The Macroeconomics of Intensive Agriculture

Timo Boppart* Patrick Kiernan† Per Krusell‡ Hannes Malmberg§

February 2022

Abstract

Developing countries have a very large share of their populations in agriculture, at the same time as agricultural labor productivity is particularly low. We take a macroeconomic approach in order to examine the agricultural sector, and its importance for GDP, along the path of development. To this end, we collect systematic measures of inputs and outputs in agricultural production around the globe. Our data show massive capital deepening and intermediate-input intensification along the development path. These patterns are in line with neoclassical forces and can account for roughly two-thirds of the agricultural labor productivity gap between the richest and the poorest countries. The effect is non-linear: agricultural TFP is similar across middle- and high-income countries and the gap between their labor productivities is entirely accounted for by input intensification. Furthermore, we show that an aggregate agricultural production function, with input substitutabilities significantly above 1, accounts very well for our stylized facts on input quantity and price ratios. On the demand side, we document that a standard non-homothetic formulation captures the expenditure share of agriculture very well in the whole cross-section of countries. We embed the results in a closed-economy general-equilibrium model with minimal distortions and show that non-agricultural TFP differences are more powerful than agricultural TFP differences in explaining income differences.

*timo.boppart@iies.su.se, Institute for International Economic Studies, University of St. Gallen, and CEPR
†kiern018@umn.edu, University of Minnesota
‡per.krusell@iies.su.se, Institute for International Economic Studies, NBER, and CEPR
§pmalmber@umn.edu, University of Minnesota

1 Introduction

The difference in income per capita between rich and poor countries is on average at least a factor of 40—an absolutely staggering gap. Income per capita is of course not directly a measure of welfare, but recent attempts to broaden the measure to include other observables—see Jones and Klenow (2013)—actually suggest that the welfare gap is even bigger, and that income per capita still accounts for the bulk of the welfare differences between rich and poor countries. In this paper we make an attempt to interpret these differences in income per capita by studying the agricultural sector from a macroeconomic perspective. The simple reason behind this is that agriculture appears to be of first-order importance: going across the development spectrum, the share of the population working in agriculture drops from around 80 percent to below one percent. That is: income in the poorest countries mostly reflects how well the agricultural sector is doing, whereas in the richest countries this sector is largely irrelevant. Figure 1 illustrates.

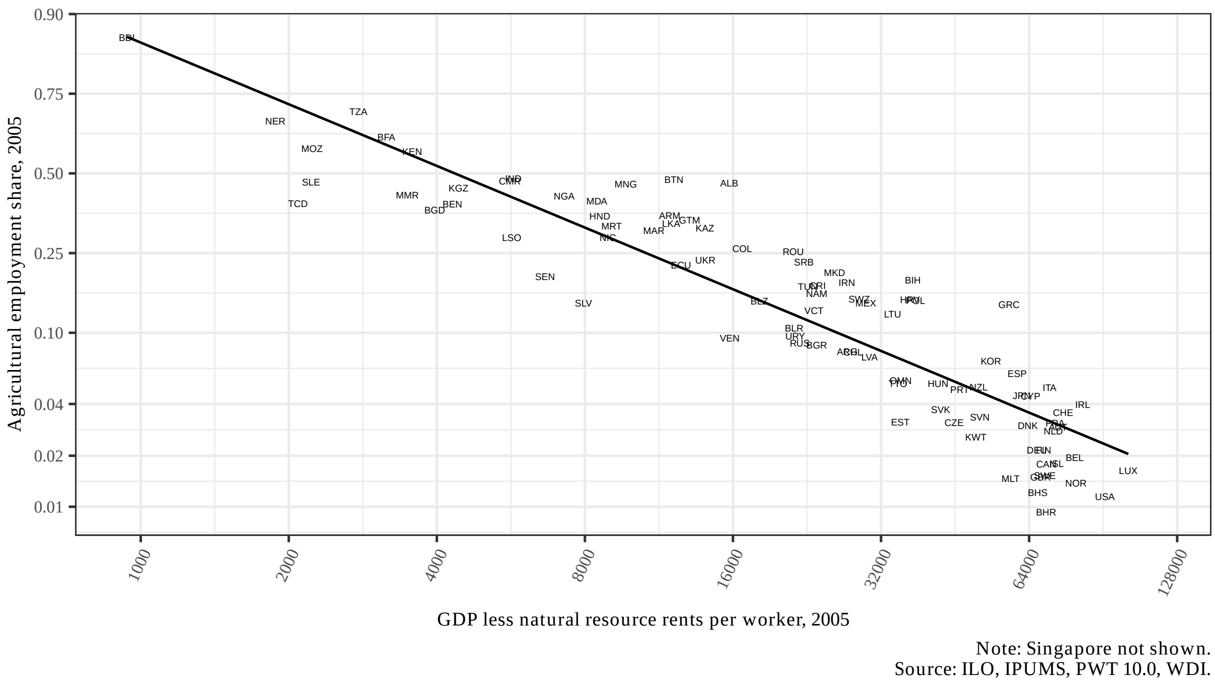

Figure 1: Agricultural employment shares across countries

The figure also shows that across the entire development spectrum there is a quite stable negative empirical relationship between the logarithms of the share of agricultural employment share and GDP per worker. Given the fact that hours worked is, on average, significantly higher in poor countries (see, e.g., Bick et al., 2018), clearly there is also a huge average productivity gap, so a related question is to what extent does productivity in the agricultural sector contribute to the overall gap. It turns out that the gap in labor productivity in agriculture between rich and poor countries is even larger—about a factor of 100. This fact, which we document below, has of course been noted before, as the following two quotes illustrate:

> *"A decomposition of aggregate labor productivity based on internationally comparable data reveals that a high share of employment and low labor productivity in agriculture are mainly responsible*

for low aggregate productivity in poor countries."
(Restuccia, Yang, and Zhu, 2008)

"... agricultural productivity is critical for understanding aggregate income differences."
(Donovan, 2020)

Motivated by these facts, this paper analyzes the agricultural sector across the development spectrum. What explains them, and the mechanisms whereby countries gradually specialize less and less in agriculture as they develop, not only appears important to understand from an intellectual point of view but also must be relevant for policymakers. In our efforts, we take a macroeconomic perspective and go through the steps that are traditionally taken in aggregate macroeconomic analysis, following Solow (1956, 1957) and many others. We collect systematic measures of outputs and inputs in agricultural production around the globe. We distinguish four input aggregates: capital, labor, intermediate goods (pesticides, fertilizers, etc.), and land. For these, we document massive capital deepening and intermediate-input intensification along the development path. The intermediate-input share of gross output, for example, goes from close to zero for the poorest country to around forty percent for the richest; the capital share goes from approximately zero to about ten percent. At the same time, the quantities of intermediate inputs and of capital, say, relative to labor, go up tremendously—by factors of 1,000 and 500, respectively. These magnitudes suggest that factor substitution is an important phenomenon, also relative to purely technological factors.

Moreover, the data we uncover have very clear neoclassical features: relative prices between inputs are systematically negatively related to relative input quantities, and in log-log space these relationships are approximately linear, though with different slope magnitudes for different input combinations. Quite strikingly, these relationships not only characterize the cross-section of countries but also the U.S. time series. These findings, in fact, motivate our approach, which is to look at the data from a macroeconomic perspective using an aggregate neoclassical production function for agriculture.

Next, we use a variant of Solow's TFP accounting—thus, we use a minimum of theory—to generate a time series that represents average agricultural TFP along the development spectrum. The endogenous forces—input intensification—account for roughly two-thirds of the agricultural labor productivity gap between the richest and the poorest countries mentioned above, with TFP accounting for the remainder. The effect is non-linear: agricultural TFP is very similar across middle- and high-income countries and the gap between their labor productivities is entirely accounted for by input intensification.

We take one more step to characterize the supply side by employing more specific assumptions: we posit a production function which is a nested CES in inputs and proceed to estimate the relevant parameters. Thus, our starting point is that a single production function for agriculture characterizes production at different levels of development, with differences only in TFP. The observed input substitution implies elasticities between labor, intermediate inputs, and capital that are significantly above unity, whereas for land a unitary elasticity is a decent approximation. The high substitution elasticities make factor intensification along the development path a central phenomenon. In fact, overall, our formulation fits the data remarkably well over the entire development spectrum. We formulate another aggregate production function for all non-agricultural goods, now only using capital and labor as inputs, and show that a Cobb-Douglas structure fits this second sector quite well, along with a TFP series for non-agriculture specified separately for consumption, investment, and intermediate goods. We find that the TFP gaps for intermediate goods are the largest, with investment goods second. Given all this, key to our account of the development path is that poor performance in the non-agricultural sectors is not just important in itself but also for low labor productivity in agriculture, as input substitution is significantly hampered.

On the demand side, we use a standard non-homothetic formulation and show that it can capture the concave shape of log agricultural consumption share very well. One drawback is that the implied income elasticity for agricultural goods – 0.35 – falls short of results from micro studies which find values around 0.6-0.7. This inconsistency does not seem reflect idiosyncracies of our particular model, but come quite directly from the data: even with a zero price elasticity giving maximum scope for substitution effect, the observed expenditure and share data requires a low income elasticity.

Equipped with a fully specified demand and supply side we formulate a full general-equilibrium theory that can be used to run counterfactuals. Our combination of assumptions on the demand and supply sides do very well at accounting for the cross-sectional facts, but they do imply significant non-linearities, especially since we make predictions across the entire development spectrum. Thus, Hulten's (1978) logic applies but only locally and very quickly the input substitution forces and structural change, with changing shares, take over and dominate quantitatively.

We assume closed economies in our general-equilibrium analysis, along with minimal distortions. This point of departure is natural but, of course, not the final word; trade is arguably important, though trade flows are of course to some extent also hindered by trade frictions, and so are distortions, especially those making a transition for workers from rural to urban areas challenging. Our general-equilibrium theory has independent predictions for the two sectors' shares of aggregate employment as well as their shares of aggregate output. All these are fit closely, reflecting the good partial equilibrium fits of production and preferences, as well as the fact that we took care to ensure that our measures were internally consistent.

Overall, as already indicated, it is differences in TFP in investment and intermediate-goods production that account for the bulk of aggregate output differences across countries. A particularly telling counterfactual is that where we raise agricultural TFP alone, from the level estimated for the very poorest country, gradually toward that of the TFP leaders in agriculture. We find important initial improvements—especially for welfare, as agricultural goods are key at low levels of wealth—but the benefits of further increases in agricultural TFP rapidly wither: the low TFP in the remaining sectors becomes key as goods demand now shifts away from agriculture, at the same time as agriculture itself suffers from lacking capital and intermediate goods.

Finally, we use our theory in an application: that of the likely effects of climate change in reducing the availability of arable land. In the IPCC reports, attempts are made to estimate by how much production will fall as a result of the projected decline in arable land. These calculations are typically based on Leontief input elasticities, and we complement their analysis by offering estimates based on our production structure. We find that even with large declines in land availability, agricultural is sustained through input intensification where factors are shifted to the agricultural sector. Thus, in equilibrium, non-agricultural consumption falls much more than agricultural consumption in response to a reduction in land supply.

The paper is organized as follows. As we view one of our main contributions to be the systematic collation of data on quantities and prices in agriculture, we begin in Section 2 by discussing our data sources. The sources are many and the quality of the data particularly important to discuss; for that reason, a long appendix accompanies this early section. The section then shows the main facts that we use as inputs into our study. Section 3 shows how we construct agricultural TFP along the development spectrum based on Solow-like assumptions. Section 4 then makes functional-form assumptions and estimates a fully parameterized production function for agriculture; we also estimate one for non-agricultural production. In Section 5, we specify a non-homothetic utility function in agricultural and non-agricultural goods and apply it to our data set. Equipped with both a supply and demand side, we then develop our general-equilibrium structure in Section 6, calibrate it—though here most of the parameters have already been pinned down earlier—and use it for our counterfactual analyses. Section 8 finally looks at land reductions due to climate change and Section 9 concludes. But before we begin let us also situate the paper in a literature context; our paper is certainly far from the only one taking the broad approach here.

1.1 Related literature

Our paper deals with two classic development questions: how development leads to structural change out of agriculture (Lewis, 1954; Kuznets 1966) and how countries solve the "food problem" (Schultz, 1953). By relying on a macroeconomic development approach and providing a quantitative general-equilibrium theory of growth and structural change, our paper is in the spirit of Gollin, Parente, and Rogerson (2002, 2007), Restuccia, Yang, and Zhu (2008), and others. In line with this literature we provide a dynamic closed-economy multi-sector model. By combining a production side with a non-homothetic demand side, our approach is similar to the one in Donovan (2018) (emphasizing the role of intermediate inputs), Chen (2020) (focusing on capital deepening), and Gottlieb and Grobovsek (2019) (considering the role of land markets). The paper closest to ours is Restuccia, Yang, and Zhu (2008), who consider labor, land, and intermediates as factors in agricultural production. We study and quantify the importance of four agricultural production factors: labor, land, intermediates, and capital.

We provide new cross-country and U.S. time-series data on factor prices and quantities in agricultural production. In a first step we use these data to back out Solow (1957) residuals, here applied across different levels of development as opposed to over time. Our residual TFP terms across countries are comparable to the results in the development accounting literature (Klenow and Rodriguez-Clare, 1997; Hall and Jones, 1999; and others). However, as in Caselli (2005) we highlight the special role of the agricultural sector and emphasize the importance of differences in relative prices of investment goods (see Hsieh and Klenow, 2007) to account for aggregate income differences. As in Caselli and Coleman (2001) we examine the implications of agricultural productivities on regional income differences, however with an application to cross-country data as opposed to U.S. states.

In backing out agricultural TFP terms our paper speaks to the origin of the observed agricultural labor productivity gaps (see Herrendorf and Schoellman (2015) and Gollin, Lagakos, and Waugh (2014)). Adamopoulos and Restuccia (2014) quantify the role of fertile land, whereas Caunedo and Keller (2021) emphasize the importance of capital upgrading.

In a second step we fit a neoclassical production function to the data on agriculutural inputs and factor prices. We show significant deviations from a Cobb-Douglas production function, with pairwise elasticities of substitution that are above one. For the factors capital and labor, Storesletten, Zhao, and Zilibotti (2019) and Chen (2020) also emphasize more substitutability in agriculture than Cobb-Douglas. Our result is also reminiscent of Herrendorf, Herrington, and Valentinyi (2015), who estimate for the post-war U.S. agricultural sector an elasticity of substitution between capital and labor of 1.58. As our data asks for a Cobb-Douglas aggregate production function over capital and labor in the non-agricultural sector, our model features structural change with different elasticities of substitution across sectors, as studied in Alvarez-Cuadrado

WDI	World Bank's World Development Indicators, multiple sub-sources
ILO	International Labour Organization; the underlying source for sectoral employment shares
PWT	Penn World Table version 10.0 database
FAO	Food and Agriculture Organization's FAOSTAT database
UN	United Nations Statistical Division
IPUMS	Minnesota Population Center's Integrated Public Use Microdata Series, multiple sub-sources
BL	Barro and Lee (2013) educational attainment dataset
CPR	Caselli, Ponticelli, and Rossi (2014) Mincerian returns estimates
WIOD	World Input-Output Database
AR	Adamopolous and Restuccia (2018) land quality data, based on GAEZ data
BACI	CEPII's Database for International Trade Analysis
LSMS	World Bank's Living Standards Measurement Studies
VDSA	ICRISAT's Village Dynamics in South Asia project
EU	European Commission's Eurostat database
USDA	U.S. Department of Agriculture National Agricultural Statistical Service's Quick Stats Database

Table 1: Data sources

et al. (2017). For the elasticity of substitution with respect to land our preferred estimate is close to Cobb-Douglas (see also Bustos, Caprettini, and Ponticelli (2016) and Leukhina and Turnovsky (2016), who emphasize elasticities of substitution with respect to land below unity).

On the preference side, non-homotheticities are inadmissable in order to be in line with Engel's law. The literature has to large extent relied on a Stone-Geary-type structure with a subsistence level of consumption (see, e.g., Caselli and Coleman (2001), Gollin, Parente, and Rogerson (2002, 2007), Restuccia, Yang, and Zhu (2008), Gottlieb and Grobovsek (2019), Chen (2020), etc.). We instead apply preferences in the price independent generalized linearity class which can generate a sustained income effect on expenditure shares (see, e.g., Boppart (2014), Eckert and Peters (2018), and Alder et al. (2022) for applications thereof in the strucutral change literature). Our specification nests the one in Gollin, Parente, and Rogerson (2002, 2007) as a special case.

Both the data and theory we provide is macroeconomic at its core. Nevertheless, it is important to relate our findings to the microeconomic evidence in the field (see Buera, Kaboski, and Townsend (2021)). The technology terms we back out suggest that there is particularly low TFP in agriculture in the very poorest countries.

2 Aggregate facts

2.1 Data sources

As mentioned, a key purpose of our paper is to put together comprehensive data series that other researchers can download and use. Given the scope of our analysis, we use a large set of different databases at different levels of aggregation. Table 1 lists these databases and other data sources, with abbreviations to facilitate the description below; the appendix contains more detailed references.

In our data construction, we aim to quality-adjust price and quantity measures whenever possible. The resulting series are projected on a development index, for which we use real GDP per worker. This projection lets us characterize the data along our dimension of interest even though not all countries have data for every measure. Real GDP is defined as nominal GDP net of natural resource rents, deflated by the PWT's PPP-adjusted GDP price deflator, where nominal GDP is obtained from the PWT and natural resource rents are obtained from the WDI. The total number of workers is taken from the PWT.

We now go through the variables that we use, beginning with the agricultural sector.

2.1.1 Agriculture: output and inputs

On the output side, we begin with the raw data from the FAO and the UN. They provide nominal agricultural gross output by country and year. The agricultural sector produces output that is also an input into the same sector (such as manure), and we use the input-output tables provided in the WIOD to net out this component and arrive at our preferred gross production measure for the agricultural sector. To obtain a measure of real gross output, we use the FAO's deflators for agricultural output.

On the input side, we start from observations of shares of agricultural revenues. Assuming total costs equal total revenues, we obtain nominal input values. For land, we use quantity data to derive prices; for intermediate inputs, we use price data to derive quantities. For capital and labor, we instead start from price and quantity measurements and then calculate shares directly.

Beginning with land, we take land shares from two different data sets. One is a directly calculated measure of land shares from EU and USDA for a large group of developed countries—land costs are average per hectare rental rates for agricultural land, which we divide by gross output per hectare to obtain our cost shares, assuming zero profits. The second is a measure for a number of poor countries in Africa and Asia: the LSMS and the VDSA provide surveys where, for farmers who rent land, the land rental rates and gross output are available. These shares are then projected onto GDP per worker.[1] The quantities of land are available from the FAO, measured in hectares. We then quality-adjust these measures using the data and calculations in AR (in turn derived from GAEZ grid data). This then also gives us land prices by dividing shares times total costs by quantities. Total costs are still to be determined below.

Turning to intermediates, we obtain shares from the UN and FAO national accounts, again netting out the within-sector inputs into agriculture using the WIOD and interpreting the resulting shares as cost shares. We then measure prices. The FAO provides some direct price estimates, but these are based on rather dated and sparse surveys. Instead, we construct price estimates using trade data from BACI: we obtain prices of a selected sample of internationally traded fertilizers and pesticides and aggregate up based on their respective weights. Given the shares and prices of intermediates, along with total agricultural costs (again to be determined), we deduce values for quantities.

For the capital input, the FAO gives a nominal measure of the capital stock used in agriculture, forestry, and fishing, which we specialize to agriculture by assuming that capital is used in proportion to the gross output shares of the respective sub-components. The price of aggregate capital is taken from the PWT and used to deflate these nominal values into real quantities of capital; here, we thus assume the same deflator for the agricultural sector as for the aggregate. We compute a user cost for each nominal unit of capital based on the aggregate nominal capital compensation, which is computed as aggregate output less the compensation to labor and land, divided by the nominal value of the capital stock.

For the labor input, we begin with employment shares in agriculture, forestry, and fishing, calculated directly for 85 countries using IPUMS microdata and taken from the ILO otherwise. As with capital,

[1] Our projections always use a quadratic function.

we specialize these shares to agriculture using the sub-sectors' gross output shares. Together with total employment from the PWT, we obtain employment by sector. We construct a measure of human-capital-adjusted employment using Mincerian returns to schooling from CPR and average years of schooling from BL. Finally, we apply an adjustment to the sectoral estimates aimed at capturing differences in efficiency units per hour between the two sectors. This efficiency-unit adjustment could be based on estimates using differences in education levels by sector, along with Mincer returns, or expert judgment (in this case Todd Schoellman). The former suggests 30 percent higher efficiency units per hour in non-agriculture versus agriculture, whereas the latter proposed a higher gap of roughly 100 percent, both fairly stable across the development dimension. Ultimately we employ the former approach, directly estimating average years of schooling by sector from IPUMS microdata. The price per efficiency unit of labor is calculated by multiplying nominal aggregate output per efficiency unit by the labor share of total compensation.

The total costs in agriculture are then computed based on knowing total costs of capital and labor, together with their shares.[2]

2.1.2 Non-agriculture

The non-agricultural sector is constructed by ensuring that nominal totals of agricultural and non-agricultural output sum to total nominal output for each country. The nominal value is then broken down into price and quantity by observing that we have already defined deflators for aggregate GDP and for agriculture, and we have also defined the nominal shares of agriculture and non-agriculture. These figures together allow us to construct a price deflator for non-agriculture.

Since we only consider two inputs in this sector—capital and labor—there is no distinction between gross output and value added. We explain above how we obtain values for nominal capital stocks and employment in agriculture; for non-agriculture, we simply subtract these amounts from aggregate nominal capital stocks and employment in the PWT, respectively. The deflator for capital is the same as that for the aggregate. The efficiency-unit adjustments and price calculations for labor are performed the same way as for agriculture.

There are three sub-components of non-agricultural output: investment, intermediates, and consumption. The relative price of investment is taken from the PWT. The relative price of intermediates, given they serve as agricultural inputs, was already constructed from trade data. The relative price of non-agricultural consumption is constructed using the PWT's price of total consumption, our price of agricultural output, and the agricultural gross output share of total consumption.

2.2 Stylized aggregate facts: agricultural production and consumption

In this section we present some key properties of the aggregates we have gathered. The facts we present motivate the approach we take in our ensuing analysis.

2.2.1 Production of agricultural goods: relative quantities and relative prices of inputs

Figure 2 displays the ratios of input quantities and their corresponding relative prices. We have 4 input categories—capital (K), labor (H, for human capital), intermediates (X), and land (L)—and so 6 different

[2]Total costs are thus total labor and capital compensation divided by one minus the sum of the cost shares of land and intermediates.

input ratios to look at. The x- and y-axes are log scales, so slopes can be thought of as reduced-form elasticities.

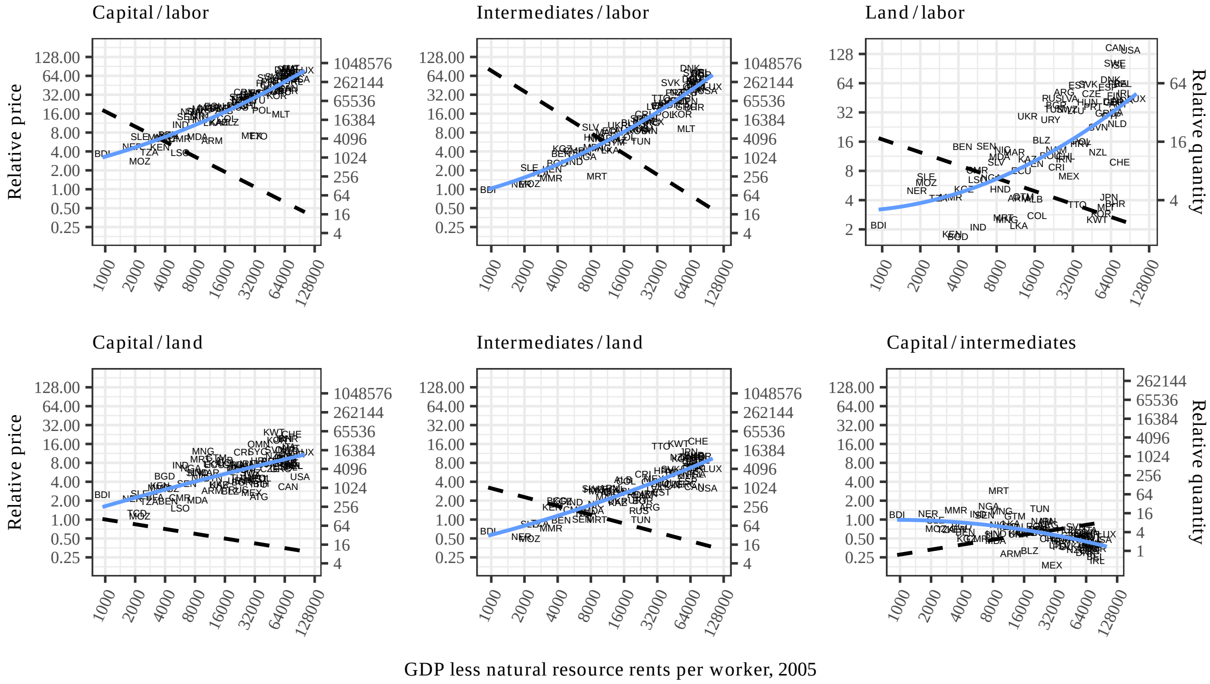

Figure 2: Agricultural inputs: quantity and price ratios

Each of the six sub-figures thus plots, with our development index (real GDP per worker) on the x-axis, both the quantity and price ratios for the two inputs in question. For quantity ratios, we plot the individual country observations along with their projection onto the development index (the solid blue line). For price ratios, to avoid cluttering, we plot only the projection (the dashed black line).

In the top left panel, we see the well-known fact that the ratio between the quantities of capital and labor inputs rises sharply along the development dimension: it goes up by a factor of roughly 1000 from the poorest economy to the richest. The ratio of the rental price of capital to the per-efficiency unit wage falls with development by a factor of around 30 over the whole range.

Similar patterns hold for the other input combinations: the price and quantity lines (i) slope in opposite ways and (ii) have very wide ranges between the two extremes of the development scale. For the land-to-labor and capital-to-intermediates ratios, the ranges of values are somewhat smaller but still sizable.

We conclude from these graphs that there is a striking, and robust, pattern involving quantity and price ratios along the development dimension: (relative) input prices in agriculture negatively co-move with their corresponding (relative) input quantities. The relationship is also surprisingly tight and involves changes in ratios by orders of magnitude. To us, this is suggestive of neoclassical forces. Where inputs are expensive, they are used less; the neoclassical production-function perspective is precisely that higher use of an input lowers its marginal product and, under competitive markets, its price. Our observations also suggest that the formulation of an aggregate production function for agriculture has a chance of providing a good account of the data, even across a broad development spectrum.

Before proceeding with the formulation of an aggregate production function, let us also look at the time-series perspective. In Figure 3 we see the same cross-sectional data as in the previous figure but now with the U.S. time series added in red; here, the data covers the period 1950–2018.[3] Nearly all of the underlying data come from the BEA, with land quantities after 1960 coming from the FAO and all other land data coming from the USDA.

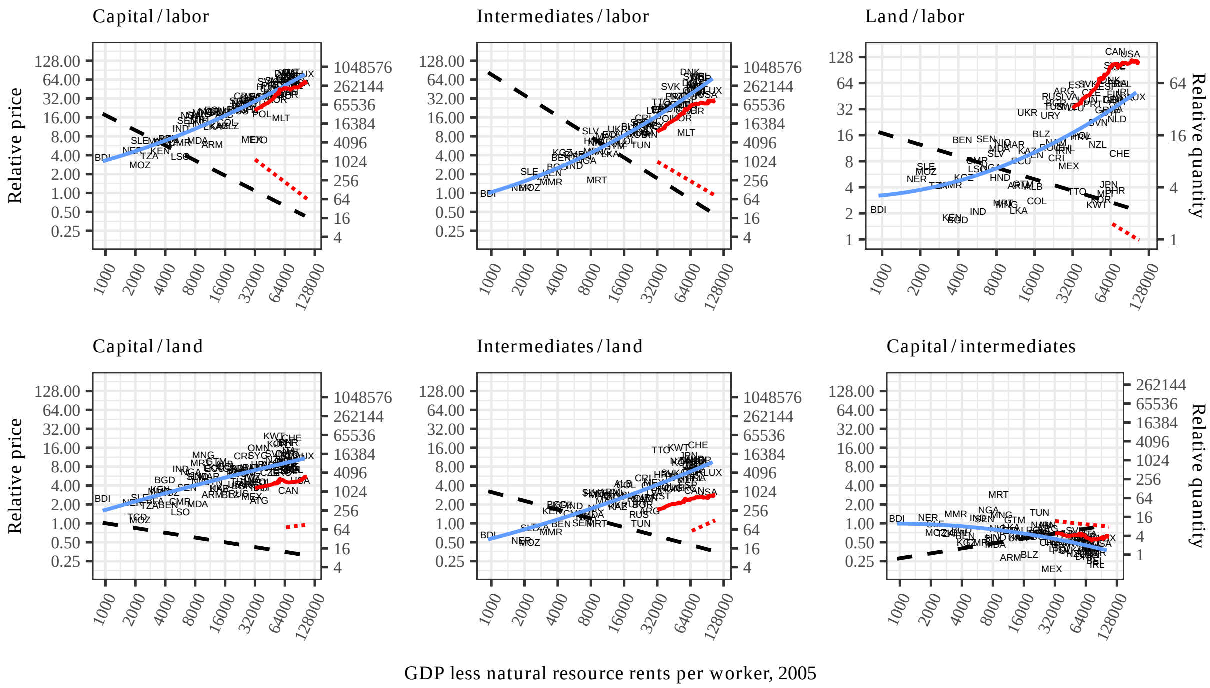

Figure 3: The cross section and the U.S. time series

The figure shows that the time-series data from the U.S. is broadly in line with the cross-sectional patterns. Although the levels appear to be somewhat off the cross-country trends for some ratios, the time-series trends closely match the (often exceptional) U.S. levels in our cross-country data. Notably, the slopes of the relationships do not appear appreciably different from those of the cross-section.

Further examination of the previous figures reveals (i) that slopes vary by input combination and (ii) that slopes of relative quantities do not sum to zero, i.e., shares are not constant along the development dimension. Figure 4 shows how shares vary. The left panel displays the input cost shares of gross agricultural output; the right panel displays shares of agricultural value added.

Beginning with the difference between the two graphs in the figure, we see that the share of intermediate inputs out of gross output is nearly zero for the poorest countries but over 50 percent for the richest. This finding, in fact, is a strong motivation for including intermediate inputs in our analysis. As discussed above, [reference:xyz Donovan etc.] also do look at the role of intermediates.

Turning to the remaining inputs, we see that, in contrast, the land share out of value added is quite constant at slightly over 20 percent; as a share of gross output, it is roughly cut in half as intermediate inputs rise from zero to 50 percent. The capital share of value added increases significantly—from less than 10 percent to around 40 percent. Finally, the labor share declines markedly, from roughly 70 percent (of both

[3] Our land price measure begins in 1990.

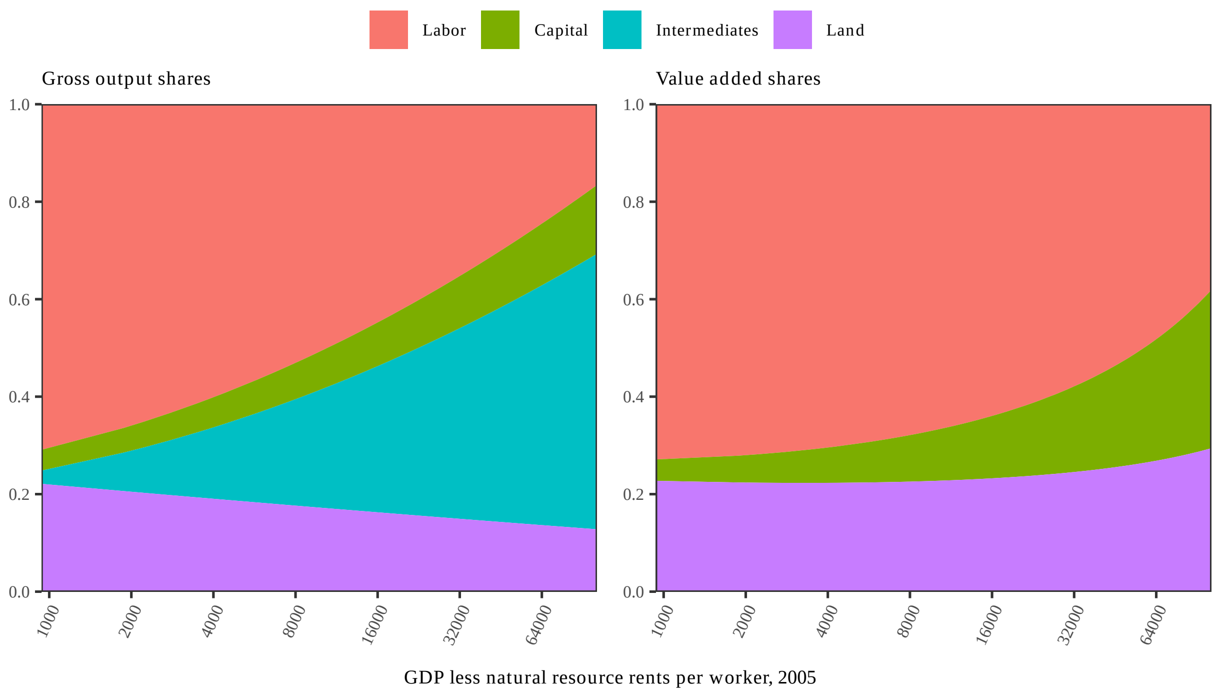

Figure 4: Input shares—gross and value added

gross output and value added) to below 40 percent of value added and half of that of gross output. Thus, in the richest countries, the labor share in agriculture is significantly below the labor share in non-agriculture. In sum, we note rather striking movements in shares. We also note that all of our four inputs have shares that account for a significant chunk of agricultural costs at least at some stage in development.

2.2.2 Consumption of agricultural goods: shares

Having studied determinants of the supply of agricultural output, we now look at the demand side. Figure 5 shows how the budget share of agricultural goods in aggregate consumption expenditures declines with development: from around 50 percent on average for the poorest economies to close to zero for the richest.

Figure 6 shows how the price of agricultural goods relative to non-agricultural consumption moves with our development index.

Here, we also have striking fact: agricultural goods become markedly less expensive relative to non-agricultural consumption moving from the poorest to the richest economies. The change in this relative price—a factor of a little over 4—is not as large as those we saw for agricultural inputs, but is still an important fact to relate to. Below, we will formulate a demand system that is consistent with the data in Figures 5 and 6. Not surprisingly, given the massive movement in the share, the theory will feature non-homotheticities in line with Engel's law.

The data in Figure 6 also give us a first hint of the relative total-factor productivities in the agricultural sector relative to the non-agricultural sector, again in the development dimension. If, namely, the production technologies in the two sectors were identical up to their TFP factors, and if markets operate without frictions so that inputs are allocated efficiently across sectors, then the observed relative price would reveal the relative

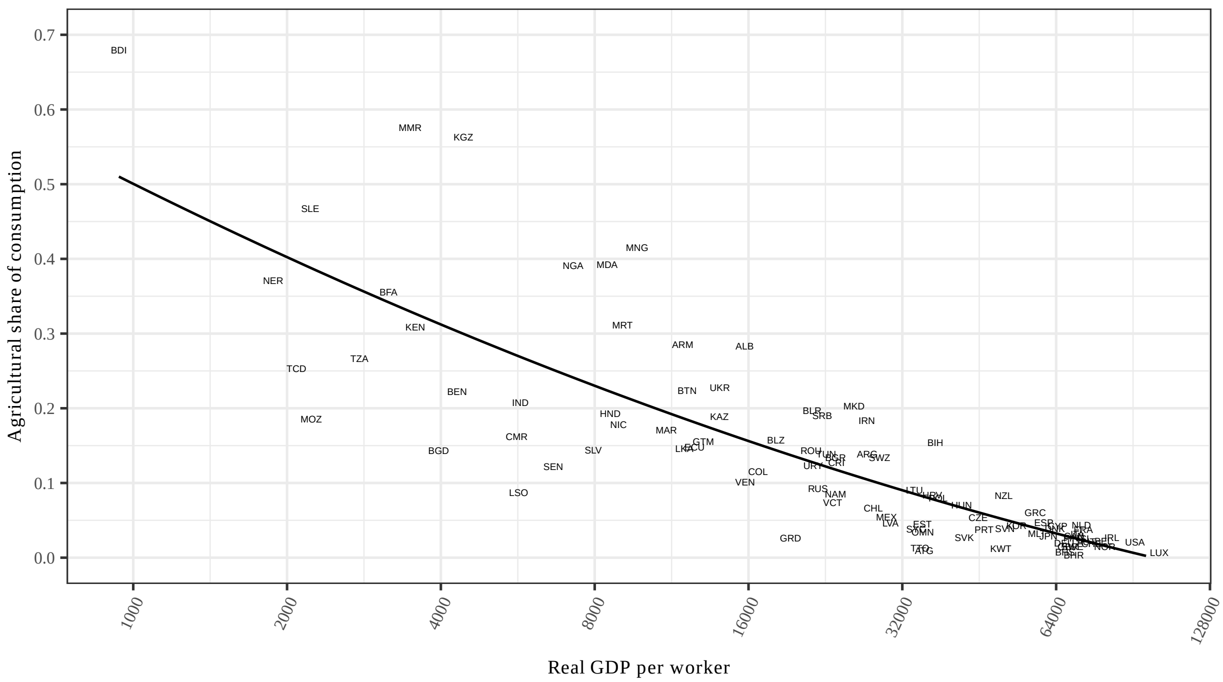

Figure 5: The share of agriculture in aggregate consumption, 2005

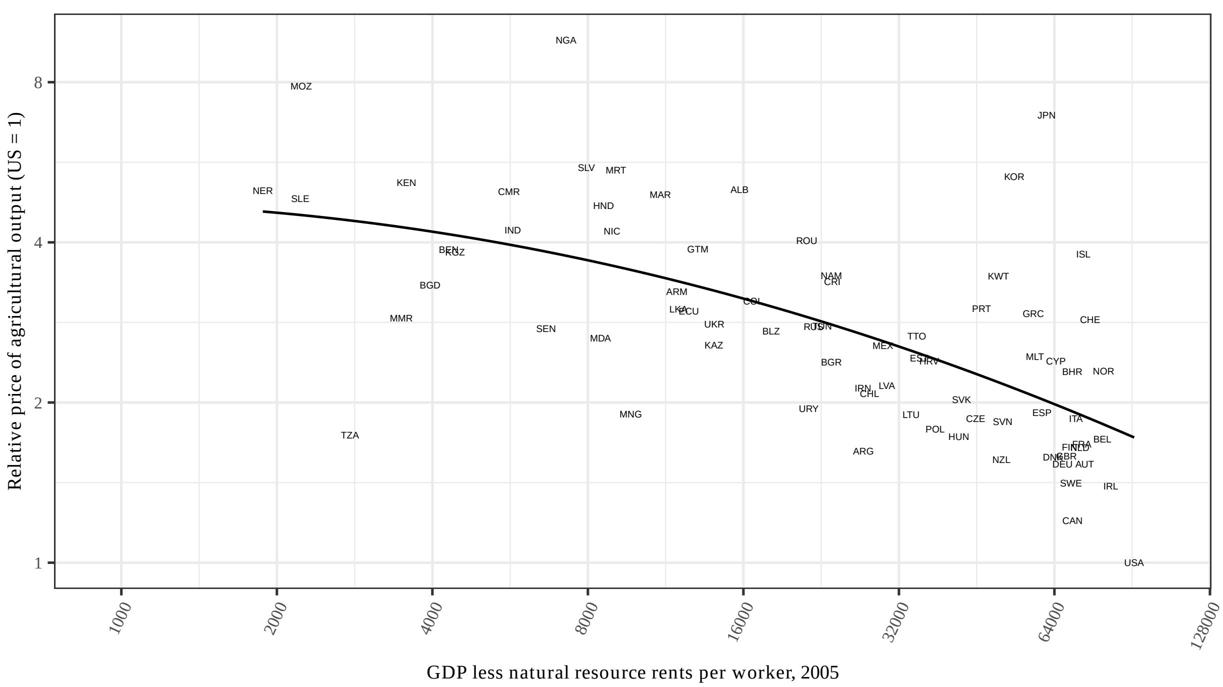

Figure 6: The relative price of agricultural goods

TFP levels in the two sectors.[4] Hence, the figure would suggest that agricultural TFP, relative to average

[4]See, e.g., Greenwood, Hercowitz, and Krusell (1997) for such an example, applied to consumption vs. investment goods.

TFP, declines with development. However: not only is this intuitively implausible, but the data on shares above strongly suggest that the production structures are very different across the two sectors. Hence, measuring how TFP in agriculture changes with development requires a more elaborate strategy, one that actually does not look at prices and that also involves much less theory than in the above-mentioned first pass.

3 Theory, I: agricultural TFP

The present section uses a minimum of theory—essentially that employed in Solow (1957), which is the basis for all growth accounting—on the agricultural sector. Thus, the core idea is to use the following, standard assumptions: (i) input (the four mentioned above) and output quantities are measurable; (ii) input prices are measurable; (iii) the production function has constant returns to scale (CRS; and it is quasi-concave, differentiable, etc.); (iv) input markets are perfectly competitive; and (v) firms maximize profits. These assumptions appeared like a good first approximation to Solow and they do here as well, given that we think of outcomes across space, and not across time.[5]

Thus, we use the production function

$$y_{a,t} = A_a f(k_{a,t}, h_{a,t}, x_{a,t}, l) \tag{1}$$

where the a index refers to agriculture and where y is output, k is capital, h (for human capital) is labor, x is the intermediate good input, and l is land. The function f thus has CRS and standard properties and A is total-factor productivity (TFP).

Similarly, for non-agricultural goods and services we assume

$$y_{n,t} = A_n g(k_{n,t}, h_{n,t}), \tag{2}$$

where the n index refers to non-agriculture. Thus, for this sector we only have capital and labor as inputs.

We also use some further assumptions, as we need to address a cross-section—unlike Solow's work, which was in the time dimension. First, notice that, in the equations above, we use green color to denote functions that we assume to be common across countries; we use blue to denote exogenous variables that differ across countries. Thus, we assume the functions f and g are not only the same across space but also across time. Second, one could imagine a richer setup where f and g had input-specific technology factors that differ across countries. As we will show, we obtain a rather good fit to the data without such factors; also, one may think that such factors are endogenous—say, through some form of directed technical change—and that f is simply the reduced form, with the TFP levels capturing an overall difference in productivity. Third, we see that while outputs and all inputs except one vary over time, land does not.

Fourth and finally, rather than estimate country-specific A terms, we "project" them onto a development index, which we take to be real GDP net natural resource rents per worker, y. We do this since we are interested in precisely how the TFPs vary in this systematic dimension and our method gets straight at this.

The same kinds of assumptions will be used later on in the present paper.

[5] The reallocation of production inputs across sectors is thus assumed to be efficient. Between the agricultural sector and the rest of the economy there can certainly be impediments to mobility, perhaps thought of as "wedges". We consider such wedges in our analysis below. However, in constructing our TFP measure for the agricultural sector (vs. the rest), the input allocation across sectors does not play a role.

To explain our procedure, let us for illustration consider the non-agricultural sector, where $y_n = A_n g(k_n, h_n)$. Taking logarithms, we obtain

$$\log y_n = \log A_n + \log g(k_n, h_n).$$

Taking total derivatives, we obtain

$$d \log y_n = d \log A_n + d \log g(k_n, h_n).$$

Notice that $d \log g(k_n, h_n) = \frac{1}{g(k_n,h_n)} \left(g_k(k_n, h_n) dk_n + g_h(k_n, h_n) dh_n \right)$. Thus, by multiplying and dividing by A_n, we see that $d \log g(k_n, h_n) = s_n^k d \log k_n + s_n^h d \log h_n$, where s_n^k and s_n^h are the cost shares of capital and labor, respectively; this comes from identifying input prices with marginal products. Hence we have

$$d \log y_n = d \log A_n + s_n^k d \log k_n + s_n^h d \log h_n.$$

The total differential here is abstract; it merely refers to a small change in the variable in question. Solow considered changes over time; e.g., he looked at $d \log k$ as the change *over time* in the capital input. We instead consider changes over our development—output per worker—dimension. With y denoting our development index, the change in log TFP at development level y is thus given by

$$d \log(A_n(y)) = d \log(y_n(y)) - s_n^k(y) d \log(k_n(y)) - s_n^h(y) d \log(h_n(y)).$$

In this equation, all variables on the right-hand side are *projections* of the data onto y, i.e., they are smooth fitted functions of the raw data. The errors for these fitted lines are small; there are, as we have seen in the above graphs, market development patterns in input quantities and shares. For concreteness, the shares plotted in Figure 4 above are the projections also used here. Finally, defining U.S. TFP to be 1, we can integrate in the y dimension to obtain the entire function $A_n(y)$.

Similarly, for the agricultural sector we obtain

$$d \log(A_a(y)) = d \log(y_a(y)) - s_a^k(y) d \log(k_a(y)) - s_a^h(y) d \log(h_a(y)) - s_a^x(y) d \log(x_a(y)) - s_a^l(y) d \log(l(y)),$$

with the same kind of notation as above.

Finally, for the production of intermediates and investment, we assume

$$y_{x,t} = A_x A_n g(k_{x,t}, h_{x,t}) \quad y_{k,t} = A_k A_n g(k_{k,t}, h_{k,t})$$

Given this structure, it follows that if markets allocate production factors efficiently—which they do if there is perfect competition—then A_x and A_k are given directly by (the inverses of) the prices, relative to consumption, of intermediates and investment, respectively.[6] The idea that relative prices reflect relative marginal costs, and therefore the inverses of relative TFPs, is intuitive and plausible. The assumption that these three different sectors—producing non-agricultural consumption, intermediates, and investment goods—have literally identical isoquants is, of course, a strong assumption. However, a richer structure cannot be implemented easily as it would require data on input use by sector, which is not available for the broad cross-section of countries in focus here. We would not expect our sector-by-sector TFP measurements

[6] This can be shown by manipulation of the first-order conditions for input choices.

to change greatly, however, had we been able to consider the more general case. In any case, this assumption is not crucial for the main results in our paper.

Having explained how we obtain all our TFP measures, we plot the results in Figure 7. The figure

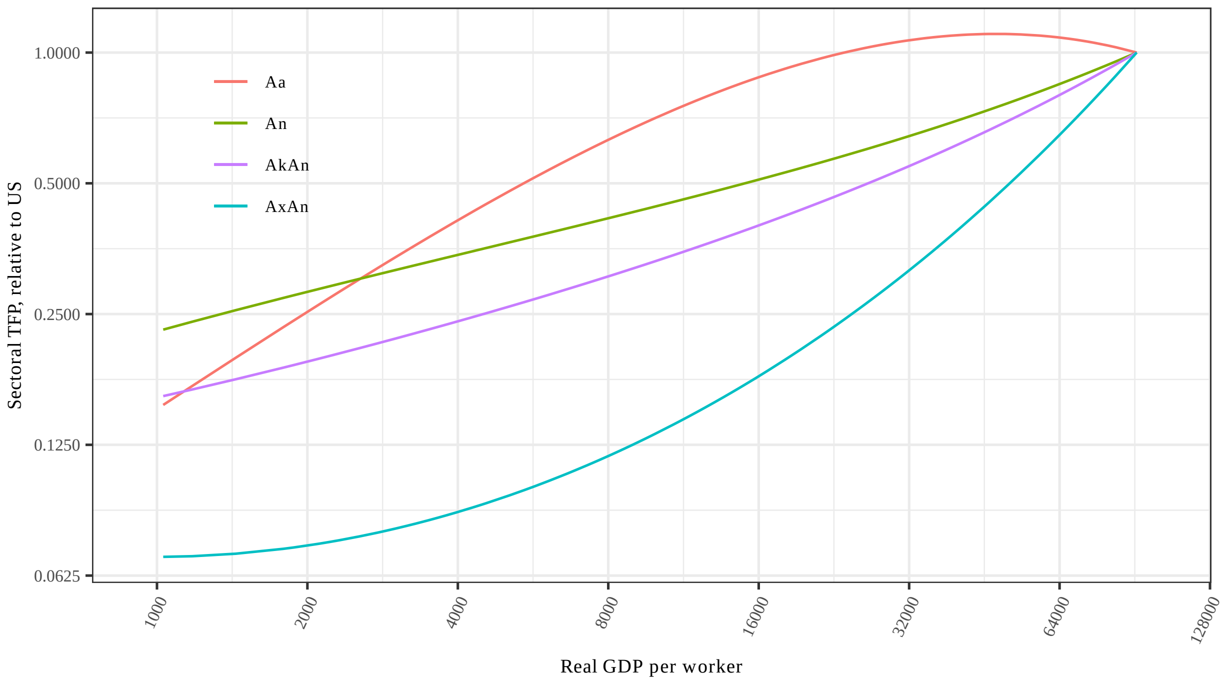

Figure 7: TFPs by sector

reveals one of our key findings: agricultural TFP in the poorest countries is a little over one-eighth of that in the U.S., it sharply increases with GDP per worker initially, and it reaches the U.S.'s level when GDP per worker is merely half of that in the U.S. (at Mexico's level, roughly); from that point on it is flat. Put differently, agricultural TFP is increasing in the level of development, but concave, to the point that the positive relationship dies out well before the development frontier. In contrast, our TFP measures for the other three sectors, which are also increasing, are convex: the largest country gaps—in percentage terms—are found among the most developed countries. These features are most striking for the production of intermediates, where the sizes of the gaps are also much larger: the TFP in the poorest countries is just 1/16th of that at the frontier. The TFP gap in the production of investment goods is at the level of that in agriculture if comparing extremes, but is otherwise significantly larger, as agricultural TFP gaps close more quickly in the development dimension. The non-agricultural consumption TFP gap is smaller at the extremes but, again, is larger than the agricultural TFP gap over most of the development range. In sum, while the agricultural TFP gaps are large in the development dimension, they are not as large as for other sectors. In other words, it appears that a "Green Revolution" has already occurred to an important degree, and that what we are waiting for is a "Non-green Revolution".

The picture looks quite different if, in contrast, one looks at the labor productivity gaps along the development dimension. Figure 8 illustrates, plotting both labor productivity and TFP for agriculture. We see that labor productivity is slightly convex, i.e., close to (log-)linear, and that its gaps are much larger. Consequently, by looking at labor productivities, it appears that a Green Revolution would be of a first-order importance. Our observation here is simply that most of the labor productivity gap in agriculture is derived from capital- and intermediate-goods intensification—not from its own TFP. Thus, while improving agricultural labor productivity is of course of first-order importance, the key to accomplishing this is perhaps to be found elsewhere: by a revolution in other sectors of the economy.[7]

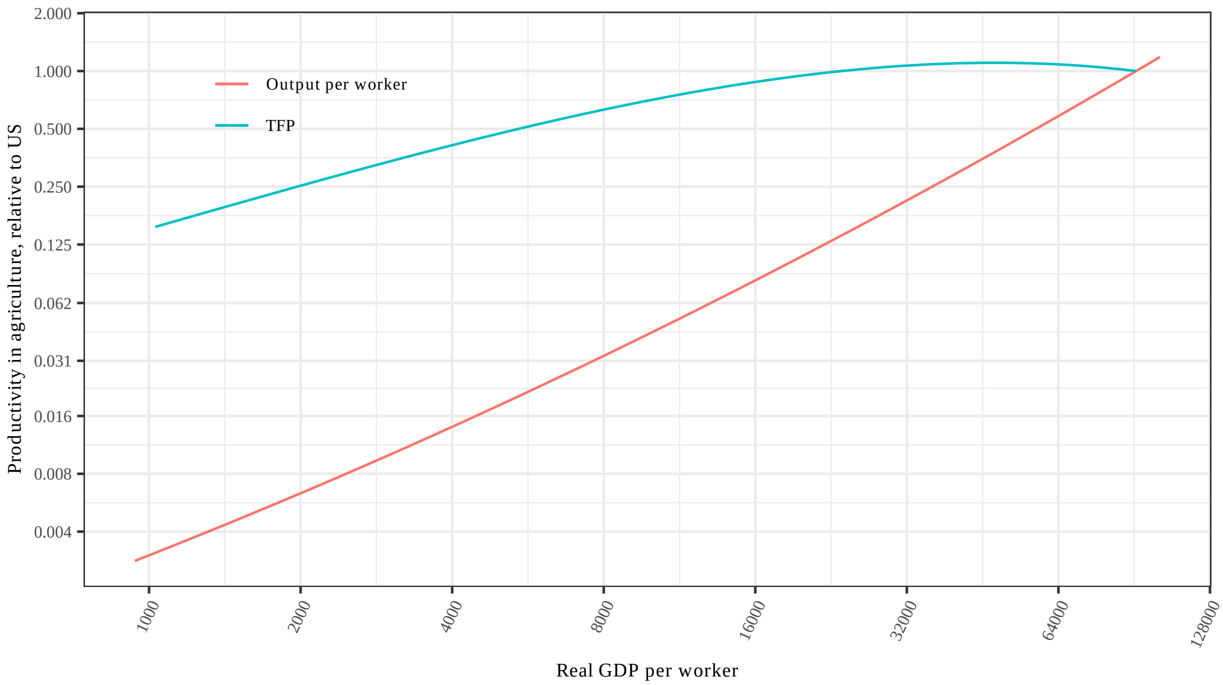

Figure 8: Labor productivity vs. TFP in agriculture

4 Theory, II: the agricultural production function

Having derived TFP series by sector, we now take one further step and make more assumptions on the shape of the production isoquants. This step allows us to measure macroeconomic input elasticities for each sector: a set of elasticities for agriculture and one for the non-agricultural sectors.[8] It thus allows us to comment on how the macroeconomic development literature has proceeded: it has used specific functional forms.

The additional assumptions we make are stronger—they involve nested CES functions (thus displaying constant elasticities within each nest)—but we will see that the fit is still remarkably good over the entire development spectrum. To us, this was a major eye-opener.

We focus on the agricultural sector, which is more non-trivial as it contains four inputs; the non-agricultural sector is straightforward and, as we shall see below, results in a CES function that is rather close to Cobb-Douglas, i.e., a function with unitary input elasticity. Thus, for agriculture, we use the following

[7] An immediate reaction here might be "why not import these goods?" We discuss international trade briefly below, and it is certainly an important topic. There is international trade in the necessary inputs but its importance seems limited; whether this is because there are fundamental costs associated with the transfer of goods across borders or institutional constraints is less clear.

[8] Recall that we assume identical isoquants between the different non-agricultural sectors.

nested structure:

$$f(h,k,x,l) = A\left[\left(\left(h^{\frac{\sigma_1-1}{\sigma_1}} + \omega_1 k^{\frac{\sigma_1-1}{\sigma_1}}\right)^{\frac{\sigma_1}{\sigma_1-1}\frac{\sigma_2-1}{\sigma_2}} + \omega_2 l^{\frac{\sigma_2-1}{\sigma_2}}\right)^{\frac{\sigma_2}{\sigma_2-1}\frac{\sigma_3-1}{\sigma_3}} + \omega_3 x^{\frac{\sigma_3-1}{\sigma_3}}\right]^{\frac{\sigma_3}{\sigma_3-1}} \tag{3}$$

Here, labor and capital are nested together with an elasticity parameter σ_1; this nest and land are further nested together with an elasticity parameter σ_2; the resulting nest, lastly, is nested with intermediates with an elasticity parameter σ_3. Clearly, this is only one out of 12 possible nesting structures. We restrict attention to nestings where capital and labor appear together, since the macroeconomic development literature most often assume such a structure (many papers assume Cobb-Douglas) and it does not appear restrictive.[9] In addition to the structure in equation (3), which we term "nesting 1", we display results on two other possible nestings given this initial restriction: "nesting 2", which switches the order between x and l above; and "nesting 3", which first nests x and l together, and then nests the h-k and x-l nests together. We will estimate the parameters for each of these three nests separately and compare their implications for pairwise elasticities of substitution between inputs.

Our estimation of the parameters—$(\sigma_1, \sigma_2, \sigma_3, \omega_1, \omega_2, \omega_3)$ in the case of equation (3)—for each nest uses the 2005 cross-section and is straightforward, since we are able to manipulate the first-order conditions into exact log-linear forms expressing ratios of input quantities as functions of ratios of input prices. More precisely, we first estimate, using OLS, a log-linear equation for the case of capital and labor letting us read off σ_1 from the slope coefficient and ω_1 from the intercept. We can then construct the h-k nest given these parameters, and given the CES structure we can also construct an exact price index for this nest in analytical form. We then repeat the procedure for the nest consisting of h-k and l, where we now have prices for both objects, and again obtain a log-linear form from simple manipulation of the first-order conditions. This step is repeated once more and we thus obtain all sought parameters. Nesting 3 is estimated slightly differently: the parameters involving the h-k nest are estimated separately from the parameters of the x-l nest, after which a CES is estimated with these two nests as inputs.[10]

Nesting 1 delivers $\sigma_1 = 1.78$. Thus, the estimate indicates that capital and labor are significantly more substitutable than under a Cobb-Douglas formulation. As a point of comparison, Herrendorf et al. (2015) estimate $\sigma_1 = 1.58$ using U.S. time-series data. We obtain $\sigma_2 = 1.34$: land's elasticity of substitution with the h-k nest is significantly closer to unity. Finally, we estimate $\sigma_3 = 1.78$, thus also indicating a fairly high elasticity of substitution. These values imply the fit shown in Figure 9.

In Figure 10 below we show the implications for the pairwise substitution elasticities of our three nestings. Given that we have more than two inputs, there is more than one way to estimate Hicksian elasticities pairwise; we use Morishima elasticities, which report slopes along isoquants for the two inputs in question keeping the other inputs fixed.[11] Some of the pairwise elasticities in the graphs are flat by construction: for

[9] Note also that some of the nestings do not allow to write the production function as a separable function of intermediates x and all the other factors, implying that the corresponding value-added production function is not well defined even under a perfect competition assumption (see Sato (1976)).

[10] More sophisticated procedures can be followed; in particular it is possible to estimate all parameters at once using GMM. A case of particular interest is that where two σs are estimated to be close to each other. If, say, in nesting 1, σ_1 and σ_2 come out to be close, one can impose the restriction that they are equal and estimate a single σ with one equation and fixed effects using all the data.

[11] The results from the Allen-Uzawa elasticity measure are reported in the Appendix XYZ. This measure does not keep the other inputs fixed but allows them to change optimally, while still remaining on the same isoquant for the inputs in question.

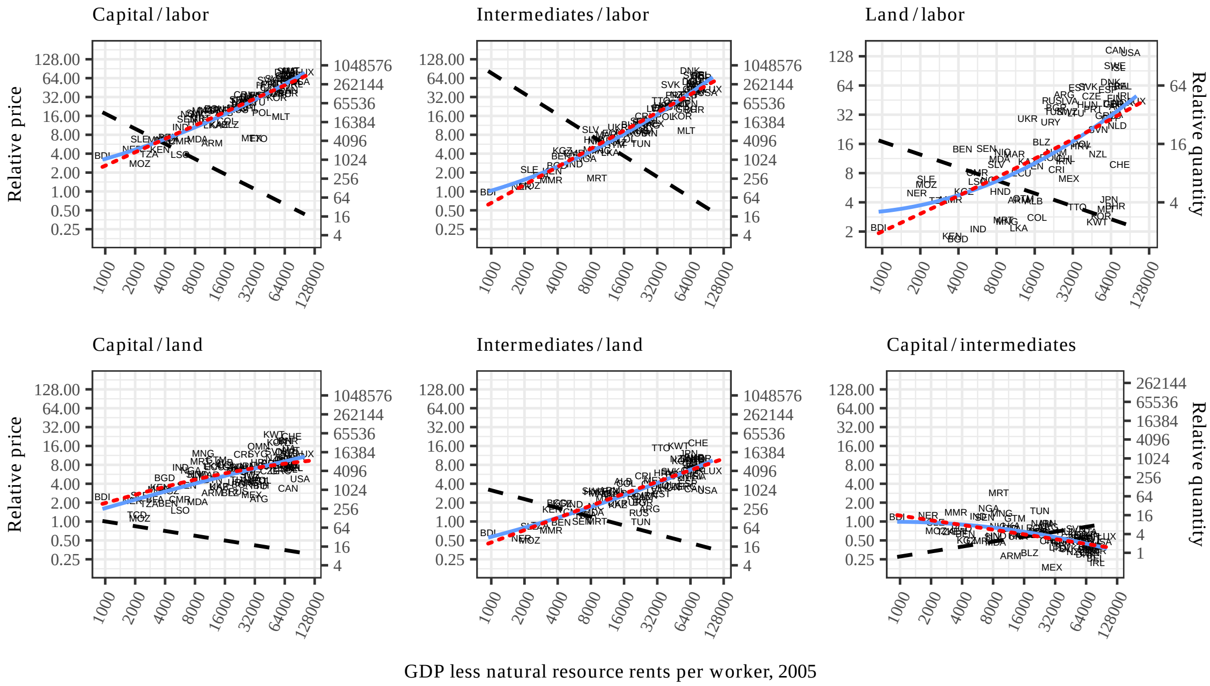

Figure 9: The fit of the nested CES

k-h in all nestings and for x-l in nesting 3. The others are nontrivial. It is quite remarkable how the pairwise elasticities do not appear to depend much at all on the level of development, going from the very poorest to the very richest country. The largest dependence is noted for the x-l elasticity implied by nesting 2 (the blue line): it moves between 1.7 and 1.25.

We also contrast our findings here with a formulation that would impose a Cobb-Douglas elasticity between capital and labor. Such an assumption can be entertained and still be consistent with the data if there is a wedge between wages in the different sectors. In particular, assume the wedge τ to be defined by

$$\tau \frac{1-\alpha}{\alpha} = \left(\frac{wh_a}{rk_a}\right).$$

On the right-hand side, the prices w and r are sector-independent. Thus, there would be no need for a wedge ($\tau = 1$) if we saw that the ratio of the capital and labor shares in agriculture were constant along the development dimension. We know, however, from Figure 4 above that they are not; in poor countries, the labor share in agriculture is much higher than in rich countries (0.7 vs. 0.2 in gross output), with capital moving in the opposite direction (0.05 vs. 0.1). Figure 11 displays how the implied wedge depends on the level of development, where we've defined the wedge to be 1 for the richest country. The wedge is higher for poor countries, reflecting a "higher wage relative to rental rate of capital" in their agricultural sectors relative to non-agriculture. This interpretation of the data has certainly been entertained in the literature and clearly has merit; the idea is that something must prevent movement from agriculture to more "modern" sectors, and various moving costs could be the culprit. However, the sheer size of the wedge we back out, in our assessment, too large to be plausible: it is almost a factor 10 for the poorest countries (and moves smoothly toward 1 as we consider higher and higher levels of development). Note also that a similarly implausible

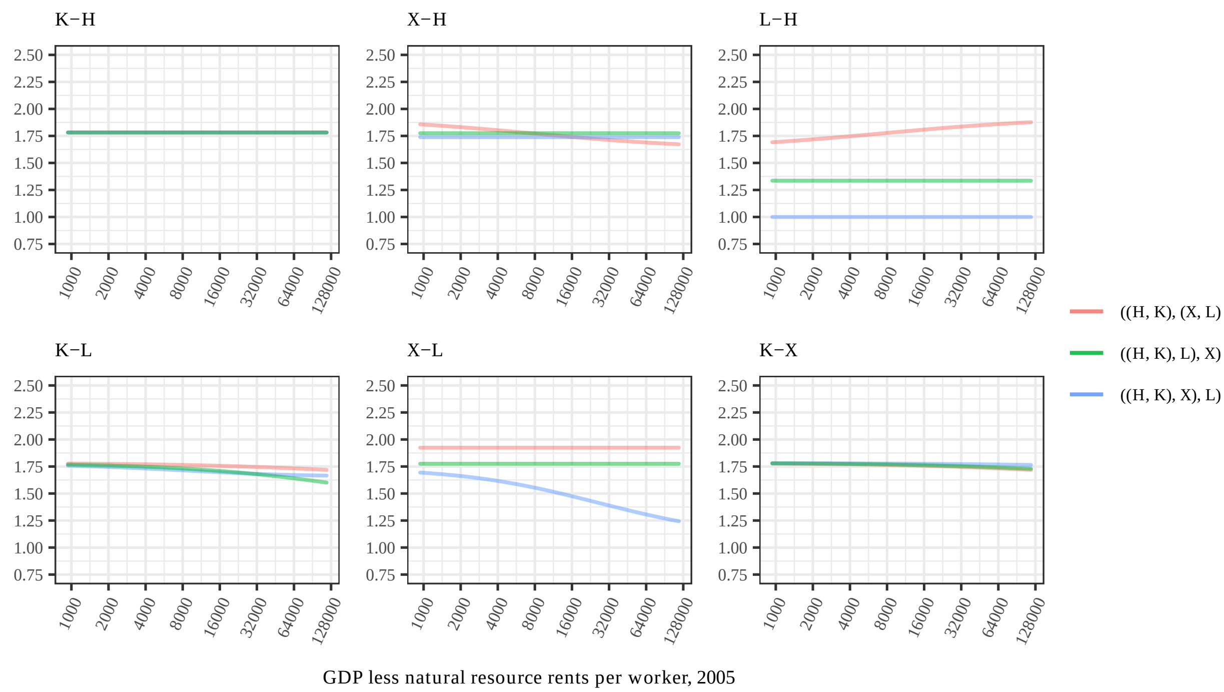

Figure 10: Pairwise input elasticities (Morishima)

wedge would result from the same type of exercise in the 1950s in the U.S. time series. Thus, we contend that while such a wedge may be important to consider, another element is needed in order to make sense of the input price and quantity data that we observe. Our aggregate nested CES does very well at that. It could be complemented with a wedge that declines as countries develop, but in our benchmark we maintain a zero wedge.

5 Theory, III: the demand side

The above sections conclude the discussion of production. With an eye toward the general-equilibrium theory in Section 6 below, which will allow us to run a number of counterfactual exercises, we now discuss consumer preferences. This discussion also has some separate insights about which preference structures can allow us to account for the facts on how budget shares change systematically in the development dimension.

As we have seen in Figure 5, the share of the total budget spent on agricultural goods is significantly higher in the poorest countries (around 50%) than in the richest countries (near 0%), so it is natural to entertain non-homothetic preferences. As in Boppart (2014), we use a price-independent generalized linearity (PIGL) specification, which is described through the following indirect utility function:

$$\max_{(c_a,c_n):p_a c_a+p_n c_n=E} u(c_a,c_n) \equiv v(p_a,p_n,E) = \frac{1}{1-\vartheta}\left(\frac{E}{p_n}\right)^{1-\vartheta} - \frac{\nu}{1-\eta}\left(\frac{p_a}{p_n}\right)^{1-\eta} - \frac{1}{1-\vartheta} + \frac{\nu}{1-\eta}. \tag{4}$$

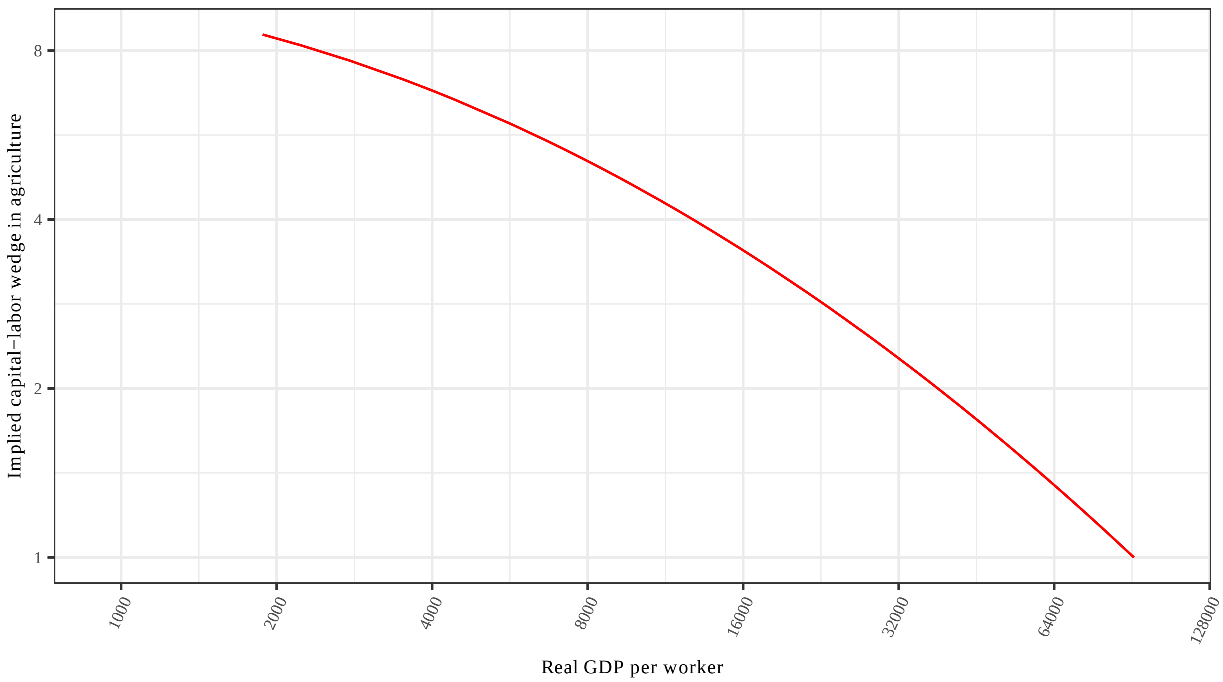

Figure 11: The "agricultural wage subsidy" implied by Cobb-Douglas production

The implied demand for the agricultural good is

$$c_a = \nu \left(\frac{E}{p_n}\right)^{\vartheta} \left(\frac{p_a}{p_n}\right)^{-\eta} = \nu e^{\vartheta} p_a^{-\eta}, \tag{5}$$

where ϑ is the (constant) expenditure elasticity of food, η is the asymptotic elasticity of substitution, ν is a share parameter, and we have used the normalization $p_n = 1$.[12] Likewise, the demand for the non-agricultural good is then

$$c_n = e - \nu e^{\vartheta} p_a^{1-\eta}. \tag{6}$$

We calibrate ϑ and η to the cross-sectional variation in agricultural consumption shares. Formally, by expressing aggregate consumption e, non-agricultural prices p_n and agricultural prices p_a as functions of GDP per worker, we construct predicted agricultural consumption shares as functions of GDP per worker. We then select ϑ and η to minimize the L^2-distance to the function that maps GDP per worker to observed agricultural consumption shares, setting the share parameter ν to ensure the curves intersect at the US income level.

Figure 12 shows the relative loss compared to the optimum as a function of ϑ and η, where a low ϑ means a strong income effect on the expenditure share and a low η means a strong substitution effect on the expenditure share. Reflecting the challenge of separating income and substitution effects in the cross-section, the figure has a ridge of equally good parameter values going in the southeast direction of higher

[12] The demand functions are straightforwardly derived using Roy's identity. By setting $\vartheta = 0$, the marginal propensity to consume agricultural goods is zero and the period utility function coincides with the case in Gollin, Parente, and Rogerson (2002, 2007) as a special case.

substitutability effects and weaker income effects. Since all maximizing ϑ are relatively small compared to micro estimates, we select the strongest complementarity and the weakest income effect in the lower-right corner of the ridge, yielding $\vartheta = 0.35$ and $\eta = 0$.

Figure 13 shows that the predicted consumption shares fit the data well across the GDP distribution. As should be expected given our calibration method, we match the overall fall in expenditure share closely, with the agricultural expenditure share falling from around 40% to around 2%. In addition, we see that the PIGL structure also matches the shape of the fall well, with a concave shape implying an accelerating fall in the log expenditure share.[13]

In terms of external validation, our estimate $\vartheta = 0.35$ implies an income elasticity of food of 0.35, which is lower than the average of 0.61 found in microeconomic estimates (Colen et al., 2018). We find this macro-micro friction interesting, because the low estimated income elasticity comes quite directly from the data. By selecting $\eta = 0$, we have maximized the scope for substitution effects by making preferences asymptotically Leontieff, and we still need a low income elasticity of food to jointly rationalize observed agricultural consumption with observed aggregate consumption. Looking ahead, we think one promising hypothesis to reconcile macro and macro is that people with higher income buy food with more value added outside of agriculture. This would imply that ultimate demand for agricultural output is less income sensitive than food demand.

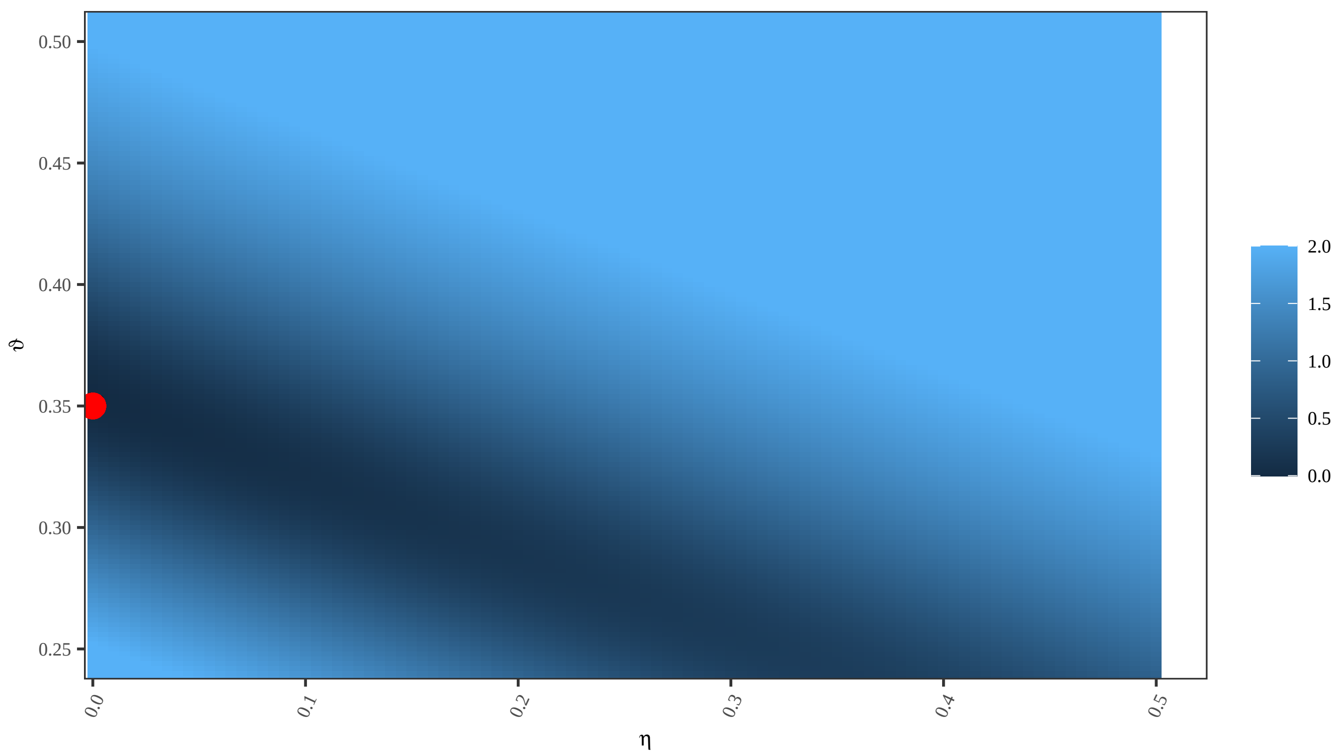

Figure 12: Log loss relative to optimum

[13]This feature is not an automatic consequence of our calibration method. For example, Stone-Geary preferences hit the overall fall well if we use the same calibration procedure, but we obtain a convex function in which the fall in the log expenditure share is decelerating with log GDP per worker.

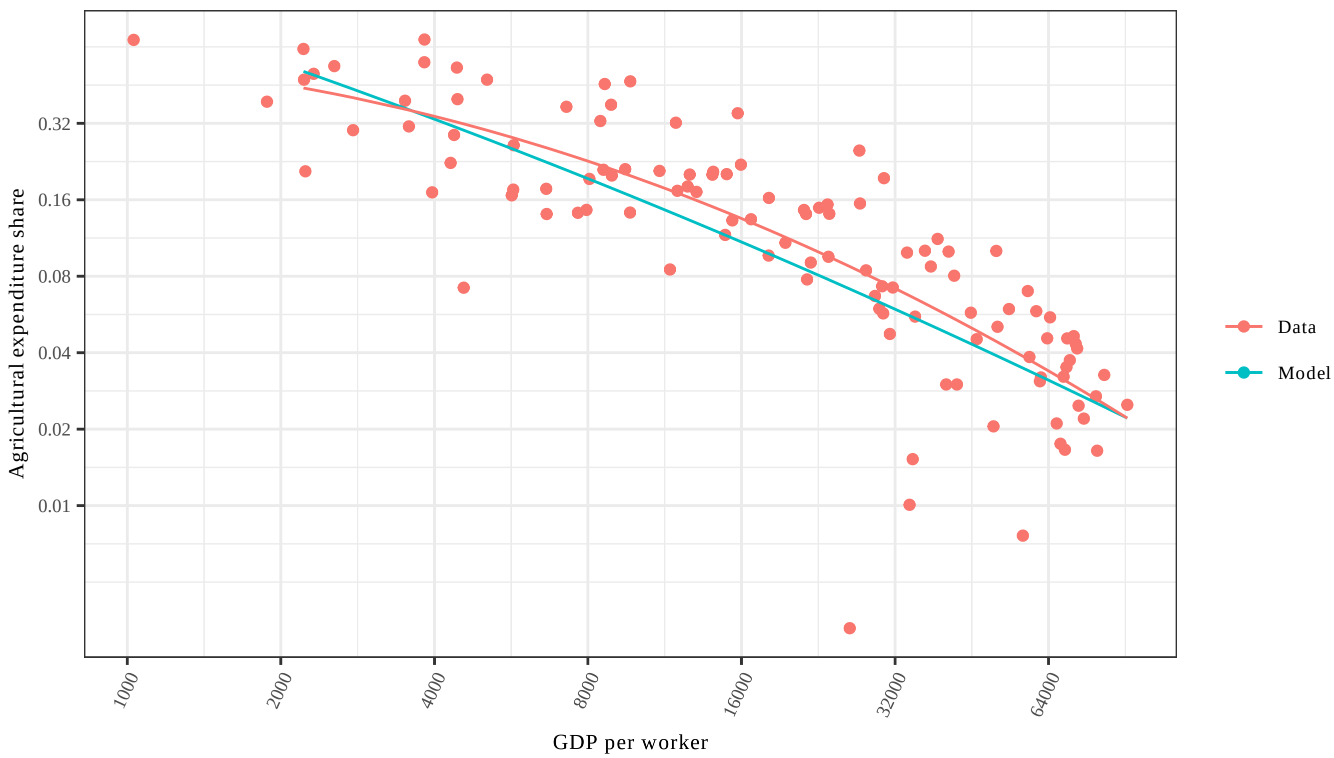

Figure 13: The share of agriculture in aggregate consumption, 2005

6 Theory, IV: general equilibrium

Next, we incorporate our findings on production functions and preferences in a general-equilibrium analysis. The key goal is to perform counterfactual exercises, such as assessing the relevant role of different exogenous inputs—such as the TFP levels in different sectors, one by one or in combinations—for output per capita and welfare. With our general-equilibrium theory we can also consider the role of wedges (such as the "subsidy" to working in agriculture).

We consider a closed economy with a representative household. Inequality is important to consider but is a topic worthy of a separate paper. Studying trade could potentially be important and we discuss trade in our concluding section. Finally, in this paper we limit attention to economies without growth, interpreting our data as representing steady states at different levels of development. Incorporating growth and interpreting the data from the perspective of ongoing development is more challenging as our preference formulation—due to its nonhomotheticity—is not consistent with exact balanced growth in the presence of capital intensity differences across sectors. Our as-yet unverified conjecture is that, since growth is a rather slow process, our quantitative results would not change significantly if we included ongoing growth.

6.1 General setup and planner's problem

In this subsection we keep the functional forms general regarding technologies in agriculture and non-agriculture and preferences over the two consumption goods; as before, we highlight these functions in green. Given that we are addressing a cross-section of countries, we will allow a number of country-specific parameters; they will, also as before, be marked in blue. The model is dynamic. Extensions involving frictions in worker mobility across sectors, ongoing technical change and growth, and international trade are then discussed in our concluding section.

The economy is populated by a representative household with preferences given by

$$\sum_{t=0}^{\infty} \beta^t u(c_{a,t}, c_{n,t}), \tag{7}$$

where $u(c_{a,t}, c_{n,t})$ denotes the period utility function defined over consumption of agricultural output c_a and non-agricultural output c_n. The parameter $\beta < 1$ captures the discount factor. The economy is endowed with H efficiency units of labor and L units of land. At any point in time, capital k_t is given from past investment decisions. Non-agricultural consumption goods are produced according to $A_n g(k_{n,t}, h_{n,t})$, where k_n and h_n denote capital and labor used in the non-agricultural sector and g is a constant-returns-to-scale function with standard regularity properties. In this paper the key focus is on how agricultural production differs from non-agricultural production. Hence we assume the same isoquants for investment as for non-agricultural consumption goods: we assume that one unit of the latter can be transformed into A_k units of the former. In agriculture, moreover, intermediates x_t are used, and we similarly assume that one unit of non-agricultural consumption can be transformed into A_x units of intermediate goods. Thus, our resource constraint for non-agricultural production reads (assuming a standard constant capital depreciation rate δ)

$$c_{n,t} + \frac{x_t}{A_x} + \frac{k_{t+1} - (1-\delta)k_t}{A_k} = A_n g(k_{n,t}, h_{n,t}). \tag{8}$$

Agricultural production takes place according to

$$c_{a,t} = A_a f(k_{a,t}, h_{a,t}, x_t, L), \tag{9}$$

where f is a constant-returns-to-scale function, also with standard properties.

In this economy a planner would choose $\{c_{a,t}, c_{n,t}, x_t, k_{t+1}, k_{a,t}, k_{n,t}, h_{a,t}, h_{n,t}\}_{t=0}^{\infty}$ to maximize (7), subject to the constraints (8), (9),

$$k_{a,t} + k_{n,t} = k_t \quad \text{and} \quad h_{a,t} + h_{n,t} = H \tag{10}$$

holding for all t.

6.2 Decentralized equilibrium

Prices are an important part of our analysis as we will use data on prices across countries. The decentralized equilibrium also allows us to entertain some factor-specific wedges τ_k (to capital accumulation) and $\tau_{h,n}$ (to labor in the non-agricultural sector). In this section we formally define the decentralized equilibrium.

We formulate a perfectly competitive equilibrium with some distortions. We choose the non-agricultural consumption good as numéraire. As the investment good and intermediate input are perfect substitutes to non-agricultural consumption under perfect competition, their relative prices are given by $1/A_k$ and $1/A_x$, respectively. We thus set these relative prices directly equal to these values and do not treat them as part of the equilibrium definition. The representative consumer solves

$$\max_{\{c_{a,t},c_{n,t},h_{n,t},k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_{a,t}, c_{n,t}) \quad \text{s.t.} \tag{11}$$

$$c_{n,t} + p_{a,t}c_{a,t} + \frac{k_{t+1}}{A_k} = (1-\tau_{h,n})w_{n,t}h_{n,t} + w_{a,t}\,(H - h_{n,t}) + \left(\frac{1-\delta}{A_k} + (1-\tau_k)r_t\right)k_t + p_{l,t}L + T_t. \tag{12}$$

Note here that we allow distortionary taxes on labor in non-agriculture and on capital, all redistributed lump-sum via T_t to the representative agent.

At each t, non-agricultural firms maximize profits according to

$$\max_{k_n,h_n} A_n g(k_n, h_n) - r_t k_n - w_{n,t}h_n \tag{13}$$

whereas agricultural firms solve

$$\max_{k_a,h_a,x,l} p_{a,t}A_a f(k_a, h_a, x, l) - r_t k_a - w_{a,t}h_a - \frac{x}{A_x} - p_{l,t}l. \tag{14}$$

Notice that all prices above—p_a, r, w_a, w_n, and p_l—are relative prices and measured in units of the non-agricultural consumption good in the same period.

An equilibrium is thus formally a sequence of quantities and prices such that:

1. quantities solve the optimization problems of consumers (11), the non-agriculture firms (13), and the agricultural firms (14), where in addition L is the optimal land choice;
2. quantities are feasible in the non-agricultural sector (8) and the agricultural sector (9);
3. capital and labor markets clear, i.e., (10) is satisfied; and
4. the government budget balances each period, i.e.,

$$T_t = \tau_{h,n}w_{n,t}h_{n,t} + \tau_k r_t k_t. \tag{15}$$

Next, we specialize our functional form assumptions and we then characterize the equilibrium and show how to solve for steady state.

6.3 Parametric forms

For utility, it suffices to recall from Section 5 that the Marshallian demands are given by

$$c_n = \frac{E}{P_n} - \nu\left(\frac{E}{P_n}\right)^{\vartheta}\left(\frac{P_a}{P_n}\right)^{1-\eta} = e - \nu e^{\vartheta}p_a^{1-\eta} \tag{16}$$

$$c_a = \nu\left(\frac{E}{P_n}\right)^{\vartheta}\left(\frac{P_a}{P_n}\right)^{-\eta} = \nu e^{\vartheta}p_a^{-\eta}. \tag{17}$$

As already explained, ν regulates the level of demand for the agricultural good and η is the asymptotic elasticity of substitution between agricultural and non-agricultural consumption goods; $\eta < 1$ implies that the two consumption goods are gross complements. The parameter ϑ is the constant expenditure elasticity of demand for agricultural consumption.[14] Hence the preferences are in line with Engel's law as long as $\vartheta < 1$, whereas preferences become homothetic with $\vartheta = 1$.[15]

As the non-agricultural production function has constant returns to scale, for some purposes it is convenient to work with its unit cost function, which is defined and given in a closed form as follows:

$$q_n(r, w_n) \equiv \min_{(k_n,h_n): A_n g(k_n,h_n)\geq 1} rk_n + w_n h_n = \frac{1}{A_n}\left(\alpha^\sigma r^{1-\sigma} + (1-\alpha)^\sigma w_n^{1-\sigma}\right)^{\frac{1}{1-\sigma}}. \tag{18}$$

This expression follows from assuming a CES production function $g(k_n, h_n) = \left(\alpha k_n^{\frac{\sigma-1}{\sigma}} + (1-\alpha) h_n^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$ translating units into c_n given the Hicks-neutral technology multiplier A_n. The parameter σ is the elasticity of substitution; α controls the weight on capital relative to labor. To obtain the cost function for investment and intermediate goods one simply premultiplies the expression in (18) by $1/A_k$ and $1/A_x$, respectively.

Turning to the agricultural sector, since $f(k_a, h_a, x, l)$ is assumed to be a nested CES form, we can also solve explicitly for cost functions, defined as

$$q_a(r, w_a, p_x, p_l) \equiv \min_{(k_a,h_a,x,l): A_a f(k_a,h_a,x,l)\geq 1} rk_a + w_a h_a + p_x x + p_l l. \tag{19}$$

Given that there are four inputs and the functional form for the cost function depends on the specific nesting, we omit the formulas here to avoid cluttering.

6.4 Solving for equilibrium

When solving for equilibrium it is helpful to split the problem up into an intratemporal and an intertemporal part. In the intratemporal part, we only consider how, given a total amount of capital k and a total amount of expenditures on the two consumption goods e, the sectoral allocation of inputs and relative prices will be determined. Stated more formally, it consists of a mapping from (k, e) to a vector $\vec{v} \equiv (c_n, c_a, w_n, w_a, r, p_l, p_a, k_n, k_a, h_n, h_a, x)$. The intertemporal part then is about finding the equilibrium sequence $\{e_t, k_{t+1}\}_{t=0}^{\infty}$. We describe these parts separately.

6.4.1 Intratemporal part: sectoral allocation and relative prices

We now collect all the relevant equilibrium conditions that will determine the vector $\vec{v}$. Since $\vec{v}$ has 12 elements we need 12 mutually independent equations. First, the consumer's demand functions must be included: (16) and (17). Then, given our choice of the non-agricultural consumption good as the numéraire,

[14] This leads to a sustained income effect whereas other non-homothetic preferences imply a varying expenditure elasticity. For example, Stone-Geary preferences imply an expenditure elasticity that asymptotes to 1 but, more importantly, converges quite quickly toward this value as income increases; hence, in our application, income effects would only be active for the very poorest countries in the sample.

[15] With $\vartheta = \eta = 1$ the preferences nest Cobb-Douglas preferences, $u(c_a, c_n) = c_a^\nu c_n^{1-\nu}$, as a special case.

perfect competition in the non-agricultural sector delivers

$$r = \alpha k_n^{-\frac{1}{\sigma}} A_n \left(\alpha k_n^{\frac{\sigma-1}{\sigma}} + (1-\alpha) h_n^{\frac{\sigma-1}{\sigma}}\right)^{\frac{1}{\sigma-1}}, \tag{20}$$

$$w_n = (1-\alpha) h_n^{-\frac{1}{\sigma}} A_n \left(\alpha k_n^{\frac{\sigma-1}{\sigma}} + (1-\alpha) h_n^{\frac{\sigma-1}{\sigma}}\right)^{\frac{1}{\sigma-1}}. \tag{21}$$

Given that this good is the numéraire, perfect competition requires its marginal cost to be one, which these two equations imply:

$$1 = \frac{1}{A_n} \left(\alpha^\sigma r^{1-\sigma} + (1-\alpha)^\sigma w_n^{1-\sigma}\right)^{\frac{1}{1-\sigma}}. \tag{22}$$

For the agricultural output price we obtain, also from perfect competition, that

$$p_a = q_a(r, w_a, p_x, p_l), \tag{23}$$

where p_x can be set to $1/A_x$. The demands for inputs in the agricultural sector follow from Shephard's lemma:

$$k_a = \frac{\partial q_a(r, w_a, p_x, p_l)}{\partial r} c_a, \tag{24}$$

$$h_a = \frac{\partial q_a(r, w_a, p_x, p_l)}{\partial w_a} c_a, \tag{25}$$

$$x = \frac{\partial q_a(r, w_a, p_x, p_l)}{\partial p_x} c_a, \tag{26}$$

$$L = \frac{\partial q_a(r, w_a, p_x, p_l)}{\partial p_l} c_a, \tag{27}$$

where p_x should be evaluated at $1/A_x$ and (27) already uses land market clearing, i.e., $L = l$. Given that q_a is the unit cost function, we can reconstruct total agricultural output from it: $A_a f(k_a, h_a, x, L)$. Inserting the above equations into this expression yields c_a, i.e., the equations above imply that the agricultural resource constraint (9) is met.

For labor supply the condition

$$w_a = w_n(1 - \tau_{h,n}) \tag{28}$$

has to hold: so long as the representative agent works in both sectors, their net-of-tax wage rates have to be equal. Finally, capital and labor market clearing requires the two equations in (10).

Collecting conditions, we have (10), (16), (17), (20), (21), and (23)–(28), which amounts to 12 equations.

This system can be reduced by recursive substitutions to produce a system of two equations in two unknowns (r and p_l); the remaining elements of the vector $\vec{v}$ are then obtained directly. To see this, note that (22) and (28) allow us to express w_n and w_a as functions of r. Substituting these expressions into (23) implies p_a as a function of the endogenous r and p_l. Together with the given e, p_a can then be inserted into the demand function for agricultural goods, (17), to deliver c_a as a function of only the endogenous r and p_l. Inserting c_a and the expression for w_a from above into the market-clearing equation for land, (27), then delivers one equation in the unknowns p_l and r. The other equation can be derived as follows. First, insert c_a and w_a into the remaining input demand relations in agriculture, (24) and (25), to deliver k_a and h_a, which consequently also become functions of p_l and r. Substituting the agricultural inputs of capital and labor into (10) then gives the values for k_n and h_n. Thus, we have k_n/h_n (also as a function of p_l and r) and, inserted into (20), this delivers the second equation in our two unknowns.

Note that this solution procedure does not involve two variables, c_n and x, that were part of the vector $\vec{v}$. These two additional variables follow immediately from (16) and (26), neither of which were used above.

6.4.2 Intertemporal decisions and the full equilibrium

The household maximizes $\sum_{t=0}^{\infty} \beta^t \left(e_t^{1-\vartheta}/(1-\vartheta) - \nu p_{a,t}^{\gamma}/\gamma - 1/(1-\vartheta) + \nu/\gamma \right)$, by choice of $\{e_t, k_{t+1}\}_{t=0}^{\infty}$ subject to the budget constraint, (12), where $c_{n,t} + p_{a,t} c_{a,t}$ on the left-hand side can be replaced by e_t. This delivers, as a first-order condition, the Euler equation

$$\left(\frac{e_{t+1}}{e_t} \right)^{1-\epsilon} = \beta \left(1 - \delta + A_k r_{t+1} (1 - \tau_k) \right). \tag{29}$$

This Euler equation (along with a transversality condition) is the key intertemporal condition, as it connects periods.

To see how the model can be solved for a transition path, consider a basic computational procedure: shooting. In the standard growth model, one would guess consumption in the first period, yielding a capital stock carried in to the second period along with an interest rate between the periods. Then the Euler equation is used to obtain consumption in the second period, and the process continues. If the so-obtained path for capital converges, a solution has been obtained; if not, the guess on initial consumption is updated. In the present model, one can proceed analogously, with a slight modification.

First, thus, select e_0, which together with k_0 delivers $\vec{v}_0$, the vector of sectoral allocations and relative prices in the first period. To obtain the capital stock for the next period, use the resource constraint for the non-agricultural consumption good: (8). This equation contains k_1 and all other variables are now given. The Euler equation, however, cannot directly be used to find e_1 in our case, since r_1 is a function of the sector-specific capital-labor ratio in period 1. However, we can simply add the Euler equation between period 0 and period 1 to the intratemporal system, where we now have 13 equations in 13 unknowns—having added the unknown e_1. At this point, the procedure is repeated, as in the standard model.

It is straightforward to show that, given the variables computed at this stage, the final equilibrium conditions are met: the budget constraints for the consumer and the government. T_t is simply defined from the government budget at t, and substituted into the consumer's budget we see, by use of the already computed equilibrium conditions, that this budget constraint holds with equality in all periods.

6.5 Steady-state equilibrium

We now solve for steady-state equilibrium. A steady state is defined as the equilibrium the economy converges to as the capital stock reaches its asymptotic level. We denote steady state variables by $\star$ superscripts.

Conceptually, the steady state is straightforward to solve for. Immediately, the Euler equation delivers the rental rate, $r^\star$. Together with the intratemporal 12-dimensional system, as well as the non-agricultural resource constraint (8), we can then solve for the 13 remaining unknowns: $e^\star$, $k^\star$, and the remaining elements of $\vec{v}^\star$. It is, however, straightforward to work this system down to one equation in one unknown: the land price. We now outline how this is accomplished.

The Euler equation directly gives us the steady-state rental rate as a function of the exogenous parameters:

$$r^\star = \frac{1/\beta - 1 + \delta}{A_k(1 - \tau_k)}. \tag{30}$$

Together with (22) and (28), this determines steady-state wages as

$$w_n^\star = \frac{A_n(1-\alpha)^{\frac{\sigma}{\sigma-1}}}{\left(1 - A_n{}^{\sigma-1}\alpha^\sigma (r^\star)^{1-\sigma}\right)^{\frac{1}{\sigma-1}}}, \tag{31}$$

$$w_a^\star = \frac{(1-\tau_{h,n})A_n(1-\alpha)^{\frac{\sigma}{\sigma-1}}}{\left(1 - A_n{}^{\sigma-1}\alpha^\sigma (r^\star)^{1-\sigma}\right)^{\frac{1}{\sigma-1}}}. \tag{32}$$

The first-order conditions from the non-agricultural firm's problem, (20) and (21), then allow us to solve for the capital-labor ratio in non-agriculture as

$$\frac{k_n^\star}{h_n^\star} = \left(\frac{\alpha w_n^\star}{(1-\alpha)r^\star}\right)^\sigma. \tag{33}$$

Since $p_x^\star$ is given by $1/A_x$ and we have $r^\star$ and $w_a^\star$, (23) determines a relationship between the agricultural price and the land price:

$$p_a^\star = q_a(r^\star, w_a^\star, 1/A_x, p_l^\star). \tag{34}$$

In the next step, we construct the consumption of agricultural goods and all the input levels in agriculture. This is achieved by first using the expression (27), which sets the demand for land equal to its supply: this pins down $c_a^\star$ as a function of $p_l^\star$. By inserting this expression for $c_a^\star$ into the rest of the agricultural input demand system, (24)–(26), we obtain $k_a^\star$, $h_a^\star$, and $x^\star$, again as functions only of the land price.

Having solved for the input levels from one sector we easily obtain those in the other, along with aggregate capital: we find $h_n^\star$, $k_n^\star$, and $k^\star$ from (10) and the use of (33). This also gives us total non-agricultural output, allowing $c_n^\star$ to be backed out from this sector's resource constraint, (8), with all variables still being a function only of $p_l^\star$. The penultimate step is to construct the steady-state expenditure level, as a function only of the land price, from $e^\star = c_n^\star + p_a^\star c_a^\star$. Finally, we insert $e^\star$ into the demand equation for agricultural goods, (17), which is our one equation in the unknown land price.

Notice that our recursive procedure allows closed-form expressions so that the only equation that needs to be solved numerically is the last equation, after which all remaining variables follow immediately. Also notice that, since we did not use any specifics of the agricultural production function, this method works no matter what this function looks like, so long as it has constant return to scale.

7 Calibration and counterfactuals

The calibration consists of the country-invariant parameters and of functions that map the development index y to country-specific parameter values. Virtually all parameters have already been assigned values. Sector-specific TFP differences were obtained in section 3; the shape of the production functions f and g were estimated in section 4, and the indirect utility function v was estimated in section 5.[16] To obtain TFP

[16]To remind the reader, we solved for TFP levels by sector and by level of development using Solow growth accounting in the cross-sectional (development) dimension, i.e., using a cross-section of countries at a point in time. The shapes of the

levels from TFP differences, we normalize TFPs $A_a(y_{us}), A_n(y_{us}), A_k(y_{us}), A_x(y_{us})$ at the US income level to fit observed price indices observed factor prices at y_{us}. For consistency with the earlier analysis, we also set the labor wedge to zero, $\tau_{n,h}(y) \equiv 0$, since this assumption was maintained when constructing the input prices in section 2.

The key remaining parameters are those regulating capital accumulation: β, δ and τ. We select these to make investment rates and rental rates consistent with observed capital stocks

$$\delta(y) = \frac{\text{Investment share of GDP}}{\text{Nominal capital-to-output ratio}}(y) \tag{35}$$

$$\frac{1}{1-\tau(y)}\left[\frac{1}{\beta} - 1 + \delta(y)\right] = \frac{\text{Capital compensation-to-GDP ratio}}{\text{Nominal capital-to-output ratio}}(y) \quad \tau(y_{us}) = 0. \tag{36}$$

The calibration in (35)-(36) ensures consistency with our previous analysis. In particular, (35) ensures that the investment share is jointly consistent with the consumption share used in the preference estimation in section 5 and the capital data used in section 2-4. In addition, (36) ensures that the model rental rate is consistent with observed capital compensation shares, which was the assumption maintained for constructing the rental rate in section 2 ($\tau_k(y_{us}) = 0$ is a normalization that pins down β).

The parameters that do not vary across countries are displayed in table 2. Figure 14 plots the value of the intertemporal parameters across countries. There are relatively small variation in the rental rates, approximately between 12% and 13.5% across countries. There is no need for large wedges τ_k ed to rationalize the data. Instead, the differences that do exist are accounted for by differences in the estimated depreciation rate, reflecting that poor countries have a relatively low investment rate relative to the size of their capital stock.

Figure 15 compares other central outcomes to the data, and shows that the fit is very good. In the appendix, we show that the good fit is a natural consequence of our calibration strategy and previous results. In particular, we show that if you start with a consistent set of price and quantity moments, you will hit the overall equilibrium provided that you individually hit the producer problem, the consumer problem, and the capital accumulation equation. This means that the good fit is ensured by the consistent measurement exercise in section 2, the good fits obtained in section 3-5, and the calibration of δ in (35) which reconciles investment and capital data.

Discussion of calibration. Our calibration strategy follows the philosophy that reasonable parameters when fitting is a better standard than good fit conditional on reasonable parameters. The good fit in our model follows from consistent moments – ensured by the consistent measurement exercise – and fitting the blocks, ensured by choosing parameters and functional forms which ensure a good fit. Thus, model fit is not the quality standard, but by comparing measured moments and parameters against external standards – something which we do continuously throughout sections 2-5.

We believe that using external parameters for validation rather than calibration offers advantages in terms of transparency. For example, consider δ, the depreciation rate of physical capital. In our calibration,

isoquants were chosen by restricting attention to nested CES functions, again fitted to the cross-section of countries. The key challenge here was to find the agricultural production function. The elasticity of substitution between capital and labor in the non-agricultural sector was estimated, using the same procedure as for agriculture, to be close to 1: $\sigma = 0.997$. The preference parameters were selected using cross-sectional data on total expenditure, relative prices of agriculture and non-agriculture, and the agricultural share of consumption.

it is selected by (35) to reconcile investment rates and capital output ratios, which implies values from 5% in poor countries to 8% from rich countries.

This depreciation rate is high relative to an external standard like the depreciation rates reported in the Penn World Table, which range from 3.5% to 4.5% from poor to rich countries. What would have happened if we hardcoded these external depreciation rates? In this case, the investment rate would be lower, giving us counterfactually high consumption expenditure, agricultural expenditure, prices of lands, and ratios of labor, capital, and intermediate inputs to land. Thus, this one calibration choice on depreciation percolates through the general equilibrium system via resource constraints and price effects.

We think that these general equilibrium effects makes it hard to interpret model outcomes when parameters are selected according to some external standard without an eye to equilibrium consistency. Instead, we think a better method to evaluate model fit is to assess the distance of calibrated parameters to these external standards. For example, in our case, we can note that our implied depreciation rates are high relative to the Penn World Table. This naturally leads to extensions – like introducing growth – that might reconcile these differences.

	Parameter	Value
Preferences:	β	0.9114
	ϑ	0.35
	η	0
	ν	3.2650
Non-agricultural production:	α	0.4748
	σ	0.9972
Agricultural production:	ω_1	0.0020
	ω_2	0.2058
	ω_3	0.0159
	σ_1	1.7822
	σ_2	1.3359
	σ_3	1.7755

Table 2: Calibrated model parameters

7.1 Counterfactual experiments

Using the calibration, we can assess the effect of increasing sectoral TFPs on the income of poor countries. We consider two experiments. In the first experiment, we increase the agricultural TFPs A_a of poor countries to the level of rich countries, leaving the non-agricultural TFPs constant. In the second experiment, we increase the three non-agricultural TFPs A_n, A_nA_k, A_nA_x of poor countries to the level of rich countries, leaving agricultural TFP constant.

Our main finding is that despite agriculture being a large sector in poor countries, increasing agricultural TFPs is much less potent in increasing aggregate income than increasing non-agricultural TFPs. The reason is that second-order effects make agricultural TFP increases partly self-defeating: the strong complementarity between agricultural and non-agricultural consumption means that the size of the agricultural sector falls rapidly as its TFP improves, meaning that subsequent TFP increases apply to a much smaller sector of the economy.

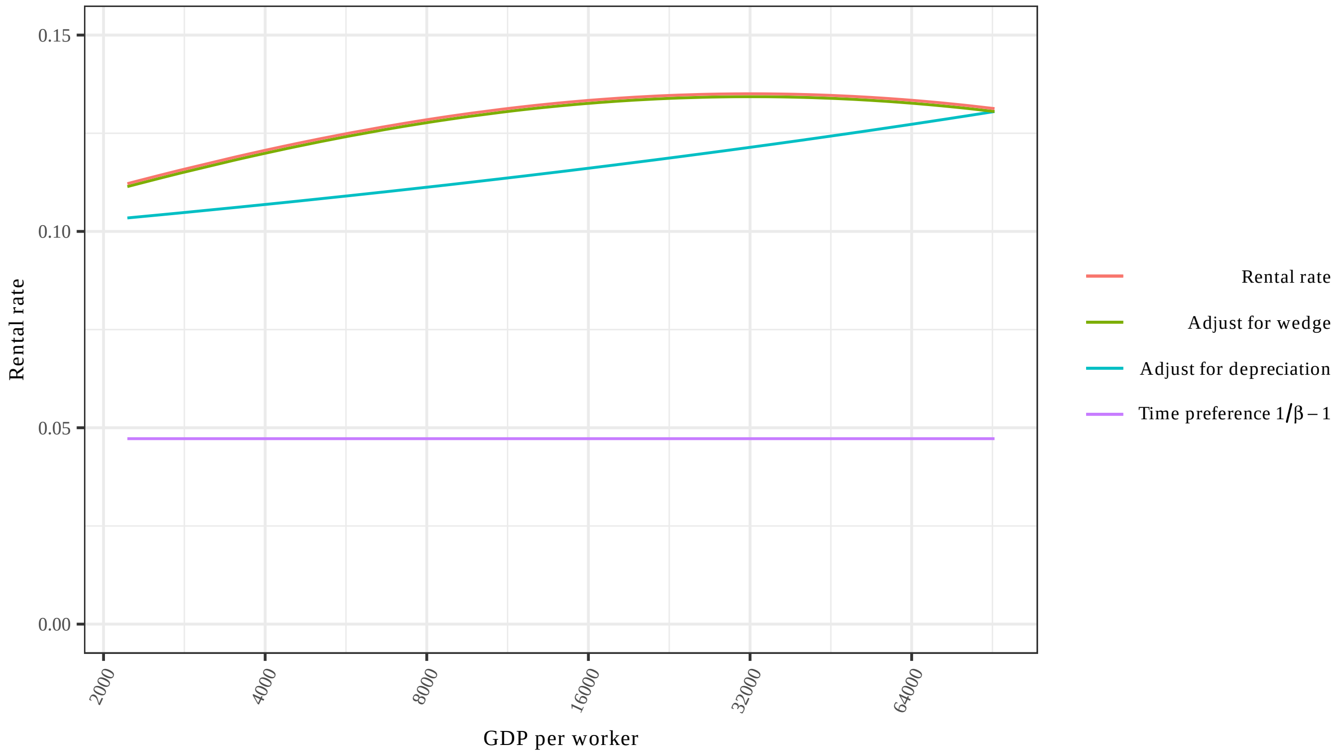

Figure 14: Sources of rental rate

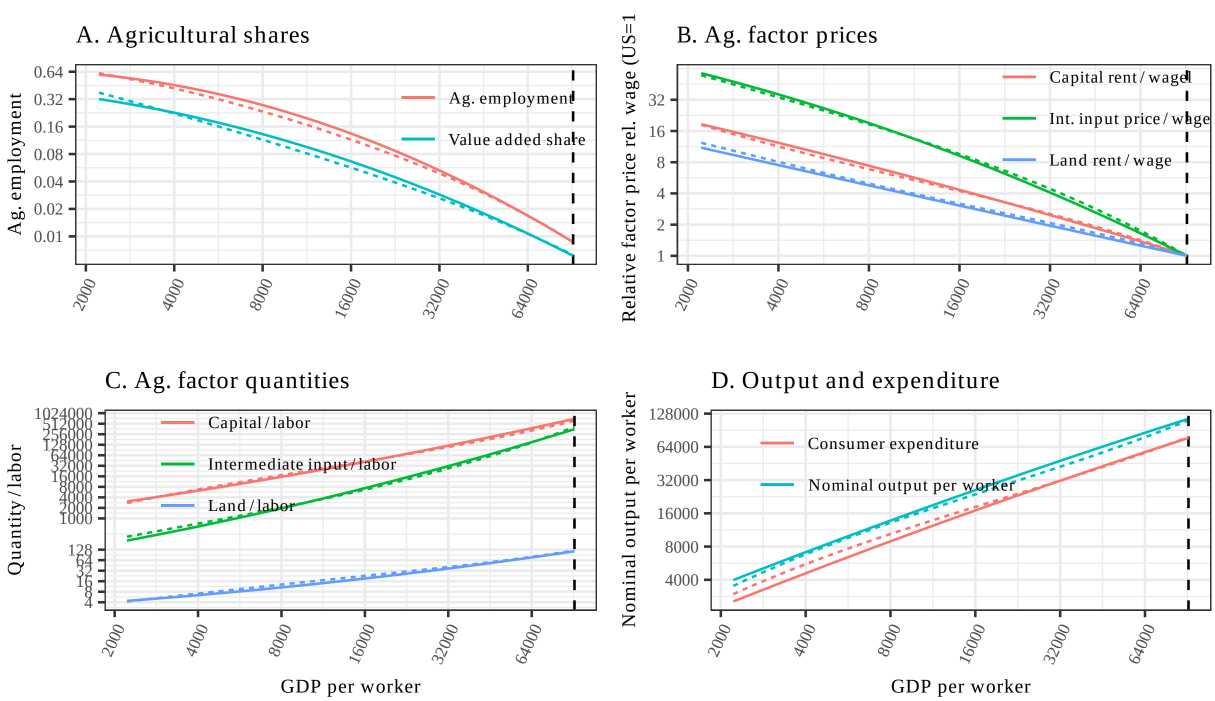

Figure 15: GE outcomes (dotted line model; solid line data)

Increasing agricultural TFP Figure 16 shows the effect of increasing agricultural TFP. The panel in the upper-left illustrates the experiment, with agricultural TFP increasing a factor of approximately 6 to match the levels of rich countries, with all other TFPs staying constant. All panels are indexed by GDP per worker on the horizontal axis, so that the results correspond to the effect of increasing agricultural TFP to the level of countries with that income level.

The upper-right graph show relative factor prices. Compared to labor, the relative prices of capital and intermediate inputs stay constant and the relative price of land falls. The constant capital and intermediate input prices reflect that they only depend on the constant non-agricultural TFPs. The falling land price reflects that structural change makes labor leave agriculture, which increases the land-to-labor ratio and requires a lower land price. The left middle panel shows how these factor price changes play out in agricultural quantities: capital-labor ratios stay flat, intermediate-labor ratios falling slightly, while land-labor ratios rise sharply together with labor productivity.

The middle-right graph shows that consumer expenditure in terms of non-agricultural good stays relatively flat, whereas the high agricultural TFP means that the relative price of agricultural goods shrinks by a factor of 8. The lower-right graph shows how this plays out in terms of structural change, with the gross output share of agriculture shrinks from 30% to less than yy%. This shrinkage primarily reflects the strong complementarity between agricultural and non-agricultural consumption, which mean that low agricultural prices translate into a low agricultural expenditure share. Also, because of stable factor intensities in agriculture, there is no compensating increase in investment and intermediate good production destined for agriculture.

The lower-right graph combine the effects into a response of GDP per workers. We see initial rise when a country with $2,000 GDP per worker increases its agricultural TFP to about that of a country with $8,000-$16,000, but the relationship is concave and even closing the full gap of agricultural TFP to the rich countries does not raise GDP per worker above $4,000. Prima facie, the small effect might be surprising given that agricultural sector is so large in poor countries, at least through the logic of Hulten (1979), which suggests that the response to sectoral productivity improvements should depend on a sector's size. The reason standard Hulten logic fails is second-order effects: as agricultural TFP increases, the size of the agricultural sector shrinks, and already when you reach the agricultural TFP of a country with $8,000 in GDP per worker, the gross output share has already fallen from 30% to 10%. Thus, while the initial increase in TFP is applied to a large sector, subsequent increases applies to a smaller and shrinking sector. This finding mirrors results from Baqaee and Farhi (2019), who show that second-order effects make a naive application of Hulten logic inappropriate in the face of large productivity shocks.

Increasing non-agricultural TFPs Figure 17 shows the effect of increasing non-agricultural TFPs. The panel in the upper-left illustrates the experiment, with agricultural TFP staying constant, while the non-agricultural TFPs increase with a factors of between 4 and 16. As found in section 3, the TFP differences are largest for intermediate input and investment goods, while the differences are smaller for non-agricultural consumption goods. As in figure 16, all panels are indexed by GDP per worker on the horizontal axis, which means that the results correspond to the effect of increasing non-agricultural TFPs to the level of countries with that income level.

Panel B shows factor prices. Mirroring the TFP patterns, there are dramatic falls in the rental rates of capital and the prices of intermediate inputs relative to labor. Due to structural change out of agriculture,

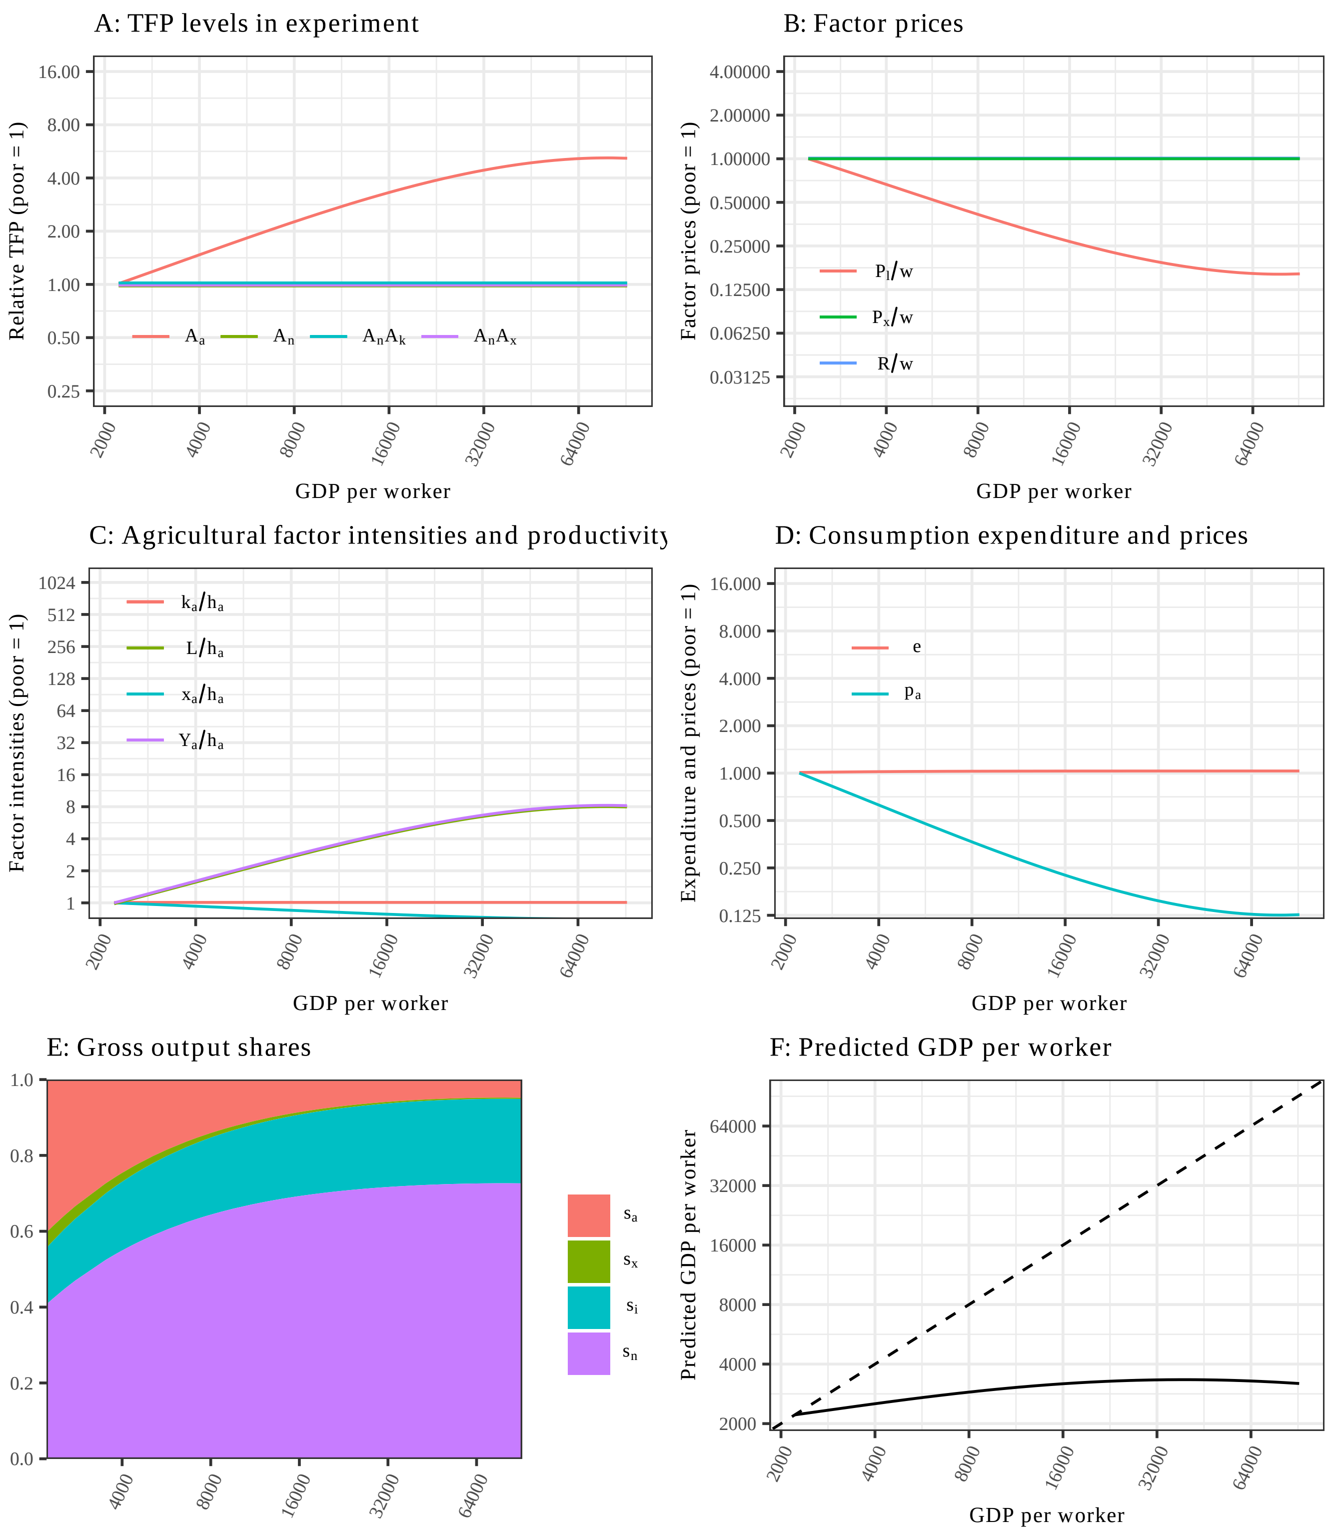

Figure 16: Effect of increasing agricultural TFP

land prices also fall relative to labor prices. Panel C shows how these factor price changes induce factor intensification in agriculture. Relative to labor, intermediate input and capital inputs grow between two and three orders of magnitude, and land per worker grows by a factor of 16. This factor intensification generates a more than 32-fold increase in labor productivity, despite agricultural TFP being constant.

Panel D and E show the consumer parameters and structural change in this experiment. From the consumer side, there are two counteracting effects on structural change: dramatically increased expenditures shrink agriculture through income effects. While increasing relative prices of agricultural goods expand agriculture due to complementarities, this effect is muted by agricultural factor intensification: relative prices only increase a factor of 2 despite relative TFPs changing a factor of 4, reflecting substitution into capital and intermediate inputs which have rapidly falling prices. The shrinkage of agriculture is further driven by expanding investment and intermediate goods production coming from factor expansion. The net effect is that the gross output share of the agricultural sector shrinks from 40% to about 10%.

Panel F shows the effect on GDP per worker. We see that there is an 11-fold increase bringing final GDP per worker to just below $32,000 dollars, closing 65% of the gap between rich and poor countries. Closing the non-agricultural TFPs is much more powerful than closing agricultural TFPs, despite the sectors being roughly the same size in poor countries and despite TFP differences being smaller in the non-agricultural consumption sector which is the dominant part of the non-agricultural economy. The key reason for the difference is that second-order effects here work in favor of big effects: as non-agricultural TFPs increase, income effects pull people out of agriculture and factor intensification simultaneously mitigate the adverse price response and move inputs into intermediate input and investment good productions.

8 An application: adaption to climate change

Our model of agriculture focuses on long-run mechanisms of adjustment in input factors and, as such, can be useful for looking at a broad variety of counterfactuals. One example is the application to climate change, a phenomenon that is expected to have highly heterogeneous effects across the globe. Available projections of "economic damages" from climate change point, in particular, toward areas of the world which are very poor and dominated by agricultural production. One such example is the Sahel region of northern Africa, an area projected to experience rapid population growth over the next several decades, roughly doubling by 2050. At the same time, this region is at severe risk of very large damages in per-capita terms (see, e.g., Krusell and Smith, 2022, who predict very high damages for large parts of the Sahel region).

Climate change brings about several features, but a combination of an increased prevalence of droughts due to higher temperature and lack of rainfall will make an increasingly large land area unsuitable for agriculture. What, then, would a reduction in land supply mean to an economy at a low level of development? Can an intensification of agriculture (i.e., increased use of other inputs) be a powerful adaptation tool in attempts to mitigate the adverse effects of land reduction? We now briefly address this question. We keep the analysis at an abstract level—we do not specifically look at the Sahel or any other region—but we do think that a fruitful way forward for research in this area would combine our setting with more country- or regions-specific detail.

In particular, our analysis stands in rather sharp contrast with available other studies. For example, when the Intergovernmental Panel on Climate Change (IPCC) projects food supply, they use various computational general equilibrium (CGE) models. One important such model is the IMAGE integrated assessment model

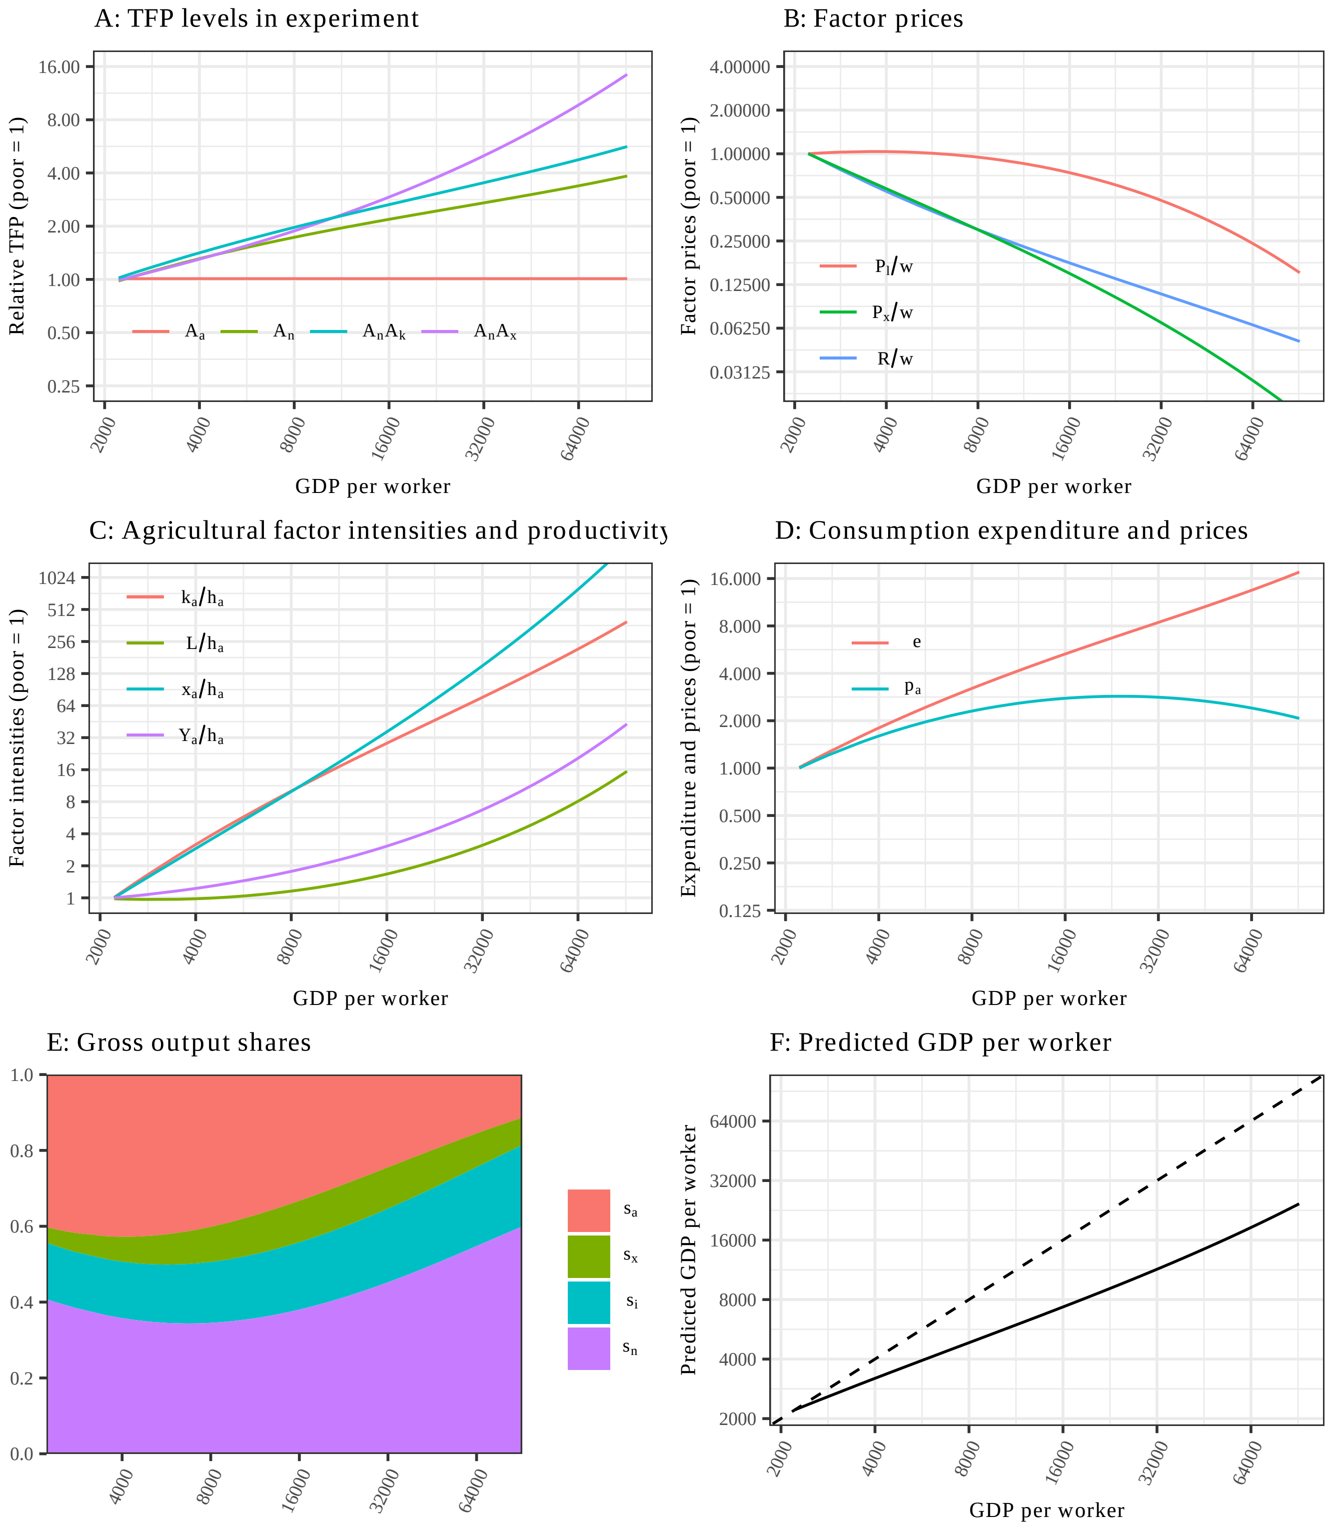

Figure 17: Effect of increasing agricultural TFP

(Stehfest et al., 2014), which has an agricultural block based on the MAGNET computational general equilibrium model (Woltjer et al., 2014). Just as in our paper, food production in MAGNET has a nested CES structure, but the default elasticities are much lower than ours: output in each good is a zero elasticity aggregate of intermediate inputs and value added, value added is an 0.1 elasticity aggregate between land and other primary inputs, and other primary inputs are an 0.64 elasticity aggregate between unskilled labor, skilled labor, capital, and natural resources.[17] While some discrepancy can be explained by MAGNET considering production-level rather than sectoral elasticities, we think that between-product substitution is unlikely to explain the full difference. Since, to our knowledge, these CGE parameters are not disciplined by cross-sectional and time series moments, it appears valuable to complement such studies with one based on our model and production-function estimates.

Applying our model to study the effects of a decrease in the land supply is straightforward. We use the calibrated parameters for the poorest county (GDP per worker of 2,200) and consider the effect of reducing its land endowment by up to 80%. Figure 18 shows results of our experiment.

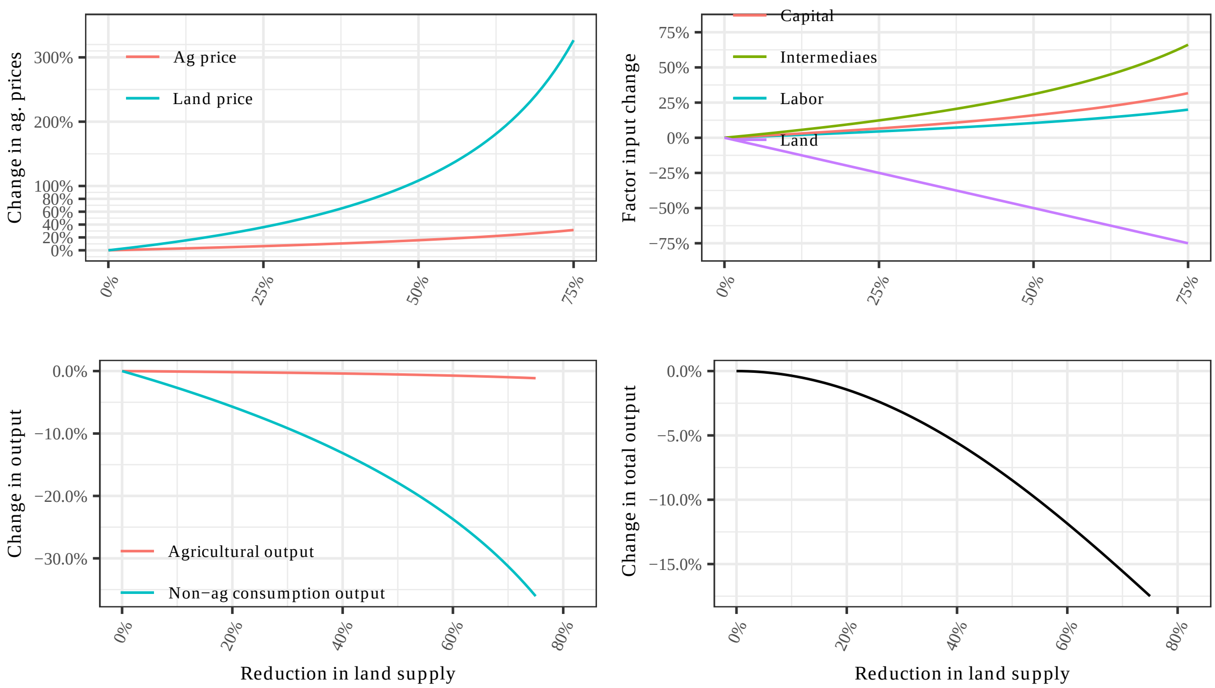

Figure 18: The effects of reduced land supply

We see from the figure how the agricultural employment share (top left), the change in factor inputs (top right), the change in agricultural output (bottom left), and the change in aggregate output per worker (bottom right) depend on the percentage amount by which the land supply is decreased (the horizontal axis). We see that land prices rise dramatically as available land shrink, but the effect on agricultural output prices is more muted since the land share is only 20%-30%. The second panel shows that as land decreases, more capital, intermediates, and labor is moved into the sector, reflecting that a higher agricultural price makes it profitable to put in more factors. The third panel shows how factor intensification implies that agricultural

[17]The MAGNET default elasticities are taken from the food column 6.5.5 in Woltjer et al. (2014).

output is actually quite steady while non-agricultural consumption falls. Intuitively, inelastic demand and strong income effects means that the agricultural sector is sustained in the price increase, and this is done at the expense of the non-agricultural sector. Total output only falls 15%.

Thus, overall, our conclusion is that input substitutability does play a major role and appears promising as an adaptation mechanism. Clearly, much more work is needed to sharpen the predictions we arrive at here, but given the large gap between our findings and those that the IPCC base their estimates on, to us this is clearly a very interesting area for future research.

9 Concluding remarks

In this paper we take an aggregate perspective on agriculture and its role in economic development. Based on a systematic collection of data from a broad cross-section of countries, we uncover some striking facts. The first facts do not require any theory but have, we think, been overlooked in this literature. The nature of these facts suggests the use of an aggregate production function. The use of an abstract and general such function generates new facts (about TFPs by sector) that are very interesting by themselves but, again, also are suggestive of taking a further step. We thus specify a functional-form class, because it allows us to draw new conclusions (about a set of factor substitution elasticities). That concludes our study of production. We also study the demand side—facts and some theory—and, like for the set of production facts, uncover robust and intuitively interpretable features: along the development dimension (measured by GDP per capita), countries differ in highly systematic ways that can be captured with a non-homothetic and yet tractable utility function. Equipped with a demand and a supply side, we then finally conduct general-equilibrium analysis, both for additional validation and for some counterfactual analyses.

The main conclusions can be summarized as follows: we uncover input intensification in agriculture as a robust and quantitatively very important channel. As countries develop, they move away from labor as an input and toward capital and intermediates (fertilizers, pesticides, etc.) that are produced in the non-agricultural sector. Overall, these changes occur as a response to sharp changes in relative prices. The forces underlying these price changes are TFP improvements, where agricultural TFP plays a minor role; instead, it is TFP in capital and intermediates that allow, via factor substitution, agricultural labor productivity to rise as countries develop. Only agricultural TFP growth, moreover, would generate strong counteracting price effects due to non-homothetic demand: consumers would move away from agricultural goods as income rises and, hence, inputs would be drawn away from the agricultural sector. Our counterfactual exercises thus emphasize that TFP improvements in non-agricultural sectors (including that in the production of non-agricultural consumption goods) are key to development, despite the very large employment fraction in agriculture (near 80%) in the poorest countries.

The above conclusions are accompanied with a simple and, we think, rather convincing overall theory of the process of development. It builds on neoclassical forces and makes use of aggregate production functions, just like the macroeconomic growth literature does (along with assumptions about consumer tastes). We are, of course, aware that such functions are not literally correct descriptions of how production takes place; heterogeneity is likely massive, especially in the poorest countries, and exact aggregation appears beyond reach. However, it is still possible that aggregate production functions are reasonable approximations. At the very least, the neoclassical features we point to are so striking that any broad-scoped analysis of development ought to relate to them. An interesting agenda forward is to build a bridge between microeconomic studies of agriculture and our analysis here. If successful, such an agenda would explain in more detail how the aggregate neoclassical facts are generated and, as a result, further help us understand the process of development and how economic policy could be designed to encourage it.

Much has, of course, been put aside as we have proceeded toward the main goals of our paper. Perhaps the most important omission is that whereas we emphasize the role and fundamental endogeneity of input factors, we treat TFPs as exogenous. They are, clearly, key to development from our perspective and need to be brought into focus.

Second, our treatment of production functions focuses on technology appearing in Hicksian, total-factor form. To us, this is the most reasonable starting point and the conclusion from our analysis is that our functional forms capture all the salient features of the data surprisingly well. Moreover, given that our focus is on the long run, we view our functional form $A_a F(k, h, x, l)$ as a reduced form of a function with input-specific technologies that can differ by level of development but are chosen endogenously ("endogenous, directed technical change") as a function of it and can be summarized in the functional form F and the TFP factor A_a; similar approaches are taken in Acemoglu (2002), Caselli and Coleman, Leon-Ledesma and Satchi (2019), and Hassler, Krusell, and Olovsson (2021). Spelling these mechanisms out would be valuable and could generate new insights, as could a more general treatment whereby factor-augmenting technology is allowed.

Third, our analysis is of the "long run" in the sense that we do not explicitly describe the development process: we think of outcomes as steady-state results of different combinations of TFPs (by sector). Our general-equilibrium analysis, in particular, does not describe countries as having growing TFPs. Clearly, this is another simplification: we take the process of accumulating capital for given TFPs to be fast enough that one obtains a good approximation with our procedure. In the presence of differences in capital intensities across sectors, interpreting our analysis as a reduced form of a balanced-growth model is not possible, since the preferences with the non-homotheticity are not consistent with exact balanced growth: the agricultural consumption share goes to zero asymptotically and in the limit the agricultural sector is negligible. It is possible to extend our analysis so that we describe countries at different stages in the path toward this asymptotically balanced path. This would be interesting and could yield new insights. We leave such an endeavor for future work while conjecturing that a slow progression of TFPs relative to capital's convergence speed would deliver results rather closely in line with those here.

Fourth, the maintained assumption in our paper is to consider countries to be closed economies. This is clearly an abstraction that we would like to relax in future work, since it appears possible to view the agricultural sector as one where a country specializes and then sells the output to a world market. As an example, New Zealand lamb and kiwis and Icelandic fish earn significant incomes for countries at high levels of development, quite unlike in our theory. However, in most of our analysis leading up to the closed-economy general-equilibrium treatment, we do not assume that trade is not possible. The separate characterizations of the production and demand sides of the economy in particular do not take a stand on the extent of openness. If one assumed countries are entirely open and that trade costs are zero, of course, prices (in the objects traded) would need to be common across all countries and in this sense would cast a doubt on the price series we use. However, our interpretation is that the differences in prices across countries do reflect trade costs. Interestingly, we do not find major differences in the returns to capital across countries, which of course contrasts wages, which are extremely different across countries (reflecting very limited labor mobility across borders). General-equilibrium analysis of trade is entirely feasible, so as to allow for "New Zealands" to emerge as an outcome of our theory. The key discipline would be to calibrate trade costs so as to generate the magnitudes of trade in agricultural goods as well as in other goods. We would find such an extension extremely interesting.

References

Colen, L., Melo, P., Abdul-Salam, Y., Roberts, D., Mary, S., and Paloma, S. G. Y. (2018). Income elasticities for food, calories and nutrients across africa: A meta-analysis. *Food Policy*, 77:116–132.

Stehfest, E., van Vuuren, D., Bouwman, L., Kram, T., et al. (2014). *Integrated assessment of global environmental change with IMAGE 3.0: Model description and policy applications.* Netherlands Environmental Assessment Agency (PBL).

Woltjer, G. B., Kuiper, M., Kavallari, A., van Meijl, H., Powell, J., Rutten, M., Shutes, L., and Tabeau, A. (2014). The magnet model: Module description. Technical report, LEI Wageningen UR.