

NBER WORKING PAPER SERIES

THE MACROECONOMICS OF INTENSIVE AGRICULTURE

Timo Boppart
Patrick Kiernan
Per Krusell
Hannes Malmberg

Working Paper 31101
<http://www.nber.org/papers/w31101>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023

Boppart thanks Swiss National Science Foundation grant 100018_200958 for financial support. Krusell and Boppart thank Riksbankens Jubileumsfond M16-0345:1. Malmberg and Kiernan thank the Heller-Hurwicz Economic Institute. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Timo Boppart, Patrick Kiernan, Per Krusell, and Hannes Malmberg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Macroeconomics of Intensive Agriculture
Timo Boppart, Patrick Kiernan, Per Krusell, and Hannes Malmberg
NBER Working Paper No. 31101
April 2023
JEL No. E0,O0,O13,Q0

ABSTRACT

Developing countries employ a very large share of their workforce in agriculture, a sector in which their labor productivity is particularly low. We take a macroeconomic approach to analyze the role of agriculture in development. We construct a new database with systematic measures of inputs and outputs of agricultural production around the globe. The data exhibits strong neoclassical features: going from poor to rich countries, capital and intermediate input prices decline dramatically relative to labor prices; concurrently, capital and intermediate input use in agriculture increases by a factor of 300–800 relative to labor. Input intensification accounts for a bit less than two thirds of the agricultural labor productivity gap between the poorest and richest countries. Our observations are well explained by an aggregate agricultural production function with input substitutabilities significantly above unity. On the demand side, standard non-homothetic preferences accurately capture how the expenditure share of agricultural goods varies across the development spectrum. We incorporate our findings into a closed-economy general-equilibrium model with minimal distortions, showing that non-agricultural TFP differences play a much more important role than agricultural TFP differences in explaining income differences.

Timo Boppart
Institute for International Economic Studies
Stockholm University
SE-106 91 Stockholm
Sweden
timo.boppart@iies.su.se

Patrick Kiernan
University of Minnesota
kiern018@umn.edu

Per Krusell
Institute for International Economic Studies
Stockholm University
106 91 STOCKHOLM
SWEDEN
and CEPR
and also NBER
per.krussell@iies.su.se

Hannes Malmberg
University of Minnesota
pmalmber@umn.edu

An online appendix is available at: <http://www.nber.org/data-appendix/w31101>

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.¹ In this paper, we interpret these differences by studying the agricultural sector from a macroeconomic perspective.

Our study is motivated by the observation that agriculture *prima facie* appears to play a central role in the development process. As we move across the development spectrum, we observe a significant decline in the share of the workforce employed in agriculture. Figure 1 illustrates how the agricultural employment share falls in GDP per capita in a stable, log-linear fashion, with agricultural employment going from approximately 80 percent in the poorest countries to less than 1 percent in the richest countries.




Figure 1: Agricultural employment shares across countries

These patterns mean that the income levels of the poorest countries are strongly shaped by the performance of the agricultural sector, whereas in the richest countries, this sector plays a relatively minor role. At the same time, the agricultural sector in poor countries performs particularly badly in terms of labor productivity. Poor countries have an agricultural labor productivity (real gross agricultural output per worker) that is about 100 times lower than in rich countries, dwarfing the already large differences in

¹Income 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 (2016)—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.

aggregate labor productivity.

The high agricultural employment shares and that low agricultural labor productivities in poor countries have put agriculture at the center of the agenda of development economists. The following 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 et al. (2008))

“... agricultural productivity is critical for understanding aggregate income differences.”

(Donovan (2021))

This paper revisits the role of agriculture in the development story. In doing so, we take a macroeconomic perspective and go through the steps that are traditionally taken in aggregate macroeconomic analysis, following Solow (1956), Solow (1957), and many others.

We begin by collecting systematic measures on 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. As we go from poor to rich countries, the quantities of capital and non-agricultural intermediates relative to labor input rise by factors of 300 and 800 respectively, and their compensation shares as a proportion of gross output rise from 3% for capital and 7% for intermediate inputs to approximately 44% and 21%, respectively. These magnitudes suggest that factor intensification in agriculture is an important phenomenon, also relative to purely technological factors.

Next, we show that the data exhibits distinct neoclassical characteristics. Specifically, we observe a systematic association between high relative input prices and low relative input quantities. Furthermore, when plotted in log-log space, these relationships are approximately linear for almost all pairs of inputs (even though slope magnitudes might vary). Strikingly, these relationships are not limited to the cross-section of countries, but also extend to the U.S. time series with similar magnitudes. These patterns motivate our approach, which is to take a macroeconomic perspective and analyze our data using a sectoral neoclassical production function for agriculture.

Starting from an arbitrary production function and a minimum of theory, we estimate agricultural TFPs across the development spectrum using a cross-sectional variant of Solow’s non-parametric TFP accounting. We find that increases in inputs together account for a bit less than two-thirds of the agricultural labor

productivity gap between the richest and the poorest countries, with agricultural TFP accounting for the remainder.

To sharpen our characterization of the supply-side, we make more specific assumptions, and posit that agriculture has a nested CES production function with Hicks-neutral TFP differences across countries. Thus, our starting point is that the agricultural sectors of various countries might exhibit different levels of productivity, but that they share a common shape of their production functions. The formulation has a strikingly good fit over the entire development spectrum: even though our estimation only uses information from the poorest and richest countries, our model accurately captures the full set of relative prices and quantities observed in between. The good fit reflects that the data has a log-linear structure of prices and quantities with respect to income levels, a structure that is well-captured by a nested CES production function. Our estimation results indicate that the elasticity between labor and capital, as well as the elasticity between labor and intermediate inputs, significantly exceeds unity. However, a unitary elasticity for land provides a reasonable approximation. These high substitution elasticities imply that factor intensification is a central phenomenon.

We also posit an aggregate production function for non-agricultural goods. This production function only uses capital and labor as inputs, and here we find a good fit with a unitary elasticity between capital and labor—that is, a Cobb-Douglas production function with capital and labor. One implication of this finding is that in the limit when agriculture shrinks, the supply side of our economy converges to a conventional Cobb-Douglas aggregator of capital and labor.

For non-agriculture, we also estimate TFP series separately for consumption, investment, and intermediate goods. We observe that TFP gaps between rich and poor countries are most pronounced in intermediate goods, followed by investment goods. The low TFPs in these sectors within poor countries have consequences beyond the non-agricultural sector. Indeed, by raising the price of intermediate and investment goods relative to labor, low TFPs in these input-producing sectors prompt substitution towards labor in agriculture, which lowers agricultural labor productivity. This phenomenon is central to our story: low labor productivity in the agricultural sector emerges as an equilibrium outcome, with low non-agricultural productivity being a central driver of low agricultural labor productivity via its knock-on effects on agricultural input intensity. The low input intensities in agriculture do not rely on distortions or inefficiencies within the agricultural sector, since when the non-agricultural sector is unproductive, low input intensity becomes the cost-minimizing production option. Conversely, if the non-agricultural sectors develop and wages rise, we expect agricultural labor productivity to improve as well.

On the demand side, we use a standard non-homothetic utility function. The implied demand system has a log-linear relationship between the agricultural expenditure share, on the one hand, and income and prices, on the other. The cross-sectional data is also approximately log-linear in these variables and hence the income and price elasticities cannot be pinned down separately. We use an income elasticity for agricultural goods at 0.35, which is the closest possible to the 0.6–0.7 range typically found in micro studies; this implies that we set the price elasticity to zero. We could alternatively use elasticity estimates straight from the micro studies, but then our model cannot match the aggregate data quite as well, which is why we opt for the lower elasticities. Fortunately, in terms of our key findings this choice does not play a quantitatively significant role.

Equipped with fully specified demand and supply sides, we formulate a complete general-equilibrium theory that can be used to run counterfactuals. The GE model is designed to match the data at the extremes of our sample of countries exactly; in particular, we demand that all the aggregates of our model (such as investment shares of output, input shares, etc.) match those in the data. For this reason, some basic parameter values need to be different across countries. We show first that agricultural revenues fall short of costs (as we measure them). The discrepancy is very minor in the poorest countries but rather large at the other end of the development spectrum. Given a model where the agricultural sector has constant returns to scale, we interpret these observations as a wedge—a subsidy—in favor of agriculture that varies with development. This wedge is not so large that it affects our key findings but we must include it in order to make our model match the aggregate data. Second, capital depreciation rates need to be slightly higher in rich than in poor countries—since investment-to-capital ratios are higher in richer economies. A comforting finding is that our GE, thus calibrated, delivers an impressive within-sample fit as well. I.e., our functional forms for technology and preferences deliver the shapes that we observe across the development spectrum for all the relevant consumption and input shares. In our general equilibrium analysis, we assume that economies are closed and that there is a minimal set of distortions. While this starting point is a natural choice, it is not the ultimate answer; trade is arguably important, though it can be hampered by trade frictions. Additionally, distortions, particularly those complicating the transition of workers from rural to urban areas, play a role.

Our main counterfactuals consist of closing TFP and human capital gaps between countries; we examine these gaps both individually and together, as the model features nonlinearities. We find that the TFP differences in investment goods production across the development spectrum are by far the most important single factor in terms of understanding output differences across countries; the TFP gaps for agricultural

and non-agricultural consumption goods are far less important. In terms of understanding agricultural labor productivity, again the central feature is input intensification that comes about from TFP improvements in capital as well as intermediates production. Intermediates production, moreover, becomes more and more important as its cost share rises. The latter point reflects another key finding in our counterfactual exercises: although analyses based on (close relatives of) the Hulten theorem are extremely useful, when changing parameters like TFPs as much as they vary across the development spectrum, there is significant structural change: expenditures and cost shares change greatly, making first-order approximations less relevant. A particularly revealing counterfactual involves gradually raising agricultural TFP from the level estimated for the poorest country towards that estimated for the richest country. This leads to significant initial improvements as agriculture is a large sector central at low income levels. However, the benefits of further increases in agricultural TFP quickly diminish: the low TFP in the remaining sectors becomes crucial as goods demand shifts away from agriculture, while agriculture itself faces challenges due to a shortage of capital and intermediate goods. We find the effects of raising human capital from the lowest to the highest level observed to have close to log-linear (and rather standard) effects. In a final counterfactual we also equalize all parameters that differ across countries, like wedges and depreciation rates, so as to check whether our main findings depend on these differences. They do not.

Finally, we use our theory in an application: we study the likely effects of climate change in reducing the availability of arable land. In the IPCC reports, attempts are made to estimate how much agricultural 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 output 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. Section 2 presents the main facts that serve as inputs for our study. Since we view one of our main contributions to be the systematic collation of data on quantities and prices, the fact presentation is preceded by a discussion of our data sources and our data construction. A detailed online appendix accompanies this part of the section. Section 3 explains the construction of agricultural and non-agricultural TFPs along the development spectrum. Section 4 makes functional form assumptions and estimates fully parameterized production functions for both agricultural and non-agricultural sectors. Section 5 introduces a non-homothetic utility function for agricultural and non-agricultural goods and applies it to our

dataset. Section 6 combines the supply and demand sides in a general-equilibrium structure, calibrates it—though most parameters have already been determined earlier—and uses it for our counterfactual analyses. Section 9 examines land reductions due to climate change, and Section 10 offers concluding remarks. However, before we begin, let us situate the paper within the relevant literature.

1.1 Related literature

Our paper addresses two classic development questions: how development leads to structural change out of agriculture (Lewis et al. (1954); Rostow (1960); Kuznets (1966)) and how countries solve the “food problem” (Schultz (1953)). Over the last three decades, research on economic development has taken important steps forward. A portion of this literature, like the present paper, adopts a macroeconomic approach, focusing on aggregates, while a substantial body of work takes a microeconomic perspective, examining microeconomic data and/or conducting randomized controlled trials (RCTs). At present, there are few papers bridging these two literatures; however, one of our primary objectives looking forward is to help provide such a bridge.² Specifically, our paper raises the question of which kinds of microeconomic structures could generate aggregate patterns that are accurately represented by the production functions we find to fit the cross-country data so well in our current analysis.

The early literature on development accounting in macroeconomics focused on aggregate production functions that characterized the economy as having a single sector (e.g., Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999)). This approach uses observations on aggregate outputs and inputs, combined with assumptions on production, to derive aggregate TFP levels for each country. Early in the literature, Caselli and Coleman II (2001), while examining convergence across U.S. states from a historical perspective, recognized the importance of separating out an agricultural sector. Moreover, Caselli (2005) highlights the importance of agriculture in understanding developing economies, specifically noting that agricultural labor productivity gaps between rich and poor countries are more than twice as large as those for aggregate labor productivity.

These points were elaborated upon in Restuccia et al. (2008) and in a more recent study by Gollin et al. (2014), which corroborate the earlier findings using improved data. Simultaneously, the significance of differences in relative prices of investment goods across countries has been underscored (see Hsieh and Klenow (2007)). Additionally, the importance of capital upgrading in agriculture has been noted (Caunedo and Keller, 2021). In sum, the literature has proposed various avenues for breaking down aggregate TFP

²This point is also emphasized in Buera et al. (2021). Interesting recent papers by Lagakos et al. (2018), Gollin et al. (2021) and Bergquist et al. (2022) are examples of how such bridges can be built.

into sectoral components.

Subsequently, several papers have developed explicit macroeconomic frameworks incorporating both agricultural and non-agricultural sectors to better understand sector-specific data and structural change out of agriculture. These papers include Gollin et al. (2002), Gollin et al. (2004), Gollin et al. (2007), Restuccia et al. (2008), and many others. We are influenced by this literature in several ways. Some of these papers highlight the role of capital in agriculture (in addition to Gollin et al. (2007) and Gollin et al. (2004) from above, especially Caunedo and Keller (2021), Storesletten et al. (2019), and Chen (2020)), emphasizing the notable differences in mechanization between rich and poor countries, thereby highlighting the significance of a well-functioning manufacturing sector for agricultural productivity. Other contributions underscore the importance of specific intermediate inputs for agriculture, such as fertilizers, pesticides, and seeds, thereby connecting productive agriculture to a well-functioning non-agricultural sector. Key papers in this area include Restuccia et al. (2008) and Donovan (2021). Yet other papers emphasize the key role played by land and misallocation thereof (in particular Adamopoulos and Restuccia (2014), Chen et al. (2022a), Chen et al. (2022b), and Gottlieb and Grobovšek (2019)), and consequently, the decreasing returns to any remaining inputs. Motivated by these studies, we employ an agricultural production function with four inputs: labor, capital, intermediates, and land.

In our paper, we devote significant effort to discussing the functional-form choices for technology and preferences. While many papers in the literature use Cobb-Douglas production functions for agricultural production (exceptions include Herrendorf et al. (2015), Alvarez-Cuadrado et al. (2017), and Chen (2020)), we study the class of nested CES functions. Our data suggest higher substitutability elasticities between inputs than Cobb-Douglas, ranging from 1.3 to 1.8 depending on the input factor pair/composite; between capital and labor, we estimate an elasticity of just under 1.8 (Herrendorf et al. (2015) estimate 1.58 in U.S. time series data).³ The degree of substitutability between labor and other inputs also influences the perceived wedge—de-facto tax—on labor allocated to non-agricultural production. Under a Cobb-Douglas structure, this wedge needs to be substantial, but with a higher elasticity between labor and other inputs, the required wedge is significantly smaller.

We interpret our nested CES function as a long-run structure, grounded in the concept of “appropriate technology” (see, e.g., Basu and Weil (1998), with a formalization as in Caselli and Coleman II (2006)).⁴ Rather than examining individual-country dynamics, we interpret the data as countries being at different

³The lowest elasticity is for land, but it is still above unity, contrasting with Bustos et al. (2016) and Leukhina and Turnovsky (2016), who argue that they are below unity.

⁴See also Costinot and Donaldson (2016) for recent evidence on the importance of crop switching for the gains from economic integration.

steady states (defined by distinct TFP pairs for agriculture and non-agriculture). This assumption is reasonable if the convergence in capital for given TFP levels occurs quickly compared to the changes in TFPs over time.

In our general-equilibrium setting, much like in many papers in the reviewed literature, we use a non-homothetic demand system to align with Engel’s law. While a significant portion of the literature relies on a Stone-Geary-type structure with a subsistence level of consumption (see, e.g., Caselli and Coleman II (2001), Gollin et al. (2002), Gollin et al. (2007), Restuccia et al. (2008), Gottlieb and Grobovšek (2019), Chen (2020), etc.), we opt for preferences within the price-independent generalized linearity class, which can generate a sustained income effect on expenditure shares (see, e.g., Boppart (2014), Eckert and Peters (2022), and Alder et al. (2022) for applications thereof in the structural change literature). Our specification nests the preferences used in Gollin et al. (2002) and Gollin et al. (2007) as a special case.

Lastly, our paper distinguishes itself from the recent agriculture-and-macro development literature in terms of data construction. We strive to collect the best available measures for all inputs and outputs, as well as their respective prices, with all constructs organized around the general-equilibrium model we subsequently analyze. For a country at any given level of development, the steady state of our model delivers model outcomes, including all shares and relative prices, that coincide with those observed in the data. We measure TFPs in agriculture using a new method, which should be seen as an extension of the approach in Solow (1957). Solow employed minimal assumptions on aggregate technology (specifically, he made no assumptions on functional form) to back out how TFP changed over time. We apply Solow’s approach across space, projecting all individual-country agricultural variables onto a development index (which we consider to be overall GDP per worker), and subsequently back out how agricultural TFP changes with development, once more without making functional-form assumptions. To estimate input elasticities, we then assume that TFP differences are Hicks-neutral and use a nested CES function with elasticities that are the same across countries.

At this point, we do not make contact with the microeconomic development literature, which of course has studied agricultural production from a variety of angles. One important strand of the literature uses a combination of theory and data to draw inferences both about production techniques and a variety of market imperfections (Besley (1995), Rosenzweig and Udry (2014), Foster and Rosenzweig (2022)). Another strand of the literature employs RCTs (Duflo et al. (2011), Bold et al. (2017), etc.) to determine what works and what doesn’t in food production. There is, moreover, some overlap between these literatures. From our viewpoint, it is crucial to note that our structure has a long-term focus—papers utilizing short-term

Table I: Data sources

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 et al. (2014) Mincerian returns estimates
WIOD	World Input-Output Database
AR	Adamopoulos and Restuccia (2022) 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

data tend to find much lower substitution elasticities across inputs—and that we interpret our aggregate production function for agriculture as a reduced form. This means that it is an aggregation based not only on crop choice and endogenous technology but also on heterogeneity across farms and market frictions. Building a bridge to these literature is, therefore, the next step (Chen (2020) represents one such attempt of micro-founding macro patterns in the context of capital). Another essential point of reference is the production-function aspect of agricultural economics (see, e.g., Binswanger (1974)); however, this literature typically has a different focus, as it often centers on specific crops/technologies and shorter time horizons.

2 Aggregate facts

2.1 Data sources

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. The database is designed to accurately represent 2005 data. Table I lists these databases and other data sources, with abbreviations to facilitate the description below; the online appendix contains a more detailed discussion of the data construction.

We use capital boldface to denote nominal variables and lowercase to denote real variables; e.g., $\mathbf{Y}$ is nominal GDP and y is real GDP. Capital letters which are not boldfaced denote certain endowments used later in our theory section, such as L (real quantity of land). Throughout the paper, we distinguish two sectors: agriculture (a) and non-agriculture (n). Our measures are constructed from a set of basic data

series, given below:

- $\mathbf{Y}$: Nominal aggregate output net of natural resource rents per worker (henceforth nominal GDP). Nominal aggregate output and the share of natural resource rents in GDP are from the WDI; employment is from the PWT.
- $\mathbf{Y}_a/\mathbf{Y}$: Nominal agricultural output as a share of GDP, with nominal agricultural output taken from the FAO and the UN.
- H : Human capital level per worker, constructed as $H = e^{s \times \text{minc}}$ with average years of schooling s from BL and Mincerian returns from CPR.
- s_h : Labor compensation share from the PWT, deflated by one minus the natural resource rent share from the WDI.
- emp_a : Share of employment in agriculture from IPUMS and the FAO.
- $\mathbf{P}_a$: Agricultural output price index from the FAO.
- $\mathbf{P}_k$: Investment price index from the PWT.
- $\mathbf{P}_c$: Consumption price index from the PWT.
- $\mathbf{P}_x$: Price index for outside-of-sector agricultural intermediate inputs. For agricultural intermediate input prices, there are unfortunately no recent aggregate data available, with the last major FAO data set being from 1985. Lacking macro data, we proxy prices using bilateral trade data from BACI: we obtain prices of a selected sample of internationally traded fertilizers and pesticides and aggregate them based on their respective weights.
- $\mathbf{K}/\mathbf{Y}, \mathbf{K}_a/\mathbf{Y}_a$: Nominal capital-output ratio in the full economy and in agriculture. The total capital stock is the PWT real capital stock multiplied by the investment good price index and the agricultural capital stock is from the FAO.⁵
- L : Quality-adjusted land per worker. Total cropland is from the FAO and land quality estimates are from AR.
- $\mathbf{I}/\mathbf{Y}$: Nominal investment-output ratio from the PWT, deflated by one minus the natural resource rent share from the WDI.

⁵For the poorest countries, we also validate the FAO capital data with microdata from the LSMS.

- $s_{a,l}$: Land input costs as shares of total agricultural revenue. Land shares are constructed from LSMS and VDSA micro data for Sub-Saharan Africa and South Asia, respectively, and are calculated using arable land and average rental rates from the EU and USDA for Europe and the U.S., respectively.
- $s_{a,x}$: Intermediate input costs as shares of total agricultural revenue. Intermediate input shares are calculated using data on intermediate input use in agriculture from the UN and the FAO, adjusted to net out within-sector intermediates using data from the WIOD.

We use a quadratic approximation to project all basic series onto a development index, which we take to be log real GDP per worker.⁶ This procedure lets us characterize the data along our dimension of interest even in cases where not all countries have data for every measure. Real GDP per worker is calculated by dividing our nominal GDP measure, that is, GDP net of national resource rents, by total employment and a price index constructed from the Penn World Table by dividing nominal with real GDP.

Derived data series. Using these data series, we construct a full set of prices and quantities for factors and outputs, as well as other variables relevant to the macroeconomy. For these derived measures, we construct approximating functions on the development index by simply combining the functions estimated for the basic measures.⁷ This strategy ensures that relevant accounting identities are preserved even under non-linear transformations.

Labor. The labor input in each sector normalized by aggregate employment is given by the share of workers in that sector times their average human capital:

$$h_a = emp_a H_a,$$

$$h_n = (1 - emp_a) H_n.$$

Average human capital levels by sector are the solutions to the following pair of equations:

$$H = emp_a H_a + (1 - emp_a) H_n,$$

$$\frac{H_n}{H_a} = gap,$$

⁶To construct real GDP, we use the PWT 10.0 variable "expenditure-side real GDP at chained PPPs" (rgdpe), deflated by the PWT employment level and multiplied by one minus the natural resource rent share of GDP from the WDI. For our projections, we choose a functional form that provides a good quadratic fit with respect to log real GDP per worker, using the logarithms of $\mathbf{Y}$, $\mathbf{Y}_a/\mathbf{Y}$, H , $\mathbf{P}_a$, $\mathbf{P}_k$, $\mathbf{P}_c$, $\mathbf{P}_x$, $\mathbf{K}/\mathbf{Y}$, $\mathbf{K}_a/\mathbf{Y}_a$, L , and $\mathbf{I}/\mathbf{Y}$, and a logit transformation of the shares s_h , emp_a , $s_{a,l}$, and $s_{a,x}$. The graphs are provided in the online appendix.

⁷For example, given a projection that expresses $\log(\mathbf{I}/\mathbf{Y})$ and $\log(\mathbf{Y})$ as quadratic functions of the development index $f_{\log(\mathbf{I}/\mathbf{Y})}(y)$ and $f_{\log(\mathbf{Y})}(y)$, we define consumption as a function of GDP as $\mathbf{C} = \mathbf{Y} - \mathbf{I} = \exp(f_{\mathbf{Y}}(y))[1 - \exp(f_{\mathbf{I}/\mathbf{Y}}(y))]$.

where *gap* is the ratio of human capital in non-agriculture vs agriculture, which we take to be a factor of 2 and stable with income level (Herrendorf and Schoellman, 2018). The average price per unit of labor input in the economy equals measured aggregate labor compensation divided by total labor input,⁸

$$\mathbf{W} = \frac{s_h \mathbf{Y}}{H}.$$

We assume that wages are the same per efficiency unit in agriculture and non-agriculture.⁹

Capital. We constructed the nominal capital inputs $\mathbf{K}$ and $\mathbf{K}_a$ from the capital-output ratios and output measures. Real capital inputs are then given by

$$k_a = \frac{\mathbf{K}_a}{\mathbf{P}_k}$$

$$k_n = \frac{\mathbf{K} - \mathbf{K}_a}{\mathbf{P}_k}$$

To obtain the (nominal) rental cost per unit of capital, we multiply the share of nominal capital values (the ratio of capital compensation to the nominal capital-output ratio) with the price index of investment goods

$$\mathbf{R} = \frac{1 - s_h - s_{a,l} \frac{\mathbf{Y}_a}{\mathbf{Y}}}{\mathbf{K}/\mathbf{Y}} \times \mathbf{P}_k.$$

Intermediate inputs. The price of intermediate inputs is one of the basic data series. The nominal and real values of intermediate inputs, respectively, are given by

$$\mathbf{X}_a = s_{a,x} \mathbf{Y}_a$$

$$x_a = \frac{\mathbf{X}_a}{\mathbf{P}_x}.$$

Land. The quality-adjusted quantity of land L is one of the basic data series. We back out land prices from total compensation to land as

$$\mathbf{P}_l = \frac{s_{a,l} \mathbf{Y}_a}{L}$$

⁸Using the PWT labor share to measure wage income introduces some degree of self-referentiality since the PWT makes assumptions about agricultural labor income to split self-employment income. For robustness, we have checked that one obtains a similar pattern by using data on self-employment shares and assuming that the wage gap between employed and self-employed people equals the wage gap between non-agricultural and agricultural employees.

⁹The human capital gap is aligned with the sectoral wage differences found in Herrendorf and Schoellman (2018), which they show can be largely explained by different levels of, and returns to, human capital across sectors. This assumption can be relaxed by introducing a mobility wedge that explaining part of the wage difference.

Equating costs and revenues. Our input cost measures in agriculture are defined independently of agricultural revenue data, which means that revenue and costs are not necessarily the same by construction. Empirically, we find that costs exceed revenues over the whole development spectrum.

Our finding is closely related to those of Herrendorf and Schoellman (2015), who note that the operator surplus in U.S. agriculture implies payments to owner-operators that are very low compared to the wages of agricultural workers. This observation is a mirror image of our findings. Indeed, one way to express our findings is that if owner-operators were paid a market wage, operator surplus would be negative.

Using U.S. data, Herrendorf and Schoellman argue that the most likely explanation is an understatement of agricultural output, both through sectoral misclassification and under-reporting of taxable income. In our data, we capture this explanation in reduced-form by introducing a subsidy τ_a that drives a wedge between agricultural revenue and agricultural costs.¹⁰ We define the costs in the agricultural sector as

$$\mathbf{Cost}_a = h_a \times \mathbf{W} + k_a \times \mathbf{R} + (s_{a,l} + s_{a,x}) \times \mathbf{Y}_a,$$

and the subsidy as

$$\tau_a = \frac{\mathbf{Cost}_a}{\mathbf{Y}_a} - 1.$$

We define agricultural consumer prices as

$$\mathbf{P}_{a,c} = \frac{\mathbf{P}_a}{1 + \tau_a},$$

recognizing that our measure $\mathbf{P}_a$ is output prices and not consumer prices. The next section shows how the subsidy varies with the level of development.

Factor compensation shares. Given the input and price measures, agricultural factor compensation shares are

$$\begin{aligned} s_{a,h} &= \frac{h_a \mathbf{W}}{\mathbf{Cost}_a} & s_{a,k} &= \frac{k_a \mathbf{R}}{\mathbf{Cost}_a} \\ s_{a,h}^{VA} &= \frac{s_{a,h}}{1 - s_{a,x}} & s_{a,k}^{VA} &= \frac{s_{a,k}}{1 - s_{a,x}} & s_{a,l}^{VA} &= \frac{s_{a,l}}{1 - s_{a,x}}, \end{aligned}$$

with the first and second rows giving shares as a proportion of gross output and value added, respectively.

Note that $s_{a,x}$ and $s_{a,l}$ were given as basic time series.

¹⁰An alternative would be to simply scale up agricultural output with a corresponding factor. We choose a subsidy approach for two reasons. First, a major part of under-reporting reflects tax evasion, and by excluding that output we keep our revenue data consistent with the one used in the construction of GDP. Second, there is extensive government support in many countries, and in these cases, subsidies might be a plausible explanation for the deviation between costs and revenues.

Compensation shares in non-agriculture are

$$s_{n,h} = \frac{h_n \mathbf{W}}{\mathbf{Y}_n},$$

$$s_{n,k} = \frac{k_n \mathbf{R}}{\mathbf{Y}_n},$$

where non-agricultural output is total output net of agricultural value added (using that GDP accounting defines output as revenue minus subsidies):

$$\mathbf{Y}_n = \mathbf{Y} - [\mathbf{Cost}_a - \mathbf{X}_a].$$

Output. Real output in the agricultural sector and real non-agricultural consumption, respectively, are defined as

$$y_a = \frac{\mathbf{Cost}_a}{\mathbf{P}_a}, \quad c_n = \frac{\mathbf{Y} - \mathbf{I} - \mathbf{Cost}_a}{\mathbf{P}_n},$$

where we use that real output in agriculture equals costs rather than revenue in the presence of the subsidy and where non-agricultural consumption is defined as total output minus total investment less consumption of agricultural goods (which we assume equals agricultural production).

Consumption shares and consumption prices. We construct total and sectoral consumption levels using the following equations (in line with how we defined non-agricultural consumption above):

$$\mathbf{C} = \mathbf{Y} - \mathbf{I},$$

$$\mathbf{C}_a = \mathbf{Y}_a(1 + \tau_a), \quad \mathbf{C}_n = \mathbf{C} - \mathbf{C}_a,$$

$$s_{a,c} = \frac{\mathbf{C}_a/(1 + \tau_a)}{\mathbf{C} - \tau_a \mathbf{Y}_a}, \quad s_{n,c} = 1 - s_{a,c},$$

where the first equation defines total nominal consumption $\mathbf{C}$ as output net of investment, and where $\mathbf{C}_a$ and $\mathbf{C}_n$ represent sectoral consumption levels for agriculture and non-agriculture, both expressed in cost terms. Last, $s_{a,c}$ and $s_{n,c}$ are the shares of agriculture and non-agriculture in total consumption, and is defined as the share of consumption expenditure taking subsidies into account.

The price index for agricultural goods $\mathbf{P}_a$ is one of our basic series. To construct the price index for non-agricultural consumption, we observe that the Divisia index for aggregate consumption goods satisfies

the following equation, for each value of our development index y (measured as GDP per capita):

$$d \log \mathbf{P}_c(y) = s_{a,c}(y) d \log \mathbf{P}_{a,c}(y) + s_{n,c}(y) d \log \mathbf{P}_n(y). \quad (1)$$

Thus, we solve for the non-agricultural prices along the development path by re-arranging this equation (1) to obtain the non-agricultural price index from the following initial value problem:

$$d \log \mathbf{P}_n(y) = \frac{d \log \mathbf{P}_c(y) - s_{a,c}(y) d \log \mathbf{P}_{a,c}(y)}{s_{n,c}(y)} \quad \log \mathbf{P}_n(y_{us}) = 0,$$

where $\log \mathbf{P}_n(y_{us}) = 0$ normalizes the non-agricultural price index to one at the U.S. income level.

2.2 Stylized aggregate facts: agricultural production and consumption

This section presents the key features of the aggregate data that we have collected. The facts provided here motivate the approach used in our subsequent 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 in agriculture. We have 4 input categories—capital (k_a), labor (h_a , for human capital), intermediates (x), and land (L)—yielding 6 different input ratios. The horizontal and vertical axes are both on log scales, so slopes can be thought of as reduced-form elasticities.

In each of the six sub-figures, the horizontal axis is our development index—real GDP per worker—and the vertical axes are the quantity and price ratios for the input pair in question, with relative prices on the left-hand scale and relative quantities on the right-hand scale. For quantity ratios, we include individual country observations along with their projection onto the development index (the solid blue line). To avoid overcrowding we only plot the projection for the price ratio (the dashed black line).

The top left panel shows the well-known fact that the ratio between the quantities of capital and labor inputs increases significantly along the development dimension: it increases by a factor of about 300 from the poorest to the richest economy. The relative price—the ratio of the rental price of capital to the per-efficiency unit wage—decreases with development by a factor of around 20.

Similar trends are observed for the other input combinations. The price and quantity lines (i) have opposite slopes, (ii) exhibit large variations across the development spectrum and (iii) the relationship are fairly log-linear. The land-to-labor and capital-to-intermediates ratios vary somewhat less, but the differences



GDP less natural resource rents per worker, 2005

Note: Relative prices (dashed black lines) are reported on the left axis. Relative quantities (data clouds and solid blue lines) are reported on the right axis. Relative prices and quantities are normalized such that the fitted value of the relative price is one for the richest country and the fitted value of the relative quantity is one for the poorest country.

Figure 2: Agricultural inputs: quantity and price ratios

are still substantial.

From these graphs, we can see that there is a striking and robust pattern in terms of quantity and price ratios along the development dimension: the relative input prices in agriculture move in an opposite direction with their corresponding relative input quantities. This relationship is tight and involves changes in ratios by orders of magnitude. To us, this is suggestive of neoclassical forces at work: where inputs are expensive, they are used less. The neoclassical production function perspective is precisely that cost-minimizing behavior induces a negative correlation between relative factor ratios and relative factor prices, as low prices of a factor imply cost-minimizing input combinations that are more intensive in that factor of production. These observations make us optimistic that an aggregate production function for agriculture can provide a good account of the data, even across very different income levels.

Before proceeding with the formulation of an aggregate production function, let us also examine the time-series perspective. In Figure 3, we see the same cross-sectional data as the previous figure, but now with the U.S. time series added in red (plotting the level of development of the U.S. at a given point in time). The data covers the period of 1950–2018 apart from land prices which begin only in 1990. 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 do not exactly align with 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 to be appreciably different from those in the cross-section.

Further examination of the previous figures reveals (i) that slopes vary by input combination and (ii) that slopes of relative quantities and prices do not sum to zero, i.e., shares are not constant along the development dimension. Figure 4 shows how the 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 about 7% for the poorest countries but about 44 percent for the richest. This finding, in fact, is a strong motivation for including intermediate inputs in our analysis. As discussed above, Restuccia et al. (2008) and Donovan (2021) also look at the role of intermediates.

Turning to the remaining inputs, we see that, in contrast, the land share out of value added is slightly decreasing in development, being a bit above 20 percent for poor countries; as a share of gross output,




Figure 4: Input shares—gross output and value added

it is cut by slightly less than one-third as intermediate inputs rise to about 44 percent.¹¹ The capital share of value added increases significantly—from less than 5 percent to almost 40 percent. Finally, the labor share declines markedly, from roughly 70 percent (of both gross output and value added) to below 50 percent of value added and about 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 of development.

2.2.2 Consumption of agricultural goods: shares and prices

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.¹²

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

¹¹For the rich countries, our land share is line with the U.S. Department of Agriculture's estimate of 15% of payments going to land services. For poor countries, the share is lower than what is found in Chen et al. (2023) or from rules of thumbs starting from, for example, sharecroppers paying 50% of output in land rent (Mundlak, 2005). See Online Appendix A.1 for a further discussion.

¹²Recall that the share of agricultural consumption is defined as the expenditure on agricultural goods relative to total consumption expenditure. This captures ultimate consumer demand for agricultural products and will, in general, be lower than the food share of expenditure that involves value added outside of agriculture.




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

Note: The agricultural consumption share is calculated as total agricultural revenue divided by total consumption expenditure. Consumption expenditure is defined as GDP net of investment and agricultural subsidy payments, with subsidy payments defined as agricultural revenue times the average subsidy rate at that income level.

our development index. Here, we also have a 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 2—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 falling share will also be partly explained by strong complementarities in consumption, together with the falling relative price of agricultural consumption.

The data in Figure 6 also provides an initial hint into the relative total-factor productivities in the agricultural sector compared to the non-agricultural sector along the development dimension. If the production technologies in the two sectors only differed in their TFPs, and if factor prices were equalized across sectors, then observed relative prices would reveal relative TFP levels. This would suggest that agricultural TFP, relative to average TFP, increases with development. However, the data on shares above strongly suggest that the production structures are vastly different across the two sectors. Therefore, measuring how TFP in agriculture changes with development requires a more comprehensive strategy, one that does not rely solely




Figure 6: The relative price of agricultural goods

on prices and also involves less theory.




Figure 7: The backed-out subsidy to the agricultural sector, in percent

The backed out subsidies in agriculture are relatively large and vary across levels of development: they are around 25% for countries at the highest level of development and close to zero at the other end of the spectrum (see Figure 7). The discrepancy between our measure of costs in the agricultural industries and their accounting revenues could potentially be given alternative interpretations (e.g., our cost estimates may be off), but we will find that the magnitude of the differences in subsidies across countries are not large enough to change our quantitative analysis more than marginally.

3 Theory, I: agricultural and non-agricultural TFP

In this section, we use minimal theory to empirically quantify total factor productivity (TFP) differences in agricultural and non-agricultural production across countries. We essentially employ a growth-accounting exercise as developed in Solow (1957) at the sectoral level but applied across the development spectrum as opposed to over time. We assume that each sector produces according to an aggregate production function with constant returns to scale over all inputs. For the agricultural sector, we assume that gross output is given by a function defined over our four studied inputs, i.e., there exists a function F that maps labor input h_a , capital k_a , intermediate goods x_a , and land l into agricultural gross output y_a . Exploiting homogeneity of degree one of the F function, we can express output per labor input as a function of k_a/h_a , x_a/h_a , and l/h_a . In the measurement Section 2 we constructed measures of all these (relative) inputs and gross output per labor input as function of the development index y . We can then write agricultural gross output in terms of these projected functions:

$$\frac{y_a}{h_a}(y) = F\left(\frac{k_a}{h_a}(y), 1, \frac{x_a}{h_a}(y), \frac{l}{h_a}(y); y\right), \quad (2)$$

where the a subscript refers to agriculture. The assumption is that there exists a sufficient statistic y , i.e., a one-dimensional development index that captures differences across countries. Consistent with this view, we will assume in our model further below that countries differ in a number of parameters and endowments but that these differences are all related to a single development index y . Importantly, y shows up in the F function in (2) as a separate argument. This captures the idea that countries at different levels of development may differ not only in terms of their factor inputs but also in terms of their “technologies”; for Solow, y was simply time, t .

For the non-agricultural sector we similarly assume there exists an aggregate production function that

allows us to write non-agricultural consumption quantities produced per labor unit as

$$\frac{y_n}{h_n}(y) = G\left(\frac{k_n}{h_n}(y), 1; y\right), \quad (3)$$

where the n subscript refers to non-agriculture. The non-agricultural production function G is specified at the value-added level and the only two primary inputs are labor h_n and capital k_n .

In spirit similar to Solow’s growth accounting we can then define a local change—in y space—in “total factor productivity” (TFP) in the two sectors as

$$\frac{\partial \log F(\cdot; y)}{\partial y} = \frac{\partial \log(y_a/h_a)(y)}{\partial y} - \epsilon_{F,k}(y) \frac{\partial \log(k_a/h_a)(y)}{\partial y} - \epsilon_{F,x}(y) \frac{\partial \log(x_a/h_a)(y)}{\partial y} - \epsilon_{F,l}(y) \frac{\partial \log(l/h_a)(y)}{\partial y} \quad (4)$$

and

$$\frac{\partial \log G(\cdot; y)}{\partial y} = \frac{\partial \log(y_n/h_n)(y)}{\partial y} - \epsilon_{G,k}(y) \frac{\partial \log(k_n/h_n)(y)}{\partial y}, \quad (5)$$

where $\epsilon_{F,z} \equiv \frac{\partial \log F}{\partial \log z}$ and $\epsilon_{G,z} \equiv \frac{\partial \log G}{\partial \log z}$ are the elasticities of F and G with respect to their inputs z . Put differently, we define a change in the respective output per labor input that is unaccounted for by the output elasticity-weighted change in inputs (again measured relative to labor input) as the residual change in TFP.

The output elasticities can be quantified by the factor cost shares that were depicted in the previous section. After normalizing the TFP levels to 1 at the U.S. level of development y , this approach allows us to back out a TFP series in y space for each sector. The TFP series is capturing a change in residual “technology” expressed in Hicks-neutral units. To see this suppose the F and G production functions take the forms

$$y_a(y) = A_a(y) f(k_a(y), h_a(y), x_a(y), l(y)) \quad (6)$$

and

$$y_n(y) = A_n(y) g(k_n(y), h_n(y)), \quad (7)$$

i.e., technologies differ across countries at different levels of development y only through a factor-neutral technology term. This is indeed a structure we will impose further below. In this case, the method above will precisely result in measures of how A_a and A_n change with y .

In the case of the agricultural sector, we have already presented the measurements of relative inputs and real gross output. For non-agriculture, we don’t have such measures separately for the production of consumption goods, intermediate goods, and investment goods. Therefore we additionally assume that the technologies to produce the non-agricultural consumption good, the intermediate input, and the investment

good share the same isoquants. Then, the capital-labor ratio equalizes across all non-agricultural industries and we can measure the relative factor input by the total non-agricultural input ratio k_n/h_n . In the example above with the Hicks-neutral technology terms, the assumption of identical isoquants in all non-agricultural industries implies that we can then write

$$y_x(y) = A_x(y)A_n(y)g(k_x(y), h_x(y)), \quad y_k(y) = A_k(y)A_n(y)g(k_k(y), h_k(y)),$$

where x stands for intermediate inputs production and k for investment goods. I.e., the three non-agricultural sectors just differ by the Hicks-neutral technology term. Then, under the assumption of competitive output markets, we do not have to redo the TFP computation for the intermediate and capital good separately. We can thus more directly back out series for A_x and A_k by observing how the relative prices of these two output goods change in y and, together with the series in A_n from above, this will yield TFP estimates for intermediates and investment goods, too. This is how we proceed and the results are shown in Figure 8.




Figure 8: TFPs by sector

The figure reveals one of our key findings: differences in agricultural TFP across countries are not systematically larger than in the other (non-agricultural) sectors. In the poorest countries, agricultural TFP is a little over one-sixth of that in the U.S. The TFP differences for the intermediates and investment sectors are larger. In intermediates, the TFP in the poorest countries is just 1/12th of that at the frontier. In

contrast, the non-agricultural consumption TFP gap between rich and poor countries is smaller than the one in agricultural TFP. In sum, while the agricultural TFP gaps are large in the development dimension, they are not as large as for other sectors.

The picture looks quite different if, in contrast, one looks at the *labor productivity gaps* along the development dimension. Figure 9 illustrates, plotting both labor productivity and TFP for agriculture. We see that labor productivity is slightly convex and that its gaps are much larger (it differs by a factor of about 100 between the richest and poorest countries). Consequently, by looking at labor productivities, it appears that closing the technological gap in agriculture would be of first-order importance. From our perspective, however, we simply note that the lion's share (a bit less than $2/3$) of the observed labor productivity gap in agriculture results from differences in capital, intermediate input, and land 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 productivity improvement in other sectors of the economy.¹³




Figure 9: Labor productivity vs. TFP in agriculture

¹³ 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.

4 Theory, II: the agricultural production function

Having estimated TFP series by sector, we now take an additional step and make more assumptions on the shape of the production isoquants. This allows us to estimate macroeconomic input elasticities for each sector: one set of elasticities for agriculture and one for the non-agricultural sectors.¹⁴ These elasticity estimates will be used in subsequent general equilibrium analysis, and they also let us comment on the specific parameter restrictions imposed in the macroeconomic development literature.

Our new assumptions are stronger than those of standard growth accounting. They involve nested CES functions—thus imposing 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 start with the agricultural sector, which is more involved due to it having four different inputs. We use the following nested structure:

$$f(h, k, x, l) = A_a \left[\left(\left(h^{\frac{\sigma_1-1}{\sigma_1}} + \omega_1 k^{\frac{\sigma_1-1}{\sigma_1}} \right)^{\frac{\sigma_1}{\sigma_1-1} \frac{\sigma_2-1}{\sigma_2}} + \omega_2 l^{\frac{\sigma_2-1}{\sigma_2}} \right)^{\frac{\sigma_2}{\sigma_2-1} \frac{\sigma_3-1}{\sigma_3}} + \omega_3 x^{\frac{\sigma_3-1}{\sigma_3}} \right]^{\frac{\sigma_3}{\sigma_3-1}}. \quad (8)$$

Here, labor and capital are nested together with an elasticity parameter σ_1 ; this nest and intermediate inputs are further nested together with an elasticity parameter σ_2 ; the resulting nest, lastly, is nested with land with an elasticity parameter σ_3 . Clearly, this is only one out of 12 possible nesting structures. We picked a nesting where labor and capital appear together, since the macroeconomic development literature most often assumes such a structure (many papers assume Cobb-Douglas) and it does not appear restrictive.¹⁵

Our estimation of the parameters— $(\sigma_1, \sigma_2, \sigma_3, \omega_1, \omega_2, \omega_3)$ in the case of equation (8)—for each nesting 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. We begin by setting the σ_1 and ω_1 in the capital-labor nest to match the capital-labor quantity and price ratios at the two extremes of our sample. Given that capital and labor are at the inner-most nest, this procedure does not involve other production function parameters.

We can then construct the h - k nest using 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

¹⁴Recall that the different non-agricultural sectors are assumed to have identical isoquants.

¹⁵Note also that some of the nestings do not allow us 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)).

simple manipulation of the first-order conditions. This step is repeated once more and we thus obtain all sought parameters. Our procedure does not by itself ensure a maximal within-sample (in the development dimension) fit; an alternative is, say, an OLS regression using all the data.¹⁶ However, the two procedures give very similar results.

The functional form in equation (8) delivers $\sigma_1 = 1.90$. 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 $\sigma_2 = 0.66$: land's elasticity of substitution with the h - k nest is significantly closer to unity. Finally, we estimate $\sigma_3 = 1.75$, thus also indicating a fairly high elasticity of substitution. These values imply the fit shown in Figure 10.




Figure 10: The fit of the nested CES

In Figure 11 below we show the implications for the pairwise substitution elasticities. 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.¹⁷ The pairwise elasticities between k - h is constant by construction. The others are nontrivial.

¹⁶More 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 σ s 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.

¹⁷An alternative measure is the Allen-Uzawa elasticity. Such a 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.

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 $k-l$ elasticity: it moves between 1.25 and 1.¹⁸




Figure 11: Pairwise input elasticities (Morishima)

We also contrast our findings here with a formulation that would impose a Cobb-Douglas elasticity between labor and capital. 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, under a Cobb-Douglas specification 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 (0.67 vs. 0.28 in gross output), with capital moving in the opposite direction (0.03 vs. 0.21). Figure 12 displays how the implied wedge depends on the level of development, where we've defined the wedge to be 1 for the richest

¹⁸Figure A4 in the online appendix compares the Morishima elasticities to an alternative nest. The implied elasticities are very similar with the biggest differences observed for the elasticities involving land.




Figure 12: The “agricultural wage subsidy” implied by Cobb-Douglas production

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 be preventing movement out of agriculture into other sectors, and various moving costs could be the culprit. However, the sheer size of the wedge we back out, in our assessment, is too large to be plausible: it is almost a factor of 16 for the poorest countries (and moves smoothly toward 1 as we consider higher and higher levels of development). Note also that a similarly implausible 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 no such 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:

$$\max_{(c_a, c_n): p_a c_a + p_n c_n = E} u(c_a, c_n) \equiv v(p_a, p_n, E) = \frac{1}{1 - \vartheta} \left(\frac{E}{p_n} \right)^{1 - \vartheta} - \frac{\nu}{1 - \eta} \left(\frac{p_a}{p_n} \right)^{1 - \eta} - \frac{1}{1 - \vartheta} + \frac{\nu}{1 - \eta}. \quad (9)$$

The implied demand for the agricultural good is

$$c_a = \nu \left(\frac{E}{p_n} \right)^{\vartheta} \left(\frac{p_a}{p_n} \right)^{-\eta} = \nu e^{\vartheta} p_a^{-\eta}, \quad (10)$$

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 $p_n = 1$.¹⁹ Likewise, the demand for the non-agricultural good is then

$$c_n = e - \nu e^{\vartheta} p_a^{1 - \eta}. \quad (11)$$

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 U.S. income level.

Figure 13 shows the relative loss compared to the optimum as a function of ϑ and η . In terms of reducing the agricultural consumption share, we obtain a more powerful income effect if there is a low expenditure elasticity ϑ , and we obtain a more powerful substitution effect if there is a low elasticity of substitution η . Reflecting the challenge of separating income and substitution effects in the cross-section, the figure has a ridge of equally good parameter values, which reflects that we explain the data either through a strong substitution effect using a low substitution elasticity η , or through a strong income effect due to a low expenditure elasticity ϑ . Since all expenditure elasticities ϑ are relatively small compared to micro estimates,

¹⁹The 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 both Gollin et al. (2002) and in Gollin et al. (2007) as a special case.

we select the strongest substitution effect and the weakest income effect in the upper-left corner of the ridge, yielding $\vartheta = 0.35$ and $\eta = 0$.

Figure 14 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 the expenditure share closely, with the agricultural expenditure share falling from around 50% 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.²⁰

In terms of external validation, our estimate of $\vartheta = 0.35$ 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). However, the typical micro estimates of the expenditure elasticity of food is at the final expenditure level (including processing, transportation and retail done by non-agricultural industries) whereas our estimate here is in terms of the implied demand at the agricultural industry level (see Herrendorf et al. (2013) for the different perspectives on structural change). Nevertheless there is a sizable gap. 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 Leontief, 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 the micro and macro estimates 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 sensitive to income than food demand.

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—e.g., 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 agricultural output).

We consider a closed economy with a representative household. Studying trade could potentially be important and we discuss trade in our concluding section. Inequality is important to consider but is a

²⁰This 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: Log loss relative to optimum

Note: The graph shows $\log \left(\frac{Loss(\vartheta, \eta)}{\min_{\vartheta, \eta} Loss(\vartheta, \eta)} \right)$ where $Loss$ is defined as the L2-distance from the predicted to the actual agricultural consumption curve given ϑ, η and a share parameter ν that is recalibrated to exactly fit the agricultural consumption share at the U.S. income level.

topic worthy of a separate paper. 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




Figure 14: The share of agriculture in aggregate consumption, 2005

Note: The country observation gives the consumer share of total agricultural consumption, where

then discussed in our concluding section.

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

$$\sum_{t=0}^{\infty} \beta^t u(c_{a,t}, c_{n,t}), \quad (12)$$

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, i.e., 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 capital depreciation rate δ)

$$c_{n,t} + \frac{x_t}{A_x} + \frac{k_{t+1} - (1 - \delta)k_t}{A_k} = A_n g(k_{n,t}, h_{n,t}). \quad (13)$$

Agricultural production takes place according to

$$c_{a,t} = A_a f(k_{a,t}, h_{a,t}, x_t, L), \quad (14)$$

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 (12), subject to the constraints (13), (14),

$$k_{a,t} + k_{n,t} = k_t \quad \text{and} \quad h_{a,t} + h_{n,t} = H \quad (15)$$

holding for all t .

6.2 Decentralized equilibrium

In this section we formally define the 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 a factor-specific wedge τ_k (to capital accumulation) and a sector-specific wedge τ_a (a subsidy to agricultural production). These two wedges are the only deviations from efficiency we consider in the baseline. We comment on other wedges in the conclusion.

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

$$\max_{\{c_{a,t}, c_{n,t}, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_{a,t}, c_{n,t}) \quad \text{s.t.} \quad (16)$$

$$c_{n,t} + p_{a,t}c_{a,t} + \frac{k_{t+1}}{A_k} = w_t H + \left(\frac{1-\delta}{A_k} + (1-\tau_k)r_t \right) k_t + p_{l,t}L + T_t. \quad (17)$$

Note here that we allow a distortionary tax on capital τ_k , redistributed lump-sum via T_t to the representative agent each period.

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

$$\max_{k_n, h_n} A_n g(k_n, h_n) - r_t k_n - w_t h_n, \quad (18)$$

whereas agricultural firms solve

$$\max_{k_a, h_a, x, l} (1 + \tau_a) p_{a,t} A_a f(k_a, h_a, x, l) - r_t k_a - w_t h_a - \frac{x}{A_x} - p_{l,t}l, \quad (19)$$

where τ_a denotes a subsidy on agricultural production.

Notice that all prices above— p_a , r , w , and p_l —are relative prices and measured in units of the non-agricultural consumption good in the same period. With this definition, p_a is the price faced by consumers.

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

1. quantities solve the optimization problems of consumers (16), the non-agriculture firms (18), and the agricultural firms (19), where in addition L is the optimal land choice;
2. quantities are feasible in the non-agricultural sector (13) and the agricultural sector (14);
3. capital and labor markets clear, i.e., (15) is satisfied; and
4. the government budget balances each period, i.e.,

$$T_t = \tau_k r_t k_t - \tau_a p_{a,t} c_{a,t}. \quad (20)$$

Next, we specialize our functional form assumptions, 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 of the utility specification we choose are given by

$$c_n = \frac{E}{p_n} - \nu \left(\frac{E}{p_n} \right)^\vartheta \left(\frac{p_a}{p_n} \right)^{1-\eta} = e - \nu e^\vartheta p_a^{1-\eta} \quad (21)$$

$$c_a = \nu \left(\frac{E}{p_n} \right)^\vartheta \left(\frac{p_a}{p_n} \right)^{-\eta} = \nu e^\vartheta p_a^{-\eta}. \quad (22)$$

As already explained, ν regulates the level of demand for the agricultural good and η is the asymptotic elasticity of substitution between the agricultural and non-agricultural consumption goods; $\eta < 1$ implies that the two consumption goods are gross complements. The parameter ϑ is the expenditure elasticity of demand for agricultural consumption, which turns out to be a constant under this formulation.²¹ Hence the preferences are in line with Engel's law as long as $\vartheta < 1$, whereas preferences become homothetic with $\vartheta = 1$.²²

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

$$q_n(r, w) \equiv \min_{(k_n, h_n): g(k_n, h_n) \geq 1} r k_n + w h_n = (\alpha^\sigma r^{1-\sigma} + (1-\alpha)^\sigma w^{1-\sigma})^{\frac{1}{1-\sigma}}. \quad (23)$$

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 functions for non-agricultural consumption, investment, and intermediate goods, one simply premultiplies the expression in (23) by $1/A_n$, $1/(A_k A_n)$ and $1/(A_x A_n)$, 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

$$q_a(r, w, p_x, p_l) \equiv \min_{(k_a, h_a, x, l): f(k_a, h_a, x, l) \geq 1} r k_a + w h_a + p_x x + p_l l. \quad (24)$$

Given that there are four inputs and the functional form for the cost function depends on the specific nesting,

²¹This 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 quantitatively relevant for the very poorest countries in the sample.

²²With $\vartheta = \eta = 1$ the preferences nest Cobb-Douglas preferences, $u(c_a, c_n) = c_a^\nu c_n^{1-\nu}$, as a special case.

we omit the formulas here to avoid cluttering. Again premultiplying $q_a(r, w, p_x, p_l)$ by $1/A_a$ gives the unit cost in agricultural production.

6.4 Solving for equilibrium

When solving for equilibrium it is helpful to split the problem up into an intratemporal part and an intertemporal part. In the intratemporal part, we only consider how, given a total amount of capital k and a total amount of consumption expenditures 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 $\vec{v} \equiv (c_n, c_a, w, r, p_l, p_a, k_n, k_a, h_n, h_a, x)$. 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 11 elements we need 11 mutually independent equations. First, the consumer's demand functions must be included: (21) and (22). Then, given our choice of the non-agricultural consumption good as the numéraire, perfect competition in the non-agricultural sector delivers

$$r = \alpha k_n^{-\frac{1}{\sigma}} A_n \left(\alpha k_n^{\frac{\sigma-1}{\sigma}} + (1-\alpha) h_n^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{\sigma-1}}, \quad (25)$$

$$w = (1-\alpha) h_n^{-\frac{1}{\sigma}} A_n \left(\alpha k_n^{\frac{\sigma-1}{\sigma}} + (1-\alpha) h_n^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{\sigma-1}}. \quad (26)$$

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

$$1 = \frac{1}{A_n} (\alpha^\sigma r^{1-\sigma} + (1-\alpha)^\sigma w^{1-\sigma})^{\frac{1}{1-\sigma}}. \quad (27)$$

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

$$(1 + \tau_a) p_a = \frac{1}{A_a} q_a(r, w, p_x, p_l), \quad (28)$$

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

$$k_a = \frac{1}{A_a} \frac{\partial q_a(r, w, p_x, p_l)}{\partial r} c_a, \quad (29)$$

$$h_a = \frac{1}{A_a} \frac{\partial q_a(r, w, p_x, p_l)}{\partial w} c_a, \quad (30)$$

$$x = \frac{1}{A_a} \frac{\partial q_a(r, w, p_x, p_l)}{\partial p_x} c_a, \quad (31)$$

$$L = \frac{1}{A_a} \frac{\partial q_a(r, w, p_x, p_l)}{\partial p_l} c_a, \quad (32)$$

where p_x should be evaluated at $1/A_x$ and (32) 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 (14) is met. Finally, capital and labor market clearing requires the two equations in (15). Collecting conditions, we have (15), (21), (22), (25), (26), and (28)–(32), which amounts to 11 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 (27) allows us to express w and as a function of r . Substituting these expressions into (28) 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, (22), to deliver c_a as a function of only the endogenous r and p_l . Inserting c_a and the expression for w from above into the market-clearing equation for land, (32), 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 into the remaining input demand relations in agriculture, (29) and (30), 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 (15) 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 (25), 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 (21) and (31), neither of which were used above.

6.4.2 Intertemporal decisions and the full equilibrium

The household maximizes $\sum_{t=0}^{\infty} \beta^t \left(e_t^{1-\vartheta} / (1-\vartheta) - \nu p_{a,t}^{1-\eta} / (1-\eta) - 1 / (1-\vartheta) + \nu / (1-\eta) \right)$, by choice of $\{e_t, k_{t+1}\}_{t=0}^{\infty}$ subject to the budget constraint, (17), 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

$$\left(\frac{e_{t+1}}{e_t}\right)^\vartheta = \beta(1 - \delta + A_k r_{t+1}(1 - \tau_k)). \quad (33)$$

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. Note, however, that our analysis below only focuses on steady states.

For a shooting algorithm adopted to our framework to solve for a transition path: first 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: (13). 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. To get around this, we can simply add the Euler equation between period 0 and period 1 to the intratemporal system, where we now have 12 equations in 12 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 ; once 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. The Euler equation immediately delivers the rental rate $r^\star$. Together with the intratemporal 12-dimensional system, as well as the non-agricultural resource constraint (13), we can then solve for the 13 remaining unknowns: $e^\star$, $k^\star$, and the remaining

elements of $\vec{v}^*$. 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:

$$r^* = \frac{1/\beta - 1 + \delta}{A_k(1 - \tau_k)}. \quad (34)$$

Together with (27), this determines the steady-state wage:

$$w^* = \frac{A_n(1 - \alpha)^{\frac{\sigma}{\sigma-1}}}{(1 - A_n^{\sigma-1}\alpha^\sigma(r^*)^{1-\sigma})^{\frac{1}{\sigma-1}}}. \quad (35)$$

The first-order conditions from the non-agricultural firm's problem, (25) and (26), then allow us to solve for the capital-labor ratio in non-agriculture as

$$\frac{k_n^*}{h_n^*} = \left(\frac{\alpha w^*}{(1 - \alpha)r^*} \right)^\sigma. \quad (36)$$

Since p_x^* is given directly by $1/A_x$, and since we already have r^* and w_a^* , (28) determines a relationship between the agricultural price and the land price:

$$(1 + \tau_a)p_a^* = \frac{1}{A_a}q_a(r^*, w^*, 1/A_x, p_l^*). \quad (37)$$

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 (32), which sets the demand for land equal to its supply: this pins down c_a^* as a function of p_l^* . By inserting this expression for c_a^* into the rest of the agricultural input demand system, (29)–(31), we obtain k_a^* , h_a^* , and x^* , 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^* , k_n^* , and k^* from (15) and the use of (36). This also gives us total non-agricultural output, allowing c_n^* to be backed out from the sector's resource constraint, (13), with all variables still being a function only of p_l^* . The penultimate step is to construct the steady-state expenditure level, as a function only of the land price, from $e^* = c_n^* + p_a^*c_a^*$. Finally, we insert e^* into the demand equation for agricultural goods, (22), 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 returns to scale.

7 Calibration

The calibration involves selecting a number of country-invariant parameters and functions that map the development index y to country-specific parameter values (the values of the blue parameters). Virtually all these parameters have already been assigned values. Sector-specific TFP differences were obtained in Section 3; the TFP differences for each consumption good were constructed with minimal theoretical assumptions (essentially those used to obtain Solow residuals), whereas the relative TFP differences for capital and intermediates production were obtained with somewhat stronger assumptions (identical isoquants as for non-agricultural consumption goods production). Next, the shape of the production functions f and g were estimated in Section 4 using cross-country variation in relative quantity ratios and relative price ratios for inputs; this procedure was motivated by interpreting the production functions as medium-run outcomes after taking appropriate technology effects (i.e., adjustments of input-specific technology factors) into account.²³ For preferences u the indirect utility function was estimated in Section 5, also based on cross-country variation, from data on total expenditure, relative prices of agriculture and non-agriculture, and the agricultural share of consumption; a comparison of our parameter estimates with those found in a separate study on a cross-section of U.S. consumers is discussed above. To obtain TFP levels from TFP differences, we normalize TFPs at the U.S. income level $A_a(y_{us})$, $A_n(y_{us})$, $A_k(y_{us})$, and $A_x(y_{us})$ to fit price indices and factor prices at y_{us} .

The key remaining parameters are those regulating capital accumulation: β , δ , and τ_k . Ideally, one would use direct observations (from micro- or macroeconomic data) for these, but absent such data we follow the quantitative business-cycle literature in requiring consistency of aggregate variables with steady-state model equations.²⁴ We thus first use an equation that says that investment, in steady state, equals capital depreciated and, second, we impose the consumer's Euler equation. These two equations jointly make

²³The shapes of the isoquants were chosen by restricting attention to nested CES functions. 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: $\sigma = 0.98$.

²⁴Note here that different models, e.g., one with explicit international trade, would lead to different parameter estimates. Note also that how we map variables in the data with model variables is also key; for example, we identify the model's consumption and investment variables with the sum of their private and public counterparts in the data.

investment rates and rental rates consistent with observed capital stocks:

$$\delta(y) = \frac{i(y)}{k(y)} = \frac{\text{nominal investment}}{\text{nominal capital}}(y) \quad (38)$$

$$\frac{1}{1 - \tau_k(y)} \left[\frac{1}{\beta} - 1 + \delta(y) \right] = A_k(y)r(y) = \frac{\text{nominal capital compensation}}{\text{nominal capital}}(y). \quad (39)$$

Equation (38) 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 Sections 2-4. Equation (39) 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. We also use $\tau_k(y_{us}) = 0$ as a normalization that pins down β .

We estimate β to equal 0.96. Figure 15 plots the value of the intertemporal parameters across countries. There are variations in the rental rates; for most of the sample, rates are very similar, except at the very top




Figure 15: Sources of rental rate

where they are higher. We see that large capital wedges τ_k are not needed 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.²⁵

We summarize the parameters that do not vary across countries in Table II, where the share parameters in the CES production functions are normalized to match relative compensation shares at the U.S. income level.²⁶

Table II: Calibrated model parameters

	Parameter	Value
<i>Preferences:</i>	β	0.96
	ϑ	0.30
	η	0.00
	ν	35.99
<i>Non-agricultural production:</i>	α	0.46
	σ	0.98
<i>Agricultural production:</i>	$\tilde{\omega}_1$	0.42
	$\tilde{\omega}_2$	0.14
	$\tilde{\omega}_3$	0.44
	σ_1	1.90
	σ_2	0.66
	σ_3	1.75

Figure 16, finally, compares other central outcomes to the data.

Some features of the GE outcomes fit the data by construction. For example, our nested CES functions obtain their elasticity and share parameters using data from the two end-points of the spectrum of countries. However, we note that the fit is remarkably good also within sample, i.e., for all intermediate values of development. The concavity of agricultural employment and the value-added share, for example, are not matched by construction. This shows that the functional forms of the commonly imposed utility and production functions do give rise to an excellent account of the cross-country data.

²⁵The depreciation rates are somewhat high 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. Had we hard-coded these external depreciation rates, we could have obtained a lower investment rate, giving us counterfactually high consumption expenditures, agricultural expenditures, land prices, and ratios of labor, capital, and intermediate inputs to land. Thus, our calibration choice on depreciation percolates through the general equilibrium system via resource constraints and price effects. We consider consistency of general equilibrium with observed outcomes, with subsequent assessment of how the parameters compare to external standards, a preferable procedure than using external standards and failing to ensure model-data consistency for all the aggregates. 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, endogenous obsolescence, or differences in the structure vs. capital mix across levels of development—that might reconcile these differences.

²⁶Formally, these parameters are obtained by using a normalized formulation in the spirit of Klump et al. (2012) where the CES is calibrated so that at the U.S. income level, relative quantities are unity, and relative prices match relative factor compensation shares. This normalization means that the share estimates are independent of measurement units; the estimated elasticities are not affected.




Figure 16: GE outcomes

Note: scales are logarithmic. Dotted lines show the model outcome and the solid lines the data. Consumption expenditure and aggregate output per worker are denoted in the amount of non-agricultural goods, with the price level of non-agricultural goods being 1 at the U.S. level of income.

8 Counterfactual experiments

We are now ready to use our calibrated model to examine a number of counterfactuals. Our primary objective is to investigate the effects of improving various technology parameters—such as TFP levels—but we also look at other country-specific factors like human capital. Before we present the results of these experiments, we briefly discuss the nature of comparative statics of our model, as its multi-sector nature and non-homothetic preferences make it somewhat more complex than a standard neoclassical growth model. The findings of our main experiments are then reported in Section 8.2. Finally, we provide a robustness analysis concerning some less central parameters that differ across countries in our calibration, such as capital depreciation rates.

8.1 The basic workings of our model

Our steady-state model is quite compact—as pointed out above, it can be reduced to one equation in the unknown land price—but the system is nonlinear and involves a nontrivial intersectoral allocation. These features make it important to draw a distinction between minor changes in parameters and major changes, such as increasing in agricultural TFP from the level of one of the poorest countries to that of the United

States. Consequently, we will differentiate between local and global elasticities, which can sometimes vary significantly, particularly when induced structural changes result in major changes in allocations across sectors. In addition, there will sometimes be important complementarities between changes in different parameters. Before examining elasticities, however, we must define the concept of aggregate output, or GDP, employed in our multi-sector economy. While this definition is somewhat arbitrary, it is designed to align as closely as possible with how output is constructed in national statistics offices.

Defining output Final demand consists of agricultural consumption c_a , non-agricultural consumption c_n , and investment δk (we assume steady-state output here). To define local changes in real GDP, we use a Divisia index

$$d \log y = s_a d \log c_a + s_n d \log c_n + s_i d \log [\delta k] \quad (40)$$

where $s_a = \frac{p_a c_a}{\tilde{y}}$, $s_n = \frac{c_n}{\tilde{y}}$, $s_i = \frac{p_k \delta k}{\tilde{y}}$ are the GDP shares of the three demand categories and where $\tilde{y} = c_n + p_a c_a + p_k \delta k$ is output in units of the non-agricultural consumption good. To extend the local definition, we note that, in any counterfactual, we can interpret the changing parameters as a function of some index $t \in [0, 1]$; e.g., as t goes from 0 to 1, A_a goes from the lowest value of this parameter observed in our cross-section to the highest. Given such an index, the outcome variables are also functions of $t \in [0, 1]$, and GDP changes satisfy

$$\log \left(\frac{y(t)}{y(0)} \right) = \int_0^t \left[s_a(t) \frac{d \log c_a(t)}{dt} + s_n(t) \frac{d \log c_n(t)}{dt} + s_i(t) \frac{d \log [\delta(t) k(t)]}{dt} \right] dt. \quad (41)$$

This method of calculating real output is consistent with how growth rates are calculated by national statistics offices that use chain weighting.²⁷

Local elasticities Beginning with local elasticities (and semi-elasticities), it is straightforward to compute the full Jacobian, i.e., the derivatives of the endogenous model variables with respect to parameters such as the TFPs; we report a subset of these in our online appendix (see Tables A1-A6). However, it is also very helpful—in particular when it comes to understanding how aggregate output is affected by a parameter—to appeal to Hulten’s theorem (Hulten (1978)), which expresses changes in output as a weighted sum of changes in parameters. Hulten’s theorem is not strictly applicable for our economy, since we have some distortions. Despite this, it turns out—from verifying that the Jacobian gives almost the same answer as Hulten’s formula—that the theorem offers a very good approximation. Thus, for the static economy and

²⁷The Penn World Table uses a different method based on various methods of averaging expenditure shares. In the online appendix, we show that our strategy yields very similar estimates in terms of GDP differences between rich and poor countries.

following Hulten’s logic, we have

$$d \log y = s_a d \log A_a + s_n d \log A_n + s_i d \log A_k A_n + s_x d \log A_x A_n + s_k d \log k + s_h d \log H + s_l d \log L, \quad (42)$$

i.e., in logs, GDP responds to all the separate TFPs, weighted by their gross output shares in aggregate GDP (their “Domar weights” s_a , s_n , s_i , and s_x), and to the basic inputs k , H , and L , weighted by their corresponding aggregate cost shares.²⁸ Thus, locally, agricultural TFP gives a big boost in poor countries, where the agricultural expenditure share is very high. For the dynamic economy, k is endogenous and it is possible to unpack the effects of changes in the TFPs and other parameters on steady-state output through how they affect steady-state $d \log k$ (Malmberg, 2023). As we will see below, these effects can significantly change the static results, and the dynamic effects differ by which parameter is changed.

Global elasticities Our focus is ultimately on global changes in parameters. For example, what if agricultural TFP could be raised from that of the poorest country to that of the richest, thus addressing the “food problem” directly? What if, instead, the ability to produce non-agricultural goods is improved by an increase in A_n from the very lowest value observed in the cross-section of countries to the very highest? Many of these changes will trigger structural change: on the demand side due to income and relative price effects, and on the supply side due to changes in relative factor prices and consequent changes in factor input mixes.

For illustration, consider an increase in agricultural TFP A_a . For a 1 percent increase, the local GDP response will be 0.55 percent (see Table A1). This includes a static effect of a bit more than 0.3 percent, reflecting agriculture’s Domar weight, as well as a steady-state capital response as labor is reallocated towards the more capital-intensive non-agricultural sector. However, a full log-point change in A_a only delivers a little less than a 0.3 log-point increase in GDP. The global effect is smaller than the local effect since an increase in A_a pushes down the agricultural price p_a . Given the inelastic demand for agricultural consumption (recall that we calibrate its price elasticity to be zero), this primarily boosts non-agricultural consumption. As a result, agriculture’s Domar weight falls, and it falls sharply; after a full log-point change in A_a , agriculture’s Domar weight falls by more than half. The structural transformation away from agriculture, whose importance falls as A_a rises, does not ultimately raise output much, however, as the productivity of the non-agricultural sector is unaffected.

²⁸The Domar weights, which sum to more than 1, are defined as: $s_a = p_a c_a / \tilde{y}$; $s_n = c_n / \tilde{y}$; $s_i = i / (\tilde{y} A_k)$, where i is the part of non-agricultural output allocated to investment; and $s_x = x / (\tilde{y} A_x)$. The factor weights, which do sum to 1, are defined as: $s_k = rk / \tilde{y}$, $s_h = wH / \tilde{y}$, and $s_l = p_l L / \tilde{y}$. Note that, for both sets of weights, $\tilde{y} = rk + wH + p_l L$.

Below, we will discuss our counterfactuals based on numerical analysis of the steady-state equilibrium system (both locally, with the use of Jacobians, and globally). Closed-form results are hard to derive except in special cases. One is the case of Cobb-Douglas (homothetic) preferences, but this case is not relevant as it entirely misses the demand channel of structural transformation. The other case, which is more relevant, is the case where intermediates and land are not inputs into agricultural production; however, while this case allows for further characterization, it does not admit full closed-form solutions either. For the numerical analysis, we present in Table III the Domar weights and cost shares evaluated at the lowest level of development—as our experiments will focus on improving the overall productivity of the economy through various alternative means.

Table III: Domar weights and cost shares at the lowest development level

s_a	s_n	s_i	s_x	s_k	s_h	s_l
0.32	0.46	0.22	0.02	0.35	0.58	0.07

8.2 Changes in key country-specific parameters

We consider a sequence of experiments. In the first group of experiments, we increase TFP levels individually. We first look at agricultural and non-agricultural TFPs, A_a and A_n . When we change the latter, we do not change the TFPs of capital or intermediates production, $A_k A_n$ and $A_x A_n$, respectively, so in these experiments we adjust A_k and A_x down by the same amount we adjust A_n up. Then in the second set of experiments, we change the TFPs of capital and intermediates; here we instead keep A_n and A_a fixed. Finally, we raise human capital. In a second class of experiments, we look at joint changes in parameters to examine possible complementarities. Finally, and more as a robustness check, we look at other parameters whose values we estimate to differ across countries (such as the depreciation rates). Throughout, our experiments consider a representative economy at the lowest level of development and then examine improvements in the parameters toward the values observed for the richest economies.

Our main finding is that, despite agriculture being a large sector in poor countries, increasing agricultural TFP is much less potent in increasing aggregate income than increasing non-agricultural TFPs, in particular that of capital. There are two main reasons for this. The first is that the gap in A_a between the poorest and richest countries is much smaller than the labor-productivity gap: it is roughly a factor 6 (as opposed to approximately 120 for labor productivity, and 90 for output per labor efficiency unit). The second reason is a nonlinearity not captured by Hulten’s theorem: structural transformation has the feature of making agricultural TFP increases partly self-defeating. In particular, the strong complementarity between

agricultural and non-agricultural consumption means that the size of the agricultural sector falls rapidly as its TFP improves, i.e., as A_a rises, it applies to a smaller and smaller part of the economy.

Increasing either agricultural or non-agricultural TFPs When we increase A_n , we decrease A_k and A_x one-for-one so as to keep $A_k A_n$ and $A_x A_n$ constant. That is, A_n corresponds to the TFP for non-agricultural *consumption goods* only. A one-percent change in A_n raises output slightly more in a static sense than does a one-percent increase in A_a : from Table III and the Hulten logic, $s_n = 0.46 > 0.32 = s_a$. However, the steady-state impact—still as a comparative-static, local elasticity, following the Hulten formula as given in equation (42)—is reversed: output now rises more when A_a is increased than when A_n is increased, 0.55 vs. 0.39 (see Table A1 for the numbers). This is because steady-state capital *decreases* when A_n is increased and *increases* when A_a is increased. A higher A_n raises income and the relative price of agriculture, and agriculture’s share of expenditures rises. This will move factors of production—labor, capital, and intermediate goods—into agriculture. This sectoral reallocation affects output precisely because the sectors have very different production technologies; in particular, in poor countries, the capital-labor ratio is much lower in the agricultural sector than in non-agriculture, by about two log-points. Given the reallocation, total capital falls. To see this, note that $k = (k_n/h_n)h_n + (k_a/h_a)h_a$ falls since h_n falls by the same amount h_a rises, while k_n/h_n and k_a/h_a are unchanged—the former is pinned down from the Euler equation by other basic parameters (not including either A_a or A_n) and the latter is pinned down by r/w (itself unchanged, as capital and labor forms a sub-nest in the agricultural production function).²⁹ Thus, the bottom line is that Schultz’s argument is very much borne out: raising TFP for agriculture (i.e., solving the food problem) is more powerful than raising TFP for non-agricultural consumption goods, with the former having a steady-state GDP elasticity almost twice that of the latter.

As for the global elasticities, changes in shares will diminish the quantitative impacts: as A_n rises, the Domar weight on non-agricultural goods falls, as there is reallocation of spending toward agriculture (the reverse of the case where agricultural TFP is increased). To see the full effect of increasing non-agricultural TFP all the way to the level in the U.S., consider Figure 17.

Panel A in the upper-left corner describes the experiment, with A_n rising, in a somewhat convex manner as a function of GDP, by a factor 5 (all axes are measured on logarithmic scales). As for factor prices, as mentioned above, r is unaffected, thus pinning down the capital-labor ratio in the non-agricultural sector and hence also pinning down the wage. The price of intermediates, $1/A_x$, relative to the wage is also

²⁹To account for the numerical values, take the case of A_n : its full output elasticity is 0.07 lower taking capital accumulation into account. This value can be computed approximately, using Hulten’s formula, as $s_k \times d \log(\delta k) / d \log A_n = 0.35 \times (-0.21)$. Note that this is almost exact, despite of the inefficiencies in our economy.




Figure 17: Effects of increasing TFP for non-agricultural consumption

unaffected. Panel B, thus, only shows the relative price of land rising somewhat. This reflects the increase in the demand for agricultural goods. Panel C shows that the increase in labor input in agriculture is the dominant movement in terms of changing its input ratios and that agricultural labor productivity falls due to land being a fixed factor. Panel D shows how the increased demand for agricultural goods raises its price,

along with total consumption expenditures $e = c_n + p_a c_a$, all of whose components rise when A_n rises. Panel E shows that structural change occurs, but in the “wrong” direction. Finally, Panel F illustrates how the effect of moving A_n up to the U.S. level has a very, very limited impact on GDP, with virtually no impact after the initial boost; in the end, the 5-fold TFP increase only raises GDP by a mere 44%, or 0.37 log points.




Figure 18: Effects of increasing agricultural TFP

Turning to agricultural TFP, consider Figure 18. As was noted earlier in the paper, agricultural TFP in the poorest economy is approximately one-sixth that in the richest (whereas for agricultural labor productivity the factor is closer to 120 for output per worker, and 90 for output per efficiency unit). Here, in Panel A, the effect is (log-)linear across the whole development range. As seen in Panel B, making agricultural goods much cheaper implies that the relative input price of land compared to labor drops precipitously, to only 3% of its initial value. This decline occurs even though agriculture gets more productive, and is driven by the significant movement of labor and capital out of agriculture, which lowers the marginal value of land. Panel C illustrates this, where we see that the land-labor ratio (and labor productivity) both rise by a factor of around 8. This is due to a demand effect: agricultural consumption demand is insensitive to its price drop, and hence this price drop spills over one-for-one to non-agricultural consumption goods. Since the two consumption shares are roughly equal in the poorest economy, a massive drop in p_a implies a massive boost to c_n , and hence a massive movement of capital and labor into non-agricultural consumption. Indeed, the price drop in p_a is massive: as can be seen in Panel D, the price falls to roughly 1/10th of its initial value. The effect of TFP on the agricultural price is more than one-for-one—a factor of 10 versus the factor 6 boost to TFP—because the land price also falls, which further lowers unit costs.³⁰ Panel D also confirms that total consumption expenditures remain roughly unchanged: c_a is almost unchanged while its price plummets, c_n rises strongly, and all these effects essentially cancel each other out. Panel E shows significant structural change in the opposite direction as in the previous counterfactual—now moving shares in the “right” direction—but again without a huge boost to output (Panel F shows that output goes up, with almost exactly the same as that generated by the rise in the TFP for non-agricultural consumption).

In sum, neither of our two counterfactuals help output appreciably. They are, in particular, both ultimately self-defeating: each parameter systematically decreases the importance of the channel through which it operates, as an increase in A_a decreases agriculture’s share s_a while an increase in A_n likewise decreases the share of non-agricultural consumption goods. It is, instead, input intensification (in capital and intermediates) that proves to be the more central driver of development. Below, we also show an interesting complementarity between A_a and A_n : the self-defeating forces are significantly diminished if these factors are moved up together.

³⁰To reconcile the numbers, note that land has close to unitary elasticity with the remaining inputs in agriculture, so that the unit cost formula approximately has p_a equal to (a constant, given that the other input prices are unaffected in this experiment, times) the land price to the power of land’s cost share, divided by A_a . This is consistent with a fall in p_a by roughly a factor of 8.

Increasing TFP for the production of capital or of intermediates We now look at the effects of (separately) changing A_k and A_x , thus, improving the TFPs of capital and intermediate goods production without a change in A_n (see Table A2). Beginning with the local output elasticities, we note from Table III above that $s_i = 0.22$ and $s_x = 0.02$, thus giving both capital and (especially) intermediates a much smaller static impact on output than for changes in A_n and A_a . In the case of capital production, however, the steady-state impact is very strong, a consequence of capital accumulation, with aggregate capital rising almost 2 percent when A_k rises by 1 percent. As a result, the local steady-state impact on output of higher TFP in capital production is 0.88 (compared to 0.55 and 0.39 for improvements in A_n and A_a , respectively). As for intermediates, the steady-state effect is mild: 0.05 percent (the doubling of the static effect reflects that more productive intermediates increases capital intensity by moving people out of agriculture).

Figures 19–20 summarize the global impacts and model mechanisms through which A_k and A_x affect the economy. As a general point, there are no strong nonlinearities on the aggregate output level in either of these cases: the structural change that is induced is relatively mild (the gross output share of agriculture does change in both cases but not as dramatically as in the previous experiments) and the effects on output inherit the shape of the TFP change itself. In contrast, there are clear nonlinearities within agriculture: as A_k and A_x increase, the cost shares of these factors within agriculture increase as well, making the local elasticity increase. As a consequence, the effects on agricultural labor productivity accelerate.³¹

Beginning with capital, we now see effects on all the relative prices of inputs: an increase in A_k lowers r one for one, raising the capital-labor ratio in non-agricultural production so that w rises (and therefore r/w falls). As a consequence of a lower r/w , the capital-labor ratio in agricultural production is raised. In fact, the capital intensity is going up more in agriculture. Thus, we see a strong force of mechanization in agriculture as a result of cheaper capital. As capital flows into agricultural production, land demand goes up, so p_l rises. More importantly, now, labor flows out of agriculture: capital is a good substitute for labor in the agricultural sector, and hence labor moves into non-agricultural production. Expenditures on both consumption goods rise and the price of agricultural goods rises. We note a shrinking share of expenditures going toward agricultural goods but the effect is not massive. Finally, we see very sizable impacts on output in panel F: more than half the gap between the poorest and the richest countries' GDP levels is closed by the improvement of TFP in capital production.

As for the production of intermediates, Figure 20 shows milder effects overall; intermediates are important,

³¹By comparing tables A2 and A5, we see that increasing A_k by 1 percent increases agricultural labor productivity in a poor country by 0.23 percent whereas this effect is 1.56 in a country at the U.S. level of development. Similarly, the effect of a 1 percent increase in A_x on agricultural labor productivity is 0.15 percent for a poor country versus 0.75 percent at the U.S. level of development.




Figure 19: Effects of the increasing TFP for capital production

but their Domar weight is small, so the effects on aggregate GDP are limited. Labor productivity in agriculture rises and the price of agricultural good falls. There is, as a result, a significant impact on the share of agricultural goods in gross expenditures (more so than when capital TFP was improved). The final effect on output also works through capital accumulation, roughly doubling the static effect. It should be noted that, though this is not visible in the figures, the cost share of intermediates in agriculture goes up significantly as A_x rises: when intermediates get cheaper, their share goes up, since they are highly substitutable with capital and labor.




Figure 20: Effects of increasing TFP for intermediates production

In sum, we see that the output effects of changes in capital-producing TFP are substantial (and the effects of intermediates' TFP less so). However, and more importantly from our perspective, we see that in the transformation of agriculture, both these specific TFP factors are critical. In panel C of the two figures, we see major factor intensification and, together, this factor intensification allows the structural transformation of agriculture to take place, thus partially accounting for the enormous difference in agricultural labor productivity between poor and rich economies.

Human capital In one-sector neoclassical models where the aggregate production function is constant returns to scale in capital and labor, an increase in the stock of human capital acts, just like in our model, through a static effect, keeping the physical capital stock fixed, and a steady-state effect, where physical capital typically adjusts one for one with human capital (so as to maintain a constant ratio of the two inputs). As a result, a doubling of the stock of human capital exactly doubles long-run output. In our economy, a similar logic of course applies, but we have two sectors, with decreasing returns (due to a fixed stock of land) in one sector. The local elasticity can, however, still be gleaned from the Hulten formula, equation (42), which says that (i) the static impact is given by the aggregate cost share of labor (s_h , which is not a parameter, unlike in the Cobb-Douglas case) but (ii) there are in addition induced effects too on steady-state output, through capital accumulation, given by the percentage change in capital times capital's cost share, s_k . Locally, thus, the approximate elasticity given by the Hulten formula is $0.58 + 0.35 \times 1.33 \approx 1.05$, very close to the 1.04 found using the Jacobian method. Note that capital increases more than one-for-one with human capital increases, reflecting that income effects induce some structural change towards the capital-intensive non-agricultural sector.

Figure 21 also shows that the global results are quite similar to the local ones, with effects very close to log-linear. Over the whole range of development, human capital rises by 0.61 log points; output rises by 0.63 log points. As the economy gets richer due to higher human capital, the share of expenditures spent on agricultural goods declines.

Combined effects As we have seen, our model has some important nonlinearities: local and global effects are, in some cases, very different. This is largely due to structural change, where in particular Domar weights—which are key for local effects—change appreciably as TFP parameters are varied. It turns out that nonlinear interactions between different TFP parameters do not appear to be major, except in one case: when A_n and A_a are raised at the same time. In particular, as we report in Appendix C.4, Table A7, a simultaneous log-point increase in these parameters give an extra output boost: the total effect on output (in log points) is 0.85, whereas it is 0.29 if only A_a is changed and 0.30 if only A_n is changed, i.e., the extra effect is 0.25 log points (28 percent). To understand this result, let us use a much simpler, static, structure that brings out the complementarities.

So suppose we keep our calibrated preferences, so that $c_a = \nu e^{0.30}$ (i.e., there are no price effects at all) and $c_n = e - p_a \nu e^{0.30}$. However, let us simplify production drastically and have production be linear in labor in both sectors, with labor productivities equal to A_a and A_n , respectively. From the production side, i.e., independently of preferences, this model implies $p_a = A_n/A_a$. Total consumption expenditures satisfy




Figure 21: Effects of increasing the stock of human capital

$e = w$ (given one unit of labor) and from profit maximization in the non-agricultural firm we obtain $w = A_n$. Thus, $e = A_n$ and we have a closed-form solution: $c_a = \nu A_n^{0.30}$ and $c_n = A_n - \nu A_n^{1.3}/A_a$. Moreover we have the share spent on agricultural goods given by $s_a = \nu A_n^{0.3}/A_a$. Thus, in this simple model e and p_a are log-linear in the TFPs but consumption levels are not.

Looking at how the simple model behaves, a simultaneous increase in these TFP parameters (of the same percentage amount) will decrease agriculture's Domar weight, responding more to A_a than to A_n ; given that they change in proportion, short of a constant we can define their constant ratio v via $v A_n \equiv A_a$ and

conclude that s_a is proportional to $A_n^{-0.7}$. Consumption of agricultural goods will respond exactly like in the case of a change in A_n only, and will thus not be a weighted average of the change under A_n and A_a separately (in the latter case, the change is zero). The consumption of non-agricultural goods will also increase and is now convex in TFP: $c_n = A_n - (\nu/v)A_n^{0.3}$. Overall, this simple model turns out to deliver close to the numerical results for global experiments in our full model, save for dynamic effects (which are also nonlinear; capital accumulation is boosted from a joint TFP increase).

Robustness and summary For completeness, we investigate how the remaining parameters that differ across countries— τ_a , τ_k , δ , and L —affect aggregate output in isolation; recall that the first three of these parameters were chosen so as to make sure our model variables match all their corresponding values in national-accounts statistics by construction.³² Table A14 in our online appendix shows the results from this robustness exercise. It reveals that if all these variables were to be changed jointly, from their levels in our poorest economy all the way to their levels in the richest economy, output would decrease by around 0.3 log points (i.e., 35 percent).³³ There would be changes to all endogenous variables (input prices, consumption levels, etc.), but these changes are not enough to be a central driver of income differences across the development spectrum.

In conclusion, our counterfactual exercises point very strongly to capital deepening and input intensification as the key to overall development. Thus, catchup in total-factor productivity for producing either agricultural consumption goods or non-agricultural is helpful, but cannot close the output gap between rich and poor more than very marginally; each of these changes are self-defeating since they generate structural change away from the good whose productivity is increased. If these TFPs are increased jointly, the self-defeating nature is to a large extent canceled, but the effect on output is still rather small.

9 An application: adaptation 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,

³²As for L , Section 9 carefully examines how significant decreases in this variable affects result, since such changes can be expected due to global warming.

³³Recall that we estimate depreciation rates to be *lower* for poor than for rich countries.

an area projected to experience rapid population growth over the next several decades, roughly doubling its population 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 Jr (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—so the purpose is more to indicate a fruitful way forward for research in this area, where one 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, in the Intergovernmental Panel on Climate Change (IPCC) projects on food supply, they use various computational general equilibrium (CGE) models. One important such model is the IMAGE integrated assessment model (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.³⁴ 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 (real GDP per worker of USD 2,000) and consider the effect of reducing its land endowment by up to 75%. Figure 22 shows results of our experiment.

We see from the figure how the land and agricultural price (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) depends 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

³⁴The MAGNET default elasticities are taken from the food column 6.5.5 in Woltjer et al. (2014).




Figure 22: The effects of reduced land supply

is more muted since the land share is only 20%–30%. The second panel shows that as land decreases, more capital, intermediates, and labor are moved into the sector, reflecting that a higher agricultural price makes it profitable to use other production factors. The third panel shows how factor intensification implies that agricultural output is actually quite steady while non-agricultural consumption falls. Intuitively, inelastic demand and strong income effects mean that the agricultural sector is sustained by the price increase, and this is occurring at the expense of the non-agricultural sector. Total output only falls 35%.

Thus, overall, our conclusion is that input substitutability does play a major role and appears promising as an adaptation mechanism.³⁵ Clearly, it must be emphasized that our analysis concerns the long run and not short-run adjustments, which of course can be much more difficult, since it appears reasonable to assume lower substitution elasticities across inputs for short-run analysis. 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, we find that there is scope for significant learning by additional work in this area.

³⁵Based on the notion that mere population growth influences land per capita, we note that between the early 1990s and 2015, the African continent saw a decrease in land per capita of over 40%. We observe, over this period, that the agricultural labor input per unit of land rose by over 40% and that agricultural productivity rose by around 10% and these observations are at least qualitatively consistent with our proposed mechanisms.

10 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, not received the attention they deserve in the literature. The nature of these facts suggests that a framework based on an aggregate production function can be quite valuable. 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 capital deepening and 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. Alone, agricultural TFP growth would generate strong counteracting price and income effects: consumers would shift their expenditures away from agricultural goods as these become cheaper 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. Therefore, a “development Marshall plan”, with a focus on building non-agricultural productivity, would both help lift the overall economies and be vital to agricultural productivity improvements.

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 can help us understand the broad patterns of the development. 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: we must, sooner or later, ask of the micro-development facts to reproduce the aggregate patterns that we uncover here. A highly 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.

We must also emphasize that, because the data we use are far from perfect, we hope that future studies will offer updates that can refine our findings. We do not think that mismeasurement errors will be so large and systematic across the development dimension that our broad findings are substantially altered, but better measurement is clearly also high on our wish list.

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 II (2001), León-Ledesma and Satchi (2019), and Hassler et al. (2021). Our approach also relate to an older tradition that has treated technology choices and institutional setups as endogenous to relative factor prices (Boserup et al., 1965; Ruttan and Hayami, 1984). Spelling out 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 TFPs growing. 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 and different capital intensities across sectors are not consistent with exact balanced growth. It is possible to extend our analysis so that we describe countries at different stages in the path toward this asymptotically balanced path (potentially using a medium-run type analysis as the one used in Buera et al. (2020)). Such an exercise 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 the last part of our paper—that involving counterfactual exercises—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. As land is a fixed input entering agricultural production, the exploitation of comparative advantages in agriculture is naturally limited (unlike for a comparative advantage in the non-agricultural sector). 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

- Acemoglu, D. (2002). Directed technical change. *The review of economic studies*, 69(4):781–809.
- Adamopoulos, T. and Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6):1667–97.
- Adamopoulos, T. and Restuccia, D. (2022). Geography and agricultural productivity: Cross-country evidence from micro plot-level data. *The Review of Economic Studies*, 89(4):1629–1653.
- Alder, S., Boppart, T., and Muller, A. (2022). A theory of structural change that can fit the data. *American Economic Journal: Macroeconomics*, 14(2):160–206.
- Alvarez-Cuadrado, F., Van Long, N., and Poschke, M. (2017). Capital–labor substitution, structural change, and growth. *Theoretical Economics*, 12(3):1229–1266.
- Barro, R. J. and Lee, J. W. (2013). A new data set of educational attainment in the world, 1950–2010. *Journal of development economics*, 104:184–198.
- Basu, S. and Weil, D. N. (1998). Appropriate technology and growth. *The Quarterly Journal of Economics*, 113(4):1025–1054.
- Bergquist, L. F., Faber, B., Fally, T., Hoelzlein, M., Miguel, E., and Rodríguez-Clare, A. (2022). Scaling agricultural policy interventions. Technical report, National Bureau of Economic Research.
- Besley, T. (1995). Property rights and investment incentives: Theory and evidence from ghana. *Journal of political Economy*, 103(5):903–937.
- Binswanger, H. P. (1974). A cost function approach to the measurement of elasticities of factor demand and elasticities of substitution. *American Journal of agricultural economics*, 56(2):377–386.
- Bold, T., Kaizzi, K. C., Svensson, J., and Yanagizawa-Drott, D. (2017). Lemon technologies and adoption: measurement, theory and evidence from agricultural markets in uganda. *The Quarterly Journal of Economics*, 132(3):1055–1100.
- Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

- Boserup, E. et al. (1965). The conditions of agricultural growth. the economics of agrarian change under population pressure. *The conditions of agricultural growth. The economics of agrarian change under population pressure*.
- Buera, F. J., Kaboski, J. P., Mestieri, M., and O'Connor, D. G. (2020). The stable transformation path. Technical report, National Bureau of Economic Research.
- Buera, F. J., Kaboski, J. P., and Townsend, R. M. (2021). From micro to macro development.
- Bustos, P., Caprettini, B., and Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from brazil. *American Economic Review*, 106(6):1320–65.
- Caselli, F. (2005). Accounting for cross-country income differences. *Handbook of economic growth*, 1:679–741.
- Caselli, F. and Coleman II, W. J. (2001). The us structural transformation and regional convergence: A reinterpretation. *Journal of political Economy*, 109(3):584–616.
- Caselli, F. and Coleman II, W. J. (2006). The world technology frontier. *American Economic Review*, 96(3):499–522.
- Caselli, F., Ponticelli, J., and Rossi, F. (2014). A new dataset on mincerian returns. *Technology differences over space and time*, pages 108–118.
- Caunedo, J. and Keller, E. (2021). Capital obsolescence and agricultural productivity. *The Quarterly Journal of Economics*, 136(1):505–561.
- Chen, C. (2020). Technology adoption, capital deepening, and international productivity differences. *Journal of Development Economics*, 143:102388.
- Chen, C., Restuccia, D., and Santaeuilà-Llopis, R. (2022a). The effects of land markets on resource allocation and agricultural productivity. *Review of Economic Dynamics*, 45:41–54.
- Chen, C., Restuccia, D., and Santaeuilà-Llopis, R. (2023). Land misallocation and productivity. *American Economic Journal: Macroeconomics*. Forthcoming.
- Chen, C., Restuccia, D., Santaeuilà-Llopis, R., et al. (2022b). Land misallocation and productivity. Technical report.
- 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.

- Costinot, A. and Donaldson, D. (2016). How large are the gains from economic integration? theory and evidence from us agriculture, 1880-1997. Technical report, National Bureau of Economic Research.
- Donovan, K. (2021). The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity. *The Review of Economic Studies*, 88(5):2275–2307.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American economic review*, 101(6):2350–90.
- Eckert, F. and Peters, M. (2022). Spatial structural change.
- Foster, A. D. and Rosenzweig, M. R. (2022). Are there too many farms in the world? labor market transaction costs, machine capacities, and optimal farm size. *Journal of Political Economy*, 130(3):636–680.
- Gollin, D., Hansen, C. W., and Wingender, A. M. (2021). Two blades of grass: The impact of the green revolution. *Journal of Political Economy*, 129(8):2344–2384.
- Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.
- Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American economic review*, 92(2):160–164.
- Gollin, D., Parente, S. L., and Rogerson, R. (2004). Farm work, home work and international productivity differences. *Review of Economic dynamics*, 7(4):827–850.
- Gollin, D., Parente, S. L., and Rogerson, R. (2007). The food problem and the evolution of international income levels. *Journal of Monetary Economics*, 54(4):1230–1255.
- Gottlieb, C. and Grobovšek, J. (2019). Communal land and agricultural productivity. *Journal of Development Economics*, 138:135–152.
- Hall, R. E. and Jones, C. I. (1999). Why do some countries produce so much more output per worker than others? *The quarterly journal of economics*, 114(1):83–116.
- Hassler, J., Krusell, P., and Olovsson, C. (2021). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, 129(11):3039–3072.
- Herrendorf, B., Herrington, C., and Valentinyi, A. (2015). Sectoral technology and structural transformation. *American Economic Journal: Macroeconomics*, 7(4):104–33.

- Herrendorf, B., Rogerson, R., and Valentinyi, A. (2013). Two perspectives on preferences and structural transformation. *American Economic Review*, 103(7):2752–2789.
- Herrendorf, B. and Schoellman, T. (2015). Why is measured productivity so low in agriculture? *Review of Economic Dynamics*, 18(4):1003–1022.
- Herrendorf, B. and Schoellman, T. (2018). Wages, human capital, and barriers to structural transformation. *American Economic Journal: Macroeconomics*, 10(2):1–23.
- Hsieh, C.-T. and Klenow, P. J. (2007). Relative prices and relative prosperity. *American Economic Review*, 97(3):562–585.
- Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 45(3):511–518.
- Jones, C. I. and Klenow, P. J. (2016). Beyond gdp? welfare across countries and time. *American Economic Review*, 106(9):2426–57.
- Klenow, P. J. and Rodriguez-Clare, A. (1997). The neoclassical revival in growth economics: Has it gone too far? *NBER macroeconomics annual*, 12:73–103.
- Klump, R., McAdam, P., and Willman, A. (2012). The normalized ces production function: theory and empirics. *Journal of Economic Surveys*, 26(5):769–799.
- Krusell, P. and Smith Jr, A. A. (2022). Climate change around the world. Technical report, National Bureau of Economic Research.
- Kuznets, S. S. (1966). *Modern economic growth*. Yale University Press.
- Lagakos, D., Mobarak, A. M., and Waugh, M. E. (2018). The welfare effects of encouraging rural-urban migration. Technical report, National Bureau of Economic Research.
- León-Ledesma, M. A. and Satchi, M. (2019). Appropriate technology and balanced growth. *The Review of Economic Studies*, 86(2):807–835.
- Leukhina, O. M. and Turnovsky, S. J. (2016). Population size effects in the structural development of england. *American Economic Journal: Macroeconomics*, 8(3):195–229.
- Lewis, W. A. et al. (1954). Economic development with unlimited supplies of labour.

- Malmberg, H. (2023). Characterizing capital amplification. Mimeo.
- Mundlak, Y. (2005). Economic growth: lessons from two centuries of american agriculture. *Journal of economic literature*, 43(4):989–1024.
- Restuccia, D., Yang, D. T., and Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of monetary economics*, 55(2):234–250.
- Rosenzweig, M. R. and Udry, C. (2014). Rainfall forecasts, weather, and wages over the agricultural production cycle. *American Economic Review*, 104(5):278–283.
- Rostow, W. (1960). The stages of economic growth: A non-communist manifesto. *Cambridge University Press*.
- Ruttan, V. W. and Hayami, Y. (1984). Toward a theory of induced institutional innovation. *The Journal of development studies*, 20(4):203–223.
- Sato, K. (1976). The meaning and measurement of the real value added index. *The Review of economics and Statistics*, pages 434–442.
- Schultz, T. W. (1953). *The economic organization of agriculture*. McGraw-Hill Book. Co.
- Solow, R. M. (1956). A contribution to the theory of economic growth. *The quarterly journal of economics*, 70(1):65–94.
- Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, pages 312–320.
- 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).
- Storesletten, K., Zhao, B., and Zilibotti, F. (2019). Business cycle during structural change: Arthur lewis’ theory from a neoclassical perspective. Technical report, National Bureau of Economic Research.
- 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.