CENTRE FOR THE STUDY OF AFRICAN ECONOMIES

CSAE Working Paper WPS/2014-21

Agricultural Technology and Structural Change*

Markus Eberhardt[a,b] Dietrich Vollrath[c†]

[a] *School of Economics, University of Nottingham, U.K.*

[b] *Centre for the Study of African Economies, Department of Economics, University of Oxford, U.K.*

[c] *Department of Economics, University of Houston, Texas, U.S.A.*

This Version: 3rd June 2014

Abstract: Using data for 128 countries we document low (high) elasticities of agricultural output with respect to labor in economies within temperate (tropical/highland) climate zones. Adopting a standard model of structural change we show that this technology heterogeneity determines the speed of structural transformation following changes in agricultural productivity and population size. Calibration exercises document shifts in sectoral labor allocation and living standards 2–3 times larger in temperate than in otherwise identical equatorial/highland regions for a given productivity shock. Eliminating technology heterogeneity can account for up to one-fifth of the observed differences in aggregate income per capita across countries.

Keywords: agricultural development, technology heterogeneity, agro-climatic environment, structural change
JEL classification: O47, O11, C23

*We thank conference and seminar participants at the World Bank 'Workshop on Agricultural and Non-Agricultural Productivity Growth and Poverty Reduction,' the 2014 Centre for the Study of African Economies conference and the Nottingham School of Economics for helpful comments and suggestions. The usual disclaimer applies. Eberhardt gratefully acknowledges the hospitality of the Centre for the Study of African Economies, University of Oxford (March 2014) and financial support from the U.K. Economic and Social Research Council [grant numbers PTA-026-27-2048 and ES/K008919/1].

†Correspondence: University of Houston, Department of Economics, 201C McElhinney Hall, Houston, TX 77204, U.S.A. Email: devollrath@uh.edu

1. INTRODUCTION

Developing countries employ a relatively large share of their workers in agriculture, and the labor productivity of those agricultural workers is only a fraction of that found in the developed world. Together, these two facts account for a significant portion of the gap in aggregate output per worker between the developing and developed world (Caselli, 2005; Restuccia, Yang, and Zhu, 2008). Explanations of these patterns generally build on the idea of the "food problem" of T. W. Schultz (1953).[1] Individuals have some subsistence requirements for food, leading to nonhomothetic preferences. Given those preferences, the allocation of labor to agriculture depends on either agricultural total factor productivity (Gollin, Parente, and Rogerson, 2007) or aggregate total factor productivity (Restuccia et al, 2008; Lagakos and Waugh, 2013). The form of the agricultural production function itself is assumed to be homogenous across countries, and it is differences in total factor productivity (TFP) that drive variation in labor allocations and output per worker.

At the same time, there is a long literature emphasizing heterogeneity in the agricultural production function across countries. Hayami and Ruttan (1970, 1985) discuss that simply allowing for differences in the level of TFP between countries is insufficient to capture underlying variation in agricultural *technology*. Technology, referring to the shape of the production function as opposed to its level, is likely to vary across different climate zones or cropping systems (e.g. wheat versus rice). Practically speaking, the elasticity of output with respect to inputs — in particular, labor — will vary between countries based on their agro-climatic conditions.

In this paper we merge this idea of technological heterogeneity with models of structural change. We first show empirically that there are significant differences in agricultural technology across countries. Following Eberhardt and Teal (2013b), we estimate agricultural production functions using a panel of 128 countries with annual data from 1961 to 2002. Our empirical model allows for technology heterogeneity across countries, as well as for cross-sectional correlation between countries and their unobserved TFP (Bai, 2009; Chudik, Pesaran, and Tosetti, 2011). We find that the elasticity of output with respect to labor varies widely by agro-climatic conditions. It is as low as 0.15 in temperate countries, rises to 0.35 in equatorial ones, and can reach as high as 0.55 in highland countries.[2]

To assess the importance of this technological heterogeneity, we set up a standard two-sector model of agriculture and non-agriculture. We show that while variation in technology — captured by the elasticity of output with respect to labor — has no effect on the qualitative predictions of the model, it does have a quantitative impact. For example, the share of labor in agriculture is always negatively related to agricultural TFP, but the magnitude of this relationship differs with the labor elasticity of agricultural production. In temperate economies, with a low elasticity, labor productivity in agriculture is highly sensitive to the number of workers. Following a TFP increase, the release of labor to non-agriculture raises the labor productivity of the remaining agricultural workers by a large amount, requiring very few workers to remain in the agricultural sector while still meeting demand. In contrast, the large labor elasticity in equatorial and highland countries means that labor productivity is relatively insensitive to the number of workers. Releasing labor to non-agriculture after a productivity improvement does not raise the average product of remaining workers much, and so more labor must remain in agriculture to meet demand. Temperate agricultural technology allows for a more rapid

[1]This concept has long been taken as a given in work on agriculture and development, see Johnston and Mellor (1961), Johnston and Kilby (1975), and Mellor (1995). Gollin (2010) provides an overview of this line of thinking, which emphasizes agricultural productivity growth as the key to economic development.

[2]We use several methods to classify a country's climate, all based on on the Köppen-Geiger classification of land within each country: *ad hoc* classifications (e.g. any country with 40% of land in temperate zones is "temperate") and more formal methods (e.g. cluster analysis based on agro-climatic distance) yield similar results. We also investigate the production function excluding livestock breeding and dropping those countries with a high share of livestock in total agricultural output value, with results robust to this alteration.

structural transformation, and hence more rapid gains in income per capita.

We calibrate our model to show the quantitative differences in development that can be attributed to technology heterogeneity in agriculture. We match the model to the experience of South Korea over the period of 1963–2005, when it shifted from having 63% of labor in agriculture to only 8%.[3] Using this, we conduct several experiments. We first examine the effect of a 20% increase in agricultural TFP on several economies that begin with 80% of their labor in agriculture, and differ only in their agricultural technology. For temperate economies, the gain to agricultural productivity reduces agriculture's share of labor by 44 percentage points (to 36%) and increases output per worker by about 75%. In contrast, in an economy with an equatorial/highland technology the agricultural share of labor falls by only 28 percentage points (to 52%) and output per worker rises by only 26%. The temperate zone agricultural economy enjoys a distinct advantage in transitioning out of agriculture in response to productivity improvements.

The situation is not universally favorable to countries with temperate technologies 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 agriculture when population rises, which acts as a decline in productivity. A 5% increase in population will take the agricultural labor share from 80% to 95% in an economy with temperate technology, while the same increase in population will only raise the labor share from 80% to 83% using equatorial/highland technology. The population increase drops real GDP per capita by 13% for temperate economies, but only by 2.5% for those in equatorial/highland regions.

Turning to cross-country comparison, we conduct a counter-factual scenario in which we eliminate agricultural technology heterogeneity across countries. Setting the labor elasticity in agriculture to 0.15 for all countries — equivalent to our estimate for temperate zones — we re-calculate the cross-country dispersion in output per worker. Without technology heterogeneity, the variance of log income per capita drops by one-fifth. The 90/10 ratio reduces from 22 to under 15. This is driven mainly by the implied shifts of labor out of agriculture, with the calibrated model indicating that if equatorial countries *could* use temperate technology, they would shift 10–30 percent of their labor force from agriculture to non-agriculture. We should be clear that this counter-factual does not suggest that equatorial or highland countries are using a sub-optimal technology. Rather, our exercise points out that their technology, while highly productive for their geographic endowment, plays a role in slowing development.

Combined, our results imply that agricultural technology represents an essential element in understanding growth and development. Countries with temperate agricultural technology can experience a more rapid structural transformation than equatorial regions following productivity improvements. If one assumes that learning by doing is particularly strong in the non-agricultural sector, as in Matsuyama (1992), then this advantage may be a source of the wide gap in development between temperate and equatorial areas of the world.

While possible, we are careful to note that our results do not imply "geographic determinism" in development. That is, equatorial countries are not doomed to be poor. Our analysis takes TFP levels as exogenously given, and ultimately those dictate the development level of a country, regardless of agricultural technology. There is nothing in our work that implies that equatorial countries necessarily will have low TFP levels in agriculture or non-agriculture, and nothing that implies the growth of productivity in those sectors is necessarily slower than in temperate countries.

The paper proceeds as follows. After reviewing the relevant literature and motivating technology

[3]We do not use the U.S. to calibrate the model, as is common in the literature, because the U.S. already had a very low share of labor in agriculture in the post-war period. South Korea also has an estimated elasticity of output with respect to labor of 0.34, close to the median in the entire sample.

differences, we proceed to the estimation of the agricultural production function across countries. Having established the variation in the labor elasticity from the estimation, we then present the model and calibration. The final section concludes.

2. LITERATURE

There is a growing literature that specifically attempts to model the shift of labor out of agriculture during the process of economic development. Models by Echevarria (1997) and Kongasmut, Rebelo, and Xie (2001) showed how this shift would result naturally given non-homothetic preferences. Alvarez-Cuadrado and Poschke (2011) use a version of these models to evaluate the relative importance of agricultural and non-agricultural productivity growth in driving structural change, while Duarte and Restuccia (2010) use the model to explain patterns of aggregate labor productivity growth for a sample of countries. Restuccia, et al. (2008) provide an explanation for the large gaps in agricultural labor productivity across countries based on a model with similar preferences. Gollin, et al. (2007) rely on the basic non-homothetic model to show that initial differences in the absolute level of agricultural productivity will determine relative income levels over time.[4]

Studies of long-run or unified growth often explicitly incorporate an agricultural sector to model the shift from Malthusian equilibria to sustained growth. Examples include Kogel and Przkawetz (2001), Hansen and Prescott (2002), Mourmouras and Rangazas (2007), Hayashi and Prescott (2008), Strulik and Weisdorf (2008), and Vollrath (2009). In each of these cases the structural shift out of agriculture is driven by non-homotheticities in demand.

A general result in this literature is that improvements in agricultural productivity (whether exogenous or endogenous) release labor from agriculture and lead to structural change. With appropriate assumptions regarding non-agricultural productivity, this structural change can be an important element in generating sustained increases in income per capita. Variation in agricultural productivity levels will, therefore, potentially lead to sizable differences in aggregate output per capita across countries (see specifically Gollin et al., 2007). However, this literature has not considered differences in agricultural *technology* and their role in determining the aggregate effect of productivity changes.

One exception is Vollrath (2011), who explicitly allows for differences in agricultural technology when studying comparative development levels prior to the Industrial Revolution. He provides evidence on persistent differences in the labor elasticity in agriculture between temperate and tropical regions, and shows that these differences influence living standards in a standard Malthusian model. As in our paper, he finds that high labor elasticity agriculture leads to poorer outcomes.

As noted in the introduction, Hayami and Ruttan (1970) highlight that differences in agricultural technology are likely to be substantial. Despite this warning, though, subsequent estimation of agricultural production functions has tended to assume a homogenous technology for all countries. Hayami and Ruttan themselves made this assumption in their own original work estimating production functions and in a later update (Hayami and Ruttan, 1985), as do Craig, Pardey, and Roseboom (1997), Mundlak (2000), and Martin and Mitra (2002). 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.

[4]Sectoral transitions may alternatively be driven by relative price differences, as in Ngai and Pissarides (2007), rather than non-homotheticies in preferences. Herrendorf, Rogerson, and Valentinyi (2013b) evaluate the two perspectives on structural change and find that non-homotheticities are mainly relevant to changes in expenditure shares, while relative price differences are partly responsible for changes in value-added shares. They do not explicitly discuss the effects on labor shares.

Although parameter heterogeneity of the type we refer to as differences in agricultural technology is very popular in the econometric literature (Pedroni, 2000, 2001; Pesaran, 2006; Pesaran and Smith, 1995), there are relatively few applied studies in the cross-country growth literature allowing for technology differences (Durlauf, Kourtellos and Minkin, 2001; Pedroni, 2007; Eberhardt, Helmers and Strauss, 2013; Eberhardt and Teal, 2013a,b). In many cases the justification for and origins of technology differences across countries are not discussed in great detail and most empirical work has not answered Durlauf et al.'s (2001) call to explain "why this parameter heterogeneity exists" (p. 935). As a result, most growth empiricists still hold with Gollin (2002) who investigated aggregate technology and referred to technology differences as "the least appealing explanation" (p. 461) for observed factor share differences across countries, arguing that "it is not clear why the relationship between inputs and outputs should suddenly shift at a national frontier" (ibid.). In the following we provide some arguments in favour of technology heterogeneity, with a particular emphasis on the case of agriculture, where the relationship between inputs and outputs could well suddenly shift at a geographical (agro-climatic) frontier.

It has long been recognized that agricultural technology, and the labor elasticity specifically, varies across climate zones and types of agricultural systems. Ruthenberg (1976) finds that lowland rice farming and upland dry cereal farming differ distinctly in their production functions. Notably, he writes that the marginal returns to labor in upland farming are "lower and decrease more rapidly with greater employment of labor" (p. 189) when compared to rice farming. This is equivalent to saying that the elasticity of output with respect to labor is lower for upland farming than for rice farming — technology heterogeneity. Conceptual arguments in favor of a heterogeneous production function point to the difference in output structure (wheat *vs* rice *vs* livestock) and in the commercialisation of agriculture (subsistence *vs* industrialised farming), both of which are functions of specialization determined substantially by agro-climatic environment. This heterogeneity in an aggregate agricultural production function implies more than simple differences in unobservable TFP across countries, but rather argues for more fundamental differences in the production function itself. These differences are apparent in the limits to technological spillover from innovations between countries with very different agro-climatic makeup and resource endowments — see Pardey, Alston, and Kang (2012), Fuglie and Nara (2013) and Block (2014) for recent contributions to the analysis of agricultural R&D and knowledge spillovers in the context of developing countries. Much of agricultural technology has to be viewed as location-specific, with attempts at direct technology transfer from one agro-climatic region to another largely doomed to failure (Ruttan, 2002; Eberhardt and Teal, 2013b). Overall, there is not a strong argument for assuming that agricultural technology is similar across countries and/or agro-climatic environments.

The three empirical studies closest to our empirical analysis are Mundlak, Larson, and Butzer (2012), Eberhardt and Teal (2013a) and Eberhardt and Teal (2013b). Building on a theoretical model developed in Mundlak (1988) the empirical applications by Mundlak, Larson, and Butzer (1999) and Mundlak, et al. (2012) argue that agricultural technology differs across countries and over time, although the main focus of the empirics is on pooled fixed effects results which assume a time-invariant, common technology. Eberhardt and Teal (2013a) show that a common factor framework provides a closer empirical representation of the Mundlak model and they relax the assumption of common technology across countries. Their results for agricultural and manufacturing sectors in 40 developing and developed countries over the 1963–1992 period show that a common technology within each sector is rejected by the data. However, sectoral technologies appear to be constant over time within any given country. Eberhardt and Teal (2013b), on which the empirical part of this paper builds, provide estimates of agricultural production functions in 128 countries that explicitly allows for heterogeneity in agricultural technology across countries. A contribution of Eberhardt and Teal (2013b) is the flexible modelling of unobserved agricultural TFP, which is made up of common and idiosyncratic components. The focus of Eberhardt and Teal (2013a, 2013b) is primarily on empirical specification and estimation, providing residual diagnostics and econometric tests of long-run equilibrium relations (cointegration) and plausibly causal interpretation of the estimation equation (weak exogeneity). Although Eberhardt and Teal (2013a) provide some graphical analysis, both studies are more concerned with *establishing* technology heterogeneity across countries, rather than investigating its underlying patterns. In this paper we explore the agro-climatic patterns to technology heterogeneity in great detail, and then establish the importance of this heterogeneity quantitatively for the relative development levels of countries.

3. DATA, EMPIRICAL STRATEGY AND RESULTS

To establish technological heterogeneity across countries and climate zones, we estimate agricultural production functions using cross-country panel techniques developed by Hashem Pesaran (2006) and a number of co-authors. These techniques allow for parameter heterogeneity across countries, which is precisely what we mean when we say "technology heterogeneity." Their methods also allow for cross-sectional dependence and flexible unobserved TFP, alleviating biases that we explain in more detail below. Once we have the estimates of the production function parameters for each country, we will show that these parameters tend to be similar for countries with similar agro-climatic conditions.

3.1 Data

The data for agricultural inputs and output are taken from the United Nation Food and Agriculture Organization's (FAO) FAOSTAT database. This contains annual data for net agricultural output (in real International \$), labor (economically active population in agriculture), capital stock (we follow convention using tractors as a proxy), livestock (in cattle-equivalents constructed from information on various groups of animals), fertilizer (aggregated across various types, expressed in volume terms) and arable and permanent crop land (in hectare) for 128 countries over the 1961–2002 period. More details, including descriptive statistics, can be found in the Data Appendix. As a robustness check we employ total net output and livestock in *value* terms from a more recent version of FAOSTAT: the countries where livestock takes up on average more than 60% of total agricultural output value were then excluded from the analysis, resulting in a sample of 106 countries. The empirical model in this case has total net output less livestock value (in logs) as dependent variable and excludes the livestock (cattle-equivalent headcount) variable.[5]

3.2 Empirical Model

We employ a common factor framework to model agricultural production in country $i = 1, \dots, N$ at time $t = 1, \dots, T$ adopting a (log-linearised) Cobb-Douglas form:

$$y_{it} = \boldsymbol{\beta}_i' \boldsymbol{x}_{it} + u_{it} \tag{1}$$

where $\boldsymbol{x}$ is a vector of observed inputs (labor, capital stock, livestock, fertilizer and land) and y is observed output (all in logarithms). Technology ($\boldsymbol{\beta}_i$) differs across countries but is constant over time.[6] We adopt a transformation of the empirical production function whereby all observed inputs

[5]Detailed results are presented in a Technical Appendix.

[6]Conceptually, $\boldsymbol{\beta}$ may well vary over time within countries. We have estimated equation (1) using various strategies (different start dates, different end dates) but found no systematic variation in $\boldsymbol{\beta}$ over time. Additionally, as will be noted below, there is no tendency for $\boldsymbol{\beta}$ to be correlated with income per capita. Hence we focus on the results that hold $\boldsymbol{\beta}$

and output are expressed in per worker terms: inclusion of the (log) labor variable then indicates the *deviation* from constant returns to scale (increasing or decreasing returns), whereas its exclusion imposes constant returns.

We now introduce the flexible structure of the common factor framework to model unobserved TFP:

$$u_{it} = \alpha_i + \gamma_i' \boldsymbol{f}_t + \varepsilon_{it} = \alpha_i + \gamma_{Si}' \boldsymbol{f}_t^S + \gamma_{Wi}' \boldsymbol{f}_t^W + \varepsilon_{it} \tag{2}$$

The first term on the right-hand side of equation (2) represents the familiar country fixed effects which capture country-specific TFP levels α_i. The second term represents a set of common factors $\boldsymbol{f}$, with country specific parameters (factor loadings) γ_i. These factors, which are orthogonal to each other, come in two flavours, as can be illustrated by the nature of their factor loadings: one type, 'strong' factors ($\boldsymbol{f}^S$), are assumed to affect *all* countries.[7] These factors can be thought of as either structural elements of technical progress (e.g. the process and evolution of agricultural innovation) or global shocks (e.g. the recent global financial crisis or with reference to agriculture the 2007/8 global food price crisis) which affect all economies in the world, but to a different extent.[8] There is assumed to be a fixed number of these 'strong' factors, in line with findings in the macro literature (Stock and Watson, 2002). A second type, 'weak' factors ($\boldsymbol{f}^W$), are assumed to affect a subsample of countries.[9] These factors represent localized effects, such as spillovers between economies in a currency union, or trade and productivity shocks to a small group of economies, possibly in close geographical proximity to each other — loosely we can think of this as spatial correlation. Following the insights of Chudik et al. (2011) our model (and perhaps more importantly our empirical implementation) can accommodate an infinity of these weak factors. The presence of common factors in the model induces correlation across countries, which is econometrically referred to as 'cross-section dependence' (Andrews, 2005). ε_{it} is white noise.

We also specify input demand in order to highlight the pervasiveness of common factors and to address possible simultaneity whereby the empirical production function estimated below may represent a misspecified labor demand or investment function. Our vector of inputs $\boldsymbol{x}$ is specified as

$$\boldsymbol{x}_{it} = \boldsymbol{\eta}_i + \boldsymbol{\Phi}_{Si}' \boldsymbol{f}_t^S + \boldsymbol{\Phi}_{Wi}' \boldsymbol{f}_t^W + \boldsymbol{\Psi}_i' \boldsymbol{g}_t + \boldsymbol{\Upsilon}_i' y_{it-1} + \boldsymbol{\epsilon}_{it} \tag{3}$$

The input equations can be seen to be functions of (some of) the same strong and weak factors which are contained in the output equation, albeit with different factor loadings.[10] This accounts for the endogeneity of production inputs, driven (in part) by unobserved TFP. If $\boldsymbol{\Upsilon}_i' \neq 0$ the above system incorporates feedbacks from output to inputs, potentially leading to biased estimates due to 'reverse causality.' The $\boldsymbol{\epsilon}_{it}$ are again white noise processes.

Finally, we allow for a general evolution of the unobserved common factors which drive all variables in the empirical model:

$$\boldsymbol{f}_t = \boldsymbol{\Pi} + \boldsymbol{\Lambda}' \boldsymbol{f}_{t-1} + \boldsymbol{E}_t \tag{4}$$

and similarly for $\boldsymbol{g}_{it}$. This setup allows for factors to be stationary $|\boldsymbol{\Lambda}_i| < 1$ or nonstationary $\boldsymbol{\Lambda}_i = 1$, thus potentially yielding random walk processes with drifts which are frequently adopted in the literature to represent the evolution of macroeconomic variables.

constant over time.

[7]Thus formally $\gamma_{Si} \neq 0 \,\forall\, i$ and as we let the number of cross-sections go to infinity the *average* of the absolute γ_{Si} converges to a positive constant.

[8]The study by Eberhardt and Teal (2013b) showed that the patterns of agricultural TFP evolution across countries is determined along lines of agro-climatic environment.

[9]Thus $\gamma_{Wi} = 0 \,\forall\, i = 1, \ldots, M$, where $M/N \to 0$ as $N \to \infty$, and as we let the number of cross-sections go to infinity the *sum* of the absolute γ_{Wi} converges to a positive constant.

[10]The additional factors $\boldsymbol{g}_t$ are present to indicate that there may be some idiosyncracy in the evolution of inputs.

We briefly illustrate the identification problem arising when the empirical model in (1) is estimated without accounting for the presence of the common factors. For simplicity we assume a single observed regressor x and a single (strong) factor f which affects both y and x — the implications do not change if we extend this setup to multiple inputs and/or factors.

$$\begin{aligned} y_{it} &= \beta_i x_{it} + \alpha_i + \gamma_i f_t + \varepsilon_{it} && (5) \\ x_{it} &= \eta_i + \phi_i f_t + \psi_i g_t + \epsilon_{it} && (6) \end{aligned}$$

Solving the regressor for the common factor f and plugging into the production function yields

$$\begin{aligned} y_{it} &= \beta_i x_{it} + \alpha_i + \gamma_i \phi_i^{-1}(x_{it} - \eta_i - \psi_i g_t - \epsilon_{it}) + \varepsilon_{it} && (7) \\ &= \underbrace{(\beta_i + \gamma_i \phi_i^{-1})}_{\varrho_i} x_{it} + \underbrace{\alpha_i + \gamma_i \phi_i^{-1}\alpha_i - \gamma_i \phi_i^{-1}\eta_i}_{\varpi_i} + \underbrace{\varepsilon_{it} - \gamma_i \phi_i^{-1}\psi_i g_t - \gamma_i \phi_i^{-1}\epsilon_{it}}_{\varsigma_{it}} && (8) \\ y_{it} &= \varrho_i x_{it} + \varpi_i + \varsigma_{it} && (9) \end{aligned}$$

Since in the standard case $\varrho_i = \beta_i + \gamma_i \phi_i^{-1} \neq \beta_i$ the slope coefficient on our regressor is not identified. In microeconometrics this is the well-known 'transmission bias' (Marschak and Andrews, 1944), whereby a productivity shock unobservable to the econometrician 'transmits' to the factor inputs choices of the firm. While in the microeconometric case the standard solution to this problem involves instrumental variables or control function estimators which require exclusion restrictions for identification, we assume in the macro panel setup that no instrumental variable z exists which is correlated with the endogenous regressor (informative instrument) but uncorrelated with the unobservables (valid instrument). This reflects the notion that unobserved factors are pervasive in the economy and represent the latent driving force behind *all* observable macroeconomic variables (Stock and Watson, 2002). We will detail in the following section how we can tackle this identification problem in the macro panel.

To summarize, our empirical framework allows for unobserved factors to affect inputs and output differentially via the factor loadings. The evolution of the factors themselves is fairly general, including nonstationarity, and the setup provides for globally common effects (strong factors) as well as more localized spillover effects (weak factors). In their impact on inputs and output these global or local shocks are however not constrained to have an identical effect in all countries. Taking an oil price hike as an example of a global shock, this setup allows for a differential impact of the shock between net oil producers and net oil consumers. A more localized shock, for instance a drought in parts of Sub-Saharan Africa, is allowed to impact some countries in the region more severely than others. The total factor productivity term u_{it} is then modelled as a fixed effect α_i and this common factor structure. Finally, and most importantly for the project we are undertaking, we allow for technology heterogeneity β_i across countries, where the patterns underlying this heterogeneity are subsequently shown to be closely linked to local agro-climatic conditions.[11]

3.3 Implementation

Our empirical implementation is focused on recent panel time series estimators which address nonstationarity, parameter heterogeneity and cross-section dependence. As was shown in the previous section the presence of unobserved common factors in this empirical model provides serious difficulties for the identification of the technology parameters β since internal instruments such as lags

[11]We also investigated the possibility of changing β_i over time (see Technical Appendix of Eberhardt and Teal, 2013b) but found relatively strong evidence in favor of time-invariant technology. Our hypothesized source of heterogeneity being time-invariant this finding is perhaps not surprising.

and instruments external to this system are all assumed functions of these factors (Andrews, 2005; Coakley et al, 2006; Phillips and Sul, 2007). The panel time series literature provides two avenues for identification: firstly, using methods developed by Jushan Bai and co-authors (Bai and Kao, 2006; Bai, Kao and Ng, 2008; Bai, 2009), the common factors can be estimated using principal component analysis and then included in the production function equation, which in turn is estimated using panel versions of the Phillips and Hanson (1990) fully modified OLS estimator.[12] This approach requires high-quality data (ideally balanced panels of moderate to large dimensions without missing observations) and crucially depends on the pre-estimation choice of the number of 'relevant' factors for the data at hand, which is difficult in the presence of both strong and weak factors (Bailey, Pesaran and Kapetanios, 2012; Pesaran, 2013). Recent theoretical results further suggest that this approach is unlikely to perform better than the cross-section average augmentation laid out below (Westerlund and Urbain, 2011).

An alternative approach by Pesaran (2006) has been shown to be extremely powerful, providing consistent results in the presence of nonstationary factors, structural breaks, and cointegration or noncointegration of the model variables (Chudik, Pesaran and Tosetti, 2011; Kapetanios, Pesaran and Yamagata, 2011; Pesaran and Tosetti, 2011), while being easy to implement, even in unbalanced panels with missing observations. The common correlated effects (CCE) estimators augment the regression equation with cross-section averages of the dependent ($\overline{y}_t$) and independent variables ($\overline{\boldsymbol{x}}_t$) to account for the presence of unobserved common factors. A further advantage of the CCE estimators over the Bai et al. approach is that although either allows for parameter heterogeneity β_i, only the former allow for explicit investigation of this heterogeneity.[13] In the Common Correlated Effects Mean Group (CMG) estimator, the individual country regression with k independent variables is specified as

$$y_{it} = a_i + \boldsymbol{\beta}_i' \boldsymbol{x}_{it} + c_{0i}\overline{y}_t + \sum_{m=1}^{k} \boldsymbol{c}_{mi}\overline{x}_{mt} + e_{it}, \tag{10}$$

estimated separately for each country i, whereupon the parameter estimates $\hat{\boldsymbol{\beta}}_i$ are averaged across countries akin to the Pesaran and Smith (1995) Mean Group approach.[14] The following lines of algebra provide the intuition why such a simple augmentation with cross-section averages can address the identification problem outlined above. Returning to our simplified empirical model from Section 3.2, we take the cross-section average of the output equation and solve for the unobserved common factor

$$\begin{aligned} \overline{y}_t &= \bar{\beta}\,\overline{x}_t + \bar{\alpha} + \bar{\gamma} f_t \\ & \Leftrightarrow f_t = \bar{\gamma}^{-1}(\overline{y}_t - \bar{\beta}\overline{x}_t - \bar{\alpha}), \end{aligned} \tag{11}$$

where the error term drops out by the standard assumption $\mathbb{E}[\varepsilon] = 0$. Plugging equation (11) back into the output model yields an empirical specification similar to the CMG model in equation (10)

$$y_{it} = \alpha_i - \gamma_i\bar{\gamma}^{-1}\bar{\alpha} + \beta_i x_{it} + \gamma_i\bar{\gamma}^{-1}\overline{y}_t - \gamma_i\bar{\gamma}^{-1}\bar{\beta}\overline{x}_t + \varepsilon_{it} \tag{12}$$

$$= a_i + \beta_i x_{it} + c_{0i}\overline{y}_t + c_{1i}\overline{x}_t + e_{it} \tag{13}$$

The cross-section averages of input x and output y thus can proxy for the unobserved common factor f, while the implementation in form of country-by-country regressions allows for a different coefficient on these averages in each country, mimicking the heterogeneous factor loadings γ_i.

[12] In practice these estimators employ iterations of these steps since the inclusion of estimated factors introduces finite sample bias into the second stage estimates.

[13] In the Bai et al. approach the estimator only provides a mean across heterogeneous technology parameters.

[14] Although $\overline{y}_t$ and e_{it} are not independent their correlation goes to zero as N becomes large.

Recent work by Chudik and Pesaran (2013) suggests that this augmentation approach works surprisingly well even once we relax the assumption that the regressors are strictly exogenous (i.e. $\Upsilon \neq 0$, inducing feedback from y to x) in moderately long T panels such as ours. We will provide evidence that the variables included in our model are indeed weakly exogeneous, i.e. that our empirical results can be interpreted as production functions and do not represent labor demand or investment equations in disguise.

In the standard CCE estimator the weights used to calculate the cross-section averages are the same for all countries ($1/N$). The economic interpretation here could be that the unobserved factors which influence productivity are common to all countries. In a simple empirical extension we apply various weight-matrices prior to taking the cross-section averages ($\overline{x}_t$ and $\overline{y}_t$) to implement alternative scenarios:[15]

(i) The 'Neighborhood Effect', whereby only the performance of contiguous neighbors to country i has a significant effect on the latter's TFP. This is a popular assumption in the regional science literature applying spatial econometric methods.

(ii) The 'Gravity Effect', where geographical distance (a proxy for climatic, soil, cultural and socio-economic differences between countries) is taken as a determinant of the unobserved heterogeneity (e.g. Frankel and Romer, 1999; Redding and Venables, 2004). For country i the observations for countries $j = 1, \ldots, N-1$ are weighted by the inverse of the distance between i and j (itself population-weighted) before computing the cross-section averages.

(iii) The 'Agro-Climatic Distance Effect', where differential agro-climatic characteristics across countries are linked to the success or failure of technology transfer: for instance, countries in the equatorial climate zone will not be able to implement some of the innovations developed for temperate climes. For each country i we weight the observations for countries $j = 1, \ldots, N-1$ by a measure for agro-climatic distance between i and j. Our measure of agro-climatic distance is computed following Jaffe (1986) adopting the Köppen-Geiger classification system. Figure 1 presents the resulting country-specific measure for the case of Kenya. In this example countries marked in green have a similar agro-climatic makeup to Kenya, whereas countries in yellow and orange are more and more different. Countries marked in red do not share any of Kenya's agro-climatic characteristics.

It is important not to overstate the power of our empirical approach: using less than 50 annual observations per country our country-specific estimates will be subject to small-sample bias. The panel time series estimator pursued here relies upon construction of *averages* of technology coefficients across countries — see Eberhardt and Teal (2013a) for a discussion of the sources of bias in individual estimates and the elimination of this bias through averaging. However, given the large size of our sample, we are able to construct fairly precise technology estimates for various sub-samples — along agro-climatic dimensions — in our data.

3.4 Empirical Results and Discussion

We present our production function estimates in Table 1. In each column (except [1]), the reported coefficients for each factor of production are the means across all countries. However, recall that we do have unique estimates of each coefficient for each country in the sample. Here we note some important properties of the global results, while in the next section we discuss the heterogeneity in parameter estimates in detail.

[15] See Appendix for details of the sources and/or construction of the various 'distance' measures.

Columns [1] and [2] report results from the pooled two-way fixed effects (2FE) and robust means from the Pesaran and Smith (1995) Mean Group (MG) estimator.[16] As the literature review in Eberhardt and Teal (2013b) suggests, the former is by far the most popular empirical estimator used in the existing literature on inter-country agricultural production functions. The performance of the MG estimator provides important insights regarding the importance of accounting for both technology heterogeneity *and* cross-section correlation, thus meriting its inclusion. Alongside these we report robust mean estimates for the four versions of the CCE estimator in columns [3] to [6]. The four versions of the CCE estimator vary in the weighting used in constructing the cross-sectional averages: simple mean, neighbors only, distance, and agro-climatic distance.

Looking across the results, it is notable that the 2FE and MG estimates imply large and statistically significant *decreasing* returns to scale (DRS) at the global level, while all CMG estimates with the exception of the distance-weighted version cannot reject constant returns (CRS), so that we only report the CMG models with CRS imposed in these cases.

There are further differences between the pooled 2FE and all other estimates, whereby the former's fertilizer coefficient is roughly double the magnitude and the land coefficient is around 50% larger — all of this translates into an almost negligible labor coefficient. Residual diagnostics reveal serious misspecification, since this model not only yields nonstationary but also highly cross-sectionally dependent residuals. Due to the significant decreasing returns the implied labor coefficient is also very modest in the MG model in column [2], for which residual diagnostics indicate cross-section dependence. Our dataset thus rejects basic assumptions for the consistency of the estimates in both of the 2FE and MG results. The various versions of the CMG estimator in contrast yield stationary and largely cross-sectionally independent residuals. Contrasting MG and CMG results emphasises the importance of dealing with cross-section dependence in order to obtain consistent estimates.

In the Appendix we provide evidence that our preferred CMG specifications (in particular, the standard and agro-climatic distance models) can be interpreted as production functions and do not suffer from bias due to reverse causality using the weak exogeneity tests of Canning and Pedroni (2008). The results strongly support the idea that we are estimating a production function; that is, our parameters capture the elasticity of output with respect to a given input.[17]

Our preferred CMG models are argued to identify the technology parameters by successfully accounting for unobserved productivity. As a further robustness check we add per capita income (in logs, taken from the Penn World Tables, Heston et al., 2009) as additional covariate in our production function model. Adopting income per capita as a proxy for productivity in country i, this specification addresses the concern that estimates of the technology coefficients would be systematically biased if our CMG approach failed to capture unobserved productivity adequately.[18] Results from this exercise are qualitatively unchanged from those in Table 1 – we investigate preferred models [3] and [6] – and if we compare labor coefficients between models with and without the the per capita income covariate we obtain correlation coefficients in excess of .9, respectively.

[16] We adopt robust regression methods (Hamilton, 1992) to compute weighted ('robust') means where they weights indicate whether the observation is close to the mass of the distribution (higher weight) or not (lower weight).

[17] Note that the distance CMG model has favorable residual diagnostics but implies large decreasing returns to scale, which appear implausible. We did however confirm that the patterns of relative magnitudes of technology (labor coefficient) estimates across climate zones, as discussed in detail in the following section, are the same if we use estimates from the distance CMG model in [5].

[18] We include $\ln y^{PWT}$ but not its cross-section average in the CMG models given that income is interpreted as a proxy for (heterogeneous) TFP.

3.5 Technology Heterogeneity Across Climate Zones

The results in Table 1 showed the mean estimated coefficients across all countries, but our primary interest is in the *heterogeneity* in those parameters across different agro-climatic zones. As we have country-level data, our choice of how to map countries into climate-zones will be relevant. We adopt a range of approaches to investigate this. While the methods will vary, the general finding will consistently be that the elasticity of output with respect to labor, $\hat{\beta}_L$, tends to be large in equatorial/highland zones and small in temperate zones.

In all cases we use the estimated labor technology coefficients from the CMG model in [6] of Table 1, which uses agro-climatic-weighted cross-section averages.[19] Our first analysis compares mean estimates for labor coefficients across four broad agro-climatic zones (full sample shares of land in the respective climate zone in parenthesis): arid (19%), temperate and cold (36%), equatorial (37%), and highland (8%) — see Appendix for details on classification.

About one third of the 128 countries in our sample have all of their arable and permanent crop land in one of these four zones; a further 40% have at least some land in two zones, almost 18% in three zones and just under 4% (5 countries) have land spread across all four zones. As we only have a single labor technology estimate for each country regardless of this within-country agro-climatic diversity we pursue a number of alternative strategies to analyse the 'zone-specific' averages. In Table 2 Panel A we begin by averaging coefficients provided the country has *any* land in the specified zone, which inevitably leads to a significant double-counting of $\hat{\beta}_L$ estimates across zones. In Panel B only those estimates are included where the country's share of arable land in the specific climate zone exceeds the zone's sample average,[20] subsequent panels of the Table impose land shares above 40%, 50% and 60%, with sample sizes in the latter reduced substantially for the highland and arid regions. This *ad hoc* approach suggests that labor coefficients are very low in the arid zone (typical countries with substantial land in this zone include countries in the Sahel, in Northern Africa and on the Arab Peninsula), around .15 in the temperate/cold (Europe, North America), around .4 in the equatorial (much of Sub-Saharan Africa) and somewhat higher at 0.56 in the highland zone (Afghanistan, Ethiopia).

In order to provide a more formal analysis and to avoid the issue of double-counting of estimates we conduct cluster analysis based on the measure of agro-climatic distance we constructed for every country pair for our regression analysis above. We carry out the cluster analysis specifying between four and six clusters — more details on the methodology can be found in the Appendix. Table 3 presents the descriptive statistics for each cluster in these exercises, where we highlight the dominant characteristics of a cluster using shading: for instance, in the four cluster case in Panel A we can identify a group of countries which is predominantly in the temperate/cold zone (92% average share of arable land for this group), one is predominantly equatorial (80%), while the remainder are combinations including these two zones (temperate/cold and arid around 40% each; equatorial and highland at 67% and 26% respectively). Moving on to the five and six cluster cases in Panels B and C maintains these four groups and adds an exclusively arid one in the former (72%), and a combination of arid and equatorial (36% and 33% respectively) in the latter. Table 4 presents average labor coefficients for each of the cluster groups for four to six clusters. We conclude from this investigation that countries in the cluster defined by arid and temperate/cold climate (e.g. Chile or Tunisia) have very low average labor coefficients, followed in magnitude by those in an exclusively temperate/cold zone, with coefficient averages significantly higher in equatorial and especially highland zones.

In a similar vein we analyse the correlation between the country-specific means of log inputs and the

[19]Results using the labor technology estimates based on the standard CMG estimator in [3] are qualitatively similar (available on request).

[20]This reveals that most countries which have land in either the equatorial or temperate/cold zones do so to a substantive extent, leading to a skew in the distribution.

country-specific estimated coefficient for that input. Our findings in Table 5 indicate very limited correlation in the preferred CMG models, which implies that there is no systematic relationship between the size of the estimated coefficient on an input and the amount of that input used. Further, there is no systematic relationship between the estimated coefficient on any input and the size of output per worker. It is also worth repeating here that while examining specifications that use different start or end dates for the data, we found no evidence that the value of $\hat{\beta}_L$ varied systematically over time for any climate group. The overall conclusion given by these regressions is that the value of $\hat{\beta}_L$ is not simply picking up differences in development levels, but represents distinct differences in agricultural technology across climate zones.

If farms are operated to maximize profits and there are perfect markets, then heterogeneity in the labor elasticity $\hat{\beta}_L$ should be reflected in heterogeneity in labor's share of output. Unfortunately, there is not good systematic evidence on labor's share in agriculture across countries. It is often presumed that labor shares (and hence the elasticity) are identical across countries, based on very limited data. A study by Jorgenson and Gollop (1992) for the U.S. identifies the agricultural elasticity on capital to be 0.21 and the elasticity for land to be 0.19. Using this, many studies make three assumptions. First, they presume that the elasticity on labor is therefore 0.60 (ignoring livestock and fertilizer). Second, they assume that U.S. agricultural markets are competitive and farmers are profit-maximizers, and therefore labor's share in output is 0.60 as well. Third, given a lack of data they assume that this share is the same in all countries. The first assumption is likely to be incorrect, as the estimated labor elasticity for the U.S. is much lower than 0.60, being close to our temperate value of 0.15.

The third assumption is typically made because few sources of labor share information exist. The limited evidence on labor's share in agricultural output has been culled from several different studies by Fuglie (2010). The labor share for the U.S. he reports is only 0.20, while other countries that fall in our temperate cluster have values such as 0.19 (several former Soviet countries), 0.23 (South Africa), and 0.30 (the U.K.). These are not only well below 60%, but are not terribly different from our average estimated elasticity for these countries of about 0.17. This can be compared to several developing countries in our equatorial cluster that have distinctly higher values: India (0.46), Indonesia (0.46), and Brazil (0.43). These values are all close to the average estimated elasticity for this 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. Given the fact that most of these countries are likely operating with something short of perfect markets, one cannot assume that labor's share and the labor elasticity should necessarily be equal.

A tentative conclusion from our analysis would be that we found evidence for a systematic pattern relating labor technology coefficients to the predominant agro-climatic environment in which their agricultural sector operates, whereas there is little evidence that the heterogeneity is largely driven by income or magnitudes of input or output. These findings are robust to the adoption of a revised production function model which excludes our measure of livestock and models agricultural output adjusted for the value of livestock output — results are presented in a Technical Appendix. In the following section we investigate the implications of technology heterogeneity for the speed and extent of structural change.

4. IMPLICATIONS OF TECHNOLOGY HETEROGENEITY

To understand the possible impact of heterogeneity in labor elasticities we use a simple model of the process of structural change and development. The model stays close to standard versions used across the literature to study the transition from agriculture to industry and overall 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.[21] The model shares its structure with that found in recent work by Duarte and Restuccia (2010), Alvarez-Cuadrado and Poschke (2011), and Herrendorf, Rogerson, and Valentinyi (2013a). Within the model, agricultural output is produced using land and labor, and agricultural technology is captured by the elasticity of output with respect to labor. As we documented in the previous section, this value varies considerably across countries and climate zones. Once we have laid out the model, we calibrate it and then use the calibrated version of the model to examine several experiments with productivity and population changes under different values for the labor elasticity of agricultural production.

4.1 Production

The log-linear agricultural production function from equation (1) can be written as

$$y_{it} = \beta_{Li} \ln L_{a,it} + \boldsymbol{\beta}_i' \boldsymbol{x}_{it} + u_{it} \tag{14}$$

where we have broken out the term involving agricultural labor (L_a) from the remaining vector of inputs ($\boldsymbol{x}^*$). u represents unobserved productivity (TFP). Note that the coefficient on labor is specific to country i, β_{Li}. Exponentiating we have

$$Y_{it} = A_{it} L_{a,it}^{\beta_{Li}} \tag{15}$$

where $A_{it} = \exp(\boldsymbol{\beta}_i' \boldsymbol{x}_{it} + u_{it})$ is the collection of other inputs and unobserved productivity. A is what we will refer to as "productivity" in agriculture, but note that it incorporates not only TFP but also the use of non-labor inputs such as fertilizers, livestock, or capital.

For ease of notation in the rest of the model we will suppress the i and t subscripts, noting that the model is meant to hold for any country i, and at all times t. This leaves us with a production function in agriculture of

$$Y_a = A L_a^{\beta_L}. \tag{16}$$

We denote agricultural output by Y_a so that it is easily distinguished from aggregate and non-agricultural output, Y and Y_n respectively. In addition to the agricultural sector, there is a non-agricultural sector that is assumed to be linear in labor, L_n, with

$$Y_n = w L_n \tag{17}$$

where w is labor productivity in that sector.

A simple adding-up constraint holds for total workers L such that

$$L = L_a + L_n. \tag{18}$$

4.2 Preferences and Individual Optimization

There are L individuals in this economy with preferences of,

$$U = \alpha \ln(c_a - \overline{c}_a) + (1 - \alpha) \ln(c_n + \overline{c}_n), \tag{19}$$

[21] This endowment is often thought of as representing some kind of non-tradable home production.

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

$$w = p_a c_a + c_n \tag{20}$$

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

$$\begin{aligned} p_a c_a &= \alpha(w - p_a\bar{c}_a + \bar{c}_n) + p_a\bar{c}_a \\ c_n &= (1-\alpha)(w - p_a\bar{c}_a + \bar{c}_n) - \bar{c}_n. \end{aligned} \tag{21}$$

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 (as in the second).

4.3 Equilibrium Allocations of Labor

We assume that labor moves freely between sectors to equalize earnings. In non-agriculture the wage is simply w, the marginal product. 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. This has been a common assumption in the development literature, often used to capture a "dual economy" feature of the economy. 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

$$w = p_a \frac{Y_a}{L_a}. \tag{22}$$

To solve for the labor allocations the last thing we need is that supply and demand are equal in both sectors, or

$$\begin{aligned} c_a L &= Y_a \\ c_n L &= Y_n. \end{aligned} \tag{23}$$

Given the static nature of the model, an equilibrium is straightforward:

Equilibrium: *Given productivity levels A and w, a 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/L and L_n/L, and a price of agricultural goods relative to non-agriculture (p_a), such that expenditures are optimal as in (21), the wage is equalized across sectors as in (22), and all labor is utilized (18).*

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 tedious but 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 decline in productivity, and the predicted responses would simply flip sign, as is typical in these models.

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.

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.

When β_L is large, the effect of a technology 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 $\beta = 1$, 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 that sector in order to meet demand.

It is easiest to see the mechanism at work by limiting individuals to consuming precisely $\overline{c}_a$ 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 $\overline{c}_a$ in agricultural goods, setting supply equal to this demand yields

$$\overline{c}_a L = A L_a^{\beta_L}. \tag{24}$$

This can be rearranged to yield the following relationship

$$\overline{c}_a = \left(\frac{L_a}{L}\right) A L_a^{\beta_L - 1}. \tag{25}$$

The per-capita demand on the left is fixed, and on the right is the share of workers in agriculture times their labor productivity. An increase in A requires that L_a must fall, as there is no increase in demand for food. How much does L_a have to fall? That depends on the size of β_L, which dictates how quickly labor productivity changes with L_a. With a low β_L, the production function in agriculture is "more curved", and labor productivity of the remaining workers rises very quickly as L_a falls. Thus a larger drop in L_a is necessary to keep the equation balanced. Re-arranging again, one can find that

$$\frac{L_a}{L} = \left(\frac{\overline{c}_a L^{1-\beta_L}}{A}\right)^{1/\beta_L}, \tag{26}$$

as an exact 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.

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. This is why our proposition does not necessarily imply any kind of "geographic determinism" in development.

What Proposition 1 shows is that structural change and development will proceed more quickly in temperate areas for a given shock to agricultural productivity. If temperate areas begin with some kind of advantage in productivity, A, then technological heterogeneity will tend to widen that gap. If temperate areas have lower productivity to begin with, then technological heterogeneity will allow them to catch up to the frontier quickly. But heterogeneity does not offer any information on whether temperate areas will have either the advantage or disadvantage to begin with.

4.4 Heterogenous Technology and Structural Change

To evaluate the importance of heterogeneity in β_L on the pace of structural change, we calibrate our model and use it to conduct several counter-factual experiments. We use data from South Korea in 1963 and 2005 to calibrate the model. We use South Korea, rather than the United States as is common in work on structural change, 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 globalising world economy. In 1963 63% of South Korea's workforce was employed in agriculture, 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.[22] 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.

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), $\overline{c}_a$ (the subsistence constraint), and $\overline{c}_n$ (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

[22] Authors' calculations using data from Timmer and de Vries (2007).

$\alpha = 0.01$, $\bar{c}_a = 0.802$, and $\bar{c}_n = 3.182$.[23] Given these values, our first experiment is designed to show the quantitative impact of differences in β_L. The question we want to ask is: for otherwise identical economies, how does the value of β_L affect their response to either an increase in agricultural productivity or the size of population?

To do this, we simulate three separate economies, with values of $\beta_L \in (0.15, 0.35, 0.55)$. These values correspond to the estimated values of the labor elasticities in temperate, equatorial, and highland clusters of countries in Table 4. We impose that each economy begins with an agricultural labor share of 80%. To do this, we adjust the initial level of population, L, so that the equilibrium outcome of the model with our calibrated parameters is precisely $L_a/L = 0.8$.[24] The variation in population size is necessary because β influences the level of output in the agricultural sector directly, in addition to affecting how fast labor can be released. Having adjusted population size, the result is that all the other interesting outcomes in the model are 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 economies in the initial set-up.

Starting from these identical initial conditions, we then impose an exogenous increase of 20% to agricultural TFP (A) for each economy. The results of these calculations can be found in Table 6. There, the equilibrium outcomes for each of the three simulated economies are shown, each relative to the baseline listed in the second column. The third column of the table is for the economy with $\beta_L = 0.15$, the temperate value. With a productivity increase of 20% in 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 0.52 or 0.60. The temperate economy is able to shift 15-23% more of the labor force out of agriculture in response to the productivity increase.

This advantage for the temperate economy shows up in all the other relevant characteristics of these economies. 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 it's 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, rising by factors of 1.6 and 1.4, respectively.

We also measure how the individual consumption baskets and an index of real income per capita respond to the improvement in agricultural productivity. Agricultural consumption rises by about 7% 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.[25]

Non-agricultural consumption, c_n, shows the real substantial effects of the productivity increase. The

[23]One may note that we calibrate by setting $L = 1$ in 1963. In this case the normalization is inconsequential because we are considering a single economy that has a single value of β_L.

[24]The actual population levels are $L = 1.18$ when $\beta_L = 0.15$, $L = 1.16$ when $\beta_L = 0.35$, and $L = 1.13$ when $\beta_L = 0.55$.

[25]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.

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 more than three-fold 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 this economy both before and after the increase in productivity, using the initial relative price of food to value goods. Recall that this initial relative price is identical in all economies, so that these are comparable real GDP numbers. The final row shows the outcomes following the increase in agricultural productivity. For $\beta_L = 0.15$, real income per capita rises by about 48%. This compares to 31% for the economy with $\beta_L = 0.35$ and 22% when $\beta_L = 0.55$. Measured income per capita rises nearly twice as much for the low-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 quantitative magnitude of the benefits. The lower is β_L, the more an economy can benefit from agricultural productivity improvements.

The logic runs in reverse, however, when there is growth in population, which acts like a decline in agricultural productivity. Again, a low-β_L economy will respond more severely to the change, but now this change will pull labor into agriculture and lower living standards. Referring back to equation (26), note that with a low elasticity, the share of workers in agriculture is more sensitive to an increase in population, L. From a Malthusian perspective, having a low β_L means that the detrimental effect of a limited supply of land is more severe.

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 6 shows the equilibrium outcomes 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. In the low-β_L economy, the average product of labor falls very quickly as new labor joins the agricultural sector, requiring even more labor to enter to meet 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 just over 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.

Table 6 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 equatorial and highland areas, are more 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.

4.5 Country-specific Growth Prospects

Variation in β_L implies that countries differ in how they respond to increases in productivity. To see more clearly the quantitative impact of β_L, we calculate two different counter-factual outcomes for a sample of 78 countries taken from Caselli (2005).[26] As a baseline for each country, we normalize $L = 1$ and $w = 1$, and solve for the level of agricultural TFP (A) that is consistent with the observed agricultural labor share (L_a/L).

Each country is assumed to have a value of β_L equal to the average value for their climate zone. There is a significant amount of heterogeneity in agricultural technologies in the sample. A total of 25 countries fall within the equatorial cluster we identified, with an average $\beta_L = 0.35$. A further 11 countries are in the equatorial/highland cluster, with a value of $\beta_L = 0.55$. The remaining 42 countries are either in the arid/temperate or the temperate cluster, all with a value of $\beta_L = 0.15$. While most of the rich countries are in the temperate cluster, there is wide variation in output per worker within each group, as will be seen in the figures to follow.

The first counter-factual calculation we make is to solve for the increase in agricultural TFP necessary to drive the agricultural labor share down to 3%, a level roughly consistent with that seen in developed nations. We then plot the ratio of counter-factual to baseline agricultural TFP against the observed output per worker. In Figure 2, we have separately denoted countries by their climate groupings. Across all climate groups, one can see that the relative agricultural TFP necessary to reach $L_a/L = 0.03$ declines with output per capita. Simply put, richer countries are closer to having only 3% of their labor force in agriculture, so smaller gains to agricultural TFP are necessary. Here, one can also see clearly the fact that output per worker and climate group are not highly correlated. There are rich and poor countries within each group.

However, there are distinct level differences between the climate groups. Consider the comparison of Mexico and Argentina. In the data, they have very similar levels of output per capita. However, Mexico is in the highland climate zone ($\beta_L = 0.55$), while Argentina is in the temperate zone ($\beta_L = 0.15$). Because of agricultural technology, in Mexico agricultural TFP must rise by a factor of 25 to reach $L_a/L = 0.03$, while in Argentina it must only rise by approximately a factor of 6. An even more stark comparison is between the poorest countries in each climate group: Ethiopia, Tanazania, and Malawi. Ethiopia, in the highland climate zone, requires a relative TFP increase of 40 to reach a 3% agricultural labor share, in Tanzania the necessary increase is half as large, and Malawi requires only a 10-fold increase in agricultural TFP. Malawi is much closer, so to speak, to reaching rich-country levels of agricultural labor, even though today it is roughly equal in output per capita to Ethiopia and Tanzania.

A second counter-factual calculation we make is to solve for the relative increase in TFP in both sectors that is necessary to double output per worker. That is, by how much do we have to scale up both A and w in order to double output per worker? Figure 3 plots the relative TFP against baseline log output per worker. As can be seen, regardless of the climate group the relative TFP necessary to double output per worker rises with current output per worker. This is because we always value agricultural output at the baseline's relative price to ensure comparability between the baseline and the counter-factual. As the baseline has a higher relative price for agricultural goods, it is not necessary to double TFP to double measured output per worker. This effect is strongest for the poorest countries, who have the largest share of agricultural workers to begin with, and experience the largest shifts of

[26]These countries are those from Caselli (2005) for which we have an estimate of their labor elasticity in agriculture.

labor in response to TFP increasing. Hence the poorest countries need a smaller relative TFP increase to double measured output per worker.

More importantly, there are distinct level differences between the climate groups because of our assumptions regarding β_L for each. In this counter-factual, Argentina (temperate) and Mexico (highland) come out essentially identical, needing TFP to increase by a factor of about 1.9 in order to double output per worker. Once countries are relatively rich, the distinctions between climate groups dissipate. However, there is still wide variation at the low end of the scale. For the three poor African economies, there are distinct differences in the TFP gains necessary to double output per worker. In Ethiopia, the value is 1.7, or a 70% increase in TFP. In Tanzania the relative gain necessary is only about 1.5 (a 50% increase), while in Malawi it is 1.4 (a 40% increase). Again, temperate economies have an advantage in development, as the low value for β_L allows labor to be released quickly from agriculture as productivity grows.

It is important to note that the distinctions between climate groups in Figures 2 and 3 are by construction. The point of the Figures is to demonstrate the quantitative effect of variation in β_L on the ability of economies to hit some counter-factual development targets. The variation seen within any given climate group is due to the imperfect empirical relationship between output per worker and the agricultural labor share in the data, and not to any additional variation we included in the model.

4.6 Cross-country Income Distribution

The prior sections show that countries will develop at different rates due to heterogeneity in agricultural technology. However, the counter-factuals shown above did not explicitly address the role that agricultural technology plays in explaining the currently observed cross-country dispersion in output per worker. Here we want to use our calibrated model to ask how much of observed variation in output per worker is attributable to agricultural technology heterogeneity. More specifically, how much variation in output per worker remains once we counter-factually eliminate heterogeneity in agricultural technology?

A question like this is standard in the development accounting literature. But the fact that we are looking explicitly at technology heterogeneity (as opposed to productivity heterogeneity) complicates the situation. We have an "apple and oranges" problem (Bernard and Jones, 1996; Martin and Mitra, 2002; Eberhardt and Teal, 2014) in comparing countries in different climate groups with different values for β_L. The value of A that we back out of a production function for a temperate country with $\beta_L = 0.15$ will not be the same as the value of A for an equatorial country with $\beta_L = 0.35$, even though they have identical agricultural labor shares, populations, and subsistence constraints. Moreover, without knowing the units of measurement, we cannot even say which country – temperate or equatorial – would have the higher value of A. Doing counter-factuals where we set productivity levels equal in these two countries are thus not meaningful (Martin and Mitra, 2002). To proceed, we need some kind of "anchor" for the value of agricultural TFP, a value of A such that all countries would achieve the same value of L_a/L, regardless of climate group.[27]

In this section we do this by assuming that $L = 1$ for all countries. Note that the model is already written under the assumption that land, X, is equal to one. So assuming that $L = 1$ is equivalent to normalizing population density to one for all countries. Under this assumption, we then have an "anchor" for agricultural TFP. With $L/X = 1$ in each country, then it is the case that if $A \to \bar{c}_a$, the agricultural labor share goes to one, regardless of climate group. Values of A higher than $\bar{c}_a$

[27] It is worth pointing out that the analysis in the prior section does not have this issue. There, each country was compared only to itself, and we were not trying to compare the agricultural TFP of one country (apple) to that of another (orange).

definitively produce values of $L_a/L < 1$ for all countries, but the exact value of L_a/L will depend on β_L. This is one choice of normalization, and other choices would yield different results in our cross-country counter-factuals, some leading to a larger role for β_L, some a smaller role. While we do not believe our choice has a demonstrable bias one direction or the other, we fully acknowledge that our choice of normalization is *ad hoc*, and so the following results must be taken as suggestive only.

With that caveat aside, we can return to the two questions regarding cross-country variation in output per worker. To begin, we first establish the relevant facts in the observable data. Table 7 shows in the first row several measures of dispersion for the sample of 78 countries from Caselli (2005). The variance in log output per worker for these countries is 1.185, and the ratio of the 90th to the 10th percentile in the distribution is about 21. Looking solely at agricultural labor productivity in the last two columns, one can see a much higher variance, and a 90/10 ratio of approximately 46. This sample captures the basic fact noted by Caselli (2005) and Restuccia, et al. (2008), namely that variation in labor productivity is much higher in agriculture than in non-agriculture.

Now we ask how much variation in output per worker remains if we eliminate heterogeneity in agricultural technology. We hold the individual TFP levels in agriculture (A) and non-agriculture (w) constant at the observed levels, and re-calculate the equilibrium allocation of labor if we set the labor elasticity in agriculture to the temperate zone value of $\beta_L = 0.15$. Using this allocation of labor, we can then re-calculate output per worker and agricultural labor productivity. This allows us to examine the distribution of both once technological heterogeneity in agriculture is eliminated. This exercise does *not* imply that switching to the temperate technology is optimal or even possible. We have no evidence to suggest that equatorial or equatorial/highland regions have deliberately chosen to ignore a higher-productivity technology. Their technologies are almost certainly optimal for their given agro-climatic conditions. However, these technologies do influence the rate of development, as we have shown above. Our counter-factual exercise is aimed at discovering just how large that influence may have been.

The second row of Table 7 shows the measures of dispersion in our counter-factual scenario. The variance of log output per worker falls to 0.996, which is only 83% of the actual value, implying that agricultural heterogeneity alone can explain almost one-fifth of the overall variation. The 90/10 ratio shows a similar reduction, from over 21 to under 15. This reduction takes place because a number of very poor countries in the equatorial and equatorial/highland clusters would be significantly richer if they instead operated a temperate zone technology. Table 8 shows the results from all 36 of the countries in those two equatorial regions. The first two columns show results for output per worker, first the actual value taken from Caselli (2005) and then our counter-factual value calculated using our calibrated model and a value of $\beta_L = 0.15$. If a country such as Angola had temperate zone agricultural technology, it would see an increase in output per worker of about 84% (from \$4,558 to \$8,405). This is a particularly dramatic gain, and other countries have smaller implied increases. Brazil's gain is about 21%, Ghana's 48%, and Uganda's 63%. For countries in the equatorial/highland region, the implied differences in output per worker are even more dramatic. In Ethiopia, income per capita essentially doubles.

Moving back to Table 7, one can see how the dispersion of agricultural labor productivity changes in our counter-factual scenario. Here, there is little to no obvious change. The variance of log labor productivity barely changes, and the 90/10 ratio drops from 46 to 44. Why does labor productivity in agriculture not show the same change as output per capita? The answer is that the variance and 90/10 ratios are a little too crude to pick up the changes in the distribution of labor productivity across countries. Figure 4 plots the kernel density of agricultural labor productivity in the actual data and in our counter-factual with temperate technology. As can be seen, there is a distinct shift, with the mass of density at log productivity of 7 (approximately \$1,100) being shifted to the right. There is thus a distinct improvement in agricultural labor productivity in the distribution. However, the variance of this new distribution is not demonstrably different from the original one, and the tails of this distribution are in essentially the same places. Hence Table 7 is missing out on this shift of mass towards higher labor productivity.

One can see this in Table 8, where columns (3) and (4) show the actual and counter-factual levels of agricultural labor productivity. In many cases, the counter-factual levels are double the actual ones, or more. For Malaysia and Mexico, agricultural labor productivity would be nearly four times larger with temperate zone technology. Despite these results, there are a number of poor countries outside the equatorial clusters and therefore unaffected by our counterfactual exercise (e.g. Malawi and Burkina Faso) that have very low productivity levels in agriculture. Hence when looking at summary statistics such as the variance or the 90/10 ratio, there does not appear to be a large change due to technology heterogeneity.

The reason that the dispersion in output per worker drops so much, although the dispersion in agricultural labor productivity does not, is because of the implied shifts in agriculture's share of labor. Figure 5 plots the density of this labor share in the actual data and in our counter-factual. As can be seen, there is a distinct shift towards lower labor shares in agriculture with temperate technology. From a bi-modal distribution with mass around 20% and 80%, the counter-factual implies a much stronger mode around 20%, and there is no longer a distinct second mode at 80%. So even if agricultural labor productivity variation does not go down much, the substantial shift of workers to non-agriculture generally raises output per worker in the equatorial countries, lowering variation across countries. In Table 8, one can see in the final two columns the implied change in the labor share in agriculture when using temperate technologies. These changes are substantial in many of the equatorial countries and generally imply between 10-20% of the labor force shifting out of agriculture if temperate technology is used. In the equatorial/highland countries, changes are even more dramatic, with declines on the order of 20-30%.

Our counter-factuals suggests that heterogeneity in agricultural *technology* may be an important element in understanding why some countries continue to lag behind the frontier. Even though they may be able to take advantage of productivity improvements in the agricultural sector, countries with high labor-elasticity technologies (i.e. equatorial regions) will not be able to shift labor into non-agricultural work rapidly, retarding their development. We think it is worth reiterating, however, that our counter-factuals do not imply that any given equatorial country would benefit from adopting a temperate zone technology, even if it were feasible. That would certainly imply some kind of shift in the level of productivity (A) — most likely a severely negative one — that we have not allowed for here. The fact that we cannot possibly implement a change in agricultural technology, though, does not imply that technology is unimportant in explaining cross-country output differences.

4.7 Trade and Agricultural Technology

We have presumed that countries are closed to trade in both agriculture and non-agriculture. The influence of β_L on the allocation of labor across sectors does not depend on this assumption. If a country is open to trade in both goods, then the labor allocation decision is separated from the consumption decision. With a given international relative price for agriculture of p_a^*, then using the wage equalization condition in (22) we have

$$L_a = \left(\frac{p_a^* A}{w}\right)^{1/(1-\beta_L)}. \tag{27}$$

With an open economy, though, note that L_a is *positively* related to agricultural productivity, A. An increase in A shifts comparative advantage towards agriculture and the economy shifts labor towards that activity. The elasticity of this relationship is $1/(1-\beta_L)$, which is increasing in β_L. In other words, the shift of labor *into* agriculture with an increase in A is stronger in the equatorial/highland clusters than in temperate zones. The explanation is similar to the closed economy case. With a large β_L, the average product of labor does not drop much as labor is moved into agriculture, so many workers are shifted into that activity when A goes up. For temperate countries with low β_L, the average product of labor declines quickly as labor shifts into agriculture, so only a few workers shift when A rises.

In a similar model to ours, Matsuyama (1992) proposed that a comparative advantage in agriculture would lead to slower growth if there were learning-by-doing effects in non-agriculture. If we take this seriously, 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. More recently, several papers have concerned themselves with the importance of trade for structural change, typically using South Korea as a primary example (Sposi, 2012; Teignier, 2012; Betts, Giri, and Verma, 2013; Uy, Yi, and 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 labor force. Heterogeneity in agricultural technology across countries does not invalidate any of this work. However, our findings do imply that their models would have very different quantitative predictions for countries that do not share the same agricultural technology as South Korea.

5. CONCLUSION

Agricultural technology — specifically the elasticity of output with respect to labor — varies widely across countries. In particular, it tends to be correlated with agro-climatic conditions: our empirical work shows that temperate areas tend to have low elasticities, around 0.15, while equatorial areas (0.35) and highland equatorial areas (0.55) have much higher values. We exploited several new techniques in panel time-series estimation to control for unobserved productivity to establish that there is heterogeneity in the technology parameters across countries.

We examine the effect of this technology heterogeneity using a standard two-sector model of the structural transformation. Calibrating the model using South Korean data from 1963–2005, we first examine the effect of productivity and population shocks on countries with different agricultural technologies. When agricultural productivity increases, low-elasticity countries in temperate zones experience large shifts of labor out of agriculture, as well as distinct increases in labor productivity in agriculture and aggregate output per capita. Although equatorial regions with high elasticities also see shifts of labor and increases in output per capita, these are quantitatively much smaller. The opposite occurs for a population shock, with temperate zones reacting more severely by shifting labor back into agriculture. In short, temperate zones are much more sensitive to shocks to productivity and population. This gives them the potential to develop much more quickly, assuming that productivity gains outpace population increase. We show that across countries, agricultural technological heterogeneity can account for about 20% of the variance in log income per capita.

Our work implies that there are significant effects of climate on development. However, we must be very clear on what precisely those effects are. Our findings are not an example of "geographic determinism", the idea that tropical 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 model. The fact that the agricultural technology varies by climate (the β_L coefficient) does not imply anything about productivity or population levels. This is consistent with our finding that the estimated β_L coefficients are only weakly correlated with income per capita in our sample: climate does not necessarily predict which countries will be rich and which will be poor.

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 due to some barrier to productivity improvements. 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 inappropriate to measure the current development trajectory of Sub-Saharan African, Central American, or South Asian countries against earlier developers in temperate regions such as South Korea or Japan. 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 impossible — the tropical countries will be slower to reach the milestones that the success stories of the past achieved.

FIGURES

Figure 1: Agro-climatic ‘distance’ — the view from Kenya

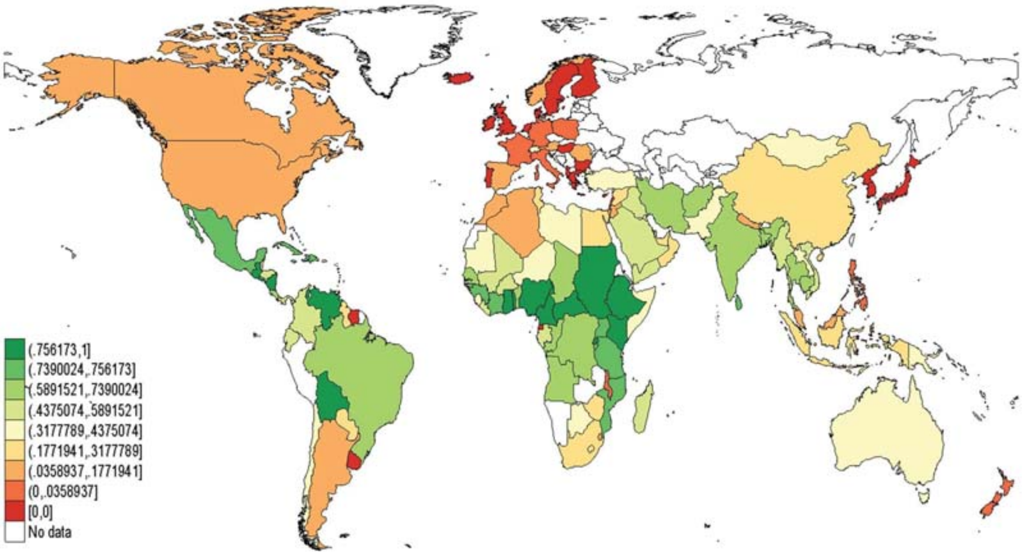

Notes: The map provides an illustrative example of the ‘agro-climatic distance’ measure.

Figure 2: Relative Agricultural TFP Increase to Reach $L_a/L = 0.03$

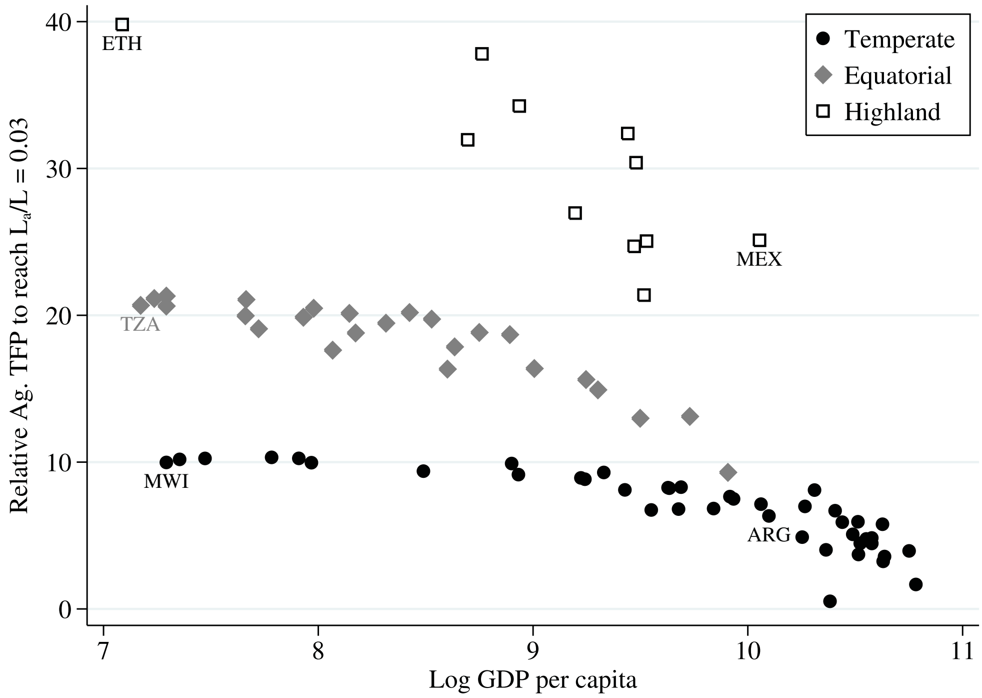

Notes: The figure shows the ratio that agricultural TFP (A) would have to increase by to reach $L_a/L = 0.03$ in each country. The 78 countries are from Caselli (2005), who provides the starting level of L_a/L and output per capita. The classification of each country into a climate zone is described in the text.

Figure 3: Relative TFP Increase to Double Output p.c.

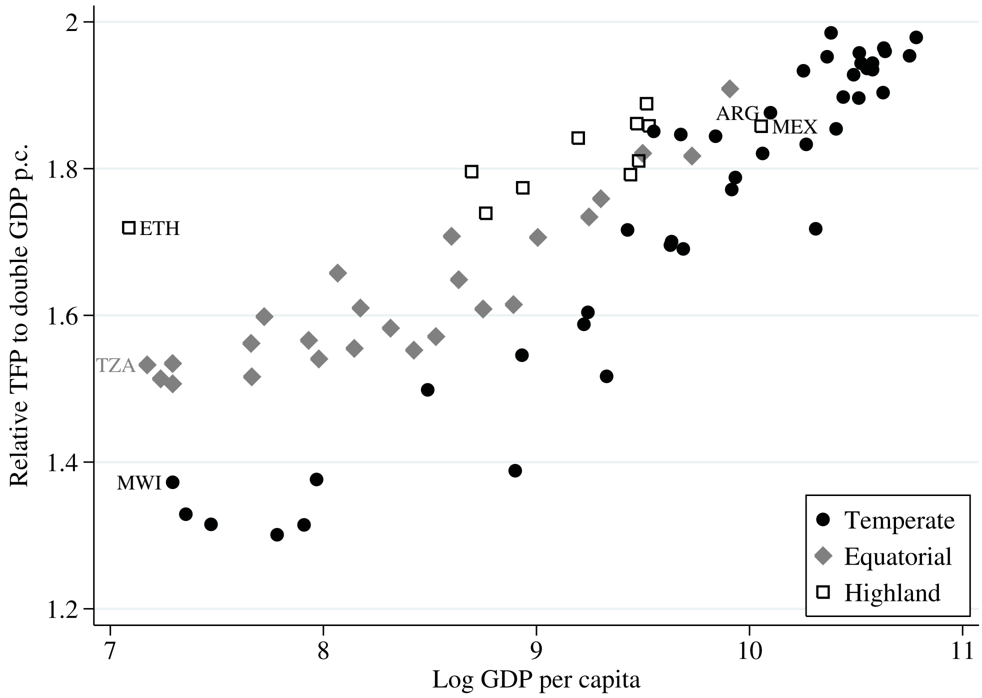

Notes: The figure shows the ratio that TFP in both sectors (A and w) would have to increase by to reach double real output per capita in each country. Output per capita is valued using the initial relative price of agricultural goods, p_a. The 78 countries are from Caselli (2005), who provides the starting level of output per capita. The classification of each country into a climate zone is described in the text.

Figure 4: Agricultural Labor Productivity — Actual and with Temperate Technology

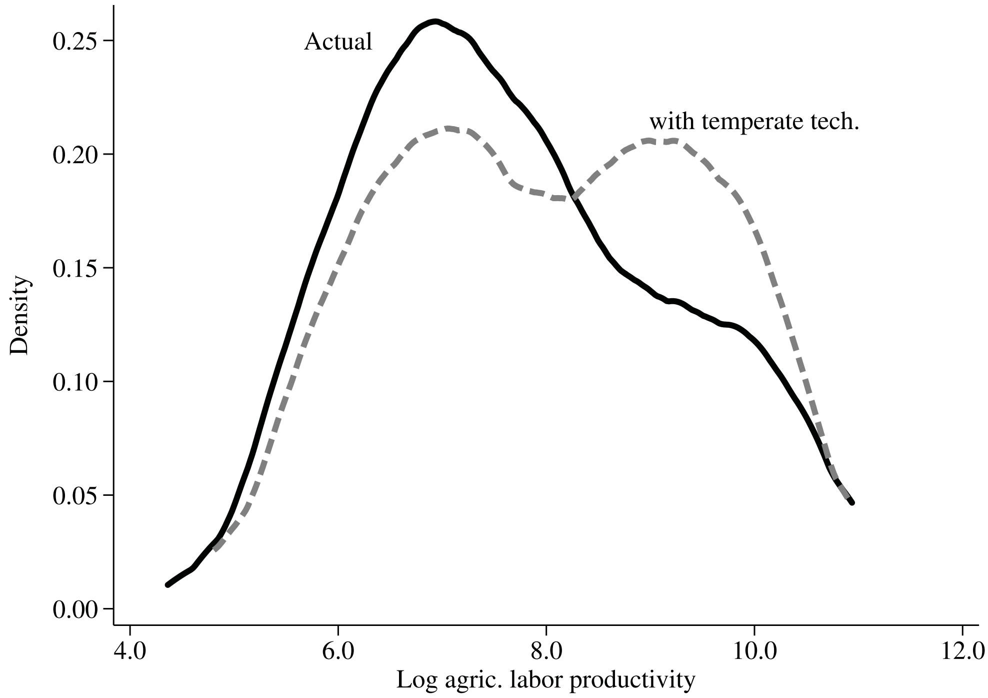

Notes: The figure shows the distribution of agricultural labor productivity for a sample of 78 countries from Caselli (2005) and the distribution of labor productivity for those same countries when they are given a temperate-zone agricultural technology with $\beta = 0.15$. The counter-factual labor productivity is calculated using our calibrated model, please see text for details. The kernel densities are calculated using an Epanechikov kernel with a bandwith of 0.56.

Figure 5: Agricultural Labor Share — Actual and with Temperate Technology

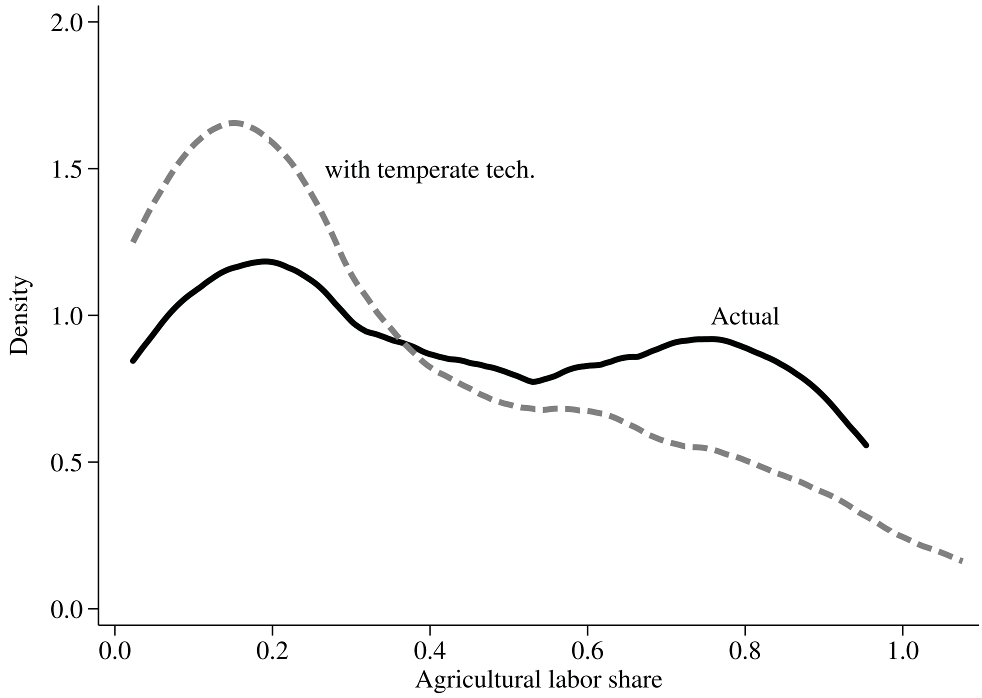

Notes: The figure shows the distribution of the agricultural labor share for a sample of 78 countries from Caselli (2005) and the distribution of labor productivity for those same countries when they are given a temperate-zone agricultural technology with $\beta = 0.15$. The counter-factual labor productivity is calculated using our calibrated model, please see text for details. The kernel densities are calculated using an Epanechikov kernel with a bandwith of 0.11.

TABLES

Table 1: Production Function Estimates

Weight matrix ♯	[1] 2FE	[2] MG	[3] CMG standard	[4] CMG neighbor	[5] CMG distance	[6] CMG agro-climate
Labor	-0.191 [10.60]**	-0.357 [2.23]*			-0.311 [2.62]**	
Tractors pw $\hat{\beta}_K$	0.058 [13.06]**	0.075 [3.31]**	0.109 [5.13]**	0.096 [4.17]**	0.078 [3.60]**	0.086 [3.82]**
Livestock pw $\hat{\beta}_{Live}$	0.358 [25.34]**	0.246 [8.07]**	0.321 [9.47]**	0.321 [8.22]**	0.278 [7.24]**	0.339 [9.97]**
Fertilizer pw $\hat{\beta}_F$	0.073 [19.87]**	0.030 [4.86]**	0.036 [5.63]**	0.035 [5.19]**	0.029 [5.11]**	0.035 [5.63]**
Land pw $\hat{\beta}_N$	0.294 [29.35]**	0.210 [2.79]**	0.201 [3.57]**	0.237 [4.14]**	0.081 [1.14]	0.190 [3.63]**
Returns to Scale ♭	DRS	DRS	CRS	CRS	DRS	CRS
Implied Avg $\hat{\beta}_L$	0.027 [2.34]*	0.082 [0.45]	0.333 [4.81]**	0.311 [4.24]**	0.223 [1.53]	0.353 [5.26]**
Avg implied $\hat{\beta}_L$		0.122 [1.06]	0.251 [4.24]**	0.283 [4.28]**	0.044 [0.53]	0.244 [4.08]**
# rejecting CRS		48			35	
$\hat{\varepsilon}$ Stationarity †	I(1)	I(0)	I(0)	I(0)	I(0)	I(0)
$\hat{\varepsilon}$ CD Test (p) ‡	9.64 (.00)	9.16 (.00)	-0.23 (0.82)	2.02 (0.04)	-0.49 (0.62)	-1.01 (0.31)
RMSE	0.148	0.066	0.059	0.060	0.053	0.060

Notes: Results for $n = 5,162$ observations from $N = 128$ countries. Estimators: 2FE – 2-way fixed effects, MG – Pesaran and Smith (1995) Mean Group, CMG – Pesaran (2006) Common Correlated Effects Mean Group. We present robust sample means for model parameters in [2]-[6]; [1] is a pooled model. * 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 based on the variance estimator following Pesaran and Smith (1995) in [2]-[6]. RMSE reports the root mean squared error.

Dependent variable – log output per worker, [1] dto. in 2FE transformation as described in Coakley et al. (2006).

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 $\hat{\beta}_L$' in a lower panel of the Table.

♭ 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.

♯ We apply different $N \times N$ weight matrices to construct the cross-section averages as described in the main text.

† Pesaran (2007) CIPS test results for residual nonstationarity: I(0) — stationary, I(1) — nonstationary. Full results available on request. ‡ Pesaran (2004) test for cross-section dependence (CD) in the residuals, H_0: no CD.

Table 2: Labor Coefficients and Climate Zones

	B Arid	C/D Temperate	A Equatorial	H Highland
Panel A: Any share of land in climatic zone				
Mean $\hat{\beta}_L$	0.185	0.137	0.397	0.240
	[0.086]**	[0.071]*	[0.089]***	[0.082]***
Sample	57	67	70	47
Panel B: Share above zone mean				
Mean $\hat{\beta}_L$	0.178	0.184	0.432	0.484
	[0.114]	[0.083]**	[0.115]***	[0.100]***
Sample	38	53	53	26
Panel C: Share above 40%				
Mean $\hat{\beta}_L$	0.103	0.161	0.419	0.357
	[0.100]	[0.086]*	[0.117]***	[0.220]
Sample	23	51	52	12
Panel D: Share above 50%				
Mean $\hat{\beta}_L$	-0.001	0.150	0.409	0.558
	[0.083]	[0.087]*	[0.124]***	[0.150]***
Sample	20	49	47	7
Panel E: Share above 60%				
Mean $\hat{\beta}_L$	0.035	0.082	0.413	0.734
	[0.115]	[0.098]	[0.131]***	[0.168]**
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. 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 — this indicates the skewness of distribution in the equatorial and temperate/cold zones. Sample indicates the number of observations from which the robust mean is constructed.

Table 3: Clusters – Descriptives

	Panel A: Four Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.059 [0.215]	0.443 [0.396]	0.403 [0.411]	0.094 [0.214]	43
Temperate/Cold	0.004 [0.015]	0.037 [0.101]	0.920 [0.141]	0.038 [0.096]	27
Equatorial	0.799 [0.231]	0.099 [0.181]	0.074 [0.139]	0.028 [0.075]	42
Equatorial & Highland	0.668 [0.307]	0.023 [0.060]	0.050 [0.193]	0.260 [0.256]	16

	Panel B: Five Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.072 [0.262]	0.285 [0.371]	0.589 [0.409]	0.053 [0.130]	28
Temperate/Cold	0.003 [0.013]	0.023 [0.065]	0.933 [0.121]	0.041 [0.099]	25
Arid	0.091 [0.140]	0.723 [0.228]	0.131 [0.214]	0.055 [0.149]	18
Equatorial	0.837 [0.202]	0.057 [0.107]	0.078 [0.142]	0.029 [0.077]	40
Equatorial & Highland	0.570 [0.357]	0.044 [0.094]	0.047 [0.188]	0.340 [0.301]	17

	Panel C: Six Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.000 [0.000]	0.146 [0.153]	0.800 [0.190]	0.054 [0.127]	15
Temperate/Cold	0.003 [0.013]	0.023 [0.065]	0.933 [0.121]	0.041 [0.099]	25
Arid	0.102 [0.145]	0.730 [0.218]	0.147 [0.223]	0.022 [0.057]	16
Arid & Equatorial	0.358 [0.462]	0.329 [0.469]	0.278 [0.426]	0.036 [0.082]	19
Equatorial	0.833 [0.202]	0.053 [0.104]	0.072 [0.138]	0.042 [0.102]	43
Equatorial & Highland	0.259 [0.283]	0.164 [0.189]	0.002 [0.008]	0.575 [0.196]	10

Notes: The table presents the mean share of land in each of the four climatic zones (denoted by their Köppen-Geiger letters) for groups of countries determined endogenously using the cluster analysis technique described in the main test. The standard deviations of these means are in brackets.

Table 4: Labor Coefficients by Clusters

Panel A: Four Clusters

Cluster	Arid & Temp/Cold	Temperate/ Cold	Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	0.143	0.166	0.320	0.555
	[0.122]	[0.078]**	[0.104]***	[0.295]*
N	43	27	42	16

Panel B: Five Clusters

Cluster	Arid & Temp/Cold	Temperate/ Cold	Arid	Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	0.011	0.166	0.183	0.382	0.537
	[0.177]	[0.084]*	[0.116]	[0.114]***	[0.236]**
N	28	25	18	40	17

Panel C: Six Clusters

Group/Cluster	Arid & Temp/Cold	Temperate/ Cold	Arid	Equatorial	Arid & Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	-0.234	0.166	0.198	0.339	0.530	0.646
	[0.220]	[0.084]*	[0.132]	[0.108]***	[0.258]*	[0.146]***
N	15	25	16	43	19	10

Notes: We present robust mean estimates (standard errors in brackets) for the labor coefficients (using estimates from Model [6] in Table 1) based on various groupings. In Panels A-C we use cluster analysis specifying 4, 5 and 6 clusters, respectively, to create our groups (see Table 3 for group descriptives).

Table 5: Correlation matrix – variable averages and CMG estimates

Variable averages	$\overline{\text{ly}}_i$	$\overline{\text{ltr}}_i$	$\overline{\text{llive}}_i$	$\overline{\text{lf}}_i$	$\overline{\text{ln}}_i$	$\hat{\beta}_i^{\text{Tr}}$	$\hat{\beta}_i^{\text{Live}}$	$\hat{\beta}_i^{\text{F}}$	$\hat{\beta}_i^{\text{N}}$
Output pw $\overline{\text{ly}}_i$	1								
Tractors pw $\overline{\text{ltr}}_i$	**0.911**	1							
Livestock pw $\overline{\text{llive}}_i$	**0.816**	**0.738**	1						
Fertilizer pw $\overline{\text{lf}}_i$	**0.902**	**0.917**	**0.695**	1					
Land pw $\overline{\text{ln}}_i$	**0.780**	**0.718**	**0.677**	**0.673**	1				
Standard CMG	$\overline{\text{ly}}_i$	$\overline{\text{ltr}}_i$	$\overline{\text{llive}}_i$	$\overline{\text{lf}}_i$	$\overline{\text{ln}}_i$	$\hat{\beta}_i^{\text{Tr}}$	$\hat{\beta}_i^{\text{Live}}$	$\hat{\beta}_i^{\text{F}}$	$\hat{\beta}_i^{\text{N}}$
$\hat{\beta}_i^{\text{Tr}}$	0.089	0.124	0.052	0.072	0.051	1			
$\hat{\beta}_i^{\text{Live}}$	0.003	-0.015	0.153	-0.051	-0.119	**-0.330**	1		
$\hat{\beta}_i^{\text{F}}$	0.115	0.123	0.075	**0.223**	0.116	-0.067	-0.119	1	
$\hat{\beta}_i^{\text{N}}$	0.105	0.139	0.076	**0.203**	0.108	**-0.203**	0.007	0.124	1
Agro-climatic CMG	$\overline{\text{ly}}_i$	$\overline{\text{ltr}}_i$	$\overline{\text{llive}}_i$	$\overline{\text{lf}}_i$	$\overline{\text{ln}}_i$	$\hat{\beta}_i^{\text{Tr}}$	$\hat{\beta}_i^{\text{Live}}$	$\hat{\beta}_i^{\text{F}}$	$\hat{\beta}_i^{\text{N}}$
$\hat{\beta}_i^{\text{Tr}}$	0.128	0.138	0.106	0.150	0.008	1			
$\hat{\beta}_i^{\text{Live}}$	0.040	0.024	0.126	-0.047	-0.007	**-0.238**	1		
$\hat{\beta}_i^{\text{F}}$	0.148	0.168	0.100	**0.282**	0.138	-0.002	**-0.218**	1	
$\hat{\beta}_i^{\text{N}}$	0.098	0.125	0.037	0.128	0.145	-0.062	-0.053	0.094	1

Notes: We correlate the country-specific variable series (means) with the standard and agro-climatic CMG technology estimates. Significant coefficients (5% level) are in bold (except for the diagonals). We employ the CRS-based estimates for the standard and agro-climate CMG respectively (Table 1, columns [3] and [6]). Coefficient estimates are for 'Tr' tractors, 'Live' livestock, 'F' fertilizer and 'N' land; 'pw' refers to variables in 'per worker' terms.

Table 6: Effects of Exogenous Productivity and Population Changes

		Equilibrium outcomes from:					
		20% Increase in Ag. TFP (A) with $\beta =$			5% Increase in Population (L) with $\beta =$		
Outcome	Baseline	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: $\alpha = 0.01$, $\bar{c}_a = 0.802$, $\bar{c}_n = 3.182$, $w = 1.0$.

Table 7: Income Dispersion, Actual and Counter-factuals

	Output per capita:		Ag. labor productivity:	
	$Var(\ln y)$	90/10 ratio	$Var(\ln Y_A/L_A)$	90/10 Ratio
Actual	1.185	21.7	2.206	46.0
Counter-factual: Set $\beta_L = 0.15$ *for all countries*				
Temperate technology	0.996	14.8	2.217	44.1

Notes: Actual data is from Caselli (2005) and covers 78 countries. Variances and 90/10 ratios are authors' calculations. Values for using temperate technology are calculated by assuming that each country has an agricultural technology with an elasticity of agricultural output w.r.t. labor of 0.15. See text for details.

Table 8: Country Data, Actual and with Temperate Technology

Country	Output p.c. Actual	Output p.c. Temp.	Ag. labor prod. Actual	Ag. labor prod. Temp.	Ag. labor share Actual	Ag. labor share Temp.
Equatorial cluster:						
Angola	4,558	8,405	354	446	0.82	0.65
Bangladesh	4,085	7,002	513	717	0.74	0.54
Bolivia	8,151	11,478	1,706	3,542	0.47	0.24
Brazil	16,808	20,299	3,853	9,742	0.28	0.13
Burundi	1,389	936	443	433	0.92	0.97
Cameroon	7,272	12,272	78	122	0.67	0.44
Costa Rica	13,340	16,053	4,072	10,523	0.27	0.12
Cote d'Ivore	5,624	8,503	1,377	2,395	0.59	0.35
Dominican Republic	10,368	14,072	2,124	4,776	0.42	0.20
El Salvador	10,960	14,444	1,746	4,196	0.38	0.17
Ghana	3,187	4,724	899	1,609	0.57	0.33
Guinea	5,053	8,947	508	683	0.77	0.58
Haiti	2,257	3,616	624	920	0.70	0.48
India	3,546	5,758	592	908	0.68	0.45
Kenya	2,918	5,252	532	642	0.85	0.71
Madagascar	2,123	3,498	682	887	0.79	0.61
Mozambique	1,469	2,701	231	272	0.87	0.74
Nigeria	2,781	4,832	451	596	0.78	0.60
Rwanda	2,129	3,898	475	526	0.92	0.83
Senegal	3,445	6,201	442	562	0.81	0.64
Sri Lanka	5,438	7,676	957	1,997	0.47	0.24
Tanzania	1,303	2,196	495	581	0.87	0.75
Thailand	6,307	10,375	858	1,309	0.68	0.45
Uganda	1,469	2,401	617	659	0.94	0.89
Venezuela	20,074	21,979	4,460	13,567	0.13	0.06
Equatorial/Highland cluster:						
Colombia	13,583	17,623	2,746	15,361	0.28	0.07
Ecuador	12,972	18,119	2,721	13,366	0.35	0.09
Ethiopia	1,196	2,589	351	590	0.79	0.48
Guatemala	12,594	22,287	1,426	4,433	0.56	0.19
Honduras	7,599	14,088	1,412	3,810	0.61	0.24
Indonesia	5,976	10,351	839	2,692	0.54	0.18
Malaysia	13,743	19,391	2,460	11,891	0.36	0.09
Mexico	23,256	32,995	2,341	11,283	0.36	0.09
Nicaragua	9,855	14,535	1,806	7,604	0.41	0.11
Papua New Guinea	6,387	13,436	1,246	2,504	0.72	0.37
Paraguay	13,090	20,982	3,820	13,648	0.50	0.15

Notes: Actual data is from Caselli (2005). The "Temp." columns refer to counter-factual values calculated using our calibrated model and assuming the country has temperate agricultural technology with $\beta = 0.15$. See text for details.

REFERENCES

[1] **Alvarez-Cuadrado, Francisco and Poschke, Markus (2011):** "Structural change out of agriculture: Labor push versus labor pull." *American Economic Journal: Macroeconomics* 3: 127–158.

[2] **Andrews, Donald W. K. (2005):** "Cross-section Regression with Common Shocks." *Econometrica* 73(5): 1551–1585.

[3] **Bai, Jushan (2009):** "Panel Data Models with Interactive Fixed Effects," *Econometrica* 77(4): 1229–79.

[4] **Bernard, Andrew, and Chad Jones (1996):** "Comparing Apples to Oranges: Productivity Convergence and Measurement Across Industries and Countries." *American Economic Review*, 86(5), pp. 1216–1238.

[5] **Betts, Caroline, Rahul Giri, and Rubina Verma (2013):** "Trade, Reform, and Structural Transformation in South Korea." Working paper.

[6] **Block, Steven A. (2014):** "The post-independence decline and rise of crop productivity in sub-Saharan Africa: measurement and explanations." *Oxford Economic Papers* 66(2): 373–396.

[7] **Canning, David, and Pedroni, Peter (2008):** "Infrastructure, long-run economic growth and causality tests for cointegrated panels." *The Manchester School*, 76(5), pp. 504–527.

[8] **Chudik, Alexander, Pesaran, M. Hashem and Tosetti, Elisa (2011):** "Weak and Strong Cross Section Dependence and Estimation of Large Panels." *Econometrics Journal*, 14(1), pp. C45–C90.

[9] **Chudik, Alexander, and Pesaran, M. Hashem (2013):** "Large Panel Data Models with Cross-Sectional Dependence: A Survey." Unpublished mimeo, University of Cambridge.

[10] **Coakley, Jerry, Fuertes, Ana-Maria and Smith, Ron P. (2006):** "Unobserved heterogeneity in panel time series models. *Computational Statistics & Data Analysis* 50(9): 2361–2380.

[11] **Craig, Barbara J, Pardey, Philip G and Roseboom, Johannes (1997):** "International Productivity Patterns: Accounting for Input Quality, Infrastructure, and Research." *American Journal of Agricultural Economics*, 79(4): 1064–1076.

[12] **Duarte, Margarida and Restuccia, Diego (2010):** "The role of the structural transformation in aggregate productivity." *Quarterly Journal of Economics* 125(1): 129–173.

[13] **Durlauf, Steven N., Kourtellos, Andros and Minkin, Artur (2001):** "The local Solow growth model." *European Economic Review*, 45(4-6): 928–940.

[14] **Eberhardt, Markus, Helmers, Christian and Strauss, Hubert (2013):** "Do Spillovers Matter When Estimating Private Returns to R&D?" *Review of Economics and Statistics*, 95(2): 436–448.

[15] **Eberhardt, Markus and Francis Teal (2013a):** "Structural Change and Cross-Country Growth Empirics." *World Bank Economic Review*, 27(2): 229–271.

[16] **Eberhardt, Markus and Francis Teal (2013b):** "No Mangoes in the Tundra: Spatial Heterogeneity in Agricultural Productivity Analysis." *Oxford Bulletin of Economics and Statistics*, 75(6): 914–939.

[17] **Eberhardt, Markus and Francis Teal (2014):** "The Magnitude of the Task Ahead: Productivity Analysis with Heterogeneous Technology." Unpublished mimeo, University of Nottingham.

[18] **Echevarria, Cristina (1997):** "Changes in sectoral composition associated with economic growth." *International Economic Review*, 38(2): 431–52.

[19] **FAO (2007):** 'FAOSTAT.' Online database, Rome: UN Food and Agriculture Organisation.

[20] **Frankel, Jeffrey A. and David Romer (1999):** "Does Trade Cause Growth?" *American Economic Review* 89(3): 379–399.

[21] **Fuglie, Keith O. (2010):** "Total Factor Productivity in the Global Agricultural Economy: Evidence from FAO Data." in *The Shifting Patterns of Agricultural Production and Productivity Worldwide*, ed. J. Alston, B. Babcock, P. Pardey. Midwest Agribusiness Trade Research and Information Center, Iowa State University, Ames, Iowa.

[22] **Fuglie, Keith O. and Rada, Nicholas E. (2013):** "Resources, Policies, and Agricultural Productivity in Sub-Saharan Africa." United States Department of Agriculture Economic Research Report Number 145.

[23] **Gallup, John L., Mellinger, Andrew D. and Sachs, Jeffrey D. (1999):** "Geography Datasets: Agricultural Measures." Available online at CID, Harvard University.

[24] **Gollin, Douglas (2002):** "Getting Income Shares Right." *Journal of Political Economy* 110(2): 458–474.

[25] **Gollin, Douglas (2010):** "Agricultural productivity and economic growth." In: Pingali, Prabhu and Evenson, Robert (eds.) *Handbook of Agricultural Economics*, Volume 4: 3825–3866. Elsevier: Amsterdam.

[26] **Gollin, Douglas, Parente, Stephen and Rogerson, Richard (2007):** "The food problem and the evolution of international income levels." *Journal of Monetary Economics* 54: 1230–1255.

[27] **Gutierrez, Luciano and 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.

[28] **Hamilton, Lawrence C. (1992):** "How Robust is Robust Regression?" *Stata Technical Bulletin* 1(2).

[29] **Hansen, Gary D. and Prescott, Edward C. (2002):** "From Malthus to Solow." *American Economic Review*, 92(4): 1205–1217.

[30] **Hayami, Yujiro and Ruttan, Vernon W. (1970):** "Agricultural productivity differences among countries." *American Economic Review*, 60(5): 895–911.

[31] **Hayami, Yujiro and Ruttan, Vernon W. (1985):** *Agricultural Development: An International Perspective*. Johns Hopkins University Press, Baltimore.

[32] **Hayashi, Fumio and Prescott, Edward C. (2008):** "The depressing effect of agricultural institutions on the prewar Japanese economy." *Journal of Political Economy*, 116(4): 573–632.

[33] **Herrendorf, Berthold, Rogerson, Richard and Valentinyi, Akos (2013a):** "Growth and structural transformation." In: Aghion, Philippe and Durlauf, Steven (eds.) *Handbook of Economic Growth*, draft chapter.

[34] **Herrendorf, Berthold, Rogerson, Richard and Valentinyi, Akos (2013b):** "Two perspectives on preferences and structural transformation." *American Economic Review*, forthcoming.

[35] **Heston, Alan, Summers, Robert and Aten, Bettina (2009):** "Penn World Table Version 6.3." Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

[36] **Jaffe, Adam (1986):** "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits and Market Value." NBER Working Paper Series No. 1815.

[37] **Johnston, Bruce F. and Kilby, Peter (1975):** *Agriculture and Structural Transformation: Economic Strategies in Late-Developing Countries*. Oxford University Press, New York, NY.

[38] **Johnston, Bruce F. and Mellor, John W. (1961):** "The role of agriculture in economic development." *American Economic Review*, 51(4): 566–93.

[39] **Jorgenson, Dale W. and Gollop, Frank M. (1992):** "Productivity Growth in U.S. Agriculture: A Postwar Perspective." *American Journal of Agricultural Economics* 74(3):745-750.

[40] **Kogel, Tomas and Prskawetz, Alexia (2001):** "Agricultural productivity growth and escape from the Malthusian trap." *Journal of Economic Growth* 6(4): 337–57.

[41] **Kongasmut, Piyabha, Rebelo, Sergio and Xie, Danyang (2001):** "Beyond Balanced Growth." *Review of Economic Studies* 68(4): 869–882.

[42] **Kottek, Markus, Grieser, Jürgen, Beck, Christoph, Rudolf, Bruno and Rubel, Franz (2006):** "World Map of the Köppen-Geiger climate classification updated.' *Meteorologische Zeitschrift* 15(3): 259–263.

[43] **Lagakos, David and Michael E. Waugh (2013):** "Selection, agriculture, and cross-country productivity differences." *American Economic Review* 103(2): 948–980.

[44] **Marschak, Jacob, and Andrews, William H. (1944):** "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12(3/4): 143–205.

[45] **Martin, Will and Mitra, Devashish (2002):** "Productivity Growth and Convergence in Agriculture versus Manufacturing." *Economic Development and Cultural Change* 49(2): 403–422.

[46] **Matsuyama, Kiminori (1992):** "Agricultural Productivity, Comparative Advantage, and Economic Growth." *Journal of Economic Theory* 58: 317-334.

[47] **Matthews, Elaine (1983):** 'Global vegetation and land use: new high resolution databases for climate studies.' *Journal of Climate and Applied Meteorology* 22: 474–487.

[48] **Mayer, Thierry and Zignago, Soledad (2006):** "Background Note on Datasets for Geodesic Distances." Centre D'Etudes Prospectives et D'Information Internationales (CEPII). Available at: www.cepii.fr/anglaisgraph/bdd/distances.

[49] **Mellor, John W. (1995):** "Introduction." In: John W. Mellor (ed.) *Agriculture on the Road to Industrialization*. Johns Hopkins University Press, Baltimore.

[50] **Mourmouras, Alex and Rangazas, Peter (2009):** "Reconciling Kuznets and Habbakuk in a unified growth theory." *Journal of Economic Growth* 14(2): 149–181.

[51] **Mundlak, Yair (1988):** "Endogenous Technology and the Measurement of Productivity." In: Susan M. Capalbo and John M. Antle (eds.) *Agricultural Productivity: Measurement and Explanation*. Washington, DC: Resources for the Future.

[52] **Mundlak, Yair (2000):** *Agriculture and Economic Growth: Theory and Measurement*. Harvard University Press, Cambridge, MA.

[53] **Mundlak, Yair, Larson, Donald F. and Butzer, Rita (1999):** "Rethinking within and between regressions: The case of agricultural production functions." *Annales D'Economie et de Statistique* 55/56: 475–501.

[54] **Mundlak, Yair, Butzer, Rita and Larson, Donald F. (2012):** "Heterogeneous technology and panel data: The case of the agricultural production function." *Journal of Development Economics* 99(1): 139–149.

[55] **Pardey, Philip G., Alston, Julian M., and Kang, Connie Chan (2012):** "Agricultural Production, Productivity and R&D over the Past Half Century: An Emerging New World Order." Paper presented at the International Association of Agricultural Economists (IAAE) Triennial Conference, Foz do Iguaçu, Brazil, August 2012.

[56] **Pedroni, Peter (2000):** "Fully modified OLS for heterogeneous cointegrated panels." In: Badi H. Baltagi (ed.) *Nonstationary panels, cointegration in panels and dynamic panels*, Elsevier: Amsterdam.

[57] **Pedroni, Peter (2001):** "Purchasing Power Parity Tests In Cointegrated Panels." *The Review of Economics and Statistics* 83(4): 727–731.

[58] **Pedroni, Peter (2007):** "Social Capital, Barriers to Production and Capital Shares: Implications for the Importance of Parameter Heterogeneity from a Nonstationary Panel Approach." *Journal of Applied Econometrics* 22(2): 429–451.

[59] **Pesaran, M. Hashem (2004):** "General Diagnostic Tests for Cross Section Dependence in Panels." Cambridge Working Papers in Economics No. 0435, University of Cambridge.

[60] **Pesaran, M. Hashem (2006):** "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure.' *Econometrica* 74(4): 967–1012.

[61] **Pesaran, M. Hashem and Smith, Ron P. (1995):** "Estimating long-run relationships from dynamic heterogeneous panels." *Journal of Econometrics* 68(1): 79–113.

[62] **Phillips, Peter and Sul, Donggyul (2007):** "Bias in Dynamic Panel Estimation with Fixed Effects, Incidental Trends and cross-sectional Dependence." *Journal of Econometrics* 137(1): 162–188.

[63] **Phillips, Peter and Hansen, Bruce E. (1990):** "Statistical inference in instrumental variables regression with I (1) processes." *The Review of Economic Studies* 57(1): 99–125.

[64] **Redding, Stephen and Anthony J. Venables (2004):** "Economic geography and international inequality." *Journal of International Economics* 62(1): 53–82.

[65] **Restuccia, Diego, Yang, Dennis and Zhu, Xiaodong (2008):** "Agriculture and aggregate productivity." *Journal of Monetary Economics*, 55(2): 234–250.

[66] **Ruthenberg, H. (1976):** *Farming Systems in the Tropics*. Clarendon Press, Oxford, UK, 2nd edition.

[67] **Ruttan, Vernon W. (2002):** "Productivity growth in world agriculture: sources and constraints." *Journal of Economic Perspectives* 16: 161–184.

[68] **Schultz, Theodore W. (1953):** *The Economic Organization of Agriculture*. McGraw-Hill, New York, NY.

[69] **Sposi, Michael (2012).** "Evolving Comparative Advantage, Structural Change, and the Composition of Trade." Working paper.

[70] **Strulik, Holger and Weisdorf, Jacob L. (2008):** "Population, food, and knowledge: A simple unified growth theory." *Journal of Economic Growth* 13(3): 195–216.

[71] **Teignier, Marc (2012).** "The Role of Trade in Structural Transformation." Working paper, Universidad de Alicante.

[72] **Timmer, Marcel P. and de Vries, Gaaitzen J.** A cross-country database for sectoral employment and productivity in Asia and Latin America, 1950–2005. GDCC Research Memorandum GD-98, August 2007.

[73] **Uy, Timothy, Yi, Kei-mu, and Jing Zhang (2013).** “Structural Change in an Open Economy.” *Journal of Monetary Economics* 60:667-682.

[74] **Vollrath, Dietrich (2009):** “The Dual Economy in Long-run Development.” *Journal of Economic Growth* 14(4): 287–312.

[75] **Vollrath, Dietrich (2011):** “The Agricultural Basis of Comparative Development,” *Journal of Economic Growth* 16(4): 343–370.

[76] **Wiebe, Keith, Soule, Meredith J., Narrod, Clare and Breneman, Vincent E. (2003):** “Resource Quality and Agricultural Productivity: A Multi-Country Comparison.” in K. Wiebe (ed.) *Land Quality, Agricultural Produtivity, and Food Security*. Northampton, MA: Edward Elger Publishing.

APPENDIX

A-1. DATA SOURCES AND 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 T 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 sector.[28] 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 asses (donkeys), buffalos, camels, cattle, chickens, ducks, horses, mules, pigs, sheep & goats and turkeys. Following convention we use a formula detailed in Hayami and Ruttan (1970) to convert the numbers for individual animal species into cattle-equivalent livestock. 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 hectare).

Descriptive statistics are presented in Table A-I. The countries in our sample are listed in Table A-II, with 'cluster affiliation' presented in Table A-III. The clusters (4, 5 or 6) were created from the country-pair Jaffe measure of agro-climatic similarity[29] described in Section A-4 below. Each country has N Jaffe measures and the grouping of countries is then computed from 'partition-clustering,' which breaks the observations into a distinct number of non-overlapping groups. More specifically, we adopt 'kmeans' clustering in Stata 11, where following the choice of number of groups to be created each observation is assigned to the group whose mean is closest. Based on this categorization the new group means are determined and the entire process is iterated until the group affiliation is stable (see Stata 11 Multivariate Statistics Manual p.85ff).

As a robustness check we drop all countries where livestock (in value terms) on average exceeds 60% of total agricultural output, adopting recently-developed FAO data for net output and livestock in 2004-2006 I\$ value terms. Figure TA-I in a Technical Appendix indicates that this cutoff point represents a natural choice given the distribution of the data. The dependent variable for these models is net output less livestock value and we drop the headcount livestock variable (described above) from these regressions.

Analysing agricultural production for a large number of countries inevitably raises concerns over data reliability. In contrast to macro data provided in for instance the Penn World Table (GGDC Groningen) FAOSTAT does not offer a data quality grade for each country, but instead labels each observation. Most output data used in our analysis carries the note '[a]ggregates may include official, semi-official or estimates.' For inputs we obtain more details, which suggest that tractor data is least reliable, with around 45% of observations estimated.[30] Thus data is far from perfect for cross-country comparison, although estimating production functions country by country and accounting for

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

[29] The multivariate correlation measure of share of arable and permanent crop land in one of 12 climate zones between each country pair.

[30] We report Penn World Table country quality grade and share of non-estimated tractor data in Table A-II.

unobserved common factors should go some way to ward against systematic over-/underreporting of variable magnitudes. We find that there is no statistically significant relationship between countries' share of tractor data estimated and the coefficients on tractor, livestock, fertilizer or land in our preferred agro-climatic CMG model, which could be viewed as evidence to that end.

Additional time-invariant data on geographical distance between countries and contiguity (neighborhood) is taken from Mayer and Zignago (2006), and data on the share of agricultural land by climatic zone from Matthews (1983), available in Gallup et al. (1999) — see Table A-IV for details.

A-2. DESCRIPTIVE STATISTICS AND SAMPLE

Table A-I: *Descriptive statistics*

Variables in untransformed level terms					
Variable	*mean*	*median*	*std. dev.*	*min.*	*max.*
logs					
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
annual growth rate					
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					
Variable	**mean**	**median**	**std. dev.**	**min.**	**max.**
logs					
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
annual growth rate					
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 net output (in I$1,000), labor (headcount), tractors (number), livestock (cattle-equivalent numbers), fertilizer (in metric tonnes) and land (in hectare) for the full sample ($n = 5,162$; $N = 128$).

Table A-II: Sample of countries and number of observations

Country	Code	Obs	PWT-Q	FAO-Q	Country	Code	Obs	FAO-Q	FAO-Q
Afghanistan	AFG	40		5%	Cambodia	KHM	33	D	25%
Angola	AGO	40	D	45%	South Korea	KOR	42	B	95%
Albania	ALB	42	C	69%	Kuwait †	KWT	24	C	43%
United Arab Emirates	ARE	31	D	23%	Lao PDR	LAO	38	D	52%
Argentina	ARG	42	B	50%	Lebanon	LBN	42	C	13%
Australia †	AUS	42	A	100%	Liberia	LBR	30	D	
Austria †	AUT	42	A	76%	Libya	LBY	42		15%
Burundi	BDI	37	C	28%	Sri Lankla	LKA	42	C	50%
Benin	BEN	42	C	13%	Lesotho †	LSO	42	D	20%
Burkina Faso	BFA	42	C	17%	Morocco	MAR	42	C	56%
Bangladesh	BGD	42	C	5%	Madagascar	MDG	42	C	17%
Bulgaria	BGR	42	C	79%	Mexico	MEX	42	C	21%
Belgium-Luxembourg †	BLX	39	A	85%	Mali	MLI	42	C	12%
Belize	BLZ	42	C	36%	Myanmar	MMR	42	D	88%
Bolivia	BOL	42	C	28%	Mongolia †	MNG	34	D	80%
Brazil	BRA	42	C	12%	Mozambique	MOZ	42	D	20%
Botswana †	BWA	42	C	33%	Mauritania †	MRT	33	C	14%
Central African Republic	CAF	42	D	22%	Malawi	MWI	42	C	63%
Canada	CAN	42	A	62%	Malaysia	MYS	42	C	51%
Switzerland †	CHE	42	A	17%	Niger	NER	34	D	63%
Chile	CHL	42	B	33%	Nigeria	NGA	42	C	5%
China	CHN	42	C	90%	Nicaragua	NIC	42	C	16%
Côte d'Ivoire	CIV	42	C	27%	Netherlands †	NLD	42	A	29%
Cameroon	CMR	42	C	32%	Norway †	NOR	42	A	71%
Congo, Republic	COG	41	C	93%	Nepal	NPL	42	C	23%
Colombia	COL	42	C	68%	New Zealand †	NZL	42	B	62%
Costa Rica	CRI	42	C	31%	Oman	OMN	30	C	38%
Cuba	CUB	42	D	57%	Pakistan	PAK	42	C	43%
Cyprus	CYP	42	D	50%	Panama	PAN	42	C	13%
Germany †	DEU	42	B	93%	Philippines	PHL	42	C	53%
Denmark †	DNK	42	A	90%	Papua New Guinea	PNG	42	D	62%
Dominican Republic	DOM	42	C	5%	Poland	POL	42	B	100%
Algeria	DZA	42	D	48%	Korea, DPR	PRK	42		50%
Ecuador	ECU	42	C	33%	Portugal	PRT	42	B	33%
Egypt	EGY	42	C	45%	Paraguay	PRY	42	C	2%
Spain	ESP	42	B	100%	Qatar †	QAT	27	C	45%
Ethiopia	ETH	42	C	25%	Romania	ROM	42	C	100%
Finland †	FIN	30	A	62%	Rwanda	RWA	34	C	24%
France	FRA	42	A	79%	Saudi Arabia	SAU	42	D	21%
Gabon	GAB	31	C		Sudan	SDN	42	D	21%
United Kingdom †	GBR	42	A	76%	Senegal	SEN	42	C	10%
Ghana	GHA	42	C	5%	Sierra Leone	SLE	42	C	35%
Guinea	GIN	41	C	10%	El Salvador	SLV	42	C	9%
Gambia	GMB	39	C	22%	Somalia †	SOM	36	D	33%
Guinea-Bissau	GNB	26	D	14%	Suriname	SUR	42	D	14%
Equatorial Guinea	GNQ	19	D		Sweden †	SWE	42	A	52%
Greece	GRC	42	B	100%	Swaziland	SWZ	42	C	74%
Guatemala	GTM	42	C	30%	Syria	SYR	42	C	100%
Guyana	GUY	42	D	6%	Chad	TCD	41	D	100%
Honduras	HND	42	C	25%	Togo	TGO	37	D	19%
Haiti	HTI	42	D	8%	Thailand	THA	42	C	45%
Hungary	HUN	42	C	79%	Trinidad & Tobago	TTO	42	C	0%
Indonesia	IDN	42	C	28%	Tunisia	TUN	42	C	33%
India	IND	42	C	83%	Turkey	TUR	42	C	100%
Ireland †	IRL	42	A	50%	Tanzania	TZA	42	C	10%
Iran	IRN	42	C	33%	Uganda	UGA	39	D	59%
Iraq	IRQ	42	D	45%	Uruguay †	URY	42	B	17%
Iceland †	ISL	42	B	79%	United States	USA	42	A	40%
Israel	ISR	42	B	83%	Venezuela	VEN	42	C	71%
Italy	ITA	42	A	100%	Vietnam	VNM	42	C	65%
Jamaica	JAM	42	C	70%	Yemen, Republic	YEM	37	D	15%
Jordan	JOR	42	C	83%	South Africa	ZAF	42	C	60%
Japan	JPN	42	A	85%	Congo, DR	ZAR	41	D	18%
Kenya	KEN	42	C	60%	Zimbabwe	ZWE	42	C	32%

Notes: The full sample contains $n = 5,162$ observations, from 1961 to 2002. † indicates countries dropped in the reduced sample (see maintext). PWT-Q reports a data quality rating for aggregate economy data from the Penn World Table project (Heston et al., 2009), where A denotes the highest and D the lowest score (http://pwt.econ.upenn.edu/Documentation/append61.pdf, Table A, column 11). FAO-Q reports the share of observations for the tractor variable which are not estimated but taken from official publications or international organisations (FAO codes: I, W, Q), which is reported for most FAO observations.

Table A-III: Country Clusters

	Country	wbcode	Number of Clusters: Four	Five	Six
1	Afghanistan	AFG	A&T/C	Eq&Hi	Eq&Hi
2	Angola	AGO	Equat	Equat	Equat
3	Albania	ALB	A&T/C	A&T/C	A&T/C
4	United Arab Emirates	ARE	A&T/C	A&T/C	A&Eq
5	Argentina	ARG	T/C	T/C	T/C
6	Australia	AUS	T/C	Arid	Arid
7	Austria	AUT	T/C	T/C	T/C
8	Burundi	BDI	Equat	Equat	Equat
9	Benin	BEN	Equat	Arid	Arid
10	Burkina Faso	BFA	A&T/C	Arid	Arid
11	Bangladesh	BGD	Equat	Equat	Equat
12	Bulgaria	BGR	T/C	T/C	T/C
13	Belgium-Luxembourg	BLX	T/C	T/C	T/C
14	Belize	BLZ	Eq&Hi	Eq&Hi	Equat
15	Bolivia	BOL	Equat	Equat	Equat
16	Brazil	BRA	Equat	Equat	Equat
17	Botswana	BWA	A&T/C	Arid	Arid
18	Central African Republic	CAF	Equat	Equat	Equat
19	Canada	CAN	A&T/C	A&T/C	A&Eq
20	Switzerland	CHE	T/C	T/C	T/C
21	Chile	CHL	A&T/C	A&T/C	A&T/C
22	China	CHN	T/C	T/C	T/C
23	Côte d'Ivoire	CIV	Equat	Equat	Equat
24	Cameroon	CMR	Equat	Equat	Equat
25	Congo, Republic	COG	Equat	Equat	Equat
26	Colombia	COL	Eq&Hi	Eq&Hi	Equat
27	Costa Rica	CRI	Equat	Equat	Equat
28	Cuba	CUB	Equat	Equat	Equat
29	Cyprus	CYP	A&T/C	A&T/C	A&T/C
30	Germany	DEU	T/C	T/C	T/C
31	Denmark	DNK	T/C	T/C	T/C
32	Dominican Republic	DOM	Equat	Equat	Equat
33	Algeria	DZA	A&T/C	A&T/C	A&T/C
34	Ecuador	ECU	Eq&Hi	Eq&Hi	Eq&Hi
35	Egypt	EGY	A&T/C	A&T/C	A&Eq
36	Spain	ESP	A&T/C	A&T/C	A&T/C
37	Ethiopia	ETH	Eq&Hi	Eq&Hi	Eq&Hi
38	Finland	FIN	A&T/C	A&T/C	A&Eq
39	France	FRA	T/C	T/C	T/C
40	Gabon	GAB	Equat	Equat	Equat
41	United Kingdom	GBR	T/C	T/C	T/C
42	Ghana	GHA	Equat	Equat	Equat
43	Guinea	GIN	Equat	Equat	Equat

	Country	wbcode	Number of Clusters: Four	Five	Six
44	Gambia	GMB	Equat	Equat	Equat
45	Guinea-Bissau	GNB	Equat	Equat	Equat
46	Equatorial Guinea	GNQ	Eq&Hi	Eq&Hi	A&Eq
47	Greece	GRC	A&T/C	A&T/C	A&T/C
48	Guatemala	GTM	Eq&Hi	Eq&Hi	Equat
49	Guyana	GUY	Eq&Hi	Eq&Hi	A&Eq
50	Honduras	HND	Eq&Hi	Eq&Hi	Eq&Hi
51	Haiti	HTI	Equat	Equat	Equat
52	Hungary	HUN	T/C	T/C	T/C
53	Indonesia	IDN	Eq&Hi	Eq&Hi	Eq&Hi
54	India	IND	Equat	Equat	Equat
55	Ireland	IRL	T/C	T/C	T/C
56	Iran	IRN	A&T/C	Arid	Eq&Hi
57	Iraq	IRQ	A&T/C	Arid	Arid
58	Iceland	ISL	T/C	T/C	T/C
59	Israel	ISR	A&T/C	A&T/C	A&T/C
60	Italy	ITA	T/C	A&T/C	A&T/C
61	Jamaica	JAM	Equat	Equat	Equat
62	Jordan	JOR	A&T/C	A&T/C	A&T/C
63	Japan	JPN	T/C	T/C	T/C
64	Kenya	KEN	Equat	Equat	Equat
65	Cambodia	KHM	Equat	Equat	Equat
66	South Korea	KOR	T/C	T/C	T/C
67	Kuwait	KWT	A&T/C	A&T/C	A&Eq
68	Lao PDR	LAO	Equat	Equat	Equat
69	Lebanon	LBN	A&T/C	A&T/C	A&T/C
70	Liberia	LBR	Equat	Equat	Equat
71	Libya	LBY	A&T/C	Arid	Arid
72	Sri Lankla	LKA	Equat	Equat	Equat
73	Lesotho	LSO	T/C	T/C	T/C
74	Morocco	MAR	A&T/C	A&T/C	A&T/C
75	Madagascar	MDG	Equat	Equat	Equat
76	Mexico	MEX	Eq&Hi	Eq&Hi	Eq&Hi
77	Mali	MLI	A&T/C	Arid	Arid
78	Myanmar	MMR	Equat	Equat	Equat
79	Mongolia	MNG	A&T/C	Arid	Arid
80	Mozambique	MOZ	Equat	Equat	Equat
81	Mauritania	MRT	A&T/C	Arid	Arid
82	Malawi	MWI	A&T/C	A&T/C	A&Eq
83	Malaysia	MYS	Eq&Hi	Eq&Hi	A&Eq
84	Niger	NER	A&T/C	Arid	Arid
85	Nigeria	NGA	Equat	Equat	Equat
86	Nicaragua	NIC	Eq&Hi	Eq&Hi	Eq&Hi

	Country	wbcode	Number of Clusters: Four	Five	Six
87	Netherlands	NLD	T/C	T/C	T/C
88	Norway	NOR	T/C	T/C	T/C
89	Nepal	NPL	A&T/C	A&T/C	A&Eq
90	New Zealand	NZL	T/C	T/C	T/C
91	Oman	OMN	A&T/C	A&T/C	A&Eq
92	Pakistan	PAK	A&T/C	Arid	A&Eq
93	Panama	PAN	Equat	Equat	Equat
94	Philippines	PHL	A&T/C	A&T/C	A&Eq
95	Papua New Guinea	PNG	Eq&Hi	Eq&Hi	A&Eq
96	Poland	POL	T/C	T/C	T/C
97	Korea, DPR	PRK	A&T/C	A&T/C	A&Eq
98	Portugal	PRT	A&T/C	A&T/C	A&T/C
99	Paraguay	PRY	Eq&Hi	Eq&Hi	A&Eq
100	Qatar	QAT	A&T/C	A&T/C	A&Eq
101	Romania	ROM	T/C	T/C	T/C
102	Rwanda	RWA	Equat	Equat	Equat
103	Saudi Arabia	SAU	A&T/C	A&T/C	Eq&Hi
104	Sudan	SDN	Equat	Arid	Arid
105	Senegal	SEN	Equat	Arid	Arid
106	Sierra Leone	SLE	Eq&Hi	Equat	Equat
107	El Salvador	SLV	Equat	Equat	Equat
108	Somalia	SOM	A&T/C	Arid	Arid
109	Suriname	SUR	Eq&Hi	Eq&Hi	A&Eq
110	Sweden	SWE	T/C	T/C	T/C
111	Swaziland	SWZ	T/C	T/C	T/C
112	Syria	SYR	A&T/C	A&T/C	A&T/C
113	Chad	TCD	A&T/C	Arid	Arid
114	Togo	TGO	Equat	Equat	Equat
115	Thailand	THA	Equat	Equat	Equat
116	Trinidad & Tobago	TTO	A&T/C	A&T/C	A&Eq
117	Tunisia	TUN	A&T/C	A&T/C	A&T/C
118	Turkey	TUR	A&T/C	A&T/C	A&T/C
119	Tanzania	TZA	Equat	Equat	Equat
120	Uganda	UGA	Equat	Equat	Equat
121	Uruguay	URY	T/C	T/C	T/C
122	United States	USA	T/C	T/C	T/C
123	Venezuela	VEN	Equat	Equat	Equat
124	Vietnam	VNM	Equat	Equat	Equat
125	Yemen, Republic	YEM	A&T/C	Eq&Hi	Eq&Hi
126	South Africa	ZAF	A&T/C	Arid	Arid
127	Congo, DR	ZAR	Equat	Equat	Equat
128	Zimbabwe	ZWE	A&T/C	Arid	Arid

Notes: This table reports the cluster affiliation from our analysis imposing four, five or six clusters, respectively. Group names (stylised): A&T/C – Arid & Temperate/Cold; T/C – Temperate/Cold; Arid; Equat – Equatorial; A&Eq – Arid & Equatorial; Eq&Hi – Equatorial & Highland.

A-3. CLIMATE ZONES

Table A-IV: 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 category — this results in 12 rather than 15 categories.

A-4. AGRO-CLIMATIC DISTANCE

Country-specific information on the share of arable land by climate zones (following the Köppen-Geiger classification) are based on Matthews (1983), available in Gallup, Mellinger and Sachs (1999), information on contiguity (neighborhood) and inverse population-weighted distance between countries is taken from CEPII (Mayer and Zignago, 2006). We construct a Jaffe (1986) measure of agro-climatic distance, using the share of cultivated land within each of twelve climatic zones (h_{im}), such that for each country i the values in the twelve zones sum up to unity ($\sum_m h_{im} = 1$). The Jaffe measure for 'agro-climatic distance' between countries i and j is then

$$\omega_{ij} = \frac{\sum_m h_{im} h_{jm}}{\left(\sum_m h_{im}^2\right)^{1/2} \left(\sum_m h_{jm}^2\right)^{1/2}}$$

A-5. WEAK EXOGENEITY TESTS

In Table A-V we provide some evidence that our preferred CMG specifications (standard and agro-climatic distance models) can be interpreted as production functions and do not suffer from bias due to reverse causality.[31]

[31] These results represent weak exogeneity tests following Canning and Pedroni (2008). An error correction specification incorporating the residuals from our empirical models in Table 1 is estimated for output and each production input. For any causal relation between x and y at least one of the coefficients on the residuals in these five equations (we impose constant returns to scale) has to be statistically significant. If this 'error correction' (EC) term is significant in the output but not the input equations, we can suggest that the inputs 'cause' output ($x \rightarrow y$). Note how the misspecified 2FE and MG models find multiple causal relationships between the variables.

Table A-V: Canning & Pedroni (2008) tests for direction of causation

2FE	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-0.97	0.00	485.2	0.00	-0.142	-7.91	$x \rightarrow y$
tractor equation	0.18	0.17	456.2	0.00	0.024	1.81	$x_{-tr}, y \rightarrow x_{tr}$
livestock equation	0.38	0.00	351.0	0.00	0.043	3.77	$x_{-live}, y \rightarrow x_{live}$
fertilizer equation	0.10	0.42	432.3	0.00	0.141	1.82	$x_{-f}, y \rightarrow x_f$
land equation	0.37	0.00	395.2	0.00	0.011	1.91	$x_{-n}, y \rightarrow x_n$
MG	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-2.93	0.00	1,612.1	0.00	-0.976	-24.00	$x \rightarrow y$
tractor equation	-0.16	0.87	274.7	0.20	-0.029	-0.98	$x_{-tr}, y \nrightarrow x_{tr}$
livestock equation	0.03	0.98	307.6	0.01	0.015	0.55	$x_{-live}, y \rightarrow x_{live}$
fertilizer equation	-0.06	0.96	257.2	0.47	-0.116	-0.85	$x_{-f}, y \nrightarrow x_f$
land equation	-0.06	0.95	286.5	0.09	-0.004	-0.33	$x_{-n}, y \rightarrow x_n$
CMG standard	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-2.09	0.04	938.4	0.00	-0.925	-18.63	x, TFP $\rightarrow y$
tractor equation	-0.08	0.94	189.3	1.00	-0.045	-1.28	x_{-tr}, y, TFP $\nrightarrow x_{tr}$
livestock equation	0.12	0.91	331.1	0.00	0.011	0.27	x_{-live}, y, TFP $\rightarrow x_{live}$
fertilizer equation	-0.04	0.97	232.7	0.69	0.023	0.14	x_{-f}, y, TFP $\nrightarrow x_f$
land equation	0.06	0.95	225.7	0.79	0.010	0.68	x_{-n}, y, TFP $\nrightarrow x_n$
CMG neighbor	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-1.88	0.06	856.0	0.00	-0.784	-16.09	x, TFP $\rightarrow y$
tractor equation	0.03	0.97	260.4	0.22	0.032	0.79	x_{-tr}, y, TFP $\nrightarrow x_{tr}$
livestock equation	0.07	0.94	305.7	0.00	0.004	0.11	x_{-live}, y, TFP $\rightarrow x_{live}$
fertilizer equation	0.16	0.87	299.1	0.01	0.179	1.02	x_{-f}, y, TFP $\rightarrow x_f$
land equation	0.07	0.95	312.3	0.00	0.002	0.11	x_{-n}, y, TFP $\rightarrow x_n$
CMG distance	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-2.04	0.04	932.4	0.00	-0.997	-21.42	x, TFP $\rightarrow y$
tractor equation	-0.12	0.90	243.8	0.49	-0.041	-1.10	x_{-tr}, y, TFP $\nrightarrow x_{tr}$
livestock equation	0.10	0.92	301.9	0.01	0.050	1.21	x_{-live}, y, TFP $\rightarrow x_{live}$
fertilizer equation	-0.02	0.98	257.5	0.26	-0.059	-0.31	x_{-f}, y, TFP $\nrightarrow x_f$
land equation	-0.12	0.91	244.2	0.48	-0.001	-0.06	x_{-n}, y, TFP $\nrightarrow x_n$
CMG agro-climate	*GM*	(p)	*Fisher*	(p)	*mean* $\hat{\lambda}_i$	t-*ratio*	*Verdict*
output equation	-2.25	0.02	1,035.5	0.00	-0.935	-20.16	x, TFP $\rightarrow y$
tractor equation	-0.02	0.98	241.8	0.53	-0.013	-0.42	x_{-tr}, y, TFP $\nrightarrow x_{tr}$
livestock equation	0.15	0.88	380.0	0.00	0.048	1.23	x_{-live}, y, TFP $\rightarrow x_{live}$
fertilizer equation	0.07	0.94	242.5	0.52	-0.004	-0.02	x_{-f}, y, TFP $\nrightarrow x_f$
land equation	-0.09	0.93	227.3	0.77	-0.001	-0.08	x_{-n}, y, TFP $\nrightarrow x_n$

Notes: We report various test statistics for the null of no long-run causal impact between sets of different variables. In each case '*variable* equation' refers to the ECM regression with 'variable' on the LHS. GM gives the group-mean average of country-specific t-ratios for the coefficient on the disequilibrium term ($\hat{\lambda}_i$) which is distributed $N(0,1)$. Fisher gives $-2\sum_i \log\pi_i$ where π_i is the probability value of the country-specific t-ratio on the disequilibrium term. The Fisher statistic is distributed $\chi^2(2N)$. The final two columns but one report the mean estimate for $\hat{\lambda}_i$ and the associated t-ratio. In the 'Verdict' column we summarise the analysis, using x_{-tr} for 'all inputs other than tractors' and $\rightarrow$ as short-hand for 'does cause' and $\nrightarrow$ for 'does not cause'. TFP is included implicitly in the CMG models via cross-section averages. For comparability we impose constant returns on all models tested.

A-6. PROOF OF PROPOSITION 1

To begin the proof, first define the following elasticities:

$$\epsilon_A^{c_a} = \frac{\partial c_a}{\partial A}\frac{A}{c_a} \tag{28}$$

$$\epsilon_A^{Y_a} = \frac{\partial Y_a/L_a}{\partial A}\frac{A}{Y_a/L_a} \tag{29}$$

$$\epsilon_A^{L_a} = \frac{\partial L_a}{\partial A}\frac{A}{L_a} \tag{30}$$

$$\epsilon_A^{p_a} = \frac{\partial p_a}{\partial A}\frac{A}{p_a}. \tag{31}$$

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

$$L_a = \left(\frac{c_a L^{1-\beta_L}}{A}\right)^{\beta_L}. \tag{32}$$

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

$$\epsilon_A^{L_a} = \frac{1}{\beta_L}\left(\epsilon_A^{c_a} - 1\right). \tag{33}$$

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

$$p_a c_a = \alpha(w - p_a\bar{c}_a + \bar{c}_n) + p_a\bar{c}_n. \tag{34}$$

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

$$c_a = \alpha\left(1 + \frac{\bar{c}_n}{w}\right)AL_a^{\beta_L - 1} + (1-\alpha)\bar{c}_a. \tag{35}$$

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

$$\epsilon_A^{c_a} = \Omega\left(1 + (\beta_L - 1)\epsilon_A^{L_a}\right) \tag{36}$$

where

$$\Omega = \frac{\alpha\left(1 + \frac{\bar{c}_n}{w}\right)}{L_a/L}. \tag{37}$$

Solving (36) with (33) yields the following expression for the elasticity of agricultural labor with respect to agricultural productivity

$$\epsilon_A^{L_a} = \frac{\Omega - 1}{\beta_L(1-\Omega) + \Omega}. \tag{38}$$

This is clearly decreasing in β_L, as claimed in the proposition. Using this result in (36) yields that

$$\epsilon_A^{c_a} = \frac{\Omega}{\beta_L(1-\Omega) + \Omega}, \tag{39}$$

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

$$\epsilon_A^{Y_a} = \frac{1}{\beta_L(1-\Omega)+\Omega}. \tag{40}$$

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.

TECHNICAL APPENDIX

TA-1. ROBUSTNESS CHECK: NON-LIVESTOCK SAMPLE

The dominance of livestock breeding in some economies may be seen to have a distorting impact on our agricultural production function estimation. As a robustness check we therefore attempt to exclude those economies from the sample where livestock rearing appears to account for a significant share of total agricultural output. In the left panel of Figure TA-I we plot the country-specific averages of livestock share in total agricultural output (in value terms) for our sample of $N = 128$ countries. The clear bi-modal nature of the distribution is driven by 22 economies with an average livestock share of more than 60% – with the exception of a small number of low income countries (e.g. Lesotho, Mongolia, Somalia) these are predominantly developed economies in the temperate or cold climate zones, including Denmark, Germany, and the United Kingdom. If we exclude these economies, the distribution of livestock share is unimodal and normal, with a mean/median just under 30% and a standard deviation of 14% (right panel of Figure TA-I).

Table TA-I provides agricultural production function estimates for a model of agricultural output net of livestock, where we have dropped the previous 'livestock' variable (the cattle-equivalent head-count). Patterns of returns to scale as well as average coefficients for capital stock and fertilizer are next to identical to the alternative specification with the full sample data in Table 1 of the maintext. Land coefficients are however substantially elevated, and with the exception of the 'distance' specification the CMG estimates for labor are considerably higher as well. Taking diagnostics into account the agro-climatic distance version of the CMG still emerges as the preferred specification.

Figure TA-I: Distribution of Livestock Share in Total Agricultural Output

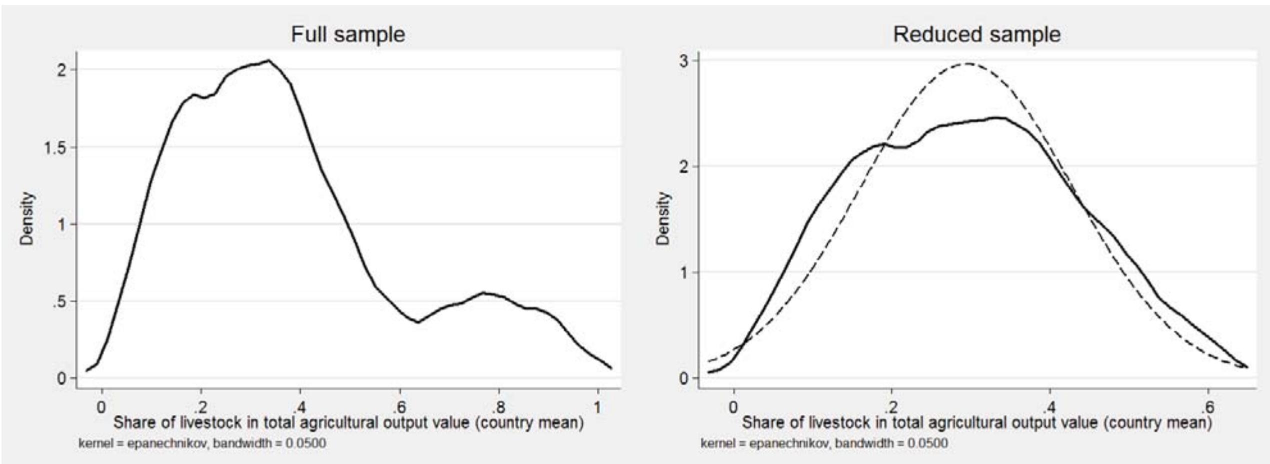

Notes: We plot the distribution of country-specific averages for the share of livestock in total agricultural output value. Left panel: full sample of $N = 128$; right panel: we drop 22 countries for which this average share exceeds 60%.

Table TA-II presents the labor coefficient analysis for the four broad climatic zones, where like in the full sample case the highland zone appears to have the highest labor coefficients, followed by the equatorial zone. In contrast to the earlier results the arid zone now has clearly higher average labor coefficients across the various scenarios than the temperate/cold zone.

We carry out cluster analysis on the agro-climatic similarity measures for the reduced sample of $N = 106$ countries and present cluster characteristics in Table TA-III and averaged labor coefficients in Table TA-IV. We pick 3 to 5 clusters – due to the one-sided reduction in the sample, which almost exclusively (15/22) dropped temperate/cold zone countries, the emerging clusters are no longer quite as clear-cut as in the full sample specification (note the presence of a 'diverse' group in the five cluster case). Average labor coefficients still broadly follow the patterns of the earlier results, with the highest labor coefficients in equatorial and highland climate zones and significantly lower coefficients in the temperate/cold and arid zones.

Table TA-I: Production Function Estimates – Reduced Sample

Weight matrix ♯	[1] 2FE	[2] MG	[3] CMG standard	[4] CMG neighbor	[5] CMG distance	[6] CMG agro-climate
Labor	-0.143	-0.502			-0.440	
	[4.63]**	[2.73]**			[2.99]**	
Tractors pw	0.072	0.062	0.120	0.098	0.075	0.093
$\hat{\beta}_K$	[13.78]**	[2.43]*	[5.27]**	[4.24]**	[2.85]**	[3.54]**
Fertilizer pw	0.097	0.041	0.043	0.050	0.034	0.043
$\hat{\beta}_F$	[20.12]**	[5.25]**	[5.15]**	[5.32]**	[4.45]**	[4.75]**
Land pw	0.565	0.307	0.417	0.391	0.230	0.368
$\hat{\beta}_N$	[16.46]**	[3.43]**	[6.26]**	[5.37]**	[2.66]**	[5.15]**
Returns to Scale ♭	DRS	DRS	CRS	CRS	DRS	CRS
Implied Avg $\hat{\beta}_L$	0.123	0.088	0.420	0.460	0.221	0.497
	[7.48]**	[0.43]	[5.92]**	[5.98]**	[1.28]	[6.48]**
Avg implied $\hat{\beta}_L$		0.136	0.356	0.396	0.225	0.431
		[0.91]	[5.15]**	[4.77]**	[2.06]*	[5.52]**
# rejecting CRS		49			34	
$\hat{\varepsilon}$ Stationarity †	I(1)	I(0)	I(0)	I(0)	I(0)	I(0)
$\hat{\varepsilon}$ CD Test (p) ‡	-2.00 (.05)	4.64 (.00)	-1.83 (.07)	3.80 (.00)	-0.01 (.99)	-1.06 (.29)
RMSE	0.187	0.074	0.070	0.071	0.062	0.070

Notes: Results for $n = 4,309$ observations from $N = 106$ countries – we omitted the countries with a dominance in livestock breeding. Estimators: 2FE – 2-way fixed effects, MG – Pesaran and Smith (1995) Mean Group, CMG – Pesaran (2006) Common Correlated Effects Mean Group. We present robust sample means for model parameters in [2]-[6]; [1] is a pooled model. * 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 based on the variance estimator following Pesaran and Smith (1995) in [2]-[6]. RMSE reports the root mean squared error.

Dependent variable – log output per worker, [1] dto. in 2FE transformation as described in Coakley et al. (2006).

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 $\hat{\beta}_L$' in a lower panel of the Table.

♭ 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.

♯ We apply different $N \times N$ weight matrices to construct the cross-section averages as described in the main text.

† Pesaran (2007) CIPS test results for residual nonstationarity: I(0) — stationary, I(1) — nonstationary. Full results available on request. ‡ Pesaran (2004) test for cross-section dependence (CD) in the residuals, H_0: no CD.

Table TA-II: Labor Coefficients and Climate Zones – Reduced Sample

	B	C/D	A	H
Panel A: Any share of land in climatic zone				
Mean $\hat{\beta}_L$	0.329	0.218	0.589	0.497
	[0.110]***	[0.114]*	[0.090]***	[0.103]***
Sample	50	50	69	41
Panel B: Share above zone mean				
Mean $\hat{\beta}_L$	0.363	0.165	0.589	0.717
	[0.137]**	[0.133]	[0.114]***	[0.102]***
Sample	34	38	51	22
Panel C: Share above 40%				
Mean $\hat{\beta}_L$	0.251	0.179	0.574	0.660
	[0.191]	[0.143]	[0.112]***	[0.120]***
Sample	16	35	52	11
Panel D: Share above 50%				
Mean $\hat{\beta}_L$	0.255	0.169	0.561	0.729
	[0.212]	[0.143]	[0.120]***	[0.125]***
Sample	15	33	47	7
Panel E: Share above 60%				
Mean $\hat{\beta}_L$	0.461	0.094	0.566	1.040
	[0.195]**	[0.153]	[0.127]***	[0.042]***
Sample	12	29	43	4

Notes: We analyse labor coefficients obtained from the analysis of the reduced sample ($N = 106$) where we exclude countries where livestock output makes up a large share of total agricultural output (in value terms) – production function results in Table TA-I. Each cell in this Table presents the robust mean of the estimated labor coefficient $\hat{\beta}_L$ given the criterion indicated. 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 — this indicates the skewness of distribution in the equatorial and temperate/cold zones. Sample indicates the number of observations from which the robust mean is constructed.

Table TA-III: Clusters – Descriptives

	Panel A: Three Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.070 [0.220]	0.285 [0.348]	0.556 [0.414]	0.089 [0.201]	50
Equatorial & Highland	0.683 [0.299]	0.034 [0.070]	0.026 [0.094]	0.257 [0.264]	15
Equatorial	0.807 [0.229]	0.097 [0.183]	0.067 [0.133]	0.029 [0.076]	41

	Panel B: Four Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.063 [0.211]	0.288 [0.361]	0.618 [0.390]	0.032 [0.080]	45
Arid & Equatorial	0.797 [0.228]	0.092 [0.177]	0.069 [0.136]	0.041 [0.100]	45
Equatorial	0.947 [0.083]	0.000 [0.000]	0.000 [0.000]	0.053 [0.083]	6
Equatorial & Highland	0.245 [0.255]	0.164 [0.189]	0.002 [0.008]	0.588 [0.169]	10

	Panel C: Five Clusters				
Climate Zone	A	B	C/D	H	*N*
Cluster					
Arid & Temperate/Cold	0.000 [0.000]	0.179 [0.199]	0.770 [0.219]	0.051 [0.124]	16
Diverse	0.177 [0.355]	0.253 [0.365]	0.442 [0.447]	0.128 [0.245]	35
Arid & Equatorial	0.293 [0.112]	0.497 [0.298]	0.075 [0.172]	0.135 [0.238]	11
Equatorial	0.874 [0.157]	0.045 [0.088]	0.056 [0.114]	0.026 [0.076]	31
Equatorial & Highland	0.795 [0.221]	0.012 [0.042]	0.045 [0.114]	0.148 [0.202]	13

Notes: The table presents the mean share of land in each of the four climatic zones (denoted by their Köppen-Geiger letters) for groups of countries determined endogenously using the cluster analysis technique described in the main test. The standard deviations of these means are in brackets.

Table TA-IV: Labor Coefficients by Clusters

Panel A: Three Clusters

Cluster	Arid & temperate/cold	Arid & Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	0.287 [0.115]**	0.482 [0.126]***	0.762 [0.135]***
N	50	41	15

Panel B: Four Clusters

Cluster	Arid & temperate/cold	Arid & Equatorial	Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	0.253 [0.118]**	0.506 [0.117]***	0.876 [0.705]	0.813 [0.123]***
N	45	45	6	10

Panel C: Five Clusters

Cluster	Arid & temperate/cold	Diverse	Arid & Equatorial	Equatorial	Equatorial & Highland
Mean $\hat{\beta}_L$	-0.225 [0.202]	0.512 [0.139]***	0.796 [0.261]**	0.382 [0.132]***	0.802 [0.182]***
N	16	35	11	31	13

Notes: We present robust mean estimates (standard errors in brackets) for the labor coefficients (using estimates from Model [6] in Table 1) based on various groupings. In Panels A-C we use cluster analysis specifying 3, 4 and 5 6 clusters, respectively, to create our groups (see Table TA-III for group descriptives).

TA-2. ROBUSTNESS CHECK: PARAMETER STABILITY

In Figure TA-II we present average labour coefficients for each of the four clusters used in the main text of the paper where the country-specific estimates of β_L are based on recursive estimates: in the left panel we estimate the agricultural production function with a sample from 1961 to 1980, compute the average β_L by agro-climatic cluster and then increase the end date of the sample by one year at a time until we reach 2002. Each line plot charts, from left to right, the evolution of the cluster-specific mean labor coefficient as the sample size is increased. In the right sample we do the same beginning with a sample from 1981 to 2002, then shifting the start year of the sample one year at a time until we reach the full sample starting in 1961 – again sample size increases from left to right. In either plot we can see that mean coefficient estimates for each cluster stabilize once the time series reaches around 30 years — it thus appears that any variation in the estimates is down to small sample imprecision rather than systematic change over time.

Figure TA-II: Average Labor Coefficient by Climate Zone: Recursive Estimates

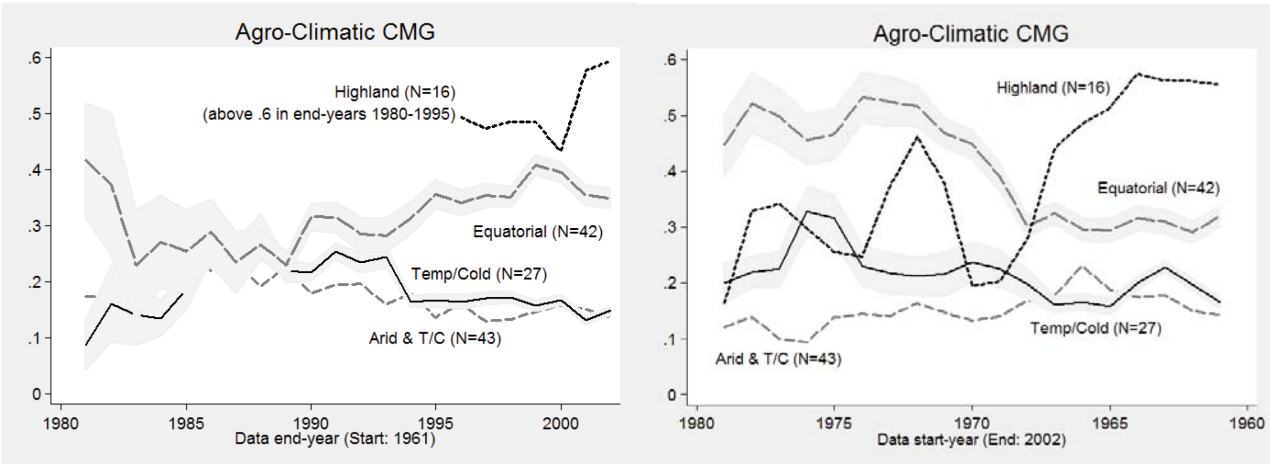

Notes: Shaded areas represent 90% confidence intervals for the robust mean estimates within each cluster, where clusters are defined as detailed in Table A-III above. In the left plot we omit the early estimates of the Highland cluster (all above .6) in order to aid presentation in the graph.