

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
Brookings Institution

November 2012

Abstract

A large agronomic literature has modeled the implications of climate change for a variety of crops and locations around the world. 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 does 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 10 crops for each of over 9 million high resolution grid cells covering the surface of the Earth, we find small adverse effects of climate change for the median country in the world. While international trade plays virtually no role in explaining the magnitude of these effects, our analysis suggests that reallocations caused by the evolution of comparative advantage within countries substantially mitigate the ill-effects of climate change.

*Authors' email addresses: costinot@mit.edu, ddonald@mit.edu and corybsmith@gmail.com. For helpful suggestions and comments we are grateful to Sam Kortum, Rob Townsend, and seminar audiences at the IIES Conference on "Climate and the Economy" and at Princeton University. Costinot and Donaldson thank the National Science Foundation (under Grant SES-1227635) for research support. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Brookings Institution.

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. 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, the impact of micro-level shocks does 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, both within and between countries, and welfare will further decrease. If climate change instead raises heterogeneity across regions, there will be more scope for 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. Yet, whether climate change will weaken or strengthen comparative advantage, both within and between countries, remains an open question.

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—crop by crop—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 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.¹

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

¹After elimination of water areas and land areas outside of the 50 major agricultural countries we include in our study (which account for over 90 percent of world crop output), our final dataset includes approximately 1.7 million grid-cells.

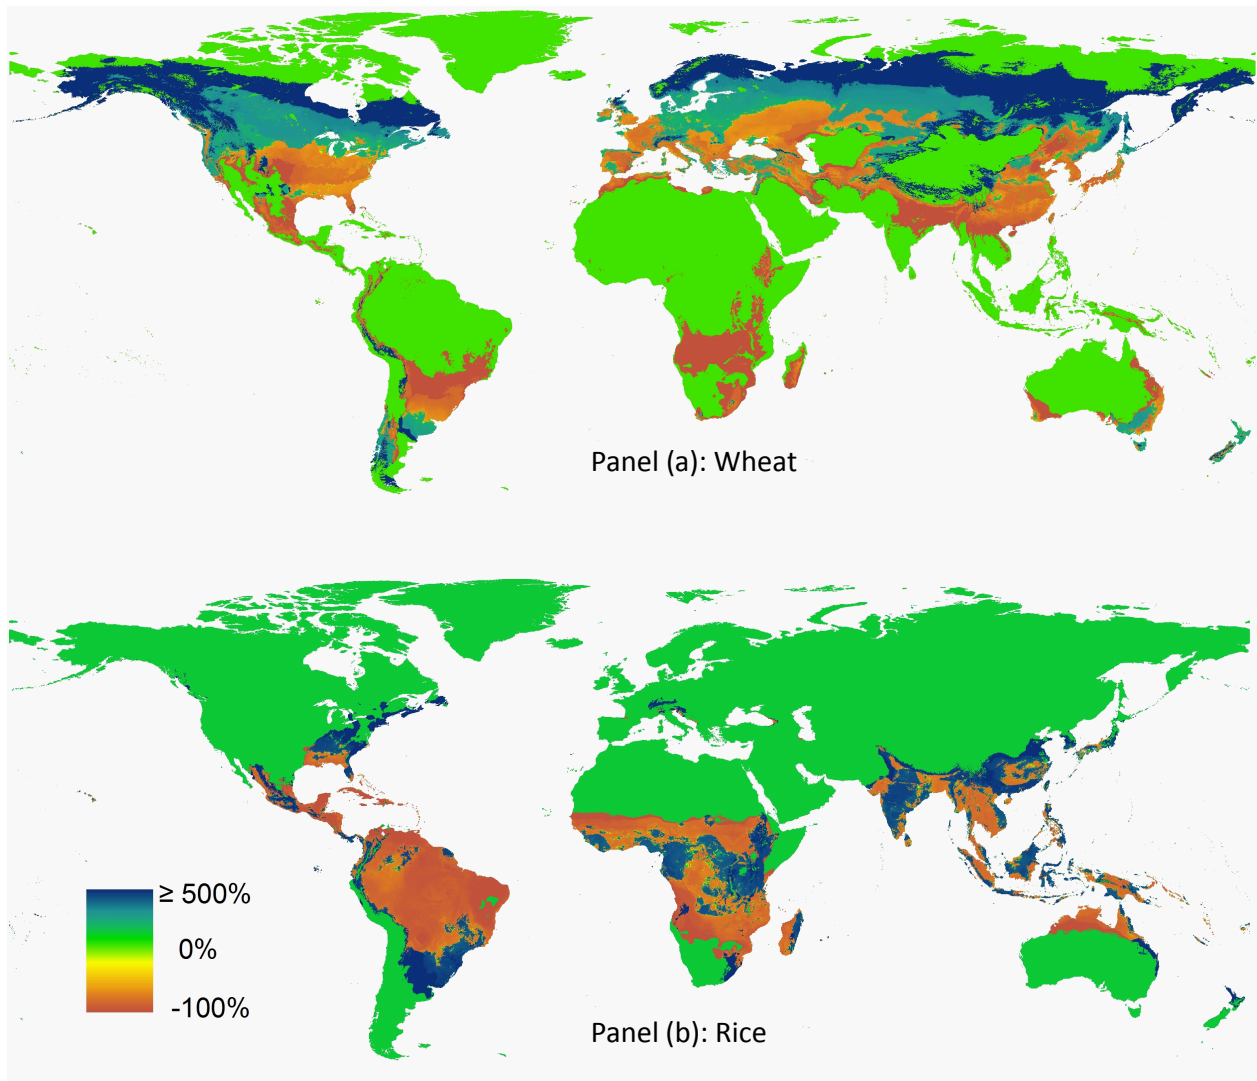


Figure 1: Percent changes in yield due to climate change in GAEZ model for wheat and rice. Large green areas are those for which yields are zero both before and after climate change.

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.

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. At the estimated parameter values, we find that the within-sample fit of our model is particularly strong for a moment closely connected to the underlying pattern of comparative advantage, revenue shares by crop and country, although this moment was not directly used in our estimation procedure.

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 countries are free to trade (subject to our estimated trade costs) and farmers are free to change their output decisions. Under this scenario, we find small adverse effects of climate change. A county like Malaysia, which experiences the median impact of climate change in our model, would see its real income decrease by 0.19% after climate change. When expressed as a share of agricultural expenditure, this still represents only a 5.25% loss.

As mentioned above, a potential reason why the adverse effects of climate change may be smaller in a global economy is because higher temperatures may lead to more heterogeneity between countries, which countries may take advantage of by trading more internationally. To explore the quantitative importance of this economic channel, we consider a second counterfactual scenario in which farmers can reallocate production, but countries are under autarky. The welfare consequences of climate change under this new scenario are of the same order of magnitude for most countries. Malaysia, for instance, would experience a welfare loss equal to 2.89% (again as a share of agricultural expenditure) if it were under autarky compared to 5.26% under trade. If anything, the adverse consequences of climate change are more severe in a world in which countries can trade.²

²Of course, this does not imply that the world economy would be better off without international trade.

Although international specialization cannot account for the small welfare effects of climate change predicted by our model, our previous results leave open the possibility that intranational reallocations caused by changes in comparative advantage may significantly affect the consequences of climate change. To shed light on this mechanism, we consider a final counterfactual scenario in which countries can trade, but farmers cannot reallocate production. Under this scenario, we find that the adverse welfare consequences of climate change are an order of magnitude larger than in the two previous scenarios. For Malaysia, the welfare loss jumps to 39.7%. This illustrates how farmers’ ability to substitute crop production in response to changes in comparative advantage—which our extremely rich micro-level dataset gives us a unique opportunity to study—may substantially mitigate the ill-effects of climate change.

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, Puzello, and Avetysyan (forthcoming) and Shapiro (2012). 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 more than a million fields around the world. 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 illustrates the connection between

There are always gains from international in the neoclassical environment that we consider. Our quantitative results merely point out that climate change tends to make such gains smaller.

trade, comparative advantage, and climate change through a simple example. Section 3 develops our theoretical framework. Section 4 presents the data that feeds into our analysis. Section 5 describes our estimation procedure, our parameter estimates, and measures of goodness of fit of the model. Section 6 then presents the results of our counterfactual simulations. Finally, Section 7 describes some robustness extensions that are in progress and Section 8 offers some concluding remarks.

2 A Simple Example

Consider an economy comprising two symmetric islands, North (N) and South (S). Each island consists of one acre of land that can be used to produce two crops, rice (R) and wheat (W). Consumers in both islands spend half of their income on rice and half of their income on wheat. Both crops are produced by a large number of farms under perfect competition. Total output of crop $k = R, W$ in island $i = N, S$ is given by $Q_i^k = A_i^k L_i^k$, where A_i^k denotes the exogenous productivity per acre for crop k in island i and L_i^k denotes the endogenous share of land allocated to that crop. The only difference between North and South is that North has a comparative advantage in wheat:

$$A_N^W/A_N^R > A_S^W/A_S^R,$$

where, by symmetry, land productivity in the two islands satisfy $A_N^W = A_S^R$ and $A_N^R = A_S^W$.

Now suppose that because of a sudden rise in temperature, land productivity goes from its initial value A_i^k to some new value $(A_i^k)'$. Would the welfare consequences of climate change in the two islands be more or less severe in a world in which North and South can trade with one another compared to a world in which they cannot?

Since the land allocation is efficient under perfect competition, we know that, up to the first-order approximation, the average log-change in real income caused by climate change under autarky and free trade, Δ_A and Δ_T , must be equal to the weighted sum of the log-change in productivity, that is

$$\begin{aligned}\Delta_A &= \frac{1}{2} \sum_{i=N,S} \sum_{k=R,W} (\lambda_i^k)_A \Delta \ln A_i^k, \\ \Delta_T &= \frac{1}{2} \sum_{i=N,S} \sum_{k=R,W} (\lambda_i^k)_T \Delta \ln A_i^k,\end{aligned}$$

where $\Delta \ln A_i^k \equiv \ln (A_i^k)' - \ln A_i^k$ denotes the log-change in productivity caused by climate change and $(\lambda_i^k)_A$ and $(\lambda_i^k)_T$ denote the share of island's i total income associated with the production of crop k in the autarky and free trade equilibria, respectively. Under autarky, the share of a crop k in revenues must be equal to its share in expenditure. Thus, $(\lambda_i^k)_A = 1/2$

for all i and k . Under free trade, the pattern of comparative advantage between North and South implies complete specialization: $(\lambda_N^W)_T = (\lambda_S^R)_T = 1$, whereas $(\lambda_N^R)_T = (\lambda_S^W)_T = 0$. Accordingly, the differential effect of climate change across the two trading regimes, $\Delta \equiv \Delta_T - \Delta_A$, can be expressed as

$$\Delta = \frac{1}{4} \left((\Delta \ln A_N^W - \Delta \ln A_N^R) - (\Delta \ln A_S^W - \Delta \ln A_S^R) \right). \quad (1)$$

According to Equation (1), trade alleviates the adverse consequences of climate change, $\Delta > 0$, if and only if it strengthens comparative advantage, in the sense that

$$\left((A_N^W)' / (A_N^R)' \right) / \left((A_S^W)' / (A_S^R)' \right) > (A_N^W / A_N^R) / (A_S^W / A_S^R).$$

Productivity shocks that are either crop-specific or island-specific in contrast, have exactly the same effects under autarky and under free trade. In words, whether trade between different locations is important (or not) to study the macro-consequences of climate change crucially depends on whether climate change affects comparative advantage.

The intuition is simple. Consider, for instance, a weakening of the pattern of comparative advantage in the form of a 10% decrease in wheat productivity in the Northern island. Such a shock would reduce *Northern* real income by 5% under autarky, whereas it would reduce *World* real income by 5% under free trade. So the negative consequences of climate change are twice as large in the latter case. The opposite would be true if we were to consider a strengthening of the pattern of comparative advantage in the form of a 10% decrease in rice productivity in the Northern island. Such a shock would also reduce Northern real income by 5% under autarky, but it would have no effect on real income under free trade. Under this alternative scenario, the negative consequences of climate change are less severe—and indeed completely alleviated—when islands are allowed to trade with one another.

The rest of our analysis aims to explore the quantitative importance of these theoretical considerations in practice. Compared to the simple model presented in this section, we will now allow for demand differences between countries, trade costs, more than two crops, and most importantly, more than two types of land in the form of 9 million grid-cells.

3 Theory

3.1 Basic Environment

We consider a world economy comprising multiple countries, indexed by $i \in \mathcal{I} \equiv \{1, \dots, I\}$. In each country, the only factors of production are fields, indexed by $f \in \mathcal{F}_i \equiv \{1, \dots, F_i\}$, 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 $k \in \mathcal{K} \equiv \{0, \dots, K\}$. Goods $1, \dots, K$ are crops, whereas good 0 will be an outside good. We think of the outside good as residential housing, services, manufacturing, forestry 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:

$$U_i = \prod_{k=0}^K (C_i^k)^{\beta_i^k}, \quad (2)$$

$$C_i^k = \left(\sum_{j=1}^I (C_{ji}^k)^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \text{ for all } k = 1, \dots, K, \quad (3)$$

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 $k = 1, \dots, K$ 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

$$Q_i^k = \sum_{f \in \mathcal{F}_i} \int_0^1 A_i^{fk}(\omega) L_i^{fk}(\omega) d\omega, \quad (4)$$

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

$$\ln A_i^{fk}(\omega) = \ln A_i^{fk} + \varepsilon_i^{fk}(\omega). \quad (5)$$

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 $k = 1, \dots, K$ 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:

$$F(\varepsilon) = \Pr \left[\varepsilon_i^{fk}(\omega) \leq \varepsilon \right] = \exp \left[-\exp(-\theta\varepsilon - \kappa) \right], \quad (6)$$

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 (5).³ 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 $k = 1, \dots, K$ 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

$$p_{ij}^k = \tau_{ij}^k p_i^k, \quad (7)$$

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 .

3.2 Competitive Equilibrium

In a competitive equilibrium, all consumers maximize their utility, all firms maximize their profits, and all markets clear. Given Equations (2), (3), and (7), utility maximization by consumers in any country i requires

$$C_i^0 = \frac{\beta_i^0 Y_i}{p_i^0}, \text{ for all } i \in \mathcal{I}, \quad (8)$$

$$C_{ji}^k = \frac{(\tau_{ji} p_j^k)^{-\sigma}}{\sum_{j'=1}^I (\tau_{j'i} p_{j'}^k)^{1-\sigma}} \beta_i^k Y_i, \text{ for all } i, j \in \mathcal{I} \text{ and } k = 1, \dots, K, \quad (9)$$

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 located in country i that are allocated to a good k . By Equation (4)-(6), we therefore have

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

The previous expression highlights in a simple manner how relative productivity differences,

³Formally, we set $\kappa \equiv \theta \ln \Gamma(\frac{\theta-1}{\theta})$, where $\Gamma(\cdot)$ denotes the Gamma function, i.e. $\Gamma(t) = \int_0^{+\infty} v^{t-1} \exp(-v) dv$ for any $t > 0$.

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

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

which, using again Equations (5) and (6), simplifies into

$$Q_i^k = \sum_{f \in \mathcal{F}_i} A_i^{fk} \left(\frac{(p_i^k A_i^{fk})^\theta}{\sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta} \right)^{(\theta-1)/\theta}. \quad (10)$$

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

$$Q_i^0 = C_i^0, \text{ for all } i \in \mathcal{I}, \quad (11)$$

$$Q_i^k = \sum_{j \in \mathcal{I}} \tau_{ij} C_{ij}^k, \text{ for all } i \in \mathcal{I} \text{ and } k = 1, \dots, K. \quad (12)$$

Let $p_i \equiv (p_i^k)_{k \in \mathcal{K}}$ denote the vector of producer prices, $Q_i \equiv (Q_i^k)_{k \in \mathcal{K}}$ denote the vector of output levels, and $C_i \equiv (C_i^0, (C_{ji}^k)_{j \in \mathcal{I}, k \neq 0})$ 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 (8)-(12) hold.*

In the remainder of this paper we use the model outlined in this section to study the global general equilibrium consequences of climate change. In Section 6 below we compute a competitive equilibrium for an economy with contemporary agricultural productivities (as given by the GAEZ data) and also for an economy with post-climate change productivities; we then compare welfare levels across these equilibria. Before describing these counterfactual computations, however, we first (in Section 4) describe the data used in our analysis and then (in Section 5) describe how we estimate the unknown parameters in our model using these data.

4 Data

We work throughout with a sample of the 50 countries and 10 crops that span the vast majority of world crop agriculture. Our countries account for over 91 percent of world crop

Table 1: Sample Description

Countries (with shares of world crop output):

Algeria (0.41%), Angola (0.34%), Argentina (1.44%), Australia (1.03%), Bangladesh (0.78%), Brazil (4.87%), Canada (1.14%), China (22.13%), Colombia (0.63%), Congo (DRC) (0.4%), Egypt (0.95%), Ethiopia (0.53%), France (1.19%), Germany (1.07%), Ghana (0.48%), Greece (0.79%), Guatemala (0.31%), India (9.21%), Indonesia (4.05%), Iran (0.82%), Italy (0.96%), Japan (2.92%), Kazakhstan (0.36%), Malawi (0.41%), Malaysia (0.59%), Mexico (0.9%), Morocco (0.37%), Myanmar (1.63%), Nigeria (3.04%), Pakistan (1.23%), Peru (0.29%), Philippines (0.82%), Poland (0.4%), Romania (0.44%), Russia (1.99%), Serbia and Montenegro (0.31%), South Africa (0.38%), South Korea (0.99%), Spain (1.14%), Sudan (0.41%), Syria (0.38%), Tanzania (0.58%), Thailand (1.28%), Turkey (2.35%), Ukraine (0.81%), United Kingdom (0.5%), United States (9.05%), Uzbekistan (0.75%), Venezuela (0.71%), Vietnam (1.92%)

Crops (with shares of world crop output):

Rice (16.83%), Maize (11.03%), Wheat (10.38%), Cotton (6.47%), Tomato (5.70%), White Potato (5.47%), Soybean (5.44%), Sugarcane (4.52%), Citrus (3.08%) and Oilpalm (2.87%)

Notes: The 50 most important (by value of all-crop output) and 10 most important (by value of all-country output) crops, which we use throughout our analysis, where 'all crop' refers to all crops covered in the FAOSTAT database. The 50 countries sum to 90.50% percent, and the 10 crops to 71.78%, of the value of all-crop, all-country world output respectively. Value of output computed as the producer price times the quantity of output, for each crop and country; in the small number of cases where the producer price is missing but output is positive we estimate the missing producer price from the fitted values of a regression of log producer price on crop- and country-specific fixed-effects. Source: Author's estimates based on FAOSTAT data.

output value, and our crops or over 72 percent. These countries and crops, along with their shares of total (i.e. all-country, all-crop) output value, are listed in Table 1.

Our analysis draws on five main types of data: (i) estimates of agricultural productivity, at each high-resolution 'field' on Earth and for each of a series of crops, in a baseline (i.e. pre-climate change) year (which we take to be 2009); (ii) similar agricultural productivity estimates, calculated in a similar manner, but using climatological scenarios that climatologists predict will obtain over the course of this century; (iii) data on actual output, producer prices and trade flows, by crop, for each country in 2009; (iv) data on total GDP by country in 2009; and (v) data on potential determinants of trade costs. We describe the sources and construction of each of these inputs here in turn.

4.1 Agricultural Productivity Estimates at Baseline

The first data source on which we draw provides estimates of average productivity during the 'baseline', or pre-climate change period, which we take to be 2009, the most recent year for which all data is available. We require a measure of A_i^{fk} in the model above, namely the productivity in crop k for a small region 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.⁴ Because this data source is non-standard we provide a lengthy description here.

Crucially, for our purposes, 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). Importantly, GAEZ handles climate conditions particularly carefully. For a given year, data on the stream of daily weather is used to predict how well a crop will fare as each date progresses. We use GAEZ output from what the GAEZ project refers to as the ‘baseline’ period, an average of runs of the GAEZ models for the daily weather records 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. As described below, GAEZ’s treatment of climate under a climate change scenario is similar to that of a historical scenario.

The second input is a set of hundreds of model parameters, each specific to crop k , that govern how a given set of growing 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 model outlined above suggests would be inappropriate (without controlling for the 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

⁴These data are available here: http://www.gaez.iiasa.ac.at/w/ctrl?_flow=Vwr&_view=Type&idAS=0&idFS=0&fieldmain=main_py_six_qdns&idPS=1e1d6e7d7ec3368cf13a68fc523d1ed4870e8b45.

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.

A ‘field’ f in our analysis corresponds to a grid cell in the GAEZ data. The size of our fields, therefore, is governed by the size of the GAEZ data grid cells, which in turn is governed by the limitations placed by the spatial resolution of the climatic data (the growing characteristic whose underlying data is most spatially coarse). Since the climatic data is available at the 5 arc-minute level, this determines the size of the GAEZ grid cells and hence the size of a field in our analysis.⁵ At the 5 arc-minute level there are 9,000,796 grid cells on Earth; after throwing out the many grid cells that lie over bodies of water or ice shelves there are just over 2.1 million grid cells on Earth, but after focusing on our sample of the 50 most agriculturally important countries, we are left with 1,722,340 grid cells.

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.⁶

Merging the GAEZ data crops to the crops used in other data sources (from the FAO) is straightforward in most cases. An exception concerns the case of rice, where GAEZ reports two versions of rice (dryland rice and wetland rice) but only the aggregate category, rice, is available in the FAO data. We hence take the maximum yield over the two rice options, within each field, as our measure of the productivity A_i^{fk} in rice (that is, 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 at which they are most productive (and the prices of each type of rice or millet is the same).

4.2 Agricultural Productivity Estimates After Climate Change

Our analysis of the impact of climate change on global agricultural markets draws on scientists’ predictions about the impact that climate change will have on crop yields around the world. In Section 6 below we refer in our model to productivity changing—for any country i , crop k and field f —from (A_i^{fk}) at baseline to $(A_i^{fk})'$ after climate change. We obtain these predictions about crop yields under an alternative climate from the GAEZ project so

⁵Many other inputs are available at the 30 arc-second grid-cell level., which is 100 times higher than 5 arc-minute resolution. The GAEZ procedure is to solve their model at this fine level and average the result at the coarser 5 arc-minute level, and publish only the latter.

⁶The CIESIN country mapping file is available 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 ties (which occurred 408 times) randomly.

that our baseline and climate change productivity estimates are computed under exactly the same maintained agronomic assumptions.

Crucially, 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 daily weather stream occurring in year t over field f is used as the input to their model; they then average over these 30 values of A_{it}^{fk} to arrive at 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 as $\left(A_i^{fk}\right)'$ —only instead of realized past weather in year t the GAEZ project uses the predicted future daily stream of weather from year t . Estimates of future daily weather series in year t (for t from 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 central prominence in the UN's IPCC programme.

Finally we note that we use the GAEZ climate change scenario in which plant carbon dioxide fertilization is assumed to be active. That is, because the heightened atmospheric carbon dioxide predicted under climate change scenarios has potentially beneficial effects on some crops, this effect is included in the GAEZ model under climate change.

4.3 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 prevailing in the baseline year. We obtain these data from the FAOSTAT program at the FAO.⁷ The FAOSTAT program aims to provide data on worldwide production and trade, by crop and country, that 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. tonnes), of crop k in country i , corresponding to 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, $X_{ij}^k = (\tau_{ij} p_i^k) C_{ij}^k$). As is standard, we obtain this variable from the imports of reporting

⁷These data are available from <http://faostat3.fao.org/home/index.html#DOWNLOAD>.

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 10 major crops in world crop agriculture. Concording crops in the output and producer 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.⁸ In the case of two crops (oilpalm 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.

4.4 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).⁹

4.5 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, of which we focus on distance. We obtain distance measures from the ‘Gravity dataset’ produced by CEPII.¹⁰ These distance measures are computed as the (geodesic) distance between national

⁸We 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. Our estimation procedure never requires us to match the overall level of a crop’s exports (which would require the ability to combine exports of the crop in both its raw and processed forms).

⁹The World Bank GDP data are available here: <http://data.worldbank.org/indicator/NY.GDP.MKTP.CD>, and Myanmar’s GDP is available here: <https://www.cia.gov/library/publications/the-world-factbook/geos/bm.html>.

¹⁰The bilateral distance data are available here: <http://www.cepii.fr/anglaisgraph/bdd/gravity.asp>.

capitals, apart from in the case of nine countries for which a second city is added (and the minimum bilateral distance is computed) because the national capital is not a major commercial city.

5 Estimation

To simulate the model described in Section 3, we require estimates of: (i) preference parameters, (β_i^k) and σ in Equations (3) and (2); (ii) technology parameters, (A_i^{fk}) and θ in Equations (5) and (6); and (iii) trade costs, (τ_{ij}^k) in Equation (7). Section 5.1 describes how we estimate each of these parameters. Section 5.2 reports our results. Section 5.3 explores the model’s fit given estimated parameters.

5.1 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 (7), we then compute the log of trade costs as

$$\ln \tau_{ij}^k = \ln p_{ij}^k - \ln p_i^k. \quad (13)$$

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 an error:

$$\ln \tau_{ij}^k = \alpha \ln d_{ij} + \delta_{ij}^k. \quad (14)$$

We then use observed trade costs for country-pairs and crops with positive trade flows, from Equation (14), to estimate α in Equation (14) 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, $A_i^{fk} = E[A_i^{fk}(\omega)]$, 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 (10), 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 $L_i^0 \equiv \sum_{f \in \mathcal{F}_i} \left(\frac{(p_i^0 A_i^0)^\theta}{\sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta} \right)^{(\theta-1)/\theta}$ 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

$$\min_{\theta} \sum_{i,k \neq 0} \left(\ln \tilde{Q}_i^k(\theta) - \ln Q_i^k \right)^2, \quad (15)$$

where $\tilde{Q}_i^k(\theta)$ is the output level predicted by our model for a given value of θ , i.e.,¹¹

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

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 $X_{ij}^k = (\tau_{ij} p_i^k) C_{ij}^k$ 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 (9) implies

$$\ln X_{ij}^k = E_i^k + M_j^k + (1 - \sigma) \ln \tau_{ij}^k + \eta_{ij}^k, \quad (16)$$

where $E_i^k \equiv (1 - \sigma) \ln p_i^k$ can be treated as an exporter fixed effect; $M_j^k = \ln(\beta_j^k Y_j) - \ln\left(\sum_{n=1}^I (\tau_{nj} p_n^k)^{1-\sigma}\right)$ can be treated as an importer fixed effect; and η_{ij}^k is the measurement error in trade flows referred to above. We obtain our estimate of σ by estimating Equation (16) using OLS.¹² To conclude, we use trade and output data to measure the share of

¹¹This procedure requires producer price data for all crops in all countries where output in that crop is produced. In 43 instances the FAOSTAT data is missing a producer price for a country-crop where output is positive. In these cases we impute the price from the fitted values of a regression of log prices on a country and crop fixed effect.

¹²In principle we can estimate a separate elasticity of substitution across varieties within each crop k (that is a separate σ^k for all k). For simplicity, and in line with the model developed in Section 3, we focus on one

Table 2: 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)	2.571	[2.378,2.666]
σ	Elasticity of substitution in demand (across varieties of a crop)	17.864	(0.772)

Notes: Parameter estimates using method described in Section 5. Standard errors for α and σ are clustered at the country level. The reported standard error for θ is the 95% confidence interval obtained from a bootstrap procedure with 340 replications.

expenditures β_i^k across goods in different countries. For each crop $k = 1, \dots, K$, 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 $k = 1, \dots, K$ 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 $\sum_{k \neq 0} \beta_i^k$.

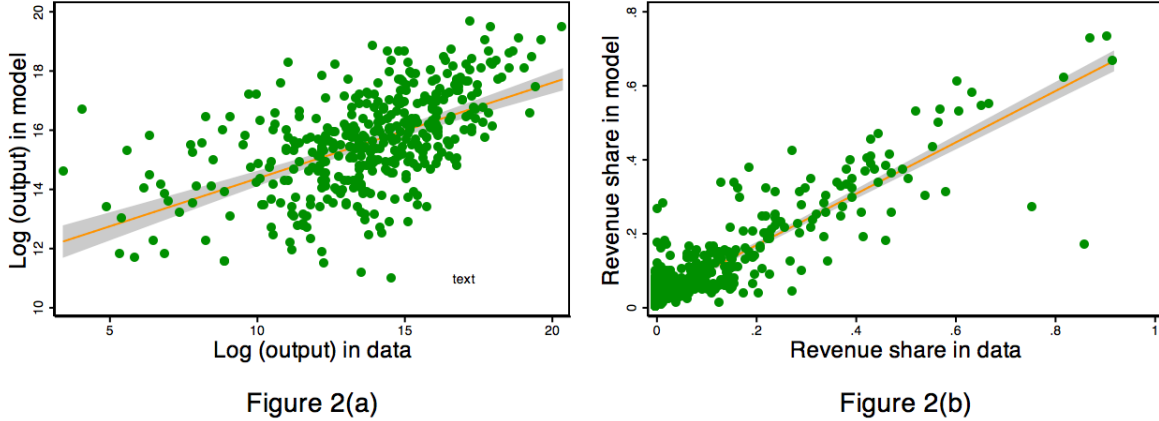
5.2 Parameter Estimates

Table 2 presents the set of parameter estimates obtained using the procedure outlined above. 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, $\theta = 2.57$, which suggests that within-field heterogeneity is substantial.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 very precisely estimated.

pooled estimate of σ that is the same across all crops.

¹³While there is no closed form for the standard error of θ we use a bootstrap procedure with 340 replications to estimate the 95% confidence interval.



Notes: Panel (a) plots log output computed in the model (y-axis) against log output in the FAO data (x-axis), across all crops and countries. Panel (b) reports the analogous comparison between revenue shares (by crop and country) in the model and the data. Best fit line and 95% confidence interval are also indicated.

5.3 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 $\theta = 2.57$. 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 ($SE = 0.037^{14}$)).

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. In Section 7, we propose an alternative estimation procedure that deals explicitly with this issue.

Figure 2 (b) illustrates that the model is considerably more successful in terms of matching relative output levels, even though this moment was not used in our estimation procedure. In Figure 2 (b) 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

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

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 (b). 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. The strong fit of our model along this dimension is reassuring since the cross-sectional variation in revenue shares is closely connected to the underlying pattern of comparative advantage.

6 Counterfactual Simulations

6.1 What are the Welfare Consequences of Climate Change?

We model climate change as a change in crop productivity from (A_i^{fk}) , as measured in the GAEZ data baseline scenario, to $(A_i^{fk})'$, as measured in the GAEZ data under the climate change scenario. All other structural parameters are held fixed at the values estimated in Section 5. Equilibrium conditions are still given by Equations (8)-(12).

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 (2) and (3), the consumer price index can be computed as

$$P_i = \prod_{k=0}^K (P_i^k)^{\beta_i^k},$$

with the component of the price index associated with crop k given by

$$P_i^k = \left(\sum_{j=1}^I (\tau_{ji} P_j^k)^{1-\sigma} \right)^{1/(1-\sigma)}.$$

Column 1 of Table 3 reports the percentage changes in real income for the median country in our sample, Malaysia, as well as the countries at the 10th and 90th percentiles, the Republic Democratic of Congo and Argentina, respectively. Appendix Table A reports the results for all 50 countries in our dataset. While the welfare consequences of climate change differ across countries—with some countries losing and some countries winning—welfare changes tend to be small in absolute values. For the median country, Malaysia, the real income loss caused by climate change is only equal to 0.19%.

An obvious reason why the welfare consequences of climate change predicted by our model are small is because the crops considered in our analysis only represent a small fraction of GDP in each country. In Column 4 of Table 3, we therefore report the exact same welfare

Table 3: Counterfactual simulation results (trade vs autarky)

	Change in real income due to climate change under scenario...		Difference between Trade (1) and Autarky (2) scenarios	Change in real income (expressed as a percentage of agricultural expenditure) due to climate change under scenario...		Difference between Trade (4) and Autarky (5) scenarios
	Trade costs, with full output reallocation	Autarky, with full output reallocation		Trade costs, with full output reallocation	Autarky, with full output reallocation	
	(1)	(2)		(4)	(5)	
World median	-0.16%	-0.06%	-0.01%	-5.2%	-2.9%	-0.2%
10th percentile	-1.89%	-1.98%	-0.17%	-35.8%	-38.3%	-5.5%
90th percentile	0.57%	0.67%	0.53%	16.4%	16.6%	4.6%

Notes: Column (1) reports the change in real income, between the model under baseline and the model under climate change, when trade costs are at the level estimated in the baseline sample (as outlined in Section 5). Column (2) reports the analogous results to column (1) for the case where trade costs are set to infinity (in both baseline and under climate change). Column (3) reports the difference between the result in columns (1) and (2), for the median, 10th and 90th percentile countries; because the identities of these countries differ across columns (1) and (2), the entries in column (3) are not equal to the difference between those in columns (1) and (2). Columns (4) through (6) report analogous results to those in columns (1) through (3) but with all numbers divided by the agricultural (for the 10 crops in our sample) expenditure share of the country in question. Source: authors' estimates as described in Sections 3, 5 and 6.

losses expressed as a share of agricultural expenditure.¹⁵ We see that, even after controlling for the size of the agricultural sector, welfare changes caused by climate change remain modest. For Malaysia, the welfare loss only goes up to 5.26%.

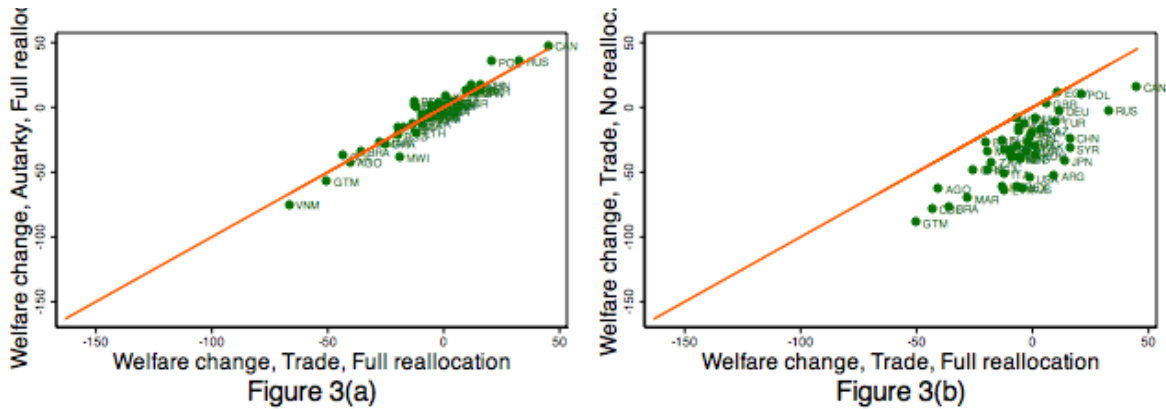
6.2 Does International Trade Matter?

At this point an interesting question is whether international trade may have contributed, through the economic channel discussed in Section 2, to alleviate the adverse consequences of climate change. Our answer is a clear no. If anything, we find that international trade aggravates the ill-effects of climate change.

Formally, we consider a second counterfactual scenario in which countries are assumed to be under autarky, i.e. trade costs are set to infinity. We recompute the equilibrium with and without climate change under this assumption. The equilibrium conditions are the same as in Section 3.2, except for the good market clearing conditions, which are now given by

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

¹⁵For the median country in our sample, the share of agricultural expenditure is 4.6%. Similarly, the world share of agricultural expenditure is 4.7%.



Notes: Panel (a) plots the real income change (% change, as a share of agricultural expenditure) due to climate change in the model under autarky with full allocation (y-axis) against the model under with full reallocation (x-axis). Panel (b) plots a similar comparison but for the model with trade and no reallocation (y-axis) against the model with trade and full reallocation (x-axis). For ease of exposition, these plots do not contain the results for Sudan or Uzbekistan (or, in panel (b), for Malawi) but these results can be found in Table A. The 45-degree line is also illustrated.

Columns 2 and 5 of Table 3 reports the welfare effects of climate change in the absence of international trade, without and with the normalization by agricultural expenditure shares, respectively. We see that for the median country, the welfare loss caused by climate change under autarky is of the same order of magnitude, albeit slightly smaller in absolute value, than the same welfare loss under trade. The median difference in the welfare loss between the two scenarios is reported in Columns 3 and 6. It is only equal to -0.01% , without the normalization by agricultural expenditure shares, and -0.2% with it. To summarize, the median country is slightly more negatively affected by climate change under trade than under autarky.

The small impact of international trade on the consequences of climate change is best seen in Figure 3 (a), which plots the welfare loss caused by climate change under autarky against the same welfare loss under trade for all countries in our dataset. Almost all observations are around the 45 degree line, suggesting that for most countries in our dataset, international trade does little to alleviate or exacerbate the consequences of climate change. Given the simple analysis of Section 2, these quantitative results suggest that comparative advantage between countries is unaffected by climate change, at least for the 10 crops in our sample.

6.3 Do Comparative-Advantage Based Reallocations Matter?

The fact that comparative advantage between countries is stable, of course, does not imply that comparative advantage within countries is stable as well. As we have illustrated in Figure 1, there is a considerable amount of within-country heterogeneity in terms of the consequences of climate change. Hence the small welfare effects of climate change uncovered in Section 6.1 may be the result of the endogenous specialization of farmers into new crops following a change in their comparative advantage (rather than merely a reflection of a small

Table 4: Counterfactual results (reallocation vs non-reallocation)

	Change in real income due to climate change under scenario...		Difference between Reallocation (1) and Non-reallocation (2) scenarios	Change in real income (expressed as a percentage of agricultural expenditure) due to climate change under scenario...		Difference between Reallocation (1) and Non-reallocation (2) scenarios
	Trade costs, with full output reallocation	Trade costs, with no output reallocation		Trade costs, with full output reallocation	Trade costs, with no output reallocation	
	(1)	(2)	(3)	(4)	(5)	(6)
World median	-0.16%	-1.58%	1.14%	-5.2%	-34.7%	30.4%
10th percentile	-1.89%	-5.07%	0.2%	-35.8%	-78.3%	10.4%
90th percentile	0.57%	-0.05%	4.0%	16.4%	-3.7%	56.6%

Notes: Column (1) reports the change in real income, between the model under baseline and the model under climate change, when trade costs are at the level estimated in the baseline sample (as outlined in Section 5) and production is allowed to reallocate between the baseline and climate change scenarios. Column (2) reports the analogous results to column (1) for the case where production is not allowed to reallocate. Column (3) reports the difference between the result in columns (1) and (2), for the median, 10th and 90th percentile countries; because the identities of these countries differ across columns (1) and (2), the entries in column (3) are not equal to the difference between those in columns (1) and (2). Columns (4) through (6) report analogous results to those in columns (1) through (3) but with all numbers divided by the agricultural (for the 10 crops in our sample) expenditure share of the country in question. Source: authors' estimates as described in Sections 3, 5 and 6.

underlying productivity shock).

To quantify this economic channel, we propose the following counterfactual scenario. We maintain the assumption that countries are allowed to trade internationally, as in Section 6.1, but we now assume that farmers cannot reallocate production across crops after climate change. Formally, we recompute the equilibrium with climate change under the assumption that the allocation of all fields to all goods in all countries is the same as in the initial equilibrium without climate change. The basic idea here is to shut down all comparative-advantage based reallocations.

Under this counterfactual scenario, total output of good k in country i is given by

$$(Q_i^k)' = \Gamma \left(\frac{\theta - 1}{\theta} \right) \exp\left(-\frac{e}{\theta}\right) \sum_{f \in \mathcal{F}_i} (A_i^{fk})' \pi_i^{fk}, \quad (17)$$

where $(A_i^{fk})'$ still denotes productivity under the climate change scenario, but $\pi_i^{fk} = \frac{(p_i^k A_i^{fk})^\theta}{\sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta}$ 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 (8), (9), (11), and (12) are unchanged.

Columns 2 and 5 of Table 4 report the welfare effects of climate change in the absence of

factor reallocation, without and with the normalization by agricultural expenditure shares, respectively. For expositional purposes, the same numbers under full reallocation, as assumed in Section 6.1, are reported in Columns 1 and 4. In contrast to the previous counterfactual scenario, we see that the welfare losses of climate change are an order of magnitude larger (in absolute value) in the absence of factor reallocation. When normalized by agricultural expenditure shares, the median difference between the welfare losses with and without factor reallocation is equal to 30.4%, as reported in Column 6.

Like in Section 6.2, the importance of comparative-advantage based reallocations is best described by plotting the welfare loss caused by climate change without reallocation against the same welfare loss with reallocation for all countries in our dataset. In Figure 3 (b), we see that almost all observations now lie below the 45 degree line. This illustrates how farmers' ability to substitute crop production in response to changes in comparative advantage—which our extremely rich micro-level dataset gives us a unique opportunity to study—may substantially mitigate the ill-effects of climate change.

The previous results are subject to an important caveat. According to the estimates presented in Table 4, agricultural markets will not be especially adversely affected by climate change, but only if farmers can and do adjust to the new climate by switching what they grow. In our model, adjustment is costless. In reality, of course, adjustment costs may be substantial, so the full-adjustment results in Column 1 of Table 4 are best interpreted as lower bounds on the costs of climate change in this setting. Equally, our results in Table 4, Column 2, are best thought of as upper bounds—at least within the confines of our exercise—on the severity of damages that can be done by climate change in agricultural markets.

7 Sensitivity Analysis

7.1 Elasticity of Substitution Between Crops

In our baseline analysis, we assume that crops enter the upper-level utility function in a Cobb-Douglas manner. The goal of this extension is to explore the sensitivity of our main results to this particular modeling choice.

To do so we generalize our analysis by considering the case of nested CES utility functions:

$$U_i = (C_i^0)^{\beta_i^0} \left(\sum_{k=1}^K \left(\frac{\beta_i^k}{1 - \beta_i^0} \right) (C_i^k)^{(\gamma-1)/\gamma} \right)^{\frac{(1-\beta_i^0)\gamma}{\gamma-1}}, \quad (18)$$

$$C_i^k = \left(\sum_{j=1}^I (C_{ji}^k)^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \text{ for all } k = 1, \dots, K, \quad (19)$$

where $\gamma \geq 1$ is the elasticity of substitution between crops and $(\beta_i^k)_{k \in \mathcal{K}}$ are exogenous pref-

erence parameters such that $\sum_{k \in \mathcal{K}} \beta_i^k = 1$. The model developed in Section 3.1 corresponds to the special case, $\gamma = 1$.

Given Equations (18) and (19), consumption levels are equal to

$$C_i^0 = \frac{\beta_i^0 Y_i}{P_i^0}, \text{ for all } i \in \mathcal{I}, \quad (20)$$

$$C_i^k = \frac{(\beta_i^k)^\gamma (P_i^k)^{-\gamma}}{\sum_{l=1}^K (\beta_i^l)^\gamma (P_i^l)^{1-\gamma}} (1 - \beta_i^0) Y_i, \text{ for all } i \in \mathcal{I} \text{ and } k = 1, \dots, K. \quad (21)$$

where $P_i^k = \left(\sum_{j=1}^I (\tau_{ji} P_j^k)^{1-\sigma} \right)^{1/(1-\sigma)}$. In the Cobb-Douglas case, $\gamma = 1$, the previous expression simplifies into $C_i^k = \beta_i^k Y_i / P_i^k$ so that β_i^k also denotes the exogenous share of expenditure on crop k in country i , as assumed in Section 3.1. The rest of our model is unchanged.

In order to estimate the new structural parameter, γ , we adopt the following strategy. Let $X_i^k \equiv P_i^k C_i^k$ denote total expenditure on crop k in country i . We first rearrange Equation (21) as

$$\ln X_i^k = D_i + (1 - \gamma) \ln P_i^k + \nu_i^k$$

where $D_i \equiv (1 - \beta_i^0) Y_i / \left(\sum_{l=1}^K (\beta_i^l)^\gamma (P_i^l)^{1-\gamma} \right)$ can be treated as a country fixed effect and $\nu_i^k \equiv \gamma \ln \beta_i^k$ reflects idiosyncratic demand shocks across crops within a country. In order to address the endogeneity between demand shocks, ν_i^k , and prices, P_i^k , we need exogenous supply shocks that are correlated with P_i^k , but uncorrelated with ν_i^k . We construct the following instrument based on the GAEZ data:

$$Z_i^k \equiv \sum_{j \in \mathcal{I}} (\tau_{ji}^k)^{1-\sigma} \ln A_i^k,$$

where $A_i^k \equiv E \left[A_i^{fk} \right]$ is the arithmetic average of productivity across fields in country i for crop k . Our exclusion restriction is that $E \left[Z_i^k \nu_i^k \right] = 0$. Our IV-estimate of γ is equal to 20.45.¹⁶

Using this new parameter estimate as well as the other estimates from Section 5, we recompute the same counterfactual scenarios as in Sections 6.1 and 6.2. Our new results, $\gamma = 20.45$, are presented in Columns 3 and 4 of Table 5. To ease comparisons, our old results, $\gamma = 1$, are presented in Columns 1 and 2. We see that, regardless of whether countries can trade or not, the welfare consequences of climate change remain very similar

¹⁶In practice, taste themselves may be endogenous to productivity. If so, one would expect our exclusion restriction to be violated with $E \left[Z_i^k \nu_i^k \right] > 0$: people living in regions in which a crop is easy to produce may develop a taste for that crop; see Atkin (Forthcoming). This would lead our IV-estimate of γ to be biased upwards. Since the main goal of this section is to explore the sensitivity of our previous results, we think that using a value of γ very different from 1 is, in any case, appealing.

Table 5: Counterfactual simulation results (robustness)

Effects of climate change on real income (expressed as percentage of agricultural expenditure) under modeling assumption and scenario...								
	Baseline		Generalized CES preferences		Unrestricted crop- and country-specific productivity adjustments to GAEZ data		Tradable outside sector	
	Trade	Autarky	Trade	Autarky	Trade	Autarky	Trade	Autarky
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
World median	-5.2%	-2.9%	-4.9%	-1.7%	-2.5%	-1.6%	-3.5%	-3.0%
10th percentile	-35.8%	-38.3%	-39.8%	-37.1%	-16.3%	-14.7%	-40.7%	-35.6%
90th percentile	16.4%	16.6%	11.1%	10.7%	3.2%	3.5%	11.8%	17.3%

Notes: Column (1) reports the change in real income, between the model under baseline and the model under climate change, when trade costs are at the level estimated in the baseline sample (as outlined in Section 5). Column (2) reports the analogous results to column (1) for the case where trade costs are set to infinity (in both baseline and under climate change). Columns (3) and (4) report analogous results to those in (1) and (2) but for a version of the model in which preferences are generalized CES (rather than Cobb-Douglas), with CES parameter $\gamma=20.45$, as described in Section 7.1. Columns (5) and (6) report analogous results to those in (1) and (2) but for a version of the model in which first we estimate crop- and country-specific unrestricted productivity adjustments to the GAEZ data, as described in Section 7.2. Columns (7) and (8) report analogous results to those in (1) and (2) but for a version of the model in which the outside sector is assumed to be freely tradable (rather than non-tradable), as described in Section 7.3. Source: authors' estimates as described in Sections 3, 5, 6 and 7.

to those predicted by our baseline model.

7.2 Productivity

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 . Except for these different measures of productivity, the model and the estimation procedures are exactly the same as in Sections 3 and 5.

In order to estimate T_i^k , we use additional data on the total amount of land allocated to crop k in country i , L_i^k . 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. Formally, we estimate θ and (T_i^k) simultaneously by solving

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

subject to:

$$L_i^k = \sum_{f \in \mathcal{F}_i} \frac{\left(p_i^k T_i^k \hat{A}_i^{fk} \right)^\theta}{\left(p_i^0 A_i^0 \right)^\theta + \sum_{l \neq 0} \left(p_i^l T_i^l \hat{A}_i^{fl} \right)^\theta}, \text{ for all } i, k \neq 0,$$

$$\tilde{Q}_i^k(\theta) = \sum_{f \in \mathcal{F}_i} T_i^k \hat{A}_i^{fk} \left(\frac{\left(p_i^k T_i^k \hat{A}_i^{fk} \right)^\theta}{\left(p_i^0 A_i^0 \right)^\theta + \sum_{l \neq 0} \left(p_i^l T_i^l \hat{A}_i^{fl} \right)^\theta} \right)^{(\theta-1)/\theta}, \text{ for all } i, k \neq 0.$$

Like in the previous extension, we use these new estimates to recompute the same counterfactual scenarios as in Sections 6.1 and 6.2. Columns 5 and 6 of Table 5 present the results. Like in our baseline analysis and in the previous extension, we find fairly small welfare losses for the median country in our sample, irrespectively of whether that country can trade internationally or not.

7.3 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 good 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 , which we normalize to one. In contrast, productivity in the outside sector, A_i^0 , is free to vary across countries.

A key difference between the present model and the one developed in Section 3 is that countries may now be net exporter or importer of agricultural goods. Under this assumption, the new good market clearing conditions become:

$$\sum_{i \in \mathcal{I}} Q_i^0 = \sum_{i \in \mathcal{I}} C_i^0,$$

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

The rest of our model is unchanged. Structural parameters can be estimated in the exact same way as in Section 5. In particular, given our choice of numeraire, A_i^0 can be computed as GDP in country i minus the total crop value divided by total acres of land allocated to

the outside sector.

Our final results for the trade and autarky counterfactual scenarios are presented in Columns 7 and 8 of Table 5. For the median country, as well as the 10th and 90th percentile countries, these new results are again very much in line with the baseline numbers presented in Section 6.

8 Concluding Remarks

A large agronomic literature has modeled the implications of climate change for a variety of crops and locations around the world. 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 find small adverse effects of climate change for the median country in the world. While international trade plays virtually no role in explaining the magnitude of these effects, our analysis suggests that reallocations caused by the evolution of comparative advantage within countries substantially mitigate the ill-effects of climate change.

References

- ATKIN, D. (Forthcoming): “Trade, Tastes, and Nutrition in India,” *American Economic Review*.
- 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.
- IPCC (2007): *Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007*. Cambridge University Press, Cambridge, United Kingdom.
- 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. (2012): “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.

Table A: Counterfactual simulation results (all countries)

	Change in real income due to climate change under scenario...			Change in real income (expressed as a percentage of agricultural expenditure) due to climate change under scenario...		
	Trade costs, with full output reallocation	Autarky, with full output reallocation	Trade costs, with no output reallocation	Trade costs, with full output reallocation	Autarky, with full output reallocation	Trade costs, with no output reallocation
	(1)	(2)	(3)	(4)	(5)	(6)
Algeria	0.0%	0.1%	-0.9%	0.0%	1.1%	-19.9%
Angola	-0.5%	-0.5%	-0.8%	-39.9%	-42.8%	-63.2%
Argentina	0.6%	0.7%	-3.1%	9.9%	12.8%	-53.4%
Australia	0.0%	0.0%	-0.6%	-3.9%	-3.7%	-63.3%
Bangladesh	-2.7%	-3.1%	-3.8%	-19.6%	-22.5%	-27.2%
Brazil	-1.1%	-1.1%	-2.5%	-35.3%	-35.0%	-78.2%
Canada	0.2%	0.2%	0.1%	45.4%	47.1%	15.0%
China	0.9%	0.9%	-1.3%	16.4%	16.5%	-24.4%
Colombia	-1.1%	-0.9%	-1.9%	-42.8%	-38.2%	-78.9%
Congo (DRC)	-1.9%	-1.8%	-4.9%	-17.0%	-16.1%	-42.7%
Egypt	0.9%	0.8%	0.9%	11.4%	10.5%	11.1%
Ethiopia	-1.4%	-2.5%	-8.1%	-11.5%	-20.0%	-64.3%
France	-0.1%	0.0%	-0.3%	-11.9%	0.8%	-32.9%
Germany	0.1%	0.1%	0.0%	12.3%	16.8%	-3.8%
Ghana	-1.7%	-2.0%	-3.3%	-25.0%	-28.8%	-48.7%
Greece	0.0%	0.0%	-0.5%	-1.5%	2.6%	-28.2%
Guatemala	-3.9%	-4.5%	-6.9%	-50.0%	-57.8%	-89.5%
India	-0.7%	-0.8%	-3.0%	-8.5%	-9.8%	-38.9%
Indonesia	-0.4%	-0.3%	-2.9%	-5.0%	-4.6%	-40.8%
Iran	0.0%	0.0%	-0.8%	-0.4%	0.0%	-24.0%
Italy	-0.1%	0.0%	-0.5%	-11.6%	-0.2%	-51.7%
Japan	0.2%	0.2%	-0.6%	14.4%	11.0%	-41.8%
Kazakhstan	0.1%	0.1%	-0.7%	3.8%	3.8%	-18.3%
South Korea	0.0%	0.2%	-1.0%	1.5%	8.1%	-37.5%
Malawi	-13.6%	-28.3%	-24.9%	-18.7%	-38.9%	-34.2%
Malaysia	-0.2%	-0.1%	-1.4%	-5.3%	-2.9%	-39.7%
Mexico	-0.1%	-0.1%	-0.9%	-6.4%	-6.0%	-62.2%
Morocco	-1.3%	-1.3%	-3.3%	-27.4%	-28.0%	-71.1%
Myanmar	-1.0%	-0.7%	-13.3%	-2.5%	-1.9%	-35.2%
Nigeria	-1.7%	-1.4%	-4.3%	-18.9%	-15.5%	-48.5%
Pakistan	0.2%	0.1%	-3.5%	1.4%	0.9%	-30.9%
Peru	-0.4%	0.2%	-2.1%	-12.4%	4.4%	-62.1%
Philippines	-0.5%	-0.3%	-2.2%	-8.9%	-6.1%	-39.5%
Poland	0.3%	0.5%	0.1%	21.2%	34.6%	9.4%
Romania	-0.2%	-0.3%	-1.0%	-6.5%	-9.3%	-30.2%
Russia	0.6%	0.6%	-0.1%	33.2%	34.8%	-3.0%
Serb. and Mont.	-1.2%	-1.8%	-3.7%	-5.2%	-7.6%	-15.6%
South Africa	-0.3%	-0.2%	-1.1%	-3.0%	-2.9%	-13.3%
Spain	-0.1%	0.0%	-0.3%	-5.1%	0.3%	-19.7%
Sudan	-1.9%	-1.8%	-4.2%	-163.3%	-158.7%	-368.8%
Syria	0.9%	0.7%	-1.7%	16.4%	12.2%	-32.4%
Tanzania	-1.1%	-1.6%	-4.0%	-8.8%	-13.1%	-32.8%
Thailand	-1.8%	-1.9%	-3.6%	-12.6%	-13.3%	-25.9%
Turkey	0.6%	0.7%	-0.7%	10.5%	12.1%	-12.3%
Ukraine	0.1%	0.1%	-0.5%	1.5%	1.2%	-9.5%
UK	0.3%	0.2%	0.1%	6.2%	3.2%	2.1%
USA	0.0%	0.0%	-0.5%	-0.8%	-3.4%	-55.4%
Uzbekistan	3.6%	1.4%	-3.6%	469.1%	181.4%	-462.1%
Venezuela	-1.9%	-2.1%	-2.8%	-6.5%	-6.8%	-9.4%
Viet-Nam	-1.7%	-1.9%	-11.3%	-66.3%	-75.8%	-449.7%

Notes: See notes to Table 3 for an explanation of columns (1)-(2), (4) and (5), and Table 4 for columns (3) and (6).