The Macroeconomics of Intensive Agriculture

Timo Boppart, Patrick Kiernan, Per Krusell, Hannes Malmberg

March, 2019

Abstract: In poor countries, the agricultural sector is large, uses few inputs per worker, and is unproductive, even relative to the local economy. We jointly study these phenomena using a neoclassical two-sector model. To inform our choice of an agricultural production function, we draw on micro and macro sources to construct a cross-country database of agricultural input use and prices. Using our framework, we separately analyze the consequences of closing the sectoral TFP gap in agriculture and non-agriculture between a poor country and the US. We find that closing the agricultural TFP gap reduces GDP per worker gap from a factor of 40 only to a factor of 14.1, whereas closing the non-agricultural TFP gap reduces the GDP per worker gap to 2.5. This finding is in stark contrast to those of simple decomposition exercises, which suggest that closing agricultural labor productivity gaps could dramatically reduce income differences. The difference emerges from general equilibrium forces. In general equilibrium, agricultural TFP improvements have diminishing marginal returns since rising incomes and falling agricultural prices reduce the importance of the agricultural sector. In contrast, non-agricultural TFP improvements have increasing marginal returns, as income effects and agricultural input intensification both work to increase the importance of the non-agricultural sector. Our results show the importance of taking a general equilibrium perspective on agricultural labor productivity when intermediate inputs and machinery are produced in other sectors. Furthermore, they show that even when agriculture is relatively unproductive, the main impediment to development may lie outside the agricultural sector.

1 Introduction

Poor countries employ most of their workforce in agriculture. At the same time, poor countries' labor productivity — relative to rich countries — is particularly low in agriculture (Restuccia et al., 2008; Gollin et al., 2014). These facts have attracted significant attention in the macroeconomic development literature, as they seem to suggest that having a large number of unproductive workers in agriculture may lie at the heart of the poverty of the developing world.[1]

However, should policy focus on agriculture just because agricultural labor productivity is low? Not necessarily. Labor productivity is not equal to total factor productivity (TFP), and there is strong evidence for dramatic factor intensification in agricultural production. In the US, the average worker in agriculture uses (in real terms) $276,000 of capital, $56,000 of non-agricultural intermediate inputs, and 73 hectares of land. In Senegal, the corresponding real quantities are $550 for capital, $172 for intermediate inputs, and 2.2 hectares for land.

Once input intensification is taken into account, agricultural labor productivity cannot be studied in isolation. If agriculture uses inputs from non-agricultural sectors, productivity improvements in these other sectors will improve agricultural labor productivity by inducing substitution away from labor through an increase in the relative price of labor. Furthermore, due to non-homothetic preferences, increases in income levels reduce the relative demand for agricultural products, which makes land relatively cheaper. This further stimulates factor intensification since labor is substituted for land.

In this paper, we analyze the macroeconomic implications of input intensification in agriculture. We start by combining data from aggregate and microeconomic sources to document how the prices and quantities of agricultural factor inputs vary between poor and rich countries. For prices, we find that the price levels of capital goods and intermediate inputs are roughly constant across countries. In contrast, land rental rates increase with income, but not as fast as labor costs. For quantities, we find that the intensification of different factors mirrors the changes in relative prices. As one moves from poorer to richer countries, all factors are used more intensely relative to labor. Capital and intermediate inputs are also used more intensely relative to land, which is consistent with the increase in the price of land relative to the prices of capital and intermediate goods.

Using the price and quantity data, we also calculate the elasticity of input intensities with respect to relative factor prices.[2] We find that capital and intermediate input intensities have a similar elasticity of -1.85 with respect to their prices relative to labor. These findings are consistent with the observations by Chen (2017) and Donovan (2016) who observe that capital use and intermediate input use grow faster in agriculture than in the overall economy. For land, we find that the use of land relative to labor has an elasticity of -1.57 with respect to the relative price of land and labor.

We use our findings to construct a gross output production function in agriculture consistent with the cross-sectional price and quantity data. We then incorporate this production function into a standard structural change model with non-homothetic preferences and a Cobb-Douglas production function in non-agriculture. Countries differ in their agricultural and non-agricultural TFP levels and are assumed to be in

[1]Macroeconomic papers that have studied the sources of low agricultural labor productivity include, among others, Restuccia et al. (2008), Adamopoulos and Restuccia (2014), Donovan (2016), Caunedo and Keller (2017), and Chen (2017).

[2]Loosely, this corresponds to an elasticity of substitution, but since the cross-section of countries features simultaneous changes in multiple relative factor prices, the cross-section does not identify elasticities of substitution unless we put more structure on the production function.

steady state.

Apart from constructing a non-agricultural production function that fits cross-sectional patterns, we use standard parameters for preferences, with the non-homotheticity calibrated from cross-sectional micro expenditure data. Doing this, we find that our model fits the movement of labor out of agriculture very well. We also find this fit is not just an artifact of our calibration of the production function, but depends crucially on having the correct micro-calibrated non-homotheticity parameter. We also show that the fit of the structural change pattern is robust: as long as the model matches the pattern of input prices and quantities, as well as the micro-calibrated non-homotheticity parameter, the model also broadly fits the pattern of structural change regardless of the other parameter choices in the model.

We use our calibrated model to conduct counterfactual experiments. To do this, we set up a benchmark poor country that is calibrated to have a GDP per worker that is 40 times lower than the US level. Using this country, we test the effect of different manipulations of its agricultural and non-agricultural productivities.

The first experiment is a simple decomposition where we examine the output effect of closing the gaps in sectoral labor productivities, separately considering the effect of closing agricultural and non-agricultural productivity gaps. This exercise ignores general equilibrium effects such as structural change, changing relative prices, and input intensification. In this exercise, we find that removing agricultural productivity differences reduces the output gap relative to the US by a factor of 25, whereas removing non-agricultural productivity differences only reduces the output gap by a factor of 17. This finding reflects that the agricultural productivity gap is larger than the non-agricultural productivity gap, 71 versus 24, and that 60% of the workforce is in agriculture in the poor country. Thus, our simple decomposition is consistent with earlier such exercises that found that agricultural productivity differences were central in explaining aggregate income differences, e.g., in Caselli (2005).

However, this conclusion does not hold in our second experiment, which is a model-based counterfactual. In this exercise, we compare the effects of closing the agricultural versus non-agricultural *TFP gaps* between the poor and the rich country. The difference in results is striking. While there still is a large effect from increasing non-agricultural TFP to US levels — the income gap falls by a factor of 16 — closing the agricultural TFP gap with the US only reduces the income gap by a factor of 3, leaving an income gap of 14 times.

When we analyze the model experiment, we find that agricultural productivity improvements fail to close income gaps for two reasons. First, non-homotheticities and an elasticity of substitution less than one imply that agricultural productivity improvements are self-defeating. As agriculture gets more productive, fewer and fewer workers are in agriculture which means that the poor country suffers more from its low productivity in non-agriculture. The falling relative price of agricultural products also means that the value added share of agriculture falls even more sharply than the employment share in agriculture. Second, since investment goods and intermediate inputs are produced in the non-agricultural sector, agriculture does not experience capital and intermediate input deepening when only agricultural TFP improves. Hence, even though the country closes the TFP gap with the rich countries, it does not fully close the labor productivity gap.

In contrast, improvements in non-agricultural TFP can close a substantial proportion of the income gaps between poor and rich countries. This is because these improvements are not self-defeating, but in fact self-reinforcing. While an increased non-agricultural TFP does lead to increasing prices of agricultural goods, which slows down structural transformation, non-agricultural TFP improvements also cause income effects and agricultural input intensification that induce structural transformation away from agriculture. In the quantitative analysis, we find that the second effect dominates. Thus, when we close poor countries' non-agricultural TFP gaps with the rich countries, they see a substantial shift of workforce and value added out of agriculture, and output grows substantially.

The findings in our paper are relevant to a number of debates in the macroeconomic development literature.

First, the paper highlights the problem with using simple decomposition exercises for policy analysis in the context of structural change. Just like in the data, our model features a relatively low physical labor productivity in agriculture for poor countries, and our decomposition exercise suggests large effects of improving agricultural labor productivity. However, despite this, the strength of the general equilibrium feedback means that the full model-based counterfactual has relatively modest effects of closing the agricultural TFP gap.

Second, the paper also provides some vindication for the historical focus on non-agriculture in development policy. These policies could seem problematic in light of the labor productivity decomposition results, which seem to suggest that poor countries are poor since they have a low agricultural productivity and many workers in agriculture. A policy interpretation of the findings is that policy should focus on productivity barriers within agriculture, and potentially on barriers to sectoral reallocation.

In contrast, our results shows that even with a low relative labor productivity in agriculture, a focus on non-agricultural TFP can bring great benefits. The key mechanism is that, for a conventionally calibrated structural change model, income effects and input intensification effects dominate the relative price effect. This means that improvements in non-agricultural TFP are enough to induce structural transformation out of agriculture. Since agricultural TFP improvements are not needed to free up workers from agriculture, non-agricultural TFP increases alone are sufficient to close a large proportion of the income gaps between poor and rich countries. In contrast, for agricultural TFP improvements, both non-homotheticities and relative price effects work to make the agricultural sector relatively less important, making agricultural TFP improvements a partly self-defeating growth strategy. Indeed, even if poor countries could manage to close their agricultural TFP gaps with the rich world, they would still be poor.

2 Data

In this section, we discuss the data used in our analysis, organized by factor input. For each input, we discuss the ideal measures to use, the challenges presented by those measures, the actual measures we use in practice, and the future improvements to our measures we plan to make.

2.1 Capital

2.1.1 Price

The ideal price measure of capital is the user cost of capital, defined as $p_k(r+\delta) - \dot{p}_k$, where p_k denotes the value of investment goods, r is the interest rate, and δ is the depreciation rate (see, e.g., Hall and Jorgenson (1967)). Therefore, in order to optimally measure capital, data is required on all three of these variables, as well as the expected change in the price of investment goods.

The price measure we use abstracts from this ideal measure by making several simplifying assumptions. Specifically, we assume that both the interest rate and the depreciation rate are constant across countries. Further, since our analysis focuses on steady-state comparisons, we impose that there is no expected change in the price of investment goods. Thus, to construct our price measure, we only require data on the relative price of investment goods, which we take directly from the Penn World Table (PWT).

2.1.2 Quantity

The ideal quantity measure of capital would be the dollar value of installed capital in agriculture, deflated by an agricultural capital goods price index. Unfortunately, to the best of our knowledge, no such widely available series exists. As a result, the measure we use is the constant 2005 US dollar value of net capital stocks in agriculture, forestry, and fishing (defined as ISIC Rev. 4 Section A), available from the FAOSTAT database provided by the Food and Agriculture Organization of the United Nations (FAO). The FAO constructs net capital stocks from gross fixed capital formation (GFCF) data using the perpetual inventory method. To calculate constant price series, the FAO applies GFCF deflator series, constructing an implicit deflator for agricultural GFCF from OECD data whenever possible and otherwise constructing an implicit GFCF deflator using total economy GFCF data from the National Accounts Estimates of Main Aggregates database provided by the United Nations Statistics Division (UNSD).

Although the FAO's estimation methodology is thoughtful, substantial data availability issues greatly limit its reliability. While officially reported statistics on agricultural GFCF are used whenever available, and subjected to extensive quality and consistency checks, for many countries and many years no data is available and the FAO instead resorts to imputation. Of the 201 countries and territories for which the FAO provides 2005 net capital stock estimates, 169 of these (or 84% of all observations) are FAO imputations. Of the 32 countries with non-imputed estimates, only one (Namibia) is in Africa. Due to this extensive imputation, any use of FAO capital stock data should be done with caution.

We adjust the FAO series such that the 2005 US capital intensity in agriculture relative to the total economy is the same as in data from the US Bureau of Economic Analysis (BEA). Formally, we take the 2005 US net capital stock from the FAO divided by our estimate of agricultural employment (detailed below) and divide this by the 2005 US aggregate capital stock divided by total employment, both from the PWT. We then calculate the same ratio using data from the BEA on nominal capital stock and hours worked by persons engaged in production in the farm sector and total economy. We then divide the BEA ratio by the FAO/PWT ratio, yielding an inflation factor of approximately 1.437 which we then apply to the FAO net capital stock series.

Going forward, we plan to improve our capital stock measure using microdata on capital used in agricultural production available from the World Bank's Living Standards Measurement Surveys (LSMS). These surveys are conducted by national statistical agencies in conjunction with the World Bank in eight African countries and collect detailed, household-level information on farm implements and machinery for agricultural production from nationally representative samples of agriculturally active households. We will use this data to estimate agricultural capital values which we will then in turn use to deflate the FAO capital stock series.

2.2 Land

2.2.1 Price

The ideal price measure of land would be the user cost of quality-adjusted agricultural land. Constructing such a measure would require a consistent quality measure of land comparable across countries and either: 1, a price series of land and the user cost of capital; or 2, a rental series of land. Unfortunately, to the best of our knowledge, no standard dataset of land prices across the world exists. For example, while the FAO provides an extensive dataset on land use across the world, they collect no information on land prices.

We base our price measure of land on rental series, specifically the average US dollar per hectare rental rate of agricultural land. These series are constructed for each country for which sufficient data is available. Currently, these countries are: Ethiopia, Malawi, Niger, Nigeria, Tanzania, Uganda, Austria, Bulgaria, Croatia, the Czech Republic, Denmark, Estonia, Finland, France, Greece, Hungary, Ireland, Latvia, Lithuania, Luxembourg, the Netherlands, Slovakia, Slovenia, Spain, Sweden, and the United States. The average rental rates for the African countries are estimated using microdata from the LSMS, which collect household-level data on agricultural land use including tract size, rental status, rent paid, and rental period. When constructing estimates from microdata, we compute the average per hectare rental rate over all tracts of land that households reported paying a positive amount of rent on. The average rental rates for the European countries are taken from Eurostat's Agriculture database, and are defined over all agricultural land. The average rental rate for the US is taken from the US Department of Agriculture, and is defined over cropland. We convert all rental rates into US dollars using exchange rates taken from the PWT. Finally, we use the cross-country relationship between average agricultural land rental rates and log GDP per worker observed over these countries to construct rental rate estimates for all other countries.

Going forward, we plan to use microdata to estimate rental rates for the two additional countries covered by the LSMS (Burkina Faso and Mali) as well as the two countries covered by ICRISAT's Village Dynamics in South Asia program (Bangladesh and India). We further plan to adjust our land rental rates for differences in land quality, incorporating climatic and soil quality information from the FAO's Global Agro-Ecological Zones (GAEZ) program.

2.2.2 Quantity

The ideal quantity measure of land would be quality-adjusted land used in agricultural production. An analysis of land quality across countries by Adamopoulos and Restuccia (2014) finds no clear evidence of systematic variation in land quality across the global income distribution.[3] Their results suggest that quality-adjusting raw land use measures will not meaningfully affect any observed cross-country relationship between land use in agriculture and log GDP per worker.

In light of these results, at this stage we use unadjusted measures of land use in agriculture. In particular, we use two land use measures provided by the FAO, both reported in hectares: cropland and agricultural land. The FAO defines cropland as "land used for cultivation of crops" and defines agricultural land as cropland plus "land used for... animal husbandry." Our baseline analysis uses the cropland measure; the

[3] Adamopoulos and Restuccia (2014) use land quality measures from the FAO's GAEZ program to compare the amount of high quality land per capita between an average rich country and an average poor country. When looking at the best quality land, they find a rich-poor ratio of 1.1, and under various alternative weighting schemes and definitions of high quality land they find ratios ranging from 0.5 to 1.6.

agricultural land measure is used to test for robustness. Both series are largely based on officially reported data and are subject to extensive quality assurance checks by the FAO.[4]

In the future, we plan to adjust these land use measures for land quality using measures derived from the FAO's GAEZ program. Although there doesn't appear to be evidence of a strong, systematic cross-country relationship between land quality and income, there do still appear to be important differences in land quality unrelated to income level. Accounting for these differences will likely yield a more accurate measure of land input in agricultural production.

2.3 Intermediate inputs

2.3.1 Price

The ideal price measure of intermediate inputs would be a price index of standard agricultural intermediate inputs. Unfortunately, to the best of our knowledge, no such price information is available for a broad set of countries. While the FAO does collect and report extensive quantity data on fertilizer use in agricultural production, they provide no information on the corresponding prices of the fertilizers.

Given this limited data availability, we assume that intermediate input prices are uncorrelated with GDP per worker. This assumption has some empirical support. Using data from Prasada Rao (1993) constructed from FAO statistics, Donovan (2016) reports evidence that the price of agricultural intermediate inputs relative to agricultural output does not vary systematically over the global income distribution, despite the substantial dispersion in these relative prices. Further, our own analysis of agricultural output prices based on FAO data finds little to no correlation between agricultural output prices and GDP per worker. Together, these two findings suggest that agricultural intermediate input prices must also be uncorrelated with GDP per worker.

We therefore use an agricultural output PPP as our price measure of agricultural intermediate inputs, which we construct using agricultural output data from the FAO. Specifically, we calculate our agricultural output PPP as the ratio of agricultural output measured in current standard local currency to agricultural output measured in constant 2005 international dollars. We further normalize our agricultural output PPP series to equal one in the US.

2.3.2 Quantity

The ideal quantity measure of intermediate inputs would be agricultural intermediate inputs from the non-agricultural sector, deflated by a price index of standard agricultural intermediate inputs. The quantity measure we use is intermediate consumption in the agricultural sector (defined as ISIC Rev. 3 Section A, Division 01) taken from the UNSD's National Accounts Official Country Data, Table 2.3. This series is close to the ideal measure, but differs in two key ways. First, the series is in nominal local currency units valued at purchaser's prices. As explained above, we lack a price index of agricultural intermediate inputs and instead use our agricultural output PPP series to deflate the UN series. Second, the UN series includes intermediate inputs from the agricultural sector as well as the non-agricultural sector. We adjust the series to account for this by multiplying it by the US share of agricultural intermediate inputs from the non-agricultural sector

[4] Over the 1991 to 2016 period, approximately 15.2% of country-year observations contained an FAO estimate for at least one of the two measures we use. We find our results are unaffected by the exclusion of such estimated values.

in 2005, which we find to be roughly 74.5%.[5]

In the future, we plan to replace our admittedly course adjustment factor with one constructed directly from country-level input-output tables available from the Global Trade Analysis Project (GTAP). The GTAP provides consistently defined, highly detailed input-output tables for a large cross-section of countries that would be ideal for estimating the share of agricultural intermediate inputs from the non-agricultural sector.

2.4 Labor

2.4.1 Price

Our analysis does not require a direct price measurement of labor, only the relative factor price of capital to labor. We estimate this relative factor price using the capital-labor ratio in the non-agricultural sector under the standard assumption that the labor share of income in this sector is 2/3.

2.4.2 Quantity

The ideal quantity measure of labor would be human-capital-adjusted actual annual hours worked in agriculture. However, in our current version, we only use agricultural employment (in millions of workers) as our labor quantity measure. Omitting hours is not likely to change our results considerably. Bick et al. (2018) report that there is little cross-country variation in hours worked per worker in the agricultural sector.[6] For now, we assume all workers have uniform human capital.

To estimate agricultural employment, we use data on total employment from the PWT and data on the agricultural share of employment from the International Labor Organization (ILO) [World Bank series ID: SL.AGR.EMPL.ZS]. We simply multiply both series together to yield our estimate of agricultural employment in each country. We compare all three of these series with equivalent series derived from microdata taken from the University of Minnesota's Integrated Public Use Microdata Series database (IPUMS) for the 66 countries for which the necessary data is available. We find the aggregate and microestimates are generally very similar for all three series, with few exceptions.[7]

3 Data findings

Figure 1 plots the six principal factor input ratios against log GDP per worker for the 63 countries in our dataset with complete data, as well as a line of best fit. The dotted line gives the relative prices between the two factors as a function of log GDP per worker. The figure makes evident that all non-labor factors of production are strongly increasing with development in per worker terms. Capital and intermediate inputs per worker rise in tandem by orders of magnitude as one moves up the global income distribution; land per

[5] This is calculated using data from the BEA's Input-Output tables on the use of commodities by sector. The data we use is at the summary level, and we define the agricultural sector as IO code 111CA, "Farms." This definition closely corresponds to NAICS subsector 111, "Crop production," and to ISIC Rev. 3 Section A, Division 01.

[6] Bick et al. (2018) compare the average weekly hours worked per worker between low, middle, and high income country groups defined over their set of "core" countries, 49 countries whose underlying hours data is explicitly actual hours worked in the production of output counted in NIPA across all jobs in a recent reference week and is collected over the entire calendar year. They find averages of 36, 38.3, and 39.7 hours worked per worker in agriculture in low, middle, and high income countries, respectively. None of these differences are found to be statistically significant.

[7] For the most part, deviations between the estimates are the result of definitional differences between the aggregate sources and the national survey the IPUMS data is sourced from.

Table 1: Slopes with respect to log GDP per worker

		$\hat{\beta}_y$	S.E.
Relative factor intensities:			
	$\log(K/L)$	1.87	0.0901
	$\log(X/L)$	1.82	0.122
	$\log(R/L)$	0.864	0.120
	$\log(K/R)$	1.01	0.119
	$\log(X/R)$	0.952	0.0961
	$\log(K/X)$	0.0578	0.112
Relative factor prices:			
	P_k/w	-0.815	0.0258
	P_r/w	-0.577	0.114

worker rises more modestly. These divergent trends between capital and intermediate inputs on the one hand and land on the other are further reflected in the second column of figures, which plot the usage of these inputs relative to each other. Unsurprisingly, the capital- and intermediate-to-land ratios both rise by nearly the same factor when moving from the lowest to highest income setting, while the capital-to-intermediate ratio is virtually flat.

The relative price series show these factor intensity patterns are consistent with a simple neoclassical story. All factor prices fall relative to labor as we go from poor to rich countries. We also find that the price of land relative to capital and intermediate inputs falls as countries get richer, which is consistent with the increased intensity of these two inputs relative to land. Furthermore, we see that capital and intermediate inputs increase at roughly the same rate across the income distribution, consistent with the lack of a trend in their prices relative to one another.

Another observation is that capital deepening and intermediate input intensification in agriculture happens along the income distribution at a faster rate than the changes of the relative prices. Without distortions — from a neoclassical point of view — this suggests that factor shares are non-constant and the elasticities of substitution between labor and capital and between labor and intermediate inputs are larger than one. As a consequence, when it comes to agriculture, assuming a Cobb-Douglas production function is questionable.

4 Model

We analyze our findings through the lens of a simple structural change model. We assume that each country is a closed economy with an agricultural and a non-agricultural sector. The non-agricultural sector, in turn, produces consumption goods, capital goods, and intermediate inputs used in agriculture. Countries differ in their sectoral TFPs, A_a and A_{na}, and our main counterfactual of interest is the steady state response of sectoral structure, labor productivity, and aggregate output to changes in A_a and A_{na}. We focus on the steady state solution for constant levels of TFP. In this, we follow standard practice in the development accounting literature by evaluating the effects of TFP changes after capital has been allowed to converge to its steady-state level.

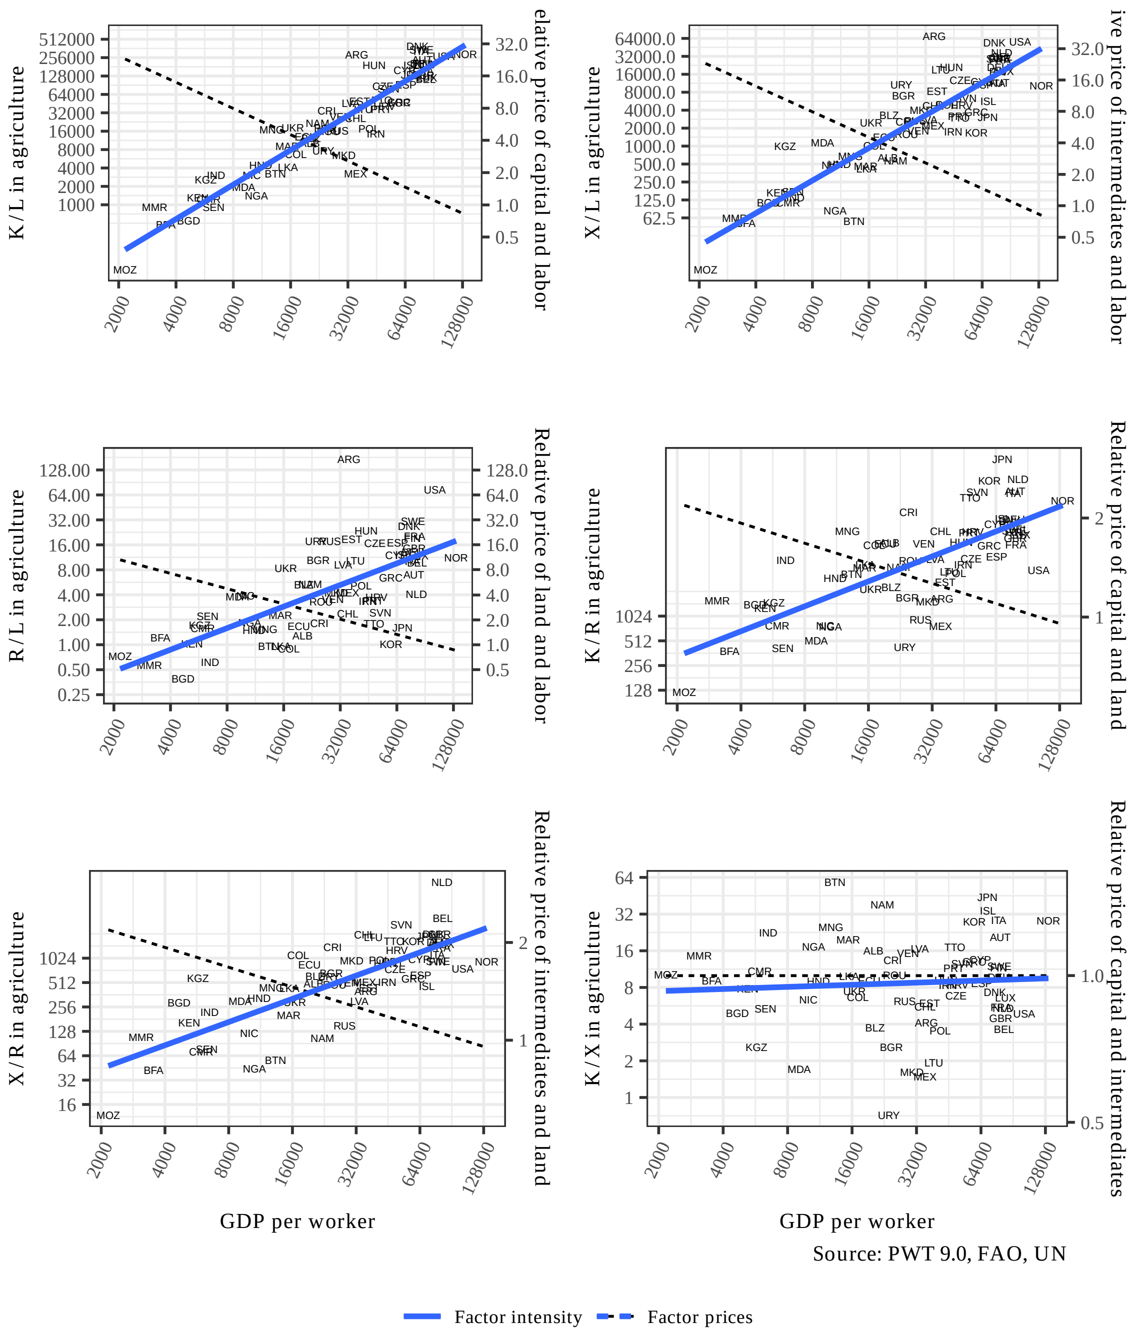

Figure 1: Factor ratios and factor prices

4.1 Households

There is a representative household that inelastically supplies one labor unit, rents out capital K_t and land R, and maximizes the preferences

$$\mathcal{U}_0 = \sum_{t=0}^{\infty} \beta^t v(e_t, p_{a,t})$$

subject to the flow budget constraint

$$e_t + P_{k,t}(K_{t+1} - (1-\delta)K_t) \leq r_t K_t + w_t + \rho_t \bar{R}.$$

$v(\cdot)$ is an instantaneous utility function given in indirect form, i.e., specified over total nominal expenditure e_t and the prices of the two consumption goods (agricultural and non-agricultural output). In the following we normalize the price of the non-agricultural output to one such that e_t should be understood as aggregate consumption expenditure *in terms of the non-agricultural good* and only the relative price of the agricultural good, $p_{a,t}$, shows up as an additional argument in the $v(\cdot)$ function. Physical capital is denoted by K_t, δ is the depreciation rate, and $P_{k,t}$ is the price of capital goods in terms of the non-agricultural good. r_t is the rental rate on capital, denoted in non-agricultural goods per unit of capital. w_t is the wage rate, ρ_t is the rental rate of land, and $\bar{R}$ is the aggregate land stock. Finally $\beta < 1$ denotes the discount factor.

The period utility function $v(\cdot)$ is of the PIGL form, as in Boppart (2014):

$$v(e_t, p_{a,t}) = \frac{1}{\epsilon}(e_t)^{\epsilon} - \frac{\nu}{\gamma}(p_{a,t})^{\gamma}.$$

The share parameter ν regulates the level of demand for the agricultural good, while the parameter γ regulates the (non-constant) elasticity of substitution between agricultural and non-agricultural consumption goods. With $\gamma > 0$ the two consumption goods are gross complements. The parameter ϵ regulates non-homotheticity of the preferences, with $1-\epsilon$ being the (constant) expenditure elasticity of demand for agricultural output. Note that if $\epsilon = 0$, the expenditure elasticity is 1 and the preferences are homothetic.

Given these preferences, the solution to the household's intratemporal problem is

$$c_{na,t} = e_t - \nu e_t^{1-\epsilon} p_{a,t}^{\gamma} \tag{1}$$

$$c_{a,t} = \nu e_t^{1-\epsilon} p_{a,t}^{\gamma-1}. \tag{2}$$

The expenditure share on agricultural goods is

$$s_{a,t} \equiv \frac{p_{a,t} c_{a,t}}{e_t} = \nu e_t^{-\epsilon} p_{a,t}^{\gamma}.$$

If preferences are non-homothetic (i.e., $\epsilon > 0$), the expenditure share is decreasing in the aggregate expenditure level e_t. Furthermore, if consumption goods are complements (i.e., $\gamma > 0$), a falling price of agricultural goods implies a falling expenditure share on agricultural goods. Thus, the preferences allow for structural change driven both from income effects and relative price effects, and hence a high agricultural productivity may push people out of agriculture for two reasons.

The household's intertemporal problem implies an Euler equation:

$$P_{k,t} e_t^{\epsilon-1} = \beta \left(P_{k,t+1}(1-\delta) + r_{t+1} \right) e_{t+1}^{\epsilon-1}. \tag{3}$$

Note that the Euler equation does not contain any term for the relative price of agricultural products, considerably simplifying the steady-state analysis.

4.2 Production side

Non-agricultural sector

We assume that the non-agricultural sector produces competitively according to the Cobb-Douglas production function,

$$Y_{na} = \phi_{na} A_{na} K_{na}^{\alpha} L_{na}^{1-\alpha},$$

where ϕ_{na} is a normalizing constant. In the following we normalize the price of the non-agricultural consumption good to 1, i.e.,

$$\phi_{na} A_{na} = \left(\frac{r_t}{\alpha} \right)^{\alpha} \left(\frac{w_t}{1-\alpha} \right)^{1-\alpha}. \tag{4}$$

As we will describe below the non-agricultural output can either be consumed or transformed in some fixed proportion into an investment good and an intermediate input for the agricultural sector. Since these transformations, however, are not one-for-one the relative prices P_k and P_x will generally differ from one.

Given this formulation, firms in the non-agricultural sector solve each period

$$\max_{K_{na}, L_{na}} \ \phi_{na} A_{na} K_{na}^{\alpha} L_{na}^{1-\alpha} - r K_{na} - w L_{na}. \tag{5}$$

The non-agricultural goods are used for consumption, the production of investment goods, and the production of intermediate goods for the agricultural sector.[8] The non-agricultural goods are transformed to consumption one-for-one. In addition, non-agricultural goods can be transformed one-for-A_K into investment goods I_t or one-for-A_X into agricultural intermediate inputs X_t. Hence we have:

$$I_t = A_K Y_I,$$
$$X_t = A_X Y_X.$$

We further assume that the productivities A_K and A_X take the following forms:

$$A_K = \phi_K A_{na}^{\kappa},$$
$$A_X = \phi_X A_{na}^{\kappa}.$$

By introducing these additional productivity terms, we permit the price of investment goods relative to consumption goods to fall as countries get richer, in line with findings in Hsieh and Klenow (2007). However, the overall TFP level of the non-agriculutral sector is still captured by the single parameter A_{na}; an increase

[8] The sector does not produce its own intermediate inputs, since we interpret the production function to be net of own-produced inputs, i.e., in terms of value-added.

in this TFP level shifts the overall productivity in the non-agriculture sector (whereas with $\kappa \neq 0$ this shift is differential across the uses of the non-agricultural output, i.e., non-agricultural consumption, investment, and intermediate inputs).

For notational convenience later, we choose the following values for the normalizing constants:

$$\phi_{na} = \left(\frac{1-\alpha}{\alpha}\right)^{\alpha}, \tag{6}$$

$$\phi_K = \frac{1/\beta - 1 + \delta}{1-\alpha}, \tag{7}$$

$$\phi_X = \frac{1}{1-\alpha}. \tag{8}$$

Under perfect competition, the prices of investment and intermediate inputs in our setting are given by:

$$P_{k,t} = A_K^{-1} = A_{na}^{-\kappa}\left(\frac{1-\alpha}{1/\beta - 1 + \delta}\right) \tag{9}$$

$$P_{x,t} = A_X^{-1} = A_{na}^{-\kappa}(1-\alpha). \tag{10}$$

Hence, these relative prices will be constant over time but differ across countries with different levels of A_{na}.

Agricultural sector

The agricultural sector has a production function

$$Y_{a,t} = A_a F(L_{a,t}, K_{a,t}, X_{a,t}, R_{a,t}),$$

with constant returns to scale in all arguments. We again assume perfect competition. This then implies in each point in time the maximization problem

$$\max_{L_a, K_a, X_a, R_a} p_a A_a F_a(L_a, K_a, X_a, R_a) - wL_a - rK_a - P_x X_a - \rho R_a, \tag{11}$$

where the factor prices are taken as given.

In our baseline specification, we assume that $F(\cdot)$ takes a nested CES form

$$F(L_{a,t}, K_{a,t}, X_{a,t}, R_{a,t}) = \left[\left(L_{a,t}^{\frac{\eta-1}{\eta}} + a_k^{1/\eta} K_{a,t}^{\frac{\eta-1}{\eta}} + a_x^{1/\eta} X_{a,t}^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}\frac{\sigma-1}{\sigma}} + a_r^{1/\sigma} R_{a,t}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{12}$$

A nested CES departs from the standard assumptions of Cobb-Douglas and CES, and is motivated by our desire to fit the cross-sectional patterns of factor input demands. In Section 3, we found that all factors of production had non-unitary elasticities with respect to prices, suggesting that a Cobb-Douglas production function would be a poor fit. Moreover, the demands of capital and intermediate inputs relative to labor are more sensitive to relative prices than the relative demand for land, suggesting that a CES production function over all arguments would also not fit the data. However, since the price sensitivity of capital and intermediate inputs were similar, a nested CES with land, labor, and capital together is more promising. However, we plan to experiment more with different nesting structures in future, for example nesting intermediate inputs and land instead.

4.2.1 Equilibrium

An equilibrium is an allocation and a set of prices such that consumers and producers optimize, and the following resource constraints and market clearing conditions are satisfied

$$L_{a,t} + L_{na,t} \leq 1 \tag{13}$$

$$K_{a,t} + K_{na,t} \leq K_t \tag{14}$$

$$R_{a,t} \leq \bar{R} \tag{15}$$

$$C_{na,t} + \frac{K_{t+1} - (1-\delta)K_t}{A_K} + \frac{X_{a,t}}{A_X} \leq Y_{na.t} \tag{16}$$

$$C_{a,t} \leq Y_{a,t} \tag{17}$$

4.3 Steady state

We look for a steady state solution where the level of the capital stock K_t converges to its long-run level given the TFP levels A_a and A_{na}. We denote steady state variables by a $\star$ superscript. The Euler equation (3) implies that

$$P_k = \beta\left((1-\delta)P_k + r\right) \Longleftrightarrow r^\star = P_k\left(\frac{1}{\beta} - 1 + \delta\right),$$

which, using (9), resolves to

$$r^\star = \frac{1-\alpha}{A_{na}^\kappa}. \tag{18}$$

Similarly, the price of agricultural intermediate inputs is

$$P_x = A_X^{-1} = \frac{1-\alpha}{A_{na}^\kappa}. \tag{19}$$

Note that if $\kappa > 0$, then a higher non-agricultural TFP A_{na} reduces the steady-state rental rate of capital and the price of intermediate inputs.

Using the price normalization equation (4), the steady-state wage level is

$$w^\star = (1-\alpha)(\phi_{na}A_{na})^{\frac{1}{1-\alpha}}(r/\alpha)^{-\alpha/(1-\alpha)} = (1-\alpha)A_{na}^{\frac{1+\alpha\kappa}{1-\alpha}}, \tag{20}$$

where we use (6) for ϕ_{na} and (18) for r. In this expression, $A_{na}^{\frac{1}{1-\alpha}}$ follows from the standard neoclassical growth model logic, with TFP increases changing wages more than one-for-one due to their indirect effect on capital accumulation. The additional $\alpha\kappa$ term is due to a higher A_{na} reducing the relative price of capital goods, thereby inducing more capital accumulation.

The previous results mean that relative factor prices satisfy

$$\frac{r\star}{w^\star} = \frac{P_x}{w^\star} = A_{na}^{-\frac{1+\kappa}{1-\alpha}},$$

and that total non-agricultural output is

$$Y_{na}^{\star} = A_{na}^{\frac{1+\alpha\kappa}{1-\alpha}} L_{na}.$$

Since $rK_{na} = \alpha Y_{na}$, total capital investment in the non-agriculture is given by

$$\delta P_k K_{na}^{\star} = \frac{\delta}{1/\beta - 1 + \delta} \alpha Y_{na}^{\star},$$

where we use (9) and note that the cost of capital investment is a share $\frac{\delta}{1/\beta-1+\delta}$ of the total cost of capital to firms, $\alpha Y_{na}^{\star}$.

We then solve for the allocation of labor across sectors that clears the consumption market given a rental rate of land ρ. Writing $k_a(\rho), x_a(\rho)$, and $r_a(\rho)$ for the factor demands in the agriculture sector relative to labor, we note that the consumption goods market clears when

$$\frac{C_{na}(\rho)}{C_a(\rho)} = \frac{e(\rho) - \nu e(\rho)^{1-\epsilon} p_a(\rho)^{\gamma}}{\nu e(\rho)^{1-\epsilon} p_a(\rho)^{\gamma-1}}, \tag{21}$$

where the aggregate consumption expenditure is

$$e(\rho) = C_{na}(\rho) + p_a(\rho) C_a(\rho). \tag{22}$$

Consumption of agricultural output is

$$C_a(\rho) = Y_a(\rho) = A_a L_a(\rho) F_a(1, k_a(\rho), x_a(\rho), r_a(\rho)), \tag{23}$$

whereas consumption in the non-agricultural sector in steady state is given by

$$C_{na}(\rho)^{\star} = [1 - L_a(\rho)^{\star}][y_{na}^{\star} - \delta P_k k_{na}^{\star}] - L_a(\rho)^{\star} \left[\delta P_k k_a(\rho)^{\star} + P_x x_a(\rho)^{\star}\right], \tag{24}$$

where $y_{na}^{\star}$ and $k_{na}^{\star}$ are the per worker output and capital stock in the non-agricultural sector in steady state.

The relative price of agricultural goods is

$$p_a(\rho) = \left[\frac{1}{A_a} \times \left(P_Q^{1-\sigma} + a_r P_r^{1-\sigma}\right)^{\frac{1}{1-\sigma}}\right], \tag{25}$$

where

$$P_Q = \left(w^{1-\eta} + a_k P_k^{1-\eta} + a_x P_x^{1-\eta}\right)^{\frac{1}{1-\eta}}.$$

Equations (21) to (25) imply that there is a unique solution for the employment share in agriculture $L_a(\rho)$ as a function of the rental rate of land ρ.

Given this function, we close the model by finding the land rental rate $\rho^{\star}$ that clears the land market

$$L_a(\rho^{\star}) r_a(\rho^{\star}) = \bar{R}.$$

5 Calibration

When calibrating the model, we construct a continuum of economies $\mathcal{Y}$ that correspond to different levels of GDP per worker. We assume that

$$\mathcal{Y} = \left[\frac{y_{US}}{40}, y_{US}\right],$$

where y_{US} is the GDP per worker in the US in 2005. Our model economies differ in their sectoral TFP levels, $A_{na}(y)$ and $A_a(y)$, with the normalization

$$A_{na}(y_{US}) = A_a(y_{US}) = 1.$$

5.1 Notation: variables as functions of GDP per worker

To fit the model, we will fit the observed cross-sectional patterns between poor and rich countries. Thus, our primary moments will be *functions* of GDP per worker. For any observable variable z, we write $z(y)$ for the value that z takes at GDP per worker level y. Similarly, for any model object Z, we write $Z(y)$ for the value taken by Z when $A = (A_{na}(y), A_a(y))$. While $Z(y)$ can always be calculated in the model, $z(y)$ will typically be a fitted value based on observations $\{(y_i, z_i)\}_{i=1}^n$, where i indexes countries.

5.2 Preference parameters

We take a standard value for the time preference, $\beta = 0.96$. The parameters in the PIGL preferences are based on the calibration done in Eckert et al. (2018). Using the CEX from 1935, they find an elasticity of agricultural consumption with respect to aggregate expenditure of 0.32. This corresponds to $\epsilon = 1 - 0.32$ in the PIGL preferences. To find γ, Eckert and Peters set $\gamma = 0.35$. Using the same values of ϵ and γ, we calibrate the share parameter ν to match the average agricultural labor share for countries at US income levels.[9]

In our counterfactual experiment, we also experiment with multiple other values of γ as a robustness check and to illustrate the effects of using different elasticities of substitution. We also recalibrate the model with $\epsilon = 0$ to evaluate the effect of non-homotheticities on structural change.

5.3 Non-agricultural production

For the non-agricultural sector, we use standard parameter values for depreciation, $\delta = 0.08$, and the capital share, $\alpha = 1/3$. The parameter κ, which regulates the relative price of capital goods and non-agricultural consumption goods, is calibrated to fit the cross-sectional relationship between GDP per worker and the relative value of the user cost of capital and wages. The calibration is conducted jointly with calibrating $A_{na}(y)$, and the exact method is described below.

[9] Formally, we regress $\log(L/L_a)$ on a second-degree polynomial of log GDP per worker, and we match the fitted value at $\log y_{US}$.

5.4 Agricultural production function

In the agricultural production function, we need to calibrate the share parameters a_k, a_x, and a_r, as well as the two elasticities, η and σ. The first-order conditions are

$$
\begin{aligned}
\frac{K_a}{L_a} &= a_k \left(\frac{P_k}{w}\right)^{-\eta} \\
\frac{X_a}{L_a} &= a_x \left(\frac{P_x}{w}\right)^{-\eta} \\
\frac{R_a}{Q} &= a_r \left(\frac{P_r}{P_Q}\right)^{-\sigma} \\
Q &= \left(L^{\frac{\eta-1}{\eta}} + a_k^{1/\eta} K^{\frac{\eta-1}{\eta}} + a_x^{1/\eta} X^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \\
P_Q &= \left(L^{1-\eta} + a_k^{1/\eta} K^{1-\eta} + a_x^{1/\eta} X^{1-\eta}\right)^{\frac{1}{1-\eta}}
\end{aligned}
$$

To calibrate the share parameters, we normalize such that relative prices compared to labor are 1 at US income levels. Since all relative prices are then 1 in the US, relative factor shares satisfy

$$
\begin{aligned}
\frac{\alpha_k(y_{US})}{\alpha_l(y_{US})} &= a_k \\
\frac{\alpha_x(y_{US})}{\alpha_l(y_{US})} &= a_x
\end{aligned}
$$

and we calibrate a_k and a_x directly using data on US factor shares.

To calibrate the elasticity η (between L, K, and X), we regress the relative prices and quantities of capital and labor on GDP per worker:

$$
\begin{aligned}
\log\left(\frac{K_{a,i}}{L_{a,i}}\right) &= \alpha_q + \beta_q \log(y_i) \\
\log\left(\frac{P_{k,i}}{w_i}\right) &= \alpha_p + \beta_p \log(y_i).
\end{aligned}
$$

We calibrate $\eta = -\hat{\beta}_q/\hat{\beta}_p$ as the ratio of the two coefficient estimates. This method ensures that the production function correctly describes the relationship between relative prices and relative quantities of capital and labor across the world income distribution. If the regression is conducted with intermediate inputs instead of with capital, the results are very similar.

To calibrate σ, we first calculate Q and P_Q as functions of GDP per worker, using the estimated share parameters, the estimated η, and the data relationships between GDP per worker and the relative prices and quantities of capital and intermediate inputs. Having done this, we then run two regressions,

$$
\begin{aligned}
\log\left(\frac{R_{a,i}}{Q_i}\right) &= \tilde{\alpha}_q + \tilde{\beta}_q \log(y_i), \\
\log\left(\frac{P_{r,i}}{P_{Q,i}}\right) &= \tilde{\alpha}_p + \tilde{\beta}_p \log(y_i),
\end{aligned}
$$

and, similar to before, we define $\sigma \equiv -\hat{\hat{\beta}}_q/\hat{\hat{\beta}}_p$. Having defined σ, we can then define a_r from

$$\frac{\alpha_r}{\alpha_l} = \frac{P_r(y_{US})R_a(y_{US})}{w(y_{US})L_a(y_{US})} = a_r \left(\frac{P_r(y_{US})}{P_Q(y_{US})}\right)^{-\sigma} \frac{Q(y_{US})}{L_a(y_{US})},$$

using that $P_r(y_{US})/w(y_{US}) = 1$ by normalization and that

$$\frac{R_a(y_{US})}{L_a(y_{US})} = \frac{R_a(y_{US})}{Q(y_{US})}\frac{Q(y_{US})}{L_a(y_{US})} = a_r \left(\frac{P_r(y_{US})}{P_Q(y_{US})}\right)^{-\sigma} \frac{Q(y_{US})}{L_a(y_{US})}$$

5.5 Agricultural TFP

To calibrate the agricultural TFP $A_a(y)$, we use the standard result that, for an arbitrary production function

$$Y = AF(x_1, \ldots, x_n),$$

we have that

$$d\log A = d\log Y - \sum_{i=1}^{n} \alpha_i d\log x_i,$$

where

$$\alpha_i = \frac{\partial F}{\partial x_i}\frac{x_i}{F}$$

is the elasticity of output with respect to input x_i, which is equal to the factor compensation share given competitive input and output markets. If F is a constant returns to scale production function and we normalize with the first factor input x_1, the corresponding expression is

$$d\log A = d\log\left(\frac{Y}{x_1}\right) - \sum_{i=2}^{n} \alpha_i d\log\left(\frac{x_i}{x_1}\right), \tag{26}$$

where we use that $\alpha_1 = 1 - \sum_{i=2}^{n} \alpha_i$.

Using (26), we can write the agricultural TFP at income level y as

$$\begin{aligned}\log A_a(y) &\doteq \log\left(\frac{A_a(y)}{A_a(y_{US})}\right) \\ &= \log\left(\frac{y_a(y)}{y_a(y_{US})}\right) - \int_{y_{US}}^{y} \sum_{z=k,x,l} \alpha_z^a(y) \frac{d\log[z_a/l_a](y)}{d\log y} d\log y\end{aligned}$$

where $y_a(y) = Y_a(y)/L_a(y)$ is the labor productivity in agriculture at income level y, $\alpha_z^a(y)$ are the factor compensation shares, and $\frac{d\log[z_a/l_a](y)}{d\log y}$ is the elasticity of the factor intensity of factor z with respect to y.

To calculate $\alpha_z^a(y)$ and $\frac{d\log[z_a/l_a](y)}{d\log y}$, we use the observed relationships between $\log y$ and factor prices to calculate the factor shares and intensities implied by our previously calibrated agricultural production function. Note that, given that the production function was calibrated using the cross-sectional factor price and quantity data, we would obtain very similar results for TFP if we had the data moments for the input intensities, instead of using the price moments in the production function.

5.6 Non-agricultural TFP

To calibrate $A_{na}(y)$, we target the model-generated GDP per worker $\hat{y}(A_a(y), A_{na}(y))$ to coincide with the actual GDP per worker.

To calculate the model-based $\hat{y}$, we use chain-weighting with value added share. Here we explicitly treat the production of capital goods and intermediate goods as separate sectors, and, thus, we have four different sectors. The differential of log chain-weighted GDP per worker is

$$d\log\hat{y} = va_{na} d\log C_{na} + va_k d\log(I_{na} + I_a) + va_x d\log X_a + va_a d\log\left(\frac{P_a Y_a - P_x X_a}{P_a^{VA}}\right),$$

where va_z is the value-added share of sector z, and $I_{na} + I_a$ is the total investment in non-agricultural and agricultural physical capital respectively. The last term gives the real value added in the agricultural sector, where P_a^{VA} is a value-added price index.

In the model, we can calculate the value added shares va_z and the real amounts of non-agricultural consumption, intermediate input production, and investment. To calculate the real value added in agriculture, we define the real value added between two neighboring GDP per worker levels y_{i+1} and y_i as

$$\log\left(\frac{VA_a^{real}(y_{i+1})}{VA_a^{real}(y_i)}\right) = \log\left(\frac{VA_a^{nom}(y_{i+1})}{VA_a^{nom}(y_i)}\right) - \log\left(\frac{P_a^{VA}(y_{i+1})}{P_a^{VA}(y_i)}\right)$$

where the value added deflator is defined by a Fisher Index:

$$\log\left(\frac{P_a^{VA}(y_{i+1})}{P_a^{VA}(y_i)}\right) = 0.5\Delta\log P_{Paasche}^{VA} + 0.5\Delta\log P_{Laspeyres}^{VA}.$$

Recall a Fisher Index is a geometric average of a Paasche and a Laspeyres index:

$$\begin{aligned}
\Delta\log P_{Paasche}^{VA} &= \log\left(\frac{P_a(y_{i+1})Y_a(y_{i+1}) - P_x(y_{i+1})X_a(y_{i+1})}{P_a(y_i)Y_a(y_{i+1}) - P_x(y_i)X_a(y_{i+1})}\right)\\
\Delta\log P_{Laspeyres}^{VA} &= \log\left(\frac{P_a(y_{i+1})Y_a(y_i) - P_x(y_{i+1})X_a(y_i)}{P_a(y_i)Y_a(y_i) - P_x(y_i)X_a(y_i)}\right).
\end{aligned}$$

We find $A_{na}(y)$ by assuming that there is a constant elasticity relationship between $\log y$ and $\log A_{na}(y)$,

$$\log A_{na}(y) = \xi(\log y - \log y_{US}),$$

and we iterate on ξ until the a regression of the model-generated $\log\hat{y}(y;\xi)$ on actual $\log y$ gives a coefficient of 1.

6 Results

Table 2 shows the calibration results. To validate the model, we test how well it explains the the structural transformation out of agriculture. The first test is how well it fits the value added share of agriculture, which is displayed in Figure 2. The solid line shows the fitted value of the value added share in agriculture in the data. The dotted line gives the fit of our model. We see that the model does a very good job in capturing structural change. Even though the model is only calibrated to fit the employment share in rich countries, it correctly predicts that the value added share of approximately 40% in the poorest economies.

Since the calibration uses observed factor intensities in agriculture, including the land-labor share, one concern is that a good fit is effectively hard-coded by the calibration of the production function. To test this, we see how the model works without non-homotheticities. Formally, we set $\epsilon = 0$ and recalibrate the model. The result is displayed as the lower dashed line in the same figure. We see that, without non-homotheticities, the model does not succeed in capturing structural transformation: it predicts that poor countries will have less than 10% of their value added in agriculture. Since ϵ was based directly on US microdata, it was unconnected to the calibration of the the agricultural production function. Hence, we see that our production function only predicts the correct structural transformation when combined with micro-founded demand parameters.

In Figure 3, we plot the structural change from the employment perspective. We see again that our model approximately captures the observed pattern of structural change, and that it fails to capture this pattern when we switch off the non-homotheticities. However, we also see our model somewhat undershoots the agricultural employment share in poor countries. One reason our model fails to fit the employment share as well as it fits the value added share is that we do not allow for any migration frictions or differences in human capital across the sectors, which results in the model underpredicting the employment share for any given value added share.

To further test the robustness of the structural transformation result, we test the effect of all combinations of the following three parameters,

$$
\begin{aligned}
\alpha &= 0.2, \mathbf{0.33}, 0.45 \\
\gamma &= -0.5, 0, \mathbf{0.35}, 0.9 \\
\epsilon &= 0, \mathbf{0.62}, 1
\end{aligned}
$$

and recalibrate the model given each combination. The bold values are our baseline parametrization. This test covers a wide range of combinations of all the relevant parameters for structural change not in the agricultural production function. The results are shown in Figure 4.

We see that the non-homotheticity parameter ϵ is the key quantitatively relevant parameter for the purpose of determining structural change. When $\epsilon = 0$ and there are no non-homotheticities, all parameter combinations underpredict the agricultural value added share in poor countries. Similarly, if $\epsilon = 1$ and the spending on agriculture is insensitive to income, we consistently overpredict the agricultural value added share in poor countries. When $\epsilon = 0.68$, all parameter combinations approximately capture the structural change in value added. Note that this is true despite us testing large changes in α and very large changes in γ.[10]

[10] We do not vary δ and β, but these can be shown to not change the structural change predictions. Since all countries have the same δ and β, changing these parameters does not affect relative prices. Furthermore, since the model is calibrated to fit at y_{US}, all changes in the levels of relative prices will be compensated by changes in the calibrated share parameters in the utility function and the agricultural production function.

Figure 2: Predicted and observed value added shares in agriculture

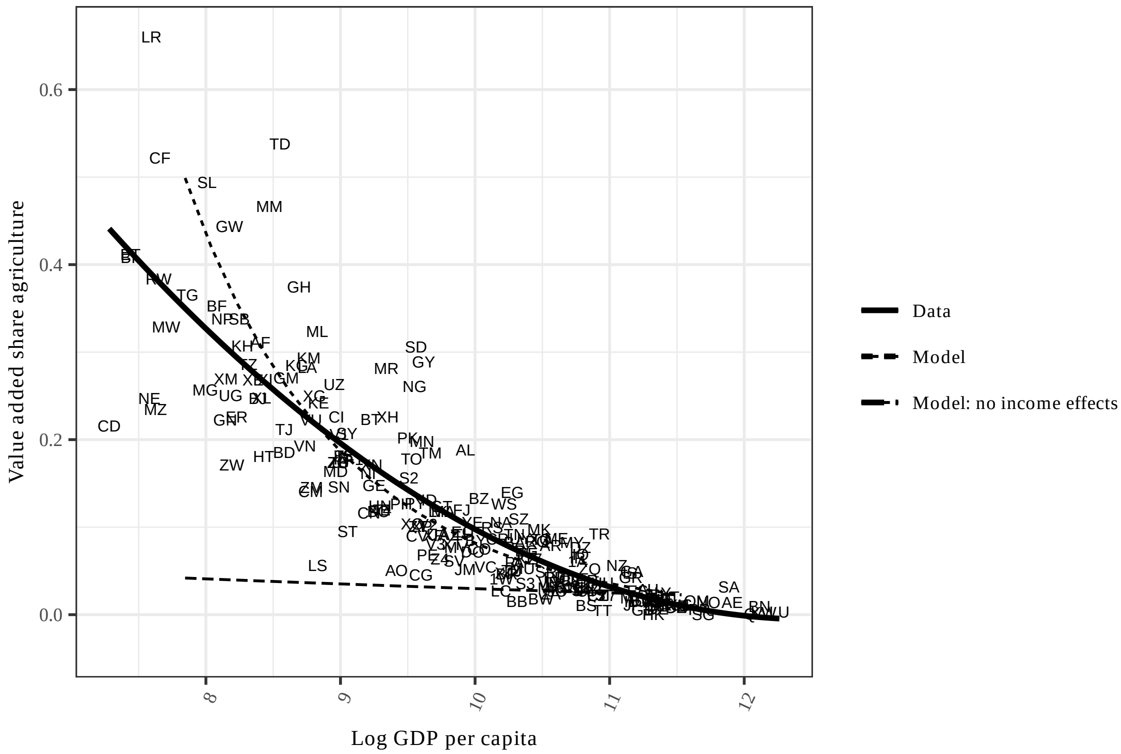

Figure 3: Predicted and observed employment shares in agriculture

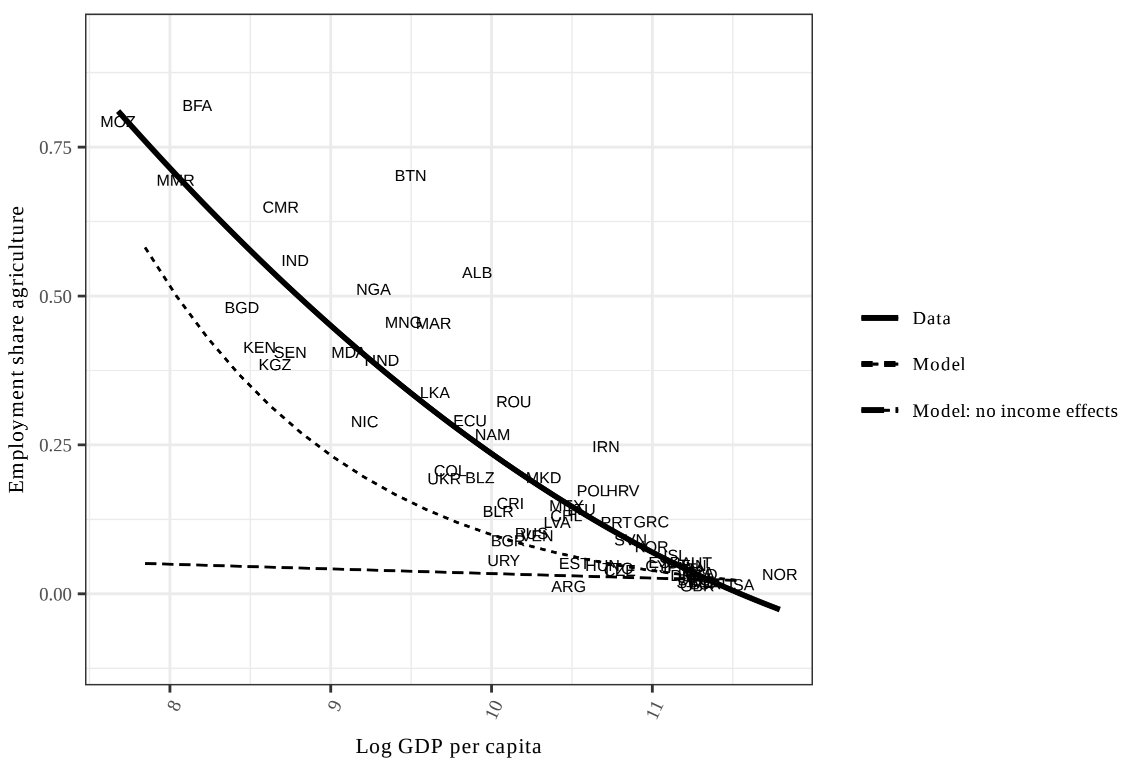

Table 2: Calibration results

	Parameter	Value
Production:		
	α	0.33
	δ	0.08
	a_k	0.18
	a_x	0.93
	a_r	0.22
	η	1.86
	σ	1.38
Preferences:		
	β	0.96
	ν	0.07
	γ	0.35
	ϵ	0.68
Non-agricultural TFP:		
	$\xi = \frac{d \log A_n}{d \log y}$	0.58

7 Counterfactual experiment

We use our calibrated model to conduct several counterfactual exercises, with the aim of comparing the effects of closing agricultural versus non-agricultural productivity gaps between poor and rich countries. The results of these exercises have potential salience in a number of key policy debates, most notably whether policies in poor countries should prioritize productivity improvements in one sector versus another. The answer to this question depends crucially on whether improvements to non-agricultural productivity induce improvements to agricultural productivity as well, something our model is particularly well equipped to assess.

Our model also allows for straightforward and direct testing of several claims regarding the importance of agriculture commonly made in the literature. We focus on one, that the particularly large labor productivity differences observed in agriculture are the proximates ource of income differences between poor and rich countries. This argument has relied on simple decomposition exercises, motivated by the compelling observation that poor countries appear to concentrate their workforce in the sector in which they are the least productive. Our experiment compares the results of this style of decomposition exercise to our model-based counterfactual exercise. We find the two exercises arrive at very different results.

7.1 Setup of counterfactual exercises

In our experiment, we analyze the effects of improving agricultural versus non-agricultural productivity in a benchmark poor country. We define this benchmark poor country to have 2.5% of the US's GDP per capita (i.e., $y_{poor}/y_{US} = 0.025$), so that the rich-poor income gap is 40. Using this benchmark poor country, we conduct two exercises: 1, a simple decomposition exercise in the spirit of Caselli (2005); and, 2, a model-based counterfactual taking into account general equilibrium forces.

In the first exercise, we closely follow the setup of Caselli (2005). We begin by calculating the poor

Figure 4: Predicted and observed employment shares for parameter perturbations

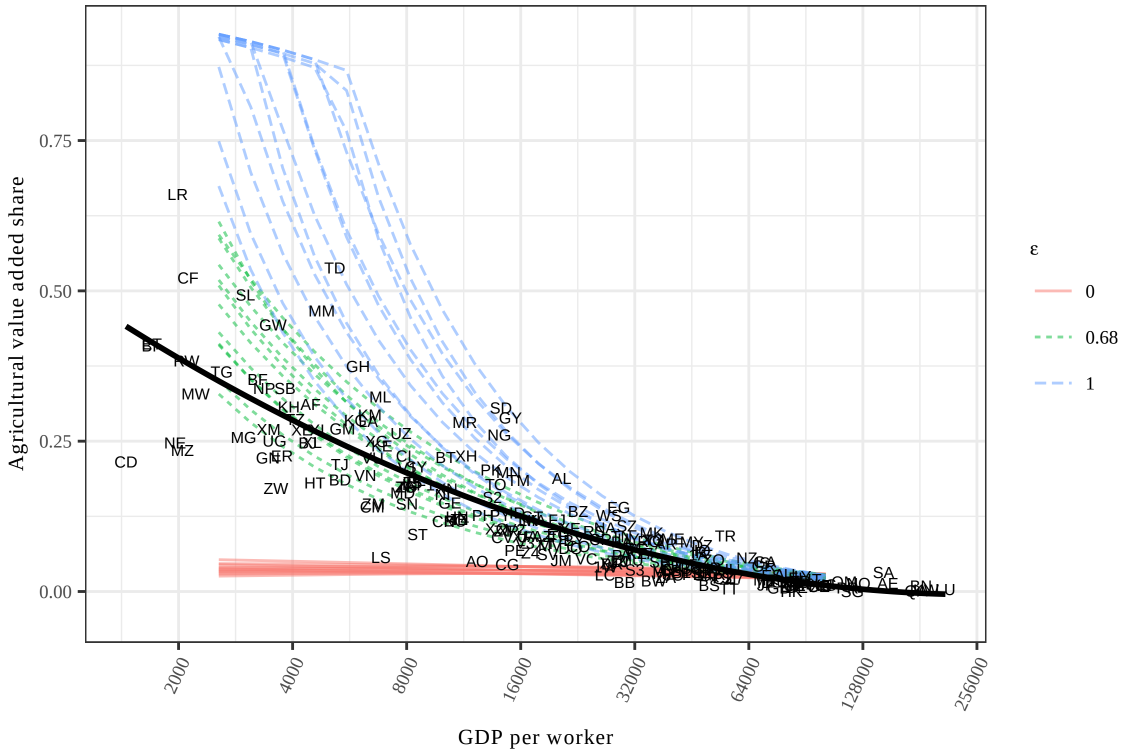

country's (real) agricultural labor productivity in the model,

$$y_a(y_{poor}) = \frac{VA_a^{real}(y_{poor})}{L_a(y_{poor})},$$

where we use a chain-weighted value-added measure of real output. We likewise follow Caselli in calculating the poor country's non-agricultural labor productivity,

$$y_{na}(y_{poor}) = \frac{y - y_a(y_{poor})L_a(y_{poor})}{1 - L_a(y_{poor})},$$

where $y_{na}(y_{poor})$ is defined such that the employment-weighted average of the sectoral labor productivities equals aggregate GDP per worker. Labor productivity in both sectors is further normalized such that both equal 1 in the US.

We then calculate two counterfactual values for GDP per worker in the poor country:

$$\hat{y}_a^{decomp} = [1 - L_a(y_{poor})]y_{na}(y_{poor}) + L_a(y_{poor})$$
$$\hat{y}_{na}^{decomp} = [1 - L_a(y_{poor})] + L_a(y_{poor})y_a(y_{poor})$$

Here, $\hat{y}_a^{decomp}$ is the counterfactual GDP per worker if agricultural labor productivity differences are eliminated, while $\hat{y}_{na}^{decomp}$ is the counterfactual GDP per worker if non-agricultural labor productivity differences are eliminated.

In the second exercise, we use a model-based counterfactual to conduct two experiments roughly equivalent to the first exercise, this time varying sectoral TFP as opposed to sectoral labor productivities. Using the model, we calculate two counterfactual values for GDP per worker in the poor country, as before:

$$\hat{y}_a^{model}(y) = \hat{y}[A_a(y), A_{na}(y_{poor})]$$
$$\hat{y}_{na}^{model}(y) = \hat{y}[A_a(y_{poor}), A_{na}(y)]$$

where $y \in [y_{poor}, y_{US}]$. Here, $\hat{y}_a^{model}(y)$ is the counterfactual GDP per worker when varying only agricultural TFP, and $\hat{y}_{na}^{model}(y)$ is the counterfactual GDP per worker when varying only non-agricultural TFP.

7.2 Results

Table 3 summarizes the results of the counterfactual experiment. Also displayed in the table are the income and sectoral productivity gaps of the benchmark poor country, which show that our model correctly features a relatively unproductive agricultural sector in the poor country. This fact is reflected in the results of the simple decomposition exercise, with the closing of the agricultural labor productivity gap reducing income disparities by moderately more than the closing of the non-agricultural gap, to a factor of just 1.6 versus 2.4. Given this exercise holds employment shares constant, and the fact that 60% of employment is in agriculture in the poor country, this result is unsurprising.

This contrasts sharply with the results of the model-based counterfactual exercise, presented towards the bottom of the table. When general equilibrium effects are taken into account, the income gains from eliminating the agricultural TFP gap are much more limited, only reducing the income gap to a factor of 14.1, while the gains from eliminating the non-agricultural TFP gap are nearly identical to before, leaving a residual income gap of just 2.5. These results can also be seen in Figure 5, which plots the counterfactual GDP per worker in the poor country from our model-based exercise. The red line shows the counterfactual GDP per worker when non-agricultural TFP is kept fixed, but where agricultural TFP has the value associated with that income level. Conversely, the blue line fixes agricultural TFP but lets non-agricultural TFP vary.

The key mechanism behind the different results of our two exercises is the endogenous response of employment and the value added share in agriculture to changes in productivity. The simple decomposition exercise, by design, holds sectoral employment shares fixed. Furthermore, it implies a dramatic increase in the value added share in the sector whose labor productivity is improved, since prices are held fixed. In contrast, the model-based exercise features a strong movement out of the agricultural sector in response to productivity changes. Thus, even as TFP strongly increases in the agricultural sector, the agricultural employment share rapidly falls, mitigating the ultimate income gains.

Figure 6 illustrates this mechanism, showing the effect of an increase in sectoral TFP on different variables. As in Figure 5, the red line plots the response to increasing agricultural TFP, holding non-agricultural TFP fixed, and the blue line plots the response to increasing non-agricultural TFP, holding agricultural TFP fixed. First, the agricultural employment share falls considerably in response to an increase in either sectoral TFP, although roughly 20% of employment remains in agriculture when only increasing agricultural TFP.

Second, even after closing the agricultural TFP gap, an agricultural labor productivity gap remains. This is the result of far lower amounts of capital and intermediate inputs being used per worker in the poor country, something that is not affected by the closing of the agricultural TFP gap. Here, the contrast with closing the

Table 3: Simple decomposition vs. model-based counterfactual

	Variable	US/Poor ratio
Poor country benchmark:		
	GDP per worker	40
	Agricultural labor productivity	71
	Non-agricultural labor productivity	24
Simple decomposition:		
	Closing agricultural gap	1.6
	Closing non-agricultural gap	2.4
Model-based counterfactual:		
	Closing agricultural gap	14.1
	Closing non-agricultural gap	2.5

non-agricultural TFP gap is very clear. Closing this gap results in strong agricultural input intensification, driving up gross output per worker even as value added per worker increases only modestly. This increased input intensification also partly explains the much faster decline of the agricultural employment share in response to non-agricultural TFP improvements. As more and more capital and intermediate input goods are used in agricultural production, the relative demand for non-agricultural goods rises, speeding up the structural change out of agriculture.

References

Adamopoulos, T. and Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6):1667–1697.

Bick, A., Fuchs-Schundeln, N., and Lagakos, D. (2018). How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications. *American Economic Review*, 108(1):170–199.

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). Chapter 9 Accounting for Cross-Country Income Differences. *Handbook of Economic Growth*, 1:679–741.

Caunedo, J. and Keller, E. (2017). Capital obsolescence and agricultural productivity.

Chen, C. (2017). Technology Adoption, Capital Deepening, and International Productivity Differences.

Donovan, K. (2016). Agricultural risk, intermediate inputs, and cross-country productivity differences.

Eckert, F., Peters, M., et al. (2018). Spatial structural change. In *2018 Meeting Papers*, number 98. Society for Economic Dynamics.

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

Figure 5: Effects of closing sectoral TFP gaps with rich countries

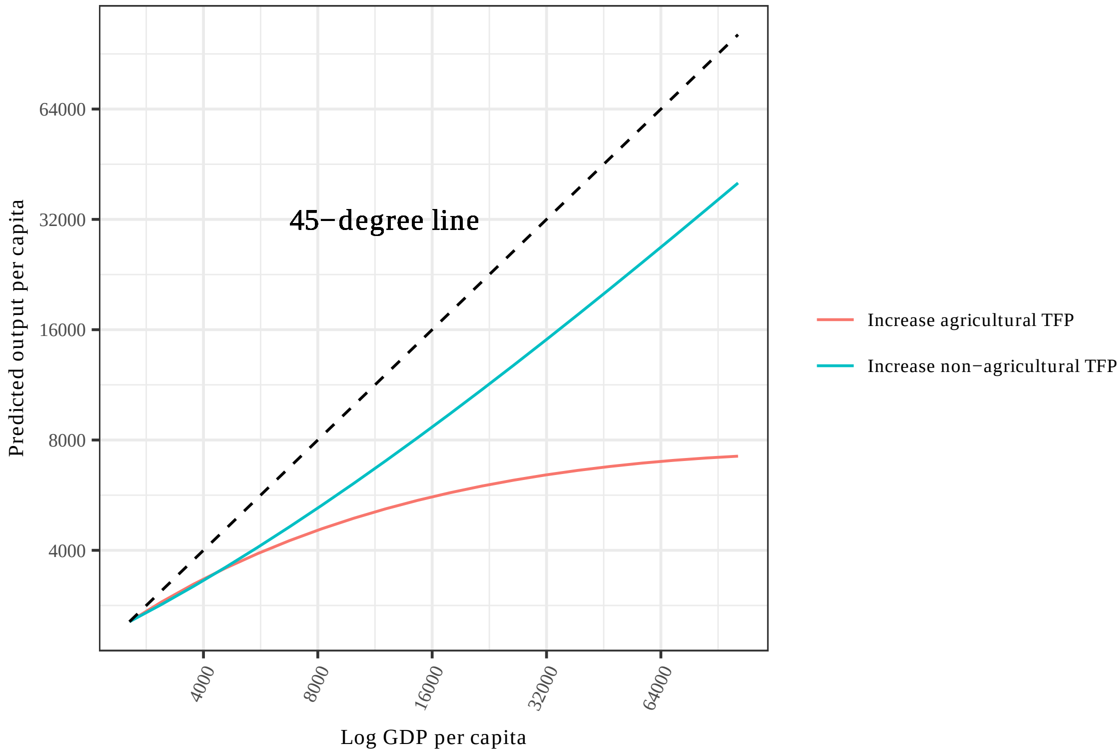

Hall, R. E. and Jorgenson, D. W. (1967). Tax Policy and Investment Behavior. *American Economic Review*, 57(3):391–414.

Hsieh, C.-T. and Klenow, P. J. (2007). Relative prices and relative prosperity. *American Economic Review*, 97(3):562–585.

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.

Figure 6: Variable changes under the two experiments