

University of Toronto
Department of Economics



Working Paper 692

Geography and Agricultural Productivity: Cross-Country
Evidence from Micro Plot-Level Data

By Tasso Adamopoulos and Diego Restuccia

March 31, 2021

Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data[†]

Tasso Adamopoulos
York University

Diego Restuccia
University of Toronto
and NBER

March 2021

ABSTRACT

We quantify the role of geography and land quality for agricultural productivity differences across countries using high-resolution micro-geography data and a spatial accounting framework. The rich spatial data provide for each cell of land covering the entire globe, the potential yield for 18 crops, which measures the maximum attainable crop output given soil quality, climate conditions, terrain topography, and a given level of cultivation inputs. While there is considerable heterogeneity in land quality across space, even within narrow geographic regions, we find that low agricultural land productivity is not due to unfavourable geographic endowments. If countries produced current crops in each cell according to potential yields, the rich-poor agricultural yield gap would virtually disappear, from 214 percent to 5 percent. We also find evidence of additional aggregate productivity gains attainable through spatial reallocation and changes in crop production.

JEL classification: O11, O14, O4.

Keywords: agriculture, land quality, productivity, spatial allocation, crop choice, cross-country.

[†]We thank Henry Kim, Yikalo Araya for valuable research assistance, and Rui Castro, Chaoran Chen, Andrew Dickens, Dave Donaldson, Doug Gollin, B. Ravikumar, Todd Schoellman, and Dietrich Vollrath for comments and suggestions. We have benefited from comments received at the Growth Conference in Montreal, the Ethiopian Economic Association Meetings in Addis Ababa, the European Economic Association Meetings in Mannheim, the International Monetary Fund, the World Bank, SED in Edinburgh, MadMac Madrid, EEA/ESEM Lisbon, and AWSOME at McMaster. Adamopoulos and Restuccia gratefully acknowledge the support from the Social Sciences and Humanities Research Council of Canada. Restuccia also thanks the support from the Canada Research Chairs program and the Bank of Canada Fellowship program. The views expressed herein are not necessarily those of the Bank of Canada and are the author's alone. Contact: Tasso Adamopoulos, aadamo@yorku.ca; Diego Restuccia, diego.restuccia@utoronto.ca

1 Introduction

Understanding the large differences in labor productivity across rich and poor countries is a fundamental issue on the research agenda in economics. It is well understood that poor countries have disproportionately low agricultural productivity and large agricultural sectors, when compared to rich countries (Gollin et al., 2002; Restuccia et al., 2008; Caselli, 2005). Why is agricultural productivity so low in poor countries? The answer to this question has important implications for poverty reduction, welfare, structural transformation, and development.

There are two possible broad explanations for the disparity in agricultural productivity across countries. First, due to varied institutions, constraints, frictions, or policies, countries make different economic choices in agriculture, affecting the level of productivity. Second, due to unfortunate endowments, featuring low land quality, rugged geography, and arid lands, some countries may have a natural disadvantage in agriculture. Understanding which of these two broad explanations is the source of low agricultural productivity across countries is essential and has dramatically different implications for policy. The vast majority of research has focused on explanations of constrained economic decisions affecting agricultural productivity. The role of land quality and geography, while often presumed and invoked in public debates on development, is largely unexplored on a systematic cross-country basis. We examine the role of geography and land quality in accounting for agricultural productivity differences across countries. While we find evidence of considerable heterogeneity in land quality, even within narrow geographic regions, our main finding is that, at the country level, differences in land quality and geography cannot explain the observed differences in agricultural productivity.

A distinctive feature of agriculture is that it is an activity that takes place across space, using location-specific inputs such as soil quality, climate conditions, and terrain topography. These inputs could matter not only for what yield is obtained for any crop cultivated, but also for what crops are ultimately cultivated in each cell (Holmes and Lee, 2012) and what cells are used for

agricultural production.

We quantify the role of geography and land quality for agricultural productivity across 162 countries using an accounting framework and high-resolution gridded micro-geography data, covering the entire globe, from the Global Agro-Ecological Zones (GAEZ) project of the Food and Agricultural Organization (FAO). A land cell in the gridded data is roughly a 10 by 10 kilometer plot, which should not be confused with plots of land operated by individual farms that in developing countries would typically be less than one hectare ([Adamopoulos and Restuccia, 2014](#)). GAEZ provides land quality attributes on each land cell in the world and, more importantly, potential yields for all the main crops, including crops not necessarily produced in the cell. The data on potential yields by crop are generated by combining *cell*-specific land quality attributes with established *crop*-specific agronomic models, for a given level of water supply and cultivation inputs. Potential yields summarize how detailed geographical attributes translate into productivity by crop. That is, differences in potential yields across cells, given inputs, represent a measure of differences in geography total factor productivity.

We develop a spatial-accounting framework that allows us to aggregate up from the cell-crop level resolution to the country level. Within a country, each of a fixed number of cells can produce any of a given number of crops. However, land cells are heterogeneous with respect to their inherent suitability in producing each crop, captured in the GAEZ data by its potential yield. We show that a country's aggregate yield, the value of total output per harvested land (\$/ha), can be expressed as a weighted average of the cell-crop yields valued at common prices, where the weights are the cell-crop land shares.¹ We use this expression of the aggregate yield to construct counterfactual yields. Our main counterfactual asks, what would aggregate yields be if each country produced each crop in each cell according to its potential yield, keeping the cell-crop land shares fixed to the

¹The measure of agricultural productivity that we focus on, given the GAEZ data, is the value of agricultural output per unit of harvested land (\$/ha), also known as land productivity or yield. While cross-country differences in aggregate yields account only for a portion of agricultural labor productivity differences, yields encapsulate the role of land quality and geography.

actual data? If the disparity in aggregate yields across countries were similar under potential yields, then this counterfactual would indicate that most of the variation in aggregate actual yields is due to geography and land quality differences.

In order to focus on the natural suitability of the land, for our baseline results, we use GAEZ's crop-cell potential yields under rainfed water supply and low-input cultivation practices (rainfed low-input scenario), which assumes subsistence based farming, labor intensive techniques, and no application of fertilizers and pesticides. We find that, if the ten percent richest and poorest countries produced crops according to their production-potential yields, the rich-poor agricultural yield gap would virtually disappear, from more than 200 percent to 5 percent. The relationship between aggregate production potential and actual yields across all countries is completely flat, implying that cross-country variation in aggregate actual yields is not due to geography and land quality variation. If the ten percent most and least land-productive countries in agriculture produced according to potential yields in each cell and each produced crop, the aggregate yield gap would shrink from a whopping 790 percent to only 42 percent.

We also find that the location of crop production within a country and crop choices within cells play important but secondary roles. Spatial reallocation raises productivity in all countries, but slightly more for low income countries. If in addition, crop mix changes cell-by-cell in each country to the highest value yield, then the aggregate rich-poor yield gap reverses to a 20 percent advantage for the poorest countries. Changes in spatial and crop choices could reduce aggregate yield gaps across the most and least land productive countries by one fifth. These findings are robust to assumptions about input use and water supply conditions. In addition, the use of irrigation and complementary inputs generates aggregate potential yield gains for all countries, particularly low income countries.

Using a standard two-sector model of agriculture and the rest of the economy, we show that the gains in productivity associated with production-potential yields across countries has important quantitative implications for structural transformation, average farm size, agricultural land productivity, and income per capita.

Our paper contributes to the growing literature studying agricultural productivity differences across countries. One branch of this literature assesses the contribution of specific factors on agricultural productivity.² Another branch focuses on measuring sectoral productivity gaps (Herrendorf and Schoellman, 2015; Gollin et al., 2014b) or agricultural productivity disparities (Prasada Rao, 1993; Restuccia et al., 2008; Gollin et al., 2014a). We instead focus on measuring the role of land quality and geography for agricultural productivity gaps using spatially explicit micro-geography data.

The literature on geography and economic development emphasizes both the direct effect of geography on income (Gallup et al., 1999; Sachs, 2003), as well as its indirect effect through the institutions channel (Acemoglu et al., 2002; Easterly and Levine, 2003; Rodrik et al., 2004). While the above literature relies on country-level data and focuses on aggregate incomes, we use detailed geo-spatial data and focus on the more direct impact of geography and land quality for agriculture, a sector where the observed geographic conditions would tend to matter the most.

It is well documented that countries with higher temperatures tend to be poorer (Nordhaus, 2006; Dell et al., 2009, 2012). The literature that studies the effect of temperature and the potential impact of climate change on agricultural productivity finds that rising temperatures lead to reductions in crop yields, particularly beyond some threshold (Schlenker and Roberts, 2009; Calzadilla et al., 2013; Burke et al., 2015; Zhao et al., 2017). The economics and agronomic literatures have also studied the impact of other geographic attributes on crop yields, such as rainfall (Jayachandran, 2006; Levine and Yang, 2014), soil quality (Cassman, 1999), and topography (Kravchenko and Bullock, 2000). While these studies isolate the impact of individual attributes, our analysis accounts for all geographic attributes that impact the biological growth of crops.

An agronomic literature estimates yield gaps for particular crops and particular regions (e.g., Lobell et al., 2009; Mueller et al., 2012; Tittone and Giller, 2013; van Ittersum et al., 2013), instead we

²Examples include low intermediate input use and misallocation of labor between agriculture and non-agriculture (Restuccia et al., 2008); poor transport infrastructure (Adamopoulos, 2011); selection (Lagakos and Waugh, 2013); misallocation of factors across farms within agriculture (Adamopoulos and Restuccia, 2014), international transport frictions (Tombe, 2015); idiosyncratic risk (Donovan, 2016), among others.

assess yield gaps using a consistent methodology on a global scale at various levels of disaggregation and we use potential yields to assess the role of land quality on productivity differences across countries. Other studies focus on cross-country differences in aggregate land quality indices and their effect on agricultural productivity (Wiebe, 2003; Wiebe et al., 2000), whereas we exploit the explicit spatial nature of the micro-geography data in GAEZ using an accounting framework. Our results are consistent with this literature, especially when controlling for agricultural practices and the level of inputs in cross-country regressions of land quality on agricultural productivity.

While the GAEZ data are increasingly used in economics (e.g., Nunn and Qian, 2011; Galor and Özak, 2016; Costinot et al., 2016; Godefroy and Lewis, 2018), we exploit the GAEZ data to study the macro-level implications of land quality endowments for cross-country differences in agricultural productivity, an issue that is paramount for understanding the foundation of poverty across the world.³

The paper proceeds as follows. The next section describes the GAEZ data and provides some measures of land quality dispersion across countries. In Section 3, we outline the spatial accounting framework and describe counterfactuals. Section 4 presents the main findings and robustness analysis. In Section 5, we study the sectoral and aggregate implications from a closing of the agricultural productivity gap between rich and poor countries. We conclude in Section 6.

2 Data

We describe the details of the data and characterize differences in land-quality attributes across countries based on this data.

³See other applications in Costinot and Donaldson (2012), Costinot and Donaldson (2016), and the survey in Donaldson and Storeygard (2016) on high-resolution spatial data in economics.


2.1 Description

We use gridded micro-geography data from the Global Agro-Ecological Zones (GAEZ) project, developed by the Food and Agricultural Organization (FAO) in collaboration with the International Institute for Applied Systems Analysis (IIASA) and aggregate cross-country income data from the Penn World Table (PWTv6.3). GAEZ is a standardized framework for the characterization of climate, soil, and terrain conditions relevant for agricultural production. GAEZ combines state-of-the-art agronomic models by crop, that account for science-based biophysical growing requirements for each crop, with high resolution spatial data on geographic attributes.

The information in GAEZ is available at the 5 arc-minute resolution. To picture it, imagine superimposing a grid of about 9 million cells or pixels covering the entire world. Figure 1 displays a grid map of the Montreal and Toronto area in Canada, based on cells of different resolutions, where the pink grid represents a 5-arc min; the blue grid a 30-arc min; and the black grid a 60-arc min. While the size of each cell is constant at 5 arc-minutes in the data, it is not constant in terms of squared kilometers or hectares, as the mapping from arc-minutes to square kilometers depends on the distance from the equator (latitude). As a rough approximation the size of each cell can be described as 10×10 kilometers.

For each cell in the grid, GAEZ reports data on the following location-specific geographic attributes that are important for agricultural production: (1) soil quality, which includes depth, fertility, drainage, texture, chemical composition; (2) climate conditions, which include temperature, sunshine hours, precipitation, humidity, and wind speed; and (3) terrain and topography, which include elevation and slope. Importantly, GAEZ calculates a potential yield for a set of crops in each cell, measured as the maximum output (in tonnes) per hectare that can be attained in the cell given the crop's production requirements, the cell's characteristics, and assumptions about input levels such as water supply conditions and cultivation practices. Therefore, differences in potential yields represent measures of total factor productivity differences of cell-level geographic attributes for each

Figure 1: Grid Resolution Example Montreal-Toronto Area



Notes: The pink grid represents a 5-arc min; the blue grid a 30-arc min; and the black grid a 60-arc min.

crop given water conditions and inputs of cultivation practices. Unfortunately, the GAEZ data do not provide information on actual amounts of inputs. Given that the attainability of potential yields depends on access to inputs, we consider two input scenarios relevant for cross-country comparisons, as detailed below. Potential yields are provided for all major crops including those not actually produced in a particular cell. Note that production statistics of crops are usually measured in fresh weight, whereas GAEZ simulated potential production is measured in dry weight. We use GAEZ standard conversion factors by crop to make the two measurements equivalent.

There are two key ingredients that go into the GAEZ estimation of potential yields for each crop in each cell. First, the detailed micro-geography characteristics on soil quality, climate, and topography outlined above for that particular cell. Second, crop-specific agronomic models that reflect each crop's biophysical requirements for growth. The parameters of the agronomic models capture how a particular set of geographic conditions maps into any given crop's yield. These parameters are based on well tested field and lab experiments by agricultural research institutes, are established

in the agronomic literature, and are updated to reflect the latest state of scientific knowledge. We stress that the agronomic model parameters are not based off a regression analysis of observed choices on outputs and inputs across countries, regions, or farms, an analysis that would be subject to serious endogeneity issues.

Potential yields are reported for alternative configurations of water supply conditions and cultivation practices. Water supply conditions include: irrigated, rainfed, and total which covers both rain fed and irrigated land. There are three levels of cultivation practices which specify input intensity use and management: (a) Low level of inputs (traditional management) assumes subsistence based farming, labor intensive techniques, no application of nutrients, chemicals, and pesticides. (b) Intermediate level of inputs (improved management) assumes partly market oriented farming, improved varieties with hand tools and/or animal traction, some mechanization, medium labor intensity, use of some fertilizer, chemicals, and pesticides. (c) High level of inputs (advanced management) assumes mainly market oriented farming, high yield variety seeds, fully mechanized with low labor intensity, optimum application of nutrients, chemicals, and pesticides as well as disease and weed control. The idea is that the resulting crop yield in each cell would depend not only on the “endowment” of land quality and geography but also on the set of complementary inputs applied by the farmer. GAEZ also reports potential yields for a baseline mixed input scenario covering both irrigated and rainfed land, and assuming a mixed level of inputs, which applies high inputs on the best quality land, intermediate inputs on moderately suitable land, and low inputs on marginal land. This classification of land suitability for agriculture is based on cell-level data on soil type, terrain-slope conditions, and climatic conditions.

In our analysis, we consider two input scenarios: (1) the low input cultivation practices scenario with rainfed water supply conditions, which is the minimum input application in GAEZ; and (2) GAEZ’s baseline scenario, which GAEZ considers as a reasonable representation of actual agricultural input application management. Throughout the paper, we refer to scenario (1) as “Low inputs,” and scenario (2) as “Mixed inputs.” Note that under each scenario we keep water supply

and cultivation practice conditions constant across all cells and all countries. This allows for a consistent quantification of potential land productivity around the world.

Potential yields in GAEZ are calculated for both average historical climate conditions (with the baseline reference period being 1960-1990), individual historical years 1901-2009, as well as projected future climate conditions based on a number of climate models. In our analysis, we use potential yields based on the average historical climate conditions, as they iron-out year-to-year idiosyncratic weather shocks.

The GAEZ database also provides at the 5 arc-minute resolution, for the year 2000, estimated data on crop choice, actual production, harvested area, and actual yield, i.e., tonnes of production per hectare (tonns/ha) by crop. The actual production data for each cell are estimated using a flexible iterative rebalancing methodology that sequentially downscales aggregate and regional agricultural production statistics (Appendix B). The actual production data at the cell level are available for all major crops. In addition, the database contains land cover data that classify land in each cell in terms of urban, cultivated, forest, grassland and woodland, water bodies, and other uses.

The data set we work with has global coverage, consisting of 162 countries. In 2000, the countries in our sample account for 87 percent of the world production of cereal in terms of acreage and 81 percent of the value of crop production.⁴ The count of grid cells (pixels) per country varies widely, from as low as 5 (Antigua and Barbuda) to as high as 421,168 (Russia). The median country in our data set consists of 2,827 cells. A complete list of the countries in our data set, along with their cell counts, and their GDP per capita (from the Penn World Table) are provided in Table A.1 in the Appendix. Our analysis focuses on 18 main crops and commodity groups, that cover the majority of produced crops across the world.⁵ In 2000, the crops that are covered in our GAEZ analysis account for 83 percent of the harvested area and 60 percent of the value of production in total crops,

⁴Based on data from FAOSTAT, available through <http://www.fao.org/faostat/en/#data>.

⁵The crops in our data set are: wheat, rice, maize, sorghum, millet, other cereals (barley, rye, oat, and other minor cereals), tubers (white potato, sweet potato), roots (cassava, yam and cocoyam), sugarcane, sugarbeets, pulses (chickpea, cowpea, dry pea, grams, pigeon-pea), soybean, sunflower, rapeseed, groundnut, oilpalm, olive, cotton.

across all countries in the world. The coverage of the crops in our analysis for the lower income countries is very similar, both in area and production (FAOSTAT). While crop production is only one component of the agricultural economy, our focus on crop production and yields is dictated by the available geo-spatial data from GAEZ. In addition, productivity can be more accurately measured and compared across countries for crops than for livestock, and land quality is a more prominent issue for crops than for livestock. Nevertheless, the crop yield is strongly correlated with the overall agricultural yield that includes all forms of agricultural production. For instance, using country-level data from FAOSTAT, the correlation between the log value of agricultural crop output per hectare of cropland and the log value of total (including livestock) agricultural output per hectare of agricultural land is 0.83, both measured in constant 2014-16 US\$.

GAEZ provides the information for each variable in raster (grid cell) files, which we work with in ArcGIS. To aggregate cell-level information to administrative units, such as regions, provinces, and countries, we use shape files from the World Borders data set of “Thematic Mapping.”⁶

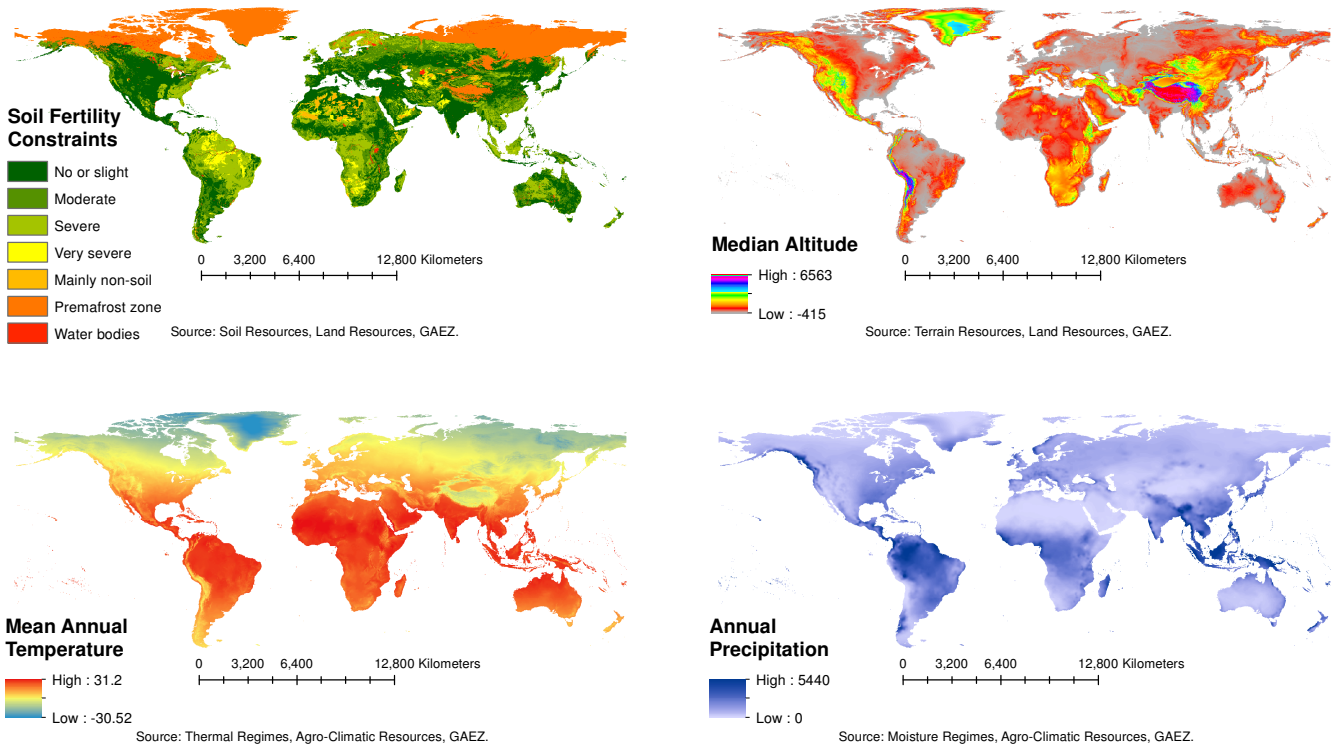
2.2 Land Characteristics across the World

We use the micro-geography data from GAEZ to illustrate the diversity of some key land quality and geographic characteristics across the world. We illustrate these characteristics in a set of maps constructed using ArcGIS for all the cells at the 5-arc minute resolution in Figure 2. The soil fertility constraint classifies the soil according to its nutrient availability, which captures soil properties such as texture (e.g., clay, silt, sand), organic carbon content, acidity (pH), and the sum of sodium, calcium, magnesium and potassium. Nutrient availability is an important indicator of soil fertility, particularly in environments with low application of intermediate inputs. The classification determines how nutrient constrained the soil in each cell, ranging from no/slightly constrained (index value 1) to very severely constrained (index value 4). Except for the premafrost zones in the north, there is considerable variation in the soil constraint around the globe, that transcends country

⁶Available through http://thematicmapping.org/downloads/world_borders.php.

borders. We also document the median altitude (in meters), the mean temperature (in degrees Celsius), and annual precipitation (in milimeters) in each cell for the whole world. Altitude is an important indicator of terrain suitability for agricultural production, as it affects solar radiation, oxygen availability as well as temperatures and moisture. The altitude varies substantially across the world, with a high of 6500 meters to a low of -415 meters. Temperature is an example of an indicator of thermal regimes, while rainfall is an example of an indicator of moisture regimes. Both thermal and moisture regimes are important measures of agro-climatic conditions and serve as key inputs into the GAEZ methodology in constructing crop-specific potential yields by cell. The maps in Figure 2 illustrate the wide diversity in these agro-climatic conditions across the world.

Figure 2: Geographic Attributes around the World



It should not be surprising that there is such wide variation in land quality characteristics across the world. Even within narrow geographic regions some locations are naturally advantaged in terms of one or more characteristics, while others are naturally disadvantaged. The importance of

a naturally advantageous geographic environment for agricultural production in a specific location is ubiquitous. However, agricultural productivity differences between the developed and developing world are often framed at the country level. As a result, we are interested in whether the land quality characteristics vary systematically across the most and least developed countries.

We examine cross-country variation in land quality attributes in Table 1 according to mean soil, terrain, and climate conditions. For soil quality conditions we report “fertility,” which captures nutrient availability and is measured as an index from 1 (unconstrained) to 4 (constrained), and “depth,” which captures rooting conditions, and is also measured as a 1 to 4 index. The terrain conditions we report are “slope,” measured as an index between 0 and 100, and “altitude” which measures mean elevation in meters. The slope of a plot is important, as it can affect for example the farming practices employed (standard mechanization can be difficult on steep irregular slopes) and the extent of topsoil erosion. The climatic conditions we report are “temperature,” measured in degrees Celsius and “precipitation,” measured in millimetres. We report the averages of these attributes across countries in the richest and poorest deciles of the 162 countries in our sample and the averages over the countries with the top and bottom deciles of the cross-country distribution of each attribute.

The main finding from Table 1 is that there is substantial variation in land quality and geographic characteristics around the globe, but that this variation is considerably more compressed across rich and poor countries. In particular, the dispersion in mean attributes, measured as the log difference between rich and poor countries for each attribute, accounts for less than one quarter of the dispersion across the world, and for most attributes less than 10 percent. This finding suggests that the mean attribute differences are not systematic across the income distribution and as a result the unconditional cross-country variation dwarfs the rich-poor variation in each of the attributes.

While we find some variation between rich and poor countries in terms of soil, terrain, and climate attributes, what matters for aggregate agricultural productivity is how differences in geographical attributes translate into productivity differences across countries. Moreover, agricultural productiv-

Table 1: Differences in Mean Geographical Attributes

	Rich 10%	Poor 10%	Top 10%	Bottom 10%
Soil Quality				
Fertility (1-4 index)	2.37	2.19	3.32	1.10
Depth (1-4 index)	2.19	1.93	3.40	1.07
Terrain Conditions				
Slope (0-100 index)	72.0	78.5	96.1	37.9
Altitude (meters)	342.8	824.0	1799.4	53.97
Climate Conditions				
Temperature ($^{\circ}C$)	12.3	23.2	27.5	2.6
Precipitation (mm)	899.6	1074.9	2515.8	123.4

Notes: Top and bottom 10% refer to the average of the highest and lowest decile in the cross-country distribution for each attribute, whereas Rich and Poor 10% refer to the average attributes of the richest and poorest decile countries in terms of real GDP per capita.

ity is the result of all geographical conditions combined and differences in a single attribute may not matter as much. For this reason, in the next section we work with potential yields by cell and crop, as a summary measure of how dispersion in geographical attributes translates into productivity differences.

3 Accounting Framework

We develop a spatial accounting framework to study the role of land quality and geographic characteristics on agricultural productivity.

3.1 The Primitives

We consider a world comprised of a fixed number of administrative units indexed by $u \in \mathcal{U} \equiv \{1, 2, \dots, U\}$. These units are countries in our analysis but in general could be lower administrative units within a country such as provinces, states, or counties. Each administrative unit u comprises a finite number G_u of grid cells (or pixels) of fixed size. We index grid cells by $g \in \mathcal{G}_u \equiv \{1, 2, \dots, G_u\}$ and aggregate cells to the country level using a mapping of cells to administrative boundaries in ArcGIS. Each grid cell can produce any of C crops, indexed by $c \in \mathcal{C} \equiv \{1, 2, \dots, C\}$.

Cells are heterogeneous with respect to their land characteristics and as a result differ in the productivity of the land across crops. In particular, a key object reported in the GAEZ data is the *potential* yield or land productivity (tonnes/ha) of each cell for crop c . We denote the potential yield of crop c in grid cell g in unit u by \hat{z}_{gu}^c . Note that for each cell g in unit u there are C such numbers, each of which reflects the inherent productivity of that cell in producing crop c under a given input scenario (water conditions and cultivation practices). In other words, the potential yield of crop c in each cell represents the maximum attainable output for that crop given inputs and as such, variation in potential yields with constant inputs reflects variation in total factor productivity of land characteristics. We note that this variation does depend on input conditions and, as such, in our quantitative analysis we consider two inputs scenarios, which consistently characterize potential yields across countries.

In practice, the land in each cell can be used for crop production or some other activity (could be agricultural such as raising livestock or non-agricultural, or some other land cover category). If a portion of the land in a cell is used for crop production, it may produce one or several specific crops which may differ from the crops in which the cell has the highest potential yield. We denote by y_{gu}^c the physical output (in tonnes) of crop c and by ℓ_{gu}^c the amount of harvested land (in hectares) of crop c for any cell g and unit u . We denote by z_{gu}^c the actual yield of crop c which is just the ratio of physical output to land (tonnes/ha) in each cell but note that our analysis does not rely

on this disaggregated actual yield nor on actual output by cell from GAEZ. Only the aggregated actual values at the country level matter for the analysis, which match the aggregate actual data by construction. For the purpose of aggregation, in any unit u and cell g , we set the amount of output and land used to zero if there is no production of a given crop c .

Similarly, for the purpose of aggregation of different crops in a location, we denote by p^c the price of each crop (in \$) which we treat as common across space and countries. Note also that the size of each vector is $C \times 1$, corresponding to the total number of crops in the GAEZ project which is 18 crops. In each cell g , all the vectors have non-zero elements only for the crops actually produced. The only vectors that have all non-zero elements for every crop are the potential yield and the price. The potential yield vector is specific to each cell g and unit u .

3.2 Aggregate Variables

We denote with upper case letters aggregate variables at the country level. We denote by L_u the amount of land used in agricultural production in country u (in hectares), given by,

$$L_u = \sum_{c \in \mathcal{C}} \sum_{g \in G_u} \ell_{gu}^c.$$

We denote by Y_u the total value of agricultural output produced (\$), given by,

$$Y_u = \sum_{c \in \mathcal{C}} \sum_{g \in G_u} p^c y_{gu}^c.$$

Given these aggregates, we define the aggregate actual yield Z_u by the ratio of aggregate value output to land (\$/ha), that is,

$$Z_u = \frac{Y_u}{L_u} = \frac{\sum_{c \in \mathcal{C}} \sum_{g \in G_u} p^c z_{gu}^c \ell_{gu}^c}{L_u} = \sum_{c \in \mathcal{C}} \sum_{g \in G_u} p^c z_{gu}^c \frac{\ell_{gu}^c}{L_u}. \quad (1)$$

The aggregate yield is a weighted average of the yields in every crop and location in a given country. Equation (1) is key in our accounting analysis as it provides the basis for assessing the role of geography and land quality on agricultural land productivity across countries. Note also that by construction of the disaggregated actual data in GAEZ, the country-level aggregate *actual* yield Y_u/L_u is consistent with aggregate actual data on output and land. The spatially disaggregated data on actual output and yields are not relevant for our country-level analysis, only the disaggregated land-use data ℓ_{gu}^c/L_u as we discuss below.

3.3 Counterfactuals

We construct a set of counterfactuals on the aggregate yield for each country u by exploiting the set of potential yields by crop at the cell level g and the spatial distribution of land use by crop across cells. All the counterfactuals involve producing crops at potential and in some cases reallocating across space and crops. Because cell-specific potential yields depend on input conditions, we construct these counterfactuals for the rainfed low input scenario as our baseline, as this involves the least human intervention, and more closely captures the natural suitability of the land. We compare our baseline results to those under the mixed input scenario.

Production potential. We assess the impact on the aggregate yield in the case of countries producing at the potential yield for each crop and each cell. We compute this counterfactual by simply using in equation (1) the potential yield \hat{z}_{gu}^c for each crop in each cell:

$$Z_u^{cf} = \sum_{c \in \mathcal{C}} \sum_{g \in G_u} p^c \hat{z}_{gu}^c \frac{\ell_{gu}^c}{L_u}.$$

In this counterfactual only the yield changes, while the weights represented by the share of cultivated land of a crop in each location are kept to the actual values ℓ_{gu}^c/L_u . Note that the construction of the aggregate production-potential yield depends only on the disaggregated data on potential yields and land used by crop.

If cross-country differences in the aggregate production-potential yield are similar to the aggregate actual yields, then production potential at the crop/cell level is an important determinant of the aggregate actual yield differences. Instead, if cross-country differences in aggregate productivity-potential yields are negligible, then geography and land quality are not important determinants of actual yield gaps across countries. A possible issue with this counterfactual may be that potential yields involve the application of inputs that are simply not available in the country. The GAEZ data does not allow us to separate from actual data at the cell or aggregate level the input conditions. As noted above however in our baseline results we consider the rainfed low-input scenario which captures the application of minimum inputs, that should be attainable by all countries and hence mitigates this concern.

Spatial potential. We assess the extent to which reallocation of agricultural production of the different crops to the most productive locations across space can increase aggregate output. This counterfactual combines production potential with a reallocation of crops across space to the most suitable locations. In particular, we reallocate production so that each crop is produced in the cells where it realizes the highest potential yields, keeping constant the amount of harvested land for that crop in the country to the actual level, i.e., $L_u^c = \sum_{g \in G_u} \ell_{gu}^c$. This allocation problem is non-trivial. Some cells within a country may exhibit higher potential productivity for all crops, while the amount of land that can be allocated to a given crop is limited. We reallocate the production of crops to cells to maximize total constant-price value output based on potential yields of the different crops, i.e., to where the relative potential return for each crop is the highest. Formally, this involves solving a large-scale linear programming problem for each country, given by,

$$\max_{\{\ell_{gu}^c\}} \sum_{c \in \mathcal{C}} \sum_{g \in G_u} p^c \hat{z}_{gu}^c \ell_{gu}^c, \quad (2)$$

subject to

$$\sum_{c \in \mathcal{C}} \ell_{gu}^c \leq L_{gu}, \quad g = 1, 2, \dots, G_u; \quad (3)$$

$$\sum_{g \in G_u} \ell_{gu}^c \leq L_u^c, \quad c = 1, 2, \dots, C; \quad (4)$$

$$\ell_{gu}^c \geq 0, \quad g = 1, 2, \dots, G_u; \quad c = 1, 2, \dots, C. \quad (5)$$

The objective is to maximize the total value of output across all cells and crops, subject to three sets of constraints. The first set of constraints restricts that land allocated to the production of the different crops cannot exceed what is available in each cell. The second set of constraints indicates that land allocated to crop c over all cells cannot exceed the total in the data. The third set of constraints allows for the possibility that not all crops are produced in all cells. Note again that the constraints only involve land allocations as the application of inputs in the GAEZ data is embedded in the different input scenarios we consider, and cannot be separated out. The low input scenario however maps more directly to the allocation problem in equations (2) to (5). Given that our focus is on the role of land quality for production, we abstract from demand. This prevents us from drawing implications about welfare and analyzing possible changes in local relative prices.

Total potential. We assess the extent to which countries may not be producing the highest yielding mix of crops in each location given their land endowment characteristics. This counterfactual involves computing the aggregate potential yield in each country by picking the crop in each location that maximizes the total value of output. Formally, we solve for ℓ_{gu} in equation (2) subject to only the constraints in equations (3) and (5). This counterfactual involves production potential, reallocation of crops across space, and changes in crop choices in order to maximize the aggregate value of agricultural output (\$). It is the allocation that generates the maximum attainable value of output in each country given the total amount of land, and the set of potential yields by cell and crop. The difference between this counterfactual and the production potential counterfactual represents the contribution to the aggregate yield of crop-mix choices and the spatial reallocation of production, whereas the difference with the spatial reallocation is the contribution of crop choice changes to the aggregate yield.

4 Results

We present the results for each counterfactual under our baseline rainfed low-input scenario for potential yields by cell and crop and show how aggregate actual and potential yields vary across countries. We then assess whether aggregate potential yields can account for agricultural land productivity (aggregate actual yield) across countries. We also present results for potential yields under the mixed input scenario and other robustness.

4.1 Baseline Results


We calculate aggregate output per hectare (aggregate actual yield) using FAO international crop prices (Geary-Khamis dollars per tonne) for the year 2000. There are substantial differences in aggregate actual yields across countries. The ratio between the ten percent richest and poorest countries in terms of GDP per capita is a factor of 3.1-fold. This dispersion in aggregate yields is consistent with cross-country estimates using micro-data sources (Gollin et al., 2014a). The aggregate actual yield varies systematically with the level of development (GDP per capita), with a correlation in logs of 0.58.

To what extent are aggregate yield differences across countries the result of differences in land quality and geography? We now address this question using our spatial accounting framework.

Production potential. We calculate the aggregate production-potential yield using equation (1) and the potential yield for each cultivated crop in each cell under the rainfed low-input scenario, keeping the crop-cell land allocation fixed to the actual one. Given that the low-input scenario has rainfed water conditions and the lowest application of inputs, it more closely captures the productivity afforded by the natural geographic and land quality endowments of each country. Figure 3 documents the aggregate production-potential yield across countries along with the aggregate actual yield. While there is substantial variation in aggregate potential yields across countries, the ratio

of the top to bottom deciles in the production-potential yield distribution is a factor of 6, the differences are not systematically related to the level of development. For instance, the disparity in the production-potential aggregate yield between Tanzania and Eritrea, two low-income countries, is roughly the same (around 3-fold) as the disparity between Belgium and Austria, two high-income countries. Unlike the aggregate actual yields, Figure 3 illustrates that there is actually a slight negative relationship between potential aggregate yields and GDP per capita, with a correlation in logs of -0.19 . Egypt and countries of the Arabian peninsula stand out as outliers due to large desert areas, which are particularly arid under pure rainfed water conditions and low input application.

Figure 3: Aggregate Actual and Production-Potential Yield across Countries



Notes: Aggregate production-potential yield under the rainfed low-input scenario.

In Panel A, Table 2, we report the production-potential yield for the weighted average of each of the richest and poorest decile countries in terms of income per capita. The results are striking. If countries produced the crops they are producing in the cells they are actually producing them but according to their potential yields, the aggregate yield disparity between rich and poor countries would drop from the actual 3.14-fold to only 1.05-fold, that is the productivity disparity would

Table 2: Counterfactual Production-Potential Yield

Panel A: All Crops (<i>country obs.</i> = 162)			
	Actual Yield	Potential Yield	Ratio
Rich 10%	739.5	237.2	0.32
Poor 10%	235.5	225.7	0.96
Ratio	3.14	1.05	1/2.99
Panel B: Wheat (<i>country obs.</i> = 110)			
	Actual Yield	Potential Yield	Ratio
Rich 10%	2.71	1.04	0.38
Poor 10%	1.07	0.66	0.62
Ratio	2.53	1.58	1/1.61
Panel C: Rice (<i>country obs.</i> = 104)			
	Actual Yield	Potential Yield	Ratio
Rich 10%	6.64	0.88	0.13
Poor 10%	1.30	0.89	0.68
Ratio	5.10	0.99	1/5.13
Panel D: Maize (<i>country obs.</i> = 142)			
	Actual Yield	Potential Yield	Ratio
Rich 10%	8.56	2.10	0.25
Poor 10%	1.31	1.31	1.00
Ratio	6.52	1.61	1/4.06

Notes: Rich and Poor refer to the weighted average of the highest and lowest decile of the real GDP per capita distribution in 2000 (PWT 6.3). Aggregate actual and potential yields are measured as total value output per hectare in international prices (GK \$/ha). Actual and potential yields by crop are measured as tonnes per hectare. The production-potential yield for each country is constructed by aggregating up from the GAEZ pixel-level information at the 5 arc-minute resolution under the low input scenario. Low inputs assumes rainfed water supply and low-input cultivation practices.

virtually disappear. In other words, with no complementary inputs and rainfed water conditions, rich countries would only attain a 5 percent higher aggregate yield. Note, that if countries were producing according to their natural endowments without complementary inputs, yields would be lower than actual yields for most countries, particularly for the high-income countries.

Our findings are consistent with an earlier literature on the role of aggregate measures of land quality for agricultural productivity across countries (Wiebe, 2003). For instance, Wiebe et al. (2000) find that good soils and climate are associated with a 13 percent increase in agricultural productivity relative to poor soil and climate. Using this data, we find that when controlling for (log) inputs such as agricultural labor, tractors, and fertilizer, the effect on agricultural productivity becomes statistically insignificant. Despite the limitations of aggregate measures of land quality, our findings with the geospatial data from GAEZ and the accounting framework, are consistent with this earlier literature, in the lack of a systematic relationship between land quality and agricultural land productivity, after controlling for other inputs.

We have used a common set of crop prices to aggregate yields in all locations and countries, however, the conclusions from the production-potential counterfactual remain when focusing on individual crops for which we can use a physical measure of land productivity. Panels B to D in Table 2 report the production-potential counterfactual for each of the three most representative crops produced across the world: wheat, rice, and maize. In each case, the yield is measured as crop output in tonnes per unit of land, a physical measure of productivity that does not require prices for aggregation. The rich-poor disparity in the actual yield differs across crops: 2.53-fold for wheat, 5.10-fold for rice, and 6.52-fold for maize. Producing these crops according to the low-input potential yields would reduce the rich-poor disparity to 1.58, 0.99, and 1.61-fold for wheat, rice, and maize. Despite differences across individual crops, the main takeaway from the production-potential counterfactual is that the disparity is substantially reduced when producing according to natural endowments, implying a limited role of land quality and geography for differences in agricultural productivity across countries.

We note that the aggregate production-potential yield summarizes a great deal of heterogeneity in land quality between cells in a country. However, we find no systematic relationship between the within-country dispersion in potential yields and the level of development, see Appendix C, Figure C.1 for a documentation of this fact.

Spatial potential. A fact of poor and developing countries is the prevalence of large rural populations, often operating at subsistence levels and facing poor infrastructure, conditions that may lead farmers to produce in locations not necessarily suitable for agricultural production (Adamopoulos, 2011; Gollin and Rogerson, 2014; Adamopoulos and Restuccia, 2014). We assess the relevance of spatial reallocation by calculating the aggregate potential yield that would result from reallocating production of crops across cultivated cells according to where they exhibit the highest relative yield value in the country, holding constant the total amount of land allocated to each crop in the country. Table 3 reports the results of this counterfactual in columns two and three (column one repeats the production-potential). Spatial reallocation has a positive effect on agricultural output for both rich and poor countries, but relatively more for poor countries. Under spatial reallocation relative to the production-potential, the aggregate yield for the poorest countries increases on average 36 percent, whereas for the richest the average increase is 22 percent. This implies a further decline in the rich-poor yield gap to 0.94, relative to the production-potential counterfactual, and a slight reversal of the rich-poor yield gap.

We also note that if spatial reallocation is guided by actual-yield differences across cells rather than potential-yield differences in our spatial counterfactual, there is more of a decline in the dispersion between rich and poor countries, from 3.14-fold in the aggregate actual yield to 2.3-fold. Nevertheless, even though there is more role for spatial reallocation under actual yields, the reduction in the disparity is much lower than producing at potential for each crop and cell which implies an aggregate production potential yield ratio of 1.05-fold.

Total potential. We now assess the effect of crop reallocation by calculating the aggregate potential yield in each country when the highest value-yielding crop is produced in every cell, holding the

Table 3: Counterfactual Aggregate Spatial and Total Potential Yields

(country obs. = 162)					
Aggregate potential yield - Low Inputs					
	Production	Spatial	Spatial/ Production	Total	Total/ Spatial
Rich 10%	237.2	288.2	1.22	363.9	1.26
Poor 10%	225.7	307.6	1.36	469.0	1.53
Ratio	1.05	0.94	1/1.11	0.78	1/1.21

Notes: Rich and Poor refer to the average of the highest and lowest decile in the real GDP per capita distribution in 2000 (PWT 6.3). “Production” refers to the production-potential counterfactual yield, “spatial” to the spatial-potential counterfactual yield, and “total” to the total-potential counterfactual yield. All yields are measured as total value output per hectare in international prices (GK \$/ha), aggregating from the GAEZ pixel-level information at the 5 arc-minute resolution under the low input scenario.

amount of land in each cell constant. This counterfactual reflects the maximum aggregate value potential yield that can be achieved via production potential across cells and crops within cells. Table 3 reports the results of this total-potential counterfactual, in columns four and five. If countries shift their crop mix to the highest yielding crops, cell-by-cell, then the aggregate yield disparity would drop between rich and poor countries from the actual 3.14-fold to 0.78-fold. Adjusting the crop mix cell-by-cell to the most suitable given each cell’s geographic characteristics, poor countries would be 22 percent more productive than the rich countries. This occurs because relative to the spatial potential, the total potential increases yields in poor countries by 53 percent, double the increase in rich countries (26 percent).

Overall, while spatial and crop reallocation both contribute to reduce the dispersion in land productivity between rich and poor countries, most of the reduction in the aggregate yield gap occurs when countries produce each crop in each cell at potential. We note that in our accounting, spatial and crop reallocation measure the increase in agricultural output associated with given geography productivity, but clearly these gains may not be efficient when taking into account other uses of land, cost and demand factors, and potential changes in relative prices. Similarly, the spatial- and

total-potential counterfactuals abstract from complementarity in production across crops, and crop rotation associated with best practices of land management, among others.

4.2 Accounting for Actual Yields


We now examine to what extent land quality and geographic differences across countries, captured by the low-input aggregate potential yields, can account for agricultural land productivity differences (aggregate actual yields).

If aggregate potential yields were roughly similar to aggregate actual yields across countries, then when plotting these variables, countries should fall around the 45 degree line. We find in contrast that the relationship between aggregate potential and actual yields is fairly flat. Figure 4 displays the aggregate production-potential yield for the rainfed low-input scenario against the aggregate actual yield. There is a weak relationship between the production-potential yield and actual yield, in fact the correlation is slightly negative, indicating a weak role for geography and land quality in accounting for agricultural land productivity differences.

In Table 4 we rank countries according to their aggregate actual yield, and show the results of the three counterfactuals for the top and bottom deciles of the actual yield distribution (rather than the rank by real GDP per capita in Tables 2 and 3). The first column shows the aggregate actual yield for the most and least land productive countries, with a disparity in actual yields of 8.91-fold. Under the production-potential counterfactual (second column), where all countries produce according to their potential yields cell-by-cell (holding constant the land and crop allocation), the disparity between the most and least productive countries declines to 1.42-fold. If countries produced according to their raw natural endowment potential the bulk of the staggering 791 percent land productivity differences would disappear. We conclude that the most land productive countries' agricultural productivity advantage is not primarily driven by favourable natural geographic endowments.

If in addition, crop production could be spatially reallocated within a country (keeping the total

Figure 4: Aggregate Production-Potential vs. Actual Yield across Countries



Notes: Aggregate production-potential yield under rainfed low-input scenario.

Table 4: Counterfactual Potential Yields by Agricultural Productivity

	<i>(country obs. = 162)</i>					
	Aggregate potential yield - Low Inputs					
	Actual	Production	Spatial	Spatial/ Production	Total	Total/ Spatial
Top 10%	1265.1	208.1	277.3	1.33	365.8	1.32
Bottom 10%	142.0	146.9	212.4	1.45	424.3	2.00
Ratio	8.91	1.42	1.31	1/1.09	0.86	1/1.52

Notes: Top and Bottom refer to the average of the highest and lowest decile of countries in terms of aggregate actual yield (agricultural productivity). “Production,” “Spatial,” and “Total” refer to the production-potential, spatial-potential, and total-potential counterfactual yields. Aggregate actual and all counterfactual potential yields are measured as total value output per hectare in international prices (GK \$/ha). Production potential yields are constructed by aggregating the GAEZ pixel-level information at the 5 arc-minute resolution under the low-input scenario.

amount of land to each crop in a country fixed), then the yield disparity between the most and least productive countries would drop further to 1.31-fold, implying that the relatively least productive countries would benefit more from such a reallocation (columns three and four in Table 4). Under the total potential counterfactual (columns five and six), if countries shift their crop mix to the highest value yielding crops, cell-by-cell, then the aggregate yield disparity between the top and bottom deciles would drop further to 0.86. In other words, by adjusting the crop mix to the most suitable given their geographic characteristics, unproductive countries would be 14 percent more productive than the top productive countries.

The potential yield gaps, under the different counterfactuals, show that the agricultural productivity differences are eliminated, in fact reversed, implying that land quality does not play a key role in explaining the actual productivity differences. We decompose the overall reduction in the yield gap between the most and least land productive countries, into the contributions of within cell-crop productivity, spatial reallocation, and crop choice, as implied by our counterfactuals. The disparity in the aggregate yield between the top and bottom 10% of countries in land productivity drops from 8.91-fold in the actual yield to 1.42-fold in the production potential counterfactual, to 1.31-fold in the spatial counterfactual, and to 0.86 in the total counterfactual using the rain and low input scenario. We can decompose the contribution of each factor (production, spatial, and total) to the decline in agricultural productivity disparity as follows:

$$\underbrace{8.91 \times \overbrace{0.16}^{\text{production}}}_{=1.42} \times \overbrace{0.92}^{\text{spatial}} \times \overbrace{0.66}^{\text{total}} = 0.86.$$

=1.31

This implies that production-potential contributes 79 percent ($\log(0.16)/\log(0.097)$) to the decline in the productivity ratio, while the spatial reallocation of production accounts for a small 4 percent and crop reallocation the remaining 17 percent.

4.3 Potential Gains under Mixed Inputs

In our analysis so far we have focused on aggregate potential yields derived under the rainfed low-input scenario from GAEZ. The main takeaway is that these potential yields do not vary systematically across countries. However, the potential yields based on the low input scenario are typically lower than the actual yields for most countries. This is expected given that they mostly capture the raw natural land endowments of countries under subsistence agricultural practices. Are there potential yield gains that countries can reap given their existing land quality and geographic characteristics? To answer this question, we use by-cell and by-crop potential yields from GAEZ under the mixed input scenario, which includes both rainfed and irrigated land, as well as a mixed level of cultivation practices and input application, that applies the highest level of inputs to the best land and the lowest level of inputs to the marginal land. GAEZ considers this as their baseline scenario as it more realistically represents actual input application approaches across the world.

We compute the aggregate counterfactual experiments under the mixed input scenario. In Table 5 we present the production-potential, spatial-potential, and total-potential yields for the 10 percent richest and poorest countries. The first column reports the aggregate actual yields. Similarly to the low-input scenario, the production-potential in column two, if countries produced their crops in the current locations but according to the mixed-input potential yields, the rich-poor disparity would drop from the actual 3.14-fold to 1.05-fold, a conclusion that is remarkably robust to the input scenario assumed. However, as column three shows, what is different under the mixed-input scenario is that aggregate production-potential yields are higher than actual yields for both rich and poor countries. Based on the mixed input production-potential yields, the potential productivity gains for the rich countries are 65 percent, whereas for the poor countries 393 percent. This implies that, conditional on their land quality, there is considerable untapped potential for poor countries, and that improvements in non-land quality factors (cultivation practices, application of complementary inputs, farm size and organization, managerial operation, among others), rather than land quality, would allow them to realize it.

Table 5: Counterfactual Potential Yields with Mixed Inputs

	(<i>country obs.</i> = 162)					
	Aggregate potential yield — Mixed inputs					
	Production	Production/ Actual	Spatial	Spatial/ Production	Total	Total/ Spatial
Rich 10%	1220.0	1.65	1446.0	1.19	2498.3	1.73
Poor 10%	1160.6	4.93	1361.1	1.17	3254.7	2.39
Ratio	1.05	1/2.99	1.06	1.02	0.77	1/1.38


Notes: Rich and Poor refer to the average of the highest and lowest decile in the real GDP per capita distribution in 2000 (PWT 6.3). “Production” refers to the production-potential counterfactual yield, “spatial” to the spatial-potential counterfactual yield, and “total” to the total-potential counterfactual yield. All yields are measured as total value output per hectare in international prices (GK \$/ha), aggregating from the GAEZ pixel-level information at the 5 arc-minute resolution under the mixed input scenario.

Figure 5 displays the production-potential yield and the actual yield for all countries in our sample. Despite the substantial variation in potential yields across countries, they do not systematically vary with the level of income per capita, as the aggregate actual yields do. Furthermore, potential yields lie above the actual yields for all countries, particularly the lower income ones. That is, conditioning on the set of crops each country produces on each plot, developing countries produce much further away from their potential than developed countries.

Spatial reallocation (columns four and five in Table 5) has a positive effect on agricultural output, but the magnitude of the effect is similar among rich and poor countries, and thus does not further reduce the disparity in agricultural productivity across countries beyond the production potential effect. Columns six and seven show the results of the total-potential counterfactual. If in addition, countries shift their crop mix to the highest yielding crops, cell-by-cell, then the aggregate yield disparity between rich and poor countries would drop to 0.77-fold with mixed inputs, very similar to the 0.78 ratio we found under low inputs. This counterfactual points to poor countries producing systematically lower yielding crops given their internal land quality characteristics.

The conclusions about the lack of a systematic correlation between the potential aggregate yield

Figure 5: Aggregate Actual and Production-Potential Yield across Countries



Notes: Aggregate production-potential yield under total-water (rainfed and irrigated) and mixed-input scenario.

and GDP per capita, and the reversal of the gap under total-potential yield remain intact between the low and mixed input scenarios. An important element in these cross-country comparisons is that the assumptions on inputs are kept the same in all countries, so the differences in the potential yields for each crop and location reflect variation in the geographical attributes of the land in each location. Interestingly, the similarity in the rich-poor potential yield disparities under the low- and mixed-input scenarios suggest that geographical endowments in poor countries are not less conducive to the use of certain inputs.

Our finding that gaps in potential-to-actual yields under the mixed input scenario are higher in developing than developed countries is consistent with findings in agronomic studies. For example, [Van Ittersum et al. \(2016\)](#) find that the average potential-to-actual yield gap across 10 Sub-Saharan African countries for the main cereals is a 5, while [Schils et al. \(2018\)](#) find that for Northern European countries, the yield gap in cereals is as low as 1.10-1.20.

4.4 Robustness

We examine the sensitivity of our main results to: (a) aggregating crops by weighing them according to their caloric content rather than using FAO international prices and (b) weighing produced crops within cells equally, rather than using the disaggregate cell-level land allocations by crop from GAEZ.

Caloric content of crops. In developing countries, a large proportion of farmers consume most of the output they produce as the amount of output produced is close to subsistence levels. For these farmers, the caloric intake from the different crops may be more relevant than international prices for the evaluation of the different crops. To examine the robustness of our main results, we consider an alternative to aggregating crops at the country level (in actual and potential yields) that uses instead the crop’s caloric content. As in Galor and Özak (2016, 2015) the caloric content of each crop, measured in kilo calories (kcal) per 100g, is obtained from the USDA’s *National Nutrient Database for Standard Reference* (Release 28). Table 6 reports the aggregate actual, the production-potential and the total-potential counterfactual yields (thousands kcal/ha) under the rainfed low-input scenario, for rich and poor countries.

Table 6: Aggregate Potential Yields using Caloric Weights

	USDA Calorie “Prices” of Crops (000s of kcal)				
	<i>(country obs. = 162)</i>				
	Actual Yield	Production	Potential Yields Production/ Actual	Total	Total/ Production
Rich 10%	18.20	5.64	0.31	12.36	2.19
Poor 10%	5.10	4.43	0.87	15.05	3.40
Ratio	3.57	1.27	1/2.82	0.82	1/1.55

Notes: Aggregate actual and potential yields are measured as total output in caloric energy (thousands of kcal) per hectare under the rainfed low-input scenario. Calorie “prices” (kcal per 100 grams) are obtained from USDA, *National Nutrient Database for Standard Reference* (Release 28). Rich and Poor refer to the highest and lowest decile of the distribution of real GDP per capita in 2000 (PWT 6.3).


We find that the main conclusions of our baseline results remain intact: there are substantial differences in actual yields across countries (3.14 in benchmark, 3.57 with calorie “prices”); producing existing crops according to potential yields removes the vast majority of differences (1.05 in our baseline, 1.27 with calorie “prices”); and allowing for crop mix and location to change reverses the aggregate yields difference between rich and poor countries (0.78 in our baseline, 0.82 with calorie “prices”).

Weighting crops within cells. In our counterfactual production-potential yield experiment we used the disaggregated land allocations by crop for each cell from GAEZ. The crop-cell-level potential yields are estimated from agronomic models, given the observed land quality and geographic conditions in each cell. The land allocation however is based on aggregate and regional data that GAEZ, through a downscaling methodology attributes to cells. We repeat the production-potential counterfactual assuming that each produced crop is equally weighed within each cell, instead of using the GAEZ land allocation weighting (the spatial-potential and total-potential counterfactuals do not rely of the disaggregated GAEZ land allocations since in these experiments the cell-level land allocations are endogenous and solved for as part of our linear programming problems). In Figure 6 we display the production-potential yield under the low-input scenario, with equal weighting of crops within cells, across countries. The actual yield is also presented in blue circles. Comparing, Figures 3 and 6, our conclusions about a weak relationship between potential yields and income per capita across countries remain intact. Hence, the lack of a key role for land quality does not hinge on the use of the GAEZ land weights.

5 Aggregate Implications

We study the macroeconomic implications from reducing the yield gap from its actual level to its potential level, modeled as a change in agricultural TFP. Under the mixed input scenario, if rich and poor countries produced according to their potential yields, the productivity increase for poor

Figure 6: Production-Potential Yield across Countries with Equal Crop Weighting



Notes: Aggregate production-potential yield with equal crop weight within each cell instead of the GAEZ actual land allocation.

countries would be 3 times that of rich countries. A stylized feature of the process of development is that increases in agricultural productivity lead to a reallocation of factors, in particular labor, from agriculture to the rest of the economy, such that consumption of agricultural goods per capita remains approximately constant (Gollin et al., 2002; Restuccia et al., 2008). What are the sectoral and aggregate implications of higher agricultural productivity in poor countries? To answer this question, we consider a standard quantitative sectoral model following the literature.

At each date there are two goods produced in sectors: agriculture and non-agriculture. Output in agriculture Y_a requires the inputs of land L (in fix supply) and labor N_a , $Y_a = A_a L^\theta N_a^{1-\theta}$, where A_a is TFP in agriculture.⁷ Output in non-agriculture just requires labor input, $Y_n = A_n N_n$. There is a fixed amount of labor N to be allocated between the two sectors, $N = N_a + N_n$. We assume that

⁷Note that we abstract from capital and intermediate inputs. Both of these two types of inputs are known to magnify the productivity and income implications and are discussed in detail in the literature (Restuccia et al., 2008; Adamopoulos and Restuccia, 2014).

in this economy there is a minimum amount of agricultural consumption goods per person \bar{a} and that after this minimum is satisfied individuals allocate their income to non-agricultural goods.⁸

Denoting per-capita labor in agriculture and land as n_a and l , and combining the demand for agricultural consumption goods and the production function above, we can solve for the share of employment in agriculture:

$$n_a = \left(\frac{\bar{a}}{A_a l^\theta} \right)^{1/(1-\theta)}. \quad (6)$$

We can then solve for labor productivity in agriculture $y_a = A_a l^\theta n_a^{-\theta}$ and average farm size $AFS = l/n_a$. Labor productivity in non-agriculture is simply $y_n = A_n$ and income per capita is given by $y = p y_a n_a + y_n (1 - n_a)$, where p is the relative price of agricultural goods. Without loss of generality we set $l = 1$.

We now proceed in three steps. (1) We calibrate a benchmark rich economy to data for the United States. We normalize productivity parameters, $A_a = A_n = 1$ and calibrate \bar{a} to the share of employment in agriculture in the United States of about 1.5 percent. This procedure implies $\bar{a} = 0.06$, $p = 0.40$, $y_a = 4.05$, and $y = 1.01$. (2) Given the preference parameter \bar{a} , we calibrate a poor economy, in particular we choose productivity parameters to jointly match a share of employment in agriculture of 70 percent and a poor-rich non-agriculture productivity ratio of one fifth. We obtain for the poor country agricultural TFP $A_a^p = 0.077$ and income per capita $y^p = 0.085$ using the calibrated benchmark relative price of agriculture. Note that the rich to poor ratio is 46.7-fold in agricultural labor productivity and average farm size and 12.2-fold in aggregate income per capita. (3) We study an increase in agricultural TFP of a factor of 3-fold in the poor economy consistent with our previous findings. The share of employment in agriculture in the poor economy falls from 70 percent to 13.5 percent and the disparities in agricultural labor productivity (average farm size) and income per capita fall from 46-fold to 9-fold and from 12.2-fold to 5-fold. The increase in agricultural TFP reduces the disparity in income per capita by half. The reason for this

⁸This assumption simplifies the analysis considerably since the total amount of agricultural goods consumed per capita is equal to \bar{a} , however, this is a close approximation to a calibration of more general preferences for agricultural and non-agricultural goods [Gollin et al. \(2002\)](#); [Restuccia et al. \(2008\)](#).

remarkable reduction in income disparity is that the increase in agricultural TFP reduces the share of employment in agriculture, increasing agricultural labor productivity and average farm size in the poor economy by a factor of 5.2-fold.

The main takeaway from this quantification is that improving agricultural productivity in poor countries unravels a substantial process of structural transformation that can go a long way in reducing the large disparities in sectoral and aggregate outcomes between rich and poor countries.

6 Conclusions

That land quality and geography matter for agricultural production at the micro-level is ubiquitous as argued by both agronomists (e.g, [Doorenbos and Kassam, 1979](#); [Steduto et al., 2012](#), GAEZ) and agricultural economists (e.g., [Sherlund et al., 2002](#); [Di Falco and Chavas, 2009](#); [Fuwa et al., 2007](#); [Jaenicke and Lengnick, 1999](#)). Using detailed micro-geography data, in this paper we quantify the macro-level consequences of land quality for agricultural productivity, measured as output per hectare (yield). In particular, we examine to what extent differences in agricultural yields across countries are the result of natural advantages/disadvantages or the result of economic choices. We find that land quality differences cannot justify the agricultural land productivity gaps between rich and poor countries. If farming practices were the same around the world then land quality would not be a constraint on farmers in poor countries. The majority of the actual yield differences would disappear if countries produced according to their potential, with a secondary role played by what crops are produced and where they are produced within the country.

Under improved agricultural practices for all countries, our analysis illustrates that there are large gaps between actual and potential yields in poor countries, much larger than in rich countries. The implication is that using existing technologies and improving allocations can increase agricultural productivity 5-fold. These seem like sizeable unrealized gains in productivity. One possibility is that the technologies agronomists treat as easily localized (in the calculation of potential yields) cannot

be profitably implemented everywhere in the developing world. The other possibility is that there are constraints that prevent the adoption of modern technologies and frictions that prevent markets from efficiently allocating resources in developing countries. More work is needed to understand the importance of each one of these explanations.

While a large body of recent work has been studying the constraints and frictions that impact agricultural productivity, with mounting evidence of their importance, much less work has been done on understanding the localization of agricultural technologies in developing countries. GAEZ has been following a variety of approaches for “ground-truthing” and verifying the results of their crop suitability analysis, but more needs to be done in terms of further validation. At the same time further research is needed to understand what factors constrain the choices of farmers in the developing world, preventing them from better exploiting their land and environmental endowments. Similarly, further work may be needed to better understand the role of non-biophysical factors such as infrastructure for marketing and distribution of agricultural products, constraints on capital and knowledge availability, as well as specific intermediate-input requirements for pest, disease and weed. We leave these important areas of research for future work.

References

- Acemoglu, D., Johnson, S., and Robinson, J. A. (2002). Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly journal of economics*, 117(4):1231–1294.
- Adamopoulos, T. (2011). Transportation costs, agricultural productivity, and cross-country income differences. *International Economic Review*, 52(2):489–521.
- Adamopoulos, T. and Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6):1667–97.
- Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.
- Calzadilla, A., Zhu, T., Rehdanz, K., Tol, R. S., and Ringler, C. (2013). Economywide impacts of climate change on agriculture in sub-saharan africa. *Ecological Economics*, 93:150–165.
- Caselli, F. (2005). Accounting for cross-country income differences. *Handbook of economic growth*, 1:679–741.
- Cassman, K. G. (1999). Ecological intensification of cereal production systems: yield potential, soil quality, and precision agriculture. *Proceedings of the National Academy of Sciences*, 96(11):5952–5959.
- Costinot, A. and Donaldson, D. (2012). Ricardo’s theory of comparative advantage: Old idea, new evidence. *The American Economic Review*, 102(3):453–458.
- Costinot, A. and Donaldson, D. (2016). How large are the gains from economic integration? theory and evidence from us agriculture, 1880-1997.
- Costinot, A., Donaldson, D., and Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1):205–248.

- Dell, M., Jones, B. F., and Olken, B. A. (2009). Temperature and income: reconciling new cross-sectional and panel estimates. *American Economic Review*, 99(2):198–204.
- Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.
- Di Falco, S. and Chavas, J.-P. (2009). On crop biodiversity, risk exposure, and food security in the highlands of ethiopia. *American Journal of Agricultural Economics*, 91(3):599–611.
- Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *The Journal of Economic Perspectives*, 30(4):171–198.
- Donovan, K. (2016). Agricultural risk, intermediate inputs, and cross-country productivity differences. Unpublished manuscript, University of Notre Dame.
- Doorenbos, J. and Kassam, A. (1979). Yield response to water. *Irrigation and drainage paper*, 33:257.
- Easterly, W. and Levine, R. (2003). Tropics, germs, and crops: how endowments influence economic development. *Journal of monetary economics*, 50(1):3–39.
- Fuwa, N., Edmonds, C., and Banik, P. (2007). Are small-scale rice farmers in eastern india really inefficient? examining the effects of microtopography on technical efficiency estimates. *Agricultural Economics*, 36(3):335–346.
- Gallup, J. L., Sachs, J. D., and Mellinger, A. D. (1999). Geography and economic development. *International regional science review*, 22(2):179–232.
- Galor, O. and Özak, Ö. (2015). Land productivity and economic development: Caloric suitability vs. agricultural suitability. working paper.
- Galor, O. and Özak, Ö. (2016). The agricultural origins of time preference. *American Economic Review*, 106(10):3064–3103.
- Godefroy, R. and Lewis, J. (2018). The impact of uncertainty in agriculture. working paper, Université de Montréal.

- Gollin, D., Lagakos, D., and Waugh, M. E. (2014a). Agricultural productivity differences across countries. *The American Economic Review: Papers and Proceedings*, 104(5):165–170.
- Gollin, D., Lagakos, D., and Waugh, M. E. (2014b). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.
- Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American Economic Review*, pages 160–164.
- Gollin, D. and Rogerson, R. (2014). Productivity, transport costs and subsistence agriculture. *Journal of Development Economics*, 107:38–48.
- Herrendorf, B. and Schoellman, T. (2015). Why is measured productivity so low in agriculture? *Review of Economic Dynamics*, 18:1003–1022.
- Holmes, T. J. and Lee, S. (2012). Economies of density versus natural advantage: Crop choice on the back forty. *Review of Economics and Statistics*, 94(1):1–19.
- Jaenicke, E. C. and Lengnick, L. L. (1999). A soil-quality index and its relationship to efficiency and productivity growth measures: two decompositions. *American Journal of Agricultural Economics*, 81(4):881–893.
- Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of political Economy*, 114(3):538–575.
- Kravchenko, A. N. and Bullock, D. G. (2000). Correlation of corn and soybean grain yield with topography and soil properties. *Agronomy Journal*, 92(1):75–83.
- Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *The American Economic Review*, 103(2):948–980.
- Levine, D. I. and Yang, D. (2014). The impact of rainfall on rice output in indonesia. Technical report, National Bureau of Economic Research.
- Lobell, D. B., Cassman, K. G., and Field, C. B. (2009). Crop yield gaps: their importance, magnitudes, and causes. *Annual review of environment and resources*, 34:179–204.

- 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).
- Mueller, N. D., Gerber, J. S., Johnston, M., Ray, D. K., Ramankutty, N., and Foley, J. A. (2012). Closing yield gaps through nutrient and water management. *Nature*, 490(7419):254.
- Nordhaus, W. D. (2006). Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, 103(10):3510–3517.
- Nunn, N. and Qian, N. (2011). The potato’s contribution to population and urbanization: Evidence from a historical experiment. *The Quarterly Journal of Economics*, 126(2):593–650.
- Prasada Rao, D. (1993). Intercountry comparisons of agricultural output and productivity. Technical report.
- Restuccia, D., Yang, D. T., and Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics*, 55(2):234–250.
- Rodrik, D., Subramanian, A., and Trebbi, F. (2004). Institutions rule: the primacy of institutions over geography and integration in economic development. *Journal of economic growth*, 9(2):131–165.
- Sachs, J. D. (2003). Institutions don’t rule: direct effects of geography on per capita income. Technical report, National Bureau of Economic Research.
- Schils, R., Olesen, J. E., Kersebaum, K.-C., Rijk, B., Oberforster, M., Kalyada, V., Khitrykau, M., Gobin, A., Kirchev, H., Manolova, V., et al. (2018). Cereal yield gaps across europe. *European Journal of Agronomy*, 101:109–120.
- Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences*, 106(37):15594–15598.

- Sherlund, S. M., Barrett, C. B., and Adesina, A. A. (2002). Smallholder technical efficiency controlling for environmental production conditions. *Journal of Development Economics*, 69(1):85–101.
- Steduto, P., Hsiao, T. C., Fereres, E., and Raes, D. (2012). *Crop yield response to water*. FAO Roma.
- Tittonell, P. and Giller, K. E. (2013). When yield gaps are poverty traps: The paradigm of ecological intensification in african smallholder agriculture. *Field Crops Research*, 143:76–90.
- Tombe, T. (2015). The missing food problem: Trade, agriculture, and international productivity differences. *American Economic Journal: Macroeconomics*, 7(3):226–258.
- van Ittersum, M. K., Cassman, K. G., Grassini, P., Wolf, J., Tittonell, P., and Hochman, Z. (2013). Yield gap analysis with local to global relevancea review. *Field Crops Research*, 143:4–17.
- Van Ittersum, M. K., Van Bussel, L. G., Wolf, J., Grassini, P., Van Wart, J., Guilpart, N., Claessens, L., de Groot, H., Wiebe, K., Mason-DCroz, D., et al. (2016). Can sub-saharan africa feed itself? *Proceedings of the National Academy of Sciences*, 113(52):14964–14969.
- Wiebe, K., Soule, M., Narrod, C., and Breneman, V. (2000). Resource quality and agricultural productivity: A multi-country comparison. USDA.
- Wiebe, K. D. (2003). *Land quality, agricultural productivity, and food security: Biophysical processes and economic choices at local, regional, and global levels*. Edward Elgar Publishing.
- Zhao, C., Liu, B., Piao, S., Wang, X., Lobell, D. B., Huang, Y., Huang, M., Yao, Y., Bassu, S., Ciais, P., et al. (2017). Temperature increase reduces global yields of major crops in four independent estimates. *Proceedings of the National Academy of Sciences*, 114(35):9326–9331.

Appendix — Not for publication

A Country Sample

Table A.1 lists all 162 countries in our data set, along with the corresponding country code, the number of cells covering the country, and the level of real GDP per capita in 2000.

Table A.1: List of Countries and Other Information

Country	Code	Cell Count	GDP per capita
Afghanistan	AFG	9000	327
Albania	ALB	444	3177
Algeria	DZA	30751	5276
Angola	AGO	14988	2901
Antigua and Barbuda	ATG	5	14522
Argentina	ARG	40080	12519
Armenia	ARM	451	4333
Australia	AUS	100208	30240
Austria	AUT	1447	31574
Azerbaijan	AZE	1311	3722
Bahamas	BHS	160	24593
Bangladesh	BGD	1759	1794
Belarus	BLR	4057	12188
Belgium	BEL	558	29693
Belize	BLZ	271	7910
Benin	BEN	1374	1336
Bhutan	BTN	523	2817
Bolivia	BOL	13284	3346
Bosnia and Herzegovina	BIH	836	5798
Botswana	BWA	7297	7219
Brazil	BRA	101847	8391
Brunei Darussalam	BRN	65	48210
Bulgaria	BGR	1754	6374

Continued on next page...

Table A.1 — *Continued from previous page*

Country	Code	Cell Count	GDP per capita
Burkina Faso	BFA	3262	1121
Burundi	BDI	312	706
Cambodia	KHM	2184	1764
Cameroon	CMR	5470	2448
Canada	CAN	244154	31471
Central African Republic	CAF	7287	918
Chad	TCD	15448	1445
Chile	CHL	11199	14309
China	CHN	136881	4076
Colombia	COL	13318	6620
Congo	COG	4032	3835
Costa Rica	CRI	609	9463
Cote d'Ivoire	CIV	3795	2761
Croatia	HRV	919	9775
Cuba	CUB	1381	7636
Cyprus	CYP	129	20275
Czech Republic	CZE	1419	16044
Democratic Republic of the Congo	ZAR	27327	312
Denmark	DNK	898	30468
Dominican Republic	DOM	598	7559
Ecuador	ECU	2996	4894
Egypt	EGY	13029	4690
El Salvador	SLV	253	5192
Equatorial Guinea	GNQ	314	8820
Eritrea	ERI	1469	668
Estonia	EST	1015	10405
Ethiopia	ETH	13365	892
Fiji	FJI	230	5784
Finland	FIN	9008	26402
France	FRA	9266	27311
Gabon	GAB	3056	8504
Gambia	GMB	132	1289

Continued on next page...

Table A.1 — *Continued from previous page*

Country	Code	Cell Count	GDP per capita
Georgia	GEO	1099	4310
Germany	GER	6608	29051
Ghana	GHA	2819	1359
Greece	GRC	1970	20708
Guatemala	GTM	1326	5530
Guinea	GIN	2908	3235
Guinea-Bissau	GNB	403	657
Guyana	GUY	2475	2457
Haiti	HTI	336	1655
Honduras	HND	1360	3062
Hungary	HUN	1590	13025
India	IND	40163	2687
Indonesia	IDN	22138	4151
Iran (Islamic Republic of)	IRN	22489	8049
Iraq	IRQ	6069	5403
Ireland	IRL	1334	31389
Israel	ISR	285	22356
Italy	ITA	4774	27142
Jamaica	JAM	135	7877
Japan	JPN	5488	28341
Jordan	JOR	1220	4329
Kazakhstan	KAZ	47485	7641
Kenya	KEN	6800	1943
Korea, Republic of	KOR	1434	18597
Kuwait	KWT	225	36146
Kyrgyzstan	KGZ	3098	3310
Lao People's Democratic Republic	LAO	2847	1777
Latvia	LVA	1371	8119
Lebanon	LBN	144	7505
Lesotho	LSO	414	1770
Liberia	LBR	1125	492
Libyan Arab Jamahiriya	LYB	21221	14674

Continued on next page...

Table A.1 — *Continued from previous page*

Country	Code	Cell Count	GDP per capita
Lithuania	LTU	1325	8566
Luxembourg	LUX	47	63392
Madagascar	MDG	7353	965
Malawi	MWI	1425	1032
Malaysia	MYS	3856	14178
Mali	MLI	15355	1108
Malta	MLT	6	19442
Mauritania	MRT	12944	2085
Mexico	MEX	25084	10339
Mongolia	MNG	26562	2008
Montenegro	MNE	214	4877
Morocco	MAR	5529	4574
Mozambique	MOZ	9647	1245
Namibia	NAM	10397	5531
Nepal	NPL	1944	1783
Netherlands	NLD	677	31927
New Zealand	NZL	4206	21437
Nicaragua	NIC	1538	2058
Niger	NER	14499	811
Nigeria	NGA	10772	1275
Norway	NOR	8617	41777
Oman	OMN	3849	23752
Pakistan	PAK	11827	2696
Panama	PAN	888	7124
Papua New Guinea	PNG	5470	2194
Paraguay	PRY	5062	4556
Peru	PER	15324	4975
Philippines	PHL	3538	3955
Poland	POL	5882	10834
Portugal	PRT	1381	19606
Puerto Rico	PRI	113	25955
Qatar	QAT	142	61389

Continued on next page...

Table A.1 — *Continued from previous page*

Country	Code	Cell Count	GDP per capita
Republic of Moldova	MDA	576	2420
Romania	ROM	3958	6151
Russia	RUS	421168	8305
Rwanda	RWA	293	994
Saudi Arabia	SAU	25034	19207
Senegal	SEN	2372	1732
Sierra Leone	SLE	863	1171
Singapore	SGP	7	35424
Slovakia	SVK	858	11844
Slovenia	SVN	341	19043
Solomon Islands	SLB	334	1318
Somalia	SOM	7490	480
South Africa	ZAF	16282	8441
Spain	ESP	7727	24945
Sri Lanka	LKA	793	4603
Sudan	SDN	30052	1546
Suriname	SUR	1706	7490
Swaziland	SWZ	228	6587
Sweden	SWE	11321	27174
Switzerland	CHE	704	34414
Syrian Arab Republic	SYR	2672	2446
Taiwan	TWN	464	21513
Tajikistan	TJK	2120	1902
Thailand	THA	6227	7058
The former Yugoslav Republic of Macedonia	MKD	396	6358
Togo	TGO	682	984
Tunisia	TUN	2186	7572
Turkey	TUR	11699	6428
Turkmenistan	TKM	7077	8716
Uganda	UGA	2834	1094
Ukraine	UKR	10587	5644

Continued on next page...

Table A.1 — *Continued from previous page*

Country	Code	Cell Count	GDP per capita
United Arab Emirates	ARE	908	38604
United Kingdom	GBR	4857	27032
United Republic of Tanzania	TZA	11088	681
United States	USA	160841	39241
Uruguay	URY	2467	11426
Uzbekistan	UZB	6960	1477
Vanuatu	VUT	157	5607
Venezuela	VEN	10758	10553
Viet Nam	VNM	3970	2407
Yemen	YEM	5148	1129
Zambia	ZMB	9045	1038
Zimbabwe	ZWE	4813	4528

Notes: The cell count of each country is from GAEZ, and refers to the number of 5-arc-minute cells covering the country. Real GDP per capita is from the PWTv6.3.

B GAEZ Data

B.1 Cell-Level Actual Data

The GAEZ methodology for estimating output and harvested land by crop for each cell in the world uses a downscaling methodology that combines aggregate and cell-level data. GAEZ first estimates for each cell, cultivated land and the split between rainfed and irrigated, using GIS land-cover datasets at the 5-arc minute resolution. The procedure ensures that the land class coverage is consistent with aggregate FAO land statistics (arable land) and land cover patterns obtained from remotely sensed data. To allocate crops to cells, GAEZ uses for each country, data on output and harvested area by crop at the national level from the FAO and at the sub-national level (regions, states, provinces, districts, counties, etcetera) from [Monfreda et al. \(2008\)](#). The downscaling procedure employs an iterative optimization algorithm that is initialized by feeding in a prior distribution of crops production allocation to cells that is based on cell-level information on the amount of cultivated land, bio-physical suitability for the production of the different crops, and socio-economic factors such as farming zone system, population density, and distance to market.

Then each iteration step determines the discrepancy between statistical totals available at the sub-national (or national) unit level and the respective totals calculated by summing harvested areas and production over cells. The magnitude of these deviations is then used to revise the land and crop allocation and to recalculate discrepancies. The process is continued until all accounting constraints are met, that is, output and harvested land sum up to the aggregates sub-nationally and nationally.

B.2 Input Assumptions for Potential Yields

Low inputs. Under this scenario GAEZ assumes a subsistence-based farming system, with traditional management, that does not necessarily produce for the market. Production is based on the use of traditional cultivars, labor intensive techniques, no mechanization, and minimum conservation measures. There is no application of nutrients, and no use of chemicals for pest or disease control. The assumed water supply under this scenario is fully rainfed farming.


Mixed inputs. Under this scenario GAEZ assumes that the highest quality land (very suitable and suitable) uses high inputs, the moderately suitable land uses an intermediate level of inputs, and the marginal land uses low inputs. Under intermediate inputs the farming system is partly market oriented involving both subsistence and commercial farming; use of improved varieties; intermediate labor intensity with hand tools and/or animal traction and some mechanization; some fertilizer application; and some chemical pest, disease, and weed control. Under high inputs there is advanced management; the farming system is primarily market oriented with commercial production; high yield varieties are used; fully mechanized; low labor intensity; optimum application of nutrients and chemical pest, disease and weed control. This scenario covers all land, both under rainfed and irrigated water supply. GAEZ considers the mixed input scenario as a “reasonable reflection of actual agricultural input and management circumstances.” (page 97, *GAEZ Model Documentation*, 2012).

C Within-country Dispersion of Potential Yields

We document a measure of the variability of land quality across countries. In particular, we compute the standard deviation of log potential yields across cells with agricultural production and report this dispersion across countries by real GDP per capita in Figure C.1. We report two measures of dispersion. The first, in panel (a), for all crops, that is the simple average of potential yields across crops within a cell; the second, in panel (b), only for maize which is the most prevalent crop produced around the world. We use the potential-yield under the rainfed low-input scenario as this

most closely reflects land quality productivity.

Figure C.1: Within-country Dispersion in Land Quality



Notes: Dispersion in potential yields across cells within a country under the rainfed low-input scenario. Panel (a) reports the simple average of yields across all crops within a cell, whereas panel (b) is only for maize in each cell.

We find that while the dispersion in land quality within a country differs quite substantially across countries, the dispersion is not systematically associated with the level of development, there is only a modest positive association. Hence, land quality endowment in poor countries is not worse than rich countries, not just in terms of averages from aggregate potential yields, but also in terms of the dispersion of land quality within the country. This finding is relevant to the extent that poor countries may not optimize on the location of production or the number of locations with agricultural production.


D Robustness and Validation

Our construction of the country-level potential yields does not rely on actual output or yields at the crop-cell level. In the case of the production-potential counterfactual, however, the aggregate potential yield uses the cell-level land allocation by crop in each country as weights. We examine the bias and reliability of the cell-level land weights, conduct our own external validation, and show the robustness of our results to alternative land weighting schemes.

Lower and middle income countries do not tend to have less detailed data. GAEZ uses the most spatially disaggregated data available from [Monfreda et al. \(2008\)](#), at the sub-national level. One

natural concern is that the prevalence of spatially disaggregated agricultural statistics that GAEZ uses is disproportionate for developed countries than for lower and middle income countries. Using sub-national data availability from [Monfreda et al. \(2008\)](#) (from their Table 2), in [Figure D.2](#) we plot the percentage of land covered by sub-national data across countries against their GDP per capita. There is no systematic relationship of sub-national data availability and the level of income. There are rich and poor countries with substantial sub-national coverage and there rich and poor countries with no sub-national coverage. We conclude that GAEZ does not have more detailed information for the richer countries than for the poorer countries, that can inherently bias their downscaling methodology of attributing national and sub-national statistics to cells across countries.

Figure D.2: Subnational Data Availability



GAEZ validation. A general limitation of cell-level estimated data is that because they require substantial amounts of crop statistics and other data at the finest level of disaggregation, there is not much “out-of-sample” data remaining in order to externally validate the results. Although, not in a systematic fashion, GAEZ has done some testing of the reliability of its actual agricultural estimates: “Tests in China and Brazil by comparing downscaled results based on statistics available on national level, with detailed sub-national statistics on county and micro region level revealed strong correlations between downscaled national statistics and county/micro-region level statistics of harvested areas, yields and crop production.” (GAEZ-FAO website, under “FAQs,” www.fao.org/nr/gaez/faqs/en/.)

Our own validation. We conduct our own validation exercise to confirm that the crops GAEZ attributes as being produced to the different cells are actually produced there. In particular, we use survey farm-level data from the World Bank’s Living Standards Measurement Study (LSMS) (<https://www.worldbank.org/en/programs/lsms/overview>) for Uganda for the earliest year with GPS data in the survey of 2009. To make the comparison we use the GPS coordinates of the locations of farms in the LSMS data to assign farm households to pixels in the GAEZ rasters (5 arc minute resolution). Given that the LSMS data are survey data, the observations are sparse with some pixels having multiple households, others very few, and most none. For the pixels for which there is any LSMS household we compute a series of dummies, one for each of the 18 GAEZ crops and for each of the LSMS and GAEZ. In the set of dummies for LSMS (GAEZ), the dummy in a pixel takes the value of 1 if the crop is produced in LSMS (GAEZ). Then we create a third dummy that takes the value of 1 for a crop in a given pixel if the crop is produced in both GAEZ and LSMS. Finally we compute the fraction of the aligned crops in total GAEZ crops for every pixel. For the median pixel 60 percent of all the crops GAEZ attributes to a cell are also confirmed to be produced according to the LSMS. If we focus only on maize, one of the most widely produced crops in Uganda 77 percent of the cells that GAEZ attributes maize production is also confirmed by the LSMS data. This is remarkable given that the LSMS data are survey data, are spatially sparse, and are 9 years later than the GAEZ data. For these reasons we could not make more direct comparisons of land allocations and yields.

Our conclusions do not change if we weigh crops within cells equally. Please refer to Robustness Section 4.4 in the main text.

Our conclusions do not change if we use crop-cell land allocations from IFPRI’s Harvest Choice project. The International Food Policy Research Institute’s (IFPRI) *Harvest Choice* project (<https://www.ifpri.org/project/harvestchoice>) is another project that uses crop statistics at the national and subnational levels across the world for the year 2005 to estimate crop yields and harvested land at the 5 arc-minute grid cell level. Their model for downscaling agricultural statistics, called “Spatial Production Allocation Model (SPAM),” uses a cross-entropy optimization approach that uses information on cropland surface, location of irrigated areas, crop suitability, rural population densities, production systems and crop prices. In this sense, even though distinct from GAEZ, *Harvest Choice*’s disaggregation methodology is similar in nature and uses similar types of disaggregate information to estimate the same cell-level resolution as GAEZ. *Harvest Choice* however uses data that reflect the year 2005 (rather than 2000 in GAEZ), and includes a finer set of crops. The underlying crop suitability surfaces that SPAM uses are from GAEZ, but uses a more updated cropland surface, and has made an effort to use more district-level agricultural statistics within countries.

While the cell-level data from *Harvest Choice* are also estimates at the 5 arc-minute resolution, and involve a five year gap from GAEZ, we compare the cell-level land allocations across crops (which determine the cell-level weights in our methodology) between GAEZ and *Harvest Choice*. The cell-by-cell correlation of harvested land for all crops between GAEZ and *Harvest Choice* for the entire world is 0.71. The same correlation for the three most popular crops, wheat, rice, and maize is 0.69. The correlation for rice alone is 0.84. While these correlations are high, we go a step further and re-compute each country’s potential yield with our methodology, aggregating GAEZ potential yields, but using the cell-level land allocations across crops from *Harvest Choice* instead (rather than the GAEZ ones). Just as in the alternative with equal weights across all produced crops above, with the Harvest Choice weights our conclusions are the same.

Figure D.3: Potential Yields with Harvest Choice Weights




Figure D.3 displays the aggregate potential yield for each country (against GDP per capita) using the cell-by-cell *Harvest Choice* weights. The cell-level potential yields from GAEZ used are for low inputs and rainfed water supply. The lack of a systematic relationship of aggregate potential yields with income is true here as well.