that cross-country poverty differences cannot be understood solely as differences in within-country redistribution; they are also a consequence of low productivity and of the sectoral organization of work.

To interpret these facts, we developed a two-sector Roy model with heterogeneous sectoral abilities, non-homothetic demand, and a sectoral mobility wedge. Calibrated to the United States in 1950 and used for cross-country quantitative experiments, the model accounts well for the joint behavior of poverty, agricultural employment, income dispersion, and the agricultural income gap. Quantitatively, economy-wide productivity emerges as the dominant force behind poverty reduction across countries: when overall productivity rises, the entire income distribution shifts upward, and workers move out of agriculture as the income elasticity of food demand falls. Agricultural productivity also plays an important role, especially by raising incomes within agriculture. By contrast, sectoral distortions have comparatively limited effects on poverty at very low levels of development, when both sectors are simply too unproductive for reallocation alone to generate large gains in living standards.

The broader implication is that cross-country differences in poverty are primarily a productivity phenomenon, magnified by the concentration of poor households in agriculture. This does not make barriers, frictions, or distributional factors unimportant; rather, it shows that their impact on poverty depends crucially on how they affect the level of productivity that determines the position of the entire income distribution. More generally, the paper highlights that the study of poverty, structural transformation, and cross-country productiv-

ity differences should not be separated. A full account of why poor households remain poor across countries requires understanding why entire economies remain unproductive and why so many workers remain in low-income agricultural activities.

References

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498.

Adamopoulos, T., Brandt, L., Leight, J., and Restuccia, D. (2022). Misallocation, selection, and productivity: A quantitative analysis with panel data from china. *Econometrica*, 90(3):1261–1282.

Adamopoulos, T. and Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6):1667–97.

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2018). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual*, 32:1–75.

Bailey, M. J. and Duquette, N. J. (2014). How johnson fought the war on poverty: The economics and politics of funding at the office of economic opportunity. *Journal of Economic History*, 74(2):351–388.

Balboni, C., Bandiera, O., Burgess, R., Ghatak, M., and Heil, A. (2022). Why do people stay poor? *The Quarterly Journal of Economics*, 137(2):785–844.

Banerjee, A. V. and Newman, A. F. (1993). Occupational choice and the process of development. *Journal of Political Economy*, 101(2):274–298.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Caselli, F. (2005). Accounting for income differences across countries. *Handbook of economic growth*, 1:679–741.

Chari, V. V., Kehoe, P. J., and McGrattan, E. (1996). The poverty of nations: a quantitative exploration.

Chen, C., Restuccia, D., and Santaeulàlia-Llopis, R. (2023). Land misallocation and productivity. *American Economic Journal: Macroeconomics*, 15(2):441–465.

Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.

Dollar, D., Kleineberg, T., and Kraay, A. (2016). Growth still is good for the poor. *European Economic Review*, 81:68–85. Model Uncertainty in Economics.

Duarte, M. and Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *The Quarterly Journal of Economics*, 125(1):129–173.

Erosa, A., Koreshkova, T., and Restuccia, D. (2010). How important is human capital? a quantitative theory assessment of world income inequality. *The Review of Economic Studies*, 77(4):1421–1449.

Galor, O. and Zeira, J. (1993). Income distribution and macroeconomics. *The Review of Economic Studies*, 60(1):35–52.

Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.

Gollin, D., Parente, S., and Rogerson, R. (2002). The Role of Agriculture in Development. *American Economic Review*, 92(2):160–164.

Gomez, M. (2025). Macro perspectives on income inequality. *Journal of Economic Perspectives*, 39(2):127–48.

Guner, N. and Ruggieri, A. (2022). Misallocation and inequality. CEPR DP 17113.

Guvenen, F., Pistaferri, L., and Violante, G. L. (2022). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics*, 13(4):1321–1360.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2014). Growth and structural transformation. In *Handbook of economic growth*, volume 2, pages 855–941. Elsevier.

Jones, C. I. (2016). The facts of economic growth. In *Handbook of macroeconomics*, volume 2, pages 3–69. Elsevier.

Klenow, P. J. and Rodriguez-Clare, A. (1997). The neoclassical revival in growth economics: Has it gone too far? *NBER macroeconomics annual*, 12:73–103.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *The Review of Economic Studies*, 68(4):869–882.

Krueger, D., Perri, F., Pistaferri, L., and Violante, G. L. (2010). Cross-sectional facts for macroeconomists. *Review of Economic dynamics*, 13(1):1–14.

Kuznetz, S. (1955). Economic growth and income inequality. *American Economic Review*, 45(1):1–28.

Lagakos, D., Moll, B., Porzio, T., Qian, N., and Schoellman, T. (2018). Life cycle wage growth across countries. *Journal of Political Economy*, 126(2):797–849.

Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *The American Economic Review*, 103(2):948–980.

McGrattan, E. R. and Schmitz Jr, J. A. (1999). Explaining cross-country income differences. *Handbook of Macroeconomics*, 1:669–737.

Milanovic, B. (2015). Global inequality of opportunity: How much of our income is determined by where we live? *Review of Economics and Statistics*, 97(2):452–460.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American Economic Review*, 97(1):429–443.

Prescott, E. C. (1998). Lawrence r. klein lecture 1997: Needed: A theory of total factor productivity. *International economic review*, pages 525–551.

Restuccia, D., Yang, D. T., and Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics*, 55(2):234–250.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2):135–146.