



The Effect of Agricultural Technology on the Speed of Development



Markus Eberhardt^{a,b}, Dietrich Vollrath^c

^a School of Economics, University of Nottingham, UK

^b Centre for the Study of African Economies, University of Oxford, UK

^c Department of Economics, University of Houston, USA

ARTICLE INFO

Article history:

Available online 8 April 2016

Key words:

agricultural development
technology heterogeneity
agro-climatic environment
structural change

SUMMARY

We examine heterogeneity in the elasticity of agricultural output with respect to labor. Employing panel data from 128 countries over a forty-year period we find distinct heterogeneity in the elasticity of agricultural output with respect to labor. This elasticity is lowest in countries with temperate and/or cold climate regions, and higher in countries including tropical or highland regions. This agricultural parameter determines the speed of structural change following changes in agricultural productivity or population. Calibration shows shifts in labor allocations and welfare will be 2–3 times larger in temperate regions than in tropical or highland regions.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

The agricultural sector is a significant portion of GDP and employment in most developing countries today, and historically made up a large fraction of both in currently developed countries. As such, agricultural production has been of interest to economists studying the process of development and growth. A prominent view of the importance of agriculture for development comes from Schultz's (1953) characterization of the “food problem”; until countries can produce a sufficient amount of food, labor is trapped in agriculture and they cannot begin the process of modern growth. This idea has been carried forward by others, such as Johnston and Mellor (1961), Johnston and Kilby (1975), and Mellor (1995). Gollin (2010) provides an overview of this line of thinking, which implies that increased agricultural productivity is a key to structural transformation and subsequent economic development in response to productivity shocks. McMillan, Rodrik, and Verduzco-Gallo (2014) argue that this structural transformation remains an important source of potential economic growth in sub-Saharan Africa.

In this paper we emphasize that the positive effects of agricultural productivity improvements—i.e., higher total factor productivity (TFP) or the increased use of non-labor inputs—depend crucially on agricultural technology. Specifically, we define agricultural technology as the elasticity of agricultural output with respect to labor. We first show that in a standard model of structural change this technology affects the pace of development. In particular, an economy with a low labor elasticity will experience larger shifts of labor between sectors, greater increases in agricul-

tural labor productivity, and greater increases in GDP per capita than an economy with a high elasticity after any kind of productivity increase in agriculture. The logic of this result rests on the importance of labor in the agricultural technology, which is captured by the elasticity. If the elasticity is low, then agricultural output is insensitive to the number of workers in that sector. A productivity increase makes it possible to release a large number of workers and still meet the demand for food. These freed-up workers are available to produce non-agricultural goods, raising GDP per capita. In contrast, a large elasticity implies that agricultural output is very sensitive to the number of workers. Even with a productivity increase, few workers can leave agriculture without decreasing production below what is demanded. Hence high-elasticity economies do not shift as many workers out of agriculture and are able to produce fewer additional non-agricultural goods in response. Real GDP per capita rises, but not by as much as seen in economies with a low elasticity.

Given this theoretical distinction, we then undertake two tasks, one empirical and one quantitative. The empirical task is to estimate agricultural production functions for a panel of 128 countries using annual data over a long time horizon. We adopt an empirical framework that allows for heterogeneity in technology as well as for common shocks to production and/or technology spillovers between countries (“cross-sectional dependence”). The common factor model framework employed is particularly suited to this type of analysis, where a primary concern is an unobserved TFP term (Bai, 2009; Chudik, Pesaran, & Tosetti, 2011). This empirical setup speaks directly to Matsuyama's (2009) criticism of empirical studies of structural change as being analyzed “under the false

assumption that each country offers an independent observation" (484).

The heterogeneous technology setup we allow in our empirical analysis takes into account Hayami and Ruttan's (1970) claim that technology differences between countries are likely to be substantial. Subsequent research—including Hayami and Ruttan's (1985) own work—has tended to ignore this claim and work with homogenous technology parameters.¹ The closest studies to our work here are Mundlak, Butzer, and Larson (2012) and Eberhardt and Teal (2013a, 2013b). These papers show that common technology across countries is rejected by the data, although technology parameters appear to be constant over time within countries. While those papers focus primarily on *establishing* the presence of technology heterogeneity, in this study we explore the agro-climatic patterns of technology heterogeneity and their implications for development.

Our results indicate that while on average the labor elasticity of agriculture (our measure of technology) is about 0.3, there is wide variation in this value based on climate type. Countries with primarily temperate climate have elasticities of approximately 0.15, while tropical and highland countries have an elasticity that ranges from 0.35 to 0.55. The relationship of the elasticity and climate zone is not an artifact of development levels, as the patterns of our country-specific estimates of the labor elasticities do not align systematically with income per capita.

The second, quantitative, task we pursue is to assess the role that agricultural technology plays in structural change and development. We calibrate our standard two-sector model of labor allocations between agriculture and non-agriculture. This type of model has been used by Duarte and Restuccia (2010), Restuccia, Yang, and Zhu (2008), and Gollin, Parente, and Rogerson (2007) among others to quantitatively study structural change and development.² We calibrate the model using data from South Korea during 1963–2005. In this period South Korea experienced a significant shift of labor out of agriculture, as well as major increases in labor productivity in both agriculture and non-agriculture. It therefore serves as a useful benchmark for evaluating the effects of productivity increases on sectoral shifts in developing countries.

Using the calibrated model we examine the effect of an increase in agricultural total factor productivity (TFP) on sectoral labor shares, consumption levels, and real GDP per capita for developing economies that begin with 80% of their workforce in agriculture. The model economies we examine are identical in their initial labor shares and consumption levels. With a labor elasticity in agriculture of 0.15, such as in temperate regions, a 20% increase in agricultural TFP reduces the agricultural labor share to under 40%, more than doubles real labor productivity in agriculture, and increases real GDP by about 50%. In contrast, an economy with a labor elasticity in agriculture of 0.55, such as in equatorial or highland regions, the same 20% increase in agricultural TFP only reduces agricultural labor to 60%, increases labor productivity in agriculture by only 37%, and increases real GDP per capita by only 22%. This is a 2–3-fold advantage for the low-elasticity economy, even though the improvement to agricultural productivity was identical.

The situation is not universally favorable to countries with low labor elasticities in agriculture, though. Just as they are able to move large amounts of labor *out* of agriculture when productivity rises, they are also forced to move large amounts of labor *in* to agri-

culture when population rises, which acts like a decline in productivity. A 5% increase in population will take the agricultural labor share from 80% to 94% in an economy with a labor elasticity of 0.15, while the same increase in population will only raise the labor share from 80% to 82% when the elasticity is 0.55. The population increase drops real GDP per capita by 15% when the elasticity is 0.15, but only by 2.5% if the elasticity is 0.55.

Agricultural technology therefore does not provide a unique indicator of development possibilities. What is crucial is the gain in agricultural productivity relative to population growth. So long as the productivity increases out-run population growth, a low-elasticity economy will develop faster than a high-elasticity one. To understand just what a dramatic difference this can make over the long run, we feed the observed changes in agricultural productivity, non-agricultural productivity, and population growth from South Korea 1963–2005 through the model using different values of the labor elasticity. For an economy with an elasticity of 0.15, this reduces the agricultural labor share from 80% to less than 5%, while raising GDP per capita by a factor of 4.6. In an economy with an elasticity of 0.55, the labor share drops from 80% to 13%, while GDP per capita only rises by a factor of 3.9. Differing solely in their agricultural technology, the economies end up with a significant difference in living standards despite having access to the same productivity improvements and experiencing the same population growth rate.

These results suggest that a "one-size-fits-all" approach to thinking about agricultural productivity and structural change is not warranted. There are significant differences in agricultural technology across countries, and these differences determine the degree to which an economy can take advantage of productivity improvements or suffer from population increases. Slow migration rates from rural to urban areas (de Brauw, Mueller, & Lee, 2014) and the subsequent inability to exploit urban agglomeration effects (Dorosh & Thurlow, 2014; Tiffen & Mary, 2003) within Africa may be the result of tropical agriculture's relatively high labor elasticity. While agricultural productivity improvements may be necessary to drive economic growth (Collier & Dercon, 2014; Diao, Hazell, & Thurlow, 2010), our results show that tropical agriculture will require a greater scale of improvements compared to temperate countries. We cannot expect that all developing countries will respond similarly to productivity-enhancing policies in agriculture, and the past experience of temperate countries may not be a useful benchmark for tropical countries today.

The paper proceeds as follows. In the next section we introduce our simple dual economy model and derive our theoretical results. Section 3 discusses the data, empirical setup, and regression results. Having established the variation in the labor elasticity from these estimation, we then present our calibration and counterfactual exercises in Section 4. The final section concludes.

2. Agricultural technology and structural change

To understand the impact of agricultural technology in contrast to agricultural productivity we use a simple model of the process of structural change and development. There are two sectors: agriculture and non-agriculture. Individuals face a subsistence constraint for agricultural goods that makes the income elasticity less than one, and they are endowed with some units of non-agricultural goods that ensure the income elasticity of these goods is greater than one. The model shares its structure with that found in recent work by Alvarez-Cuadrado and Poschke (2011), Duarte and Restuccia (2010), Herrendorf, Rogerson, and Valentinyi (2014). Within the model, agriculture is produced using land and labor, and the agricultural technology is captured by the elasticity of agricultural output with respect to labor. The model will show that the

¹ Examples include Craig, Pardey, and Roseboom (1997), Martin and Mitra (2002), and Mundlak (2000). Gutierrez and Gutierrez (2003) and Wiebe, Soule, Narrod, and Breneman (2003) all split their samples by geographical region, generally confirming different agricultural technology across regions.

² This literature almost invariably uses the assumption that agricultural technology is homogenous across countries, with Vollrath (2011) as one exception.

response of the agricultural labor share to a shock in productivity depends on agricultural technology. In particular, economies with low labor elasticities will experience more rapid structural change than economies with high elasticities.

(a) Production

The production function for agriculture is

$$Y_a = AL_a^{\beta_L}, \quad (1)$$

where L_a is labor employed in agriculture, and β_L is the elasticity of output with respect to that labor. This parameter will be the focus of our analysis, and we will show below in the empirical section that this parameter varies widely across different types of crop systems.³

Production in the non-agricultural sector is assumed to be linear in labor, L_n , with

$$Y_n = wL_n, \quad (2)$$

where w is labor productivity in that sector.⁴ Finally, an adding-up constraint holds for total workers L such that

$$L = L_a + L_n. \quad (3)$$

(b) Preferences and individual optimization

There are L individuals in this economy with preferences of

$$U = \alpha \ln(c_a - \bar{c}_a) + (1 - \alpha) \ln(c_n + \bar{c}_n), \quad (4)$$

where c_a and c_n are consumption of agricultural goods and non-agricultural goods, respectively. The terms \bar{c}_a and \bar{c}_n are subsistence/endowment terms. By including them we will be able to capture the regularities in the data regarding expenditure shares on agriculture and non-agriculture. The subsistence constraint for food will ensure that the income elasticity is less than one, while the endowment of non-agricultural goods makes the income elasticity for those goods greater than one.

Income will be made up solely of a wage, w , earned by individuals from providing labor. Their budget constraint is thus

$$w = p_a c_a + c_n, \quad (5)$$

where p_a is the price of food relative to non-agricultural goods, and non-agricultural goods are the numeraire. Utility optimization given this constraint is straightforward, and leads to the following expressions for expenditures on the two goods

$$\begin{aligned} p_a c_a &= \alpha(w - p_a \bar{c}_a + \bar{c}_n) + p_a \bar{c}_a \\ c_n &= (1 - \alpha)(w - p_a \bar{c}_a + \bar{c}_n) - \bar{c}_n. \end{aligned} \quad (6)$$

Individuals will spend a fraction of their surplus income (given in the parentheses) equal to the weight in the utility function on the two goods. This is modified by the amount they must additionally spend on subsistence agricultural goods (in the first equation), and the lower amount they must spend on non-agricultural goods (in the second). To be clear, the value of c_a is consumption of pur-

chased agricultural goods, although only $c_a - \bar{c}_a$ of that consumption provides utility. Similarly, the value of c_n is consumption of purchased non-agricultural goods, while $c_n + \bar{c}_n$ is the total amount of consumption of non-agricultural goods that provide utility.

(c) Equilibrium allocations of labor

We assume that agricultural productivity is sufficiently high that the economy can meet the subsistence constraint for all individuals. Specifically, we presume that

$$A > \bar{c}_a L^{1-\beta_L}. \quad (7)$$

Thus, if all workers are engaged in agriculture, they produce enough to provide more than \bar{c}_a to each individual. This assumption simply ensures that there is a meaningful allocation of workers to agriculture, or $L_a < L$.

We further assume that labor moves freely between sectors to equalize earnings. In agriculture we assume that wages are equal to the average product of labor, implying that there are no rents to land-owners or the owners of other factors. If there were perfect factor markets, the wage in agriculture would be $p_a \beta_L Y_a / L_a$, and the size of β_L would influence the allocation of labor through wage equalization as well as through the shape of the production function. The assumption that labor earns its average product in agriculture removes the direct effect of labor's share of output on labor allocation, allowing us to focus solely on the effect of β_L on the curvature of the production function. Hence we have that

$$w = p_a \frac{Y_a}{L_a}. \quad (8)$$

Finally, to solve for the labor allocations we need that supply and demand are equal in both sectors, or

$$\begin{aligned} c_a L &= Y_a \\ c_n L &= Y_n. \end{aligned} \quad (9)$$

Given the static nature of the model it is straightforward to determine an equilibrium:

Equilibrium: Given productivity levels A and w , total population L , a value of β_L , and subsistence terms \bar{c}_a and \bar{c}_n , an equilibrium in the model consists of allocations of labor L_a and L_n , and a price of agricultural goods relative to non-agriculture (p_a), such that expenditures are optimal as in (6), the wage is equalized across sectors as in (8), and all labor is utilized as in (3).

The comparative statics in equilibrium are straightforward. The model is a simplified version of the ones found in Duarte and Restuccia (2010) and Alvarez-Cuadrado and Poschke (2011). As such, all the standard results follow. In particular, it is straightforward to show that in response to an increase in agricultural productivity, A , the following changes take place,

- Agricultural labor declines: $\frac{\partial L_a}{\partial A} \frac{A}{L_a} < 0$
- Agricultural consumption increases: $\frac{\partial c_a}{\partial A} \frac{A}{c_a} > 0$
- Agricultural labor productivity rises: $\frac{\partial Y_a/L_a}{\partial A} \frac{A}{Y_a/L_a} > 0$
- Relative price of agriculture falls: $\frac{\partial p_a}{\partial A} \frac{A}{p_a} < 0$

If we instead considered an increase in population, L , then this acts essentially as a reduction in productivity, and the predicted responses would simply flip sign.

Agricultural technology (i.e., the value of β_L) does not change these qualitative predictions. However, β_L does have a distinct effect on the *magnitude* of the changes following a change in productivity. The following proposition establishes this formally.

³ The productivity term A captures both total factor productivity and the effects of any other inputs used in agriculture. Assuming that the production function is Cobb-Douglas over labor and a set of k other inputs, then we could write A as $A = TFP_a \prod_{i=1}^k X_i^{\beta_i}$. If we further assumed constant returns to scale, as will be indicated in the empirical work, then we are assuming that $\sum_{i=1}^k \beta_i = 1 - \beta_L$.

⁴ The assumption that non-agricultural output is linear with respect to labor is not crucial. One can think of our production function for Y_n as representing a more general Cobb-Douglas case with $Y_n = TFP_n K^\alpha L_n^{1-\alpha}$, but where either (a) the economy is open to capital flows and so the K/Y_n ratio is fixed by the world rate of return on capital, or (b) the economy is in steady state where the K/Y_n ratio is fixed by the time preferences of individuals in the economy. In either case, $w = TFP_n (K/Y_n)^{\alpha/(1-\alpha)}$.

Proposition 1. *The absolute values of the elasticities of the responses to an increase in agricultural productivity, A , are decreasing in the size of β_L . In particular,*

- $\left| \frac{\partial L_a}{\partial A} \frac{A}{L_a} \right|$ falls as β_L rises
- $\left| \frac{\partial c_a}{\partial A} \frac{A}{c_a} \right|$ falls as β_L rises
- $\left| \frac{\partial Y_a/L_a}{\partial A} \frac{A}{Y_a/L_a} \right|$ falls as β_L rises
- $\left| \frac{\partial p_a}{\partial A} \frac{A}{p_a} \right|$ falls as β_L rises

Proof. See Appendix. \square

When β_L is large, the effect of a productivity shock is muted compared to when β_L is small. Given our empirical results, this implies that temperate areas (low β_L) will experience faster structural change and increases in labor productivity in response to productivity improvements when compared to equatorial or highland zones (high β_L).

What is driving this result? The value of β_L dictates how the labor productivity in agriculture is related to the number of agricultural workers. Large values of β_L imply that the labor productivity is relatively insensitive to the number of workers. Consider the extremes. If $\beta_L = 1$, then labor productivity is simply equal to A . If A rises, then labor productivity rises by the same proportion, allowing some labor to be released to non-agriculture while still meeting demand for food. In contrast, consider the situation when β_L approaches zero, and labor productivity approaches A/L_a . Now, when A rises labor productivity rises directly, but also rises due to the decline in L_a . This extra boost to labor productivity means that even fewer individuals are required in the agriculture sector to meet demand.

It is easiest to see the mechanism at work by limiting individuals to consuming precisely \bar{c}_a at all times. As will be seen in the simulations, the implied changes in c_a are very small, and hence this is not a particularly strong assumption. With everyone consuming precisely \bar{c}_a in agricultural goods, setting $\bar{c}_a L = Y_a$ yields

$$\frac{L_a}{L} = \left(\frac{\bar{c}_a L^{1-\beta_L}}{A} \right)^{1/\beta_L}, \quad (10)$$

as a solution. Note that the effect of a change in A is to lower L_a/L , regardless of the value of β_L . However, the elasticity of this effect is $1/\beta_L$: as β_L falls, this elasticity becomes larger.⁵

It is important to note here that the results in Proposition 1 refer to the elasticity of changes with respect to changes in A . They do not imply anything about the absolute level of L_a/L , Y_a/L_a , or other quantities. The level of A still dictates the distribution of labor across sectors, and it is quite possible for a country with a large value of β_L to also have a large value for A .

(d) The open economy

Thus far we have presumed a closed economy. If the economy is open to trade, then the qualitative predictions of the model change, as higher agricultural productivity would lead to *more* workers in agriculture. However, the quantitative change in L_a would still depend on β_L . To see this, note that for a given world price, p_a^* , using (8) the equilibrium allocation of labor to agriculture is

$$L_a = \left(\frac{p_a^* A}{w} \right)^{1/(1-\beta_L)}, \quad (11)$$

⁵ The elasticity with respect to A is made holding L constant. One can also see from Eqn. (10) that the elasticity of L_a/L with respect to L falls with β_L .

regardless of the utility function for agriculture and non-agriculture.

Examining this allocation rule, we see that the effect of productivity and/or world price changes depend on β_L . The elasticity of L_a with respect to A is $1/(1 - \beta_L) > 0$, which is increasing in β_L . In our setting, this shift of labor has no ultimate effect on real output per worker, as that is pinned down by w .⁶ Despite the change in the sign of the effect of A on the labor allocation, the strength of the effect remains dependent on β_L . In the quantitative results below, we will examine the strength of this effect in an open economy, and will find it is substantial.

There could be long-run effects of differences in β_L if productivity growth (in either sector) depends on the size of that sector. In a model similar to ours, Matsuyama (1992) proposed that comparative advantage in agriculture would lead to slower growth if there were learning-by-doing effects in *non*-agriculture. If this were the case, then having a large β_L value would be particularly bad for future growth as it leads to an even greater specialization in agriculture when the economy is open.⁷

3. Data, empirical strategy, and results

Theoretically, then, agricultural technology in the form of β will influence the pace of development. However, this is only an interesting theoretical issue if the value of β in our data does indeed vary across countries. In this section we establish that there is significant variation across countries in this technology parameter.

To do this, we estimate agricultural production function using cross-country panel techniques that explicitly accommodate heterogeneity in coefficients across countries. Heterogeneity in the estimated coefficients of the production function is precisely what we mean by “technology heterogeneity”. Once we have estimates of the production function coefficients for each country, we show that those coefficients vary systematically with climate.

(a) Data

The United Nation Food and Agriculture Organization’s (FAO) FAOSTAT database is the principle source for our agricultural production data. This contains annual data for net agricultural output (in thousands of real International \$), labor, capital stock (using tractors as a proxy), livestock, fertilizer, and arable and permanent crop land for 128 countries over the 1961–2002 period. Country-specific information on the share of arable land by climate zones (adopting the Köppen-Geiger classification) are taken from Matthews (1983). We provide more details, including descriptive statistics, in Appendix. We carried out panel unit root tests (results available on request) which could not reject the null that these data constitute nonstationary processes. We also investigated the cross-section correlation properties of the data (results available on request) and found strong evidence for a ‘factor structure’ in the productive inputs and output—the nature of this property is discussed in the following subsection.

⁶ Recall from above that a constant w is consistent with either an economy open to capital flows and/or the steady state of an economy with endogenous savings. Even if non-agricultural output per worker, w , rises temporarily because of the shift of labor into agriculture, this will be offset by declines in capital induced by international capital movements or the lower return to capital in non-agriculture.

⁷ Several recent papers have concerned themselves with the importance of trade for structural change, typically using South Korea as example (Betts, Giri, & Verma, 2013; Sposi, 2015; Uy, Yi, & Zhang, 2013). In these papers a combination of shifts in world prices, trade costs, and productivity are used to match the observed structural transformation of the Korean economy. Heterogeneity in agricultural technology across countries does not invalidate this work. However, our findings imply that these models would have very different quantitative predictions for countries that do not share the same agricultural technology as Korea.

(b) Empirical model

Our model of production aims to capture a number of important features: firstly, we want to allow for endogenous production inputs, an issue which has been the primary focus of the literature on cross-country empirical analysis; secondly, we want to model TFP in a very flexible manner, such that it can contain both country-specific and common elements (i.e., spillovers), while also not being constrained to linearity; thirdly, time series properties of macro panel data, most notably the potential for nonstationary inputs and/or TFP, need to be accommodated; finally, and this represents the main focus of our study, we want to allow for technology heterogeneity across countries following an unspecified pattern (i.e. we do not assume that technology is identical within continents or other *ad hoc* groupings).

In order to capture these characteristics we adopt the following (log-linearized) Cobb–Douglas production function for country $i = 1, \dots, N$ at time $t = 1, \dots, T$

$$Y_{it} = \beta_i' X_{it} + u_{it}, \quad (12)$$

where X is the set of inputs and Y is output (all in logarithms). In practice we use these observable variables in per worker terms, since then the inclusion of the labor variable provides a test for constant returns to scale.⁸ Technology (β_i) differs across countries but is constant over time.⁹ The vector β_i contains a coefficient for each input, including labor, which is the input we are ultimately interested in.¹⁰

Total factor productivity (u), is unobserved, and this is the source of identification issues: if X_{it} is correlated with u then we will receive biased estimates of β . By adopting a common factor framework, we will be able to account for country/year variation in productivity. Formally, let

$$u_{it} = \alpha_i + \gamma_i' f_t + \varepsilon_{it}. \quad (13)$$

TFP consists of (i) a country-specific TFP-level term α_i which captures time-invariant aspects; (ii) a country-time-specific element in the form of a vector of common factors f interacted with country-specific parameters γ_i ; and (iii) remaining country-time variation in ε_{it} that is assumed to be uncorrelated with X_{it} .

The country-specific term α_i is handled by including country-level fixed effects. The common factors term deserves further explanation. f_t is a vector with an arbitrary number of factors that represent “shocks” which affect all countries, although their response to these factors is country-specific, which is what γ_i denotes. Shocks here can refer to events like the global financial crisis, but also secular processes such as the economic rise of China or the evolution of knowledge spillovers in agriculture. These individual factors are orthogonal to each other, so that it is easy to see that we can obtain a highly idiosyncratic “trend” for country i if we combine a small number of these factors and country-specific parameter.

So far we have only given some structure to the unobservable TFP evolution. Where does the endogeneity problem come into the equation? Inputs are deemed endogenous because we believe

inputs are a function of TFP, the measure of ignorance of everything explicitly omitted from our production function. Our empirical model mimics this by specifying that (some of) the same common factors f_t making up TFP also affect the inputs: u_{it} and X_{it} are correlated.

To estimate (12) we use recent panel time series estimators, in particular the common correlated effects estimation approach introduced in Pesaran (2006). Its implementation is very simple, building on the basic insight that since the unobserved factors (f_t) are common to all countries we can use information from the entire panel to proxy for them. And if we have sound proxies for them in the regression, then we no longer have an identification problem for β , since the unobservable error term is now no longer u_{it} but ε_{it} . The interested reader is referred to Eberhardt and Teal (2013b) or Eberhardt, Helmers and Strauss (2013) for a deeper discussion of the technique, but the intuition is straightforward. By including the cross-section averages of the dependent (\bar{Y}_t) and all independent variables (\bar{X}_t) in our actual estimation, we will implicitly be controlling for the unobserved factors f_t . This leaves us merely with the problem of the country-specific impact of the common factors, which is simply addressed by estimating the agricultural production function separately for each country. The technique has been shown to be extremely powerful in this type of panel setting, in that it can provide consistent estimates of β_i or its cross-country average even if the factors are nonstationary, if there are structural breaks in the factors, and if there is cointegration or non-cointegration between the model variables.¹¹

In the Pesaran (2006) Common Correlated Effects Mean Group (CMG) estimator, the individual country production function with k inputs is then specified as

$$Y_{it} = a_i + b_i' X_{it} + c_{0i} \bar{Y}_t + \sum_{m=1}^k c_{mi} \bar{X}_{mt} + e_{it}. \quad (14)$$

The country fixed effect a_i handles the country-specific level of TFP, α_i . The terms involving \bar{Y}_t and \bar{X}_{mt} are capturing the influence of the unobserved factors f_t , and each country is allowed to have its unique coefficients responding to these factors.

Eqn. (14) is estimated country by country using OLS. At that point we have individual estimates of β_i for each country. If we were interested only in some form of average agricultural technology we can compute the mean of these β_i values across all countries.¹² Alternatively, we can look at the average value of β_i across sub-groups of countries with similar climate conditions, and see if the average value for temperate countries differs from equatorial countries, as an example. As we will document below, there do appear to be clear differences across climate zones in the average value of β_i . It is important to note that we do not really put a lot of faith into each individual β_i , since each of these is computed from a small sample with very few degrees of freedom. Put simply, each estimate is a noisy signal of the true parameter value, and averaging across small groups of countries boosts this signal and reduces the noise.¹³ Averaging across alternative groups will enable us to indicate the central tendencies in technology parameters we summarize below.

Our analysis in the following section compares and contrasts empirical results from a variety of different estimators which make different assumptions about the nature of the underlying

⁸ A statistically significant positive (negative) labor coefficient indicates increasing (decreasing) returns to scale, insignificance points to constant returns which can then be imposed by dropping the labor variable from the model.

⁹ We empirically tested this assumption adopting recursive regression models. While average estimates by climate zone change over time as we add further years to the regression sample, we can show that these changes are driven by sample size, rather than systematic differences over time (results available on request).

¹⁰ To be more precise, the labor coefficient is backed out from the estimates on all other inputs (in per worker terms) and the returns to scale estimate—see footnote 8 above.

¹¹ An alternative would be to use some sort of instrument for the endogenous independent variables, either lagged values or an external instrument. However, if these common factors are present, and each individual country responds differently to them, then the use of any of instruments is problematic as they may be functions of the same common factors (Andrews, 2005; Coakley, Fuertes, & Smith, 2006).

¹² This follows the practice of the Pesaran and Smith (1995) Mean Group estimator. Although \bar{Y}_t and e_{it} are not independent their correlation goes to zero as N becomes large.

¹³ Eberhardt and Teal (2013b) provide a more formal exposition.

Table 1
Production function estimates

	[1] POLS	[2] 2FE	[3] MG	[4] CDMG	[5] CMG
Labor	−0.059	−0.191	−0.357		
$\hat{\beta}_L$	[17.58]**	[10.60]**	[2.23]*		
Tractors pw	0.131	0.058	0.075	0.072	0.109
$\hat{\beta}_K$	[24.22]**	[13.06]**	[3.31]**	[3.18]**	[5.23]**
Livestock pw	0.219	0.358	0.246	0.327	0.321
$\hat{\beta}_{Live}$	[28.73]**	[25.34]**	[8.07]**	[8.28]**	[9.47]**
Fertilizer pw	0.169	0.073	0.030	0.043	0.036
$\hat{\beta}_F$	[28.19]**	[19.87]**	[4.86]**	[5.70]**	[5.63]**
Land pw	0.253	0.294	0.210	0.329	0.201
$\hat{\beta}_N$	[30.55]**	[13.05]**	[2.79]**	[7.10]**	[3.57]**
Returns to Scale ^a	DRS	DRS	DRS	CRS	CRS
Implied Avg $\hat{\beta}_L$	0.169	0.027	0.082	0.229	0.333
Avg Implied $\hat{\beta}_L$	n/a	n/a	0.122	0.157	0.251
$\hat{\epsilon}$ Stationarity ^b	I(0)/I(1)	I(1)	I(0)	I(0)	I(0)
$\hat{\epsilon}$ CD Test (p) ^c	−2.49 (.01)	9.64 (.00)	9.16 (.00)	3.59 (.00)	−0.23 (.82)
RMSE	0.433	0.148	0.066	0.081	0.059

Notes. Results for $n = 5,162$ observations from $N = 128$ countries. Estimators: POLS—Pooled OLS (with year fixed effects), 2FE—2-way fixed effects, MG—Pesaran and Smith (1995) Mean Group, CDMG—MG using data in deviation from cross-section means, CMG—Pesaran (2006) Common Correlated Effects Mean Group. We present robust sample means for model parameters in [3]–[5]; [1] and [2] are pooled models. * and ** indicate statistical significance at the 5% and 1% level respectively. Terms in square brackets are absolute t -statistics based on standard heteroskedasticity-robust standard errors in [1] and [2], and based on the variance estimator following Pesaran and Smith (1995) in [3]–[5]. RMSE reports the root mean squared error.

Dependent variable: log output per worker, [2] dto. in 2FE transformation as described in Coakley et al. (2006), [4] dto. in deviation from cross-section means.

Independent variables: all variables are in log per worker terms with the exception of ‘Labor,’ for which the coefficient estimate indicates deviation from constant returns (formally: $\hat{\beta}_L + \hat{\beta}_K + \hat{\beta}_{Live} + \hat{\beta}_F + \hat{\beta}_N - 1$). This is not the technology coefficient on labor, which is reported under ‘Implied Avg $\hat{\beta}_L$ ’ and ‘Avg Implied $\hat{\beta}_L$ ’ in a lower panel of the Table.

All empirical analysis was conducted in Stata using the `multipurt`, `xted` and `xtmg` commands written by Markus Eberhardt and freely available at SSC (“findit” command) within Stata on a computer linked to the internet.

^a The implied returns to scale are labeled decreasing (DRS) if the coefficient on Labor is negative significant, and constant (CRS) if this coefficient is insignificant. We present models with CRS imposed if the unrestricted model cannot reject this specification. For heterogeneous models we present two robust mean estimates for the implied labor coefficient: the first is computed from the averages presented in the upper panel of the Table, while the second is constructed at the country level and then averaged.

^b Pesaran (2007) CIPS test results for residual nonstationarity: I(0)—stationary, I(1)—nonstationary. Full results available on request.

^c Pesaran (2015) test for weak cross-section dependence (CD) in the residuals, H_0 : no strong CD.

productive relationship.¹⁴ We consider (i) pooled estimators, assuming common technology with flexible but common TFP evolution over time ($\alpha_i = \alpha$ and $\beta_i = \beta$: POLS), or assuming TFP-level differences across countries (α_i and $\beta_i = \beta$: 2FE); (ii) a heterogeneous estimator assuming country-specific but constant TFP growth (β_i and $\gamma_i f_t = \gamma_i t$: MG); (iii) a heterogeneous estimator allowing for flexible TFP growth over time but assuming this is common to all countries (β_i and $\gamma_i f_t = \gamma_i t$: CDMG); and (iv) a heterogeneous estimator allowing for flexible TFP growth over time which can differ across countries (β_i and $\gamma_i f_t$: CMG).

Our empirics here are not of the “treatment/effect” type. We are not testing the robustness of a specific empirical regularity in the data (e.g., the statistical significance and sign of a specific variable) across different models. Our approach is more in line with time series forecasting, where the goal is find a model that best explains the variation in the data. As such, deciding which model best fits the data is based on formal diagnostic testing. Specifically, since nonstationarity is a concern in these data, we investigate the regression residuals using panel unit root tests, where nonstationarity of residuals implies that the empirical results are potentially spurious. We also carry out tests for residual cross-section dependence: if we reject independence¹⁵ this is an indication that our

identification strategy for the technology parameters is likely flawed, as the remaining unobserved variation in TFP does not conform to the assumption that it is i.i.d.: we are likely not to have captured all the factors, such that the residuals are not white noise (ϵ_{it}).

(c) Empirical results and diagnostics

We start by providing production function estimates of the common technology in Table 1. In columns [1]–[4] we present several alternative estimators of the agricultural production function to show how they compare diagnostically to our preferred CMG estimator in column [5]. Columns [1] and [2] present results from simple pooled regressions. POLS ignores the common factors and country fixed effects, includes only year fixed effects, and assumes that all countries have the same $\beta_i = \beta$. 2FE also ignores the common factors, but includes both country and year fixed effects. However, it also assumes that $\beta_i = \beta$.

What is immediately apparent is that the estimated coefficient on labor is negative for both the POLS and 2FE models. As noted above, we are estimating the model with all variables in per-worker terms, while also including labor as a separate regressor. Thus the negative coefficient on labor implies that there are decreasing returns to scale at the global level—a common finding in the literature reviewed in Eberhardt and Teal (2013a). This does not imply that this is a reasonable description of agricultural production. As can be seen in a lower panel of the Table, residuals remain cross-sectionally dependent and are further likely nonstationary. In other words, there are reasons to suspect that the results (on labor in particular) are spurious.

In column [3] we present the Pesaran and Smith (1995) Mean Group estimator. This allows for country-specific β_i . The common

¹⁴ Abbreviations for the estimators implemented: POLS—Pooled OLS (with year fixed effects), 2FE—2-way fixed effects (equivalent to country and year fixed effects), MG—Pesaran and Smith (1995) Mean Group (with country-specific trends), CDMG—MG using data in deviation from cross-section means, CMG—Pesaran (2006) Common Correlated Effects Mean Group.

¹⁵ More precisely we adopt the Pesaran (2015) CD test for weak dependence, with strong dependence the alternative. Akin to heteroskedasticity the former is not a problem for estimation, merely for inference; however, since the standard errors for our CCEMG estimates are computed non-parametrically (i.e. not from the standard errors of $\hat{\beta}_i$) this does not cause any additional problem.

factor term is limited to $\gamma'f_t = \gamma_t t$, which implies that each country has a unique growth rate for TFP, but that this growth rate is constant over time. What can be seen in the table is that the labor coefficient here is again negative, implying decreasing returns to scale at the global level. While the error terms are stationary in this estimation, they are still cross-sectionally dependent, indicating that this model does not suffice to capture the agricultural production process in our sample.

The CDMG estimate in column [4] also allows for country-specific coefficients, but assumes that the common factor is such that $\gamma'f_t = \gamma f_t$. In this case, the common factor f_t is allowed to vary arbitrarily over time, but each country is constrained to respond to this identically. In this case, our original estimates indicated that the labor coefficient was not significantly different from zero (available upon request), which is consistent with constant returns to scale. What is reported in the table is an estimate where we constrain the model to be constant returns to scale.

The implied values for $\hat{\beta}_L$ in this case are found in the second panel. As we have individual estimates of all coefficients for each country, there are two ways to find the average $\hat{\beta}_L$. We can get the implied average $\hat{\beta}_L$ from taking one minus the reported coefficients on other inputs reported in column [4], giving us 0.229. Alternatively, we could find $\hat{\beta}_{Li}$ for each country as one minus the sum of their individual estimates of the coefficients on the other inputs, and then average those, which yields 0.157 as the common labor coefficient. These estimates are distinctly higher than what we find in the POLS, 2FE, or MG estimates. But while the CDMG estimates have stationary residuals, there is still cross-sectional dependence indicated in the residuals, implying that it does not capture the full heterogeneity in the data.

This leaves us with our preferred estimates in column [5], which allow for the most flexibility. As discussed in the prior section, we allow for heterogeneity in β_i , as well as allowing each country to have a unique reaction to the set of arbitrary common factors, $\gamma'f_t$. The CMG estimates also cannot reject constant returns to scale at the global level, and hence we report the constrained

model in the table. Note that residuals are both stationary, and unlike the other models presented, we can reject cross-sectional dependence. The CMG model appears to have captured the heterogeneity present in the data. Further, note that the root mean square error of the CMG model is lower than any other specification, indicating a better fit to the data.

Finally, the CMG results imply a much more substantial average coefficient for labor than any of the other models, either 0.333 or 0.251 depending on how one constructs the average. Recall, though, that these estimates of $\hat{\beta}_L$ are averages across all countries in our sample, who each have an individual $\hat{\beta}_{Li}$ estimate. Table 1 is useful to establish that the CMG estimates have the best diagnostic properties of available specifications, and thus as we proceed we will use the CMG estimates exclusively. But the main purpose of the empirical analysis conducted was not to analyse average estimates across the whole sample, but to provide some insights into the patterns of technology heterogeneity across countries, which we turn to in the following section.

(d) Labor coefficients by agro-climatic zone

To show the heterogeneity across agro-climatic zone in labor coefficients, we compare the mean estimates of the labor coefficients for four different groups of countries. The groups are defined based on their broad agro-climatic characteristics (full sample shares of land in the respective climate zone in parenthesis, see Appendix for classification): arid (19%), temperate and cold (36%), equatorial (37%), and highland (8%). As most countries have agricultural land covering more than one of these types of climate zones, there is no single way to assign countries to an individual group.

We thus use different cutoff points to analyze the agro-climatic zone averages of the labor coefficient. In all cases the average labor coefficient is the outlier-robust mean following Hamilton (1992). In Table 2 Panel A we average coefficients for those countries with any land in the specified zone, which inevitably leads to a significant double-counting. Despite that, one can see that the labor coefficients

Table 2
Labor coefficients and climate zones

	Arid	Temperate/Cold	Equatorial	Highland
<i>Panel A: Any share of arable land in climatic zone indicated</i>				
Mean $\hat{\beta}_L$	0.183	0.116	0.405	0.265
	[0.093]*	[0.071]	[0.082]***	[0.099]**
Sample	57	67	70	47
<i>Panel B: Share above zone mean</i>				
Mean $\hat{\beta}_L$	0.144	0.179	0.502	0.342
	[0.110]	[0.075]**	[0.101]***	[0.141]**
Sample	38	53	53	26
<i>Panel C: Share above 40%</i>				
Mean $\hat{\beta}_L$	0.088	0.154	0.500	0.209
	[0.124]	[0.077]*	[0.104]***	[0.243]
Sample	23	51	52	12
<i>Panel D: Share above 50%</i>				
Mean $\hat{\beta}_L$	0.068	0.155	0.492	0.312
	[0.143]	[0.079]*	[0.109]***	[0.329]
Sample	20	49	47	7
<i>Panel E: Share above 60%</i>				
Mean $\hat{\beta}_L$	0.087	0.104	0.506	0.801
	[0.136]	[0.091]	[0.117]***	[0.303]*
Sample	17	43	43	4

Notes. Each cell in this Table presents the robust mean of the estimated labor coefficient $\hat{\beta}_L$ given the criterion indicated. These averages are based on the CMG estimates, column [5] of Table 1. All criteria refer to the share of agricultural land in the respective climatic zones (denoted by their Köppen–Geiger letters), e.g. each robust mean computed in Panel B only includes countries which have a share of land above the full sample mean in the specific climatic zone. Sample indicates the number of observations from which the robust mean is constructed.

are quite low in both arid (0.183) and temperate countries (0.116), while they are significantly higher in equatorial (0.405) and highland (0.265) ones. In Panel B the averages are constructed from countries where the share of agricultural land in the specified zone is higher than the overall sample mean share. The pattern here is similar, with arid and temperate countries have lower average labor coefficients than equatorial or highland countries.

In panels C, D, and E, we impose *ad hoc* cutoffs of 40%, 50%, and 60% for the share of land in a climate zone to be classified into that group. This limits the number of countries that count as arid or highland, but leaves the numbers of temperate and equatorial countries similar. Based on this simple descriptive analysis we observe a pattern whereby labor coefficients are very low in the arid zone (typically countries located in the Sahel, in Northern Africa or on the Arab Peninsula), around 0.15 in the temperate/cold (Europe, North America), around 0.5 in the equatorial (much of Sub-Saharan Africa) and around 0.3 (Afghanistan, Ethiopia) in the highland zone.

What Table 2 establishes is that the labor elasticity of agriculture varies across agro-climatic zone. As seen in the prior section, the labor elasticity plays a quantitative role in determining the speed at which countries move from agriculture to non-agriculture. The question now is whether the variation in β_L implies an economically significant quantitative effect on development. We turn to that question in the following section.

Prior to the quantitative work, we want to note address several possible concerns with the wide variation we find in the labor elasticity of agriculture. This may seem at odds with evidence of labor shares in output, which are generally assumed to be around 0.60. First, aggregate labor shares in that range are perfectly consistent with low agricultural labor shares, as the aggregate labor share is a value-added weighted average of the labor shares across sectors. Agriculture accounts for a relatively small fraction of value-added even in developing countries, and hence agricultural labor shares of 0.15 or 0.35 would imply aggregate labor shares of around 0.55–0.59 depending on the specific country.

Second, what limited evidence we have on agricultural labor shares is consistent with our estimates. Fuglie (2010) documents that the labor share for the U.S. is only 0.20, while other countries in temperate areas have values such as 0.19 (several former Soviet countries), 0.23 (South Africa), and 0.30 (the U.K.). These are not vastly different from our average estimated elasticity for these countries of about 0.15. This can be compared to several developing countries in our equatorial cluster that Fuglie documents as having distinctly higher values for labor's share in agriculture: India (0.46), Indonesia (0.46), and Brazil (0.43). These values are all close to the average estimated elasticity for the equatorial group of around 0.35. In short, there does seem to be evidence of heterogeneity in labor shares, and it roughly matches the heterogeneity in elasticities we estimate.

Finally, the labor elasticity of agricultural output with respect to total agricultural labor—which is relevant for development—is not necessarily identical to the elasticity of output with respect to labor at the farm level using a specific agricultural technique. As Hayami and Ruttan (1985) discussed, what is relevant is the “meta-production function” for agriculture, an envelope of the available techniques, and that is what we are estimating here. Our estimates tell us the percent increase in output associated with a percent increase in labor, allowing farmers to change agricultural techniques. For studying long-run growth and development, this is the more appropriate elasticity to focus on.

4. Implications of technology heterogeneity

Theoretically, we know that agricultural technology influences the pace of development. Empirically, we established that there

was significant variation across countries in the estimated values of β_L . Most importantly, we saw that the labor elasticity (which we will refer to as β_L from here on out) varies by climate zone. The purpose of this section is to establish if the variation in β_L is quantitatively important for development.

(a) Agricultural technology and structural change

We use a calibrated version of our model to compare the equilibrium outcomes following exogenous increases in agricultural technology across values of β_L . To calibrate the model we use data from South Korea in 1963 and 2005. We use South Korea as we wanted to ensure that our calibration captures the experience of countries undergoing the early stages of structural change while at the same time anchoring our analysis in the post-WWII period of an increasingly globalizing world economy. In 1963 South Korea's workforce was 63% agricultural, but by 2005 this share had fallen to only 8%. Over the same period agricultural output per worker went up by a factor of 7.4, non-agricultural output per worker went up by a factor of 3.5, and total population rose by a factor of 1.8.¹⁶ It is an example of a relatively underdeveloped country that experienced rapid structural transformation, and we feel it provides a useful benchmark for evaluating other developing countries. We use the closed economy version of the model, despite Korea's known expansion of trade, as the closed economy model still tracks the broad trends in South Korea more closely than an open economy version. In particular, given the increase in agricultural TFP we find in South Korea, an open economy model would have predicted a nearly complete specialization in agriculture by South Korea, which obviously is not the case.

Setting the initial values of A , w , and L all equal to one, we find values for the utility parameters α (the long-run preference weight on agriculture), \bar{c}_a (the subsistence constraint), and \bar{c}_n (the endowment of non-agricultural consumption) that deliver the observed drop in the agricultural labor share from 63% to 8% given the observed changes in labor productivity in the two sectors. The resulting values are $\alpha = 0.01$, $\bar{c}_a = 0.802$, and $\bar{c}_n = 3.182$.

Given these values, we construct three separate simulated economies, varying in their values of β_L . We use the values $\beta \in (0.15, 0.35, 0.55)$ as these capture the rough span of values found in Table 2. The value of 0.15 is near the average for countries in temperate/cold climate zones. The values of 0.35 and 0.55 span the range of average values found for Highland and Equatorial countries.

We set the value of $A = 1$ and $w = 1$ in each economy initially as a normalization. To ensure a meaningful comparison, we impose that each economy begins with an agricultural labor share of 80%, corresponding to an under-developed economy. To do this we adjust the initial level of population, L , so that the equilibrium outcome of the model using our calibrated parameters is precisely $L_a/L = 0.8$. This has the consequence that all of the interesting outcomes in the model are also identical across the different economies. In particular, the relative price of food (p_a), utility, agricultural output per worker (Y_a/L_a), and consumption of the two goods (c_a and c_n) are all identical across the three economies in the initial set-up.

Starting from this identical initial condition, we then impose an exogenous increase of 20% to agricultural TFP (A) to each economy. The results of these calculations can be found in Table 3. There, the outcomes for each of the three simulated economies are shown, each relative to the baseline listed in the second column. The first column of the simulations is for the economy with $\beta_L = 0.15$, the temperate climate value. With a productivity increase of 20% in

¹⁶ Authors' calculations using data from Timmer and de Vries (2007).

Table 3
Effects of exogenous productivity and population changes—closed economies

Outcome	Baseline	Equilibrium outcomes from					
		20% Increase in Ag. TFP (A) with $\beta =$			5% Increase in Population (L) with $\beta =$		
		0.15	0.35	0.55	0.15	0.35	0.55
Ag. labor share (L_a/L)	0.8	0.369	0.518	0.595	0.944	0.852	0.823
Ag. relative price (p_a)	1.0	0.432	0.629	0.729	1.189	1.068	1.031
Ag. labor productivity (Y_a/L_a)	1.0	2.314	1.591	1.371	0.841	0.936	0.970
Ag. consumption p.c. (c_a)	1.0	1.069	1.030	1.019	0.992	0.997	0.998
Non-ag. consumption p.c. (c_n)	1.0	3.153	2.408	2.026	0.281	0.740	0.882
Real income p.c. (y)	1.0	1.485	1.306	1.221	0.849	0.945	0.975

Notes. The table shows the equilibrium outcomes for each of the listed measures. The “Baseline” is the starting equilibrium for each of the six simulations. The results of (1) increasing agricultural TFP by 20% and (2) increasing population size by 5% are shown for three sample economies. These sample economies vary only in the size of their β , the elasticity of agricultural output with respect to labor. The initial level of agricultural TFP is equal to one in each simulation. The initial size of the population, L , is set so that the baseline values are identical (L equals 1.18, 1.16, or 1.13 for $\beta = 0.15, 0.35, 0.55$, respectively). Real income per capita is calculated using the initial relative price of agricultural goods, which is identical for all three economies. The parameter values used in the calculations are as follows: $\alpha = 0.01$, $\bar{c}_a = 0.802$, $\bar{c}_n = 3.182$, $w = 1.0$.

agriculture, this leads to a drop in the agricultural labor share from 0.80 to 0.37. Compare this to the values for economies with $\beta_L = 0.35$ or $\beta_L = 0.55$, where the labor share in agriculture falls only to 0.518 or 0.595. An additional 20% of the labor force is able to move into non-agricultural following this productivity increase when the labor elasticity is as low as $\beta_L = 0.15$.

In all three economies, the relative price of agricultural goods declines, but note that the drop is most pronounced when the value of β_L is only 0.15. The price is only 43% of its original level, compared to the 73% when $\beta_L = 0.55$. Labor productivity in agriculture also rises for all three economies, consistent with the drop in the relative price. The size of the increase, though, is again quite sensitive to the value of β_L . For $\beta_L = 0.15$, labor productivity in agriculture is roughly 2.3 times higher than prior to the productivity increase. The sizeable jump is due not only to the productivity improvement, but the exit of workers from that sector, which drives up the average product of the remaining ones. While labor productivity in agriculture rises in the economies with $\beta_L = 0.35$ and $\beta_L = 0.55$, the gain is not as substantial, with only 59% and 37% increases, respectively.

The gain in labor productivity gets smaller as β_L rises because β_L determines the speed at which the marginal product of labor rises when labor exits. With a high value of β_L , the agricultural production function is close to linear with respect to labor. As labor leaves that sector following the productivity increase, the marginal product of labor does not rise very much. As the average product of labor is proportional to the marginal product in the Cobb–Douglas setting, this means that the average product does not rise much either. For an economy with a low $\beta_L = 0.15$, there is more curvature in the production function, and as labor exits agriculture the marginal product of the remaining labor rises very quickly. This drives up the average product as well.

We also measure how the individual consumption baskets and an index of real income per capita respond to the agricultural productivity improvement. Agricultural consumption rises by about 6.9% in the economy with $\beta_L = 0.15$, which is 2–3 times larger than the gains seen in the economies with higher values of β_L . While the low- β_L economy has a larger gain in agricultural consumption, in no case is the change very sizable. This occurs because of the nature of the utility function used to model these economies. The marginal utility of c_a is very high at low levels of consumption, but declines very quickly as more c_a is consumed. The majority of the new purchasing power made available from the productivity increase is spent on non-agricultural goods.¹⁷

Non-agricultural consumption, c_n , shows the real substantial effects of the productivity increase. The flow of workers into non-agriculture raises output in that sector, allowing for a substantial increase in consumption of those goods. For the economy with $\beta_L = 0.15$, there is a threefold increase in consumption of non-agricultural goods. While the gains are large (more than doubling) for the other simulated economies, their increase in c_n is not nearly as big as when $\beta_L = 0.15$.

Finally, we calculate a constant-price GDP (y) for each economy both before and after the increase in productivity, using the initial relative price to value goods. As this initial price is identical in the three economies, the output per capita numbers are comparable across the three economies. The final row shows the increase following the improvement in agricultural productivity. For $\beta_L = 0.15$, real income per capita rises by about 49%. This compares to only 31% for the economy with $\beta_L = 0.35$ and 22% when $\beta_L = 0.55$. Measured income per capita rises twice as much for the lowest elasticity economy compared to the others, despite the three economies sharing an identical starting point and an identical shock to agricultural productivity. The low-elasticity economy is able to generate such a large gain because of its ability to release labor easily from the agricultural sector, ramping up non-agricultural production.

Note that agricultural productivity improvements always benefits an economy. Our point is simply that the nature of agricultural technology (β_L) is crucial in determining the magnitude of the benefits. The lower is β_L , the more an economy can benefit from agricultural productivity improvements. An alternative way of seeing the influence of β_L on development can be found in Figures 1 and 2. In the first, for the three different calibrated economies we plot the labor share in agriculture over a range of values for agricultural productivity. The economies differ only in their labor elasticity, β_L , and are again calibrated so that they have identical outcomes when $A = 1$. As agricultural productivity increases the labor share declines in all three, however, the reduction is clearly faster when $\beta_L = 0.15$ than in the other economies. Labor is able to leave agriculture in this low-elasticity economy with a smaller “push” from agricultural productivity than in high-elasticity economies.

Similarly, Figure 2 shows that our measure of real income per capita increases faster with improvements in agricultural productivity when $\beta_L = 0.15$ than when β_L is a higher value. The ability of the low- β_L economy to shed labor from the agricultural sector makes non-agricultural output more widely available, raising living standards faster than in the alternative economies. The persistent advantage of the low- β_L economy holds even as productivity rises to 3.5 times the baseline level, equivalent to over twenty-five years of 5% per annum growth in agricultural productivity.

¹⁷ This is a common feature of models of structural change, and is justified mainly by the fact that the models are capable of matching the movements in expenditure on agricultural and non-agricultural goods quite closely.




Figure 1. Agricultural Productivity and Labor Share. *Notes.* The figure shows the share of labor in agriculture for given levels of agricultural productivity (A) in the calibrated model. The three lines denote the relationship in an economy using the noted value for β_L , the labor elasticity of agricultural production. The calibration for each economy is constructed so that the agricultural labor share is exactly equal to 0.80 when $A = 1$.




Figure 2. Agricultural Productivity and Real Income per capita. *Notes.* The figure shows an index of real income per capita for given levels of agricultural productivity (A) in the calibrated model. The three lines denote real income per capita relative to the baseline value of 100 for an economy using the noted value for β_L , the labor elasticity of agricultural production. The calibration for each economy is constructed so that they have an identical real income per capita when $A = 1$.

(b) Population growth and structural change

While economies with a large β_L do not benefit as much from productivity increases, they do have a distinct advantage over economies with a low β_L . Referring back to Eqn. (10), note that with a high elasticity, the share of workers in agriculture is less sensitive to an increase in population, L . From a Malthusian perspective, having a high β means that the detrimental effect of a limited supply of land is less severe. In this sense, a high- β_L economy has a distinct advantage over one with a low β_L .

To make these effects concrete, we perform another simulation using our three model economies. This time we hold agricultural productivity constant but raise the size of the population by 5%. Table 3 shows the results of these experiments in the final three columns. For the economy with $\beta_L = 0.15$, the share of labor in agriculture rises from 80% to 94%. Compared to this, the economy with $\beta_L = 0.35$ only sees an increase to 85%, and for $\beta_L = 0.55$ the

increase is only to 82%. The low- β_L economy is at a distinct disadvantage here, shifting a much larger fraction of its workforce back into agriculture to meet food demand.

The effects on agricultural productivity, consumption levels, and income per capita are also more severe when $\beta_L = 0.15$ than for higher values. The low-elasticity economy experiences a drop in labor productivity in agriculture to only 84% of its prior level, and the loss of non-agricultural workers lowers the consumption of those goods to nearly one-fourth of the prior value. This leads to a decline in real income per capita to about 85% of its value before the population increase.

Compared to this, the economies with higher β_L values fare better. For $\beta_L = 0.55$, agricultural labor productivity stays at about 97% of the level before the population shift, and non-agricultural consumption remains at about 88% of the earlier value. Real income per capita falls only to 97.5% of the baseline value. The economy with $\beta_L = 0.35$ is similarly better off than the low-elasticity economy.

The reason for the advantage of the high- β_L economies is the exact opposite logic of that applied above regarding productivity increases. With a high value for β_L , agricultural output is very sensitive to the number of workers. When population increases, more workers are necessary, but it only takes a few to increase output by enough to meet the new demand. It is now the low β_L economy that lags behind, as it must shift large amounts of labor into the agricultural sector to try and drive up output to meet demand.

Table 3 shows that agricultural technology does not have a monotonic relationship with development. Temperate zones with low levels of β_L will be able to take advantage of agricultural productivity improvements more aggressively, but are also more susceptible to declining living standards through population growth. High β_L zones, such as tropical and highland areas, are better able to adapt to population growth, but will not be able to reap the benefits of productivity improvements in the same manner as temperate areas.

(c) Structural change in an open economy

As noted above, the direction of change in L_a from a shock to A is different in an open economy, as labor is shifted *into* the more productivity agricultural sector. Here we show simulations that use the same structure as those pursued to this point, but where the relative international price of agricultural goods, p_a^* is fixed. The initial conditions are the same as used previously. In particular, the labor share in agriculture is set to be 80% in all economies.

In Table 4 we simulate the effect of a 10% increase in agricultural TFP and, separately, a 10% increase in non-agricultural TFP. Focusing first on the agricultural TFP increase, one can see that the labor share in agriculture after this shock is higher in the $\beta_L = 0.55$ economies, at about 98%, as compared to just under 90% for the $\beta_L = 0.15$ economy. In an open economy, the response of labor shares to shocks is increasing with the size of β_L as there is no internal adjustment of the relative price of agricultural goods. Despite the large shift of labor into agriculture, the $\beta_L = 0.55$ economy is not any better (or worse) off than the others after the shock. Given that we are holding the non-agricultural output per worker constant in this simulation, *all three* economies have identical agricultural labor productivity and identical real income per capita after the shock. Agricultural labor productivity is unchanged as the increase due to A is offset exactly by the influx of workers in order to equalize wages across sectors. Real income per capita is unchanged as this is pinned down by the non-agricultural output per worker.

The second part of Table 4 shows the results of a shock to non-agricultural TFP, w . Here again, the high- β_L economies react more

Table 4
Effects of exogenous productivity and population changes—open economies

Outcome	Baseline	Equilibrium outcomes from					
		10% Increase in Ag. TFP (A) with $\beta =$			10% Increase in Non-ag. TFP (w) with $\beta =$		
		0.15	0.35	0.55	0.15	0.35	0.55
Ag. labor share (L_a/L)	0.8	0.895	0.926	0.989	0.715	0.691	0.647
Ag. labor productivity (Y_a/L_a)	1.0	1.000	1.000	1.000	1.100	1.100	1.100
Real income p.c. (y)	1.0	1.000	1.000	1.000	1.100	1.100	1.100

Notes. The table shows the equilibrium outcomes for each of the listed measures using the model with free trade of both agricultural and non-agricultural goods. The “Baseline” is the starting equilibrium for each of the six simulations. The results of (1) increasing agricultural TFP by 10% and (2) increasing non-agricultural TFP by 10% are shown for three sample economies. These sample economies vary only in the size of their β , the elasticity of agricultural output with respect to labor. The initial level of agricultural TFP and non-agricultural TFP is equal to one in each simulation. The relative world price of agricultural goods is set to $p_a^* = 0.5$. The initial size of the population, L , is set so that the baseline values are identical (L equals 0.55, 0.43, 0.27 for $\beta = 0.15, 0.35, 0.55$, respectively). Real income per capita is calculated using world price of agricultural goods. The parameter values used in the calculations are as follows: $\alpha = 0.01$, $\bar{c}_a = 0.802$, $\bar{c}_n = 3.182$.

Table 5
Long-run development outcomes

Outcome	Baseline	Equilibrium after 270% increase in A 250% increase in w 80% increase in L with $\beta =$		
		0.15	0.35	0.55
Ag. labor share (L_a/L)	0.800	0.046	0.083	0.131
Ag. relative price (p_a)	100.0	12.5	29.4	51.9
Ag. labor productivity (Y_a/L_a)	100.0	2804.9	1191.0	675.0
Ag. consumption p.c. (c_a)	100.0	161.7	123.2	110.9
Non-ag. consumption p.c. (c_n)	100.0	1669.3	1605.2	1520.0
Real income p.c. (y)	100.0	463.2	419.6	392.7

Notes. The table shows the equilibrium outcomes for each of the listed measures. The “Baseline” is the starting equilibrium for each of the three simulations. For each simulation, agricultural productivity, non-agricultural productivity, and population exogenously grow by an amount equivalent to that experienced by South Korea in the period 1963–2005. The simulation economies vary only in the size of their β , the elasticity of agricultural output with respect to labor. The initial level of agricultural TFP is equal to one in each simulation. The initial size of the population, L , is set so that the baseline values are identical (L equals 1.51, 1.22, or 1.07 for $\beta = 0.15, 0.35, 0.55$, respectively). Real income per capita is calculated using the final period relative price of agricultural goods. The parameter values used in the calculations are as follows: $\alpha = 0.01$, $\bar{c}_a = 0.802$, $\bar{c}_n = 3.182$, $w = 1.0$.

strongly to the shock. The agricultural labor share falls to 0.65 when $\beta_L = 0.55$, but only to 0.72 when $\beta_L = 0.15$. However, similar to the agricultural shock, all three economies experience identical effects on agricultural labor productivity and real income per capita. In each case, there is a 10% increase in both, matching the 10% increase in non-agricultural labor productivity. The higher labor productivity in agriculture is due to the shift of labor to the non-agricultural sector. With $\beta_L = 0.15$, only a few workers need be released to non-agriculture before the average product is raised enough to equalize wages across the sectors, which is why the shift of labor is much smaller in this case. For the high- β_L economies, more labor must leave agriculture to raise average products there, and so they see higher shifts of labor when reaching the 10% gain in productivity and output per capita.

Open economies reinforce the point that β_L introduces significant heterogeneity into the response of developing economies to shocks. In the open economy case, this causes larger labor shifts in the high- β_L economies, where the opposite occurs in closed economies. However, the more severe shifts of labor in open economies are not associated with heterogeneity in agricultural labor productivity or income per capita. Trade limits the sensitivity of living standards and labor productivity to shocks, even as it raises the sensitivity of labor shares.

(d) Long-run development and agricultural technology

To grasp the long-run implications of differences in agricultural technology we consider a final experiment with the calibrated closed-economy model. We allow for each economy to experience

the same growth in productivity and population that occurred in South Korea from 1963 to 2005, and compare the outcomes. In that period, agricultural productivity (A) rose by a factor of 3.7, while non-agricultural labor productivity (w) increased by a factor of 3.5. Population nearly doubled, rising by 80% in those forty-two years.¹⁸ The productivity increases are substantial, suggesting that places with relatively low values of β_L would benefit greatly. However, the large increase in population acts to offset this somewhat, restricting the impact of the productivity changes. The value of using our calibrated model is that we can see which effect was stronger.

Table 5 shows the results of this simple experiment. As can be seen in the first column under the results, for an economy with $\beta_L = 0.15$, the labor share in agriculture drops to less than 5% after the South Korea-like changes. Agricultural labor productivity rises by a factor of 28, which in part leads to an increase of 61% in agricultural good consumption. This is dwarfed by a 17-fold increase in the consumption of non-agricultural goods, as the vast shift of labor out of agriculture allows for a massive expansion of that sector. At constant prices, these changes lead to a 4.6-fold increase in GDP in the economy with $\beta_L = 0.15$.

Compared to this, the effects in economies with $\beta_L = 0.35$ and $\beta_L = 0.55$ are not as impressive, although they are quite dramatic on their own. Consider the economy with the value of $\beta_L = 0.55$. The agricultural labor share falls, but only to about 13%, remaining

¹⁸ Authors' calculations using data from Timmer and de Vries (2007). The growth in agricultural productivity is the growth in A necessary to raise agricultural output per worker by the observed factor of 7.4, given the observed change in the number of agricultural workers. It includes not only shifts in total factor productivity, but any increases in fertilizer use and capital use during this period.

more than twice as large as that seen in the low- β_L economy. Labor productivity in agriculture rises by a factor of 6.75, a major improvement, but under a fourth as great as that seen in the low- β_L situation. Constant-price GDP also rises by a factor of under 4, indicating a massive increase in living standards. However, even though they started out in exactly the same situation and experienced exactly the same changes to productivity and population, the $\beta_L = 0.55$ economy now has living standards that lag those in the $\beta_L = 0.15$ economy.

The increases in productivity experienced were large enough to offset the increased population growth, and the economy with a low β_L is able to reap the benefits to a greater extent. The low elasticity of output with respect to labor means that the economy is able to shed labor from agriculture quickly, as production is not that dependent on labor. These freed-up workers are available for the non-agricultural sector, where they produce a massive increase in goods available.

5. Discussion and conclusion

Our simulated model shows that structural change is highly sensitive to the elasticity of agricultural output with respect to labor, β_L in our terminology. The empirical work documented that this elasticity does in fact vary widely across countries. In particular, the value of β_L is closely correlated with the agro-climatic zone that dominates within a country. Summarizing our results, we find that temperate/cold climate zones have low labor elasticities in agriculture, on the order of 0.15. Compared to this, equatorial and highland zones have elasticities ranging from 0.35 to 0.55.

Our results imply that there are significant effects of agricultural technology on development. However, we must be very clear on what precisely those effects are. Our findings are not an example of “geographic determinism,” the notion that tropical or highland countries are somehow doomed to be poor by their climate. Recall that the effect of labor elasticity (and hence climate) differs for agricultural productivity and population growth. Fundamentally, it is the relative strength of these two forces that determine a country's living standards in our simple model. The fact that the agricultural technology (the β_L coefficient) varies by climate does not imply anything about productivity or population levels varying by climate.

Where climate does appear to play an important role is in the speed of structural change, contingent on productivity levels. As we showed in the model, the size of β_L dictates the importance of labor for agricultural production. In low- β_L countries of the temperate zone, labor is relatively unimportant, and hence can be released easily to non-agricultural work following a productivity improvement. Hence temperate countries will be able to transition quickly from agriculture to industry given a favorable productivity shock. In contrast, equatorial or tropical countries have high β_L values, and hence labor is particularly important in agricultural production. Labor is released less readily to non-agricultural work, and so the transition from agriculture to industry will be slower for a given productivity shock.

Our work thus offers both an optimistic and pessimistic outlook on development. The optimistic viewpoint is that the relatively slow structural change in equatorial regions is not (only) due to some barrier to productivity improvements and that these places may well be adopting new technologies and techniques in agriculture at rates similar to their peers in other regions. The pessimistic aspect is that the inherent agricultural technology will continue to cause equatorial regions to move through the process of structural change at a slower pace than seen in temperate countries. In that sense it would be unfair to compare the current development trajectory of Sub-Saharan African, Central American, or South Asian

countries to earlier developers in temperate regions such as Japan or South Korea. Even given the same impetus of productivity improvements and slower population growth, the tropical regions will take longer to release labor from agriculture to alternative uses. Short of fundamentally altering the underlying agricultural technology—which may be biologically infeasible—the tropical countries will be slower to reach the milestones that the success stories of the past achieved.

Acknowledgement

Eberhardt gratefully acknowledges financial support from the UK Economic and Social Research Council [grant numbers PTA-026-27-2048 and ES/K008919/1].

Appendix A. Data construction

The principal data source for our empirical analysis is the Food and Agriculture Organisation's FAOSTAT database (FAO, 2007), from which we obtain annual observations for agricultural net output, economically active labor force in agriculture, number of tractors used in agriculture, arable and permanent crop land and fertilizer use in 128 countries from 1961 to 2002. The total number of observations is 5,162 with an average country series of 40.3. Real agricultural net output (in thousand International \$) is based on all crops and livestock products originating in each country. Intermediate primary inputs of agricultural origin are deducted, including fodder and seed. The quantities for each commodity are weighted by the respective 1999–2001 average international commodity prices and then summed for each year by country. The prices are in international dollars, derived using a Geary–Khamis formula for the agricultural

Table 6
Descriptive statistics

Variable	Mean	Median	Std. dev.	Min.	Max.
Variables in untransformed level terms					
<i>logs</i>					
Output	14.24	14.24	1.71	8.07	19.57
Labor	14.01	14.09	1.84	8.01	20.05
Tractors	9.01	8.87	2.79	0.69	15.51
Livestock	14.90	14.92	1.71	8.80	19.51
Fertilizer	10.82	10.97	2.69	1.61	17.49
Land	14.69	14.78	1.80	6.91	19.07
<i>annual growth rate</i>					
Output	2.3%	2.4%	8.8%	−83.0%	87.6%
Labor	0.3%	0.8%	2.6%	−28.8%	28.8%
Tractors	4.4%	2.0%	9.9%	−121.8%	138.6%
Livestock	1.4%	1.6%	6.4%	−93.3%	182.9%
Fertilizer	5.6%	3.5%	40.1%	−626.3%	393.2%
Land	0.8%	0.1%	3.6%	−41.8%	79.0%
Variables in per worker terms					
<i>logs</i>					
Output	0.23	−0.03	1.42	−2.22	4.00
Tractors	−5.00	−4.97	3.01	−13.67	0.68
Livestock	0.89	0.81	1.38	−2.77	4.63
Fertilizer	−3.19	−2.87	2.67	−11.56	1.95
Land	0.68	0.67	1.15	−2.20	4.95
<i>Annual growth rate</i>					
Output	2.0%	2.0%	9.0%	−80.3%	109.9%
Tractors	4.1%	2.1%	10.1%	−120.2%	136.5%
Livestock	1.2%	1.2%	6.6%	−93.5%	182.9%
Fertilizer	5.4%	4.2%	40.0%	−627.8%	390.8%
Land	0.5%	0.0%	4.1%	−43.0%	81.6%

Notes. We report the descriptive statistics for output (in \$1,000), labor (headcount), tractors (number), livestock (cattle-equivalent numbers), fertilizer (in metric tonnes) and land (in hectare) for the full sample ($n = 5,162$; $N = 128$).

Table 7
Climate Zones following Köppen–Geiger

A	Equatorial climates	Af	Equatorial rainforest, fully humid
		Am	Equatorial monsoon
		As	Equatorial savannah with dry summer
		Aw	Equatorial savannah with dry winter
B	Arid climates	Bs	Steppe climate
		Bw	Desert climate
C	Warm temperate climates	Cf	Warm temperate climate, fully humid
		Cs	Warm temperate climate with dry summer
		Cw	Warm temperate climate with dry winter
D	Snow climates	Df	Snow climate, fully humid
		Ds	Snow climate with dry summer
		Dw	Snow climate with dry winter
E	Polar climates	Ef	Frost climate
		Et	Tundra climate
H	Highland climate	Above 2,500m elevation	

Notes. This classification is taken from [Kottek et al. \(2006\)](#). The Highland category was added after the creation of the Köppen–Geiger classification, with an elevation cut-off of 2,500m suggested in a number of online databases. The [Matthews \(1983\)](#) data has a marginally different classification where As and Ds are not classified and the two polar climates are combined to a single H category—this results in 12 rather than 15 categories.

sector.¹⁹ The labor variable represents the annual time series for total economically active population in agriculture. For capital stock in agriculture we follow a common convention and use total number of agricultural tractors in use as a proxy. The livestock variable is constructed from the data for individual groups of animals. Following convention we use a conversion detailed in [Hayami and Ruttan \(1970\)](#) to convert the numbers for these individual species into the livestock variable. The fertilizer variable represents agricultural fertilizer consumed in metric tons, which includes “crude” and “manufactured” fertilizers. The land variable represents arable and permanent crop land (in 1,000 hectare). Descriptive statistics are presented in [Table 6](#).

Additional time-invariant data on the share of agricultural land by climatic zone from [Matthews \(1983\)](#), available in [Gallup, Mellinger, and Sachs \(1999\)](#)—see [Table 7](#) for details.

Appendix B. Descriptive statistics

Appendix C. Climate zones

Appendix D. Proof of Proposition 1

To begin the proof, first define the following elasticities:

$$\epsilon_A^{c_a} = \frac{\partial c_a}{\partial A} \frac{A}{c_a} \quad (15)$$

$$\epsilon_A^{Y_a} = \frac{\partial Y_a/L_a}{\partial A} \frac{A}{Y_a/L_a} \quad (16)$$

$$\epsilon_A^{L_a} = \frac{\partial L_a}{\partial A} \frac{A}{L_a} \quad (17)$$

$$\epsilon_A^{p_a} = \frac{\partial p_a}{\partial A} \frac{A}{p_a} \quad (18)$$

Market clearing requires that $c_a L = Y_a$, which can be re-written using the production function $Y_a = AL_a^{\beta_L}$ as

$$L_a = \left(\frac{c_a L^{1-\beta_L}}{A} \right)^{\beta_L} \quad (19)$$

Taking the derivative of L_a with respect to A , and allowing for the fact that c_a is a function of A as well, with some manipulation we have

$$\epsilon_A^{L_a} = \frac{1}{\beta_L} (\epsilon_A^{c_a} - 1). \quad (20)$$

Holding that result for a moment, consider that expenditure on agricultural goods must satisfy

$$p_a c_a = \alpha(w - p_a \bar{c}_a + \bar{c}_n) + p_a \bar{c}_n. \quad (21)$$

Rearranging this and using the labor-market clearing condition that $w = p_a Y_a / L_a$ as well as the production function for Y_a yields

$$c_a = \alpha \left(1 + \frac{\bar{c}_n}{w} \right) AL_a^{\beta_L-1} + (1 - \alpha) \bar{c}_a. \quad (22)$$

Taking the derivative of c_a with respect to A , accounting for the fact that L_a is a function of A yields after some manipulation

$$\epsilon_A^{c_a} = \Omega (1 + (\beta_L - 1) \epsilon_A^{L_a}), \quad (23)$$

where

$$\Omega = \frac{\alpha \left(1 + \frac{\bar{c}_n}{w} \right)}{L_a/L}. \quad (24)$$

Solving (23) with (20) yields the following expression for the elasticity of agricultural labor with respect to agricultural productivity

$$\epsilon_A^{L_a} = \frac{\Omega - 1}{\beta_L(1 - \Omega) + \Omega}. \quad (25)$$

This is clearly decreasing in β_L , as claimed in the proposition. Using this result in (23) yields that

$$\epsilon_A^{c_a} = \frac{\Omega}{\beta_L(1 - \Omega) + \Omega}, \quad (26)$$

and this is also decreasing in β_L , matching the second claim of the proposition. From market clearing we have that $Y_a/L_a = (c_a L)/L_a$, which indicates that $\epsilon_A^{Y_a} = \epsilon_A^{c_a} - \epsilon_A^{L_a}$. That makes it straightforward to see that

$$\epsilon_A^{Y_a} = \frac{1}{\beta_L(1 - \Omega) + \Omega}. \quad (27)$$

This is declining in β_L , as claimed in the proposition. Finally, the labor market clearing condition $w = p_a Y_a / L_a$ shows that $\epsilon_A^{p_a} = -\epsilon_A^{Y_a}$, as w is held constant. Thus the absolute value of $\epsilon_A^{p_a}$ is decreasing in β_L , as claimed, given that $\epsilon_A^{Y_a}$ is declining in β_L .

References

- Alvarez-Cuadrado, F., & Poschke, M. (2011). Structural change out of agriculture: Labor push versus labor pull. *American Economic Journal: Macroeconomics*, 3, 127–158.
- Andrews, D. W. K. (2005). Cross-section regression with common shocks. *Econometrica*, 73(5), 1551–1585.
- Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4), 1229–1279.
- Betts, C., Giri, R., & Verma, R. (July 2013). *Trade, reform, and structural transformation in South Korea* (MPRA Paper 49540). Germany: University Library of Munich.
- Chudik, A., Pesaran, M. H., & Tosetti, E. (2011). Weak and strong cross section dependence and estimation of large panels. *Econometrics Journal*, 14(1), C45–C90.

¹⁹ Refer to the Technical Appendix to [Restuccia et al. \(2008\)](#), available on Restuccia's website.

- Coakley, J., Fuertes, A.-M., & Smith, R. P. (2006). Unobserved heterogeneity in panel time series models. *Computational Statistics & Data Analysis*, 50(9), 2361–2380.
- Collier, P., & Dercon, S. (2014). African agriculture in 50 years: Smallholders in a rapidly changing world? *World Development*, 63, 92–101.
- Craig, B. J., Pardey, P. G., & Roseboom, Johannes (1997). International productivity patterns: Accounting for input quality, infrastructure, and research. *American Journal of Agricultural Economics*, 79(4), 1064–1076.
- de Brauw, A., Mueller, V., & Lee, H. L. (2014). The role of rural–urban migration in the structural transformation of Sub-Saharan Africa. *World Development*, 63, 33–42.
- Diao, X., Hazell, P., & Thurlow, J. (2010). The role of agriculture in african development. *World Development*, 38(10), 1375–1383.
- Dorosh, P., & Thurlow, J. (2014). Can cities or towns drive african development? Economywide analysis for Ethiopia and Uganda. *World Development*, 63, 113–123.
- Duarte, M., & Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *Quarterly Journal of Economics*, 125(1), 129–173.
- Eberhardt, M., & Teal, F. (2013a). No mangoes in the tundra: Spatial heterogeneity in agricultural productivity analysis. *Oxford Bulletin of Economics and Statistics*, 75(6), 914–939.
- Eberhardt, M., & Teal, F. (2013b). Structural change and cross-country growth empirics. *World Bank Economic Review*, 27(2), 229–271.
- Eberhardt, M., Helmers, C., & Strauss, H. (2013). Do spillovers matter when estimating private returns to R&D? *The Review of Economics and Statistics*, 95, 436–448.
- FAO (2007). *FAOSTAT. Online database*. Rome: UN Food and Agriculture Organisation.
- Fuglie, K. (2010). Total factor productivity in the global agricultural economy: Evidence from fao data. In P. Pardey, J. Alston, & B. Babcock (Eds.), *The shifting patterns of agricultural production and productivity worldwide* (pp. 63–95). Ames, Iowa: Midwest Agribusiness Trade and Research Information Center.
- Gallup, J. L., Mellinger, A. D., & Sachs, J. D. (1999). *Geography datasets: Agricultural measures*. Harvard University (available online at CID).
- Gollin, D. (2010). Agricultural productivity and economic growth. In P. Pingali & R. Evenson (Eds.), *Handbook of agricultural economics* (Vol. 4, pp. 3825–3866). Amsterdam: Elsevier.
- Gollin, D., Parente, S., & Rogerson, R. (2007). The food problem and the evolution of international income levels. *Journal of Monetary Economics*, 54, 1230–1255.
- Gutierrez, L., & Gutierrez, M. M. (2003). International R&D spillovers and productivity growth in the agricultural sector. A panel cointegration approach. *European Review of Agriculture Economics*, 30(3), 281–303.
- Hamilton, L. C. (1992). How robust is robust regression? *Stata Technical Bulletin*, 1(2).
- Hayami, Y., & Ruttan, V. W. (1970). Agricultural productivity differences among countries. *American Economic Review*, 60(5), 895–911.
- Hayami, Y., & Ruttan, V. W. (1985). *Agricultural development: An international perspective*. Baltimore: Johns Hopkins University Press.
- Herrendorf, B., Rogerson, R. & Valentinyi, A. (2014). Growth and structural transformation. In: Aghion, P. & Durlauf, S. (Eds.) *Handbook of economic growth*, Vol. 2. Elsevier.
- Johnston, B. F., & Kilby, P. (1975). *Agriculture and structural transformation: economic strategies in late-developing countries*. New York, NY: Oxford University Press.
- Johnston, B. F., & Mellor, J. W. (1961). The role of agriculture in economic development. *American Economic Review*, 51(4), 566–593.
- Kottek, M., Grieser, J., Beck, C., Rudolf, B., & Rubel, F. (2006). World map of the Köppen–Geiger climate classification updated. *Meteorologische Zeitschrift*, 15(3), 259–263.
- Martin, W., & Mitra, D. (2002). Productivity growth and convergence in agriculture versus manufacturing. *Economic Development and Cultural Change*, 49(2), 403–422.
- Matsuyama, K. (1992). Agricultural productivity, comparative advantage, and economic growth. *Journal of Economic Theory*, 58(2), 317–334.
- Matsuyama, K. (2009). Structural change in an interdependent world: A global view of manufacturing decline. *Journal of the European Economic Association*, 7(2–3), 478–486.
- Matthews, E. (1983). Global vegetation and land use: New high resolution databases for climate studies. *Journal of Climate and Applied Meteorology*, 22, 474–487.
- McMillan, M., Rodrik, D., & Verduzco-Gallo, Í. (2014). Globalization, structural change, and productivity growth, with an update on Africa. *World Development*, 63, 11–32.
- Mellor, J. W. (1995). Introduction. In J. W. Mellor (Ed.), *Agriculture on the road to industrialization*. Baltimore: Johns Hopkins University Press.
- Mundlak, Y. (2000). *Agriculture and economic growth: Theory and measurement*. Cambridge, MA: Harvard University Press.
- Mundlak, Y., Butzer, R., & Larson, D. F. (2012). Heterogeneous technology and panel data: The case of the agricultural production function. *Journal of Development Economics*, 99(1), 139–149.
- Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967–1012.
- Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, 22(2), 265–312.
- Pesaran, M. H. (2015). Testing Weak cross-sectional dependence in large panels. *Econometrics Reviews*, 34(6–10), 1089–1117.
- Pesaran, M. H., & Smith, R. P. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1), 79–113.
- Restuccia, D., Yang, D., & Zhu, X. (2008). Agriculture and aggregate productivity. *Journal of Monetary Economics*, 55(2), 234–250.
- Schultz, T. W. (1953). *The economic organization of agriculture*. New York, NY: McGraw-Hill.
- Sposi, M. J. (2015). Evolving comparative advantage, sectoral linkages, and structural change. *Globalization and Monetary Policy Institute Working Paper 231*. Federal Reserve Bank of Dallas.
- Tiffen, M. (2003). Transition in Sub-Saharan Africa: Agriculture, urbanization, and growth. *World Development*, 31(8), 1343–1366.
- Timmer, M. P., & de Vries, G. J. (2007). A cross-country database for sectoral employment and productivity in Asia and Latin America, 1950–2005. *GDCC Research Memorandum GD-98*, August 2007.
- Uy, T., Yi, K.-M., & Zhang, J. (2013). Structural change in an open economy. *Journal of Monetary Economics*, 60(6), 667–682.
- Vollrath, D. (2011). The agricultural basis of comparative development. *Journal of Economic Growth*, 16(4), 343–370.
- Wiebe, K., Soule, M. J., Narrod, C., & Breneman, V. E. (2003). Resource quality and agricultural productivity: A multi-country comparison. In K. Wiebe (Ed.), *Land quality, agricultural productivity, and food security*. Northampton, MA: Edward Elgar Publishing.