

Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World

Arnaud Costinot

Massachusetts Institute of Technology and National Bureau of Economic Research

Dave Donaldson

Stanford University and National Bureau of Economic Research

Cory Smith

Massachusetts Institute of Technology

A large agronomic literature models 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. Using an extremely rich micro-level data set that contains information about the productivity—both before and after climate change—of each of 10 crops for each of 1.7 million fields covering the surface of the earth, we find that the impact of climate change on these agricultural markets would amount to a 0.26 percent reduction in global GDP when trade and production patterns are allowed to adjust. Since the value of output in our 10 crops is equal to 1.8 percent of world GDP, this corresponds to about one-sixth of total crop value.

For helpful suggestions and comments we are grateful to Sam Kortum, Rob Townsend, Ivan Werning, and seminar audiences at the Institute for International Economic Studies conference on “Climate and the Economy,” Princeton University, Bocconi, the University of Munich, University of British Columbia, and University of California, Berkeley. Moya Chin provided excellent research assistance. Costinot and Donaldson thank the National Science Foundation (under grant SES-1227635) for research support.

Electronically published January 12, 2016

[*Journal of Political Economy*, 2016, vol. 124, no. 1]

© 2016 by The University of Chicago. All rights reserved. 0022-3808/2016/12401-0006\$10.00

I. 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, chap. 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 depends not only on their average level but also on their dispersion over space. If climate change affects all crops in all countries in a uniform manner, then there is no room for farmers to adjust what they grow or for countries to adjust what they import and export. If climate change instead has a differential effect on crop yields both within and between countries, then adjustments through production and trade patterns may significantly dampen the adverse consequences of climate change. For instance, a country may stop producing a crop whose yields have fallen and import it in exchange for another crop whose yields have remained constant at home. 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 affect 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 data set on agricultural productivity: the Food and Agriculture Organization's Global Agro-Ecological Zones (GAEZ) data set. This data set 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 1.7 million high-resolution grid cells covering the surface of the earth. The GAEZ data set is available both under contemporary growing conditions and under a climate change scenario used by the UN's Intergovernmental Panel on Climate Change (IPCC).¹ By comparing productivity for a given crop under the two scenarios at each of our 1.7 million grid cells, we can therefore directly observe the evolution of comparative advantage across space, as predicted by climatologists and agronomists.

¹ The GAEZ data set reports post-climate change yield predictions for a total of 11 distinct climate change scenarios. The results in our baseline analysis, which we describe in this introduction, are based on estimates from the Hadley CM3 AIFI model. Section VII.A presents results for all other scenarios. Data are provided as supplementary material online.

A sample of the GAEZ predictions can be seen in figure 1. Here we plot, for each grid cell around the world, 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.

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, in turn, which productivity changes are relevant and which ones are not; (ii) how shocks to the

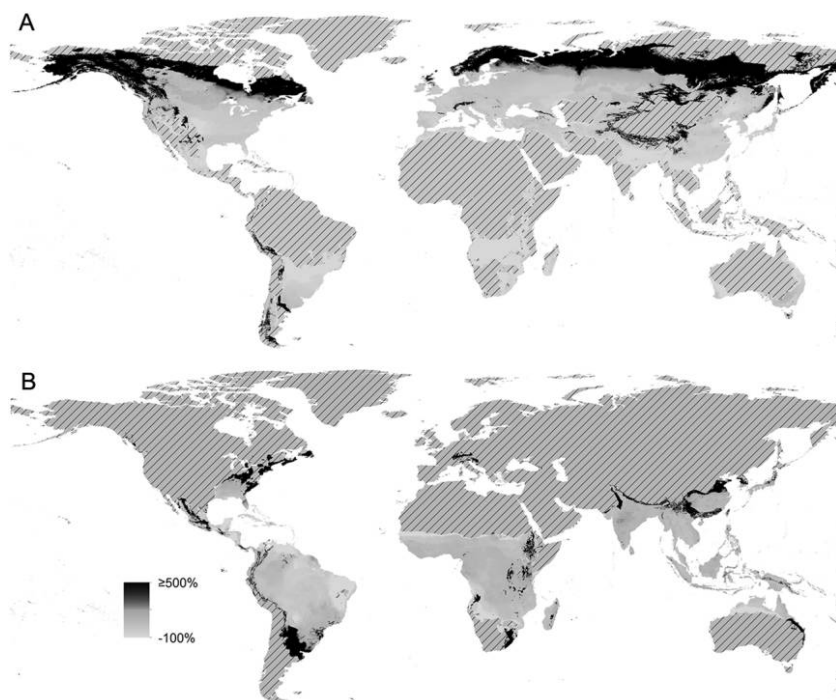


FIG. 1.—Predicted yield changes. Percentage changes in yield due to climate change in the GAEZ model. Areas with diagonal stripes indicate regions for which predicted yields are zero both before and after climate change: *A*, wheat; *B*, rice. Color version available as an online enhancement.

supply of crops affect prices around the world; and (iii) how changes in productivity and prices map into consumption and 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 1.7 million grid cells in the GAEZ data. In this model, comparative advantage, that is, relative productivity differences across crops and fields, determines the pattern of specialization within and between countries, whereas trade costs determine the level of integration of local agricultural markets.

Besides the highly detailed GAEZ data, the quantitative predictions of our model depend on three key elasticities: (i) the elasticity of substitution between different varieties of the same crop; (ii) the elasticity of substitution between different crops; and (iii) the extent of within-field heterogeneity in productivity, which is unobserved in the GAEZ data. These three parameters can be separately estimated using trade, output, and price data in a transparent manner. At the estimated parameter values, we find that the within-sample fit of our model for output levels, land use, and trade flows is good, including for moments that were 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 no-climate change scenario and explore three counterfactual scenarios. In our first scenario, we study the consequences of climate change—that is, 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 very heterogeneous effects across countries, with some countries like Malawi experiencing dramatic welfare losses. Overall, climate change amounts to a 0.26 percent decrease in world GDP. Since the value of output in our 10 crops is equal to 1.8 percent of world GDP, this corresponds to about one-sixth of total crop value.

As mentioned above, a potential source of adjustment to climate change is the ability of farmers to produce different crops. To shed light on this mechanism, we consider a second counterfactual scenario in which countries can trade, but farmers cannot reallocate production within each field. Under this scenario, we find that the adverse welfare consequences of climate change are significantly larger than in the previous scenario. For the world as a whole, the loss would be three times as large: 0.78 percent of GDP. This illustrates how farmers’ ability to substitute crop production in response to changes in their comparative advantage—which our micro-level data set gives us a unique opportunity to study—may substantially mitigate the ill effects of climate change.

Another potential source of adjustment is the ability of countries to change what they trade with the rest of the world. To explore the quantitative importance of this economic channel, we consider a final counterfactual scenario in which farmers can reallocate production, but export patterns—that is, shares of crop output exported to the rest of the world—are held fixed before and after climate change. In contrast to the previous counterfactual scenario, the welfare consequences of climate change in this case, a 0.27 percent loss in world GDP, remain very similar to those obtained under full adjustment, suggesting that international trade may play only a minor role in alleviating the consequences of climate change.

There is, of course, a great deal of uncertainty about future climate change and how it will affect crop yields at various locations. In the final part of our paper, we explore the sensitivity of our counterfactual results to different assumptions about future climatic conditions as well as the contemporary growing conditions used in the GAEZ data. We also discuss the robustness of our results to alternative assumptions about technology, preferences, and trade costs. Not surprisingly, the large uncertainty about future crop yields leads to large uncertainty over the welfare consequences of climate change. Interestingly, however, the relative importance of production adjustments relative to trade adjustments remains of similar magnitude across all our robustness checks.

The literature on international trade and climate change is large and varied, though mostly based on computational general equilibrium 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 et al. 2013; Shapiro 2013). 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, that is, 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 et al. 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; 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 microconsequences 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 these rich microdata into a general equilibrium model 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 changes in production and trade patterns.

Finally, our analysis is also related to our earlier work (Costinot and Donaldson 2011, 2012), which also uses the GAEZ data to quantify the gains from economic integration in US agricultural markets from 1880 to 2000 as well as to test the predictions of the Ricardian model. Theoretically, the present paper departs from our earlier work by introducing productivity heterogeneity within each field, which simplifies the problem of assigning fields to crops in a competitive equilibrium. Empirically, of course, neither of our previous papers uses the post-climate change GAEZ predictions, which are at the core of the present analysis.

The rest of this paper is organized as follows. Section II illustrates how the evolution of comparative advantage may mitigate the consequences of climate change in agricultural markets through a simple example. Section III develops our theoretical framework. Section IV presents the data that feed into our analysis. Section V describes our estimation procedure, our parameter estimates, and measures of goodness of fit of the model. Section VI then presents the results of our counterfactual simulations. Section VII explores the sensitivity of our results, and Section VIII offers some concluding remarks.

II. A Simple Example

Consider an agrarian economy comprising two symmetric islands, North and South. Each island consists of two fields, East and West, that can be used to produce two crops, wheat and rice. All fields have the same size and the same wheat yields and rice yields per acre. All consumers have the same preferences. They spend half of their income on wheat and half of their income on rice. Hence, there exists a competitive equilibrium, which we assume prevails, such that the relative price of the two crops is one, western fields produce wheat, whereas eastern fields produce rice on both islands.²

Because of climate change, agronomists expect land productivity to go down in the two islands, but differentially so across crops and fields. In the southern island, rice yields will be unchanged, whereas wheat yields will decrease by 50 percent in the eastern field and go down to zero in the western field. In the northern island, the situation is the exact opposite. While wheat yields will be unchanged, rice yields will decrease by 50 percent in the western field and go down to zero in the eastern field.

² Of course, if crop yields are exactly the same across all fields, there exist other competitive equilibria. None of this discussion hinges on the multiplicity of equilibria. Formally, one should think of the equilibrium that we focus on as capturing a situation in which eastern fields have a comparative advantage in rice, albeit an arbitrarily small one.

The basic questions that we are interested in are as follows: How concerned should inhabitants of the two islands be about these dire forecasts? By how much will their standards of living go down after climate change? To what extent can the two islands adapt to climate change by producing or trading different crops?

In the present context, answers to these three questions are straightforward. After climate change, the northern island has a comparative advantage in wheat, whereas the southern island has a comparative advantage in rice. Thus in the new competitive equilibrium, all northern fields will produce wheat—which will then be exported to South—whereas all southern fields will produce rice—which will then be exported to North. As a result, world production and world consumption will be exactly the same as before climate change. Here, in spite of causing large negative productivity shocks that range from -50 percent to -100 percent, climate change will have no effects on living standards.

The two islands are unaffected by climate change for two reasons. First, fields can change what they produce. Second, islands can change what they export and import. In a counterfactual scenario in which fields can change what they produce but islands are under autarky, both islands would be forced to produce both crops. Therefore, negative shocks to rice yields in the North and wheat yields in the South would necessarily lower living standards. The same observation applies to a counterfactual scenario in which islands are free to trade, but fields cannot change what they produce. The extent of the welfare losses in both counterfactual scenarios will be our measures of the importance of international trade and production reallocation for alleviating the adverse consequences of climate change.

Needless to say, the previous example is very special. In particular, it assumes that climate change creates comparative advantage across fields when there was none. Thus, by specializing and trading more, islands are able to dampen (fully) the adverse consequences of climate change. If climate change were to weaken comparative advantage, then islands may adapt instead by diversifying production and trading less. In both cases, however, the key observation is that, to the extent that climate change affects comparative advantage across regions of the world, the welfare consequences of climate change may crucially depend on the ability of a country to change its production and trade patterns.

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 and trade costs between countries, more than two crops, and, most importantly, more than two types of land in the form of more than a million fields.

III. Theory

A. Basic Environment

We consider a world economy comprising multiple countries, indexed by $i \in \mathcal{I} \equiv \{1, \dots, I\}$. In each country there are two factors of production, labor and land, that can be used to produce multiple crops, indexed by $k \in \mathcal{K} \equiv \{1, \dots, K\}$, and an outside good, which we think of as a composite of manufactured goods and services. Labor is homogeneous, perfectly mobile within a country, and immobile across countries. The term N_i denotes the total endowment of labor and w_i denotes the wage in country i . Land comes in the form of heterogeneous fields, indexed by $f \in \mathcal{F}_i \equiv \{1, \dots, F_i\}$, each comprising a continuum of heterogeneous parcels, indexed by $\omega \in [0, 1]$. All fields correspond to a 5-arc-minute grid cell in our data set, and there are 1.7 million such grid cells around the world. Owing to the curvature of the earth, grid cells at different latitudes cover different areas. We let s_i^f denote the area in hectares of field f in country i .

Preferences.—In each country i there is a representative agent who derives utility from consuming the outside good, C_i^0 , and a composite of all crops, C_i :

$$U_i = C_i^0 + \beta_i \ln C_i. \quad (1)$$

Since the upper-level utility function in equation (1) is quasi-linear, there are no income effects. The total demand for crops depends only on a country-specific demand shifter, $\beta_i \geq 0$. Since the crops on which we focus account for a small fraction of consumers' expenditure around the world, we view the absence of income effects as a minor limitation of our analysis.³

Aggregate crop consumption, C_i , depends on the consumption of each crop, C_i^k , which itself depends on the consumption of varieties from different origins, C_{ji}^k :

$$C_i = \left[\sum_{k \in \mathcal{K}} (\beta_i^k)^{1/\kappa} (C_i^k)^{(\kappa-1)/\kappa} \right]^{\kappa/(\kappa-1)}, \quad (2)$$

³ The assumption of log preferences, which implies a unit price elasticity, is more restrictive. It rules out scenarios in which consumers may change their total expenditure on crops in response to climate change. If one thinks, for example, of the outside good as "health," one may expect situations in which consumers may want to spend a higher fraction of their income on the outside good after climate change, perhaps because climate change destroys rare species that are critical inputs in the pharmaceutical industry. Under such circumstances, the welfare costs of productivity changes in agriculture would be smaller than those computed in our paper.

$$C_i^k = \left[\sum_{j \in \mathcal{I}} (\beta_{ji}^k)^{1/\sigma} (C_{ji}^k)^{(\sigma-1)/\sigma} \right]^{\sigma/(\sigma-1)}, \quad (3)$$

where $\kappa > 0$ denotes the elasticity of substitution between different crops (e.g., wheat vs. corn) and $\sigma > 0$ denotes the elasticity of substitution between different varieties of a given crop (e.g., French vs. US wheat). The last preference parameters, $\beta_i^k \geq 0$ and $\beta_{ji}^k \geq 0$, are crop- and country-specific and crop origin- and destination-specific demand shocks, respectively.

Although crops are fairly homogeneous goods, the assumption that crops enter utility through a general constant elasticity of substitution (CES) aggregator rather than as perfect substitutes is convenient from a computational standpoint. It avoids the need to determine, for each crop, whether a country is a net exporter and who its trading partners are. Here, all countries export each crop that they produce to all other countries (as long as $\beta_{ji}^k > 0$). Of course, the degree of substitutability between crops from different countries is, ultimately, an empirical question, which we tackle in Section V.A.

Technology.—The outside good is produced under constant returns to scale using labor only. The term $A_i^0 > 0$ denotes labor productivity in country i 's outside sector. In the agricultural sector, we adopt a parsimonious representation of technology designed to keep the model transparent yet flexible enough to harness the vast amount of micro-level data collected by agronomists. We assume that labor and parcels of land are perfect complements in the production of each crop. By combining $L_i^{fk}(\omega)$ hectares of parcel ω with $N_i^{fk}(\omega)$ workers, a representative firm can produce

$$Q_i^{fk}(\omega) = A_i^{fk}(\omega) \min\{L_i^{fk}(\omega), N_i^{fk}(\omega)/\nu_i^f(\omega)\}, \quad (4)$$

where $A_i^{fk}(\omega) \geq 0$ denotes the total factor productivity (TFP) of parcel ω in field f if allocated to crop k in country i and $\nu_i^f(\omega) > 0$ measures the labor intensity of the associated production process. In the spirit of Eaton and Kortum (2002), we assume that TFP and labor intensity are independently drawn for each (i, f, ω) from a Fréchet distribution:

$$\begin{aligned} & \Pr\{A_i^{f1}(\omega) \leq a^1, \dots, A_i^{fK}(\omega) \leq a^K, \nu_i^f(\omega) \leq \nu\} \\ &= \exp \left\{ -\gamma \left[\sum_{k \in \mathcal{K}} (a^k/A_i^{fk})^{-\theta} + (\nu/\nu_i)^{-\theta} \right] \right\}, \end{aligned}$$

where $\theta > 1$ measures the extent of technological heterogeneity within each field and the constant γ is set such that $A_i^{fk} = E[A_i^{fk}(\omega)]$ and $\nu_i = E[\nu_i^{fk}(\omega)]$.⁴ The term $A_i^{fk} \geq 0$ measures the comparative and absolute ad-

⁴ Formally, we set $\gamma \equiv \Gamma((\theta-1)/\theta)^{-\theta}$, where $\Gamma(\cdot)$ denotes the gamma function; i.e., $\Gamma(t) = \int_0^{+\infty} u^{t-1} \exp(-u) du$ for any $t > 0$. We relax the assumption that the extent of technological heterogeneity θ is constant across countries in Sec. VII.

vantage of a field in producing particular goods. The GAEZ project data give us direct information about $A_i^k \geq 0$ for all crops and all locations as a function of global temperatures, which will be the core inputs in our quantitative exercise. Finally, since we do not have access to disaggregated data on labor intensity, we require average labor intensity $\nu_i > 0$ to be identical across crops and fields, though agriculture is allowed to be more labor intensive in some countries than in others.

Market structure and trade costs.—All markets are perfectly competitive. The outside good is freely traded. In the rest of our analysis we use it as our numeraire. In contrast, international trade in crops $k \in \mathcal{K}$ may be subject to iceberg trade costs. In order to sell one unit of a good in country j , firms from country i must ship τ_{ij}^k units. Nonarbitrage therefore requires the price of a crop k produced in country i and sold in country j to be equal to

$$p_{ij}^k = \tau_{ij}^k p_i^k, \quad (5)$$

where p_i^k denotes the local price of the domestic variety of crop k in country i .

B. Competitive Equilibrium

In a competitive equilibrium, all consumers maximize their utility, all firms maximize their profits, and all markets clear.

Utility maximization.—Given equations (1), (2), (3), and (5), utility maximization by the representative agent in each country requires that

$$C_{ji}^k = \beta_i \frac{\beta_i^k (P_i^k)^{1-\kappa}}{\sum_{l \in \mathcal{K}} \beta_i^l (P_i^l)^{1-\kappa}} \frac{\beta_{ji}^k (\tau_{ji}^k p_j^k)^{-\sigma}}{\sum_{n \in \mathcal{I}} \beta_{ni}^k (\tau_{ni}^k p_n^k)^{1-\sigma}} \quad \text{for all } i, j \in \mathcal{I}, k \in \mathcal{K}, \quad (6)$$

where

$$P_i^k \equiv \left[\sum_{n \in \mathcal{I}} \beta_{ni}^k (\tau_{ni}^k p_n^k)^{1-\sigma} \right]^{1/(1-\sigma)}$$

denotes the CES price index associated with crop k in country i .

Profit maximization.—In the outside sector, profit maximization requires that $w_i = A_i^0$ whenever the outside good is produced. Throughout this paper, we assume that labor endowments, N_i , are large enough for the outside good to be produced in all countries. Thus we can use A_i^0 in place of the wage w_i and treat it as an exogenous parameter from now on.

In the agricultural sector, profit maximization requires that all parcels of land are (i) allocated to the crop that maximizes the value of their marginal product if such value is greater than the wage bill associated

with operating that parcel or (ii) left unused if the maximum value of their marginal product is less than the wage bill. Thus given equation (4), the land allocation can be solved as a simple discrete choice problem. Let

$$\pi_i^{fk} \equiv \Pr\{p_i^k A_i^{fk}(\omega) = \max\{A_i^0 v_i^f(\omega), p_i^1 A_i^{f1}(\omega), \dots, p_i^K A_i^{fK}(\omega)\}\}$$

denote the probability that a parcel ω of a field f located in country i is allocated to crop k . Since there is a continuum of parcels within each field, π_i^{fk} also corresponds to the share of parcels allocated to that crop.

Given our distributional assumptions, standard algebra implies

$$\pi_i^{fk} = \frac{(p_i^k A_i^{fk})^\theta}{\alpha_i^\theta + \sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta} \quad \text{for all } f \in \mathcal{F}_i, i \in \mathcal{I}, \text{ and } k \in \mathcal{K}, \quad (7)$$

where $\alpha_i \equiv A_i^0 v_i$ parameterizes cross-country differences in labor costs, because of differences in either wages or labor intensity. The higher α_i is, the more costly it is to hire workers to produce crops, and the smaller the share of a field f allocated to any given crop k . Likewise, the higher the average value of the marginal product of land, $p_i^k A_i^{fk}$, the higher the share of field f allocated to crop k . In our model, the extent of technological heterogeneity, θ , determines the elasticity of the relative supply of land to various crops. When θ is higher, parcels are more homogeneous within a field, which makes the supply of land more sensitive to changes in prices, p_i^k , or productivity, A_i^{fk} .

Let $Q_i^k \equiv \sum_{f \in \mathcal{F}_i} \int_0^1 Q_i^{fk}(\omega) d\omega$ denote that the total output of crop k in country i . By equation (4) and the law of iterated expectations, we must have

$$\begin{aligned} Q_i^k &= \sum_{f \in \mathcal{F}_i} s_i^f \pi_i^{fk} E[A_i^{fk}(\omega) | p_i^k A_i^{fk}(\omega)] \\ &= \max\{A_i^0 v_i^f(\omega), p_i^1 A_i^{f1}(\omega), \dots, p_i^K A_i^{fK}(\omega)\}. \end{aligned}$$

Given our distributional assumptions, one can also check that

$$\begin{aligned} E[A_i^{fk}(\omega) | p_i^k A_i^{fk}(\omega)] &= \max\{A_i^0 v_i^f(\omega), p_i^1 A_i^{f1}(\omega), \dots, p_i^K A_i^{fK}(\omega)\} \\ &= A_i^{fk} \times (\pi_i^{fk})^{-1/\theta}. \end{aligned}$$

Note that because of the endogenous selection of fields into crops, the average productivity conditional on a crop being produced is strictly greater than the unconditional average, $A_i^{fk} \times (\pi_i^{fk})^{-1/\theta} > A_i^{fk}$.

Combining the two previous expressions with equation (7), we obtain the following expression for the supply of crop k in country i :

$$Q_i^k = \sum_{f \in \mathcal{F}_i} s_i^f A_i^{fk} \left[\frac{(p_i^k A_i^{fk})^\theta}{(\alpha_i)^\theta + \sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta} \right]^{(\theta-1)/\theta} \quad \text{for all } i \in \mathcal{I} \text{ and } k \in \mathcal{K}. \quad (8)$$

Market clearing.—Since trade in crops is subject to iceberg trade costs, market clearing for all varieties of all crops requires

$$Q_i^k = \sum_{j \in \mathcal{I}} \tau_{ij}^k C_{ij}^k \quad \text{for all } i \in \mathcal{I} \text{ and } k \in \mathcal{K}. \quad (9)$$

As discussed above, parcels of land may remain idle if the value of their marginal product is below the labor cost required to produce on these parcels. Thus, by construction, land demand is weakly less than land supply at all locations. Finally, under the assumption that the outside good is produced in all countries, the amount of labor demanded by the outside sector adjusts to guarantee labor market clearing at the wage equal to A_i^0 .

In the rest of this paper we formally define a competitive equilibrium as follows.

DEFINITION 1. A competitive equilibrium is a vector of consumption, (C_{ji}^k) , output, (Q_i^k) , and prices, (p_i^k) , such that equations (6), (8), and (9) hold.

In the remainder of this paper we will use the model outlined in this section to study the global consequences of climate change. In Section VI below, we will compute competitive equilibria for economies with contemporary agricultural productivities, compute competitive equilibria for counterfactual economies with post-climate change productivities, and then compare welfare levels across equilibria. Before presenting these counterfactual simulations, however, we need to describe the data used in our analysis as well as how we estimate the unknown structural parameters of our model using these data. This is the object of Sections IV and V, respectively.

IV. 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 89.7 percent of world crop output value and our crops for 71.1 percent. These countries and crops, along with their shares of total output across all countries and all crops, are listed in table 1.

Our analysis draws on three main types of data: (i) pre-climate change estimates of agricultural productivity, at each high-resolution field on earth for each of our 10 crops; (ii) similar post-climate change estimates of agricultural productivity; and (iii) data on actual output, land use, prices, and trade flows, by crop, for each country in 2009.

A. Pre-Climate Change Estimates of Agricultural Productivity

The first data source on which we draw provides estimates of average productivity during the pre-climate change period. We require a mea-

TABLE 1
SAMPLE DESCRIPTION

Country	Share of World Output	Country	Share of World Output	Crop	Share of World Output
Algeria	.4%	Korea, South	1.0%	Banana	3.4%
Argentina	1.2%	Malawi	.5%	Soybean	5.6%
Australia	.9%	Malaysia	.8%	Cotton	3.2%
Bangladesh	.8%	Mexico	.9%	Sugarcane	4.7%
Brazil	4.8%	Morocco	.4%	Maize	11.4%
Burma	1.7%	Netherlands	.4%	Tomato	5.9%
Cameroon	.4%	Nigeria	3.3%	Oil palm	3.2%
Canada	1.3%	Pakistan	1.1%	Wheat	10.7%
China	21.2%	Philippines	.8%	Rice	17.3%
Colombia	.7%	Poland	.4%	White potato	5.6%
Cote d'Ivoire	.4%	Romania	.5%		
D.R. Congo	.5%	Russia	2.0%		
Ecuador	.4%	South Africa	.4%		
Egypt	.9%	Spain	.8%		
Ethiopia	.5%	Sudan	.4%		
France	1.3%	Tanzania	.6%		
Germany	1.2%	Thailand	1.3%		
Ghana	.7%	Turkey	1.8%		
Greece	.4%	Uganda	.7%		
India	9.0%	Ukraine	.9%		
Indonesia	4.1%	United Kingdom	.6%		
Iran	.8%	United States	9.3%		
Italy	.8%	Uzbekistan	.6%		
Japan	3.0%	Venezuela	.7%		
Kazakhstan	.4%	Viet Nam	2.0%		

NOTE.—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 GAEZ data set. The 50 countries sum to 89.7 percent, and the 10 crops to 71.1 percent, of the value of all-crop, all-country world output, respectively. Value of output is computed as the local price p_i^k times the quantity of output Q_i^k , for each crop and country.

sure 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 GAEZ project, which is organized under the auspices of the Food and Agriculture Organization (FAO) and the International Institute for Applied Systems Analysis. Because this data source is nonstandard, 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 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 such as 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 are 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.

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 exposure, 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) with "rain-fed" water supply. We explore the sensitivity of our results to this assumption in Section VII.C.

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 coarseness of climatic data. Since the climatic data are available at the 5-arc-minute level, this determines the size of the GAEZ grid cells and hence the area of a field in our analysis.⁵ At the 5-arc-minute level there are around 9 million grid cells on earth. After throwing out the many grid cells that lie over bodies of water

⁵ Many other inputs are available for 30-arc-second grid cells, which correspond to 1/100 of a 5-arc-minute grid cell. 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.

or ice shelves, there are just over 2.2 million grid cells on earth. Focusing on our sample of the 50 most agriculturally important countries, we are left with 1,683,193 grid cells. The area of these fields ranges from 8,586.3 hectares (at the most equatorial fields) to 1,040.2 hectares (at the most polar fields).

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 boundaries (for boundaries in 2008) file available from the World Borders Dataset (Thematic Mapping).⁶ At each of the 1,683,193 grid cells, the GAEZ data set reports predicted yields in tons per hectare for each of our 10 crops.

Merging the GAEZ data crops to the crops used in other data sources 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. In this case we use the maximum yield over the two rice options, within each field, as our measure of the productivity A_i^{fk} in rice (i.e., our measure of $A_i^{fk=\text{rice}} = \max\{A_i^{fk=\text{dryland rice}}, A_i^{fk=\text{wetland rice}}\}$).

B. Post-Climate Change Estimates of Agricultural Productivity

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 VI below we let $(A_i^{fk})'$ denote post-climate change productivity for any country i , crop k , and field f in our counterfactual scenarios. In order to obtain estimates of $(A_i^{fk})'$, we rely again on the agronomic predictions from the GAEZ project.

There are two changes that the GAEZ project implements when computing post-climate change estimates $(A_i^{fk})'$ rather than pre-climate change estimates A_i^{fk} . The first and most important one concerns the weather that prevails at field f in country i in each scenario. Like pre-climate change estimates, post-climate change estimates are based on an average of runs of the GAEZ models over a 30-year period, only instead of realized past weather, the GAEZ project uses the predicted future daily stream of weather from 2071 to 2100. Estimates of future daily weather series are themselves obtained from an average of runs of a general circulation model (GCM) under a particular Special Report on Emission Scenarios (SRES) from the IPCC program.

The IPCC program comprises a total of 23 GCM and 40 SRES. Each SRES corresponds to a different "narrative storyline" about the potential

⁶ Most grid cells have a unique matching to countries. For grid cells that overlap with country borders, we assign grid cells to countries on the basis of the country that occupies the largest share of each grid cell.

evolution of various demographic, social, economic, technological, and environment variables. SRES are broken down into four main “families”—A1, A2, B1, and B2—which vary primarily in terms of their predicted rates of population and economic growth. Each SRES then maps the previous variables into predictions about future greenhouse gas and sulfur emissions around the world, with large variations between scenarios. A GCM developed by independent teams of climatologists, in turn, maps future gas emissions into future climatic conditions, which can finally be fed into the GAEZ project’s agronomic model.

Out of all possible GCM-SRES combinations, GAEZ post-climate change productivity estimates are available for 11 GCM-SRES pairs. In our baseline analysis, we use those obtained from the Hadley CM3 A1FI model. The A1FI scenario describes a future world of very rapid economic growth, global population that peaks in midcentury and declines thereafter, and the rapid introduction of new and more efficient fossil-intensive technologies. Section VII.A discusses the sensitivity of our results to the use of alternative GCM-SRES combinations available in the GAEZ data set. As we will see, results based on the Hadley CM3 A1FI model provide an upper bound on the welfare consequences of climate change in all our counterfactual scenarios.

The second change that the GAEZ project implements when computing $(A_i^k)'$ rather than A_i^k relates to atmospheric carbon dioxide levels. GAEZ post-climate change estimates are available under two types of modeling assumptions: with and without active plant carbon dioxide fertilization. In the former case, changes in carbon emissions due to climate change are allowed to have a direct effect on crop yields. In the latter case, they are not. We use estimates allowing for plant carbon dioxide fertilization throughout our analysis.

C. *Agricultural Output, Land Use, Price, and Trade Flow Data*

An essential aspect of our analysis is the ability to estimate all of the unknown parameters in our model in a theoretically consistent manner. This estimation procedure—described below—requires data on actual output, land use, prices, and trade flows prevailing in a baseline year. In all cases we use data from 2009, the most recent year for which data from all sources were available.

We obtain data on agricultural output, land use, and prices from the FAOSTAT program at the FAO. The first variable, Q_i^k , measures output of crop k in country i , expressed in tons.⁷ The second variable, L_i , mea-

⁷ FAOSTAT reports output in tons of “fresh weight.” In contrast, the GAEZ project reports estimates of “dry-weight” matter per hectare, excluding the water content of a crop. We use the dry-weight conversion table provided by the GAEZ project to convert the original GAEZ yields to fresh-weight yields.

sures the total amount of land in country i allocated to the 10 crops covered by our analysis, expressed in hectares. And the third variable, p_i^k , measures the producer price of crop k in country i . We treat this variable as the counterpart of the local price of crop k in country i , for both consumers and producers, in our model. Whenever it is missing, we impute the price p_i^k from the fitted values of a regression of log prices on a country and crop fixed effect.

For the purposes of measuring trade flows, X_{ij}^k , we use UN Comtrade data because of their denser coverage relative to FAOSTAT trade data. Trade flow X_{ij}^k measures the total value of exports of crop k from country i to country j , expressed in US dollars. As is standard, we obtain the value of bilateral trade flows from the imports of reporting countries, that is, the countries that collected the data underlying the trade flow in question, in contrast to the partner country in any trade flow. Whenever imports are reported missing or zero but exports are not, we use the value reported by the exporter instead.

To concord FAOSTAT crops to those in the Comtrade data we use a concordance table available from FAOSTAT. The only exception is oil palm, for which the raw version of the crop, the oil palm fruit, is never traded because the crop is typically immediately processed into palm oil and oil palm kernels. We therefore compute the value of oil palm exports as the sum of exports in palm oil and oil palm kernels.

V. Estimation

To simulate the model described in Section III, we require estimates of demand- and supply-side parameters. Sections V.A and V.B describe how we obtain these. Section V.C then explores the model's fit given estimated parameters.

A. Demand

We proceed in three steps, each corresponding to a different level of our nested demand system. First, we use data on producer prices, p_i^k , and bilateral trade flows at the crop level, X_{ij}^k , to estimate the elasticity of substitution σ between different varieties of a given crop as well as a composite of the lower-level demand shifters, β_{ij}^k , and trade costs, τ_{ij}^k . Second, we use the previous estimates to construct price indices at the crop level, P_j^k , and combine them with data on crop expenditures, $X_j^k \equiv \sum_{i \in \mathcal{I}} X_{ij}^k$, to estimate the elasticity of substitution κ between crops as well as the mid-level demand shifters, β_j^k . Finally, we use data on total crop expenditures, $X_j \equiv \sum_{k \in \mathcal{K}} X_j^k$, to estimate the upper-level demand shifters, β_j .

Step 1: By definition, the value of exports of crop k from country i to country j is given by $X_{ij}^k = (\tau_{ij}^k p_i^k) C_{ij}^k$. By equation (6), we therefore have

$$X_{ij}^k = \beta_j \frac{\beta_j^k (P_j^k)^{1-\kappa}}{\sum_{l \in \mathcal{K}} \beta_j^l (P_j^l)^{1-\kappa}} \frac{\beta_{ij}^k (\tau_{ij}^k p_i^k)^{1-\sigma}}{\sum_{n \in \mathcal{I}} \beta_{nj}^k (\tau_{nj}^k p_n^k)^{1-\sigma}} \quad \text{for all } i, j \in \mathcal{I} \text{ and } k \in \mathcal{K}. \quad (10)$$

When estimating the lower level of our demand system, we consider the cases of zero and nonzero trade flows separately. If $X_{ij}^k = 0$, we simply set $\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma} = 0$. If $X_{ij}^k > 0$, we take logs and rearrange equation (10) as

$$\ln(X_{ij}^k / X_j^k) = M_j^k + (1 - \sigma) \ln p_i^k + \varepsilon_{ij}^k, \quad (11)$$

where the first term on the right-hand side,

$$M_j^k \equiv -\ln \left[\sum_{n \in \mathcal{I}: X_{nj}^k > 0} \beta_{nj}^k (\tau_{nj}^k p_n^k)^{1-\sigma} \right],$$

can be treated as an importer and crop fixed effect and the final term, $\varepsilon_{ij}^k \equiv \ln[\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma}]$, reflects idiosyncratic demand shocks across varieties of different crops as well as trade costs. Without loss of generality, we normalize these shocks such that

$$\sum_{i \in \mathcal{I}: X_{ij}^k > 0} \varepsilon_{ij}^k = 0. \quad (12)$$

Equilibrium crop prices, of course, depend on demand shocks, ε_{ij}^k . To address the endogeneity of crop prices, p_i^k , in equation (11), we need exogenous supply shocks that are correlated with p_i^k but uncorrelated with ε_{ij}^k . We construct the following instrument based on the GAEZ data,

$$Z_i^k \equiv \ln \left(\frac{1}{F_{i \in \mathcal{F}_i}} \sum A_i^k \right),$$

which corresponds to the log of the arithmetic average of crop k 's yields across all fields in country i . Our exclusion restriction is that $E[Z_i^k \varepsilon_{ij}^k] = 0$.

The resulting instrumental variable (IV) estimate of σ is reported in table 2. We estimate $\sigma = 5.40$ (with a standard error of 1.26 when clustered at the crop importer and crop exporter levels), and (as also reported in table 2) the instrument both is strong, by conventional standards, and has the expected negative sign (such that higher productivity leads to lower

TABLE 2
PARAMETER ESTIMATES

Parameter	Estimate (SE)	Method	Specification	Observations	First-Stage (<i>F</i> -Statistic)
σ	5.40 (1.26)	IV	$\ln(X_i^k/X_j^k) = M_i^k + (1 - \sigma) \ln p_i^k + \varepsilon_{ij}^k$, $\ln(1/F_i \sum_{j \in \mathcal{F}_i} A_j^k)$ as IV for $\ln p_i^k$	4,120	-.06 (19.64)
κ	2.82 (.54)	IV	$\ln(X_i^k/X_j^k) = M_j^k + (1 - \kappa) \ln p_j^k + \varepsilon_{ij}^k$, $\ln(1/F_j \sum_{i \in \mathcal{F}_j} A_i^k)$ as IV for $\ln p_j^k$	408	-.12 (22.30)
θ	2.46 [2.28, 2.62]	NLS	$\min_{\theta, (\alpha_i)} \sum_k [\ln Q_i^k(\theta, \alpha_i) - \ln Q_i^k]^2$ s.t. $L_i(\theta, \alpha_i) = L_i \forall i$	379	NA

NOTE.—In the first row (estimates of σ) the standard error (SE) and first-stage *F*-statistic are clustered at the crop exporter and crop importer levels. In the second row (estimates of κ) the SE and first-stage *F*-statistic are clustered at the importer level. In the third row (estimates of θ) the 95 percent confidence interval is estimated via a bootstrap procedure with replacement at the country level

prices).⁸ Although our estimation strategy is very different, this is in line with the estimates of Broda and Weinstein (2006), whose median elasticity for our crops is equal to 2.6.⁹

Having estimated σ , we then solve for the remaining $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma}$ as residuals. Namely, we find $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma}$ for all $i, j \in \mathcal{I}$ and $k \in \mathcal{K}$ for which $X_{ij}^k > 0$ so that equations (11) and (12) simultaneously hold for all crops and countries. Although this estimation procedure does not allow us to identify separately lower-level demand shifters, β_{ij}^k , from trade costs, τ_{ij}^k , the composite shock, $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma}$, is all we need to construct equilibria in Section VI.¹⁰

Step 2: The second step of our procedure is similar to the first one: the price index, P_j^k , plays the role of the individual crop price, p_i^k , whereas crop expenditure, X_j^k , plays the role of bilateral trade flows, X_{ij}^k .

For all crops and countries with positive expenditure, $X_j^k > 0$, we can again use equation (10) and take logs to get

$$\ln(X_j^k/X_j) = M_j + (1 - \kappa) \ln P_j^k + \varepsilon_j^k, \quad (13)$$

where the first term on the right-hand side,

$$M_j \equiv -\ln \left[\sum_{l \in \mathcal{K}: X_j^l > 0} \beta_j^l (P_j^l)^{1-\kappa} \right],$$

can now be treated as an importer fixed effect, and the final term, $\varepsilon_j^k \equiv \ln(\beta_j^k)$, reflects idiosyncratic demand shocks across crops.¹¹ Without loss of generality, we again normalize these shocks such that

⁸ In principle, one could estimate a separate elasticity of substitution across varieties within each crop k , i.e., a separate σ^k for all k . When doing so, we cannot reject (at standard levels of statistical significance) the null that all σ^k 's are equal (the p -value of this test is .17). This remains true when we expand our sample to one that includes the 20 most important crops worldwide, in which case we estimate that $\sigma = 5.13$ (standard error [SE] = 1.29). Similarly, we cannot reject (at standard levels) the null that there is a separate σ_j for each importing country j (the p -value of this test is .08). For simplicity, and in line with the model developed in Sec. III, we therefore focus on one pooled estimate of σ that is the same across all crops and importers.

⁹ This is excluding sugarcane and cotton, whose GAEZ codes we have not been able to match to the 5-digit standard international trade classification codes used in Broda and Weinstein (2006).

¹⁰ Formally, using eq. (6), the good market equilibrium condition (9) can be rearranged in value terms as

$$p_i^k Q_i^k = \sum_{j \in \mathcal{I}} \beta_j \frac{\beta_j^k (P_j^k)^{1-\kappa}}{\sum_{l \in \mathcal{K}} \beta_j^l (P_j^l)^{1-\kappa}} \frac{\beta_{ij}^k (\tau_{ij}^k p_i^k)^{1-\sigma}}{\sum_{n \in \mathcal{I}} \beta_{nj}^k (\tau_{nj}^k p_n^k)^{1-\sigma}}.$$

This depends only on β_{ij}^k and τ_{ij}^k through the composite shock, $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma}$.

¹¹ There are two out of 500 crop and country pairs in our data set that have zero expenditures, i.e., $X_j^k = 0$. For any such pair (k, j) , step 1 of our procedure already implies that $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma} = 0$ for all exporters i . Thus total demand for these two pairs, X_j^k , is necessarily zero, independently of the value of β_j^k .

$$\sum_{k \in \mathcal{K}: X_j^k > 0} \varepsilon_j^k = 0. \quad (14)$$

The key difference between step 1 and step 2 is that rather than observe P_j^k directly in the data, we construct it as

$$P_j^k = \left[\sum_{i \in \mathcal{I}} \beta_{ij}^k (\tau_{ij}^k p_i^k)^{1-\sigma} \right]^{1/(1-\sigma)}$$

using data on crop prices, p_i^k , as well as our estimates of σ and $\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma}$.

To estimate κ , we still need to address the endogeneity between demand shocks, ε_j^k , and prices, P_j^k , at that higher level of aggregation. To do so, we now propose to instrument P_j^k with Z_j^k ; that is, we impose the following exclusion restriction: $E[Z_j^k \varepsilon_j^k] = 0$. As reported in table 2, our IV estimate is $\kappa = 2.82$ (SE = 0.54, clustered at the importer level). As in step 1, once the elasticity of substitution, κ , is known, we can solve for β_j^k for all $j \in \mathcal{I}$ and $k \in \mathcal{K}$ such that $X_j^k > 0$ as residuals using equations (13) and (14).

Step 3: The final step of our procedure is the simplest. Since we have assumed log preferences at the upper level, the demand shifters can be read directly from data on total expenditure across crops. Specifically, equation (10) implies that $\beta_j = X_j$ for all $j \in \mathcal{I}$.¹²

B. Supply

Since the productivity of fields across crops, A_i^k , is directly observable in the GAEZ data, the technological parameters that need to be estimated are the extent of technological heterogeneity, θ , as well as the country-specific labor cost shifters, α_i . To emphasize the dependence of our supply-side predictions on these unknown parameters, we let $Q_i^k(\theta, \alpha_i)$ denote predicted output of crop k in country i as a function of (θ, α_i) ,

$$Q_i^k(\theta, \alpha_i) \equiv \sum_{f \in \mathcal{F}_i} s_i^f A_i^k \left[\frac{(p_i^k A_i^k)^\theta}{\alpha_i^\theta + \sum_{l \in \mathcal{K}} (p_i^l A_i^l)^\theta} \right]^{(\theta-1)/\theta}. \quad (15)$$

¹² In principle, one could have assumed a more general upper-level utility function,

$$U_i = C_i^0 + \frac{\alpha_i \zeta}{\zeta - 1} C_i^{(\zeta-1)/\zeta},$$

where the own-price elasticity $\zeta > 0$ is allowed to be different from one. Under this assumption, one could estimate ζ in the same way as we have estimated σ and κ . We have experimented with such an approach, but given the small number of observations at that level of aggregation (i.e., 50), we obtain very imprecise IV estimates of ζ . By necessity, we therefore assume log preferences for our baseline analysis, subject to the caveat discussed in fn. 3.

Similarly, we let $L_i(\theta, \alpha_i)$ denote predicted land allocated to all crops in country i as a function of (θ, α_i) ,

$$L_i(\theta, \alpha_i) \equiv \sum_{k \in \mathcal{K}} \sum_{f \in \mathcal{F}_i} s_i^f \left[\frac{(p_i^k A_i^f)^{\theta}}{\alpha_i^{\theta} + \sum_{l \in \mathcal{K}} (p_i^l A_i^l)^{\theta}} \right].$$

This is equal to the sum across all crops and all fields of the total field size, s_i^f , times the share of parcels allocated to a crop, π_i^{fk} .

To estimate (θ, α_i) , we proceed as follows. Conditional on θ , we look for a labor cost shifter, α_i , such that the total amount of land allocated to crops predicted by the model, $L_i(\theta, \alpha_i)$, exactly matches the total amount of land allocated to crops in the data, $L_{i,}$ country by country. Since the only role of α_i in our model is to determine the share of agricultural land, we believe that this is the most natural moment to estimate this parameter.¹³ Given a vector of labor cost shifters (α_i) for all countries in our data set, we then search for the extent of heterogeneity θ such that output predicted by the model, $Q_i^k(\theta, \alpha_i)$, best matches output observed in the data, Q_i^k . Formally, we use nonlinear least squares (NLS) to estimate θ and α_i for all $i \in \mathcal{I}$ as the solution of

$$\min_{\theta, (\alpha_i)} \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} [\ln Q_i^k(\theta, \alpha_i) - \ln Q_i^k]^2$$

subject to

$$L_i(\theta, \alpha_i) = L_i \quad \text{for all } i \in \mathcal{I}.$$

Table 2 presents our parameter estimate of θ . We find that $\theta = 2.46$ with a 95 percent confidence interval given by $[2.28, 2.62]$.¹⁴ Since θ is inversely related to the within-field, within-crop productivity dispersion in agriculture, this low estimate of θ suggests that unobserved land heterogeneity is substantial.¹⁵

¹³ For one country in our data set, Egypt, the maximum amount of land that can be allocated to crops according to the GAEZ data, i.e., the total area with positive crop yields in at least one crop, is strictly less than the total amount of land allocated to crops in the FAO data. This reflects the fact that we use the same GAEZ scenario for all countries in our data set and that this scenario may be a poor description of Egypt's growing conditions. In Sec. VII.C, we explore the sensitivity of our counterfactual results to assigning different countries to different GAEZ scenarios. For now, we deal with this issue in the simplest possible way by setting $\alpha_{\text{Egypt}} = 0$ and increasing uniformly the size of all Egyptian fields with positive crop yields so that $L_{\text{Egypt}}(\theta, 0) = L_{\text{Egypt}}$.

¹⁴ While there is no closed form for the standard error of θ , we use a bootstrap procedure, with replacement at the country level, with 400 replications to estimate the 95 percent confidence interval.

¹⁵ There is also a large amount of variation between countries in our estimates of the labor cost shifters, α_i , though no systematic relationship between α_i and country i 's level of development. This is consistent with the observation that poor countries have lower wages, which tends to push labor costs down, but also more labor-intensive technologies, which tends to push labor costs up.

The previous estimation procedure differs from our demand-side estimation in two important ways. First, although our theory predicts that revenue for each field and crop takes a simple CES form, $p_i^k Q_i^{fk} \propto (p_i^k A_i^{fk})^\theta$, we do not have output disaggregated at the field level, Q_i^{fk} . If we did, the revenue elasticity at the field level, which is governed by θ , could be estimated using a simple linear regression of $\ln(p_i^k Q_i^{fk})$ on $\ln p_i^k$ and field fixed effects, the same way we estimated σ and κ in Section V.A. Given an estimate of θ , we could then infer α_i using the values of the field fixed effects in country i . Here instead, we need to aggregate output at the country level for each crop, which breaks down the log linearity of our model and leads us to use NLS.

Second, we do not have an instrument to deal with idiosyncratic supply shocks, that is, discrepancies between true productivity and GAEZ estimates of crop yields at the field level. This leads us to adopt the extreme assumption, implicit in the estimation procedure above, that any difference between predicted and observed output comes from measurement error in output rather than productivity levels. With access to data on field-level output, one could deal with measurement error in productivity by regressing $\ln(p_i^k A_i^{fk})$ on $\ln(p_i^k Q_i^{fk})$ under the polar assumption that output is measured without errors. This is the approach followed by Sotelo (2013). Reassuringly, using Peruvian data on crop output disaggregated at the district level, he finds an estimate of θ equal to 2.06, very close to ours.

C. Goodness of Fit

Before we go on to consider how our model behaves under the counterfactual scenario of new agricultural productivities brought about by climate change, it is natural to ask how well the model fits the data within sample. Here, we report several moments—related to output, land use, and trade flows—as generated by the model at the previously estimated parameters and compare them to the same moments in the data.

Output.—Figure 2 compares log output in the model, as given by equation (15), and in the data at our preferred estimates of θ and (α_i) . In panel *a*, we use the crop prices observed in the raw data to compute output levels. Hence our demand estimates play no role here. Panel *a* merely plots the fit of the moment that we use to estimate θ . Panel *b*, in contrast, uses the crop prices predicted by the model, that is, the prices that solve the good market-clearing condition (9) given our estimates of demand and supply parameters.

In both cases, there is a positive and statistically significant correlation between the model and the data. In panel *b*, for instance, a regression of log observed output on log predicted output, with a constant, yields a coefficient estimate of 0.67 with a standard error (clustered at the country




FIG. 2.—Output fit. Observed output (panel *a*) versus predicted output (panel *b*), across all countries *i* and crops *k*, where predicted output in panel *a* is calculated on the basis of FAOSTAT producer prices while that in panel *b* is calculated on the basis of equilibrium model prices. The 45-degree line is also shown.

level) of 0.06. While the fit of the model in terms of log output 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. In particular, the estimated constant in the regression illustrated in panel *b* is equal to 4.07, implying that predicted output is considerably lower 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, at least in part, from our assumption that agricultural technologies around the world differ only in terms of their labor intensity. In Section VII.C, we propose an alternative estimation procedure that deals explicitly with this issue.

While figure 2 focuses only on nonzero observations, the fit of the model along the extensive margin—that is, whether a crop is produced or not—is also relatively strong. For example, the correlation between a dummy variable for a nonzero in the data, $Q_i^k > 0$, and that for a nonzero in the model, $Q_i^k(\theta, \alpha_i) > 0$, is .71 when using predicted prices and .66 when using observed prices.

Figure 3 illustrates that the model is more successful in terms of matching relative revenue levels, even though this moment was not used in our estimation procedure. In panels *a* and *b* we plot the revenue share predicted by the model using observed and predicted prices, respectively,




FIG. 3.—Revenue share fit. Observed revenue shares (panel *a*) versus predicted revenue shares (*b*), across all countries *i* and crops *k*, where the predicted revenue share in panel *a* is calculated on the basis of FAOSTAT producer prices while that in panel *b* is calculated on the basis of equilibrium model prices. The 45-degree line is also shown.

against the equivalent revenue share in the FAO data for each crop and country. The line of best fit in panel *b* has a slope coefficient of 0.81 (SE = 0.07, clustered at the country level) and an estimated constant of -0.77 ; the R^2 for this regression is .52. The 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.¹⁶

Land use.—Since we will model climate change as a counterfactual shock to land productivity, a key question is whether our model is able to replicate the observed allocation of land across crops and countries. Figure 4 sheds light on this issue by plotting land shares allocated to a given crop in a given country in the model, that is, $\sum_{f \in F_i} s_i^f \pi_i^{fk} / L_i$, against its value in the data. As before, panel *a* uses observed prices to compute π_i^{fk} , whereas panel *b* uses predicted prices. We find that our model can also match the observed land allocation, especially when predicted prices are used. In panel *b*, a regression of the predicted land shares on the observed land shares, with a constant, yields a coefficient estimate of 0.47 with SE = 0.05 (clustered at the country level) and $R^2 = .27$.

Trade flows.—Finally, since one of our main goals is to assess whether international trade may help alleviate the consequences of climate change,

¹⁶ Figure 3 suggests that predicted revenue shares are typically lower than observed revenue shares, especially when using observed prices (panel *a*). This is to be expected because the model predicts fewer zeroes than in the data; i.e., conditional on both observed and predicted revenue being positive, the predicted value will be lower.




FIG. 4.—Land use fit. *a*, land share fit at observed prices; *b*, land share fit at predicted prices. Observed versus predicted land shares, across all countries i and crops k , where the predicted land share in panel *a* is calculated on the basis of FAOSTAT producer prices while that in panel *b* is calculated on the basis of equilibrium model prices. The 45-degree line is also shown.

it is important to check that our model also does a reasonable job at matching the observed cross section of international trade flows. Figure 5 plots bilateral trade flows in the model, X_{ij}^k , using predicted prices against the same trade flows observed in the data. Note that since our model allows for unobserved trade costs and demand shocks, the only reason why the model does not perfectly match trade data is that predicted prices do not perfectly match observed prices (which itself comes from the fact that output levels in the model do not perfectly match output levels in the data). If we were to use observed prices to compute predicted trade flows, all observations would be, trivially, on the 45-degree line. In this sense the goodness of fit for trade flows illustrated in figure 5 also provides indirect evidence for how observed prices compare to model-predicted prices.

A regression of the predicted trade flows on the observed trade flows, with a constant, yields a coefficient estimate of 0.78 with $SE = 0.02$ (clustered at the importer level) and $R^2 = .74$. The main discrepancy between theory and data comes from the fact that the model predicts an excess of very small bilateral trade flows, that is, values for which the log is less than zero, with no counterparts in the data. This is to be expected because, in the trade data, all bilateral trade flows are reported as integer values of 2009 US dollars (which presumably reflects rounding procedures), so the lowest nonzero observation is US\$1. Since figure 5 reports log trade flows, these very small numbers lead to average log trade flows that are smaller on average than those observed in the data




FIG. 5.—Bilateral trade fit. Observed bilateral trade versus predicted bilateral trade, across all exporters i , importers j , and crops k , where predicted trade is calculated on the basis of equilibrium model prices. The 45-degree line is also shown.

(although by construction total expenditure on crops is the same in the model and in the data).

VI. Counterfactual Simulations

In the previous section, we have shown that our model of international trade fits output, land, and trade data reasonably well, including along some dimensions that were not directly targeted in the estimation procedures of Sections V.A and V.B. We now use our estimated parameters to simulate our model and quantify the various channels through which climate change affects agricultural markets in a global economy. To get a sense of the magnitudes involved, Section VI.A first describes the average productivity shocks caused by climate change for all countries in our data set. Section VI.B then uses the full structure of the model to quantify the overall welfare impact of climate change if production and

trade patterns can fully adjust. Finally, Sections VI.C and VI.D explore the extent to which each of these two margins of adjustments helps alleviate the adverse consequences of climate change.

A. First-Order Approach

As discussed in Section IV.B, we model climate change as a shock to 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 V.

A simple way to aggregate the consequences of climate change is to compute for each country in our data set the following weighted average of the micro-level productivity shocks that it is expected to trigger:

$$\Delta A_i = \sum_{k \in \mathcal{K}} \sum_{f \in \mathcal{F}_i} \rho_i^{fk} \Delta A_i^{fk}, \quad (16)$$

where $\rho_i^{fk} \equiv p_i^k Q_i^{fk} / Y_i$ denotes the initial share of country i 's GDP associated with the production of crop k in field f and $\Delta A_i^{fk} \equiv [(A_i^{fk})' - A_i^{fk}] / A_i^{fk}$ denotes the percentage change in productivity for that particular crop and location.

The previous formula is attractive for two reasons. First, abstracting from terms-of-trade movements, ΔA_i provides a first-order approximation to the welfare impact of climate change in country i , expressed as a percentage change of country i 's GDP.¹⁷ Second, equation (16) relies as little

¹⁷ The argument can be sketched as follows. Let $r_i(p, A)$ and $e_i(p, u_i)$ denote the revenue and expenditure functions in country i . One can express these two functions as

$$\begin{aligned} r_i(p, A) &= \max_{L_i^{fk}(\omega)} \left\{ N_i A_i^0 + \sum_k \sum_f \int_0^1 \left[p_i^k A_i^{fk}(\omega) - \sum_k \sum_f \int_0^1 A_i^0 v_i^f(\omega) \right] L_i^{fk}(\omega) d\omega \right\}, \\ e_i(p, u_i) &= \min_{C_{ji}^k} \left\{ u_i + \sum_j \sum_k p_{ji}^k C_{ji}^k - \alpha_i \ln C_i \right\}. \end{aligned}$$

Now consider a change in crop productivity, holding crop prices fixed. The envelope theorem implies

$$\begin{aligned} dr_i(p, A) &= \sum_k \sum_f \int_0^1 p_i^k dA_i^{fk}(\omega) L_i^{fk}(\omega) d\omega \\ &= \sum_k \sum_f \int_0^1 p_i^k A_i^{fk}(\omega) L_i^{fk}(\omega) \left[\frac{dA_i^{fk}(\omega)}{A_i^{fk}(\omega)} \right] d\omega, \\ de_i(p, u_i) &= 0. \end{aligned}$$

Using the fact that

$$\frac{dA_i^{fk}(\omega)}{A_i^{fk}(\omega)} = \frac{dA_i^{fk}}{A_i^{fk}},$$

as possible on the structure of the model. Crop prices, p_i^k , country's GDP, Y_i , and productivity shocks, ΔA_i^k , are directly observable in the data.¹⁸ Thus we only need the structure of the model to go from observed revenue shares at the country level, $\rho_i^k \equiv p_i^k Q_i^k / Y_i$, to revenue shares at the field level, ρ_i^{fk} , which, unfortunately, are not observable in the data. We now describe how we obtain those.

To predict revenue shares at the field level, ρ_i^{fk} , we first use our model to compute the share of crop k 's output associated with field f in country i , $\phi_i^{fk} \equiv Q_i^{fk} / Q_i^k$. The same algebra used to derive equation (8) in Section III implies

$$\phi_i^{fk} = \frac{(A_i^k)^\theta [\alpha_i^\theta + \sum_{l \in \mathcal{K}} (p_i^l A_i^{fl})^\theta]^{(1-\theta)/\theta}}{\sum_{g \in \mathcal{F}_i} (A_i^g)^\theta [\alpha_i^\theta + \sum_{l \in \mathcal{K}} (p_i^l A_i^{gl})^\theta]^{(1-\theta)/\theta}},$$

which can be computed using our estimates of α_i and θ as well as the GAEZ productivity estimates and price data. Given a value of ϕ_i^{fk} , we then use the identity $\rho_i^{fk} = \rho_i^k \times \phi_i^{fk}$ to obtain revenue shares at the field level for all crops and countries.¹⁹

Column 2 of table 3 reports our estimates of ΔA_i for the 50 countries in our data set. For the most affected country, Malawi, ΔA_i corresponds to a whopping 38 percent decrease in total GDP. For most other countries, the consequences of climate change are an order of magnitude smaller but remain nontrivial. For the median country, Indonesia, ΔA_i corresponds to a 1.26 percent decrease in total GDP.²⁰

An obvious reason why the welfare consequences of climate change predicted by our model may appear modest for most countries is that the crops considered in our analysis represent only a small fraction of GDP in each country. To get a sense of the role played by the size of the agricultural sector, column 1 reports the total value of crop output divided by GDP for each of the 50 countries in our data set. Not surprisingly, the

this implies

$$\frac{dr_i(p, A) - de_i(p, u_i)}{r_i(p, A)} = \sum_k \sum_f \rho_i^k \frac{dA_i^{fk}}{A_i^{fk}},$$

which is the counterpart of eq. (16) above.

¹⁸ We obtain GDP data in 2009 for all countries from the World Bank except Myanmar, whose GDP is taken from the International Monetary Fund.

¹⁹ Alternatively, one could directly use the predictions of the model in terms of output levels, Q_i^{fk} , and compute revenue shares at the field level as $\rho_i^{fk} = p_i^k Q_i^{fk} / Y_i$. This would lead to different revenue shares at both the field and country levels because our model does not perfectly match output levels in the data. Since our goal here is to rely on the structure of the model as little as possible, we find it preferable to use predicted shares, ϕ_i^{fk} , which allow us to match perfectly observed revenue shares, ρ_i^k , at the country level.

²⁰ Since there is an even number of countries in our data set, 50, the "median country" refers to the country such that 24 countries have larger losses and 25 countries have smaller losses in absolute value.

TABLE 3
BASELINE COUNTERFACTUAL RESULTS

COUNTRY	CROP OUTPUT AS % OF TOTAL GDP (1)	WEIGHTED AVERAGE PRODUCTIVITY CHANGE (ΔA_i) (2)	WELFARE CHANGE (as % of Total GDP)		
			Full Adjustment (ΔW_i) (3)	No Production Adjustment (ΔW_i^p) (4)	No Trade Adjustment (ΔW_i^T) (5)
Algeria	2.51	-.88	-.94	-1.60	-.99
Argentina	4.24	.20	.54	.46	.61
Australia	.79	-.23	-.08	-.07	-.07
Bangladesh	12.36	-3.44	-8.18	-12.37	-8.19
Brazil	3.19	-1.94	-1.47	-2.77	-1.57
Burma	49.81	-5.96	-11.17	-18.20	-11.09
Cameroon	9.38	-3.07	-3.18	-4.40	-3.21
Canada	.60	.26	.59	.47	.63
China	4.56	-.46	.17	-1.81	.19
Colombia	2.36	-1.47	-1.75	-3.85	-1.80
D.R. Congo	15.35	-7.04	-10.70	-19.91	-10.71
Cote d'Ivoire	7.30	-2.52	-2.79	-5.19	-3.00
Ecuador	8.26	-4.61	-.24	-1.08	-.16
Egypt	5.25	-1.96	-.34	-2.33	-.97
Ethiopia	11.93	-3.75	-.66	-4.97	-.66
France	.33	-.08	-.03	-.02	-.04
Germany	.17	.01	.02	-.01	.03
Ghana	10.72	-6.37	-10.63	-17.40	-10.64
Greece	1.19	-.13	.00	-.05	.00
India	7.41	-1.84	-1.65	-2.71	-1.65
Indonesia	8.12	-1.26	-2.53	-5.17	-2.71
Iran	2.46	-.94	-.43	-.99	-.44
Italy	.38	-.07	-.03	-.08	-.03
Japan	.62	.04	.04	-.09	.04
Kazakhstan	3.48	-.49	-.75	-2.27	-.78
Korea, South	1.40	.17	-.33	-.68	-.34
Malawi	74.92	-38.41	-49.07	-82.36	-49.25
Malaysia	4.78	-.35	-.55	-1.10	-.58
Mexico	1.06	-.39	-.38	-.85	-.40
Morocco	3.27	-1.57	-1.87	-2.97	-1.92
Netherlands	.33	.08	.15	.02	.13
Nigeria	7.92	-2.59	-5.34	-8.68	-5.39
Pakistan	8.51	-1.74	-1.34	-2.31	-1.27
Philippines	5.65	-1.32	-1.89	-4.04	-1.94
Poland	.66	.06	.14	.18	.15
Romania	2.80	-.60	-.50	-.57	-.51
Russia	1.62	-.03	1.06	.10	1.07
South Africa	1.24	-.54	-.52	-.95	-.52
Spain	.43	-.09	.01	-.12	-.02
Sudan	3.89	-3.35	-6.84	-10.82	-7.55
Tanzania	19.56	-5.91	-5.43	-9.70	-5.43
Thailand	6.05	-2.61	-7.33	-12.02	-7.72
Turkey	2.85	-.30	.15	-.19	.14
Uganda	33.83	-3.82	-3.95	-7.61	-3.96
Ukraine	5.58	-.54	-.19	.15	-.11
United Kingdom	.16	.00	.19	.05	.21
United States	.83	-.14	-.08	-.12	-.08

TABLE 3 (Continued)

COUNTRY	CROP OUTPUT AS % OF TOTAL GDP (1)	WEIGHTED AVERAGE PRODUCTIVITY CHANGE (ΔA_i) (2)	WELFARE CHANGE (as % of Total GDP)		
			Full Adjustment (ΔW_i) (3)	No Production Adjustment (ΔW_i^p) (4)	No Trade Adjustment (ΔW_i^T) (5)
Uzbekistan	20.76	-2.22	1.43	-1.05	1.43
Venezuela	2.50	-2.24	-3.76	-6.65	-4.01
Viet Nam	22.50	-4.16	-6.53	-11.52	-6.80
World total	1.80	-.34	-.26	-.78	-.27

correlation between this ratio and ΔA_i is high, $-.85$. Malawi, in particular, which was an outlier in terms of the aggregate productivity shock, ΔA_p , is also an outlier in terms of its crop output to GDP ratio. For Indonesia, the value of crop output divided by GDP is 8.12 percent. So in output terms the estimated effect of climate change, -1.26 percent of GDP, corresponds to a 16 percent decrease in total crop value.

The last row of table 3 reports the weighted sum of the productivity shocks across countries,

$$\Delta A_W \equiv \sum_{i \in \mathcal{I}} \left(\frac{Y_i}{Y_W} \right) \Delta A_i,$$

where Y_W is world GDP. An appealing feature of this statistic is that it is not affected by terms-of-trade effects, which cancel out when summing up across countries. According to this formula, climate change for the world as a whole amounts to a 0.34 percent decrease in world GDP.²¹ In the next subsection, we demonstrate that the previous number provides a very good approximation of the welfare consequences of climate change, at least at the world level, when both production and trade patterns are allowed to adjust.

B. Full Adjustment

To measure the overall welfare consequences of climate change under full adjustment, we now solve for competitive equilibria before and after

²¹ An alternative way to approximate the aggregate consequences of climate change without using at all the structure of our model would be to compute $\Delta \bar{A}_i = \sum_{k \in \mathcal{K}} \rho_i^k \Delta A_i^k$, where $\Delta \bar{A}_i^k$ denotes the arithmetic average of ΔA_i^{fk} across all fields f within country i for crop k . At the world level, such an approach would lead to

$$\Delta \bar{A}_W \equiv \sum_{i \in \mathcal{I}} \left(\frac{Y_i}{Y_W} \right) \Delta \bar{A}_i = -0.21 \text{ percent}$$

of world GDP, slightly smaller in absolute value than the average productivity shock, ΔA_W , obtained using the predictions of our model at the field level.

climate change. The first equilibrium is characterized by equations (6), (8), and (9) with crop productivity, A_i^{fk} , as measured in the pre-climate change GAEZ data and all other structural parameters as described in Section V. The second equilibrium is characterized by the same equations and structural parameters, except for crop productivity, which is now set to $(A_i^{fk})'$, as measured in the post-climate change GAEZ data.²²

Given quasi-linear preferences, welfare changes can be simply computed as changes in social surplus, again expressed as a fraction of GDP in the initial equilibrium:

$$\Delta W_i = \frac{(Y_i)' - Y_i + (\beta_i \ln C_i - P_i C_i)' - (\beta_i \ln C_i - P_i C_i)}{Y_i}, \quad (17)$$

where primes denote variables associated with the counterfactual equilibrium. The term ΔW_i corresponds to both the equivalent variation and the compensating variation associated with a given counterfactual scenario expressed as a fraction of country i 's GDP in the initial equilibrium.²³ Column 3 of table 3 reports ΔW_i for all countries in our data set. While the welfare consequences of climate change differ across countries, welfare losses tend to have the same order of magnitude as predicted by the first-order approximation given in Section VI.A. For Malawi, the welfare loss is now equal to 49.1 percent, whereas for the median country, Malaysia in this counterfactual exercise, the welfare loss is equal to 0.55 percent. Part of the discrepancy between the first-order approximation and the predictions of the full model comes from terms-of-trade effects. Indeed, we see that for the world as a whole, the welfare change $\Delta W_w \equiv \sum_{i \in \mathcal{I}} (Y_i/Y_w) \Delta W_i$ is now equal to 0.26 percent of world GDP, close to the 0.34 percent obtained in Section VI.A. Since the value of output in our 10 crops is equal to 1.8 percent of world GDP, this corresponds to about one-sixth of total crop value.

C. No Production Adjustment

Taking the previous numbers as our benchmark, we now explore the extent to which changes in patterns of specialization at the field level can help countries adapt to and mitigate the consequences of climate change.

²² Our counterfactual exercise therefore rules out endogenous technology adoption in the agricultural sector in response to climate change. Consumption and factor allocations are free to vary in response to TFP shocks from A_i^{fk} to $(A_i^{fk})'$, but production functions remain given by eq. (8). Although adding more margins of adjustment would further mitigate the consequences of climate change, it is not clear how these extra margins would affect (if at all) the relative importance of changes in production and trade patterns studied in the next two subsections.

²³ Since the total value of crops predicted by our model is strictly lower than observed GDP for all countries in our data set, we implicitly set labor productivity, A_i^0 , such that country i 's GDP in the initial equilibrium exactly matches GDP in the data.

To do so, we ask the following counterfactual question. Imagine a world economy in which preferences and technology are as described in Section III, but factor allocation cannot adjust to changes in productivity from A_i^{fk} to $(A_i^{fk})'$. How much larger would have been the welfare losses, ΔW_i^P , around the world?

We construct our counterfactual equilibrium under climate change without production adjustment as follows. On the supply side, we assume that each parcel ω produces the same crop (if any) before and after climate change. Thus total output of crop k in country i in the counterfactual equilibrium is given by

$$\begin{aligned}(Q_i^k)' &= \sum_{f \in \mathcal{F}_i} s_i^f \pi_i^{fk} E[(A_i^{fk})'(\omega) | p_i^k A_i^{fk}(\omega)] \\ &= \max\{A_i^0 \nu_i^f(\omega), p_i^1 A_i^{f1}(\omega), \dots, p_i^K A_i^{fK}(\omega)\},\end{aligned}$$

where π_i^{fk} is still given by equation (7). Using the fact that $(A_i^{fk})'(\omega) = (A_i^{fk})' \times [A_i^{fk}(\omega)/A_i^{fk}]$, this can be rearranged as

$$(Q_i^k)' = \sum_{f \in \mathcal{F}_i} s_i^f (A_i^{fk})' (\pi_i^{fk})^{(\theta-1)/\theta} \quad \text{for all } i \in \mathcal{I} \text{ and } k \in \mathcal{K}. \quad (18)$$

Compared to equation (8), which was used to compute output levels under climate change in Section VI.B, the only difference is that π_i^{fk} corresponds to parcel shares in the initial equilibrium, not the optimal parcel allocation given the new productivity levels, $(A_i^{fk})'$. All other equilibrium conditions are unchanged: the vector of consumption satisfies equation (6) at the new crop prices and supply equals demand for all varieties of all crops, as stated in equation (9).

Column 4 of table 3 reports the welfare effects of climate change in the absence of factor reallocation, ΔW_i^P . For the most affected country, Malawi, the adverse consequences of climate change under this counterfactual scenario, -82 percent, are almost twice as large as those reported in Section VI.B. For the world as a whole, a similar picture emerges: ΔW_W^P is equal to -0.78 percent, three times the loss obtained when production is allowed to adjust at the field level.²⁴

If the productivity shocks caused by climate change had been uniform across crops and fields, and hence comparative advantage stable, such

²⁴ Using the same logic as in fn. 17, one can check that the number presented in Sec. VI.A, $\Delta A_W = -0.34$ percent, remains a first-order approximation to the welfare change without production adjustments, $\Delta W_W^P = -0.78$ percent. Nevertheless, we see that the first-order approximation is much less accurate under this counterfactual scenario. This reflects larger changes in the price of crops in response to climate change when factors cannot reallocate across fields. The same price effects also explain why social surplus at the world level is not a linear function of productivity, A_i^{fk} , and, in turn, why ΔA_W does not offer a lower bound on welfare changes at the world level.




FIG. 6.—Comparison of simulated welfare changes. Panel *a* plots the real income change (as a percentage of total GDP) when allowing for full adjustment (*x*-axis) and when allowing for no change in production land shares (*y*-axis). Panel *b* does the same but when allowing for no change in trade shares (*y*-axis). In both cases, for expositional clarity, Malawi is omitted (but can be found in table 3). The 45-degree line is also shown.

a margin of adjustment would have been irrelevant. But this is not the case in our data set. As figure 1 already suggested, there is a considerable amount of heterogeneity across space in terms of the consequences of climate change. The importance of factor reallocations is best described by plotting, as in figure 6, the welfare loss caused by climate change without reallocation against the same welfare loss with reallocation for all countries in our data set. In panel *a*, almost all observations lie well below the 45-degree line.²⁵ This illustrates how farmers' ability to substitute crop production in response to changes in comparative advantage—which our extremely rich micro-level data set gives us a unique opportunity to study—may substantially mitigate the ill effects of climate change.

D. No Trade Adjustment

To conclude our quantitative analysis, we turn to the consequences of international trade. Given the ability of farmers to switch production at the field level, our goal is to explore the extent to which adjustments to trade patterns may further alleviate the adverse consequences of climate change. Since international trade flows are the difference between production and consumption, this last counterfactual exercise will also indirectly shed light on the importance of adjustments on the demand side.

²⁵ Individual countries may be better off in the counterfactual equilibrium without adjustments because of strong, positive terms-of-trade effects. In fig. 6, we have excluded Malawi, which is an outlier in terms of welfare changes with or without adjustments.

We follow the same basic approach as in Section VI.C. We consider a world economy in which preferences and technology are as described in Section III, but international trade flows cannot adjust in response to productivity changes from A_i^k to $(A_i^k)'$. Formally, we assume that in addition to the good market-clearing condition (9), consumption and output in this new counterfactual equilibrium, $(C_{ij}^k)'$ and $(Q_i^k)'$, satisfy

$$\left[\sum_{j \neq i} \tau_{ij}^k (C_{ij}^k)' \right] / (Q_i^k)' = \left(\sum_{j \neq i} \tau_{ij}^k C_{ij}^k \right) / Q_i^k \quad \text{for all } i \in \mathcal{I} \text{ and } k \in \mathcal{K}. \quad (19)$$

Equation (19) requires total exports of a crop k in country i to remain a constant fraction of its output.²⁶ If country i is under autarky in the initial equilibrium, equation (19) merely requires country i to remain under autarky, as in the simple example of Section II.

For equation (19) to hold in a competitive equilibrium under the new productivity levels, $(A_i^k)'$, we relax equation (5) and assume that counterfactual crop prices satisfy

$$(p_{ij}^k)' = \delta_i^k \tau_{ij}^k (p_i^k)' \quad \text{for all } i \neq j \in \mathcal{I} \text{ and } k \in \mathcal{K}, \quad (20)$$

where the wedge δ_i^k corresponds to the shadow price of the quantity restriction (19). It is equivalent to an ad valorem trade tax. A value of δ_i^k strictly greater than one corresponds to an export tax on crop k in country i , whereas a value of δ_i^k strictly lower than one corresponds to an export subsidy. Combining equations (6) and (20), we can rearrange crop consumption as

$$(C_{ji}^k)' = \beta_i \frac{\beta_i^k ((P_i^k)')^{1-\kappa}}{\sum_{i \in \mathcal{K}} \beta_i^k ((P_i^k)')^{1-\kappa}} \frac{\beta_{ji}^k (\delta_{ji}^k \tau_{ji}^k (p_j^k)')^{-\sigma}}{\sum_{n \in \mathcal{I}} \beta_{ni}^k (\delta_{ni}^k \tau_{ni}^k (p_n^k)')^{1-\sigma}} \quad (21)$$

$$\text{for all } i, j \in \mathcal{I} \text{ and } k \in \mathcal{K},$$

where we use the convention $\delta_{ji}^k \equiv \delta_j^k$ if $i \neq j$ and $\delta_{ji}^k \equiv 1$ otherwise. All other equilibrium conditions are unchanged so that a counterfactual equilibrium can now be defined as consumption, $(C_{ji}^k)'$, output, $(Q_i^k)'$, prices, $(p_i^k)'$, and wedges, (δ_i^k) , such that equations (8), (9), (19), and (21) hold.²⁷

²⁶ Formally, eq. (19) applies only to crops and countries such that $Q_i^k, (Q_i^k)' > 0$. If $(Q_i^k)' = 0$, we do not impose any constraint since there are no exports to speak of in the counterfactual equilibrium. If $(Q_i^k)' > 0$ but $Q_i^k = 0$, there may be exports in the counterfactual equilibrium in spite of exports being zero in the initial equilibrium. In this case, we assume that exports remain equal to zero in the counterfactual equilibrium.

²⁷ In the few cases in which a country i produces a crop k but does not consume it, the wedge δ_i^k is not uniquely determined since country i will always export 100 percent of its output in the counterfactual equilibrium, regardless of the value of δ_i^k . In such cases, we simply set δ_i^k to one, which is the optimal value from the world point of view.

Column 5 of table 3 reports the welfare effects of climate change in the absence of trade adjustments, ΔW_i^T .²⁸ We see that for the median country in our baseline scenario, Malaysia, the welfare loss caused by climate change, 0.58 percent, has the same order of magnitude as the welfare loss in Section VI.B, albeit slightly smaller because of negative terms-of-trade effects. Similarly, at the aggregate level, the welfare loss at the world level is now equal to 0.27 percent, very close to the 0.26 percent found in the model with full adjustment.

As in Section VI.C, the small impact of international trade on the consequences of climate change is best seen in figure 6, panel *b*, which plots the welfare loss caused by climate change without trade adjustment against the same welfare loss with trade adjustment for all countries in our data set. Almost all observations are now around the 45-degree line, suggesting that for most countries in our data set, changes in the pattern of exports across crops do little to alleviate the consequences of climate change.

In summary, our various counterfactual exercises suggest that climate change will lower world GDP by 0.26 percent, but only if farmers can and do adjust to the new climate by switching what they grow. In an environment in which such production adjustments are prohibited (or too costly to take place), the decrease in world GDP would be three times as large (−0.78 percent). In contrast, trade adjustments, at least as formalized in this subsection, appear to play a minor role. Prohibiting changes in export shares across countries in response to climate change raises the world GDP loss only from 0.26 percent to 0.27 percent. We now investigate the sensitivity of these conclusions.

VII. Sensitivity Analysis

We split our sensitivity analysis into two parts. The first part maintains the same assumptions and estimated parameters as in our baseline analysis but uses different forecasts of productivity after climate change, $(A_i^{jk})'$, based on alternative general circulation models and emission scenarios. In short, the economic environment is the same, but the nature of climate change itself is allowed to vary. The second part instead uses the GAEZ estimates associated with the Hadley CM3 A1FI model but relaxes auxiliary assumptions about demand, supply, and trade costs. For expositional purposes, we report only welfare changes for the world as a whole, expressed as a fraction of world GDP.

²⁸ For all countries, ΔW_i^T is inclusive of country *i*'s tax revenues in the counterfactual equilibrium, which may be positive or negative depending on whether countries tax or subsidize exports. Ignoring tax revenues has very little effect on our results.

A. *Future Climatic Conditions*

In our baseline analysis, we use post-climate change GAEZ productivity estimates based on the predictions of the Hadley CM3 A1FI model. As discussed in Nakicenovic and Swart (2000), there is a substantial amount of uncertainty surrounding various emission scenarios, with no clear “most likely” or “central” scenario. The goal of this first extension is to explore the sensitivity of our counterfactual results to alternative climatic predictions.

In addition to the predictions of the Hadley CM3 A1FI model, the GAEZ estimates are available under 10 other GCM-SRES combinations: the Hadley CM3 model, under scenarios B1, B2, and A2; the MPI ECHAM4 model, under scenarios B2 and A2; the CSIRO Mk2 model, under scenarios B1, B2, and A2; and the CCCma CGCM2, under scenarios B2 and A2. Each of these general circulation models has been developed independently by a team of climatologists from a different country and then combined with the emission scenarios from the IPCC program. Compared to scenario A1FI, scenario B1 features the same pattern of population growth, with peaks in midcentury, but with rapid changes toward clean and resource-efficient technologies. In contrast, scenarios A2 and B2 assume that global population is continuously increasing, but with slower and faster technological change and economic growth, respectively, than under scenarios A1 and B1.

Table 4 reports the world welfare changes under the future climatic conditions predicted by all available combinations of general circulation

TABLE 4
COUNTERFACTUAL RESULTS UNDER ALTERNATIVE FUTURE CLIMATIC CONDITIONS

CLIMATE CHANGE SCENARIO	WORLD TOTAL WEIGHTED AVERAGE	WORLD TOTAL WELFARE CHANGE (as % of Total GDP)		
	PRODUCTIVITY CHANGE	Full	No Output	No Trade
	(ΔA_w) (1)	Adjustment (ΔW_w) (2)	Adjustment (ΔW_w^p) (3)	Adjustment (ΔW_w^T) (4)
Hadley CM3 A1FI (baseline)	-.34	-.26	-.78	-.27
Hadley CM3 B1	-.13	-.10	-.32	-.10
Hadley CM3 B2	-.20	-.17	-.44	-.17
Hadley CM3 A2	-.32	-.24	-.65	-.25
MPI ECHAM4 B2	-.06	.01	-.23	.01
MPI ECHAM4 A2	-.08	.10	-.30	.09
CSIRO Mk2 B1	-.12	-.00	-.22	-.01
CSIRO Mk2 B2	-.15	-.04	-.29	-.05
CSIRO Mk2 A2	-.18	-.04	-.39	-.05
CCCma CGCM2 B2	.06	.15	-.11	.15
CCCma CGCM2 A2	-.16	-.03	-.34	-.04

models and emission scenarios, with each row corresponding to a distinct combination. Except for the new post-climate change productivity levels, $(A_i^{fk})'$, all assumptions and estimated parameters are the same as in previous sections. The first row corresponds to the predictions of the Hadley CM3 A1FI model, as already presented in table 3. Table 4 makes it clear that there is a large amount of uncertainty over how future climatic conditions will affect agricultural productivity. As can be seen in column 1, the average change in crop productivity at the world level, ΔA_w , ranges from -0.34 percent according to the Hadley CM3 A1FI model to $+0.06$ percent according to the CCCma CGCM2 B2 model. Interestingly, most of this variation comes from the use of different general circulation models rather than alternative assumptions on population and economic growth under different emission scenarios, A1FI, A2, B1, or B2.

As one might have expected, differences in predicted aggregate productivity shocks, ΔA_w , then map into differences in predicted welfare changes under our different counterfactual scenarios, ΔW_w , ΔW_w^P , and ΔW_w^T . Welfare losses are the largest under the Hadley CM3 A1FI model and the lowest—they are actually gains—under the CCCma CGCM2 B2 model. In terms of relative welfare losses across counterfactual scenarios, however, the picture that emerges is the same as in Section VI. Without production adjustments, welfare losses are much larger. In contrast, trade adjustments have very little effect.

B. Extensive Margin of Trade

In our baseline analysis, estimated demand parameters are such that if a country does not buy a crop from a particular country, then it will never buy that crop from that country in any counterfactual scenario. As discussed in Section V.A, we set $\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma} = 0$. This implies that the set of exporting countries from which an importing country sources each crop cannot adjust in response to climate change. While consistent with the theoretical framework presented in Section III, this feature of our estimation procedure mechanically reduces the importance of international trade as a margin of adjustment since changes at the extensive margin are de facto ruled out.

Here, we propose to interpret zero trade flows less literally than in Section V.A by treating them instead as random missing observations. Intuitively, this assumption should now lead us to overestimate the magnitude of changes in trade flows in response to climate change since zero trade flows are very likely to have lower demand and higher trade costs than nonzero trade flows in practice. Formally, we start from the estimates, $\varepsilon_{ij}^k \equiv \ln[\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma}]$, of the composite of lower-level demand shifters and

TABLE 5
COUNTERFACTUAL RESULTS UNDER ALTERNATIVE MODELING ASSUMPTIONS

MODELING ASSUMPTION	WORLD TOTAL WEIGHTED AVERAGE PRODUCTIVITY CHANGE (ΔA_W) (1)	WORLD TOTAL WELFARE CHANGE (as % of Total GDP)		
		Full Adjustment (ΔW_W) (2)	No Production Adjustment (ΔW_W^p) (3)	No Trade Adjustment (ΔW_W^t) (4)
Baseline	-.34	-.26	-.78	-.27
Treatment of zero trade flows	-.34	-.21	-.74	-.23
Heterogeneous nonland inputs	-.33	-.18	-.70	-.20
Heterogeneous θ	-.34	-.24	-.77	-.25
Intranational trade costs	-.33	-.27	-.78	-.28

trade costs obtained in Section V.A for nonzero trade flows. We then project these residuals on a crop and exporter fixed effect, d_i^k , and a vector of trade barriers, x_{ij} , suggested by the previous gravity literature:

$$\epsilon_{ij}^k = d_i^k + \chi \cdot x_{ij} + \epsilon_{ij}^k. \tag{22}$$

The first term, d_i^k , aims to capture systematic differences in quality across varieties. The second term, x_{ij} , includes the (log) physical distance between the national capitals of countries i and j ; a dummy that is equal to one if $i = j$ and zero otherwise in order to allow for home bias; and a vector of dummy variables for whether countries i and j share a common border, a common language, or colonial ties or whether they are part of a regional trade agreement.²⁹ The final term, ϵ_{ij}^k , captures all other sources of trade costs and demand differences.

Row 2 of table 5 reports world welfare change results under the assumption that $\beta_{ij}^k(\tau_{ij}^k)^{1-\sigma} = \exp(d_i^k + \chi \cdot x_{ij})$ for all zero trade flows, where d_i^k and χ are obtained by estimating equation (22) using ordinary least squares (OLS).³⁰ To ease comparisons, row 1 again reports our baseline counterfactual results, as described in table 3. Not surprisingly, changes

²⁹ We obtain data on geographic barriers from the “Gravity Dataset” produced by the Centre d’Études Prospectives et d’Informations Internationales. Distance measures are computed as the (geodesic) distance between national capitals. We obtain very similar results if we use the population-weighted average distance among each country pair’s largest cities instead.

³⁰ The crop and exporter fixed effect d_i^k is missing whenever country i does not produce crop k . In such instances, we use the arithmetic average of the crop and exporter fixed effects over all crops produced by country i as our proxy for d_i^k .

in trade flows at the extensive margin imply that the welfare loss under full adjustment is now smaller than before, with slightly larger gains from trade adjustments. Here ΔW_w and ΔW_w^T are equal to -0.21 percent and -0.23 percent, respectively, compared to -0.26 percent and -0.27 percent in Section VI. Similarly, the impact of climate change without production adjustment is now equal to -0.74 percent rather than -0.78 percent. Overall, though, we find that the treatment of zero trade flows has very little effect on our quantitative results.

C. Technological Differences

The GAEZ productivity estimates used in our baseline analysis allow for an unusually rich amount of heterogeneity, both before and after climate change. To harness this large amount of micro-level data, we have chosen to use a parsimonious representation of technology that otherwise severely restricts technological differences across countries. In practice, one may imagine that beyond land and labor, developed and developing countries use very different amounts of fertilizers and other complementary inputs. Similarly, one may imagine that the extent of productivity heterogeneity not measured in the GAEZ data—and parameterized by θ in our model—may be very different as well in developed and developing countries.³¹

To explore the quantitative importance of the first observation, we take advantage of the fact that, as mentioned in Section IV.A, GAEZ productivity estimates are also available under different scenarios regarding the application of complementary inputs. In our baseline results, we have used the scenario referred to as “high inputs” with “rain-fed” water supply for all countries in our data set. We now instead assume that productivity estimates for developing countries before and after climate change, A_i^h and $(A_i^h)'$, are given by the scenario referred to as “low inputs” with “rain-fed” water supply. Productivity estimates for developed countries as well as other estimated parameters are unchanged.³² The results of our counterfactual exercises under this alternative assumption are reported in row 3 of table 5. Compared to our baseline results, welfare losses caused by climate change are consistently smaller in all counterfactual scenarios. This is consistent with the observation that developing countries tend to be more negatively affected by climate change and that

³¹ Indeed, a test of the null hypothesis that the parameter θ is equal across countries rejects the null at the 5 percent level.

³² Developing countries correspond to low- and lower-middle-income countries according to the World Bank in 2009, whereas developed countries correspond to upper-middle- and high-income countries. There are 21 developing countries and 29 developed countries in our data set.

switching from high inputs to low inputs tends to reduce agricultural output in these countries. Finally, note that although welfare losses are smaller in all scenarios, relative differences across scenarios are similar to those presented in Section VI. As before, production adjustments appear to matter much more than trade adjustments.

To explore the quantitative importance of the second observation, we again break down countries in our data set into two groups, developed and developing, and estimate a separate shape parameter for the distribution of within-field, within-crop productivity, $\theta_{\text{Developing}}$ and $\theta_{\text{Developed}}$, using the exact same procedure as in Section V.B for each subsample. We find $\theta_{\text{Developing}} = 2.71$ and $\theta_{\text{Developed}} = 2.29$, suggesting more unobserved heterogeneity on average among developed countries. The implications of this alternative source of heterogeneity across countries for our counterfactual results are reported in row 4 of table 5. Since both $\theta_{\text{Developing}}$ and $\theta_{\text{Developed}}$ remain close to our baseline estimate of $\theta = 2.46$, there is very little difference between these new results and those presented in Section VI. Allowing for different degrees of heterogeneity between developed and developing countries leads to changes in ΔW_W , ΔW_W^P , and ΔW_W^T smaller than 0.02 percent under our three counterfactual scenarios.

D. Intranational Trade Costs

Our baseline model does not impose any restriction on the structure of iceberg trade costs between countries. They are isomorphic to demand shifters that can be simply recovered as residuals in the demand equation (11). In contrast, our baseline model rules out trade costs within countries. Crops are implicitly assumed to be shipped freely from the fields in which they are produced to the cities in which they are consumed or the ports from which they are exported. The goal of this final extension is to explore the robustness of our results to the introduction of intranational trade costs.

Compared to the baseline model presented in Section III, we assume that in any given country, crops must be shipped from the field in which they are produced to the national capital in which they will be either consumed locally or traded internationally. Crucially, intranational shipments are now also subject to iceberg trade costs, τ_i^f , which depend on the distance between field f and the national capital of country i . This new model is therefore identical to the baseline model but for the fact that the relevant, trade cost-adjusted productivity of field f for crop k in country i is now equal to A_i^k / τ_i^f .

To revisit the results of Section VI, we therefore need measures only of τ_i^f . The simplest way to obtain such measures is to use the OLS estimate, χ_{dist} , of the effect of log distance on the composite of demand shifters and trade costs, $\varepsilon_{ij}^k \equiv \ln[\beta_{ij}^k (\tau_{ij}^k)^{1-\sigma}]$, computed in Section VII.B. In our

model, the parameter χ_{dist} measures the elasticity of trade flows, X_{ij}^k , with respect to distance. Our OLS estimate implies an elasticity of trade flows with respect to physical distance equal to $\chi_{\text{dist}} = -0.83$, very much in line with existing estimates in the literature (see Disdier and Head 2008).

Under the assumption that international and intranational trade costs vary with distance in the same way, we can then approximate (log) intranational trade costs by

$$\ln \tau_i^f = \chi_{\text{dist}} \ln(\text{dist}_i^f) / (1 - \sigma), \quad (23)$$

where dist_i^f denotes the physical distance between field f and the national capital of country i and $\sigma = 5.40$ corresponds to the IV estimate obtained in Section V.A.³³ Row 5 of table 5 reports the results of our counterfactual exercises in the presence of intranational trade costs as measured by equation (23). These numbers again point toward similar effects of climate change on world GDP under full adjustment, -0.27 percent, with changes in production patterns playing a more important role than changes in trade patterns. In a counterfactual scenario without production adjustments, the welfare loss associated with climate change would again be -0.78 percent of world GDP compared to -0.28 percent in a counterfactual scenario without trade adjustments. The robustness of our results to the introduction of intranational trade costs suggests that there is little correlation in practice between the extent of productivity shocks across fields and their distance to the capital city.

VIII. 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 depends not only 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 data set that con-

³³ If international trade costs take the general form $\ln \tau_{ij}^k = \tau_0 + \tau \ln(\text{dist}_{ij})$, where τ_0 is a constant, then intranational trade costs should be approximated by $\ln \tau_i^f = \tau_0 + \tau \ln(\text{dist}_i^f)$. Since the OLS estimation of eq. (22) does not allow us to separately identify τ_0 , we have omitted this constant term from eq. (23) above. It should be clear, however, that none of our counterfactual results are affected by this normalization. Given our assumption of log preferences, a change in τ_0 would act as a Hicks-neutral TFP shock in agriculture, which would affect crop prices and welfare levels uniformly but would leave welfare changes between equilibria unchanged.

tains agronomists' estimates about the productivity—both before and after climate change—of each crop for each of 1.7 million fields 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 agronomists' calculations is the climatic data that enter their models, which are drawn from leading climatological models of climate change.

Using a general equilibrium model of trade among these 1.7 million fields, we find that the impact of climate change on these agricultural markets would amount to a 0.26 percent reduction in global GDP when trade and production patterns are allowed to adjust. While trade adjustments play little role in explaining the magnitude of these effects, our analysis suggests that production adjustments caused by the evolution of comparative advantage substantially mitigate the ill effects of climate change.

References

- Babiker, M. H. 2005. "Climate Change Policy, Market Structure, and Carbon Leakage." *J. Internat. Econ.* 65:421–45.
- Broda, C., and D. Weinstein. 2006. "Globalization and the Gains from Variety." *Q.J.E.* 121 (2): 541–85.
- Costinot, A., and D. Donaldson. 2011. "How Large Are the Gains from Economic Integration? Theory and Evidence from U.S. Agriculture, 1880–2002." Manuscript, Massachusetts Inst. Tech.
- . 2012. "Ricardo's Theory of Comparative Advantage: Old Idea, New Evidence." *A.E.R. Papers and Proc.* 102 (3): 453–58.
- Cristea, A., D. Hummels, L. Puzello, and M. Avetysyan. 2013. "The Contribution of International Transport to Global Greenhouse Gas Emissions." *J. Environmental Econ. and Management* 65 (1): 153–73.
- Disdier, A.-C., and K. Head. 2008. "The Puzzling Persistence of the Distance Effect on Bilateral Trade." *Rev. Econ. and Statis.* 90 (1): 37–48.
- Eaton, J., and S. Kortum. 2002. "Technology, Geography and Trade." *Econometrica* 70 (5): 1741–79.
- Elliott, J., I. Foster, S. Kortum, T. Munson, F. P. Cervantes, and D. Weisbach. 2010. "Trade and Carbon Taxes." *A.E.R. Papers and Proc.* 100 (2): 465–69.
- Felder, S., and T. Rutherford. 1993. "Unilateral CO₂ Reductions and Carbon Leakage: The Consequences of Trade in Oil and Basic Materials." *J. Environmental Econ. and Management* 25:162–76.
- Hertel, T., and T. Randhir. 2000. "Trade Liberalization as a Vehicle for Adapting to Global Warming." *Agriculture and Resource Econ. Rev.* 29 (2): 1–14.
- IPCC (Intergovernmental Panel on Climate Change). 2007. *Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007*. Cambridge: Cambridge Univ. Press.
- Nakicenovic, N., and R. Swart, eds. 2000. *Special Report on Emissions Scenarios: A Special Report of Working Group III of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge Univ. Press.
- Reilly, J., and N. Hohmann. 1993. "Climate Change and Agriculture: The Role of International Trade." *A.E.R. Papers and Proc.* 83 (2): 306–12.

- Rosenzweig, C., and M. Parry. 1994. "Potential Impact of Climate Change on World Food Supply." *Nature* 367:133–38.
- Shapiro, J. 2013. "Trade, CO₂, and the Environment." Manuscript, Yale Univ.
- Sotelo, S. 2013. "Trade Frictions and Agricultural Productivity: Theory and Evidence from Peru." Manuscript, Univ. Chicago.
- Tsigas, M., G. Friswold, and B. Kuhn. 1997. *Global Climate Change and Agriculture, Global Trade Analysis: Modeling and Applications*. Cambridge: Cambridge Univ. Press.