Economica

Economica (2020) **87**, 901–937
doi:10.1111/ecca.12335

The Role of Land in Temperate and Tropical Agriculture

By T. Ryan Johnson† and Dietrich Vollrath‡

†*Washington University* ‡*University of Houston*

Final version received 11 December 2019.

We document differences in the elasticity of agricultural output with respect to land in temperate and tropical regions. We estimate this elasticity from the relationship of rural labour/land ratios and agroclimatic constraints using global district-level data. We find that the elasticity in temperate areas (0.285) is higher than in the tropics (0.126), and that this is not an artefact of the level of development. The land elasticity influences the degree of decreasing returns to labour and capital in agriculture, and thus how sensitive living standards are to shocks in productivity and population. Evidence from the postwar mortality transition supports this prediction.

Introduction

Agricultural production relies on the use of a finite (or inelastically supplied) resource, namely land. But that reliance on land need not be identical in different locations. To be specific, the elasticity of agricultural output with respect to land may differ by climate or the type of crops suitable for production. This land elasticity is relevant to any study of growth and development that includes an agricultural sector, as with the mild assumption of constant returns to scale, one minus the land elasticity tells us how sensitive agricultural output is to the use of non-land inputs like capital and labour. This in turn determines how many non-land inputs move out of (or into) agriculture in response to shocks to productivity and population. Differences in the land elasticity by crop or climate thus imply differences in the reactions of economies to shocks, with implications for studies of comparative development, structural change, Malthusian stagnation, the take-off to sustained growth, and long-run growth with finite resources.[1]

In this paper, we estimate the land elasticity, and show that it varies across different agricultural regions and climate types. Estimating the parameter(s) of an agricultural production function is not straightforward, for the standard reasons that total factor productivity and some inputs may be unobserved. To address these issues, we first develop a method for estimating the aggregate land elasticity using the relationship between the labour/land ratio in agriculture and the potential agroclimatic yield across small geographic units (e.g. second-level districts within states/provinces). The methodology relies on the mobility of labour between districts within states, as well as between agricultural and non-agricultural uses. We show that this mobility is supported in the data, first by reviewing recent research on this subject, second by providing evidence drawn from the Demographic and Health Surveys, and third by documenting the small size (in terms of population and area) of districts.

Given mobility, our method does not require us to identify exactly what the inputs are beyond land and labour, avoiding mismeasurement issues. We use agroclimatic yield data to give us a source of exogenous variation in productivity, and combine that with measures of district-level development (e.g. night lights, road density and urbanization) to control for other unobservable elements of agricultural productivity. Most relevant, our estimates are made using only within-state variation across districts, meaning that unobservable variation in productivity across states, as well as across countries, is excluded from the estimates. This means that our framework is robust to arbitrary distortions (e.g. taxes or subsidies) of agricultural and factor prices at the state level.

We assemble data at the district level for rural labour/land ratios in the year 2000, and combine those with a measure of potential agroclimatic yield in districts built from the data of Galor and Özak (2016). As in their work, our measure is built on constraints plausibly unaffected by human activity (e.g. soil quality and length of growing season) from the Global Agro-Ecological Zone (GAEZ) project (Food and Agriculture Organization 2012), combined with information on the calorie contents of various crops. Grid-cell potential caloric yields are aggregated to the district level to serve as our measure of agroclimatic yield.[2]

In the end, we have a dataset of 28,475 districts, coming from 2282 states in 151 countries. We then divide districts into 'temperate' and 'tropical' regions based on their agroclimatic characteristics. In our baseline, we make this division based on the types of crops that can be grown within a district. The temperate region includes districts that can grow crops such as wheat, barley and rye, while the tropical region includes districts that can grow crops such as paddy rice, cassava and pearl millet. We also divide districts based on their frost-free days (e.g. tropical areas are frost-free all year round, while temperate areas are not), or by their Köppen–Geiger climate classification (Kottek *et al.* 2006), and our results are similar. Regardless of the definition, our assignment is made at the district level and we do not assume that agriculture has a homogeneous land elasticity within a country.

Our baseline estimate is that the land elasticity is 0.285 in temperate districts. In contrast, our baseline estimate of the land elasticity is only 0.126 for tropical districts. The difference is statistically significant, and is robust to the exclusion of districts that contain large urban areas, districts that are large relative to their state, or districts from any developed country. Further, the results are consistent if we use alternative measures of rural labour/land ratios, alternative measures of the potential agroclimatic yield, or alternative measures of the area of agricultural land used within a district. In all cases, the aggregate land elasticity in temperate districts is approximately 0.16 higher than in tropical districts, and the difference is statistically significant.[3] As the measure of agroclimatic yield that we use is based on staple crops, our results should be interpreted as differences in the land elasticity in staple crop production.[4]

Relative to the existing literature, our approach to estimating the aggregate land elasticity has several advantages. The standard approach has been to use country-level panel data (Hayami and Ruttan 1970, 1985; Craig *et al.* 1997; Martin and Mitra 2001; Mundlak 2000; Mundlak *et al.* 2012; Eberhardt and Teal 2013) to estimate agricultural production functions, with a common set of coefficients across countries for each input, including land. Issues arise with unobserved productivity, the measurement of non-land inputs, and the assumption that coefficients are common to all countries. Some have examined heterogeneity in these coefficients (Gutierrez and Gutierrez 2003; Wiebe *et al.* 2003) by continent, while others have attempted to estimate country-level coefficients using factor analysis to address unobserved productivity (Eberhardt and Teal 2013; Eberhardt and Vollrath 2018). Compared to this, our district-level data allow us to control for unobserved country and state-level effects, and the use of agroclimatic yield data gives us an explicit measure of productivity.[5]

As may be apparent, we are estimating the elasticity of not a *farm*-level production function, but rather an *aggregate* production function. Farm-level estimates of the land elasticity would not necessarily be informative about the aggregate production function, given that those estimates would refer to farmers using a given technique, while the aggregate function can be thought of as an envelope across techniques available to farmers (Hayami and Ruttan 1970).[6] The aggregate land elasticity is a useful parameter for studying the role of the agricultural sector and its interaction with other sectors at the macro level, as we discuss below, while farm-level elasticities would be useful for studying farm-level policies or outcomes within the agricultural sector itself. This distinction explains one of the limitations of our study, which is that we cannot use our results to identify why the aggregate land elasticity differs between temperate and tropical regions. An explanation would require details on the interaction of farmers with biological production functions for specific crops that are beyond the scope of this paper.

With that caveat in mind, we show in the second half of the paper that the aggregate land elasticity is central to any study that looks at the relationship of agriculture to non-agriculture, and the variation we have identified between temperate and tropical regions has implications for development. The intuition is that the land elasticity dictates—given an assumption of constant returns to scale—the degree of decreasing returns to scale for labour and capital in agriculture. A large land elasticity implies more severe decreasing returns, and in response to shocks to productivity or population, this means more severe movements of those factors into or out of agriculture. Temperate areas therefore have exaggerated responses to shocks relative to tropical areas. This is a benefit to temperate areas when shocks are positive (e.g. higher total factor productivity (TFP) or lower population growth), but a burden in the face of negative shocks (e.g. lower TFP or higher population growth).

In the last part of the paper we confirm these predictions by using data from Acemoglu and Johnson (2007) to examine the consequence of population shocks arising from the epidemiological transition after the Second World War. The shock to mortality was negatively correlated with GDP per capita, and GDP per worker, across all developing countries. But we find that the size of that negative correlation was three times larger for countries with temperate land elasticities compared to countries with tropical land elasticities, consistent with our intuition. The difference in correlation is statistically significant, and holds whether we measure the population shock in terms of mortality or life expectancy.

At a broader level, variation in the land elasticity may be relevant for the study of historical and contemporary development. For any given positive shock to productivity (or negative shock to population growth), areas with temperate land elasticities experience more urbanization and faster growth in living standards, whatever the fundamental driver of those shocks: institutions, geography or culture.[7] This may help explain why it was that western Europe, with a high aggregate land elasticity, diverged from Asia, with a low aggregate land elasticity, even though western Europe did not have an advantage in technological or institutional improvements.[8] It may also help explain why the tropical areas of Central America and Sub-Saharan Africa, with relatively low land elasticities, lagged behind other areas following decolonization.[9]

The paper proceeds as follows. Section I describes the basic nature of the district-level data that we use, including evidence on mobility across these districts, which informs our estimation. Section II shows how we estimate the land elasticity, and which assumptions about mobility are required for identification. Section III presents the data and results, while Section IV discusses the aggregate implications of variation in land elasticity. Section V concludes.

I. District-level Characteristics

It will be useful to first establish the characteristics of the districts that we use as our units of observations, and show that populations are mobile across district boundaries. This will inform our method for identifying the land elasticity.

A district, as the term is used in our paper, is a second-level administrative unit within a country, regardless of the terminology used. It is thus part of a first-level administrative unit, which we call a state. In India our terminology matches the local terminology. For example, the district of Kadapa lies within the state of Andhra Pradesh. For the USA, a 'district' is referred to as a county. Marathon County, in the state of Wisconsin, is an example of a district in our data. In Nigeria, a 'district' is a local government area, which is part of a state. Thus Demsa, in the state of Adamawa, is a district in our data. In total, we have 28,475 districts, coming from 2282 states in 151 countries.[10]

These districts tend to be small, both in absolute terms and relative to the states in which they reside. Table 1 shows summary statistics on population from the Global Rural–Urban Mapping Project (GRUMP) (Center for International Earth Science Information Network (CIESIN) *et al.* 2011) for districts in Panel A. The mean population of a district is 105,800 people, although the median district has only 22,600. For our empirical work, the rural population will be crucial. The rural population of a district is even smaller, with a mean value of 75,800 and a median of only 16,000.[11]

By nature, urban population is concentrated into small areas, so the distribution of urban populations across districts is skewed within states. The average urban population of a district is around 29,900, but the median district has an urban population of zero. Thus for many districts, the urban share of district population is zero, and the average share is only about one-fifth (0.19). At the other extreme, we do have some districts with a large percentage of urban population, with a share of 0.67 at the 90th percentile of our sample.

As a proportion of their state, most districts are also quite small. The average district represents only 5% of its state population, with a median of only 2%. For most states, one district often represents the majority of state population, and that almost invariably contains an urban area. The median district has about 0% of the state urban population. A similar finding holds for absolute area, where the average district represents only about 6% of total state area and the median district is only 2% of state area. The median district encompasses only 53,000 hectares, or 530 square kilometres. That represents a square of only about 23 kilometres on each side.

The districts in our data are small in absolute and relative terms, thus it seems reasonable to guess that workers are mobile between districts. Recent work by Young (2013) and Hicks *et al.* (2017) confirms that. Those studies find that first, within developing countries there is significant movement of workers between urban and rural areas on a regular basis, either in the universe of Demographic and Health Surveys (DHS) studies (Young 2013) or for a set of longitudinal studies (Hicks *et al.* 2017). Both studies find rural-to-urban movement, but also substantial urban-to-rural movement. Second, consistent with economic intuition, this movement is associated with an equalization of the wage per unit of human capital across urban and rural areas.[12] Combining the first finding with our summary statistics showing the concentration of urban population in a handful of districts, the implication is that there must be movement of people between districts within any given state. There may be more extensive movement of people between states themselves, but for our empirical setting, movement between districts within a state is most relevant.

TABLE 1
SUMMARY STATISTICS FOR DISTRICT-LEVEL DATA, 2000

			Percentiles:				
	Mean	SD	10th	25th	50th	75th	90th
Panel A: Population and area							
Total population (000s)	105.8	480.2	3.7	8.4	22.6	60.9	149.1
Rural population (000s)	75.8	357.5	3.2	6.5	16.0	40.8	102.8
Urban population (000s)	29.9	167.0	0.0	0.0	0.0	12.9	52.9
Urban share of district population	0.19	0.27	0.00	0.00	0.00	0.35	0.67
Share of state population	0.05	0.09	0.00	0.00	0.02	0.06	0.14
Share of state urban population	0.04	0.13	0.00	0.00	0.00	0.01	0.08
Share of state area	0.06	0.09	0.00	0.01	0.02	0.07	0.15
Total area (000s ha)	179.8	475.7	8.7	17.1	53.3	160.4	394.5
Panel B: Labour/land ratios, yields, and other controls							
Labour/land (persons/ha)	0.76	1.19	0.05	0.13	0.34	0.81	1.92
Caloric yield (million cal/ha)	10.83	4.83	5.01	7.20	10.68	13.83	16.93
Log light density	−2.97	2.92	−6.43	−3.93	−2.63	−1.03	0.20
Road density (km per sq. km)	0.40	0.53	0.06	0.11	0.22	0.47	0.90
Share of roads, highway	0.03	0.09	0.00	0.00	0.00	0.00	0.08
Share of roads, primary	0.15	0.20	0.00	0.00	0.09	0.22	0.39
Share of roads, secondary	0.33	0.30	0.00	0.08	0.24	0.51	0.82
Slope index	70.78	24.18	33.64	52.63	77.69	91.88	97.05
Distance (km) to city of 100,000	61.20	70.57	6.89	16.95	40.33	78.50	137.32

Notes
These are summary statistics for districts used in the regression analysis. There are a total of 28,475 observations for each variable (these come from 2282 states in 151 countries). All population data are derived from GRUMP (CIESIN *et al.* 2011). Districts are defined by the Global Administrative Areas Project (2019), and correspond to second-level administrative areas within countries (e.g. counties). Caloric yield A_{is}^{GAEZ} is calculated by the authors using data from Galor and Özak (2016). Rural labour/land ratio L_{Ais}/X_{is} is calculated by the authors using data from Goldewijk *et al.* (2011) for rural population. Both caloric yield and rural labour/land ratio were trimmed at the 99th and 1st percentiles of their raw data prior to calculating the summary statistics in this table. Log mean light density is derived from the Global Radiance Calibrated Nighttime Lights data provided by NOAA/NGDC, as in Henderson *et al.* (2016). Road density and the share of roads by type are from Meijer *et al.* (2018). The slope index is from Food and Agriculture Organization (2012). Distance from nearest city of 100,000 population is the authors' calculation using centroids of districts (see text).

To illustrate the amount of migration within developing countries, we use data from the DHS (distributed by ICF, 1986–2017) similar to Young (2013). Table 2 shows summary statistics on migration taken from 86 separate surveys in the DHS (Panel A), or 68 surveys (Panel B). The numbers reported in the table are summary statistics of survey level averages. Thus the first row of figures shows that, on average across the 86 surveys, 49% of respondents report moving at some point in their life. Even at the 10th percentile, 32% of individuals report moving at some point. If we restrict ourselves to people who moved within the last 5 years, the average across all surveys is still 21%. If we limit ourselves instead to only those who are 25–50 years old—prime working age—then the percentage that have moved at some point is 54%, and the 10th percentile is 35%.

Of course, that movement may not reflect a large geographic change, or a change from rural to urban ares. In Panel B of Table 2, we look at the self-reported origin of those who moved into either urban or rural areas. Of those who moved into urban areas, on average across surveys 40% reported being from the 'country', with a 10th percentile

TABLE 2
SUMMARY STATISTICS ON MIGRATION FROM THE DHS

	Mean	SD	Percentiles: 10th	25th	50th	75th	90th
Panel A: Share moving measured by:							
All movers/all individuals	0.49	0.14	0.32	0.42	0.50	0.59	0.67
Movers in last 5 years/all individuals	0.21	0.09	0.11	0.15	0.21	0.26	0.33
Movers aged 25–50/all aged 25–50	0.54	0.15	0.35	0.43	0.55	0.66	0.73
Panel B: Share of movers to location by self-reported origin:							
To urban areas from country	0.40	0.15	0.19	0.32	0.41	0.51	0.59
To rural areas from city or town	0.17	0.12	0.03	0.07	0.14	0.27	0.35

Notes
This table shows the overall prevalence of migration (Panel A), and the prevalence of migration between different areas (Panel B). These are summary statistics of shares calculated from individual surveys in the DHS. Panel A is based on 86 surveys representing 43 countries. Panel B is based on 68 surveys representing 34 countries.

across surveys of 19% and a 90th percentile of 59%, indicating that urban in-migrants were not simply arriving from other cities or towns. In the last row, we look instead at the percentage of movers to rural areas who reported being from either 'city' or 'town'. On average, 17% came from those areas. These numbers are lower, consistent with Young (2013), and in part reflect the general drift of urbanization over time.

Overall, what the district-level statistics and the migration data indicate is that there is substantial movement of workers across districts within states. There may in addition be movement of people between states within a given country, but as noted, that is not something on which we will rely in our empirical setting. We focused on developing countries in this section as concerns about frictions in the movement of people at the district level may be most pronounced for them. Our results are all robust to excluding developed nations, or excluding specific cases (e.g. China) that have particular migration restrictions.

II. IDENTIFYING THE AGGREGATE LAND ELASTICITY

Given the information on the small size of districts relative to their states, the evidence on the movement of people between urban and rural areas (and thus across districts in many cases), and the finding that this movement is associated with equalized wages between areas, we use this to build up an identification strategy for the land elasticity. We will be using variation in the labour/land ratio and agricultural productivity across districts *within* states to identify this elasticity, where the movement of workers across districts *within* states will allow us to eliminate several confounders by using state fixed effects. To do this, we will be deriving a labour demand function for the agricultural sector of a district, given some assumptions about the production function in both agriculture and non-agriculture.

Production and optimization

Consider district i located in a state s. There are a total of L_{is} workers in the district, who work in either the agricultural (A) or non-agricultural (N) sector, so that

$L_{is} = L_{Ais} + L_{Nis}$. There is also a total amount of capital K_{is}, which is also used in both sectors, so that $K_{is} = K_{Ais} + K_{Nis}$.[13]

Let the agricultural production function for a district be given by

$$Y_{Ais} = A_{Ais} X_{is}^{\beta} \left(K_{Ais}^{\phi} L_{Ais}^{1-\phi} \right)^{1-\beta}, \tag{1}$$

where A_{Ais} is total factor productivity and X_{is} is land. The land elasticity that we are interested in estimating is β.[14] We assume that agricultural operators in a district try to maximize profits, and take as given the wage of agricultural workers (w_{Ais}), the rental rate of agricultural capital (r_{Ais}), and the (state-level) price of agriculture goods relative to non-agricultural goods (p_{As}) facing them.

In addition, we allow for a revenue-wedge or price-wedge in each district, τ_{Ais}, that acts like a tax (or a subsidy if τ_{Ais} is negative) to producers in district i. An example of this wedge would be transportation costs, so that remote districts receive a lower price net of transport for their output. The profits of the agricultural sector in district i are therefore

$$\pi_{Ai} = (1 - \tau_{Asi}) p_{As} Y_{Ais} - w_{Ais} L_{Ais} - r_{Ais} K_{Ais}. \tag{2}$$

The first-order conditions of the profit maximization problem are

$$\begin{cases} w_{Ais} = (1-\phi)(1-\beta)(1-\tau_{Asi}) p_{As} \frac{Y_{Ais}}{L_{Ais}}, \\ r_{Ais} = \phi(1-\beta)(1-\tau_{Asi}) p_{As} \frac{Y_{Ais}}{K_{Ais}}. \end{cases} \tag{3}$$

Given these two conditions, the agricultural capital/labour ratio used in the district will be

$$\frac{K_{Ais}}{L_{Ais}} = \frac{\phi}{1-\phi} \frac{w_{Ais}}{r_{Ais}}. \tag{4}$$

Non-agricultural production in the district is given by

$$Y_{Nis} = A_{Nis} K_{Nis}^{\phi} L_{Nis}^{1-\phi}, \tag{5}$$

and the non-agricultural operators are also profit maximizers, who take the wage of non-agricultural workers (w_{Nis}) and the rental rate of non-agricultural capital (r_{Nis}) as given. They also face an arbitrary revenue-wedge or price-wedge of τ_{Nsi}, where again transport costs would be a natural interpretation. Their profits are

$$\pi_{Ni} = (1 - \tau_{Nsi}) Y_{Nis} - w_{Nis} L_{Nis} - r_{Nis} K_{Nis}, \tag{6}$$

where

$$w_{Nis} = (1-\phi)(1-\tau_{Nsi}) \frac{Y_{Nis}}{L_{Nis}},$$
$$r_{Nis} = \phi(1-\tau_{Nsi}) \frac{Y_{Nis}}{K_{Nis}},$$

which will result in a capital/labour ratio in non-agriculture of

$$(7) \qquad \frac{K_{Nis}}{L_{Nis}} = \frac{\phi}{1-\phi}\frac{w_{Nis}}{r_{Nis}}.$$

Mobility and the labour/land ratio

At this point we appeal to the characteristics of districts and the migration evidence from the prior section to apply several assumptions. First, based on the evidence from Young (2013), Hicks *et al*. (2017) and Table 2, we assume that labour is mobile across districts within a state such that $w_{Ais} = w_{As}$. Second, based on the sources just cited, we assume that labour is mobile between the agricultural and non-agricultural sectors, so that $w_{Ais} = w_{Nis}$ within any given district.

The last assumption that we make is that capital is also mobile *within* a district between agriculture and non-agriculture, such that $r_{Ais} = r_{Nis}$, although it need not be mobile across districts. We do not provide direct evidence for this assumption, as with the migration data. Making this assumption (or dropping it) would change the nature of controls that we need to include in our regressions, and we show later in the paper that for a subset of districts for which certain agricultural-specific capital controls are available, our results hold.[15]

With the assumptions that capital and labour are moving between agricultural and non-agricultural sectors within districts, and given equations (4) and (7), it will be the case that

$$\frac{K_{Ais}}{L_{Ais}} = \frac{K_{Nis}}{L_{Nis}} = \frac{K_{is}}{L_{is}},$$

or that the capital/labour ratio in both sectors will be equal to the aggregate capital/ labour ratio of a given district.

Incorporating that into the first-order condition for agricultural labour in equation (3), substituting in for the production function in equation (1), and applying the assumptions $w_{Ais} = w_{As}$, we arrive at

$$(8) \qquad w_{As} = (1-\phi)(1-\beta)(1-\tau_{Asi})p_{As}A_{Ais}\left(\frac{X_{is}}{L_{Ais}}\right)^{\beta}\left(\frac{K_{is}}{L_{is}}\right)^{\phi(1-\beta)},$$

which represents a labour demand curve relating w_{As} to L_{Ais}, where productivity, the capital/labour ratio, the wedge and the relative price of agricultural output act as demand curve shifters.

Taking logs of equation (8) and rearranging gives

$$(9) \qquad \ln A_{Ais} = \beta\ln\left(\frac{L_{Ais}}{X_{is}}\right) - \phi(1-\beta)\ln\left(\frac{K_{is}}{L_{is}}\right) - \ln(1-\tau_{Asi}) - \ln\left(\frac{w_{As}}{p_{As}}\right) + \ln(1-\phi)(1-\beta).$$

Examining equation (9), there is a linear relationship between (log) productivity and the (log) labour/land ratio in agriculture, and the coefficient on the labour/land ratio is simply β. In the next section we will describe how we measure agricultural productivity A_{Ais}, but for the moment take that as given. In principle, we should be able to use the relationship of (log) agricultural productivity and (log) labour/land ratios across districts to identify the value of β in a regression.

Doing this requires that we account for the additional terms in equation (9). The first is the (log) capital/labour ratio in the district, which would independently influence the labour/land ratio by affecting the productivity of workers. For this term we will introduce controls into our regressions, such as the density of night lights and/or measures of real assets owned by households from the DHS.[16] The second additional term is the (log) of the agricultural wedge term $1-\tau_{Asi}$. As noted, this could represent differences in transportation costs, and in our regressions we introduce road density, ruggedness and distance to major cities as controls for those costs.

The final term in equation (9) is state-specific but not district-specific. As such, it can be accounted for by state fixed effects. The real wage w_{As}/p_{As} is common to all districts, given the movement of labour and the relatively small size of districts within the state economy. This is not assumed to be a competitive equilibrium real wage, and it can contain any arbitrary distortion to the relative price of agricultural goods or wages at the state level. This allows for heterogeneity in wages and distortions across *states* within a given country in our empirical work. This further implies that country-level distortions to agricultural prices (e.g. tariffs, subsidies or taxes) do not bias our results in any way.

The main threat to our identification of β comes from the possibility that the real wage w_{As}/p_{As} is not equalized across districts within states. In that case, equation (9) holds for each individual district, but without an explicit way to control for the real wage, we would have a built-in bias of our estimate towards zero, as the labour/land ratio and the unique real wage within a district (an omitted variable) would be negatively related by definition. As we have argued, the evidence on migration and the small size of districts would support the assumption that the real wage is state-specific, and thus this bias is not present. We will also, as part of our robustness checks, exclude districts from our regressions that one may worry do not conform to the assumption of a common real wage (e.g. large urban districts), and all the results go through.

III. Estimates of the Aggregate Land Elasticity

To build our actual estimation equation, we need to specify one final thing, the measurement of agricultural productivity A_{Ais}. To do this, we rely on the work of Galor and Özak (2016), which it itself built on the Global Agro-Ecological Zone (GAEZ) project of the Food and Agriculture Organization (2012). We describe the GAEZ data in detail below, but consider it to be a noisy measure of true agroclimatic productivity. We thus break down agricultural productivity as

$$\ln A_{Ais} = \ln A_{As} + \ln A_{Ais}^{\mathrm{GAEZ}} + \varepsilon_{is},$$

where $\ln A_{As}$ captures the state-specific level of non-agroclimatic productivity (e.g. culture or institutions), while $\ln A_{Ais}^{\mathrm{GAEZ}}$ captures the agroclimatic elements of productivity, and ε_{is} represents noise in this measure of true agroclimatic conditions. In short, we assume that the GAEZ project did not make systematic errors in measuring agroclimatic productivity.[17]

We combine this relationship for agricultural productivity with the relationship in equation (9) to form our estimation specification, which includes an additional subscript g to account for the fact that we will be running this regression for a specific geographic region (e.g. tropical or temperate):

$$\ln A_{Aisg}^{\mathrm{GAEZ}} = \alpha_g + \beta_g \ln\left(\frac{L_{Aisg}}{X_{isg}}\right) + \gamma_s + \delta_g' \mathbf{Z}_{isg} + \varepsilon_{isg}, \tag{10}$$

where i denotes a district (e.g. Saoguan) in state s (e.g. Guangdong in China), which is part of a geographic region g. As can be seen, the coefficient β_g is unique to a geographic region. We will assign districts to a geographic region based on some physical characteristic (e.g. temperate climate), and all districts within that geographic region will be assumed to have an identical value for β_g. Our hypothesis is that the values of β_g vary with geographic characteristics, and over the course of the empirical work we will document that there are differences in β_g between geographic regions.

The term α_g is a constant. The value γ_s is the state fixed effect, and it picks up the real wage w_{As}/p_{As} as well as the state-specific level of non-agroclimatic productivity A_{As}. The term $\mathbf{Z}_{isg}$ is the set of controls that we use to proxy for the district capital/labour ratio K_{is}/L_{is} and the district-specific transportation costs τ_{Asi}. The δ_g' are the coefficients on those controls.

Standard errors ε_{isg} is a noise term, and we allow that it may be spatially autocorrelated. To account for this in our estimation, we use Conley standard errors. For any given district i, the error term of any other district that has a centroid (latitude/longitude) within 500 km of the centroid (latitude/longitude) of district i is allowed to have a non-zero covariance with ε_{isg}. The covariance of all other districts outside that 500 km window is presumed to be zero. Allowing the weight on the covariance to decay with distance from the centroid of district i does not change the results in a material way. We also experimented with other windows (1000 km, 2000 km), but we obtain similar standard errors using 500 km and report those.

Hypothesis testing We will be estimating equation (10) for geographic regions g. The typical significance test of estimated coefficients, with a null hypothesis that $\beta_g = 0$, is a test of whether the land elasticity is zero in region g. As will be seen in the results, we can reject this null hypothesis in all subsamples.

What is more relevant is whether the β_g that we estimate for one geographic region is statistically different from the β_g that we estimate using a different region. We choose one region (e.g. temperate) to be a reference region, and then test the estimated $\widehat{\beta}_g$ values for a different region (e.g. tropical) against $\widehat{\beta}_{\mathrm{Ref}}$. In practice, this is implemented as a simple interaction regression, where $I(\mathrm{Ref})$ is an indicator variable for inclusion in the reference region. The specification is

$$\begin{aligned} \ln A_{isg}^{\mathrm{GAEZ}} =& \alpha_g + \beta_g \ln\left(\frac{L_{Aisg}}{X_{isg}}\right) + (\beta_{\mathrm{Ref}} - \beta_g) \ln\left(\frac{L_{Aisg}}{X_{isg}}\right) \times I(\mathrm{Ref}) \\ & + \gamma_s + \delta_g' \mathbf{Z}_{isg} + (\delta_{\mathrm{Ref}}' - \delta_g') \mathbf{Z}_{isg} \times I(\mathrm{Ref}) + \varepsilon_{isg}. \end{aligned} \tag{11}$$

We then perform a statistical test with the null hypothesis of $H_0 : (\beta_{\text{Ref}} - \beta_g) = 0$ using the results of this interaction regression. Rejecting this hypothesis indicates that β_{Ref} and β_g are statistically different, and for our purposes this is the hypothesis of interest.

District population, productivity and other data

Population The underlying population data come from GRUMP (CIESIN *et al.* 2011), and are provided at a 30 arc-second (approximately 1 km) grid-cell resolution. This project provides counts of total population as well as urban and rural populations for each cell, and is an extension of the Gridded Population of the World data.[18] Our baseline uses their population counts from 2000, the latest year available.

Because the cell population counts may be allocated from higher-level data (e.g. subnational population counts) the grid-cell level counts are inappropriate for our purposes. If we use the grid-cell population data, then we could be estimating their algorithm and not the relationship of labour/land and productivity. Therefore we use their data only at the level of districts. We overlay second-level political boundary data from the Global Administrative Areas Project (2019) on top of the GRUMP grid-cell data, and use this to rebuild the population count data for each district.

The estimation in equation (10) requires data on *agricultural* population, and GRUMP provides a measure of *rural* population. There is no perfect overlap of these two sets, but in the absence of any way of measuring the number of agricultural workers, we use the rural data as a proxy. After the main results, we discuss several alternative sources of data (e.g. the International Public-Use Microdata Series) to control for agricultural workers. As part of our controls, we also use data on the urbanization rate within districts as well as their (log) total population. This can be recovered from GRUMP using their counts of total population (rural plus urban) and urban population.

To deal with outliers, we calculate the labour/land ratio for each district. We then discard all observations above the 99th percentile and below the 1st percentile. We also exclude all districts with fewer than 100 total rural residents, again to avoid outliers. Regressions including these observations do not change the results. Summary statistics for the remaining data on the labour/land ratio can be round in Panel B of Table 1. For our entire sample, which covers 28,475 districts for the year 2000, there are 0.73 rural residents per hectare. The percentile distribution of this is shown as well, ranging from only 0.04 per hectare at the 10th percentile to 1.86 at the 90th.

Inherent agricultural productivity We rely on the work of Galor and Özak (2016) to provide our measure of agricultural productivity A_{isg}^{GAEZ}. The authors form a measure of the potential caloric yield at a grid-cell level, combining crop yield information from the GAEZ project with nutritional information on those crops. As argued by Galor and Özak (2016), the caloric suitability index is more informative for analysis of agricultural productivity than raw tonnes of output, as it relates to the nutritional needs of humans. We address the use of calories to compare crops in the robustness section below, and this is not driving our results.

For our purposes, we use the crop-specific data underlying the Galor and Özak (2016) index, restricting ourselves to primary staple crops.[19] Those authors provide details of the construction of this data, but we can provide a summary. For each grid-cell, we calculate the total potential calories that each crop will provide, given the potential production from the GAEZ project (Food and Agriculture Organization 2012) combined with information on calories per tonne for each crop. Within each cell, we then identify the maximum amount of calories that are possible across the different crops. Finally, for a given district, one can sum up those maximum calories to arrive at A_{isg}^{GAEZ}.

After we calculate A_{isg}^{GAEZ} for each district, we discard values above the 99th percentile and below the 1st percentile from that total sample, to avoid outliers. Our results are not sensitive to this trimming. Summary statistics for A_{isg}^{GAEZ} in the remaining districts can be found in Table 1 in the second row of Panel B, reported in millions of calories per hectare. The mean is 10.65 million calories per hectare. At the 10th percentile of the trimmed distribution, the caloric yield is only 4.98 million calories per hectare, while it is four times higher at the 90th percentile, at around 17.03 million calories per hectare. The maximum caloric yield in our sample is 32.64 million calories, while the lowest is only 0.48 million calories.

Crop suitability As a way of creating geographic regions of districts based on crop types, we use 'crop suitability indices', which are also from the GAEZ project (Food and Agriculture Organization 2012) and are provided for each grid-cell on a scale of 0–100. The GAEZ crop suitability indices are used to divide districts based on the types of crops that they produce, but we continue to use our A_{isg}^{GAEZ} to measure actual productivity as the suitability indices are not a measure of potential output.

The GAEZ suitability index depends on climate conditions (precipitation, temperature, evapotranspiration), soil (acidity, nutrient availability) and terrain (slope). For districts of a country, we construct an overall suitability index as a weighted (by area) sum of the grid-cell suitability indices. Given that the grid-cell suitability measures run from 0 to 100, our aggregated index for each district also runs from 0 to 100.

Land area Our measure of land area X_{isg} is the total land area of a district, without adjusting for cultivated area. We will thus be estimating the elasticity of output with respect to the *possible* stock of land. Choosing to not crop certain plots is akin to choosing to apply zero labour or capital to those plots. We discuss after the main results that our estimates do not differ if we use information on cultivated area in place of total land.

Night-time lights We follow Henderson *et al*. (2016) and use the Global Radiance Calibrated Nighttime Lights data provided by NOAA/NGDC, described in Elvidge *et al*. (1999), and reported at 30 arc-second degree resolution. This dataset contains more detail on low levels of light emissions (thus capturing detail for undeveloped areas), and avoids most top-coding of areas saturated by light (thus capturing more detail in developed areas). To match the data that we use on population, we use the dataset from 2000, and create district-level measures of night-time light density by averaging across the pixels contained within each district.

We adjust for the fact that the lights data are reported with zero values, which is part of an adjustment from NOAA/NGDC to account for possible noise in pixels that report very small amounts of light. Similar to Henderson *et al*. (2016), for any district that has a raw value of zero for night lights, we replace that with the minimum positive value found in the rest of the sample of districts. This prevents us from understating light density in those districts. Once this adjustment is made, we take logs of the average lights in a district. Summary statistics for the final night lights data can be found in Panel B of Table 1.

Road density and distance to cities Meijer *et al.* (2018) provide the Global Roads Inventory Project (GRIP) dataset, which is a gridded map of the world containing information on road density (km per square km), broken down by type of road (highway, primary roads, secondary roads, tertiary roads, local roads). We use their data and aggregate to the district level to find (log) road density for each district, along with the percentage of all roads in a district that are coded as being highways, primary roads or secondary roads. In addition to road density, we calculate the (great circle) distance in kilometres from the centroid of each district to the closest city of 100,000 or more residents. For districts that contain such a city, the distance is zero. We do not restrict ourselves to searching for the closest city within a given state; the closest city is allowed to be in a neighbouring state. Summary statistics for all road and distance variables can be found in Table 1.

Defining temperate and tropical regions

Our primary distinction of a region g is as either temperate or tropical. There is no definitive way of assigning districts to either temperate or tropical regions, so we pursue several possibilities. Regardless of the assignment rule, it is worth reiterating that it is applied at the district level, and countries (and states) are not assumed to be homogeneous.

- *By crop suitability* The first way of denoting temperate and tropical is through the types of crops capable of being grown, as this depends on the overall agroclimatic characteristics. Here we define *temperate* districts as those that have any grid-cells suitable for barley, buckwheat, rye, oats, wheat or white potatoes, but have precisely *zero* grid-cells suitable for any of cassava, cowpeas, paddy rice, pearl millet, sweet potato and yams. Suitability for any crop is taken from the GAEZ project. The *tropical* districts are those that have any grid-cells suitable for cassava, cowpeas, paddy rice, pearl millet, sweet potato or yams, but precisely *zero* grid-cells suitable for barley, buckwheat, rye, oats, wheat and white potatoes.[20] In total, we have 8774 districts classified as temperate using crop suitability, and 7018 classified as tropical. There are 12,388 districts that are suitable for *both* types of crops, meaning that they contain both temperate and tropical grid-cells, or they have grid-cells reported by the GAEZ project as suitable for both types of crops. These mixed districts are excluded from our baseline analysis, but we return to them later in the paper.
- *By frost-free days* Rather than crop suitability, which combines several climate characteristics, we can narrow the assignment down to a single characteristic, frost-free days. Frost plays a role in agriculture through culling various micro-organisms related to plant disease and the mineralization of organic matter (Masters and McMillan 2001), and its presence or absence can be a useful indicator. We define *temperate* districts as those that have fewer than 365 frost-free days, meaning that they experience at least one frost day during the year, on average. We define *tropical* districts as those with 365 frost-free days, meaning they do not experience any frost, on average. This gives us 14,242 temperate districts, and 14,233 tropical districts, for total coverage of our sample.[21] Data on frost-free days are from the GAEZ project.
- *By Köppen–Geiger climate zones* A final classification is to use direct climate characteristics. We use the Köppen–Geiger scheme to assign 9956 districts as *temperate* and another 9731 as *tropical*.[22] This broad classification also does not result in exclusive assignment, and there are 446 districts that qualify as *both* temperate and

tropical, as their land area is split across both definitions. Excluding or including those districts with an overlap has no effect on our results.

Our results are not contingent on the choice of definition for temperate/tropical, as will be shown below. For much of the paper we will focus on the first definition, based on crop suitability. Figure 1 shows how grid-cells across the world are coded as temperate, tropical, suitable for both types of crops, or unsuitable for either type (e.g. deserts or polar regions).[23] While the temperate area covers much of North America and Eurasia, as expected, there are several pockets of temperate areas around the world. Central Mexico, the spine of South America, the Tigris/Euphrates watershed, an area roughly corresponding to Manchuria, and a few pockets in East Africa all fall in our temperate region. We will use districts from these areas to show that our estimated land elasticities are robust to excluding the developed countries in Europe and North America from the estimation. The tropical area runs in a zone around the equator, as one would expect, and areas suitable for both types of crops tend to exist in between temperate and tropical areas. Given the broad range of these crop classifications and the small sizes of districts, nearly all districts are found to be homogeneous with respect to their classification as temperate or tropical.

Turning to district-level characteristics, Figure 2 shows the density plots of (log) rural labour/land for the two regions. One can see that rural labour/land tends to be higher in tropical districts, with a peak between 0.33 rural residents per hectare (i.e. log value -1) and 1.0 rural residents per hectare (i.e. log value 0). In comparison, while there are a few districts in the temperate group with densities as high as 1.0 rural residents per hectare, the peak is around 0.33 rural residents per hectare (i.e. log value -1).

There is a similar distinction in the density plots of caloric yield A_{isg}^{GAEZ} for districts in the tropical and temperate groups. Figure 3 shows these plots, and the tropical districts have a strong peak at around 12–15 million calories per hectare, while the peak for temperate districts is closer to 5 million calories, although the tail of the temperate distribution runs as high as for tropical districts. This reflects both inherent agroclimatic productivity differences and the fact that the calories per tonne of the crops defining the tropical districts (e.g. cassava, wet rice) are much higher than the calories per tonne defining temperate districts (e.g. barley, wheat). We discuss below that the calories per tonne values for each crop cannot explain our results.

The two plots in Figure 3 capture the raw information about rural labour/land and calories per hectare, but note that the distinctions in medians and modes between temperate and tropical districts are immaterial to our estimation. We will be using the district-level variation in rural labour/land and caloric yield only *within* states, and only for districts that share a common definition of temperate or tropical. Hence the differences in the distributions seen in Figures 2 and 3 are not driving our results.

Estimates for temperate and tropical regions

Table 3 shows the baseline estimates of β_g for our temperate and tropical regions. In column (1) of Panel A, one can see that the estimate of β_g for temperate districts is 0.285, while in column (2) the estimate of β_g for tropical districts is 0.126, a difference of approximately 0.16. Below these estimates are two hypothesis tests. The first row tests the hypothesis that the true β_g is equal to zero, and in both samples we reject this at below 0.1% significance. The second row tests the hypothesis that the β_g from the tropical

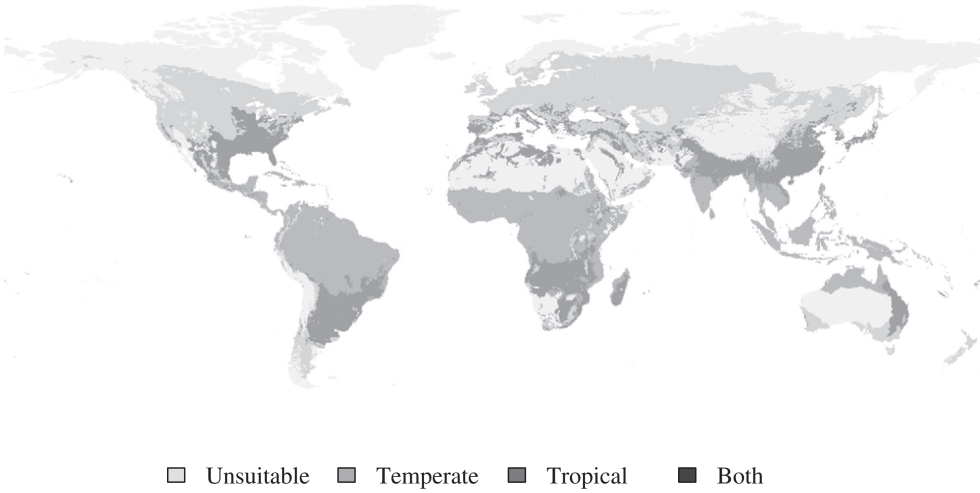

FIGURE 1. Map of geographic regions.
Notes: The figure shows the classification of each pixel into geographic regions, along with the baseline estimated land elasticity. 'Temperate' pixels are those that are capable of growing the temperate crops but not tropical crops (see text). 'Tropical' pixels can grow tropical crops, but not temperate crops. 'Both' pixels are capable of growing both temperate and tropical crops. 'Unsuitable' pixels can grow neither kind of crops. Crop suitability is assessed using the GAEZ project (Food and Agriculture Organization 2012). These pixels are what the assignment of districts to 'Temperate' or 'Tropical' regions is based on.

region is equal to the β_g from the temperate region. We can reject that null hypothesis at 0.1%.

Figure 4 plots the residual relationship of log caloric yield and log rural labour/land found from columns (1) and (2) of Table 3, controlling for state fixed effects, log light density, log population, the urban percentage in a district, road density and the share of roads of different types, distance from a major city (more than 100,000), and a log slope index. Given the large number of observations, we plot the average values of the residuals for 50 different quantiles of our data to make the figure legible. As these are residuals, the values of rural labour/land and caloric yield are all centred around zero.[24] The difference in the slopes of the lines for tropical and temperate districts implies a difference in the values of the land elasticity β_g, and as Table 3 indicates, that difference is statistically significant. The additional value of Figure 4 is that it allows us to assess our linearity assumption and judge if there are outliers perhaps driving the results. Overall, the linearity assumption appears solid. Excluding the observations at the very highest or very lowest labour/land ratios does not alter our main results.

Returning to Table 3, the remainder of Panel A shows variations on our baseline result using different definitions of temperate and tropical districts. In columns (3) and (4), we use the definition of temperate and tropical based on the number of frost-free days. The results are similar to our baseline, with an estimated β_g of 0.262 for temperate districts, but only 0.130 for tropical ones. The gap here is about the same as our baseline results from columns (1) and (2), and is significant at 0.4%. Columns (5) and (6) use the Köppen–Geiger definition of temperate and tropical regions. Here the results are similar to those using the crop suitability definition: β_g is estimated to be 0.272 in temperate districts, and only 0.117 in tropical ones, for a difference of about 0.16 that is again statistically significant at 0.1%. Our results are not sensitive to the exact definition of temperate/tropical.

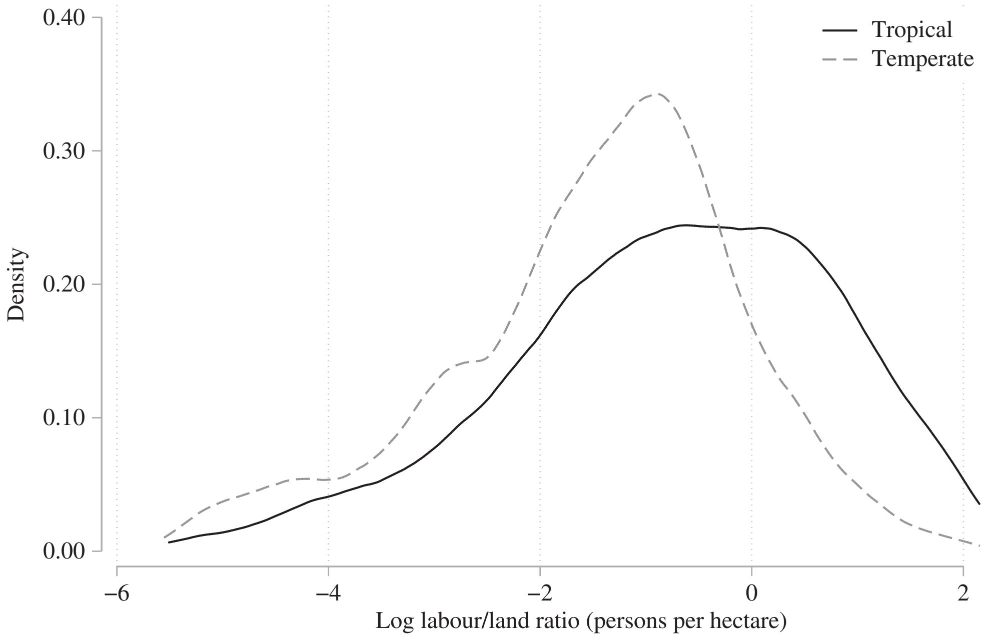

FIGURE 2. Density plot of log rural labour/land ratios (L_{Ais}/X_{is}), by crop type, 2000.
Notes: Kernel density plot, Epanechnikov kernel, of the (log) rural labour/land ratio L_{Aisc}/X_{isc} at the district level, calculated by the authors using data from Center for International Earth Science Information Network (CIESIN) for rural population. 'Temperate' includes districts that are suitable for growing barley, buckwheat, oats, rye, wheat and white potatoes, but have zero suitability for cassava, cowpeas, pearl millet, sweet potato, wet rice and yams. 'Tropical' includes districts suitable for the latter set of crops, but zero suitability for the former.

Panel B of Table 3 provides an initial set of robustness checks on the results. In all regressions in Panel B, the definition of temperate versus tropical region is based on crop suitability, as in the first two columns of Panel A. In Panel B, columns (1) and (2) exclude any district with a reported urban population greater than 50,000 people. The worry is that highly urbanized districts may operate a different type of agricultural technology and/or may skew the labour/land ratio near them (perhaps due to definitions of urban areas), and that our original results were affected by this. As can be seen from Table 3, however, the distinction in β_g grows to 0.300 for temperate districts while remaining at 0.126 for tropical districts, which is an absolute difference of almost 0.18. This difference is again significant.

Columns (3) and (4) of Panel B of Table 3 exclude districts that have a total population greater than 5% of their state total, which again eliminates large urban areas but also eliminates any districts that may happen to be relatively large with respect to their state. The results conform to those in columns (1) and (2), with a temperate estimate of β_g equal to 0.296, and a tropical estimate of 0.124, a difference that is statistically significant at less than 0.1%.

Finally, columns (5) and (6) of Panel B of Table 3 exclude both Europe (including Russia west of the Urals) and North America from the samples to address the worry that these areas may use types of agricultural technologies different to those in other places at lower development levels.[25] The earlier findings still hold, with an estimated tropical β_g of 0.124 compared to 0.298 for temperate districts. The difference is significant at less than 0.1%.

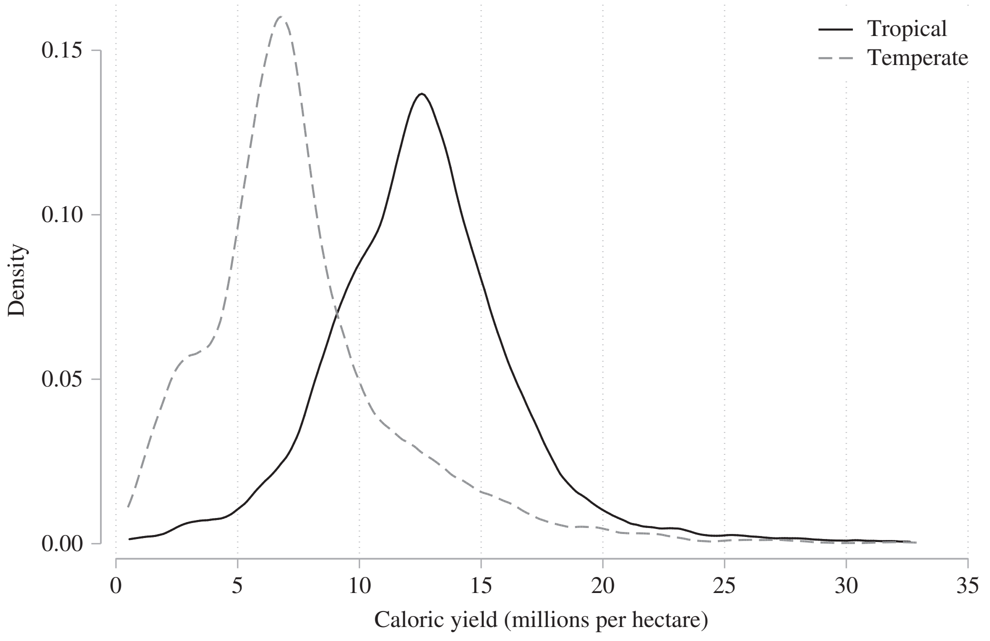

FIGURE 3. Density plot of caloric yield (A_{is}^{GAEZ}), by crop type.
Notes: Kernel density plot, Epanechnikov kernel, of the caloric yield A_{isc} at the district level, calculated by the authors using data from Galor and Özak (2016); see text for details. This measure sums the maximum calories available per grid-cell within a district, then divides by total area of the district. 'Temperate' includes districts that are suitable for growing barley, buckwheat, oats, rye, wheat and white potatoes, but have zero suitability for cassava, cowpeas, pearl millet, sweet potato, wet rice and yams. 'Tropical' includes districts suitable for the latter set of crops, but zero suitability for the former.

Robustness checks

Rural labour/land data Panel A of Table 4 shows results using different sources for the rural population data L_{Ai}. In columns (1) and (2), we re-estimate the values of β_g for temperate and tropical regions using population data from GRUMP, but from 1990. The results of 0.288 for temperate and 0.126 for tropical are almost identical results to the 2000 data. In columns (3) and (4), we show that our results are not driven by using GRUMP as the data source. We use the HYDE 3.1 database (Goldewijk *et al*. 2011) for 2000. Again, the results conform to our baseline, with 0.241 for temperate areas and 0.117 for tropical areas.

In columns (5) and (6) of Panel A of Table 4, we turn to the International Public-Use Microdata Series (IPUMS) database (Minnesota Population Center 2017) to extract individual level data for 39 countries that have geographic identifiers at the subnational level. Using this we can accomplish two things. We can find direct information on the number of people living within a given geographic area as opposed to relying on GRUMP. Because of the limited country coverage of the IPUMS, and because the 'districts' used by the IPUMS are larger than our baseline, we end up with only 3519 observations.[26] Nevertheless, in columns (5) and (6) the results are consistent with our baseline, although shifted down in both cases. The temperate elasticity is estimated to be 0.190, while the tropical elasticity is only 0.017. The gap using the IPUMS is similar in size to our baseline.

TABLE 3
ESTIMATES OF LAND ELASTICITY β_g BY AGRICULTURAL TYPE, 2000

	Temperate (1)	Tropical (2)	Temperate (3)	Tropical (4)	Temperate (5)	Tropical (6)
Panel A: Regions defined by:						
	Crop suitability		Frost days		Köppen–Geiger	
Log labour/land ratio (β_g)	0.285	0.126	0.262	0.130	0.272	0.117
	(0.043)	(0.023)	(0.044)	(0.014)	(0.044)	(0.019)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.000	0.000
p-value $\beta_g = \beta_{\text{Temp}}$		0.001		0.004		0.001
Countries	72	67	88	96	81	72
Observations	8416	6731	13,811	13,879	9287	9457
R-squared (no FEs)	0.19	0.15	0.16	0.13	0.18	0.14
Panel B: With other restrictions (using crop suitability to define temperate/tropical):						
	Urban pop. < 50K		Pop. share < 0.05		Excl. Europe/N. America	
Log labour/land ratio (β_g)	0.300	0.126	0.296	0.124	0.298	0.124
	(0.045)	(0.024)	(0.048)	(0.025)	(0.045)	(0.023)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.000	0.000
p-value $\beta_g = \beta_{\text{Temp}}$		0.001		0.001		0.001
Countries	68	66	56	36	17	62
Observations	7529	6192	6429	4071	813	6676
R-squared (no FEs)	0.20	0.16	0.20	0.17	0.15	0.11

Notes
Dependent variable in both panels: log caloric yield (A_{isg}^{GAEZ}).
Conley standard errors, adjusted for spatial autocorrelation with a cut-off distance of 500 km, are shown in parentheses. All regressions include state fixed effects, a constant, and controls for the district urbanization rate, log density of district night-time lights, log total population, log road density, share of roads of different types, distance to nearest city of 100,000 people, and a log slope index. The coefficient estimate on rural population labour/land indicates the value of β_g; see equation (10). Rural population is from the GRUMP database (CIESIN *et al.* 2011), and caloric yield is the authors' calculations based on data from Galor and Özak (2016). Inclusion of districts in the regression is based on the listed criteria, either crop suitability, the number of frost-free days, or Köppen–Geiger climate zones. See text for details of how temperate and tropical regions are defined in each case. In Panel B, the columns include districts with fewer than 50,000 urban residents, include districts that contain less than 5% of state population, or exclude districts from any country in Europe (including Russia west of the Urals) or North America.

The second use for the IPUMS is for information on occupation and/or industry. This allows us to distinguish agricultural workers from rural residents. The measure of L_{Ai} in columns (5) and (6) of Panel A of Table 4 is based on those who report agriculture as their industry of employment. An additional reassurance for our baseline results is that the IPUMS shows that the correlation of rural residents with the number of agricultural workers is 0.91 and significant at less than 1%. Our baseline GRUMP data on rural residents are not making systematic errors in measuring agricultural worker labour/land ratios.

Land area Our baseline results measure land X_i in a district as the total area, as this represents the stock of *possible* agricultural land. However, we can restrict ourselves to looking at the labour/land ratio of agricultural workers on actual cultivated land. We use

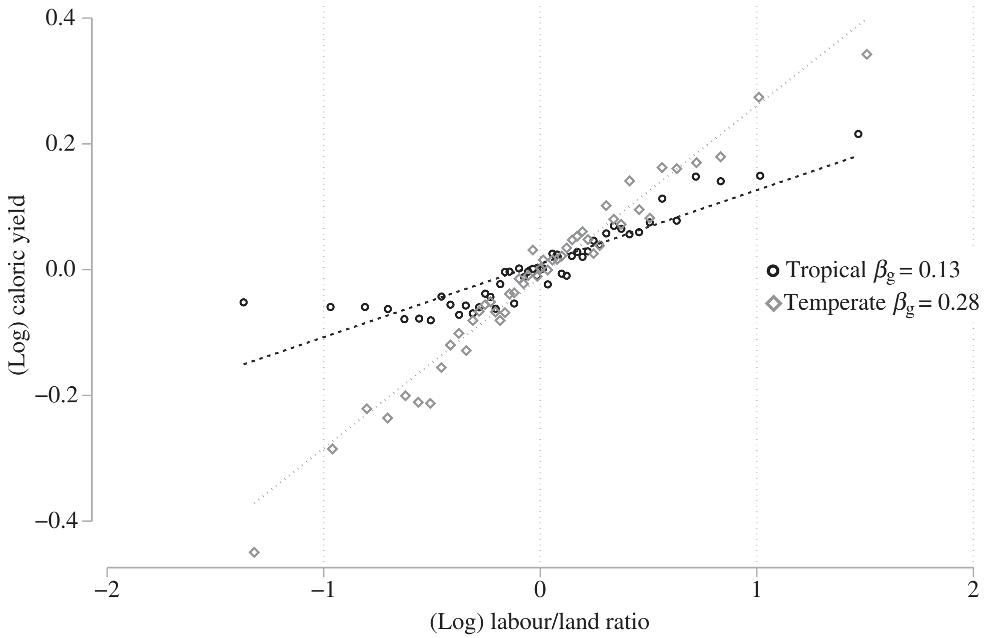

FIGURE 4. Residual relationship of caloric yield (A_{is}^{GAEZ}) and rural labour/land ratios.
Notes: Plotted are the quantile averages of both log caloric yield and log rural labour/land ratio for each sample, temperate and tropical. 50 quantiles are used in each sample. The quantiles are taken from the residuals of caloric yield and rural labour/land ratio after controlling for log light density, urban percentage in 2000, log total population, distance to a city of 100,000 people, road density (km per square km), percentage of roads as highways, primary and secondary roads, the (log) slope and state fixed effects. Linear fits are shown, and the estimated slopes are in the legend. The binscatter command from Stata was used to prepare the figure. 'Temperate' includes districts that are suitable for growing barley, buckwheat, oats, rye, wheat and white potatoes, but have zero suitability for cassava, cowpeas, pearl millet, sweet potato, wet rice and yams. 'Tropical' includes districts suitable for the latter set of crops, but zero suitability for the former.

the GAEZ project to build a measure of the area of cultivated land in a given district as X_i^C. Our baseline rural labour/land ratio can thus be written as $\ln(L_{Ai}/X_i) = \ln(L_{Ai}/X_i^C) + \ln(X_i^C/X_i)$. The first term on the right is the (log) ratio of agricultural workers per cultivated land, while the second term is the (log) share of cultivated land in total land area. We can include both of the right-hand-side terms as controls in our regressions, and recover the estimate of β_g from the coefficient on $\ln(L_{Ai}/X_i^C)$, labour/land measured per unit of cultivated land. In columns (1) and (2) of Panel B of Table 4, we present results using cultivated land to measure rural labour/land ratios. Again, the results are consistent with our baseline (0.278 for temperate areas and 0.126 for tropical areas).

Cash crops and livestock Our measure of A_{isg}^{GAEZ} is based on staple crops, as opposed to cash crops (e.g. cocoa) or livestock production. A particular problem would be if some districts within a state focus on cash crops or livestock, while other districts focus on staple crops. The differences in labour/land ratios between these districts may not be related to our measure of staple crop productivity A_{isg}^{GAEZ}, and thus our estimate of β_g could be biased.

TABLE 4
ESTIMATES OF LAND ELASTICITY β_g, ADDITIONAL ROBUSTNESS CHECKS

	Temperate (1)	Tropical (2)	Temperate (3)	Tropical (4)	Temperate (5)	Tropical (6)
Panel A: Different rural population sources:	GRUMP 1990		HYDE 2000		IPUMS (various)	
Log labour/land ratio (β_g)	0.288	0.126	0.241	0.117	0.190	0.017
	(0.042)	(0.023)	(0.014)	(0.018)	(0.087)	(0.017)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.029	0.326
p-value $\beta_g = \beta_{\text{Temp}}$		0.001		0.000		0.040
Countries	72	67	72	68	22	24
Observations	8416	6731	8170	6465	1103	2416
R-squared (no FEs)	0.19	0.15	0.20	0.17	0.08	0.07
Panel B: Different land assumptions (with GRUMP labour/land ratio):	Cultivated area		Cash crops < 10% area		Pasture < 50% area	
Log labour/land ratio (β_g)	0.278	0.126	0.257	0.108	0.286	0.137
	(0.044)	(0.024)	(0.050)	(0.032)	(0.044)	(0.025)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.001	0.000	0.000
p-value $\beta_g = \beta_{\text{Temp}}$		0.002		0.011		0.003
Countries	72	66	55	37	70	64
Observations	8382	6694	5679	2356	7582	5692
R-squared (no FEs)	0.18	0.14	0.18	0.15	0.19	0.16

Notes
Dependent variable in both panels: log caloric yield (A_{isg}^{GAEZ}).
Temperate and tropical samples are defined by the suitability measures described in Table 3. Conley standard errors, adjusted for spatial autocorrelation with a cut-off distance of 500 km, are shown in parentheses. All regressions include state fixed effects, a constant, and controls for the district urbanization rate, log density of district night-time lights, log total population, log road density, share of roads of different types, distance to nearest city of 100,000 people, and a log slope index. The coefficient estimate on rural population labour/land indicates the value of β_g; see equation (10). Caloric yield is the authors' calculations based on data from Galor and Özak (2016). In Panel A, the population data used to define rural labour/land differ based on the heading in the table (see text for details). In Panel B, the first set of results uses rural population (from GRUMP) relative to cultivated land area (as opposed to actual land area) to measure labour/land ratios. The second set drops districts that have less than 10% of their area in cash crops (see text for a list of those crops), and the third set shows results for districts that have less than 50% of their area as pasture land.

To address this, we draw in additional data on land use to eliminate districts that are heavy cash crop or livestock producers. In columns (3) and (4) of Panel B of Table 4, we drop any district that has more than 10% of its harvested area coming from cash crops. Data on the harvested area are from Monfreda *et al.* (2008).[27] The estimated β_g in temperate areas, 0.257, remains larger than the estimate for tropical areas, 0.108, and that difference remains significant. Columns (5) and (6) drop any districts that have more than 50% of their area devoted to pasture, using data from Ramankutty *et al.* (2008). Again, the temperate and tropical estimates are around our prior estimates, 0.286 and 0.137, respectively.

Productivity data Another concern with the existing results is that they are reliant on the specific caloric suitability index A_{isg}^{GAEZ} that we derived. In particular, we used the underlying data from the GAEZ project for 'low-input, rain-fed' agriculture to construct this index, matching Galor and Özak (2016). This could overstate the variation in 'true' productivity (A_{isg}) across districts within states, because it ignores the possibility that inherently low-productivity districts can adopt the use of fertilizer and/or irrigation to raise their productivity. If A_{isg}^{GAEZ} overstates the variation in productivity across districts, then we may be overstating the size of β_g. If, for some reason, this problem is pronounced in temperate areas, then this could explain our finding that temperate areas have high β_g values. Alternatively, A_{isg}^{GAEZ} may understate variation in A_{isg} if irrigation or modern inputs allow some districts to increase their total factor productivity relative to others. If this is true in tropical regions, then we would be underestimating β_g for tropical areas.

To address these concerns, in Table 5, we show results where we reconstruct the index A_{isg}^{GAEZ} using different underlying data on productivity from the GAEZ project. In columns (1) and (2), for example, we use their 'medium-input, irrigated' estimates of productivity to derive A_{isg}^{GAEZ}, and then re-run our regressions. As can be seen, temperate and tropical β_g estimates fall slightly relative to our baseline (0.254 for temperate and 0.120 for tropical). But the gap remains 0.13, and is significant at 1.4%.

In columns (3) and (4) of Table 5, we do a similar exercise, but now use the 'high-input, rain-fed' productivity data from the GAEZ project to construct A_{isg}^{GAEZ}. Here the results are nearly identical to our baseline (0.286 for temperate and 0.132 for tropical). Columns (5) and (6) use the 'high-input, irrigated' productivity data to construct A_{isg}^{GAEZ}, and the results are similar to when we use the irrigated productivity measures from columns (1) and (2). The estimated effects (0.253 for temperate and 0.120 for tropical) are again slightly smaller, but remain significantly different at 1.4%.

Another potential issue with the construction of A_{isg}^{GAEZ}, regardless of the choice of inputs and water use, is that it relies on the calorie content of different crops to make them comparable to one another. It could be that the calorie counts used by Galor and Özak (2016) that we adopt are incorrect. Or perhaps calories are an imperfect way of comparing crops, and we should be using something like relative prices. We address this by using the individual crop-level measures of raw productivity (in tonnes, rather than calories) from the GAEZ project as our measure of A_{isg}^{GAEZ}. For temperate regions, for example, we run separate regressions using the raw potential barley yield as our measure of A_{isg}^{GAEZ}, and then do so for buckwheat, then oats, and so on. We do similar regressions for tropical areas with raw yields of the tropical crops. The full results are available in the Online Appendix.

In all cases, the estimated size of β_g using the individual crop raw potential yields gives us nearly identical results to what we find in our baseline using the caloric suitability index. The consistency of the results using separate crop-specific raw potential yields shows that weighting crop yield by calorie counts to aggregate them together is not important to our results. Further, this consistency across crops also implies that *any* weighting scheme to compare the value of crops (e.g. prices) would also yield similar results for β_g as our baseline.

Demographic and asset controls

The state fixed effects and controls for night lights, urban share, total population and road density may not control fully for district-level variation in the capital/labour ratio or transport costs, in particular due to differences in the characteristics of the population in a district (e.g. education) and the availability of capital (e.g. livestock or the presence of electrical service). To assess if this is biasing our results, we run the same regressions for a limited subsample of districts for which we can assemble detailed data on demographics and assets.

TABLE 5
ESTIMATES OF LAND ELASTICITY β_g, ALTERNATIVE PRODUCTIVITY MEASURES

	Temperate (1)	Tropical (2)	Temperate (3)	Tropical (4)	Temperate (5)	Tropical (6)
Caloric yield based on GAEZ input/water use:	Medium/irrigated		High/rain-fed		High/irrigated	
Log labour/land ratio (β_g)	0.254	0.120	0.286	0.132	0.253	0.120
	(0.050)	(0.022)	(0.045)	(0.025)	(0.050)	(0.022)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.000	0.000
p-value $\beta_g = \beta_{\text{Temp}}$		0.014		0.003		0.014
Countries	72	67	72	67	72	67
Observations	8416	6731	8389	6719	8416	6731
R-squared (no FEs)	0.17	0.14	0.18	0.14	0.17	0.14

Notes
Dependent variable: log caloric yield (A_{isg}^{GAEZ}).
Temperate and tropical samples are defined by the suitability measures described in Table 3. Conley standard errors, adjusted for spatial autocorrelation with a cut-off distance of 500 km, are shown in parentheses. All regressions include state fixed effects, a constant, and controls for the district urbanization rate, log density of district night-time lights, log total population, log road density, share of roads of different types, distance to nearest city of 100,000 people, and a log slope index. The coefficient estimate on rural population labour/land indicates the value of β_g; see equation (10). In Panel A, the construction of the A_{isg}^{GAEZ} caloric suitability yield differs across the columns. In columns (1) and (2), the yield is derived from the underlying GAEZ medium input irrigated data, and the following columns use the high input rain-fed data, or the high input irrigated data, as indicated.

We use the DHS, which provide individual and household level data in a consistent manner across a wide range of developing countries. Many of these surveys contain geographic information systems (GIS) information on the latitude and longitude of the surveyed clusters (e.g. a village), which allows us to identify which clusters are located within which districts. For those surveys with GIS data, we create district-level aggregate demographic and asset measures.[28] With the DHS data, this gives us a sample of 1581 districts, of which 290 are part of our temperate region, and 1291 are part of our tropical region. Details on the countries from which these districts are drawn are available in the Online Appendix.

Table 6 shows the results of estimating β_g for temperate and tropical regions. Columns (1) and (2) are limited to those districts that have DHS data, but these data are not included as controls in these regressions. The results here, with a land elasticity of 0.375 for temperate districts and 0.104 for tropical districts, have a larger spread than in our baseline, but the pattern is consistent. In columns (3) and (4) we repeat these regressions but now include the DHS demographic data. The results are nearly identical, with a small drop in the tropical estimate to 0.102. Columns (5) and (6) include the DHS demographic and asset data, and again the results are nearly identical. The consistency as we add DHS controls provides some reassurance that the main findings are not due to unobserved district-level variation in the composition of the labour force or availability of capital.

Production function specification

For expositional purposes, we used a Cobb–Douglas production function in equation (1), which implies that the land elasticity does not vary with the endowment of land, labour or other inputs. This need not be the case, of course, and the different results on β_g for temperate and tropical areas may reflect not a fundamental difference in the production function, but rather a difference in those endowments. In particular, we know from Figure 2 that tropical areas have higher rural densities than temperate areas. If the elasticity of substitution between land and labour were more than 1, then higher rural labour/land ratios would be associated with a lower land elasticity (and a higher labour elasticity).[29]

We do not believe that this can explain our results. While tropical areas on average have a higher labour/land ratio, there are numerous examples of low-density tropical areas (parts of Central and South America, areas in Sub-Saharan Africa). In the Online Appendix we estimate a separate value of β_g for each state in the temperate and tropical regions containing 10 or more districts. We can then plot the values of β_g against the labour/land ratio and there is no systematic relationship. We can also drop districts that have very high (or very low) absolute labour/land ratios, and find similar results.

Comparison to factor shares

A possible point of comparison for our estimates of β_g is the factor share of land in agricultural output. With competitive markets for *all* inputs to agriculture, the factor share of land should be equal to the elasticity β_g. There is variation in these factor shares across countries, but they are not always consistent with our estimates. Fuglie (2010) reports factor share estimates for a set of countries, finding shares between 0.22 and 0.33 for land and structures. The inclusion of structures muddies the comparison with our estimate of β_g. Nevertheless, he reports land shares between 0.22 and 0.25 for India,

TABLE 6
ESTIMATES OF LAND ELASTICITY β_g, WITH DHS DISTRICT CONTROLS

	Temperate (1)	Tropical (2)	Temperate (3)	Tropical (4)	Temperate (5)	Tropical (6)
Log labour/land ratio (β_g)	0.375 (0.081)	0.104 (0.020)	0.375 (0.080)	0.102 (0.020)	0.374 (0.083)	0.110 (0.021)
Demographic controls	No	No	Yes	Yes	Yes	Yes
Asset controls	No	No	No	No	Yes	Yes
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.000	0.000
p-value $\beta_g = \beta_{\text{Temp}}$		0.000		0.000		0.000
Countries	15	29	15	29	15	29
Observations	290	1291	290	1291	290	1291
R-squared	0.82	0.82	0.82	0.82	0.83	0.82

Notes
Dependent variable: log caloric yield (A_{isg}^{GAEZ}).
Temperate and tropical samples are defined by the suitability measures described in Table 3. Conley standard errors, adjusted for spatial autocorrelation with a cut-off distance of 500 km, are shown in parentheses. All regressions include state fixed effects, a constant, and controls for the district urbanization rate, log density of district night-time lights, log total population, log road density, share of roads of different types, distance to nearest city of 100,000 people, and a log slope index. The coefficient estimate on rural population labour/land indicates the value of β_g; see equation (10). The districts included in these regressions have villages/clusters that took part in the DHS. Using DHS data, the columns include district-level means or medians of demographic variables (e.g. household head education and age) and asset variables (e.g. household ownership of cattle or use of electricity); see text for details of the precise controls.

Brazil and Indonesia. There is substantial heterogeneity within each of these countries in climate and crop type, but our estimates would suggest values of β_g around 0.12–0.13, based on the prevalence of tropical agriculture. The factor share of land and structures for China is 0.22, which is difficult to compare to our results given the heterogeneity in climate zones within China.

Reported factor shares for land and structures in the USA (0.19) and former Soviet Union (0.21–0.26) are lower than our β_g estimates for areas using temperate agriculture, although both of those countries also contain heterogeneity in climate zones. A study by Jorgenson and Gollop (1992) reports a land share of 0.21 for the USA, which is under our temperate estimate. Fuglie (2010) reports a factor share of 0.17 for land and structures in the UK, which is also lower. However, Clark (2002) reports long-run factor shares of land for England, and that share is between 0.30 and 0.36 for several centuries, somewhat higher than our estimated β_g for temperate areas. Hayami *et al*. (1979) provide longer-run estimates of land shares for several east Asian economies, finding estimates between 0.3 and 0.5 for Taiwan, Japan, Korea and the Philippines from the late 1800s until the middle of the 20th century. These numbers cannot be directly compared to our β_g estimates, as much of Japan and Korea, and all of Taiwan, are excluded from our analysis because they are suitable for *both* temperate and tropical crops, as we have defined them.

Comparing to land shares thus provides mixed results. Nevertheless, we think that there is information in our estimates. Our estimates are built using the assumption of mobility of labour between districts, but are robust to arbitrary distortions to wages between agriculture and non-agriculture, or arbitrary distortions in the relative price of agriculture (which could include market power in either sector). In contrast, for factor shares to be good estimates of the elasticities, it would have to be that returns are equalized across districts *and* there are no distortions or frictions in the state-wide factor markets, so that factor shares are in fact identical to elasticities. There is no obvious reason to think that those assumptions about perfect factor markets conditions hold. Furthermore, the factor share data are an aggregation from a snapshot of *farm*-level payments to land, but as noted before, the farm-level production function may not be equivalent to the aggregate production function that we are estimating. It is not clear that the factor share data cited should be privileged in terms of their relevance for the question at hand.

Districts suitable for both kinds of agriculture

To illustrate the difference in land elasticities, we have focused on temperate and tropical regions, and those were defined in a stark way. Our definition left out a substantial number of districts that have some suitability for *both*. There are 12,388 districts that the GAEZ project reports as being capable of growing at least some temperate and tropical crops.

For comparison, we estimated the value of β_g for this mixed group. The results are in Table 7. Column (1) shows the result with the same controls as we used in our baseline regressions. The estimated elasticity is 0.167, which falls between the 0.126 estimate for tropical districts and the 0.285 for temperate areas. We feel that this provides some assurance that our estimates are picking up a realistic difference in the land elasticity between temperate and tropical districts. The difference between the mixed elasticity and the temperate elasticity is significant at 0.1%, while the difference with the tropical elasticity is significant at only 15.5%.

TABLE 7
ESTIMATES OF LAND ELASTICITY β_g FOR MIXED REGION, 2000

Specification defined by:	Baseline (1)	Urban pop. < 50K (2)	Excl. Europe/ N. America (3)	Land X_{is} cult. area (4)	Cash crops < 10% area (5)	GAEZ high input (6)
Log labour/ land ratio (β_g)	0.167	0.175	0.121	0.164	0.183	0.176
	(0.022)	(0.024)	(0.016)	(0.023)	(0.030)	(0.024)
p-value $\beta_g = 0$	0.000	0.000	0.000	0.000	0.000	0.000
p-value $\beta_g = \beta_g^{\text{Temp}}$	0.001	0.001	0.005	0.002	0.111	0.005
p-value $\beta_g = \beta_g^{\text{Trop}}$	0.155	0.111	0.975	0.183	0.104	0.143
Countries	106	99	55	105	71	105
Observations	12,388	10,800	6301	12,342	6010	12,321
R-squared (no FEs)	0.13	0.13	0.09	0.12	0.14	0.11

Notes
Dependent variable: log caloric yield (A_{isg}^{GAEZ}).
For all regressions, the sample includes districts that are suitable for both temperate and tropical crops, as defined in the text. Conley standard errors, adjusted for spatial autocorrelation with a cut-off distance of 500 km, are shown in parentheses. All regressions include state fixed effects, a constant, and controls for the district urbanization rate, log density of district night-time lights, log total population, log road density, share of roads of different types, distance to nearest city of 50,000 people, and a log slope index. The coefficient estimate on rural population labour/land indicates the value of β_g; see equation (10). Rural population is from the GRUMP database (CIESIN *et al.* 2011), and caloric yield is the authors' calculations based on the data from Galor and Özak (2016). Inclusion of districts in the regression is based on the listed criteria. Column (4) uses cultivated land (rather than total land) to measure the labour/land ratio. Column (5) excludes districts that have more than 10% of their total land used for cash crops. Column (6) measures the caloric yield with the GAEZ high input measure of agricultural potential (as opposed to the low input baseline).

The remaining columns of Table 7 show the estimated β_g under different robustness checks. Excluding urban populations of less than 50,000 people raises the estimate slightly to 0.175, while excluding Europe and North America lowers it to 0.121, close to the baseline tropical average. Using cultivated land in place of total land results in a value of 0.164, while excluding any district with more than 10% of its area under cash crops gives an estimate of 0.183. If we use the GAEZ high-input productivity numbers to measure A_{isg}^{GAEZ}, then we find a value of 0.176 as well. There appears to be a consistent case that the land elasticity for these mixed areas lies in between the tropical and temperate values, and centres around 0.17.

Looking across the columns of Table 7, one can see that in each case there is a significant difference between the mixed elasticity and the temperate elasticity, with the only possible exception being in column (5) where the p-value reaches 11%. On the other hand, the difference between the mixed elasticity and the tropical elasticity is not always statistically significant, with p-values ranging up to 97.5%. In general, mixed districts are similar to tropical districts, while temperate districts appear to be an outlier relative to these two regions.

Aggregate land elasticities

We can combine our full set of estimates to demonstrate how the land elasticity varies across the world at the country level. Each district in a country is assigned the baseline land elasticity associated with the geographic region to which it was assigned in our regression analysis. Temperate districts thus receive a value of 0.285, tropical districts a value of 0.126, and districts capable of growing both types of crops receive the value of 0.167. Districts incapable of growing any of those crops (e.g. deserts or polar areas) are excluded from the aggregation.

The aggregate land elasticity is a weighted average of the district-level elasticities, with the weights based on the total potential calories that can be produced by a district relative to the country as a whole. Those potential calories are built as in Galor and Özak (2016).[30] The formula for country c is

$$\beta_c = \sum_{i \in I_c} \frac{\mathrm{cal}_{ic}}{\sum_{j \in I_c} \mathrm{cal}_{jc}} \beta_{ic}, \tag{12}$$

where I_c is the set of districts in country c, cal_{ic} are the potential calories in district i in country c, and β_{ic} is the land elasticity of district i from country c.

Table 8 shows the results of this aggregation.[31] Many countries have all their districts with identical geographic classifications, so their aggregate value is identical to one of our baseline values. Burkina Faso is entirely tropical, so the implied elasticity is exactly 0.126, while Denmark is entirely temperate, so has a value of 0.285. Regardless, one can see the variation across countries that does exist. Some of the more interesting entries involve countries that are heterogeneous in climate. Brazil has an aggregate elasticity of 0.139, China a value of 0.172, and the USA a value of 0.200. What does become clear is that there are few countries that are entirely temperate, and they are almost all exclusively in northern Europe. Given that, it is worth reiterating that our estimates of these elasticities did not use any cross-country, or even cross-state, variation. The variation seen in Table 8 is not an artefact of the level of development but rather represents something distinct about land elasticities by geographic area.

IV. Implications of Variation in Land Elasticities

The importance of the land elasticity for development comes from a combination of the low-income elasticity for agricultural goods, and the (relatively) fixed nature of land. Some combination of these two features is part and parcel of nearly every description of the structural transformation out of agriculture (Kogel and Prskawetz 2001; Gollin *et al.* 2007; Restuccia *et al.* 2008; Gollin 2010; Vollrath 2011; Alvarez-Cuadrado and Poschke 2011; Herrendorf *et al.* 2014; Duarte and Restuccia 2010). To see the logic involved, consider a very simplified model where L people have a fixed demand for agricultural goods of $\overline{c}_A$, and only land and labour are involved in production.[32] Equating demand and supply, we have

$$\overline{c}_A L = A_A X^\beta L_A^{1-\beta}, \tag{13}$$

which can be solved for the share of labour in agriculture,

$$\frac{L_A}{L} = \left(\frac{\overline{c}_A L^\beta}{A_A X^\beta} \right)^{1/(1-\beta)}. \tag{14}$$

TABLE 8
COUNTRY-LEVEL AGGREGATE LAND ELASTICITY ESTIMATES

Country	β	Country	β	Country	β	Country	β
Afghanistan	0.202	Egypt	0.173	Lithuania	0.285	Sao Tome	0.133
Albania	0.180	El Salvador	0.129	Luxembourg	0.285	Senegal	0.126
Algeria	0.212	Eq. Guinea	0.128	Macao	0.167	Serbia	0.202
American Samoa	0.126	Eritrea	0.153	Madagascar	0.165	Sierra Leone	0.133
Angola	0.165	Estonia	0.285	Malawi	0.161	Slovakia	0.274
Argentina	0.179	Ethiopia	0.164	Malaysia	0.137	Slovenia	0.276
Australia	0.176	Fiji	0.145	Mali	0.126	Solomon Islands	0.125
Austria	0.285	Finland	0.285	Martinique	0.126	Somalia	0.132
Azerbaijan	0.183	France	0.245	Mauritania	0.133	South Africa	0.167
Bangladesh	0.166	French Guiana	0.126	Mexico	0.165	South Korea	0.189
Belarus	0.285	Gabon	0.126	Mongolia	0.285	South Sudan	0.128
Belgium	0.285	Gambia	0.126	Morocco	0.171	Spain	0.177
Benin	0.126	Georgia	0.188	Mozambique	0.159	Sri Lanka	0.128
Bhutan	0.205	Germany	0.285	Myanmar	0.161	Sudan	0.134
Bolivia	0.157	Ghana	0.126	Namibia	0.167	Suriname	0.126
Bosnia	0.197	Greece	0.167	Netherlands	0.285	Swaziland	0.171
Botswana	0.167	Guadeloupe	0.126	New Caledonia	0.164	Sweden	0.285
Brazil	0.140	Guatemala	0.157	New Zealand	0.279	Switzerland	0.281
Brunei	0.126	Guinea	0.139	Nicaragua	0.131	Syria	0.200
Bulgaria	0.204	Guinea-Bissau	0.126	Niger	0.131	Taiwan	0.167
Burkina Faso	0.126	Guyana	0.134	Nigeria	0.128	Tajikistan	0.198
Burundi	0.181	Haiti	0.142	North Korea	0.266	Tanzania	0.157
C. African Rep.	0.126	Honduras	0.151	Norway	0.285	Thailand	0.138
Cambodia	0.128	Hungary	0.213	Oman	0.172	Timor-Leste	0.129
Cameroon	0.139	India	0.157	Pakistan	0.170	Togo	0.126
Canada	0.283	Indonesia	0.141	Palestine	0.182	Tunisia	0.167
Chad	0.127	Iran	0.195	Panama	0.134	Turkey	0.218
Chile	0.260	Iraq	0.172	Papua N. G.	0.154	Uganda	0.141
China	0.173	Isle of Man	0.285	Paraguay	0.165	Ukraine	0.278
Colombia	0.141	Italy	0.167	Peru	0.156	United Kingdom	0.284
Costa Rica	0.145	Japan	0.192	Philippines	0.131	United States	0.203
Cote d'Ivoire	0.126	Jordan	0.239	Poland	0.285	Uruguay	0.167
Croatia	0.212	Kazakhstan	0.280	Portugal	0.178	Uzbekistan	0.251
Cuba	0.128	Kenya	0.157	Rep. of Congo	0.127	Vanuatu	0.130
Czech Republic	0.285	Kosovo	0.247	Reunion	0.169	Venezuela	0.145
D. R. Congo	0.141	Kyrgyzstan	0.275	Romania	0.243	Vietnam	0.154
Denmark	0.285	Laos	0.160	Russia	0.278	Virgin Islands, US	0.126
Djibouti	0.126	Latvia	0.285	Rwanda	0.210	Zambia	0.167
Dominican Rep.	0.144	Lebanon	0.184	Samoa	0.132	Zimbabwe	0.167
Ecuador	0.159	Liberia	0.126				

Notes
This table reports the aggregated value of the land elasticity β for each country. The aggregate value is a weighted average of district land elasticities with tropical districts (0.126), temperate districts (0.285), and mixed districts (0.167) that can grow both tropical and temperate crops. The weights in the average are the maximum calories that can be produced in a district relative to the maximum calories that can be produced by all districts in the country.

Note that the sensitivity of L_A/L to productivity (A_A) and population (L) depends on the size of β. The larger the land elasticity, the *more* sensitive is L_A/L to both of these terms. This is because β dictates the degree of decreasing returns to scale for labour in agriculture, with a larger land elasticity implying more severe decreasing returns. Hence larger movements of labour into or out of agriculture are necessary to keep the supply of agricultural goods equal to demand. Given that non-agricultural labour is the alternative use for labour, this means that changes in non-agricultural employment (and hence urbanization to some extent) also depend on the land elasticity.

The variation in the land elasticity that we found between temperate and tropical regions has a significant effect on how these places respond to technological improvements or to population growth. Temperate and tropical areas starting out with identical living standards and shares of workers in agriculture could end up far different over time even if they faced the same trend growth in productivity. The temperate area would have a larger fraction of workers in non-agriculture (and plausibly a higher urbanization rate and living standards) than the tropical area. If there were agglomeration effects in urban areas, or demographic effects of a declining agricultural labour share, then any initial advantage conveyed on a temperate area would be exaggerated.

That said, it is not the case that temperate areas with high land elasticities must necessarily have an advantage in development. The high land elasticity also makes temperate areas more sensitive to *negative* shocks to productivity, or to increases in population. Tropical regions with low land elasticities would thus be able to survive poor weather or unexpected population increases with a smaller effect on the agricultural labour share (and plausibly on urbanization and living standards). Low land elasticity would have allowed tropical regions to be more resilient in the face of shocks compared to temperate regions, but perhaps at the cost of long-run development.

Evidence from the epidemiological transition

The epidemiological transition that occurred following the Second World War provides a useful context in which to assess whether the variation in β that we documented has effects consistent with the intuition presented in the previous subsection. Acemoglu and Johnson (2007) collected mortality rate data from the postwar period for a set of 15 infectious diseases (e.g. tuberculosis and malaria). They argue that the epidemiological transition formed an exogenous shock to population health, and therefore population size, in developing countries, and use it to identify the causal impact of health on living standards. We can use the same empirical setting to ask whether the impact of these plausibly exogenous health interventions *differs* based on whether countries had a high β value or a low β value. Based on our simple explanation above, we would expect that living standards in places with the high β should be more sensitive to these mortality shocks than places with low β values. In particular, given that this is a positive shock to population size, our expectation is that high β places will experience a more severe *negative* shock to living standards.

To test this, we first restrict ourselves to the low- and middle-income sample from Acemoglu and Johnson (2007). We make this restriction because rich countries, regardless of their value of β, are not going to be affected by the decreasing returns in the agricultural sector to any meaningful degree, given their low agricultural labour share to begin with.[33] With those low- and middle-income countries, we then assign them to either a 'tropical' or a 'temperate' group. The assignment is based on the estimated β for each country found in Table 8. Those with a value of β less than or equal to 0.20 are classified as tropical (which in practice includes any area with a significant amount of mixed land), while those with a value of β above 0.20 are assigned to the temperate category. We make the cut-off relatively high so that we can isolate the true temperate countries in the Acemoglu and Johnson dataset. Thus 30 countries are classified as tropical, and only 5 as temperate.

For both tropical and temperate groups, we use the original data from Acemoglu and Johnson (2007) to run panel regressions with the specification

$$y_{it} = \alpha + \theta x_{it} + \gamma_i + \delta_t + \varepsilon_{it}, \tag{15}$$

where y_{it} is one of three different dependent variables (log GDP per capita, log GDP per worker, or log population), and x_{it} is one of three different independent variables (mortality rates, log life expectancy, or log population). θ captures the effect of the independent variable on y_{it}, and we will compare the value of θ across samples that differ based on whether they have low land elasticities or high land elasticities. γ_i and δ_t are country and decade fixed effects, while ε_{it} is the error term. Each country has up to eight decadal observations, running from 1930 to 2000, but the panel is not balanced.[34]

Table 9 presents the results. In Panel A, the explanatory x_{it} variable is the original mortality instrument from Acemoglu and Johnson (2007), which measures the mortality rate from the 15 infectious diseases that were affected by the interventions following the Second World War. In columns (1) and (2), we show the effect of mortality rates on (log) GDP per capita. As can be seen, the estimated coefficient for tropical countries (0.403) in column (1) is smaller than the estimate for temperate countries (1.226) in column (2). Below these estimates are two hypothesis tests. First, we see the test that the effect size is zero, $\theta = 0$. We can reject zero effect for both temperate and tropical countries. The hypothesis that θ is identical for the two samples has a p-value of less than 0.1%. The higher value for temperate countries is consistent with the intuition of the previous subsection, where a high land elasticity is associated with higher sensitivity to shocks.

Columns (3) and (4) of Panel A of Table 9 repeat this test, but now using (log) GDP per worker as the dependent variable. The effect of mortality is estimated to be almost three times larger for temperate as for tropical countries. This difference is significant at less than 0.1%, and shows that high land elasticity countries are more sensitive to population shocks than low land elasticity countries. These columns show that mortality shocks affected the average output of each *worker*, and the effect on per capita GDP did not arise because of short-run changes in the age structure of the economy.

Columns (5) and (6) of Panel A of Table 9 show the effect of the mortality shocks on population size. In tropical countries, the effect of mortality on population was estimated to be smaller than in temperate countries (-0.255 versus -1.442). Thus it may be that the temperate countries were hit by a larger shock to their population due to the epidemiological transition, perhaps acting as part of the explanation for their stronger response to the mortality changes. The results in columns (1)–(4) should be interpreted with that caveat in mind.

Panel B of Table 9 repeats the regressions, but now uses life expectancy itself as the explanatory variable x_{it}, matching the original work of Acemoglu and Johnson (2007). Whether looking at GDP per capita (columns (1) and (2)) or GDP per worker (columns (3) and (4)), we have large and statistically significant differences in the estimated effects of life expectancy between tropical and temperate countries. For both areas, the effect of

TABLE 9
PANEL ESTIMATES OF EFFECT OF POPULATION CHANGE, BY LAND ELASTICITY

	Log GDP per capita		Log GDP per worker		Log population	
	Temperate	Tropical	Temperate	Tropical	Temperate	Tropical
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
Panel A:						
Mortality rate	0.403	1.226	0.418	1.273	−0.255	−1.442
	(0.151)	(0.132)	(0.158)	(0.140)	(0.108)	(0.193)
p-value $\theta = 0$	0.008	0.000	0.009	0.000	0.019	0.000
p-value $\theta = \theta^{\text{Trop}}$		0.000		0.000		0.000
Countries	30	5	30	5	30	5
Observations	238	40	238	40	238	40
Panel B:						
Log life expectancy	−0.743	−2.180	−0.703	−2.324	1.383	2.870
	(0.290)	(0.171)	(0.285)	(0.187)	(0.165)	(0.248)
p-value $\theta = 0$	0.011	0.000	0.014	0.000	0.000	0.000
p-value $\theta = \theta^{\text{Below}}$		0.000		0.000		0.000
Countries	30	5	30	5	30	5
Observations	224	40	224	40	224	40

Notes
Robust standard errors are reported in parentheses. All regressions include both year fixed effects and country fixed effects. Countries are assigned as either 'Temperate' or 'Tropical' based on their estimated β from Table 8. Those with $\beta < 0.20$ are classified as tropical, while those with $\beta \geq 0.20$ are classified as temperate. The mortality rate used as an explanatory variable in Panel A is the mortality rate from 15 infectious diseases, as documented by Acemoglu and Johnson (2007). All data on GDP per capita, GDP per worker, population and life expectancy are also taken from those authors' dataset. The p-value of $\theta = \theta^{\text{Trop}}$ is from a test that the estimated coefficient for temperate countries is identical to the estimated coefficient for tropical countries.

higher life expectancy is negative for GDP per capita and GDP per worker, but is more severe for the temperate countries.[35] The difference is again significant at less than 0.1%.

In columns (5) and (6) of Panel B of Table 9, the effect of life expectancy on population size is positive in both sets of countries, with a smaller estimated effect size in tropical countries. As in Panel A, it is possible that there was simply a larger population shock in temperate areas in response to the health interventions.

The evidence in Table 9 shows that the variation in land elasticity (β) that we identified in the main part of the paper has effects consistent with those predicted by the intuition earlier in this section. Given the differentials that we estimated in the effect of the epidemiological transition, the land elasticity appears to have non-trivial implications for development, making temperate areas more sensitive to shocks than tropical areas.

V. CONCLUSION

The role that land plays in agricultural production is relevant to any study of agriculture and development. We estimated the elasticity of aggregate agricultural production with respect to land, and found that it differed significantly between temperate and tropical regions of the world.

Our estimates are made by looking at the relationship between agricultural worker labour/land ratios and potential agroclimatic yield at the district level (e.g. second-level administrative units) from 154 countries. Our methodology lets us use the district variation within states to identify the land elasticity, and avoids the need to specify or measure other inputs directly. This also avoids comparing countries—or even states—at different levels of development. The method rests mainly on the assumption that labour is mobile across districts within states, something that recent evidence supports.

Our baseline finding, that the land elasticity in temperate areas is about 0.285 while it is only 0.126 in tropical areas, is robust to different ways of measuring rural labour/land ratios and potential yield, and robust to alternative definitions of what constitutes tropical versus temperate areas. What our estimation technique does not provide is a way of identifying *why* the aggregate elasticities vary so much between tropical and temperate areas, and whether that is due to biological requirements of certain crops, or the constraints imposed by aspects of the climate itself.

These estimates are for the aggregate land elasticity, and as such are informative for research that studies the role of the aggregate agricultural sector in development, whether that is related to structural change in developing countries today, or related to historical development in standard Malthusian settings. This aggregate land elasticity, regardless of the setting, is an important parameter in determining the sensitivity of income per capita and the share of labour in agriculture to shocks in population growth or productivity. The larger the land elasticity, the more sensitive an economy is to those shocks. We confirmed this prediction by showing that in response to the epidemiological transition following the Second World War, countries with larger land elasticities did see more severe changes in their GDP per worker and GDP per capita.

More generally, we contribute to the understanding of relative development levels in tropical and temperate areas of the world. By making temperate areas more sensitive to shocks, a high aggregate land elasticity allowed them to leverage positive shocks to productivity (e.g. technological improvements) and population growth (e.g. the demographic transition) to accelerate their growth relative to tropical areas. Slower development in tropical regions—either historically or in the current era—may reflect in part differences in the size of the aggregate land elasticity, rather than any deficiency in productivity or population growth.

ACKNOWLEDGMENTS

We thank Tasso Adamopoulos, Francesco Caselli, Areendam Chanda, Martin Fiszbein, Oded Galor, Remi Jedwab, Nippe Lagerlöf, Debin Ma, Stelios Michalopolous, Nathan Nunn, Ömer Özak, Stephen Smith, Enrico Spolaore, Joachim Voth and David Weil, as well as seminar participants at LSU, York University, London School of Economics, the Department of Agricultural Economics at Texas A&M University, the Brown Conference on ‘Deep-rooted determinants of development’, George Washington University, and the University of Houston brown bag series for their comments. This paper previously circulated under the title ‘How tight are Malthusian constraints?’. All errors remain our own.

NOTES

1. Agriculture and land feature in stories of divergence across global regions (Kogel and Prskawetz 2001; Galor and Mountford 2008; Vollrath 2011; Voigtländer and Voth 2013a,b; Cervellati and Sunde 2015). On structural change, see Gollin *et al*. (2007), Restuccia *et al*. (2008), Weil and Wilde (2009), Gollin (2010), and Eberhardt and Vollrath (2018). For Malthusian stagnation, see Ashraf and Galor (2011) for a baseline model, and Galor (2011) for a review of major contributions to the literature on the take-off to growth (Galor and Weil 2000; Galor and Moav 2002; Hansen and Prescott 2002; Doepke 2004; Cervellati and

Sunde 2005; Lagerlöf 2006; Crafts and Mills 2009; Strulik and Weisdorf 2008). On the relevance of resources for long-run growth, see Peretto and Valente (2015).

2. There are two studies that also study the spatial distribution of labour at the global level in some capacity. The first is Motamed *et al*. (2014), which examines the growth of urbanization at the grid-cell level over the last two thousand years. The second is Henderson *et al*. (2016), which examines the spatial distribution of economic activity at the grid-cell level using night lights.
3. These results are consistent with the work of Ruthenberg (1976) and Bray (1994), who discuss the inherent differences in the response of tropical crops (rice, in particular) to the application of labour. They both cite the relatively *high* elasticity of output with respect to labour in tropical agriculture, which is consistent with a low elasticity of output with respect to land.
4. We show as part of our robustness checks that our results hold if districts that are large livestock or cash crop producers are included or excluded from the regressions. The Online Appendix also contains explicit estimates of the land elasticity for districts that are major cash crop producers, and their elasticities tend to be slightly higher than the tropical value, apart from tea producers, which have an elasticity near the temperate estimate.
5. Firm-level methods (Olley and Pakes 1996; Levinsohn and Petrin 2003) are not applicable in our context because we do not have a panel of district data, nor do we have full data on the inputs used at the district level.
6. More general treatments of this idea can be found in Houthakker (1955) and Jones (2005). In short, the farm-level land elasticities may not be informative on the aggregate land elasticity, and farm-level production functions may well take on forms (i.e. Leontief versus Cobb–Douglas) different from the aggregate function.
7. It would be impossible to summarize or cite all the research on comparative development. Several useful reviews of this literature can be found in Acemoglu *et al*. (2005), Nunn (2009), Galor (2011), Spolaore and Wacziarg (2013), and Vries (2013).
8. The divergence of China, and the lower Yangtze region in particular, from north-western Europe is the subject of a large literature. Pomeranz (2000) is the standard starting point, while Allen *et al*. (2011), Huang (2002), Ma (2013), Lee *et al*. (2002), and Broadberry and Gupta (2006) are a brief selection of relevant papers.
9. Our work is related to several recent studies on the role of geography and/or inherent agricultural productivity in development (Olsson and Hibbs 2005; Ashraf and Galor 2011; Nunn and Qian 2011; Nunn and Puga 2012; Michalopoulos 2012; Alesina *et al*. 2013; Cook 2014a,b; Fenske 2014; Alsan 2015; Ashraf and Michalopoulos 2015; Dalgaard *et al*. 2015; Galor and Özak 2016; Litina 2016; Andersen *et al*. 2016; Frankema and Papaioannou 2017). Unlike those papers, ours does not propose a direct causal relationship between geography and development, but rather suggests that *any* proposed causal impact has differential effects based on the size of the land elasticity.
10. This covers nearly every country in the world, apart from Libya and Saudi Arabia, for which usable maps of the district level were not available.
11. There are a handful of very high population districts, of course, representing large urban areas. In our data, 1040 districts have populations above one million people. But that represents less than 3% of all districts. Our results exclude these high population districts.
12. Note that this does not imply equalization of total *earnings*, given differences in average human capital per person in the two areas. But both papers establish that conditional on measures of human capital, there do not appear to be any significant wage premia simply for living in urban areas.
13. This 'capital' can be thought of as a combination term capturing a set of non-labour inputs to agriculture.
14. While we have written the function here as Cobb–Douglas, this is solely for ease of exposition; the analysis does not require this. In the Online Appendix, we show that one could use a general constant returns to scale function to derive a similar estimation equation.
15. In the Online Appendix, we show how our empirical specification would change when dropping this assumption. We also show that the empirical setting is not dependent on the assumption of homogeneous labour.
16. If one were to assume that capital was also mobile across districts within a state, then these controls would not be necessary. The capital/labour ratio would then be identical across districts, and state fixed effects would absorb the capital/labour term.
17. In the Online Appendix, we explain in more detail why it makes sense to assume that A_{Ais} and A_{Ais}^{GAEZ} are related with an elasticity of 1, as in our specification.
18. Links to the raw files for population, and all other data used in this paper, along with code to build our datasets and replicate all regressions, can be found at https://github.com/dvollrath/Crops (accessed 21 December 2019).
19. We use the low-input, rain-fed indices of caloric yield provided by Galor and Özak (2016) in our baseline specification. Our results are robust to using different assumptions on inputs and water use. The specific crops included in our calculation are alfalfa, banana, barley, buckwheat, cassava, chickpea, cowpea, drypea, flax, foxtail millet, greengram, groundnut, indica rice, maize, oat, pearl millet, phaseolus bean, pigeon pea, rye, sorghum, soybean, spring wheat, sweet potato, rape, wet/paddy rice, wheat, winter wheat, white potato and yams.

20. We have experimented with alternative sets of crops to define the regions, without any material change to our results.
21. There are reasons to believe that frost may raise the productivity *level* of agriculture by killing off pests and organisms that mineralize organic matter, but this difference in productivity does *not* have anything to do with our results. Our estimates of β_g are made within-state for districts that have the *same* frost characteristics, and are not based on any comparison of frost versus frost-free districts.
22. The Köppen–Geiger scheme has several levels. For temperate, we use districts that have any land in their climate class 'C' (warm temperate) or 'D' (snow), and also having any land in their temperature class 'b' (warm summer) or 'c' (cool summer). For tropical, we use districts that have any land in their climate class 'A' (equatorial). There are no temperature subdivisions within the equatorial class. There are also precipitation classifications, but we do not use those for either temperate or tropical assignment. Pixel-level data on Köppen–Geiger classification is from Kottek *et al.* (2006).
23. We suppressed the district-level borders from the map as they are so small that it becomes something of a mess, and prevents one from seeing the information about crop types.
24. Using the quantiles still gives an accurate indication of the relationships in the data. See Chetty *et al.* (2013) for an explanation and example of this kind of figure.
25. Advanced economies with modern farming like Japan and South Korea are already excluded from our regressions by how we define tropical and temperate areas, given that they are capable of growing both kinds of crops.
26. Because district-level boundaries can change over time, the IPUMS aggregates to the largest possible units that are stable over time, which means fewer districts. This also means that there are far fewer districts within any given state (and in some cases even states are aggregated), so we use country-level fixed effects with the IPUMS regressions, rather than state-level.
27. The cash crops that we consider are bananas, cocoa, coffee, cotton, jute, palm oil, rubber, sunflowers, tea, tobacco, sugar beets and sugar cane. Estimates of the land elasticity for districts that are substantial producers of these cash crops are provided in the Online Appendix.
28. On the demographic side, we have the median, and 10th and 90th percentiles of the household head's age, years of education, and typical number of household residents. On the asset side, we have the fraction of households with the following: toilet, electricity, television, refrigerator, improved flooring, any agricultural land, a bank account, any cattle, any draft animals, and any sheep. Some surveys contain measures of the amount of agricultural land, as well as counts of livestock, but there are too few of these to do a comparison across temperate and tropical regions.
29. Work by Wilde (2012) indicates that the elasticity of substitution is *less* than 1, using historical information from the UK.
30. Conceptually, the weights should properly be based on the share of real output produced by each district. In the absence of real agricultural output data at the district level, we use the potential calories as a proxy.
31. This does not encompass every country in the world, as several do not have second-level districts to use in the aggregation. An alternative is to aggregate up from the pixel level, which produces similar values for the reported countries, and increases the total number of countries reported. Those results are in the Online Appendix.
32. In the Online Appendix we present a richer two-sector model that allows for more nuanced income and substitution effects in the demand for agriculture, capital in the production function, and an endogenous relative price of agricultural goods that demonstrates the same conclusion that we describe here.
33. We could expand the data to include up to 45 countries in some regressions where we have sufficient mortality and GDP data. To create comparable samples across all of our regressions, we limit ourselves to the 35 countries with full data.
34. Rather than separating countries into two groups based on β and comparing θ between them, an alternative specification would be to interact β_i with x_{it}, as in $y_{it} = \alpha + \theta_0 x_{it} + \theta_1 \beta_i \times x_{it} + \gamma_i + \delta_t + \varepsilon_{it}$. In this case, the estimated value of θ_1 would indicate how the effect of x_{it} differs with the size of β. Doing this produces results consistent with those presented in Table 9.
35. Whether changes in health, as proxied by life expectancy, are in fact positive or negative in the long run for development is beyond the scope of this paper, and the original findings of Acemoglu and Johnson (2007) are debated (Bloom *et al.* 2014).

REFERENCES

Acemoglu, D. and Johnson, S. (2007). Disease and development: the effect of life expectancy on economic growth. *Journal of Political Economy*, **115**(6), 925–85.

———, ——— and Robinson, J. A. (2005). Institutions and fundamental cause of long-run growth. In P. Aghion and S. Durlauf (eds), *Handbook of Economic Growth*, Vol. 1. Amsterdam: North-Holland, pp. 385–472.

Alesina, A., Giuliano, P. and Nunn, N. (2013). On the origins of gender roles: women and the plough. *Quarterly Journal of Economics*, **128**(2), 469–530.

Allen, R. C., Bassino, J.-P., Ma, D., Moll-Murata, C. and Van Zanden, J. L. (2011). Wages, prices, and living standards in China, 1738–1925: in comparison with Europe, Japan, and India. *Economic History Review*, **64**(S1), 8–38.

Alsan, M. (2015). The effect of the tsetse fly on African development. *American Economic Review*, **105**(1), 382–410.

Alvarez-Cuadrado, F. and Poschke, M. (2011). Structural change out of agriculture: labor push versus labor pull. *American Economic Journal: Macroeconomics*, **3**, 127–58.

Andersen, T. B., Dalgaard, C.-J. and Selaya, P. (2016). Climate and the emergence of global income differences. *Review of Economic Studies*, **83**(4), 1334–63.

Ashraf, Q. and Galor, O. (2011). Dynamics and stagnation in the Malthusian epoch. *American Economic Review*, **101**(5), 2003–41.

——— and Michalopoulos, S. (2015). Climatic fluctuations and the diffusion of agriculture. *Review of Economics and Statistics*, **97**(3), 589–609.

Bloom, D. E., Canning, D. and Fink, G. (2014). Disease and development revisited. *Journal of Political Economy*, **122**(6), 1355–66.

Bray, F. (1994). *The Rice Economies, Technology and Development in Asian Societies*. Berkeley, CA: University of California Press.

Broadberry, S. and Gupta, B. (2006). The early modern great divergence: wages, prices and economic development in Europe and Asia, 1500–1800. *Economic History Review*, **59**(1), 2–31.

Center for International Earth Science Information Network (CIESIN), International Food Policy Research Institute (IFPRI), World Bank and Centro Internacional de Agricultura Tropical (CIAT) (2011). Global Rural–Urban Mapping Project, Version 1.

Cervellati, M. and Sunde, U. (2005). Human capital formation, life expectancy, and the process of development. *American Economic Review*, **95**(5), 1653–72.

——— and ——— (2015). The economic and demographic transition, mortality, and comparative development. *American Economic Journal: Macroeconomics*, **7**(3), 189–225.

Chetty, R., Friedman, J. N. and Saez, E. (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *American Economic Review*, **103**(7), 2683–721.

Clark, G. (2002). The Agricultural Revolution and the Industrial Revolution. UC-Davis Working Paper.

Cook, C. J. (2014a). Potatoes, milk, and the Old World population boom. *Journal of Development Economics*, **110**(C), 123–38.

——— (2014b). The role of lactase persistence in precolonial development. *Journal of Economic Growth*, **19**(4), 369–406.

Crafts, N. and Mills, T. C. (2009). From Malthus to Solow: how did the Malthusian economy really evolve? *Journal of Macroeconomics,* **31**(1), 68–93.

Craig, B. J., Pardey, P. G. and Roseboom, J. (1997). International productivity patterns: accounting for input quality, infrastructure, and research. *American Journal of Agricultural Economics*, **79**(4), 1064–76.

Dalgaard, C.-J., Knudsen, A. S. B. and Selaya, P. (2015). The bounty of the sea and long-run development. CESifo Working Paper no. 5547.

Doepke, M. (2004). Accounting for fertility decline during the transition to growth. *Journal of Economic Growth*, **9**(3), 347–83.

Duarte, M. and Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *Quarterly Journal of Economics*, **125**(1), 129–73.

Eberhardt, M. and Teal, F. (2013). No mangos in the tundra: spatial heterogeneity in agricultural productivity analysis. *Oxford Bulletin of Economics and Statistics*, **75**(6), 914–39.

——— and Vollrath, D. (2018). The effect of agricultural technology on the speed of development. *World Development*, **109**, 483–96.

Elvidge, C. D., Baugh, K. E., Dietz, J. B., Bland, T., Sutton, P. C. and Kroehl, H. W. (1999). Radiance calibration of DMSP-OLS low-light imaging data of human settlements. *Remote Sensing of Environment*, **68** (1), 77–88.

Fenske, J. (2014). Ecology, trade, and states in pre-colonial Africa. *Journal of the European Economic Association*, **12**(3), 612–40.

Food and Agriculture Organization (2012). Global Agro-Ecological Zones; available online at www.fao.org/nr/GAEZ (accessed 22 December 2019).

Frankema, E. and Papaioannou, K. (2017). Rainfall patterns and human settlement in tropical Africa and Asia compared. Did African farmers face greater insecurity? CEPR Discussion Paper.

Fuglie, K. (2010). Total factor productivity in the global agricultural economy: evidence from FAO data. In P. Pardey, J. Alston and B. Babcock (eds), *The Shifting Patterns of Agricultural Production and Productivity Worldwide* Ames, IA: Midwest Agribusiness Trade and Research Information Center, pp. 63–95.

GALOR, O. (2011). *Unified Growth Theory*. Princeton, NJ: Princeton University Press.

——— and MOAV, O. (2002). Natural Selection and the origin of economic growth. *Quarterly Journal of Economics*, **117**(4), 1133–91.

——— and MOUNTFORD, A. (2008). Trading population for productivity: theory and evidence. *Review of Economic Studies*, **75**(4), 1143–79.

——— and ÖZAK, Ö. (2016). The agricultural origins of time preference. *American Economic Review*, **106**(10), 3064–103.

——— and WEIL, D. N. (2000). Population, technology, and growth: from Malthusian stagnation to the demographic transition and beyond. *American Economic Review*, **90**(4), 806–28.

GLOBAL ADMINISTRATIVE AREAS PROJECT (2019). GADM Database; available online at https://gadm.org/index.html (accessed 22 December 2019).

GOLDEWIJK, K. K., BEUSEN, A., VAN DRECHT, G. and DE VOS, M. (2011). The HYDE 3.1 spatially explicit database of human-induced global land-use change over the past 12,000 years. *Global Ecology and Biogeography*, **20**(1), 73–86.

GOLLIN, D. (2010). Agricultural productivity and economic growth. In P. Pingali and R. Evenson (eds), *Handbook of Agricultural Economics*, Vol. 4. Amsterdam: Elsevier, pp. 3825–66.

———, PARENTE, S. and ROGERSON, R. (2007). The food problem and the evolution of international income levels. *Journal of Monetary Economics*, **54**, 1230–55.

GUTIERREZ, L. and GUTIERREZ, M. M. (2003). International R&D spillovers and productivity growth in the agricultural sector. A panel cointegration approach. *European Review of Agricultural Economics*, **30**(3), 281–303.

HANSEN, G. D. and PRESCOTT, E. C. (2002). From Malthus to Solow. *American Economic Review*, **92**(4), 1205–17.

HAYAMI, Y. and RUTTAN, V. W. (1970). Agricultural productivity differences among countries. *American Economic Review*, **60**(5), 895–911.

——— and ——— (1985). *Agricultural Development: An International Perspective*. Baltimore, MD: Johns Hopkins University Press.

———, ——— and SOUTHWORTH, H. M. (1979). *Agricultural Growth in Japan, Taiwan, Korea, and the Philippines*. Honolulu, HI: East-West Center.

HENDERSON, J. V., SQUIRES, T. L., STOREYGARD, A. and WEIL, D. N. (2016). The global spatial distribution of economic activity: nature, history, and the role of trade. NBER Working Paper no. 22145.

HERRENDORF, B., ROGERSON, R. and VALENTINYI, Á. (2014). Growth and structural transformation. In P. Aghion and S. Durlauf (eds), *Handbook of Economic Growth*, Vol. 2. Kidlington: Elsevier, pp. 855–941.

HICKS, J. H., KLEEMANS, M., LI, N. Y. and MIGUEL, E. (2017). Reevaluating agricultural productivity gaps with longitudinal microdata. NBER Working Paper no. 23253.

HOUTHAKKER, H. S. (1955). The Pareto distribution and the Cobb–Douglas production function in activity analysis. *Review of Economic Studies*, **23**(1), 27–31.

HUANG, P. C. C. (2002). Development or involution in eighteenth-century Britain and China? A review of Kenneth Pomeranz's *The Great Divergence: China, Europe, and the Making of the Modern World Economy*. *Journal of Asian Studies*, **61**(2), 501–38.

JONES, C. I. (2005). The shape of the production function and the direction of technical change. *Quarterly Journal of Economics*, **120**(2), 517–49.

JORGENSON, D. and GOLLOP, F. (1992). Productivity growth in U.S. agriculture: a postwar perspective. *American Journal of Agricultural Economics*, **74**(3), 745–50.

KOGEL, T. and PRSKAWETZ, A. (2001). Agricultural productivity growth and escape from the Malthusian trap. *Journal of Economic Growth*, **6**(4), 337–57.

KOTTEK, M., GRIESER, J., BECK, C., RUDOLF, B. and RUBEL, F. (2006). World map of the Köppen–Geiger climate classification updated. *Meteorologische Zeitschrift*, **15**(3), 259–63.

LAGERLÖF, N.-P. (2006). The Galor–Weil model revisited: a quantitative exercise. *Review of Economic Dynamics*, **9**(1), 116–42.

LEE, J., CAMPBELL, C. and FENG, W. (2002). Positive check or Chinese checks? *Journal of Asian Studies*, **61**(2), 591–607.

LEVINSOHN, J. and PETRIN, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, **70**(2), 317–41.

LITINA, A. (2016). Natural land productivity, cooperation and comparative development. *Journal of Economic Growth*, **21**(4), 351–408.

MA, D. (2013). State capacity and great divergence, the case of Qing China (1644–1911). *Eurasian Geography and Economics*, **54**(5–6), 484–99.

MARTIN, W. and MITRA, D. (2001). Productivity growth and convergence in agriculture versus manufacturing. *Economic Development and Cultural Change*, **49**(2), 403–22.

MASTERS, W. A. and MCMILLAN, M. S. (2001). Climate and scale in economic growth. *Journal of Economic Growth*, **6**(3), 167–86.

MEIJER, J. R., HUIJBREGTS, M. A. J., SCHOTTEN, K. C. G. J. and SCHIPPER, A. M. (2018). Global patterns of current and future road infrastructure. *Environmental Research Letters*, **13**(6), 1–10.

MICHALOPOULOS, S. (2012). The origins of ethnolinguistic diversity. *American Economic Review*, **102**(4), 1508–39.

MINNESOTA POPULATION CENTER (2017). Integrated Public Use Microdata Series (IPUMS), International: Version 6.5. Minneapolis, MN: University of Minnesota.

MONFREDA, C., RAMANKUTTY, N. and FOLEY, J. A. (2008). Farming the planet: 2. Geographic distribution of crop areas, yields, physiological types, and net primary production in the year 2000. *Global Biogeochemical Cycles*, **22**(1), 1–19.

MOTAMED, M. J., FLORAX, R. J. G. M. and MASTERS, W. A. (2014). Agriculture, transportation and the timing of urbanization: global analysis at the grid cell level. *Journal of Economic Growth*, **19**(3), 339–68.

MUNDLAK, Y. (2000). *Agriculture and Economic Growth: Theory and Measurement*. Cambridge, MA: Harvard University Press.

———, BUTZER, R. and LARSON, D. F. (2012). Heterogeneous technology and panel data: the case of the agricultural production function. *Journal of Development Economics*, **99**(1), 139–49.

NUNN, N. (2009). The importance of history for economic development. *Annual Review of Economics*, **1**, 65–92.

——— and PUGA, D. (2012). Ruggedness: the blessing of bad geography in Africa. *Review of Economics and Statistics*, **94**(1), 20–36.

——— and QIAN, N. (2011). The potato's contribution to population and urbanization: evidence from a historical experiment. *Quarterly Journal of Economics*, **126**(2), 593–650.

OLLEY, G. S. and PAKES, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, **64**(6), 1263–97.

OLSSON, O. and HIBBS, D. (2005). Biogeography and long-run economic development. *European Economic Review*, **49**(4), 909–38.

PERETTO, P. and VALENTE, S. (2015). Growth on a finite planet: resources, technology and population in the long run. *Journal of Economic Growth*, **20**(3), 305–31.

POMERANZ, K. (2000). *The Great Divergence*. Princeton, NJ: Princeton University Press.

RAMANKUTTY, N., EVAN, A. T., MONFREDA, C. and FOLEY, J. A. (2008). Farming the planet: 1. Geographic distribution of global agricultural lands in the year 2000. *Global Biogeochemical Cycles*, **22**(1), 1–19.

RESTUCCIA, D., YANG, D. and ZHU, X. (2008). Agriculture and aggregate productivity. *Journal of Monetary Economics*, **55**(2), 234–50.

RUTHENBERG, H. (1976). *Farming Systems in the Tropics*. Oxford: Clarendon Press.

SPOLAORE, E. and WACZIARG, R. (2013). How deep are the roots of economic development? *Journal of Economic Literature,* **51**(2), 325–69.

STRULIK, H. and WEISDORF, J. L. (2008). Population, food, and knowledge: a simple unified growth theory. *Journal of Economic Growth*, **13**(3), 195–216.

VOIGTLÄNDER, N. and VOTH, H.-J. (2013a). How the West invented fertility restriction. *American Economic Review*, **103**(6), 2227–64.

——— and ——— (2013b). The three horsemen of riches: plague, war, and urbanization in early modern Europe. *Review of Economic Studies*, **80**(2), 774–811.

VOLLRATH, D. (2011). The agricultural basis of comparative development. *Journal of Economic Growth*, **16**, 343–70.

VRIES, P. (2013). *Escaping Poverty: The Origins of Modern Economic Growth*. Vienna: Vienna University Press.

WEIL, D. N. and WILDE, J. (2009). How relevant Is Malthus for economic development today? *American Economic Review: Papers and Proceedings*, **99**(2), 255–60.

WIEBE, K., SOULE, M. J., NARROD, C. and BRENEMAN, V. E. (2003). Resource quality and agricultural productivity: a multi-country comparison. In K. Wiebe (ed.), *Land Quality, Agricultural Productivity, and Food Security*. Northampton, MA: Edward Elgar Publishing.

WILDE, J. (2012). How substitutable are fixed factors in production? Evidence from preindustrial England. University Library of Munich, MPRA Paper no. 39278.

YOUNG, A. (2013). Inequality, the urban–rural gap and migration. *Quarterly Journal of Economics*, **128**(4), 1727–85.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

A.1. General version of empirical setup
A.2. Specification with capital immobile between sectors
A.3. Specification with a wage/rent wedge included
A.4. Specification with different capital elasticities
A.5. Multiple labor types
A.6. Explicit two-sector model
A.7. Adding Malthusian fertility responses
A.8. GAEZ productivity measures and A_{Ais}
A.9. Demographic and Health Survey Data
A.10. Alternative Population Data
A.11. Labor/land and β_g
A.12. Alternative measure of A_{isg}^{GAEZ}
A.13. Climate zone results
A.14. Figures
A.15. Robustness tables