Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from a 9 Million-Field Partition of the Earth

Arnaud Costinot
MIT and NBER

Dave Donaldson
MIT and NBER

Cory Smith
MIT

September 2012

Abstract

A large agronomic literature has modeled the implications of such climate change for crop yields, crop by crop and location by location. These studies document the harm that climate change is expected to inflict on a specific crop at a specific location. The goal of the present paper is to quantify the macro-level consequences of these micro-level shocks. Our analysis builds on the simple observation that in a globalized world, the impact of micro-level shocks do not only depend on their average level, but also on their dispersion over space, i.e. how they affect comparative advantage. Using an extremely rich micro-level dataset that contains information about the productivity—both before and after climate change—of each of 39 crops for each of over 9 million high resolution grid cells covering the surface of the Earth, we find that international trade, even with reasonable levels of trade costs, will substantially mitigate the ill-effects of climate change on agricultural living standards in the average country.

1 Introduction

The warmer climates predicted by climatological models portend a grim future for many biological systems, such as agricultural plant life, on which human welfare depends. But just how much will living standards suffer as plants wilt in a hotter world? A large agronomic literature has modeled the implications of such climate change for crop yields, crop by crop and location by location (see IPCC, 2007, Chapter 5 for a review). These studies document the harm that may be inflicted on a specific crop at a specific location. The goal of our paper is to quantify the macro-level consequences of these micro-level shocks.

Our analysis builds on the simple observation that in a globalized world—which is the world we inhabit—the impact of micro-level shocks do not only depend on their average level, but also on their dispersion over space. If climate change makes regions of the world more homogeneous in terms of their agricultural productivity, there will be less trade and welfare will further decrease. If climate change instead raises heterogeneity across regions, there will be more scope for international trade, which will dampen the adverse consequences of climate change. In short, the macro-consequences of climate change in a global economy are inherently related to how it affects comparative advantage across regions of the world.

To shed light on the relationship between climate change and comparative advantage, we take advantage of an extremely rich micro-level dataset on agricultural productivity: the Food and Agriculture Organization's (FAO) Global Agro-Ecological Zones (GAEZ) dataset. This dataset uses agronomic models and high resolution data on geographic characteristics such as soil, topography, elevation and, crucially, climatic conditions to predict the yield that would be obtainable at over 9 million high resolution grid cells covering the surface of the Earth. The GAEZ dataset is available both under contemporary growing conditions and under a climate change scenario similar to those used by the UN's Intergovernmental Panel on Climate Change (IPCC). By comparing productivity for a given crop under the two scenarios at each of our 9 million grid-cells, we can therefore directly observe the evolution of comparative advantage across space, as predicted by climatologists and agronomists.

A sample of the GAEZ predictions can be seen in Figure 1. Here we plot, for each grid cell on Earth, the predicted percentage change in productivity associated with climate change for two of the world's most important crops: wheat (panel (a)) and rice (panel (b)). As is clear, there exists a great deal of heterogeneity in the effects of climate change both across crops and over space—many regions see a differential productivity change in wheat and rice, and this relative productivity change is different from that of other regions. Further, the contours of the effects of climate change on rice and wheat appear not to reflect country borders. Within-country heterogeneity is a central feature of these data..

There is a more systematic way to see that an understanding of the impact of climate change in agriculture requires an understanding of what will happen to comparative advantage. A simple regression demonstrates the fact that, at least according to the GAEZ estimates, climate change will distort comparative advantage considerably around the world. If we regress the log change in each grid-cell's yield (i.e. before and after climate change) on a country and a crop fixed effect, the R-squared from this regression is only 50 percent. Put differently, half of the effects of climate change will preserve the existing pattern of agricultural comparative advantage around the world (i.e. yields will evolve in a manner that is constant within countries and/or within crops) but half of the effects of climate change will alter this pattern. It is the latter half of these effects of climate change that are the focus of the present paper.

To go beyond the evolution of comparative advantage documented in the agronomic GAEZ data and quantify the economic macro-consequences of climate change, we need an economic model of agricultural markets that can predict: (*i*) where crops are produced and consumed (despite the presence of trade frictions), and in turn, which productivity changes are relevant and which ones are not; (*ii*) how shocks to the supply of crops affect prices around the world; and (*iii*) how changes in productivity and prices map into welfare changes. We propose a perfectly competitive model of trade in which each country consists of a large number of 'fields' with heterogeneous productivity across multiple crops. These are the theoretical counterparts of the 9 million grid-cells in the GAEZ data. In this model, comparative advantage, i.e. relative productivity differences across crops and fields, determines the pattern of specialization within and between countries. Finally, international trade is subject to iceberg trade costs whose magnitude pins down the level of integration of local agricultural markets.

Besides the highly detailed GAEZ data, our quantitative model depends only on a small number of parameters: (*i*) the elasticity of substitution between crops from different countries, which is the equivalent of the Armington elasticity in standard Computational General Equilibrium (CGE) models; (*ii*) the extent of within-field heterogeneity in productivity, which is unobserved in the GAEZ data; and (*iii*) the elasticity of trade costs with respect to distance, which we assume is the sole determinant of iceberg trade costs. These three parameters can be separately estimated using trade, output, and price data in a straightforward and transparent manner.

Armed with these three parameters and the detailed knowledge of the pattern of comparative advantage across fields and crops around the world, we simulate our model under the baseline no-climate change scenario and explore three counterfactual scenarios. In our first scenario, we study the consequences of climate change—i.e., a change in the GAEZ productivity from contemporary growing conditions to climate change conditions—under the assumption that farmers are free to change their output decisions and countries are free

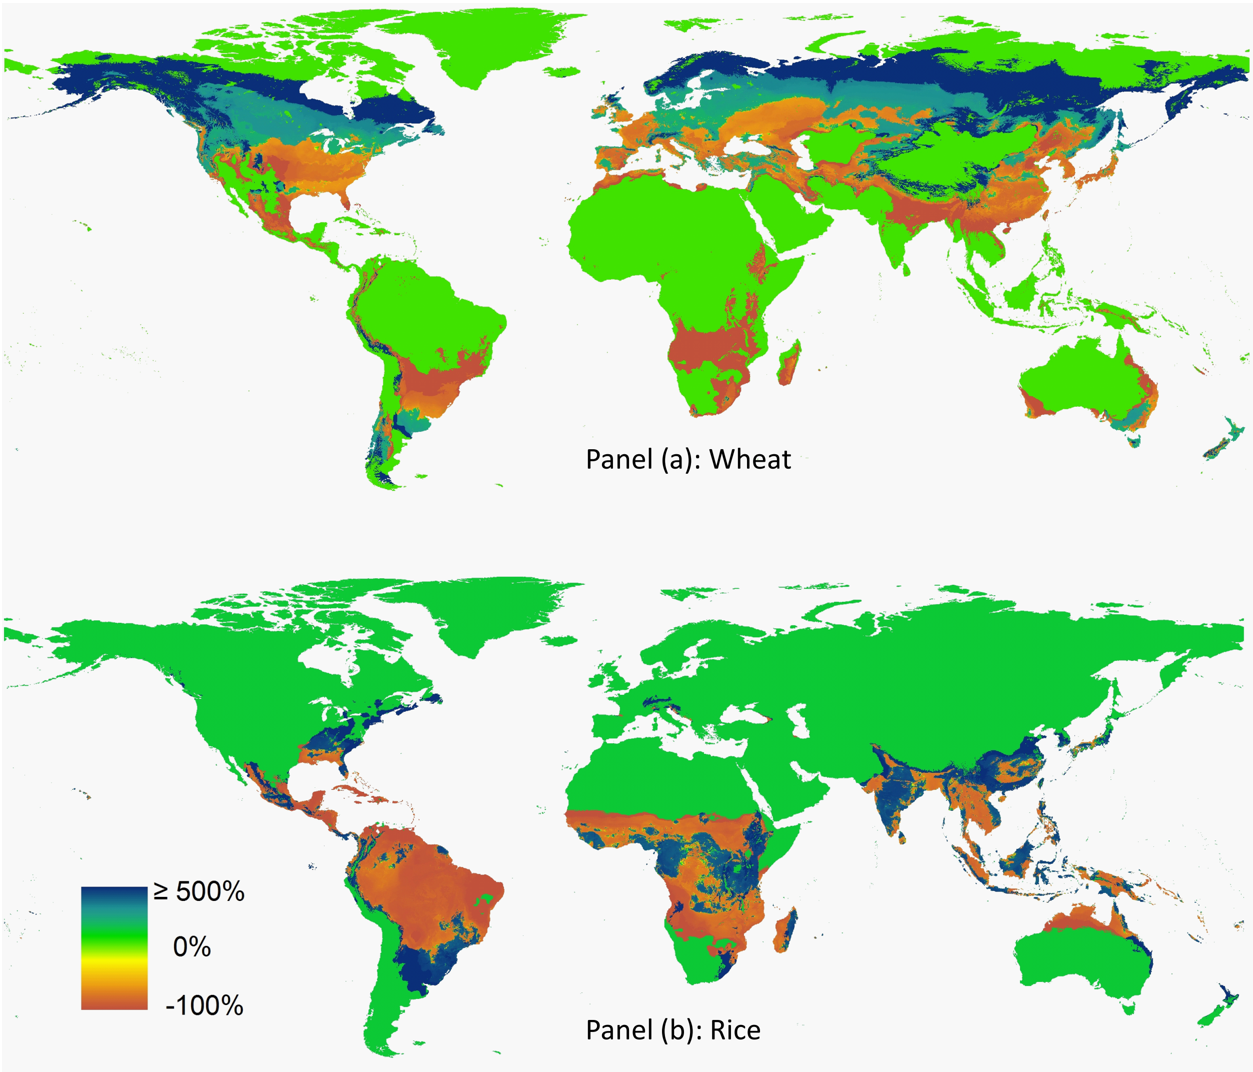

Figure 1: Percent changes in yield due to climate change in GAEZ model for wheat and rice. Large green areas are those for which yields are zero both before and after climate change.

to trade (subject to our estimated trade costs). We then contrast the welfare implications of climate change under this scenario to those in a counterfactual world in which countries can trade, but farmers cannot reallocate productions, and conversely, a counterfactual world in which farmers can reallocate production, but countries are under autarky. This allows us to quantify how much changes in comparative advantage within and between countries, respectively, may offset or magnify the effects of a hotter climate. Ultimately, while we find that the negative effects of climate change will be substantial for most countries (e.g. the effect is approximately 13 percent of expenditure on agricultural goods in the world average country), these negative effects would be much worse if fields could not change their what they grow (38 percent loss) or if countries could not trade at all (19 percent loss).

The literature on international trade and climate change is large and varied, though mostly based on Computational General Equilibrium (CGE) models. A first group of papers focuses on the direct impact of international trade on the level of carbon emissions caused by international transportation; see e.g. Cristea, Hummels, Puzzello, and Avetysyan (forthcoming) and Shapiro (12). A key insight is that although international transportation negatively affects the environment, the associated welfare consequences are an order of magnitude smaller than the gains from international trade. A second group of papers focuses on the issue of carbon leakages, i.e. the idea that if only a subset of countries tax carbon emissions, the level of emissions of nontaxing countries is likely to go up; see Felder and Rutherford (1993), Babiker (2005), Elliott, Foster, Kortum, Munson, Cervantes, and Weisbach (2010).

More closely related to this paper are studies on international trade and adaptation in agriculture; see Reilly and Hohmann (1993), Rosenzweig and Parry (1994), Tsigas, Friswold, and Kuhn (1997) and Hertel and Randhir (2000). The main difference between previous papers and the present analysis lies in the level of disaggregation at which we observe the micro-consequences of climate changes—while the existing literature works with country averages, we aggregate up in a theoretically consistent manner from the full set of nine million fields around the world. (We document in Section 4 below how an analysis based on national averages performs significantly worse at matching the data within sample.) By feeding this rich micro-data into a general equilibrium in which comparative advantage determines the pattern of specialization, both within and across countries, we are then able to study, quantify and compare the gains from adaptation to climate change through local and international specialization. Finally, our analysis is related to Costinot and Donaldson (2011) who also use the GAEZ data to quantify the gains from economic integration in U.S. agricultural markets from 1880 to 2000.

The rest of this paper is organized as follows. Section 2 introduces our theoretical framework. Section 3 describes the data that feeds into our analysis. Section 4 describes our empirical parameter estimation procedure and presents our parameter estimates (as well as measures of goodness of fit of the model). Section 5 then presents the results of our counterfactual simulations. Finally, Section 6 describes some robustness extensions that are in progress and Section 7 offers some concluding remarks.

2 Theory

2.1 Basic Environment

We consider a world economy comprising multiple countries, indexed by $i \in \mathcal{I} \equiv \{1, ..., I\}$. In each country, the only factors of production are fields, indexed by $f \in \mathcal{F}_i \equiv \{1, ..., F_i\}$, each comprising a continuum of heterogeneous parcels of land, indexed by $\omega \in [0,1]$. We think of land as equipped land, i.e. land plus physical capital and labor, though we abstract from the allocation of physical capital and labor across fields. All fields have the same size, which we normalize to one. In our dataset, the size of a field is equal to 5 arc-minute grid-cell and there are 9 million such grid-cells on Earth.

Fields can be used to produce multiple goods indexed by $k \in \mathcal{K} \equiv \{0, ..., K\}$. Goods $1, ..., K$ are crops, whereas good 0 will be an outside good. We think of the outside good as residential housing, forestry, manufacturing, or any agricultural activity (such as livestock production) that does not correspond to the crops included in our dataset.

There is a representative agent in each country i whose preferences can be represented by a two-level utility function:

$$U_i = \prod_{k=0}^{K} \left(C_i^k\right)^{\beta_i^k}, \tag{1}$$

$$C_i^k = \left(\sum_{j=1}^{I} \left(C_{ji}^k\right)^{(\sigma-1)/\sigma}\right)^{\sigma/(\sigma-1)}, \text{ for all } k = 1, ..., K, \tag{2}$$

where $\beta_i^k \geq 0$ denotes exogenous expenditure shares, with $\sum_{k=0}^{K} \beta_i^k = 1$; $\sigma > 0$ denotes the elasticity of substitution between crops from different origins, e.g. French versus U.S. wheat; C_{ji}^k denotes the consumption in country i of crop $k = 1, ..., K$ produced in country j, with C_i^k the aggregate consumption of crop k in country i; and C_i^0 denotes the aggregate consumption of the outside good in country i.

Parcels of land are perfect substitutes in the production of each good, but vary in their exogenously-given productivity per acre, $A_i^{fk}(\omega) \geq 0$. Total output Q_i^k of good k in country i is given by

$$Q_i^k = \sum_{f \in \mathcal{F}_i} \int_0^1 A_i^{fk}(\omega)\, L_i^{fk}(\omega)\, d\omega, \tag{3}$$

where $L_i^{fk}(\omega) \geq 0$ denotes the endogenous number of acres of parcel ω in field f allocated to good k in country i. For all goods $k \in \mathcal{K}$, we assume that the productivity of each parcel can be expressed as

$$\ln A_i^{fk}(\omega) = \ln A_i^{fk} + \varepsilon_i^{fk}(\omega). \tag{4}$$

The first term, $A_i^{fk} > 0$, is a common productivity shifter of all parcels in field f. It measures the comparative and absolute advantage of a field in producing particular goods. The GAEZ project data give us direct information about A_i^{fk} for all crops $k = 1, ..., K$ as a function of global temperatures, which will be the core inputs in our quantitative exercise. $\varepsilon_i^{fk}(\omega)$ reflects unobserved within-field heterogeneity in productivity across parcels. In line with Eaton and Kortum (2002), we assume that $\varepsilon_i^{fk}(\omega)$ is independently drawn for each

(i, f, k, ω) from a Gumbel distribution:

$$F(\varepsilon) = \Pr\left[\varepsilon_i^{fk}(\omega) \leq \varepsilon\right] = \exp\left[-\exp(-\theta\varepsilon - \kappa)\right], \quad (5)$$

where $\theta > 1$ measures the extent of within-field heterogeneity and the constant κ is set such that $A_i^{fk} = E\left[A_i^{fk}(\omega)\right]$ in Equation (4).[1] Finally, since we do not have disaggregated productivity data in the outside sector, we assume that in all countries $i \in \mathcal{I}$, all fields $f \in \mathcal{F}_i$ have the same productivity in the outside sector, $A_i^{f0} = A_i^0$, which we normalize to one in all countries.

All markets are perfectly competitive. International trade in crops $k = 1, ..., K$ is subject to iceberg trade costs. In order to sell one unit of a good in country j, firms from country i must ship $\tau_{ij}^k \geq 1$ units, with $\tau_{ii}^k = 1$. Non-arbitrage therefore requires the price of a crop k produced in country i and sold in country j to be equal to

$$p_{ij}^k = \tau_{ij}^k p_i^k, \quad (6)$$

where p_i^k is the producer of farm-gate price of crop k in country i. The outside good, by contrast, is not traded. In line with the previous notation, we denote by p_i^0 the price of the outside good in country i.

2.2 Competitive Equilibrium

In a competitive equilibrium, all consumers maximize their utility, all firms maximize their profits, and all markets clear. Given Equations (1), (2), and (6), utility maximization by consumers in any country i requires

$$C_i^0 = \frac{\beta_i^0 Y_i}{p_i^0}, \text{ for all } i \in \mathcal{I} \quad (7)$$

$$C_{ji}^k = \frac{\left(\tau_{ji} p_j^k\right)^{-\sigma}}{\sum_{j'=1}^{I}\left(\tau_{j'i} p_{j'}^k\right)^{1-\sigma}} \beta_i^k Y_i, \text{ for all } i, j \in \mathcal{I} \text{ and } k = 1, ..., K, \quad (8)$$

where $Y_i \equiv \sum_{k \in \mathcal{K}} p_i^k Q_i^k$ denotes total income in country i.

Profit maximization requires that all parcels of land are allocated to the good that maximizes the value of their marginal product. Let π_i^{fk} denote the share of parcels in a field f

[1] Formally, we set $\kappa \equiv \theta \ln \Gamma\left(\frac{\theta-1}{\theta}\right)$, where $\Gamma(\cdot)$ denotes the Gamma function, i.e. $\Gamma(t) = \int_0^{+\infty} v^{t-1}\exp(-v)dv$ for any $t > 0$.

located in country i that are allocated to a good k. By Equation (3)-(5), we therefore have

$$\pi_i^{fk} = \Pr\left\{\frac{A_i^{fk}(\omega)}{A_i^{fl}(\omega)} > \frac{p_i^l}{p_i^k} \text{ if } l \neq k\right\} = \frac{\left(p_i^k A_i^{fk}\right)^\theta}{\sum_{l\in\mathcal{K}}\left(p_i^l A_i^{fl}\right)^\theta}.$$

The previous expression highlights in a simple manner how relative productivity differences, i.e. comparative advantage, determines factor allocation in this economy.

Given factor allocation, total output for good k in country i can be expressed as

$$Q_i^k = \sum_{f\in\mathcal{F}_i} E\left[A_i^{fk}(\omega)\,|p_i^k A_i^{fk}(\omega) = \max_{l\in\mathcal{K}} p_i^l A_i^{fl}(\omega)\right]\left(\frac{\left(p_i^k A_i^{fk}\right)^\theta}{\sum_{l\in\mathcal{K}}\left(p_i^l A_i^{fl}\right)^\theta}\right),$$

which, using again Equations (4) and (5), simplifies into

$$Q_i^k = \sum_{f\in\mathcal{F}_i} A_i^{fk}\left(\frac{\left(p_i^k A_i^{fk}\right)^\theta}{\sum_{l\in\mathcal{K}}\left(p_i^l A_i^{fl}\right)^\theta}\right)^{(\theta-1)/\theta}. \tag{9}$$

Finally, good market clearing requires that the supply of each good is equal to its demand:

$$Q_i^0 = C_i^0, \text{ for all } i \in \mathcal{I}, \tag{10}$$

$$Q_i^k = \sum_{j\in\mathcal{I}} \tau_{ij} C_{ij}^k, \text{ for all } i \in \mathcal{I} \text{ and } k = 1, ..., K. \tag{11}$$

Let $p_i \equiv \left(p_i^k\right)_{k\in\mathcal{K}}$ denote the vector of producer prices, $Q_i \equiv \left(Q_i^k\right)_{k\in\mathcal{K}}$ denote the vector of output levels, and $C_i \equiv \left(C_i^0, \left(C_{ji}^k\right)_{j\in\mathcal{I},k\neq 0}\right)$ denote the vector of consumption levels in country i. In the rest of this paper we formally define a competitive equilibrium as follows.

Definition 1 *A competitive equilibrium is a set of producer prices,* $(p_i)_{i\in\mathcal{I}}$*, output levels,* $(Q_i)_{i\in\mathcal{I}}$*, and consumption levels,* $(C_i)_{i\in\mathcal{I}}$*, such that Equations* (7)-(11) *hold.*

3 Description of Data

Our analysis draws on four main types of data: (*i*) estimates of agricultural productivity, at each high-resolution land parcel on Earth and for each of a series of crops, for a baseline (i.e. pre-climate change) scenario; (*ii*) similar agricultural productivity estimates, calculated in a similar manner, but for a climate change scenario; (*iii*) data on actual output, producer prices and trade flows, by crop, for each country; (*iv*) data on total GDP by country; and

(v) data on various potential determinants of trade costs. We describe the sources and construction of each of these inputs here in turn.

3.1 Agricultural Productivity Estimates at Baseline

The first data source that draw on provides estimates of average productivity during the 'baseline', or pre-climate change, period. We require a measure of A_i^{fk} in the model above, namely the productivity for crop k of a small parcel of land (which we refer to as a 'field', f) in country i. We obtain these measures from the Global Agro-Ecological Zones (GAEZ) project, which is organized under the auspices of the Food and Agriculture Organization (FAO) and the IIASA.[2] Because this data source is non-standard we provide a lengthy description here.

Crucially the GAEZ productivity estimates are available for each field f regardless of whether field f is actually growing crop k. The GAEZ project provides these estimates by drawing on state-of-the-art agronomic models of how each crop k will fare in the growing conditions available at field f. The primary goal of the GAEZ project is to inform farmers and government agencies about optimal crop choice (for given prices) in any given location on Earth—that is, to help farmers to know how productive they would be at crops they are not currently growing.

Three inputs enter the GAEZ project's agronomic model. The first input is a long vector of attributes describing the growing characteristics at field f. These characteristics include eight different soil types and conditions, elevation, average land gradient, and climatic variables (based on rainfall, temperature, humidity, wind speed and sun exposure). The size of a field f in the GAEZ data is governed by the limitations placed by the spatial resolution of the climatic data, which is the land characteristic whose underlying data is most coarse. Since the climatic data is available at the 5 arc-minute level, this governs the size of field in our analysis.[3] At the 5 arc-minute level there are 9,000,796 grid cells on Earth; we are left with 2,114,956 fields within the 187 countries in our sample after throwing out grid cells that lie over water. The second input is a set of hundreds of model parameters, each specific to crop k, that govern how a given set of land characteristics map into the yield of crop k according to the GAEZ project's agronomic model. The parameters used by GAEZ are an aggregation of such parameters found in the agronomic literature and each is estimated through the use of field experiments at agricultural research stations. They are not estimated through the use of any sort of statistical procedure that compares outputs to inputs across a population of farmers without the absence of experimental control—a procedure that the

[2] We accessed these data here: http://www.gaez.iiasa.ac.at/w/ctrl?_flow=Vwr&_view=Type&idAS=0&idFS=0&fieldmain=main_py_six_qdns&idPS=1e1d6e7d7ec3368cf13a68fc523d1ed4870e8b45.

[3] Many other inputs are available at the 30 arc-second grid-cell level.

model outlined above suggests would be inappropriate (without controlling for endogenous sorting of fields into crops based on prevailing prices). The third and final input into the GAEZ model is a set of assumptions about the extent to which complementary inputs (such as irrigation, fertilizers, machinery and labor) are applied to the growing of crop k at field f. Naturally, farmers' decisions about how to grow their crops and what complementary inputs to apply affect crop yields in addition to the land characteristics (such as sunlight) over which farmers have relatively little control. For this reason the GAEZ project constructs different sets of productivity predictions for different scenarios regarding the application of complementary inputs. In the results presented here we use the scenario referred to as 'high inputs' (in which modern machinery, etc., are assumed to be available in the GAEZ agronomic model if that is deemed useful) with 'rain-fed' water supply.

The GAEZ data are made available as gridded machine-readable files. We map each grid cell to the country i in which it is located by using a country-to-grid cell mapping available as part of the Global Poverty Dataset produced by CIESIN at Columbia University.[4] The GAEZ data are available for all countries in the world apart from those that are extremely small (roughly speaking, smaller than a field).[5]

The GAEZ data are produced for each of 43 crops. Of these, two crops (gram and jatropha) are not available in the FAO data (on output, prices and trade flows) described below, so we drop these. And two pairs of crops (dryland rice and wetland rice, as well as pearl millet and foxtail millet) are only available in the FAO data as aggregates (i.e. rice and millet, respectively) so we take the maximum yield over each pair, within each field, as our measure of the productivity A_i^{fk} of these aggregates (for example, our measure of $A_i^{fk=rice} = \max\{A_i^{fk=drylandrice}, A_i^{fk=wetlandrice}\}$); this implicitly assumes that farmers are using the type of rice or millet at which they are most productive (and the prices of each type of rice or millet is the same). We are then left with 39 crops that concord precisely with those crops in the FAO data.[6]

Finally, we obtain these GAEZ productivity estimates for what the GAEZ project refers to as the 'baseline' period, an average of runs of the GAEZ models for the weather observed in each year from 1961 to 1990. This has the attraction of averaging, in a coherent manner, over the idiosyncrasies of any given year's weather. (An alternative would be to pick the

[4] We accessed the CIESIN country mapping file here: http://sedac.ciesin.columbia.edu/povmap/ds_global.jsp. The CIESIN file is at a finer (2.5 arc-minute) level than the GAEZ data (5 arc-minute level). We therefore assign a field f (ie a grid cell in the GAEZ data) to the country i that has the largest number of CIESIN grid cells within a GAEZ grid cell, breaking the small number of ties randomly.

[5] For computational ease we work, for now, with the fifty largest (in terms of total world crop revenue) countries in the world. These account for roughly 92 percent of world agricultural output.

[6] While in principle the analysis here could be based on all 39 crops, for computational ease we work with the ten most important (in terms of total world revenue) crops. These comprise over 90 percent of world agricultural revenue.

GAEZ data output from one particular year but the most recent available year is 2000 and the FAO data we use below is from 2009.)

3.2 Agricultural Productivity Estimates After Climate Change

Our analysis of the impact of climate change on global agricultural markets draws naturally on scientists' predictions about the impact that climate change will have on crop yields around the world. In Section 5 below we refer in our model to productivity changing—for any country i, crop k and field f—from $\left(A_i^{fk}\right)$ at baseline to $\left(A_i^{fk}\right)'$ after climate change. We obtain these predictions about crop yields under an alternative climate from the GAEZ project so that our baseline and climate change productivity estimates are computed under exactly the same maintained agronomic assumptions.

The only change that the GAEZ project implements when computing post-climate change productivity estimates $\left(A_i^{fk}\right)'$ rather than baseline productivity estimates $\left(A_i^{fk}\right)$ concerns the weather that prevails at field f in country i in each scenario. As described above, when computing baseline productivity estimates the GAEZ project obtains a separate A_{it}^{fk} for each year t from 1961 to 1990 when the weather from each year is used as the input to their model; they then average over these 30 values of A_i^{fk}. A similar procedure is used when the GAEZ project computes post-climate change productivity estimates—the average over a separate A_{it}^{fk} for each year t from 2071 to 2100 is reported—only instead of realized past weather in year t the GAEZ project uses the predicted future weather from year t. Estimates of future weather in year t (for t =2071 to 2100) are obtained from an average of runs of a global circulation model (GCM) of the sort used by climatologists to predict the nature of climate change. While the GAEZ estimates are available for a range of different GCMs, we use that of the Hadley CM3 A1FI model because of its prominence in the UN's IPCC programme.[7]

Finally we note that we use the GAEZ climate change scenario in which plant carbon dioxide fertilization is assumed to be active.

3.3 Agricultural Output, Price, and Trade Flow Data

An essential aspect of our analysis is the ability to estimate all of the unknown parameters in our model, at baseline, in a manner that is consistent with our model. This estimation procedure—described below—requires data on actual output, producer prices and trade flows at baseline. We obtain these data from the FAOSTAT program at the FAO.[8] The FAOSTAT program aims to provide data on worldwide production and trade, by crop and country, that

[7] In ongoing work we aim to compare the output of our model of global agricultural markets across different climate change models and prediction horizons.

[8] We accessed these data from http://faostat3.fao.org/home/index.html#DOWNLOAD.

is both consistent and complete. We use four variables from FAOSTAT in our analysis. The first variable we use is the output, in physical units (i.e. tonness), of crop k in country i, denoted by Q_i^k in the model above. The second variable, which we denote by p_i^k, is the producer price (i.e. the price paid to producers, after taxes and subsidies) of crop k in country i. The third variable is the total value of exports of crop k from country i to country j, denoted by X_{ij}^k below (in the notation introduced above, $X_{ij}^k = \left(\tau_{ij} p_i^k\right) C_{ij}^k$). We obtain this variable from the imports of reporting countries (the country that collected the data underlying the trade flow in question, in contrast to the partner country in any trade flow) in the FAOSTAT data. Finally, the fourth variable that we use is the landed (or CIF) price of crop k sent from country i to country j, which we obtain from the unit value (i.e. the total reported value of a trade flow divided by the total reported quantity traded) associated with imports as reported by reporting countries who report imports in CIF terms. (For example, if j is a reporting country then the value of its imports of crop k from country i, denoted X_{ij}^k, is in CIF terms.)

As described above, we work with the 39 crops that concord between the FAO and GAEZ data (which cover over 98% of world crop output, according to the FAO). Concording crops in the output and price data to crops in the GAEZ data is straightforward since both treat crop products only in their pre-processed forms. Concording crops in the trade data to crops in the GAEZ data, however, is more involved because for some crops the traded product is primarily a processed version of the pre-processed (or 'raw') output of the crop. In the majority of cases countries trade some quantity of both the processed and the raw product of a given crop; in these cases we work only with the trade in the raw product.[9] In the case of two crops (oil palm and cotton) there is very little trade in the raw version of the crop but the FAO provides conversion factors to convert the processed version of a crop into its raw crop equivalent quantity. And in one case (cassava) there is trade only in the processed version and no conversion factor is available. We therefore drop this crop from our analysis when it involves trade flows.

Finally, we work with the 187 countries that are reported in both the GAEZ data and the FAO data. This spans the vast majority of world agricultural output (since only very minor countries are omitted from the GAEZ or FAO data).

[9] We do this in order to estimate model parameters—the strength of determinants of trade flows and the elasticity of substitution across varieties of a crop, ie σ—using data that is as relevant as possible to the crops in their raw form. These model parameters are never estimated using moments that require us to match the overall level of a crop's exports (ie the exports of both its raw and processed forms).

3.4 Non-Agricultural GDP Data

In order to estimate the value marginal product of the outside sector in each country (i.e. $p_i^0 A_i^0$) we require data on the total value of GDP in the entire economy (in 2009, the same year as the FAO data from above). We obtain this from the World Bank (with the exception of Myanmar, whose GDP data we obtained from the CIA World Factbook).[10]

3.5 Data on Determinants of Trade Costs

A central component of the model introduced above concerns trade costs—that is, the frictions that impede trade between countries. We follow the extensive gravity literature and model trade costs as a function of observed (potential) determinants of trade costs such as distance.[11] We obtain distance measures from the 'Gravity dataset' produced by CEPII.[12]

4 Model Parameter Estimation

To simulate the model described in Section 2, we need estimates of: (i) preference parameters, $\left(\beta_i^k\right)$ and σ in Equations (2) and (1); (ii) technology parameters, $\left(A_i^{fk}\right)$ and θ in Equations (4) and (5); and (iii) trade costs, $\left(\tau_{ij}^k\right)$ in Equation (6). Section 4.1 describes how we estimate each of these parameters. Section 4.2 reports our results. Section 4.3 explores the model's fit given estimated parameters.

4.1 Estimation Procedure

We proceed in three steps.

Step 1: Trade Costs. We use price data to estimate trade costs, τ_{ij}^k. In our dataset, if a country i exports a crop k to another country j, we observe both the producer price in country i, p_i, as well as the unit values of crop k shipped from country i into j, which we use as a proxy for the consumer price of that crop in country j, p_{ij}^k. Using Equation (6), we then compute the log of trade costs as

$$\ln \tau_{ij}^k = \ln p_{ij}^k - \ln p_i^k. \tag{12}$$

[10] We accessed the World Bank GDP data here: http://data.worldbank.org/indicator/NY.GDP.MKTP.CD, and Myanmar's GDP here: https://www.cia.gov/library/publications/the-world-factbook/geos/bm.html.

[11] While it is straightforward to extend the set of determinants of trade costs—for example, to include an indicator for whether trading partners share a common language as in Eaton and Kortum (2002)—for simplicity we focus for now on distance as the sole determiant of trade costs. In future work we aim to relax this constraint.

[12] We access this data from http://www.cepii.fr/anglaisgraph/bdd/gravity.asp.

Many country-pairs and crops in our dataset, however, have zero trade flows. In this case, trade costs are not directly observable. To get around this issue, we assume that trade costs are a log-linear function of distance between countries plus some noise:

$$\ln \tau_{ij}^{k} = \alpha \ln d_{ij} + \delta_{ij}^{k}. \tag{13}$$

We then use observed trade costs for country-pairs and crops with positive trade flows, from Equation (13), to estimate α in Equation (13) by Ordinary Least Squares (OLS). In all subsequent sections, we use $\alpha \ln d_{ij}$ as our preferred measure of trade costs between country i and country j for all crops (whether trade flows are zero or not).

Step 2: Technology. We use output, price, and land data to estimate the extent of within-field heterogeneity θ. Since the productivity of fields across crops, $A_i^{fk} = E\left[A_i^{fk}(\omega)\right]$, is directly observable in the GAEZ data, this is the key technological parameter that needs to be estimated. The basic idea is to find θ such that the output levels predicted by the model, Equation (9), best fits the output levels observed in the data. The only issue is that in order to compute output levels predicted by the model, we need estimates of $p_i^k A_i^{fk}$. For crops, productivity and prices are directly observable, but for the outside good they are not. To infer $p_i^0 A_i^{f0} = p_i^0 A_i^0$, we use the fact that according to our model, the value of output in the outside sector is equal to

$$p_i^0 Q_i^0 = p_i^0 A_i^0 L_i^0,$$

where $L_i^0 \equiv \sum_{f \in \mathcal{F}_i} \left(\frac{\left(p_i^0 A_i^0\right)^{\theta}}{\sum_{l \in \mathcal{K}} \left(p_i^l A_i^{fl}\right)^{\theta}} \right)^{(\theta-1)/\theta}$ is equal to the amount of land allocated to the outside sector. In our model, $p_i^0 Q_i^0$ is also equal to total income in country i minus the total value of crops produced in that country, $\sum_{k \neq 0} p_i^k Q_i^k$. So we can measure $p_i^0 A_i^0$ as GDP in country i minus the total crop value divided by total acres of land allocated to the outside sector, which are all observable in the data. Given $p_i^0 A_i^0$, as well as data on output, Q_i^k, crop prices, p_i^k, and fields productivity, A_i^{fk}, we use Non-Linear Least Squares to estimate θ as the solution of

$$\min_{\theta} \sum_{i,k \neq 0} \left(\ln \tilde{Q}_i^k(\theta) - \ln Q_i^k \right)^2, \tag{14}$$

where $\tilde{Q}_i^k(\theta)$ is the output level predicted by our model for a given value of θ, i.e.,

$$\tilde{Q}_i^k(\theta) = \sum_{f \in \mathcal{F}_i} A_i^{fk} \left(\frac{\left(p_i^k A_i^{fk}\right)^{\theta}}{\sum_{l \in \mathcal{K}} \left(p_i^l A_i^{fl}\right)^{\theta}} \right)^{(\theta-1)/\theta}.$$

Step 3: Preferences. We start by using trade data and our estimates of trade costs to estimate the elasticity of substitution σ between crops from different countries. Let $X_{ij}^k = \left(\tau_{ij} p_i^k\right) C_{ij}^k$ denote the value of exports of crop k from country i to country j. We assume that trade flows are observed with measurement error so that Equation (8) implies

$$\ln X_{ij}^k = E_i^k + M_j^k + (1 - \sigma) \ln \tau_{ij}^k + \eta_{ij}^k, \tag{15}$$

where $E_i^k \equiv (1 - \sigma) \ln p_i^k$ can be treated as an exporter fixed effect; $M_j^k = \ln\left(\beta_j^k Y_j\right) - \ln\left(\sum_{n=1}^{I} \left(\tau_{nj} p_n^k\right)^{1-\sigma}\right)$ can be treated as an importer fixed effect; and η_{ji}^k is the measurement error in trade flows referred to above. We obtain our estimate of σ by estimating Equation (15) using OLS. In principle we can estimate a separate elasticity of substitution across varieties within each crop k (that is a separate σ^k for all k) but for simplicity we first focus on one pooled estimate of σ that is the same across all crops. To conclude, we use trade and output data to measure the share of expenditures β_i^k across goods in different countries. For each crop $k = 1, ..., K$, we compute total expenditure S_i^k on crop k in country i as $\sum_{j \in \mathcal{I}} X_{ji}^k$, where total imports, $\sum_{j \neq i} X_{ji}^k$, are directly observable in the data and the value of domestic consumption, X_{ii}^k, is computed as the value of output minus exports, $p_i^k Q_i^k - \sum_{j \neq i} X_{ij}^k$. Given total expenditures for all crops $k = 1, ..., K$ and countries, we can compute β_i^k as the ratio of S_i^k over GDP in country i. Expenditure shares on the outside good are then given by one minus $\sum_{k \neq 0} \beta_i^k$.

4.2 Estimation Results

We now discuss a preliminary set of parameter estimates obtained using the procedure outlined above. To reduce computational complexity in the counterfactual simulations below we focus on a limited set of the 50 largest countries and 10 largest crops, in terms of value of output. Because of the skewness of economic activity in agriculture, across countries and crops, this dataset still spans over ninety percent of world agricultural GDP and trade. We also scale down the number of fields by a factor of 4 (that is, each grid cell is now 2 times larger in width and height).

Based on this augmented set of data, we obtain for the model's three parameters, presented in Table 1. The first parameter estimate is α, which governs the elasticity of trade costs with respect to distance. This is approximately $\alpha = 0.11$ ($SE = 0.003$), which is in line with standard estimates in the empirical trade literature but estimated using different methodologies and based on a manufacturing sample. The second parameter estimate is θ, which governs the within-field, within-crop productivity heterogeneity in agriculture (through its inverse effect on the dispersion of the productivity distribution). Within any given field it is the elasticity of relative supply (across any two crops) to relative prices. We find that, approximately, $\theta = 2.39$, which suggests that within-field heterogeneity is sub-

Table 1: Parameter estimates

Parameter	Description	Parameter estimate	Parameter standard error
α	Elasticity of trade costs with respect to distance	0.110	(0.003)
θ	Within-field heterogeneity dispersion (and within-field elasticity of substitution in supply)	2.391	N/A
σ	Elasticity of substitution in demand (across varieties of a crop)	17.864	(0.772)

Notes: Parameter estimates using method described in Section 4. Standard errors for θ to be computed in future work using a bootstrap procedure.

stantial.[13] Finally, the third parameter we estimate is σ, the elasticity of substitution across varieties of a crop (within any given crop). We find that $\sigma = 17.86$ ($SE = 0.77$), which is a very high elasticity of substitution, perhaps to be expected for the case of relatively homogenous agricultural goods.

In short we find it reassuring that these parameter estimates are of plausible magnitudes and that, where standard errors are available, they are very precisely estimated.

4.3 Model Fit

It is natural to ask, before we go on to considering how our model behaves under the counterfactual scenario of new agricultural productivities brought about by climate change, how well the model fits the data within sample. Figure 2 (a) plots the fit of the moment that we use to estimate θ, namely a comparison between log output in the model and in the data at our preferred estimate of $\theta = 2.39$. There is a positive and statistically significant correlation between the model and the data (a regression of the former on the latter, with a constant, yields a coefficient estimate of 0.311 ($SE = 0.037$[14]). For comparison Figure 2 (b) plots the analogous figure but for a naive case where we allow all fields within each country to be identical and to have yields, in each crop, equal to the country-specific average yield (and where we re-estimate θ based on this new productivity data). As the comparison

[13] We do not yet know the standard deviation of our estimate of θ. While there is no closed form for this standard error in future work we will use a bootstrap procedure to estimate it.

[14] This and all other standard errors referred to in this section are clustered at the country level.

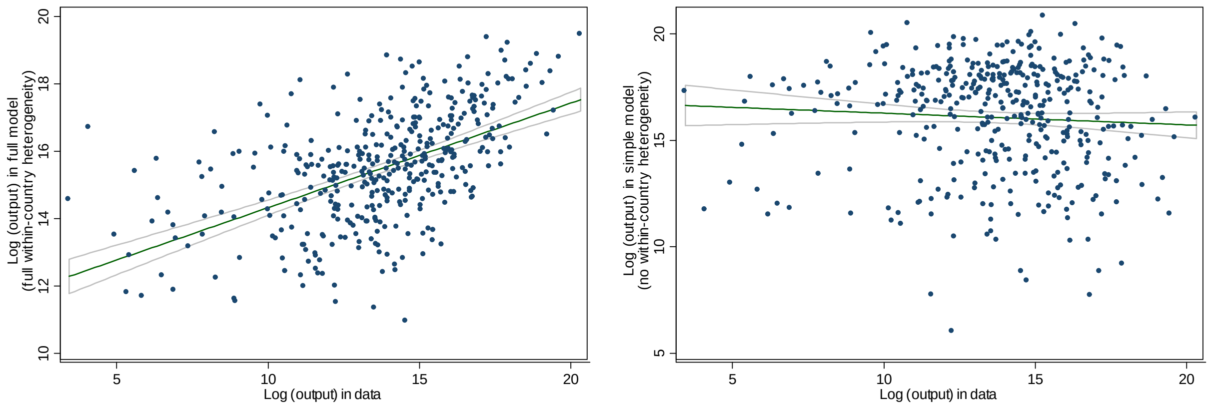

Figure 2(a) Figure 2(b)

Figure 2: Comparison between log output computed in the model (y-axis) and FAO data (x-axis), across all crops and countries. Panel (a) reports log output calculated using the full model outlined in Section 2 and the full GAEZ data (hence there is significant within-country heterogeneity). Panel (b) reports log output calculated using the full model outlined in Section 2 but based on simplified GAEZ data (with each field's productivity replaced with its country's average in the GAEZ data, removing all within-country heterogeneity). Best fit line and 95% confidence interval are also indicated.

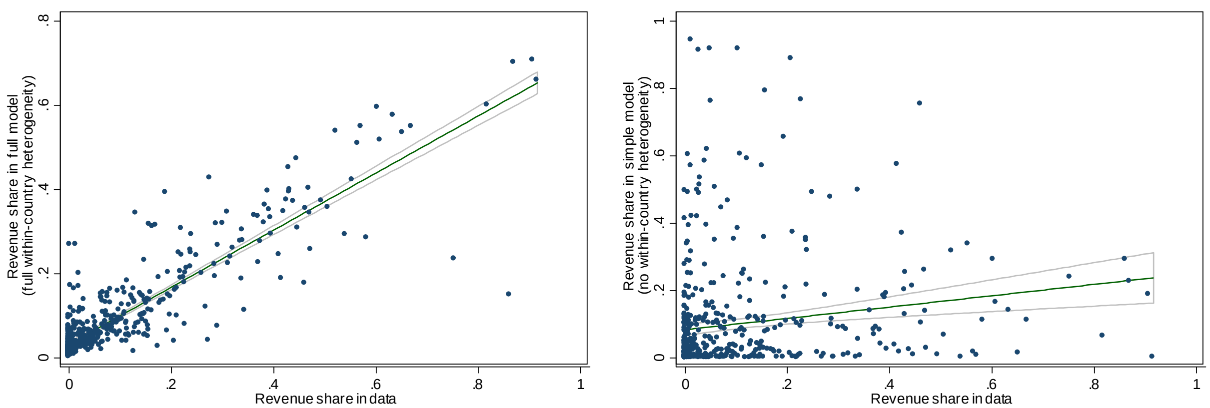

Figure 3(a) Figure 3(b)

Figure 3: Comparison between revenue shares computed in the model (y-axis) and FAO data (x-axis), across all crops and countries. Panel (a) reports revenue shares calculated using the full model outlined in Section 2 and the full GAEZ data (hence there is significant within-country heterogeneity). Panel (b) reports revenue shares calculated using the full model outlined in Section 2 but based on simplified GAEZ data (with each field's productivity replaced with its country's average in the GAEZ data, removing all within-country heterogeneity). Best fit line and 95% confidence interval are also indicated.

between Figures 2 (a) and 2 (b) makes clear, the fit of our model improves substantially when intra-national productivity heterogeneity is allowed for; that is, a 'representative field' analysis based on the GAEZ data would be an abject failure. (A regression of log model output on log actual output using a naive, no within-country heterogeneity model based on averaging GAEZ data delivers a coefficient of -0.055 ($SE = 0.058$).

While the fit of the model in terms of log output, illustrated in Figure 2 (a), indicates that this model is capable of capturing, with some accuracy, the pattern of international specialization, it is also clear that the absolute level of output in the model does not fit that in the data particularly well. (For example, the estimated constant in the regression illustrated in Figure 2 (a) is equal to 11.2, implying that predicted output is considerably higher than actual output.) This is presumably not a first-order concern given that our analysis focuses on changes in output due to climate change, rather than any absolute level of output, it is likely that this inability to match output levels stems from our assumption that agricultural technologies do not differ around the world. (See Section 6 below for a lengthier discussion of this and a proposed solution that is in progress.)

Figure 3(a), however, illustrates how the fit of the model is considerably more successful in terms of matching relative crop production, i.e. the pattern of specialization that is at the heart of comparative advantage and therefore at the heart of our analysis here. In Figure 3 (a) we plot the predicted revenue share, for each crop and country, as predicted by the model, against the equivalent revenue share in the FAO data. (That is, in the case of the model revenue shares on the y-axis we use the model's equilibrium price and quantity for a crop and divide by the model's total revenue amongst all crops. The data revenue shares on the x-axis are computed analogously but using only the FAO price and quantity data.) Here the fit is considerably better than when evaluating log output in Figure 2 (a). The line of best fit has a slope coefficient of 0.678 ($SE = 0.042$) and an estimated constant of 0.032 ($SE = 0.004$).; the R-squared for this regression is 0.76. Again by way of comparison we plot (in Figure 3 (b)) the equivalent fit for a simpler model in which we ignore all within-country heterogeneity in the GAEZ data. Here the slope coefficient is just 0.168 ($SE = 0.044$) and the R-squared is 0.03..We see once again that allowing the wtihin-country data to speak to the model is vital in improving its fit to the real world.

Finally, Table 2 reports the coefficient estimates from regressions such as those illustrated in Figures 2 (a) and 3 (a), for these and other variables in the FAO data. We look at, respectively, revenue shares (as in Figure 3 (a)), log revenues, log output (as in Figure 2 (a)), land shares, and log producer prices. A pattern emerges that is similar to that discussed above. Variables in levels correlate well between the data and the model (that is, the estimated slope coefficient is positive and very precisely estimated) but the absolute level is off (that is, the estimated constant is far from zero.) But what is encouraging here, given our focus on changes due to climate change (not the absolute level of output before or after climate change), and given our focus on comparative advantage, is how the two variables based on shares—revenue shares in column (1) and land shares in column (4)—agree between the model and the data.

5 Counterfactual Simulations

5.1 Consequences of Climate Change

We model climate change as a change in crop productivity from $\left(A_i^{fk}\right)$, as measured in the GAEZ data baseline scenario, to $\left(A_i^{fk}\right)'$, as measured in the GAEZ data under the climate

Table 2: Model fit

	depdendent variable: variable 'X' in model, where variable 'X' is…				
	revenue share	log (revenue)	log (output)	land share	log (producer price)
	(1)	(2)	(3)	(4)	(5)
Variable 'X' in FAO data	0.6781***	0.3153***	0.3105***	1.0035***	0.1954***
	(0.0417)	(0.0387)	(0.0367)	(0.0223)	(0.0185)
Constant	0.0322***	17.2779***	11.2136***	-0.0003	6.8681***
	(0.0042)	(0.7674)	(0.5552)	(0.0020)	(0.1048)
Number of observations	500	406	406	550	500
R-squared	0.786	0.291	0.288	0.947	0.099

Notes: Each column is a regression of variable 'X', as computed in the full model (outlined in Section 2) using the full GAEZ data (with full within-country heterogeneity), on variable 'X' in the FAO data. Variable 'X' varies from column to column. In column (1) variable 'X' is the revenue share for each crop in total crop GDP for each country. In column (2) variable 'X' is log(revenue) for each non-zero crop in each country. In column (3) variable 'X' is log(output) for each non-zero crop in each country. In column (4) variable 'X' is the share of land devoted to each crop in each country. And in column (5) variable 'X' is the log (producer price). Standard errors are clustered at the country level.

change scenario. All other structural parameters are held fixed at the values estimated in Section 4. Equilibrium conditions are still given by Equations (7)-(11).

We focus on changes in real income, $W_i \equiv Y_i/P_i$, where $Y_i \equiv \sum_{k\in\mathcal{K}} p_i^k Q_i^k$ denotes total income in country i and P_i denotes the consumer price index. Given our preference structure, Equations (1) and (2), the consumer price index can be computed as $P_i = \prod_{k=0}^{K} \left(P_i^k\right)^{\beta_i^k}$, with the component of the price index associated with crop k given by $P_i^k = \left(\sum_{j=1}^{I} \left(\tau_{ji} p_j^k\right)^{1-\sigma}\right)^{1/(1-\sigma)}$. The second column of Table 3 reports the percentage changes in real income for the (unweighted) world average. (Appendix Table A reports the results from all 50 countries used in our preliminary analysis.) The average world country (in our sample of the 50 largest agricultural countries) will be hurt by climate change (in a real income sense) though naturally there are exceptions. The world average estimated real income loss is approximately 13 percent of agricultural expenditure (which, as reported in column (1), is 7.75 percent of all expenditure, for the world average).

5.2 Gains from Local Specialization

To quantify the importance local specialization—i.e., how changes in comparative advantage at the field level affect the consequences of climate change—we recompute the equilibrium with climate change under the assumption that the allocation of all fields to all goods in all

Table 3: Counterfactual simulation results

	Share of Expenditure on Crops in Sample	Change in real income (as share of crop expenditure) due to climate change under… Trade costs at baseline, full output adjustment	Trade costs at baseline, no land share adjustment	Autarky, full land share adjustment	Gains (as share of crop expenditure) due to… Local Specialization = (2)-(3)	International Specialization = (2)-(4)
	(1)	(2)	(3)	(4)	(5)	(6)
World Average	7.8%	-13.0%	-37.6%	-18.7%	24.6%	5.7%

Notes: Column (1) reports the share (of total expenditure) in each country spent on the crops in our sample. Column (2) reports the change in real income (expressed as a share of expenditure on crops in our sample), between the model under baseline and the model under climate change, when trade costs are at the level estimated in the baseline sample (as discussed in Section 4). Column (3) reports the analogous result to column (2) only for the case in which fields cannot change the share of land allocated to each good. Column (4) reports the analogous result to column (2) only for the case where trade costs are set to infinity (in both baseline and under climate change). Columns (5) and (6) report differences between columns (2) and (3), and (2) and (4), respectively.

countries is the same as in the in initial equilibrium without climate change. This implies that total output of good k in country i in the counterfactual equilibrium is given by

$$\left(Q_i^k\right)' = \Gamma\left(\frac{\theta-1}{\theta}\right)\exp(-\frac{e}{\theta})\sum_{f\in\mathcal{F}_i}\left(A_i^{fk}\right)'\pi_i^{fk}, \tag{16}$$

where $\left(A_i^{fk}\right)'$ still denotes productivity under the climate change scenario, but $\pi_i^{fk} = \frac{\left(p_i^k A_i^{fk}\right)^\theta}{\sum_{l\in\mathcal{K}}\left(p_i^l A_i^{fl}\right)^\theta}$ corresponds to the share of parcels in a field f located in country i that are allocated to a good k in the initial equilibrium without climate change. The other equilibrium conditions (7), (8), (10), and (11) are unchanged.

Column (3) of Table 3 reports the change in real income (again relative to total agricultural expenditure) for the average country from climate change under this scenario. Here the ill-effects of climate change are, as to be expected, considerably worse than in the full-adjustment scenario reported in column (3). Column (5) reports what we refer to as the gains from local specialization, i.e., the difference between the full local adjustment scenario (column (2)) and the no-local adjustment scenario (column (3)). These are positive and substantial.

5.3 Gains from International Specialization

To quantify the importance of international specialization—i.e., how the possibility to exploit changes in comparative advantage across countries through international trade affects the consequences of climate change—we recompute the equilibrium with and without climate change under autarky. The equilibrium conditions are the same as in Section 2.2, except for the good market clearing conditions, which are now given by

$$Q_i^k = C_i^k, \text{ for all } i \in \mathcal{I} \text{ and } k \in \mathcal{K}.$$

The fourth column of Table 3 reports changes in real income caused by climate change in the absence of international. We see that the reductions to real income due to climate change, under autarky, can be substantial. The difference between column 2 and column 4 (reported in column 6) captures what we refer to as the gains from international specialization. At the (unweighted) world average level, we estimate that the ill-effects of climate change are equal to approximately 19 percent of total agricultural expenditure in autarky, but only to approximately 13 percent under restricted trade. That is, the world average country is affected by climate change to an extent that is about one-third less severe when it is able to trade (even with significant international trade costs) than when it is under autarky. But it is important to note that, due to adverse terms-of-trade effects, not all countries benefit from restricted trade, as shown in Appendix Table A.

6 Sensitivity Analysis (in progress)

We now describe a number of extensions to the analysis presented above that are under preparation.

6.1 Productivity Measures

In our baseline analysis, we assume that GAEZ data perfectly predict productivity across crops and fields. In practice, they do not. In order to take the imperfect fit of the GAEZ data into account, we now assume that

$$A_i^{fk} = \hat{A}_i^{fk} \times T_i^k,$$

where $\hat{A}_i^{fk}$ is the measure in the GAEZ data of the productivity field f in country i if it were to produce crop k and T_i^k is some technological shock—unpredicted by agronomists—that affects the productivity of all fields in country i for crop k. In order to estimate T_i^k, the basic idea is to use data on land shares. If a crop receives a relatively higher land share than what our model predicts given observed prices and productivity in the GAEZ data, then its true productivity, and hence its T_i^k, must be relatively higher than the true productivity of other crops.

6.2 Outside Good

In our baseline analysis, we assume that the outside good is non-tradable. While this assumption seems reasonable if one interprets the outside as residential housing or services, it is less so if one interprets it as forestry or manufacturing. In this subsection, we explore the polar case in which the outside sector is assumed to be freely traded around the world at a common price, p^0, though differences in productivity across countries, A_i^0, may now affect the value of the marginal product of land around the world.

6.3 Substitution Between Crops

In our baseline analysis, we assume that crops enter the upper-level utility function in a Cobb-Douglas manner. We now consider the case nested of CES utility functions:

$$U_i = \left(C_i^0\right)^{\beta_i^0} \left(\sum_{k=1}^{K} \left(\frac{\beta_i^k}{1-\beta_i^0}\right) \left(C_i^k\right)^{(\gamma-1)/\gamma}\right)^{\frac{\left(1-\beta_i^0\right)\gamma}{1-\gamma}}, \tag{17}$$

$$C_i^k = \left(\sum_{j=1}^{I} \left(C_{ji}^k\right)^{(\sigma-1)/\sigma}\right)^{\sigma/(\sigma-1)}, \text{ for all } k = 1, ..., K, \tag{18}$$

where $\gamma > 1$ is the elasticity of substitution between crops. We plan to estimate γ through non-linear least squares in the same way as we have estimated θ on the supply-side of our model in Section 4.1.

7 Concluding Remarks

A large agronomic literature has modeled the implications of such climate change for crop yields, crop by crop and location by location. These studies document the harm that may be inflicted on a specific crop at a specific location. The goal of this paper has been to move beyond these micro-level studies and aggregate them together into a coherent, macro-level understanding of how climate change will affect agricultural markets.

Aggregating micro-level impacts in a globalized world means that impacts depend on the simple economics of comparative advantage—that is, the impact of micro-level shocks do not only depend on their average level, but also on their dispersion over space. To measure the impact of climate change at the micro-level we draw on an extremely rich dataset that contains agronomist's estimates about the productivity—both before and after climate change—of each crop for each of over 9 million high resolution grid cells covering the surface of the Earth. Crucially, the same agronomic model is used to generate both the pre-climate change and post-climate estimates; all that changes in the agronomist's calculations is the climatic data that enters their models, which is drawn from leading climatological models of climate change.

Using a general equilibrium model of trade among these 9 million grid cells we have found that international trade, even when subject to reasonable levels of trade costs (designed to match the data within-sample), will substantially mitigate the ill-effects of climate change on agricultural living standards in most countries. Put simply, this is because, according to agronomic estimates, climate change is expected to harm agricultural plants but to do so in a manner that exaggerates agricultural yield dispersion across space.

References

Babiker, M. H. (2005): "Climate Change Policy, Market Structure, and Carbon Leakage," *Journal of International Economics*, 65, 421–445.

Costinot, A., and D. Donaldson (2011): "How Large Are the Gains from Economic Integration? Theory and Evidence from U.S. Agriculture, 1880-2002," *mimeo MIT*.

Cristea, A., D. Hummels, L. Puzzello, and M. Avetysyan (forthcoming): "The Contribution of International Transport to Global Greenhouse Gas Emissions," *Journal of Environmental Economics and Management*.

Eaton, J., and S. Kortum (2002): "Technology, Geography and Trade," *Econometrica*, 70(5), 1741–1779.

Elliott, J., I. Foster, S. Kortum, T. Munson, F. P. Cervantes, and D. Weisbach (2010): "Trade and Carbon Taxes," *American Economic Review Papers and Proceedings*, 100(2), 465–469.

Felder, S., and T. Rutherford (1993): "Unilateral CO2 Reductions and Carbon Leakage: the Consequences of Trade in Oil and Basic Materials," *Journal of Environmental Economics and Management*, 25, 162–176.

Hertel, T., and T. Randhir (2000): "Trade Liberalization as a Vehicle for Adapting to Global Warming," *Agriculture and Resource Economics Review*, 29(2), 1–14.

Reilly, J., and N. Hohmann (1993): "Climate Change and Agriculture: The Role of International Trade," *American Economic Review Papers and Proceedings*, 83(2), 306–312.

Rosenzweig, C., and M. Parry (1994): "Potential Impact of Climate Change on World Food Supply," *Nature*, 367(133-138).

Shapiro, J. (12): "Trade, Oil, and the Environment," *mimeo*.

Tsigas, M., G. Friswold, and B. Kuhn (1997): *Global Climate Change and Agriculture*, Global Trade Analysis: Modeling and Applications. Cambridge University Press.

Appendix Table A: Counterfactual simulation results by country

	Change in real income (as share of crop expenditure) due to climate change under…			Gains (as share of crop expenditure) due to…	
	Trade costs at baseline, full output adjustment	Trade costs at baseline, no output adjustment	Autarky, full output adjustment	Local Specialization = (1)-(2)	International Specialization = (1)-(3)
Country	(1)	(2)	(3)	(4)	(5)
Algeria	-0.7%	-1.1%	-0.6%	0.43%	-0.05%
Angola	-0.7%	-0.8%	-0.8%	0.13%	0.09%
Argentina	0.1%	-1.3%	0.3%	1.45%	-0.15%
Australia	-0.1%	-0.3%	-0.1%	0.21%	0.00%
Bangladesh	-3.8%	-4.3%	-4.6%	0.57%	0.79%
Brazil	-1.8%	-2.3%	-1.8%	0.53%	0.01%
Canada	0.2%	0.1%	0.2%	0.06%	-0.01%
China	0.4%	-0.4%	0.4%	0.82%	0.00%
Colombia	-1.4%	-1.8%	-1.3%	0.39%	-0.14%
Congo, the DRC	-2.9%	-4.0%	-2.8%	1.09%	-0.09%
Egypt	-0.2%	-0.2%	-0.3%	0.01%	0.03%
Ethiopia	-1.9%	-4.7%	-3.3%	2.76%	1.33%
France	-0.2%	-0.3%	-0.1%	0.07%	-0.08%
Germany	0.0%	-0.1%	0.0%	0.06%	-0.01%
Ghana	-2.3%	-2.9%	-2.5%	0.65%	0.27%
Greece	-0.3%	-0.5%	-0.2%	0.20%	-0.07%
Guatemala	-4.6%	-5.9%	-5.1%	1.28%	0.54%
India	-1.2%	-2.0%	-1.3%	0.86%	0.16%
Indonesia	-0.6%	-1.5%	-0.6%	0.83%	-0.02%
Iran	-0.5%	-0.8%	-0.4%	0.35%	-0.01%
Italy	-0.3%	-0.4%	-0.2%	0.17%	-0.10%
Japan	0.2%	-0.1%	0.1%	0.31%	0.07%
Kazakhstan	-0.4%	-0.8%	-0.4%	0.38%	0.01%
South Korea	-0.1%	-0.5%	0.1%	0.43%	-0.19%
Malawi	-19.3%	-25.6%	-38.8%	6.37%	19.57%
Malaysia	-0.3%	-0.7%	-0.3%	0.42%	-0.07%
Mexico	-0.2%	-0.5%	-0.2%	0.31%	0.00%
Morocco	-1.7%	-2.6%	-1.7%	0.92%	0.04%
Myanmar	-2.6%	-6.6%	-2.1%	3.97%	-0.47%
Nigeria	-2.5%	-3.4%	-2.1%	0.93%	-0.42%
Pakistan	-0.6%	-2.2%	-0.7%	1.61%	0.09%
Peru	-0.7%	-1.5%	-0.1%	0.80%	-0.58%
Philippines	-0.8%	-1.3%	-0.6%	0.55%	-0.14%
Poland	0.1%	0.0%	0.3%	0.12%	-0.17%
Romania	-0.4%	-0.7%	-0.5%	0.29%	0.14%
Russia	0.4%	0.1%	0.4%	0.27%	-0.03%
Serbia And Montenegro	-0.4%	-0.8%	-0.4%	0.36%	-0.01%
South Africa	-0.3%	-0.4%	-0.2%	0.09%	-0.08%
Spain	-2.4%	-3.5%	-2.3%	1.06%	-0.05%
Sudan	-0.2%	-1.2%	-0.5%	1.03%	0.33%
Syria	-1.5%	-2.6%	-2.3%	1.10%	0.85%
Tanzania	-2.2%	-2.9%	-2.2%	0.69%	0.04%
Thailand	0.4%	-0.1%	0.5%	0.48%	-0.11%
Turkey	-0.4%	-0.6%	-0.4%	0.23%	0.02%
Ukraine	0.3%	0.3%	0.1%	0.08%	0.19%
United Kingdom	0.0%	-0.2%	0.0%	0.18%	0.03%
United States	2.1%	-1.4%	-0.1%	3.52%	2.18%
Uzbekistan	-2.4%	-2.8%	-2.5%	0.40%	0.10%
Venezuela	-2.3%	-5.4%	-2.4%	3.11%	0.14%
Vietnam	-1.3%	-2.2%	-2.0%	0.91%	0.68%