Climate Change, The Food Problem, and the Challenge of Adaptation through Sectoral Reallocation

Ishan Nath *

Princeton University

Current Draft: July 23, 2021

Abstract

This paper combines local temperature treatment effects with a quantitative macroeconomic model to assess the potential for global reallocation between agricultural and non-agricultural production to reduce the costs of climate change. First, I use firm-level panel data from a wide range of countries to show that extreme heat reduces productivity less in manufacturing and services than in agriculture, implying that hot countries could achieve large potential gains by adapting to climate change through increasing imports of food and shifting labor toward manufacturing. To investigate the likelihood that such gains will be realized, I embed the empirically estimated productivity effects in a model of sectoral specialization and trade covering 158 countries. Simulations suggest that climate change does little to alter the geography of agricultural production, however, as high trade barriers in developing countries temper the influence of shifting comparative advantage. Instead, climate change accentuates the existing pattern, known as "the food problem," in which poor countries specialize heavily in relatively low productivity agricultural production to meet subsistence consumer needs. The productivity effects of climate change reduce welfare by 6-10% for the poorest quartile of the world with existing trade barriers, but by 70% less in an alternative policy counterfactual that moves low-income countries to OECD levels of trade openness.

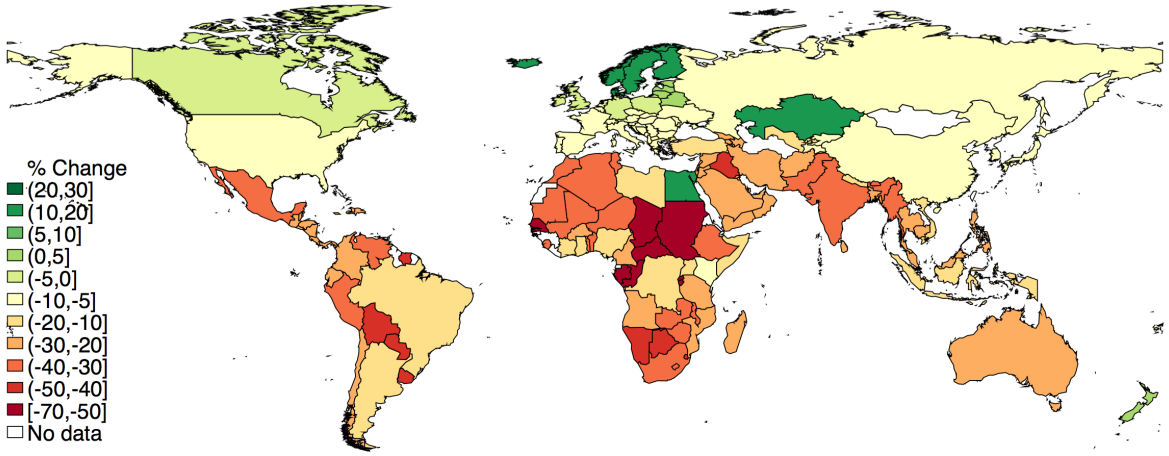
*I thank my advisers, Michael Greenstone, Chang-Tai Hsieh, and Pete Klenow for their guidance and mentorship, and for sharing access to data. I also appreciate helpful comments from Martí Mestieri, Hannes Malmberg, Levi Crews, Will Rafey, Allan Hsiao, Tamma Carleton, Michael Dinerstein, Alessandra Voena, Manasi Deshpande, Lauren Bergquist, Steve Cicala, Ryan Kellogg, Eyal Frank, and seminar participants at the University of Chicago, the Columbia University Young Economists Symposium, the NBER Agriculture and Trade Conference, the UCLA Climate Adaptation Symposium, and the NBER Risks in Agricultural Supply Chains Conference. I thank Henry Zhang, Nick Tsivanidis, and Cian Ruane for help processing and cleaning the datasets, Steve Mohr and Theodor Kulczycki for guidance on coding and computing, and the Chicago Booth Initiative on Global Markets (IGM), and Jennifer Williams and Peggy Eppink in particular, for providing access to the Amadeus dataset from Bureau van Dijk. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. All errors and omissions are my own.

Contact information - inath33@gmail.com. Note that Princeton affiliation officially begins 9/1/2021.

1 Introduction

Existing research suggests that climate change will cause large and heterogeneous changes in agricultural productivity across the world during the 21st century. Figure 1 displays a representative example from this literature:[1]

Figure 1: Cline (2007) Projected Impact of Climate Change on Agricultural Productivity, 2080-2099



Notes: Figure shows the projected change in revenue per acre from producing grains, vegetables, fruits, and livestock according to analysis by Cline (2007).

The pattern in Figure 1 suggests large potential gains from shifting the geography of agricultural production. While climate change projects to reduce agricultural productivity by 30-60% in hot, largely agrarian, regions such as Sub-Saharan Africa and South Asia, colder locations such as Canada and northern Europe can expect neutral or positive effects. These forecasts imply that sectoral reallocation could substantially curtail climate change damages if agricultural activity moves toward temperate climates while tropical regions shift production to other sectors. Conversely, if the general equilibrium forces that cause equatorial regions to specialize heavily in agriculture at present persist, or strengthen, the gains from this channel of adaptation will be limited.

Two key elements of sectoral allocation complicate the idea that the changes in Figure 1 will push agriculture away from the equator. First, these estimates show changes in the absolute advantage of agriculture, whereas comparative advantage across sectors drives international trade. Ricardian models of trade will only predict that Canada will export more food and India will import more food if the *relative* productivity of agriculture rises in Canada and falls in India. Given

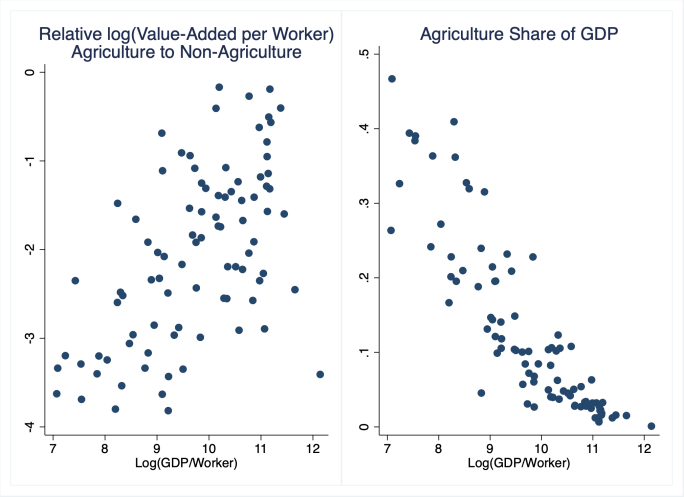
[1] I explain the methods used in Cline (2007) more in Section 7.1. The findings are broadly consistent with a large body of economics research on the impacts of climate change in agriculture, which includes Mendelsohn, Nordhaus and Shaw (1994), Deschenes and Greenstone (2007), Schlenker and Roberts (2009), and Schlenker and Lobell (2010), among many others.

existing evidence that temperature also affects non-agricultural productivity, the change in comparative advantage is not immediately clear.[2]

Second, comparative advantage does not exclusively, or even primarily, determine sectoral specialization. Figure 2 shows that poor countries have high agricultural GDP shares despite much lower relative value-added per worker in agriculture compared to non-agriculture. Lagakos and Waugh (2013) calculate that, adjusting for prices, the gap in aggregate output per worker between the 90th to 10th percentile of the world's income distribution is 45 to 1 in agriculture, but just 4 to 1 in non-agriculture. Yet agriculture's share of employment averages 65% in 10th percentile countries and only 3% in 90th percentile countries. Projecting the effects of climate change on sectoral reallocation requires accounting for the forces that drive poor countries to specialize in agriculture despite low absolute and relative productivity, a fact which the literature on the general equilibrium effects of climate change has not yet confronted.

This paper addresses these challenges by integrating local temperature treatment effect estimates and a quantitative macroeconomic model to evaluate the impact of climate change on sectoral reallocation and global welfare. First, to project changes in agricultural comparative advantage, I provide the first global micro estimates of the impact of climate change on productivity in manufacturing and services using nationally representative firm-level panel data from 17 countries that cover over half the world's population and represent nearly the full range of temperatures and income levels. Using methods developed by Carleton et al. (2020), I use the data to estimate plausibly causal treatment effects of extreme temperatures on output-per-worker, and account for firm-level adaptation by allowing these treat-

Figure 2: Comparative Advantage and Specialization in Agriculture



Notes: Figure shows data from Tombe (2015) that adjusts for prices for the global cross-section in 2005. Poor countries specialize heavily in agriculture despite low productivity relative to other sectors.

[2]This evidence includes work by Zhang, Deschenes, Meng and Zhang (2018) and Somanathan, Somanathan, Sudarshan, Tewari et al. (2015).

ment effects to vary with income and expectations of temperature.

Second, I construct a global open economy model of structural transformation that explains the existing distribution of sectoral specialization as a function of sector-level productivities. The model incorporates two key features of consumer preferences - nonhomothetic preferences and low substitutability across sectors - that explain the high agricultural share of consumption in poor countries with high relative prices for food. Gollin, Parente and Rogerson (2007) refer to the macro-development effects of these subsistence requirements as "the food problem," which drives developing countries to specialize in a relatively low-productivity sector because people need food to survive. In principle, imports could meet consumer needs for food, but in practice this channel is weak in developing countries. The average person in the poorest quartile of the world consumes 91% domestically produced food, compared with 45% in the richest quartile. In these relatively closed economies, high agricultural production and labor shares follow from the high consumption shares necessary for people with low incomes to meet subsistence requirements for food.

Thus, the model shows that two competing effects govern the response of sectoral specialization to climate change. If the trade effect dominates, then countries can dampen the effect of falling agricultural productivity by shifting production to other sectors; exporting more manufactured goods and importing more food. To the extent that climate change exacerbates 'the food problem' by reducing agricultural productivity, however, the general equilibrium response could pull labor into the sector suffering large declines in productivity, limiting the welfare gains from reallocation.[3] To quantify the relative strengths of these mechanisms, I estimate the model to match data on income levels, trade flows, and sectoral GDP shares for 158 countries covering over 99.9% of global GDP. I then embed the empirically estimated projected impacts of climate change on productivity in agriculture, manufacturing, and services into the model, and conduct counterfactual simulations that calculate the effects of climate change on sectoral specialization, trade, prices, GDP, and welfare.

The paper has four key findings. First, I find that extreme temperatures have substantial effects on non-agricultural productivity, but with strong evidence of adaptation in rich countries and to temperatures with which agents are accustomed. In poor countries with moderate climates, an extreme day with daily maximum temperature of 40°C or -5°C reduces annual output-per-worker by up to 0.4%, approximately the equivalent of one full working day.[4] Effects are about half

[3]While rural-urban migration within countries plays a key implicit role in my analysis of sectoral reallocation, I do not incorporate international migration in the model. The results imply that climate change raises the income gap between the richest and poorest quartile of the world from a factor of 38 to a factor of 41, implying an increase in the shadow value of migration but also fewer resources with which to pay the costs of migration for people in poor countries. I leave it to future work to estimate the response to these changes.

[4]I find similar effects for manufacturing and services firms, though I lack data coverage for

as large in middle-income countries, and smaller still in those places that experience given extremes more frequently. The effects of extreme days in rich countries are negligible, with some evidence of mild effects from unexpected extremes caused by hot days in cold places and cold days in hot places.[5] I combine these estimates of predicted temperature sensitivity with global climate model predictions of future temperatures to project the country-level effects of climate change on manufacturing and services productivity. The effects of climate change on non-agricultural productivity are non-trivial (5-15%) in some poor countries, but are generally small relative to estimated effects in agriculture. Thus, the change in the global comparative advantage of agriculture is qualitatively similar to its change in absolute advantage.

Second, model simulations suggest that 'the food problem' drives sectoral reallocation more than comparative advantage. Climate change raises the agriculture share of GDP by 2.8 percentage points in the poorest quartile of the world even as agricultural productivity decreases sharply. While net exports of agriculture rise in colder countries and net imports generally rise in hot regions, most developing countries are not sufficiently open to trade for this adjustment to play a primary role in sectoral specialization. As a consequence of economic activity in many countries shifting toward the sector where productivity falls most, global GDP as measured by standard price indices declines by 12% more, and 52% more for the poorest quartile of the world, than in a naïve calculation with fixed sectoral shares.

Third, the productivity effects of climate change cause a net decline in global welfare, with losses heavily concentrated in poor countries. The equivalent variation willingness-to-pay to avoid these impacts is between 1.5-2.7% of global GDP, and between 6.2-10% of income for the average person in the poorest quartile of the world.[6]

Fourth, simulations suggest that trade does little to cushion the effects of climate change under current policy, but that lowering trade barriers could substantially reduce climate damages. The aggregate global willingness-to-pay to avoid climate change is only 7.4% lower at estimated levels of trade costs than in autarky, largely because those countries most susceptible to climate change also tend to be least open to trade. However, in an alternative scenario that assigns all countries to 90th percentile levels of trade openness, the global willingness-to-pay to avoid climate change falls by 31%.[7] Trade liberalization is particularly valuable for climate change adaptation in poor countries. The welfare costs of climate change for

services firms in poor countries where the effects of temperature are most detectable.

[5] I find evidence that firms in rich countries mitigate the effect of extreme temperatures on labor productivity through costly adaptation investments such as higher energy expenditures. I use a revealed preference method to infer the magnitude of these costs and account for them in calculating the welfare cost of climate change.

[6] Note that these estimates do not include other costs of climate change beyond the scope of this paper, such as hurricanes, sea-level rise, or the health effects of extreme temperatures.

[7] Note that this calculation nets out the gains from trade unrelated to climate change adaptation.

the poorest quartile of the world are 68% lower in the low trade cost counterfactual because greater tradability makes specialization more responsive to comparative advantage and eases the constraints of the food problem.

This paper relates to several literatures on climate change and macroeconomic development. The two most similar papers are Costinot, Donaldson and Smith (2016), who examine reallocation across crops but do not consider income effects or cross-sector reallocation, and Desmet and Rossi-Hansberg (2015), who primarily focus on the important role for international migration in climate change adaptation. The latter paper includes changes in the global distribution of sectoral specialization in the model, but does not attempt to incorporate realistic trade costs or the importance of 'the food problem' in the analysis. This paper is the first to consider the effects of climate change on structural transformation.

My empirical work on temperature and productivity builds on country-level estimates produced by Somanathan, Somanathan, Sudarshan, Tewari et al. (2015) and Zhang, Deschenes, Meng and Zhang (2018) in India and China. The model builds on the central insight of Matsuyama (1992) about structural transformation in an open-economy setting and incorporates features from several recent related papers including Tombe (2015), Uy, Yi and Zhang (2013), and Teignier (2018). I also use a nonhomothetic CES specification for consumer preferences from Comin, Lashkari and Mestieri (2015). My model counterfactuals relate to empirical work by Colmer (2018) and Liu, Shamdasani and Taraz (2020) that examines the local relationship between temperature and sectoral reallocation in Indian districts. This research finds that adverse weather shocks drive labor out of agriculture under some conditions, but raise the agriculture share of employment in remote locations with weak road networks, consistent with my model predictions about tradability and the food problem. Finally, some of my results about the role of trade and the spatial correlation of shocks relate to the work of Dingel, Meng and Hsiang (2019).

More broadly, this paper advances the frontier of methods in climate change economics by embedding credible empirical estimates into a general equilibrium model. To date, the climate impacts literature has followed two primary tracks: macroeconomic models in the spirit of Nordhaus (1992) and partial equilibrium econometric estimates such as Deschenes and Greenstone (2007). The former grouping facilitates conclusions about policy and welfare at a global scale, but generally adopts a stylized approach to quantification. In contrast, the latter branch establishes precise causal relationships between weather and specific outcome variables, but employs identification strategies that necessarily hold constant cross-sector and cross-national interactions relevant to future projections. This paper lays out a unifying approach that combines each of these strengths: quantitative theory with empirical estimates that map directly into model parameters, allowing the researcher to evaluate climate damages, adaptation, and policy counterfactuals in a framework that captures equilibrium behavior and welfare.

The paper is structured as follows. Sections 2, 3, and 4 describe the data, empirical strategy, and results for the estimation of the relationship between temper-

ature and non-agricultural productivity. Section 5 lays out the model. Section 6 explains the model estimation and describes the model's success in fitting the data. Section 7 contains the counterfactual model simulations. Section 8 provides additional country-level panel regression evidence on the impact of agriculture-biased productivity shocks on sectoral reallocation. Section 9 discusses implications for policy and Section 10 concludes.

2 Data

Firm Data

I assemble a panel of firm-level microdata with broad global coverage to estimate the relationship between temperature and productivity in manufacturing and services. Table 1 lists the countries and years included in the dataset as well as the data source for each country. The data combines surveys administered by national governments with data acquired from the Amadeus database maintained by Bureau van Dijk (BVD). BVD is a private company owned by Moody's Analytics that collects and distributes firm-level financial information from around the world. They collect data both by acquiring administrative data directly from national business registers and by conducting their own surveys.

I conduct analysis on countries for which I am able to obtain nationally representative panels. This includes government surveys from India, Colombia, Indonesia, China, and the United States, and Amadeus data from twelve European countries with mandatory filing requirements according to BVD documentation.[8] Bloom, Draca and Van Reenen (2016) report that the data in most of these European countries contains nearly the full population of public and private firms.[9] Gopinath, Kalemli-Özcan, Karabarbounis and Villegas-Sanchez (2017) also use data from Amadeus and Alfaro and Chen (2018) use data from Orbis, a related firm dataset produced by BVD.

The sample covers both manufacturing and services firms in developed and developing countries. While the government surveys cover only manufacturing firms, the BVD data covers the entire spectrum of 2-digit industries. I report results for the pooled sample of all firms, separately for manufacturing firms, and separately for services firms, though the latter subset lacks developing country coverage.[10] BVD also reports additional branch locations and subsidiary ownership for many firms. I drop all firms that list subsidiaries or additional branches so that reported

[8]Importantly, the online version of the Amadeus database does not maintain accurate historical records. Thus, I download the data directly from the 2005, 2010, and 2015 vintages (CDs). Each Amadeus vintage contains 10 years of historical data for each firm. I match firms across years using BVD's unique firm identification number, and drop a small subset of observations with inconsistent data across vintages for the same firm-year.

[9]Denmark, Ireland, and Portugal also have mandatory reporting requirements, but were unavailable to me due to data licensing restrictions and missing or outdated geographic identifiers.

[10]I drop firms marked mining, construction, utilities, and agriculture, though results are very similar when including these firms in the pooled sample.

Table 1: Global Firm-Level Panel Microdata

Country	Data Source	Dataset	Years
Austria	Bureau Van Dijk	Amadeus	1995-2014
Belgium	Bureau Van Dijk	Amadeus	1995-2014
China	National Bureau of Statistics	Chinese Industrial Survey	2003-2012
Colombia	National Administrative Department of Statistics (DANE)	Annual Manufacturing Survey	1977-1991
Finland	Bureau Van Dijk	Amadeus	1995-2014
France	Bureau Van Dijk	Amadeus	1995-2014
Germany	Bureau Van Dijk	Amadeus	1995-2014
Greece	Bureau Van Dijk	Amadeus	1995-2014
India	Central Statistical Office	Annual Survey of Industries	1985-2007
Indonesia	Badan Pusat Statistik	Annual Manufacturing Survey	1975-1995
Italy	Bureau Van Dijk	Amadeus	1995-2014
Norway	Bureau Van Dijk	Amadeus	1995-2014
Spain	Bureau Van Dijk	Amadeus	1995-2014
Sweden	Bureau Van Dijk	Amadeus	1995-2014
Switzerland	Bureau Van Dijk	Amadeus	1995-2014
United Kingdom	Bureau Van Dijk	Amadeus	1995-2014
United States	Census Bureau	Annual Survey of Manufacturers, Census of Manufacturers	1976-2014

Notes: Data includes nationally representative samples of firm-level data on revenue and number of employees, with varying coverage of capital stock (tangible fixed assets). Amadeus data includes both manufacturing and services firms.

firm output aligns as closely as possible to the measure of temperature exposure at the main location. I also drop firms containing fewer than three observations and those with missing data for revenue or number of employees.

In total, the sample includes 17 countries that cover 59.4% of the world's manufacturing output and 51.1% of the global population.[11] The dataset also spans virtually the full range of climate and income levels in the global cross-section. According to the Penn World Tables, PPP-adjusted GDP per capita in the sample ranges from $1,137 in India in 1985 to $64,274 in Norway in 2014, which covers the 3rd to the 99th percentile of the global population in 2014. Similarly, country-level average daily maximum temperature in the sample ranges from 8.5 C°in Norway to 31.5 C°in India, covering the 1st to the 90th percentile of global population-weighted long-run temperature. Thus, to the extent that income and average tem-

[11] I cannot include the United States in the main pooled specification because I can only access the data at a secure government facility. I also exclude the data from China from the main specification for data quality reasons explained in Section 4.

perature predict adaptation to extreme temperatures, the data contains information about the full range of heterogeneity in the global temperature-productivity relationship.

Climate Data

I use temperature data from Version 3 of the Global Meteorological Forcing Dataset (GMFD) produced at Princeton University. The data covers the entire world at a 0.25°by 0.25°grid for the years 1948-2016. GMFD is a reanalysis dataset that reconstructs historical temperature using a combination of observational data and local climate models. Following Graff Zivin and Neidell (2014), I use daily maximum temperature as the variable of interest to best approximate the temperature people experience during working hours.

I match firm and climate data at the county level. The government surveys provide county location for each firm directly. The BVD data provides city name and zip code, which I match to the county level using GeoPostcodes, a global geocoding dataset provided by GeoData Limited.[12] I apply nonlinear transformations to the GMFD temperature variable at the pixel level, and then average across pixels to the county level weighting by population.[13]

Other Data

I use purchasing power parity adjusted GDP per capita data from the Penn World Tables as a measure of the income level of each country-year in the sample.

3 Empirical Strategy

I start by laying out the following three objectives for the empirical results to allow me to quantify the effects of climate change on sectoral reallocation and aggregate productivity. First, I need to estimate the causal effect of temperature on productivity in manufacturing and services. Second, I need to estimate the heterogeneity in that relationship such that I can predict the response to temperature for every country in the world. The model counterfactuals in Section 7 require an estimate of the response of manufacturing productivity to temperature in Algeria without having data from Algeria. Third, the estimates should incorporate the benefits and costs of adaptation. Future projections should reflect that the effects of a given temperature realization will likely change as countries grow richer, firms improve technology, and agents adjust expectations to the shifting distribution of temperatures. To quantify the effects of climate change in Section 7, I need to make projections not just for Algeria today, but for future Algerian firms experiencing climate change in 2080.

[12] GeoData Limited estimates that their latitude and longitude coordinates for the center of each zip code are precise to within 100 meters. I independently verify a subset of observations in each country to ensure accuracy. I also hand-code a small number (under 1%) of unmerged observations using city name, and drop those unmerged observations for which the city name is non-unique within a country.

[13] For some countries, the administrative unit to which I aggregate is more comparable to a town than a county.

3.1 Conceptual Framework

To motivate the estimation strategy I start with a version of the production function from Burnside, Eichenbaum and Rebelo (1993) with variable labor effort:

$$Y = AK^{\alpha}(e * L)^{1-\alpha} \text{ with } 0 \leq e \leq 1 \tag{1}$$

The parameter e governs effective units of labor input. Intuitively, temperature could affect e through several channels. Extreme temperatures could cause illness or physical fatigue, impair cognitive function, or increase the disutility of labor such that workers reduce effort or minutes spent working.[14]

Rearranging the production function in terms of output per worker and taking logs gives:

$$ln\left(\frac{Y}{L}\right) = ln(e) + \left(\frac{1}{1-\alpha}\right)ln(A) + \left(\frac{\alpha}{1-\alpha}\right)ln\left(\frac{K}{Y}\right) \tag{2}$$

Equation 2 provides the basis for using output per worker as the dependent variable in the main specification. The change in output per worker equals the change in e when the firm's technology and capital-to-output ratio stay constant. To gain further insight into the firm's optimal response to climate conditions, I model worker effort as a function of exposure to extreme heat (cooling degree days, or CDD), extreme cold (heating degree days, or HDD), and adaptation investments b_h and b_c:

$$\begin{gathered} e^* = 1 - CDD * g_h(b_h) - HDD * g_c(b_c) \\ g \geq 0, g' < 0, g'' > 0 \end{gathered} \tag{3}$$

In this framework, the firm has access to separate technologies that mitigate the impact of extreme heat and extreme cold on worker effort with diminishing returns in each.[15] The first order conditions for a profit-maximizing firm yield the following expression for the firm's optimal investment in hot weather adaptation b_h:

$$-g_h'(b_h) = \frac{c_h * e}{p * MPL * L * CDD} \tag{4}$$

[14]The health effects of extreme temperatures have been widely documented, including in Deschênes and Greenstone (2011). Several laboratory experiments, including Seppanen, Fisk and Lei (2006) find evidence of reduced worker cognitive functioning. Graff Zivin and Neidell (2014) use time-use surveys to show that people work fewer minutes per day in the presence of extreme temperatures.

[15]Zhang, Deschenes, Meng and Zhang (2018) mention that capital equipment could also perform poorly in extreme temperature conditions. If so, augmenting the production function with variable effective capital utilization, u, as in Burnside and Eichenbaum (1996), would capture this effect. In that case, the interpretation in Equation 2 would be that the reduction in $\frac{Y}{L}$ was attributable to a combination of declines in e and u.

Since g is convex in b_h, Equation 4 predicts that firm adaptation investments will be increasing in the firm's exposure to extreme heat (CDD), the marginal product of labor, the firm's labor input, and the output price, and decreasing in the cost of the adaptive technology, c_h, and the level of worker effort.[16] Thus, the firm's optimal condition predicts that worker effort will be less sensitive to temperature at more productive firms with more expected exposure to extreme temperatures, but that this reduced sensitivity comes at a cost.

To capture this heterogeneity, the empirical strategy focuses on modeling output per worker, and consequently e, as a function of temperature realizations, the level of productivity, and expectations over the distribution of temperature. By measuring the effects of climate change on e, I can use the estimates to project the change in the sector-by-country aggregate productivity parameters that govern average output per worker in the model introduced in Section 5.

3.2 Causal Effect of Temperature

Following the framework outlined in Deryugina and Hsiang (2014), I start by noting that workers experience daily realizations of weather. San Francisco and Washington D.C. have similar annual temperatures, but very different exposure to extremes. To capture this logic, I treat daily output as a function of temperature on day d, $Y_d = f(T_d)$. To aggregate to annual output, the level of my data, I sum daily outputs along with functions of daily temperature, $f(T_d)$, across all days experienced by firm i in year t:

$$Y_{it} = \sum_{d=1}^{365} Y_{id} = \sum_{d=1}^{365} f(T_{id}) = F(T)_{it} \tag{5}$$

Thus, I treat nonlinear transformations of daily temperature summed over the year as the primary independent variable of interest. Using annual data also has the important advantage of allowing for intertemporal substitution of labor. If workers produce less due to extreme temperatures on Tuesday but produce extra on Saturday instead, annual data captures the effects of temperature net of this reallocation.

For parsimony, my main specification uses a piecewise linear functional form for temperature, where output is allowed to vary linearly with daily maximum temperature above 30°C (CDD) and below 5°C (HDD):

$$f(T) = \begin{cases} \beta_1(5 - T_{max}) & \text{if } T_{max} < 5 \\ 0 & \text{if } 0 \leq T_{max} \leq 30 \\ \beta_2(T_{max} - 30) & \text{if } T_{max} > 30 \end{cases} \tag{6}$$

[16] Optimal adaptation investment is decreasing in the level of worker effort because there are concave returns to effort.

This formulation allows cold and hot temperatures to have separately estimated effects, β_1 and β_2, on labor productivity. I also conduct robustness checks with more flexible functional forms such as a polynomial of degree four and bins of daily maximum temperature.

Following other work in the climate impacts literature, I isolate the causal impact of temperature by exploiting interannual variation in weather. In line with the framework outlined in Section 3.1 the main specification models log output per worker at firm i in year t as a function of the vector of temperature effects, $\boldsymbol{\beta}$:

$$ln\left(\frac{Y_{it}}{L_{it}}\right) = \boldsymbol{\beta} F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \tag{7}$$

I control for permanent firm-specific features such as technology and management with firm fixed effects, δ_i, and for unobserved aggregate shocks such as technological progress and recessions with region (country or state) by year fixed effects, κ_{rt}. I cluster standard errors at the firm and county-by-year level to account for both serial and spatial correlation. Equation 7 allows for estimating the average treatment effect of temperature realizations, which fulfills part of the purpose of this section.

3.3 Heterogeneity and Adaptation

Following the strategy of Carleton et al. (2018), I allow for heterogeneity in the effect of temperature on output per worker by interacting the vector of temperature coefficients with income and long-run average temperature. This setup follows from the prediction in Section 3.1 that more productive firms in high-income countries and those that expect to experience extremes more frequently will be better adapted. I specify the interacted regression as follows:

$$\begin{aligned} ln\left(\frac{Y_{it}}{L_{it}}\right) = \boldsymbol{\beta} F(T)_{it} + \boldsymbol{\gamma_1} ln(GDPpc)_{rt} \times F(T)_{it} \\ + \boldsymbol{\gamma_2} TMEAN_i \times F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \end{aligned} \tag{8}$$

The interaction variables in Equation 8 are country-level annual GDP per capita and long-run average daily maximum temperature in the county containing firm i.[17]

Estimating Equation 8 allows me to predict the treatment effects of extreme cold, β_1, and extreme heat, β_2, as a function of two factors - income and average climate. While there are certainly other variables that affect temperature sensitivity, this parsimonious specification makes it feasible to predict the treatment effects in any country for which I have data on GDP per capita and average temperature.

[17] I use country-level income because reliable global data on subnational income is unavailable. Average temperature is calculated as a 40-year average in the county of firm i, which is the same geographic scale at which contemporaneous temperature is measured.

Given the existence of this data for the full range of countries in the global cross-section, as well as of readily available plausible future projections of temperature change and economic growth, this approach allows me to project the effects of temperature both across space and over time. In line with the goals for this section, the interacted model allows me to predict temperature sensitivity in Algeria today and in Algeria in 2080.

The coefficients on the interaction terms in Equation 8 are identified using cross-sectional, rather than panel, variation, but the identification assumption is also weaker. Estimating the main causal effect of temperature relies on the standard identification assumption - that the independent variable of interest is uncorrelated with omitted variables that affect output per worker conditional on the set of controls. For the interaction variables, however, I am interested in how income and climate *predict* temperature sensitivity, rather than in isolating their specific causal effect. Thus, the identification assumption is not that income and climate are uncorrelated with omitted variables affecting temperature sensitivity, but rather that this correlation remains constant across space and over time. Indeed, the aim is to use income and average climate as a proxy for the full set of underlying mechanisms that govern adaptation. The cross-sectional approach will produce valid predictions if the effects of temperature realizations on output per worker in parts of the world with income levels and average temperatures similar to India are similar to the effects measured in India.[18]

Allowing the treatment effects of temperature to vary with long-run conditions also bridges the gap between weather and climate. A primary concern with using weather variation to inform estimates of the costs of climate change is that the estimated treatment effects may change as agents adjust their expectations in the long-run. I address this concern by explicitly modeling the treatment effects as a function of those expectations, as represented by long-run average temperature. In my formulation, climate is a distribution of temperatures and weather is a draw from that distribution. By allowing the treatment effect of a draw to depend on the distribution, the estimates for the effects of each draw remain valid as the distribution shifts. Intuitively, a hot day in Toronto could be more harmful than a hot day in Texas because it is more unexpected, but becomes less so as Toronto warms and its agents adapt. I capture this effect by assigning Toronto the estimated treatment effect of Texas once it heats up to that level.

4 Empirical Results

4.1 Main Regression Results

Table 2 contains the main results from estimating Equations 7 and 8. Column 1 displays the treatment effect of extreme temperatures for the average unit of output

[18] Empirical estimation of adaptation in the climate impacts literature broadly relies heavily on cross-sectional variation because of the inherent difficulty in finding quasi-experimental variation in long-run conditions.

in the countries in the sample by weighting observations by country-level GDP and the inverse of each country dataset's sample size. While the estimated average treatment effects show that the effects of temperature are statistically different from zero, the magnitude of these coefficients is far too small to be economically meaningful. The estimates in Column 1 imply that a day with maximum temperature of either -5°C or 40°C would reduce annual output per worker by just 0.03% relative to a day in the moderate range of 5°C to 30°C.[19]

Table 2: Effects of Daily Temperature on Annual Revenue per Worker

	(1) Revenue/Worker	(2) Revenue/Worker	(3) Revenue	(4) Employment	(5) Revenue/Worker
TMax-30	-0.0000311	-0.00119	-0.00250	-0.00131	-0.00100
	(-2.29)	(-4.73)	(-6.80)	(-5.25)	(-4.03)
5-TMax	-0.0000315	-0.000956	-0.00180	-0.000842	-0.000452
	(-2.15)	(-2.15)	(-2.91)	(-1.92)	(-2.07)
(TMax-30) X log(GDPpc)		0.0000715	0.000178	0.000107	0.0000595
		(4.07)	(6.79)	(6.06)	(3.65)
(TMax-30) X $\overline{\text{TMax}}$		0.0000186	0.0000334	0.0000148	0.0000160
		(4.85)	(6.24)	(3.93)	(3.96)
(5-TMax) X log(GDPpc)		0.0000898	0.000167	0.0000769	0.0000416
		(2.14)	(2.85)	(1.85)	(2.02)
(5-TMax) X $\overline{\text{TMax}}$		-0.00000292	0.00000212	0.00000504	0.000000703
		(-1.54)	(0.93)	(2.85)	(0.59)
N	4125776	4125776	4125776	4125776	17938084
Manufacturing	X	X	X	X	X
Services					X
Firm FE	X	X	X	X	X
Country X Year FE	X	X	X	X	X
Inverse Sample Size Weights	X				
GDP Weights	X				
Countries Included	15	15	15	15	15

Notes: t-statistics in parentheses. Dependent variables all in logs. Standard errors are two-way clustered at the firm and county-by-year level. Column 1 shows the coefficients from estimating Equation 7 and Columns 2-5 show the results from Equation 8. Outcome variables come from the data sources listed in Table 1 and temperature data is from GMFD. Countries included are Austria, Belgium, Colombia, Finland, France, Germany, Greece, India, Indonesia, Italy, Norway, Spain, Sweden, Switzerland, and the United Kingdom. Section 4.3 shows results for the United States and Appendix C shows results for China.

[19]Note that the data contains information on revenues rather than physical output. The results can be interpreted as the effect on physical labor productivity under the assumption that firms are price-takers in the output market and the local shocks used to identify the effects do not affect

Column 2 in Table 2 shows substantial heterogeneity in the effects of temperature on annual output per worker. Note that I do not weight observations in the regressions in which I model heterogeneity explicitly because the aim is to understand how the treatment effect varies across the full range of the interaction variables. The unweighted regression with differential sample sizes in different places also effectively allows areas with more data, and consequently more precise estimates of the effect of temperature, to contribute more to estimating the interaction terms.

The main effects of temperature in the unweighted interacted regression in Column 2 are large, negative, and precisely estimated, though the magnitudes cannot be interpreted without considering the interaction terms. The coefficients on both interaction terms for log GDP per capita are large and positive, indicating that richer countries are insulated from the effects of both extreme heat and cold. Consistent with intuition about adaptation to long-run conditions, the coefficient on the interaction term for average long run temperature is positive for hot extremes and negative for cold extremes, indicating that places are less susceptible to temperatures which they experience more frequently. All four interaction coefficients on income and average temperature are consistent with the predictions from Equation 4 - more productive firms with more exposure to given extremes invest more in adaptation.

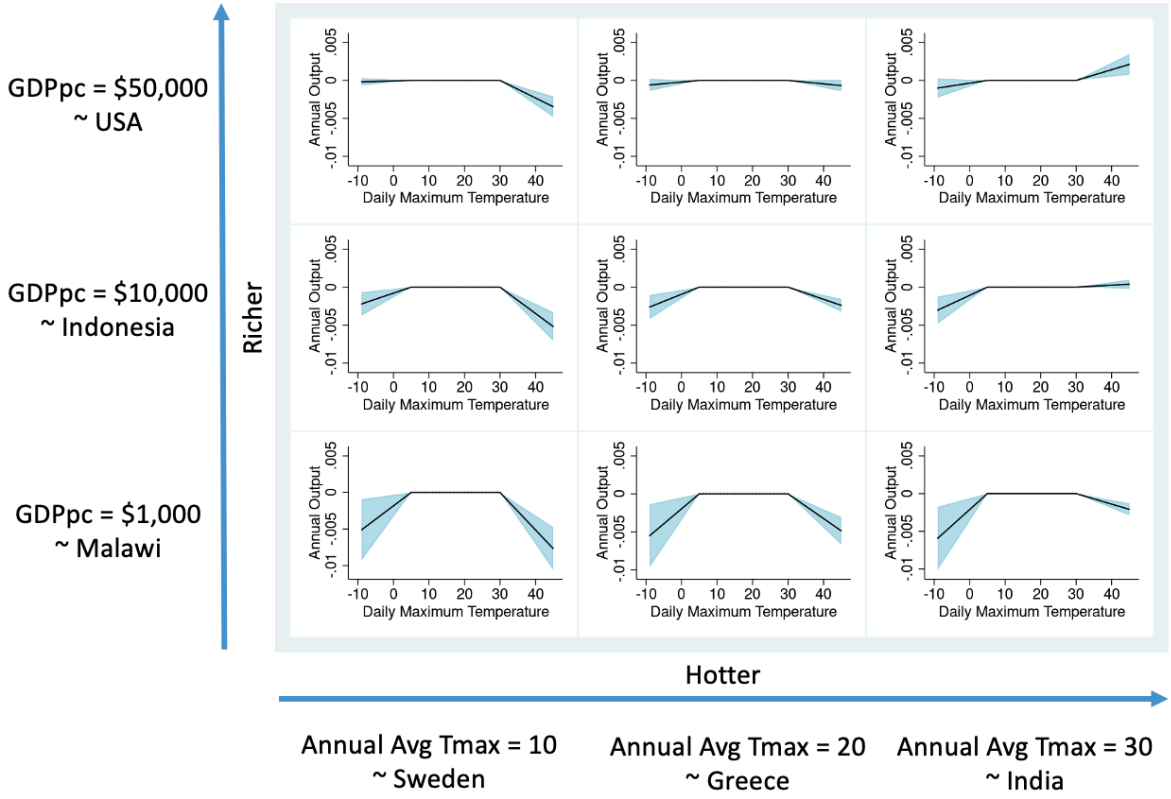
Figure 3 shows the predicted effects of temperature from Column 2 of Table 2 at points across the distribution of observed income and climate levels in the world. Consistent with the results of the GDP-weighted regression in Column 1, the graphs show that temperature has little effect on productivity in rich countries (top row), with some effects from hot days in cold, rich places (top left cell) and mild effects from cold days in hot, rich places (top right cell).

Conversely, extreme temperatures have very large effects on productivity in poor countries (bottom row). Experiencing one day at -5°C or 40°C in a poor country with moderate long-run temperatures (bottom middle cell) reduces annual output per worker by about 0.4%. In a working year consisting of 50 work weeks of 5 days each, this is equivalent to each worker reducing production on that day to zero with no compensating substitution to other days. These effects in poor countries imply potentially large productivity costs from climate change in hot parts of the world in the absence of adaptation. In parts of Sub-Saharan Africa, climate change projections imply an increase in extreme heat on the order of moving 100 days per year from 30°C to 40°C by 2080, which would suggest substantial declines in manufacturing productivity in poor countries.

Columns 3 and 4 of Table 2 separately estimate the effects of temperature on revenue and employment. The effects of both hot days and cold days on revenue are substantially larger than those on revenue per worker because firms adjust employment in response to extreme temperatures. As shown in Appendix Figures A-1

national or global product market prices.

Figure 3: Predicted Heterogeneous Response of Annual Manufacturing Revenue per Worker to Daily Maximum Temperature



Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 2 of Table 2. 95% confidence intervals are shown in blue.

and A-2, which again evaluate the predicted coefficients throughout the covariate space, these effects also primarily manifest only in poor countries. This finding is consistent with the firm's first order condition in the framework laid out in Section 3.1 - firms should be expected to reduce labor input in response to the fall in the marginal product of labor driven by a decline in e. However, it is perhaps surprising that firms in the sample do not face adjustment costs large enough to dissuade this adjustment in response to the short-run variation used to identify these effects.

Column 5 of Table 2 shows the effects of temperature on a pooled sample of manufacturing and services firms. The effects are very similar to the sample of only manufacturing firms in both magnitude and patterns of adaptation, with the exception of the finding that colder countries are less vulnerable to extremely cold temperatures. The sample size increases substantially in this specification because many of the firms in the data are services firms, though I do not have any services coverage in low-income countries.

4.2 Robustness

I conduct robustness checks with different ways to specify the functional forms of temperature. Appendix Figures A-3 and A-4 show the predicted effects from the main specification in Column 2 of Table 2 using bins and a polynomial of degree four in daily maximum temperature, respectively. The results are qualitatively very similar to the main specification.

I also show robustness to including more stringent state-by-year, rather than country-by-year, fixed effects. The results are very similar for specifications that use all the data (pooling manufacturing and services firms) with more flexible functional forms such as bins or a polynomial of degree four. These two specifications are shown in Appendix Figures A-6 and A-7. These results are sensitive to functional form, however. The more parsimonious functional forms with a single parameter each governing the response to cold days and hot days show muted effects, particularly in the specification with manufacturing firms only. This is consistent with the fact that considerably less variation in temperature realizations remains within states in a given year, so more data and flexible estimation is necessary to recover the underlying pattern.

Figure A-8 shows robustness to including controls for capital. While the standard errors for this specification are somewhat larger because I lack data on capital for approximately a quarter of the observations in the main specification, the pattern of predicted effects is very similar.

4.3 U.S. Results

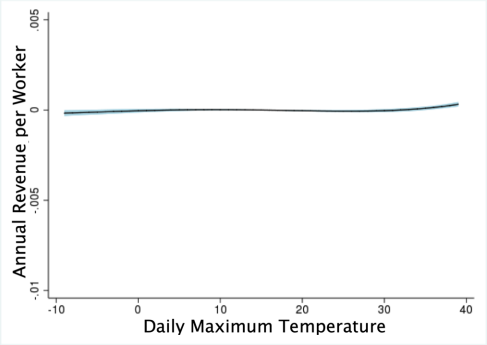
In this section, I use separate estimates of the effect of extreme temperatures on manufacturing in the United States to externally validate the results in Section 4.1.[20] Predictions using the global interacted regression suggest that temperature has a negligible effect on manufacturing revenue per worker in rich, temperate countries such as the U.S. (see the top middle cell of Figure 3). Figure 4 shows the corresponding estimate for the treatment effect of temperature in the United States using data from the U.S. Census Bureau.

Consistent with predictions from global data in Figure 3, I find a precisely estimated null effect of temperature on output-per-worker in the U.S.[21] The U.S. data also includes information on other inputs that I lack in the global sample, allowing me to directly observe some of the adaptation costs incurred by U.S. firms. Appendix Figure A-13 shows that the average U.S. plant increases expenditures on electricity and other fuels by several thousand dollars for each extremely hot and

[20] The results in Section 4.1 do not include data from the United States due to physical constraints on data access. Plant-level manufacturing data from the United States Census Bureau must be analyzed at restricted access Federal Statistical Research Data Centers (RDC).

[21] The result displayed in Figure 4 uses a polynomial of degree four in daily maximum temperature, but the null result is robust to choice of functional form. Appendix Table A-2 shows a range of specifications, all of which are consistent with a null effect on output and employment.

Figure 4: Estimated Response of U.S. Annual Manufacturing Revenue per Worker to Daily Maximum Temperature



Notes: Figure shows the response of annual revenue per worker to daily maximum temperature estimated using Equation 7 with a polynomial of degree four. 95% confidence interval is shown in blue. Outcome variable data comes from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD. Standard errors are two-way clustered at the firm and county-by-year level.

cold day, presumably for cooling and heating expenses.[22] These expenditures are small in the context of U.S. plant size, however, such that temperature still has a null effect on revenue total factor productivity, which accounts for expenditures on energy and materials, as shown in Figure A-12.

4.4 Projected Global Sensitivity to Extreme Temperatures

To connect the regression results from this section with the model presented in Section 5, I predict the effects of temperature in all 158 countries for which I estimate the model. Figure 5a shows the predicted effects of a day with maximum temperature of 40°C on annual manufacturing revenue per worker and Figure 5b shows the effect of a -5°C day. Consistent with intuition about adaptation and the results displayed in Figure 3, poor countries and those which experience given temperatures less frequently are more susceptible to extreme realizations.[23]

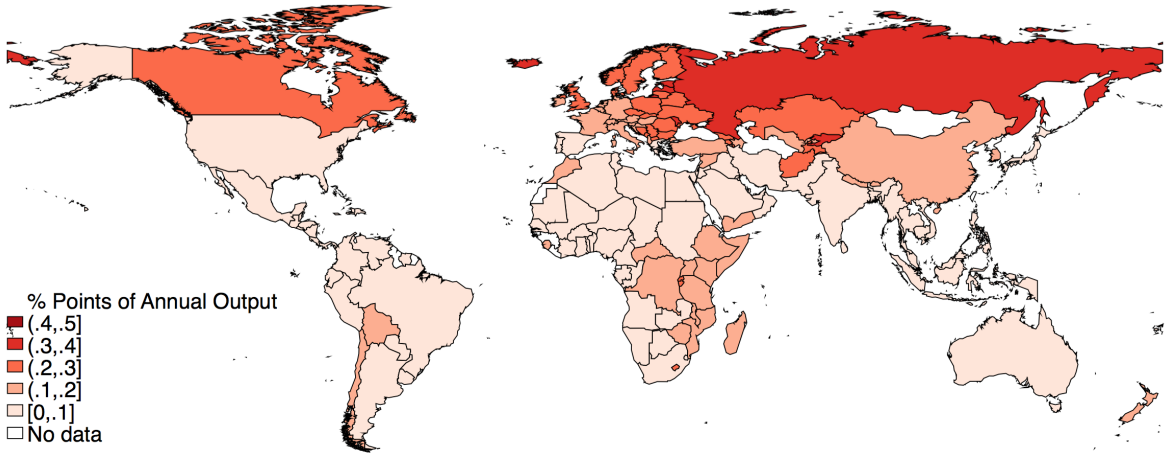
Projecting the impacts of climate change also requires accounting for adaptation by adjusting the temperature sensitivities shown in Figures 5a and 5b to projected changes in long-run average temperature. The firm's optimal adaptation decision in Equation 4 implies that firms will increase investment in protection from extreme heat as the climate warms. I account for the benefits of these investments by reevaluating predicted heat sensitivity at projected end-of-century

[22]Total energy expenditures are defined as the sum of electricity expenditures and the cost of other fuels. Full results for this outcome variable are shown in Appendix Table A-3.

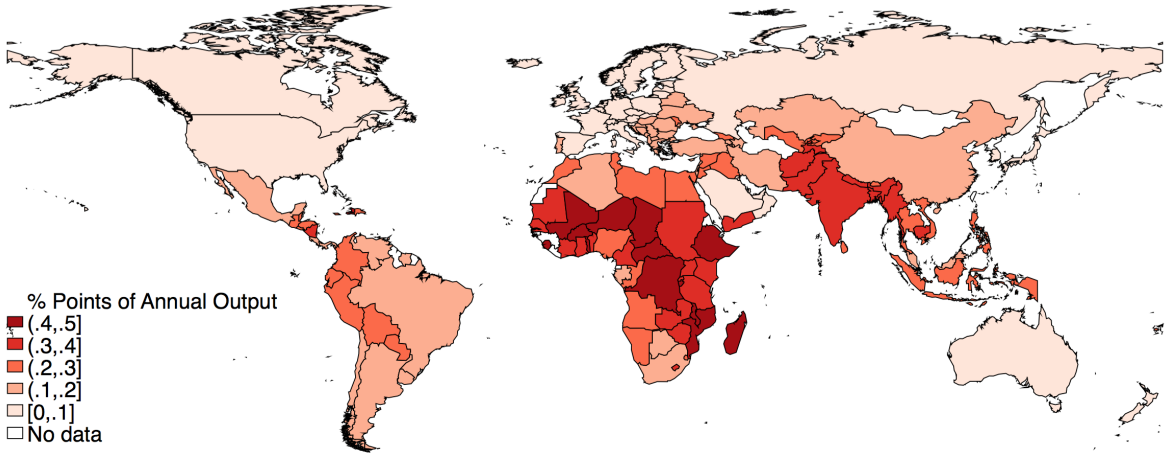
[23]Note that following Carleton et al. (2018), these predictions define full adaptation as productivity that is invariant to temperature, and thus do not allow the effect of extreme temperatures to go above zero. The effects of extreme temperatures are weakly negative in the range of incomes and climates in the sample used for estimation, and I maintain this pattern as incomes and temperatures go out of sample.

Figure 5: Predicted Effect of Extreme Temperatures on Annual Manufacturing Revenue per Worker

(a) 40°C Day



(b) -5°C Day



Notes: Maps show predicted annual percentage point loss in revenue per worker from a 40°C day and -5°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's GDP per capita and long-run average temperature.

temperatures in Appendix Figure A-17.[24] The results show noticeably muted effects when allowing for expectations to adjust to future temperatures. The mean global damage from a 40°C day is about 34% lower when evaluated at future temperatures (0.067% of annual revenues versus 0.1%).

The adaptation benefits of adjusting to extreme heat come at a cost. If it were

[24]End-of-century temperature projections are the 30-year average of annual average maximum temperature from the climate model predictions used in Section 7.1. In Section 7.6 I also allow for economic growth to make countries richer in the future, further reducing their temperature sensitivity.

costless to protect production from extreme heat, no firms would show effects of temperature on productivity. Instead, the results show that firms which experience given extremes infrequently find it optimal to invest less in adaptation, implying that the costs they would incur to achieve a marginal reduction in temperature sensitivity exceed the benefits. I leverage this intuition combined with the firm's first order conditions in Section 3.1 to infer a revealed preference measure of these adaptation costs following methods developed in Carleton et al. (2018). Appendix D covers the details of this calculation.

Quantifying the aggregate productivity consequences of climate change also requires projecting temperature sensitivity in services. I make projections for services using the pooled sample of manufacturing and services firms due to the lack of services data coverage in poor countries.[25] This choice follows from the estimated strong gradient of temperature sensitivity with respect to income but very similar coefficients between the manufacturing only and manufacturing/services specifications in Columns 2 and 5 of Table 2.[26] Intuitively, the results suggest that manufacturing firms in India are a better proxy for services firms in India than services firms in Germany would be. Appendix Figures A-20 and A-21 show predicted current global sensitivity to hot and cold days in services using results from the pooled regression. I follow the same procedure to account for future adaptation benefits and costs as in manufacturing.

Overall, the results in this section allow me to predict the sensitivity of non-agricultural firm output per worker to extreme temperatures in every country in the world in the present and future. I use these results to project the impact of climate change on global comparative advantage between agriculture and manufacturing in Section 7.1, and to simulate the corresponding changes in sectoral allocation and aggregate productivity.

5 Model

This section lays out a static general equilibrium model of global production, consumption, and trade in agriculture, manufacturing, and services to analyze how changes in sectoral productivity affect sectoral specialization, trade flows, aggregate productivity, and welfare. I show that the model makes ambiguous predictions about how reductions in agricultural productivity affect the employment share of agriculture, and that openness to trade is a key determinant of the reallocation that

[25] I show results for regressions using only services firms in Appendix Figures A-9, A-10, and A-11. The results for extreme heat with more flexible functional forms such as a fourth degree polynomial are qualitatively similar to those of the pooled manufacturing and services regression, but these specifications are sensitive to functional form. Furthermore, the predictions in poor countries are extrapolating far out of the sample, which only includes European firms in a narrow range of high income levels.

[26] A formal test shows that coefficients for manufacturing and services firms in the pooled regression have statistically indistinguishable responses to extreme heat and marginally significant evidence that services firms are less susceptible than manufacturing firms to extreme cold.

follows from asymmetric sectoral productivity shocks.

5.1 Model Ingredients

Consumption

Following the demand system specified in Comin, Lashkari and Mestieri (2015), consumers in each country gain utility from final goods in each of the three sectors - agriculture, manufacturing, and services - according to the following implicitly defined utility function:

$$\Omega_a^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_a}{\sigma}} C_{ak}^{\frac{\sigma-1}{\sigma}} + \Omega_m^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_m}{\sigma}} C_{mk}^{\frac{\sigma-1}{\sigma}} + \Omega_s^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_s}{\sigma}} C_{sk}^{\frac{\sigma-1}{\sigma}} = 1 \tag{9}$$

Here, $\{\epsilon_a, \epsilon_m, \epsilon_s\}$ are utility elasticities for each sector that allow for nonhomothetic preferences, $\{\Omega_a, \Omega_m, \Omega_s\}$ are sectoral taste parameters, and σ is the cross-sector elasticity of substitution. I choose this nonhomothetic CES preference specification because it can closely match the observed pattern of smooth structural transformation out of agriculture.[27]

Households consume their full wage, w, which varies at the level of country k. The aggregate budget constraint, summed across the country-level population, L_k, equates income to total expenditures across the three sectors:

$$P_{ak}C_{ak} + P_{mk}C_{mk} + P_{sk}C_{sk} = w_k L_k \tag{10}$$

Solving the consumer's problem gives the following expression for the expenditure share, ω_{jk}, in sector j in country k:

$$\omega_{jk} = \frac{P_{jk}C_{jk}}{w_k L_k} = \Omega_j \left(\frac{P_{jk}}{P_k}\right)^{1-\sigma} \left(\frac{w_k L_k}{P_k}\right)^{\epsilon_j - (1-\sigma)} \tag{11}$$

where the average cost index, $P_k = \frac{w_k L_k}{U_k}$, satisfies:

$$P_k = \left[\sum_{j\in\{a,m,s\}} \left(\Omega_j P_{jk}^{1-\sigma}\right)^{\frac{1-\sigma}{\epsilon_j}} \left(\omega_{jk}(w_k L_k)^{1-\sigma}\right)^{\frac{\epsilon_j-(1-\sigma)}{\epsilon_j}}\right]^{\frac{1}{1-\sigma}} \tag{12}$$

The household's expenditure function for achieving utility U_k at a given vector of sectoral prices is as follows:

$$E(U_k|P_{ak}, P_{mk}, P_{sk}) = \left[\sum_{j\in\{a,m,s\}} \Omega_j U_k^{\epsilon_j} P_{jk}^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{13}$$

[27] Nonhomothetic CES preferences improve model fit substantially compared to using generalized Stone-Geary preferences, another common specification used to represent nonhomotheticity in the structural transformation literature, particularly in middle income countries. I show robustness to using Stone-Geary preferences in Appendix F.1.

Production

The final good in sector j in country k is a CES composite of intermediate varieties indexed by i, where $\overline{y}_{ijk}$ represents the final goods producer's demand for variety i from the country from which it is sourced:[28]

$$Y_{jk} = \left(\int_0^1 \overline{y}_{ijk}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}} \tag{14}$$

Intermediate goods producers for each variety in each country receive a productivity draw, z_{ijk}, drawn from a Frechet distribution with sector-specific shape parameter θ_j and sector-country specific start value Z_{jk}. The production function for intermediate goods is linear in labor:

$$y_{ijk} = z_{ijk} * l_{ijk} \tag{15}$$

$$z_{ijk} \sim F_{jk} \text{ where } F_{jk}(z) = exp(-Z_{jk} z^{-\theta_j})$$

$$\text{and } Z_{jk} = f(\mu_{jk}, T_{jk}, E(T_{jk})) \tag{16}$$

The sector-country specific aggregate productivity parameters, Z_{jk}, connect the model to the empirical results in Section 4. In particular, I allow Z_{jk} to be a function of temperature realizations, T_{jk}, expectations over temperature, $E(T_{jk})$, and a vector, μ_{jk}, of country-sector specific features such as technology, institutions, and human capital. In making future projections in Section 7, climate change enters the model by perturbing the vector of Z_{jk} with empirically estimated productivity impacts that vary at the country-sector level.

Trade

The trade portion of the model follows Eaton and Kortum (2002). When selling to foreign countries, intermediate goods producers face an iceberg trade cost, τ_{jkn}, that varies by sector, j, exporter country, k, and importer country, n. So, intuitively, shipping food from Canada to Malawi incurs a different trade cost than shipping food from Malawi to Canada, and manufactured goods shipped between Canada and Malawi have two separate trade costs of their own. Services are nontradable.

Intermediate goods producers price at marginal cost. Since labor is the only input, the price of a domestically produced good in country k is given by $p_{ijk} = \frac{w_k}{z_{ijk}}$. When selling to foreign country n and incurring the cost of trade, the intermediate goods producer in country k prices as follows:

$$p_{ijk} = \frac{\tau_{jkn} w_k}{z_{ijk}} \tag{17}$$

This representation of trade incorporates Ricardian comparative advantage both within and across sectors. A producer's ability to sell competitively priced exports

[28] The final good is non-tradable and only used in consumption so that $C_{jk} = Y_{jk}$.

depends both on their productivity and on the domestic wage. Low productivity countries will have low wages in equilibrium, so their relatively productive producers will be able to export their products even if their absolute productivity is low. Thus, relative productivity between sectors is the key determinant of net imports and exports.

The final goods producer sources each variety from the lowest-priced producer. The sectoral final goods prices are given by the CES price index of all intermediate varieties used in that sector:

$$P_{jk} = \left(\int_0^1 p_{ijk}^{1-\eta} di \right)^{\frac{1}{1-\eta}} \tag{18}$$

The final goods producer's demand function for variety i as a function of these prices is given by:

$$\overline{y}_{ijk} = \left(\frac{p_{ijk}}{P_{jk}} \right)^{-\eta} Y_{jk} \tag{19}$$

Intuitively, the price of the final good in agriculture, P_{ak}, can be thought of as a price index for the complete basket of food items while the price of each individual variety, p_{iak}, is the price of one particular food. η is the elasticity of substitution between varieties.

Market-Clearing

The model has two market-clearing conditions. First, total income in country k is the sum of all domestic and foreign sales in all three sectors.

$$w_k L_k = \sum_{j \in \{a,m,s\}} \left(\pi_{jkk} P_{jk} C_{jk} + \sum_{n \neq k}^{N} \pi_{jkn} P_{jn} C_{jn} \right) \tag{20}$$

Here, π_{jkn} is the share of varieties from sector j consumed in country n that country k produces. So country k receives income both from its production share of domestic consumption in sector j, and from the share of consumption in every foreign country comprised of its exports. Since consumption equals income in each country, this condition also ensures that trade balances.

The second market-clearing condition concerns the labor market. The total labor force is allocated across the three sectors:

$$L_k = L_{ak} + L_{mk} + L_{sk} \tag{21}$$

In autarky, market-clearing requires that income equals expenditures in each sector, $P_{jk} C_{jk} = w_k L_{jk}$, which means that the employment share, l_{jk}, equals the expenditure share, ω_{jk}. In the open-economy case, the employment share equals the production share of revenues in each sector, incorporating net exports, which

yields the following equation:[29]

$$l_{jk} = \pi_{jkk}\omega_{jk} + \sum_{n \neq k} \pi_{jkn}\omega_{jn}\frac{w_n L_n}{w_k L_k} \tag{22}$$

This condition illustrates the importance of both domestic consumer preferences and international trade in determining the sectoral allocation of labor. Intuitively, Equation 22 says that if country k has agricultural consumption worth 30% of spending and agricultural net exports worth 10% of GDP, then 40% of its labor force will be in agriculture.

Equilibrium

For a given set of preference parameters, $\{L_k\}$, $\{Z_{jk}\}$, and $\{\tau_{jkn}\}$, equilibrium is given by a set of wages $\{w_k\}$, variety level prices $\{p_{ijk}\}$ and demand $\{\overline{y}_{ijk}\}$, final goods prices $\{P_{jk}\}$ and demand $\{C_{jk}\}$, average cost indices $\{P_k\}$, expenditure shares $\{\omega_{jk}\}$, and trade shares $\{\pi_{jkn}\}$ such that consumers and producers optimize (Equations 10, 11, 12, 14, 17, 18, and 19 hold) and trade balances (Equation 20 holds).

Willingness-To-Pay and Aggregate Productivity

I calculate the willingness-to-pay (WTP) to avoid climate change productivity impacts as equivalent variation (EV) using the nonhomothetic measure of utility. In particular, equivalent variation is defined as the change in nominal income from the original level, w_k^0, that would leave the agent able to achieve post-shock utility, U_k^1, at the pre-shock vector of prices, $\{P_{ak}^0, P_{mk}^0, P_{sk}^0\}$. Since EV is negative when the agent becomes worse off, willingness-to-pay has the opposite sign:

$$WTP_k = -EV_k = -\left[E(U_k^1; P_{ak}^0, P_{mk}^0, P_{sk}^0) - w_k^0\right] \tag{23}$$

I also quantify the impact of sectoral productivity changes on measured GDP by using a standard Törnqvist (1936) price index that uses sectoral expenditure shares from before and after the shock, (ω_{jk}^0 and ω_{jk}^1), to construct an aggregate price index with which to deflate nominal income:

$$P_k^T = \prod P_{jk}^{(\omega_{jk}^0 + \omega_{jk}^1)/2} \longrightarrow GDP_k = \frac{w_k L_k}{P_k^T} \tag{24}$$

This calculation captures the logic of Baqaee and Farhi (2019), who extend Hulten (1978) to show that the aggregate productivity impact of a sectoral shock is given by the weighted average of the pre and post-shock sectoral shares. While aggregate productivity does not capture welfare in this setting with non-homotheticities, showing how it varies across counterfactuals helps illustrate the key mechanism in the model: trade barriers and subsistence requirements for food play a critical role in determining the welfare consequences of climate change by forcing labor toward agriculture rather than allowing it to reallocate to less affected sectors.

[29] This equation is also derived in Uy, Yi and Zhang (2013).

5.2 Comparative Statics

I now use the model to characterize the factors that influence sectoral reallocation in response to climate change. Consider a country that suffers an agriculture-biased reduction in aggregate productivity, consistent with projections for hot parts of the world made in Section 7. To see how the employment share in agriculture changes in Equation 22, I start with a version of Equation 11 for agriculture's expenditure share expressed in logs:

$$log(\omega_{ak}) = log(\Omega_a) + \underbrace{(1-\sigma)log\left(\frac{p_{ak}}{P_k}\right)}_{\text{Substitution Effect}} + \underbrace{(\epsilon_a - (1-\sigma))log\left(\frac{w_k}{P_k}\right)}_{\text{Income Effect}} \tag{25}$$

The agriculture-biased reduction in productivity has two effects that appear in Equation 25.[30] First, the reduction in productivity drives down the equilibrium real wage $(\frac{w_k}{P_k})$, making consumers poorer. If $(\epsilon_a - (1-\sigma)) < 0$, as is the case with the parameter estimates presented in Section 6, then the reduction in real wage drives up the expenditure share on food, ω_{ak}. This is the effect of nonhomotheticity. Food is a larger share of consumption for poorer people, so climate change tends to drive up the share of agricultural consumption by making people poorer.

Second, the relative decline in agricultural productivity will increase the domestic price of agricultural goods relative to the aggregate price index $(\frac{p_{ak}}{P_k})$.[31] If $\sigma < 1$, as is also the case in Section 6, then the rising relative price of agricultural goods raises the expenditure share on agriculture. Intuitively, if food is not substitutable with other consumption, then its relative quantity falls less than the relative price rises, and the share of spending on food goes up. This is the same logic that underlies Baumol's cost disease (Baumol and Bowen, 1966), a theory that endeavors to explain why low-substitutability service sectors with relatively low productivity growth, such as health care and education, tend to rise as a share of expenditures over time.

Together, nonhomotheticity and low substitutability at the sector level combine to push up the expenditure share on agriculture in response to declines in agricultural productivity. The macro-development literature on structural transformation (see, for instance, Gollin, Parente and Rogerson (2007)) refers to these features of consumer preferences as 'the food problem' - the explanation given to the large share of the labor force in agriculture in most developing countries despite very low absolute and relative productivity.

These features of the model also explain why its predictions about the protec-

[30] This equation also appears in Comin, Lashkari and Mestieri (2015). They estimate that nonhomotheticities (the income effect) account for about 75% of observed historical structural transformation, with changes in relative prices (the substitution effect) accounting for the rest.

[31] In a closed economy, relative sectoral prices are exactly proportional to sectoral productivities. In an open economy, the domestic relative price of agriculture responds to domestic agricultural productivity in proportion to the domestically produced share of consumption.

tive effects of reallocation diverge from those of Costinot, Donaldson and Smith (2016). Their paper finds that reallocating production across crops reduces the aggregate damages from climate change by two-thirds. To capture reallocation at the crop level, their model has no income effects and high substitutability across products.[32] This specification makes sense for capturing reallocation across crops, but does not generalize to the cross-sector case where income effects become important and the elasticity of substitution is very low. Intuitively, if the productivity of corn falls markedly relative to the productivity of wheat, consumers can respond by eating more wheat. If the productivity of producing food falls relative to the productivity of manufacturing, however, consumers cannot subsist by eating more manufactured goods.

In contrast to the food problem, the Ricardian comparative advantage effects of falling relative productivity in agriculture will tend to push labor into other sectors. Returning to Equation 22, shifting comparative advantage away from agriculture will tend to push up food imports (π_{akk} falls for country k) and push down food exports (π_{akn} falls). Equation 26 captures the horserace between the food problem and international trade that drives general equilibrium sectoral reallocation in response to climate change.[33]

$$l_{ak} = \underbrace{\pi_{akk}}_{\downarrow} \underbrace{\omega_{ak}}_{\uparrow} + \underbrace{\sum_{n \neq k}^{N} \pi_{akn} \omega_{an} \frac{w_n L_n}{w_k L_k}}_{\downarrow} \tag{26}$$

In autarky, falling relative agricultural productivity would drive up the employment share in agriculture. In an economy with costless trade, climate change would dramatically shift the global geography of agricultural production and trade flows, substantially limiting the impact on welfare. To quantify the relative strength of these effects in practice, I need to estimate the parameters of the model and simulate the general equilibrium response to the estimated impacts of climate change on productivity at the country-sector level.

6 Model Calibration

I solve the model presented in Section 5 by computing the equilibrium numerically. By simulating the model in levels and explicitly computing all moments – wages, sectoral consumption and output shares, sectoral bilateral trade flows, and sectoral price indices – for all countries, I am able to estimate key model parameters using simulated method of moments (SMM), assess the quality of the model's fit across a wide range of empirical moments, and disentangle competing mechanisms in the

[32] They estimate an elasticity of substitution of 5.4 across varieties of the same crop and 2.8 across crops. I estimate an elasticity of 0.27 between sectors.

[33] The importance of trade for promoting structural transformation out of agriculture has been previously emphasized by Tombe (2015), Teignier (2018), and Uy, Yi and Zhang (2013).

counterfactual simulations.[34]

6.1 Parameter Estimates

I infer model parameters from the data using a combination of calibration and estimation. Table 3 shows a list of the data sources and target moments corresponding to each parameter. I choose sector-country productivities, non-homothetic CES preference parameters, and bilateral trade costs to match sectoral value-added per worker, GDP shares, and bilateral trade flows for 158 countries. Appendix E.1 contains details about data construction and the SMM procedure.

Table 4 displays the estimates of the preference parameters for the nonhomothetic CES utility specification. The values are estimated using a simulated annealing procedure to minimize the sum of squared errors between simulated and empirical sectoral GDP shares across all countries, conditional on all other parameters.[35] Intuitively, each preference parameter corresponds to a key feature of the sectoral share data used in the estimation. The utility elasticities, ϵ_a, ϵ_m, and ϵ_s, that govern

Table 3: Model Parameters and Target Moments

Parameters	Data Moment	Data Source
σ	Sectoral GDP Shares	World Bank
$\Omega_a, \Omega_m, \Omega_s$	Sectoral GDP Shares	World Bank
$\epsilon_a, \epsilon_m, \epsilon_s$	Sectoral GDP Shares	World Bank
θ_a, θ_m	Calibrated from Tombe (2015)	
τ_{jkn}	Trade Flows	UN Comtrade
Z_{jk}	Sectoral Value-Added per Worker	World Bank
L_k	Population	World Bank

Notes: Table shows the data sources for moments targeted in the simulated method of moments procedure to estimate parameters for the model presented in Section 5. Data is for the global cross-section in 2011, accessed from the World Bank Databank.

[34] Solving the model in changes following Dekle, Eaton and Kortum (2007) is not feasible for the key results presented in Section 7.6 because doing so would require data that does not exist for the future baseline global equilibrium. In particular, implementing the climate change counterfactuals in relative changes would require information about the future levels of sectoral expenditure shares and trade flows without climate change. By simulating the model in levels, both the baseline (no climate change) and counterfactual (with climate change) values of all endogenous moments can be computed by allowing the productivity parameters to evolve using projections of future baseline economic growth and the estimated sectoral impacts of climate change.

[35] Given that the simulations incorporate trade, sectoral GDP shares translate directly into sectoral expenditure shares once net exports are subtracted. Thus, production shares from the data can be used to infer the parameters that govern consumption shares.

Table 4: Consumption Parameter Estimates

Parameter	Description	Estimate
σ	Cross-Sector Elasticity of Substitution	0.27
		(0.21)
ϵ_a	Agriculture Utility Elasticity	0.29
		(0.39)
ϵ_m	Manufacturing Utility Elasticity	1.00
		(0.27)
ϵ_s	Services Utility Elasticity	1.15
		(0.41)
Ω_a	Agriculture Taste Parameter	11.73
		(0.51)
Ω_m	Manufacturing Taste Parameter	3.70
		(0.35)
Ω_s	Services Taste Parameter	10
		(-)

Notes: Parameters estimated using simulated method of moments. Standard errors in parentheses calculated following Gourieroux, Monfort and Renault (1993) with derivatives simulated numerically. Ω_s is normalized to 10 as only relative values of Ω_j affect consumer choices. Within sector elasticity of substitution across varieties, η, is calibrated to 1.

non-homotheticity in the model are inferred from the pattern by which sectoral shares vary with income across countries. The sectoral taste parameters, Ω_a, Ω_m, and Ω_s, follow from the average level of each sector's shares across countries. Finally, σ is inferred from the degree to which sectoral shares vary as a function of relative prices, conditional on income.[36]

Two points about these estimates are worth noting. First, I estimate a cross-sector elasticity of substitution, $\sigma = 0.27$, of substantially less than one, indicating that the expenditure share in a sector sharply increases with its relative price. My estimate of σ to target the global cross-section of sectoral shares matches up well with that of Comin, Lashkari and Mestieri (2015), who use various historical panel datasets to estimate σ between 0.2 and 0.6.[37]

[36]While I do not use the data on prices directly in estimation, sectoral relative prices in the model follow from domestic relative productivities and trade costs, which are inferred from sectoral value-added per worker and observed trade flows, respectively. Figure 7 shows that the pattern of relative prices in the model is similar to that of the data.

[37]I estimate consumption parameters separately rather than calibrating to the values they estimate for three reasons. First, their paper does not present estimates of Ω_a, Ω_m, and Ω_s. Second, point estimates vary somewhat across specifications in their analysis so independent estimation provides an alternative to arbitrary designation of a preferred specification. Third, the fact that

Second, I estimate that $\epsilon_a - (1-\sigma) = -0.44$, which implies from Equation 25 that the consumption share of agriculture is strongly diminishing in real income. The estimates of ϵ_a, ϵ_m, and ϵ_s also align closely with those of Comin, Lashkari and Mestieri (2015). In particular, the estimates in Table 4 imply sectoral income elasticities of 0.48 for agriculture, 0.98 for manufacturing, and 1.09 for services, whereas those presented in their paper range from 0.37 to 0.56 for agriculture, 0.83 to 1.03 for manufacturing, and 1.14 to 1.20 for services.[38] Thus, referring back to Equation 25, the parameter estimates imply clearly that an agriculture-biased decline in productivity will raise the expenditure share of agriculture through both the income and substitution effect.

6.2 Model Fit

The model closely matches the features of the data most relevant to the counterfactual simulations of the impacts of climate change. Table 5 summarizes the correlation between key simulated moments in the model and their empirical counterparts. The simulations match the income level of each country almost exactly

Table 5: Summary of Model Fit

	(1) Data log(GDP per capita)	(2) Data Ag Share of GDP	(3) Data π_{akk} (Ag Domestic Production Share)
Simulated log(GDP per capita)	1.006		
	(0.00251)		
Simulated Ag Share of GDP		0.866	
		(0.0563)	
Simulated π_{akk} (Ag Domestic Production Share)			1.009
			(0.0392)
Observations	158	158	158
R^2	0.999	0.603	0.809

Notes: Table shows the results from regressing empirical moments in the data on their simulated counterparts. Data on nominal income levels and the agriculture share of GDP are from the World Bank. Data on the domestically produced share of expenditures in agriculture is constructed using Comtrade data. A coefficient of 1 with $R^2 = 1$ would constitute a perfect fit. The fit for other moments in the model is displayed in Appendix E.2 Figures A-27, A-28, and A-29.

using indirect inference to match sectoral shares in the global cross-section produces very similar estimates to those that they generate from historical panel data using instruments for prices and expenditures lends an additional dimension of support to the parameter values used in the analysis.

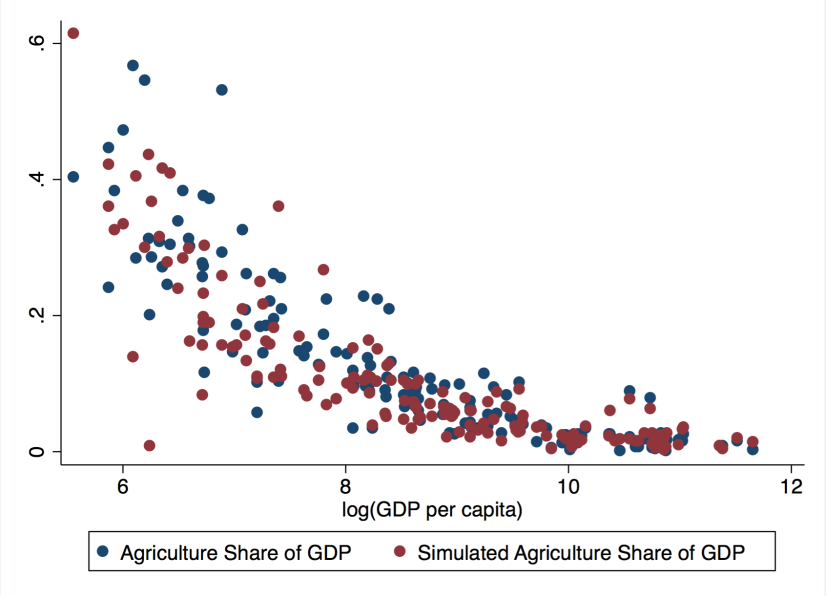
[38] The formula for the income elasticity in sector j in the non-homothetic CES specification is given by $\sigma + (1-\sigma) \times \epsilon_j \times \sum_{j \in J} \omega_j \epsilon_j$ where ω_j is the expenditure share. I report income elasticities for the expenditure shares of the average country in my sample.

through the calibration of the country-level aggregate productivity parameters. Similarly, the simulations closely match the domestic production share of agricultural consumption since I choose exporter-importer-sector-specific trade costs, τ_{jkn}, to match all observed bilateral trade flows. As shown in Appendix Figure A-30, most developing countries import little of their food. In the data, the average person in the poorest quartile of the world consumes 91.3% domestically produced food (89.4% in the simulation) compared to 45.1% in the richest quartile (52.4% in the simulation). I present suggestive evidence on some of the underlying causes of these high barriers to trade in poor countries in Section 9.

The model also explains most of the variation in the global agriculture share of GDP. I slightly under-predict agricultural shares on average, but overall the model explains 60.3% of the variation in the data. This is a relatively strong fit considering that only the seven free parameters in Table 4 were chosen to match 316 target moments consisting of GDP shares for agriculture, manufacturing, and services in 158 countries. As shown in Figure 6, the nonhomothetic CES demand specification enables the simulation to closely mirror the smooth decline of agricultural GDP with log income per capita.[39]

The model also reproduces the general pattern of high relative prices for agri-

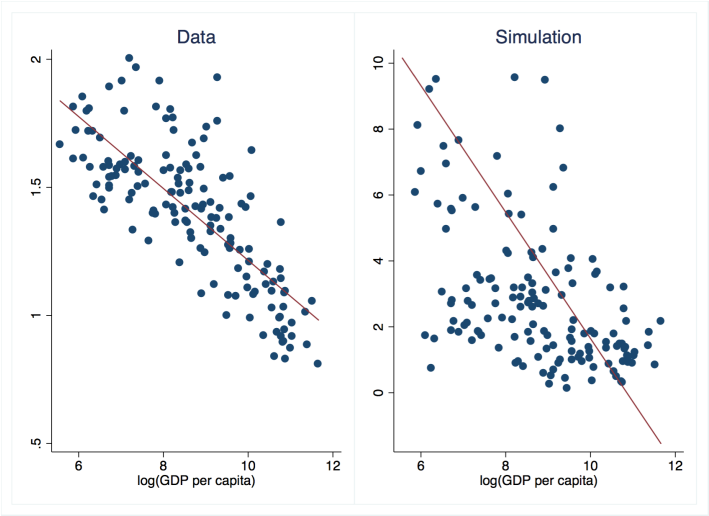
Figure 6: Agriculture Share of GDP - Data vs. Simulation



Notes: Graph shows the fit of simulated agriculture share of GDP in the model to data from the World Bank. The simulation explains over 60% of the variation in the data, and reproduces the smooth pattern of non-homotheticity observed in the empirical relationship between agricultural shares and income.

[39]For comparison, the best fit using a Stone-Geary utility specification has an R^2 of 0.43 and predominantly underpredicts the agriculture share as shown in Appendix Figure A-31.

Figure 7: Relative Price of Food - Data vs. Simulation



The graph on the left shows the ratio of a country-level food price index to an aggregate price index using data from the International Comparison Program. The graph on the right shows an analogous moment in the model - the ratio of the aggregate agricultural and manufacturing price indices, P_a and P_m. The model reproduces the empirical relationship that poor countries tend to have higher relative prices for food - a moment I do not target in the estimation.

cultural consumption in poor countries - a moment I do not target in the estimation. In Figure 7, I compare the simulated pattern of the relative price of agricultural and manufacturing consumption, P_{ak} and P_{mk}, to an empirical analogue constructed using aggregate sectoral price indices from the World Bank's International Comparison Program. While the simulated and empirical price indices have different units that prevent direct comparison, they share the same pattern of high relative prices for food in developing countries with low relative agricultural productivity.

Overall, the model matches the existing global pattern of sectoral specialization through a combination of consumer preferences and barriers to trade. Low incomes and the high relative price of food drive up agriculture's share of expenditures in poor countries through the nonhomotheticity and low elasticity of substitution in the preference specification. High estimated trade costs chosen to rationalize observed trade flows tightly link domestic consumption to domestic production, causing many developing countries to specialize in agriculture despite its low productivity.[40] In the next section, I use the model to investigate projected sectoral reallocation in response to climate change and its welfare consequences.

[40] As discussed in Section 5, this explanation is consistent with the work of Tombe (2015), Gollin, Parente and Rogerson (2007), and the broader literature on structural transformation.

7 Model Counterfactuals

This section uses the estimated model to project the impacts of climate change on trade flows, sectoral specialization, prices, GDP, and welfare.

7.1 Estimated Productivity Impacts

I start by projecting the impacts of climate change on country-sector level productivity. For agricultural productivity effects, I use the estimates from Cline (2007) displayed in Figure 1. As I explain in Appendix E.3, I choose these projections for two reasons. First, they use globally representative data to produce results broadly consistent with the literature on climate and agricultural production. Second, they represent the most comprehensive available source of global impact estimates that account carefully for adaptation and crop-switching within the agricultural sector.

To project the impact of climate change on productivity in manufacturing and services, I combine the country-sector specific temperature sensitivities estimated in Section 4.4 with projections of the future distribution of temperature in 2080-2099.[41] I obtain future temperature predictions from the CSIRO-MK-3.6.0 model produced by Jeffrey et al. (2013), one of the climate models used by Cline (2007), for consistency with the projected changes in agricultural productivity.[42] The projected changes in manufacturing and services productivity are shown in Figure 8a and Appendix Figure A-24 respectively. Figure 8b brings together the estimated impacts on agricultural productivity from Cline (2007) with my estimates of the change in manufacturing productivity to show the change in the relative productivity of the model's tradable sectors for every country.

The pattern in Figure 8b shows clearly that climate change shifts comparative advantage in agriculture toward colder countries far from the equator on average. While the negative effects of climate change on manufacturing productivity are concentrated in similar parts of the world to agricultural productivity, they are generally smaller in magnitude. Every country in Africa, South Asia, and Latin America (with the exception of Egypt) has larger estimated productivity losses in agriculture than manufacturing. Thus, to the extent that specialization follows Ricardian comparative advantage, we would expect to see agricultural production move toward colder places away from the equator in response to climate change.

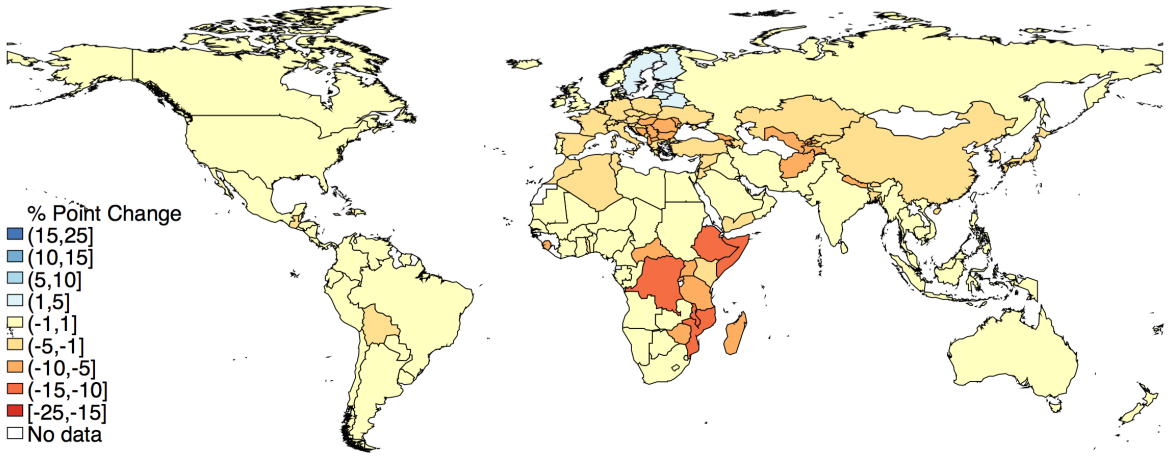
Next, I integrate these empirically estimated productivity changes into the model by applying them to the sector-country specific aggregate productivity Z_{jk} and recalculating equilibrium wages, prices, and trade flows.

[41] I use the estimates that allow for firms to adjust adaptation investments to their end-of-century temperatures. I account for the costs of this adaptation in Section 7.6.

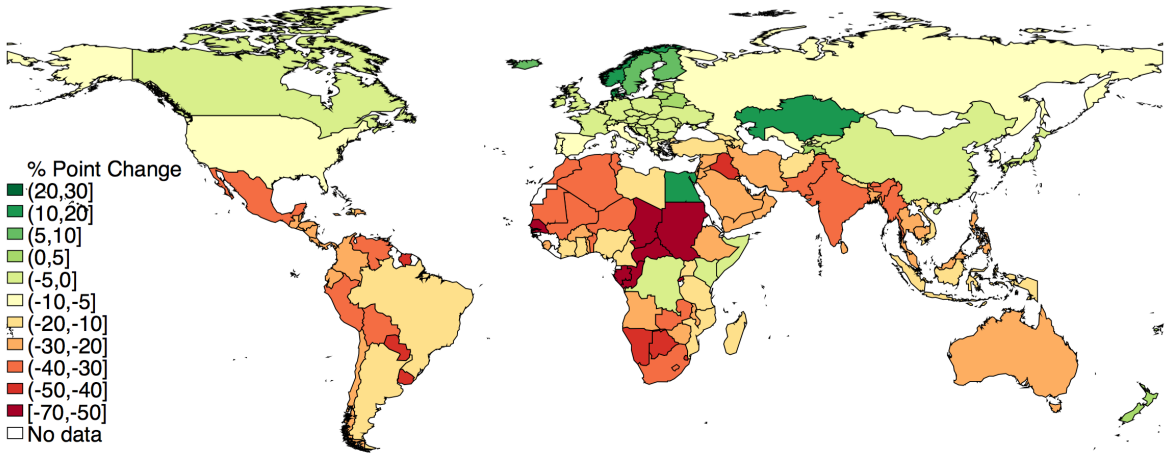
[42] The estimates from the interacted model in Section 4 give me an estimate of the reduction in annual manufacturing and services output per worker for each degree-day above 30°C and below 5°C. The CSIRO model projections give me population-weighted change in degree-days above 30°C and below 5°C for every country in the world in 2080-2099, which are shown in Appendix Figures A-22 and A-23. I multiply the country-level coefficients by the projected changes in hot and cold temperatures to get the impacts shown here.

Figure 8: Projected Impact of Climate Change on Productivity

(a) Manufacturing



(b) Agriculture Relative to Manufacturing

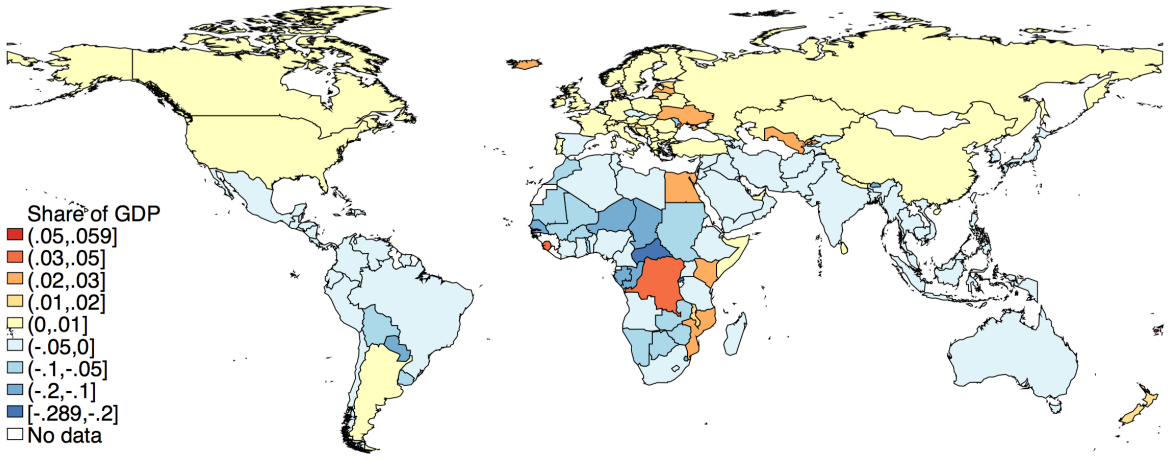


Notes: Panel A shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 2 of Table 2 at each country's income and end-of-century long-run average temperature. Panel B shows the change in agricultural productivity from Cline (2007) minus my estimate of the change in manufacturing productivity, shown above, in percentage points.

7.2 Comparative Advantage and Trade

Figure 9 shows the projected equilibrium change in agricultural net exports in response to climate change. Consistent with the estimated change in comparative advantage, the predominant pattern is that hotter countries experiencing large declines in agricultural productivity import more food, while cooler countries with neutral or improving agricultural productivity export more food. For instance, Den-

Figure 9: Projected Impact of Climate Change on Agricultural Net Exports



Notes: Map shows model simulations of the change in agricultural net exports as a share of GDP driven by the effects of climate change on sector-level productivity and comparative advantage shown in Figure 8b.

mark and Canada roughly double agricultural net exports, from 1.9% to 3.8% and 0.5% to 1.2% of GDP respectively. Conversely, most of Sub-Saharan Africa and South Asia increase imports of food. The few exceptions to this finding are those hot countries for whom the change in agricultural productivity is not large relative to the change in manufacturing productivity, particularly in relation to their close trading partners.

The magnitudes of the projected change in trade flows are generally modest as a share of the economy. No country increases agricultural net exports by more than 6% of GDP, and only 12 out of 158 countries decrease agricultural net exports by more than 10% of GDP.

7.3 Sectoral Reallocation

As shown in Section 5.2, the change in trade flows is only a partial summary of the change in sectoral specialization. Agriculture's share of GDP (and consequently the labor force) depends on both the change in net exports and the change in the expenditure share on food. I reproduce Equation 26 summarizing labor reallocation in response to an agriculture-biased decline in productivity here for convenience:

$$l_{ak} = \underbrace{\pi_{akk}}_{\downarrow}\underbrace{\omega_{ak}}_{\uparrow} + \underbrace{\sum_{n=1}^{N} \pi_{akn}\omega_{an}\frac{w_n L_n}{w_k L_k}}_{\downarrow}$$

The change in net exports shown in Figure 9 captures the first and third effects in the above equation. Given the strong nonhomotheticity and low cross-sector substitutability implied by the estimates of ϵ_a and σ in Section 6, the change in the agriculture expenditure share, ω_{ak}, is also likely to be substantial.

The horserace between these two competing effects - comparative advantage and 'the food problem' - that govern sectoral reallocation in response to climate change plays a critical role in the aggregate productivity and welfare consequences. As discussed in Section 5, the simple logic formalized by Baqaee and Farhi (2019) is that production moving toward the sector suffering a larger decline in productivity exacerbates the aggregate consequences of a given shock.

I decompose the competing effects of climate change on the agriculture share of GDP by running separate counterfactuals with and without trade. In autarky, the change in a sector's relative price equals the change in that sector's productivity. Thus, I start by applying country-sector level price changes equal to the inverse of the projected change in productivity and calculating the change in expenditure shares. This gives me the change in ω_{ak}, which in autarky equals the change in agriculture's share of GDP. In contrast, the standard counterfactual incorporating trade gives me the full effect of both types of reallocation. Table 6 displays the baseline, autarky counterfactual, and trade-inclusive counterfactual agriculture shares of GDP for a selection of countries.

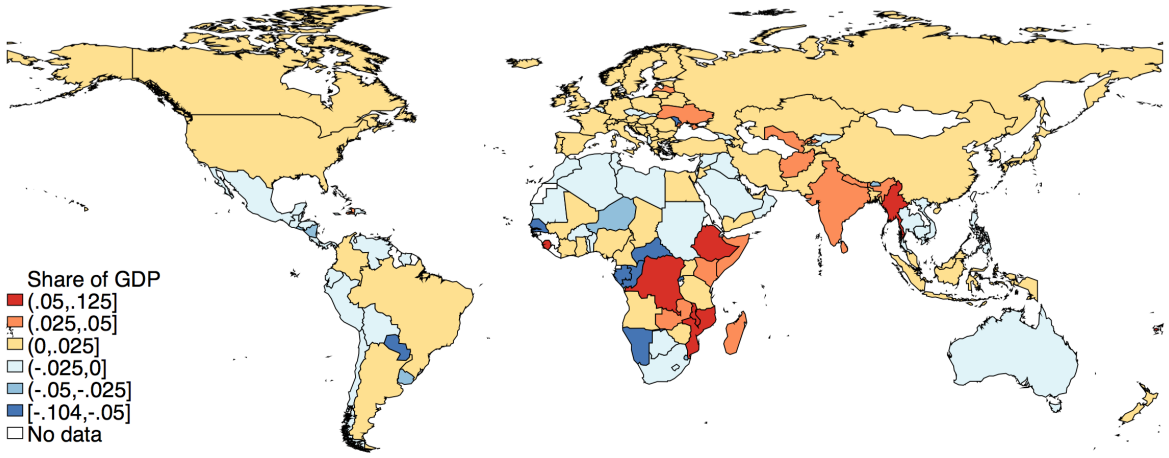
The results in Table 6 show that the consumption response and trade response both have substantial effects on specialization in agriculture, with significant heterogeneity across countries. In Ethiopia and India the 'food problem' effect dominates and the agriculture share of GDP rises in response to climate change despite

Table 6: Counterfactual Agriculture GDP Shares - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	.065	.067	.07
Canada	-.022	-.007	.019	.019	.026
China	-.072	-.036	.064	.068	.074
Denmark	.109	.006	.033	.032	.051
Ethiopia	-.313	-.102	.359	.437	.409
India	-.381	0	.161	.224	.194
Kenya	-.054	-.044	.156	.16	.185
Mozambique	-.217	-.104	.367	.426	.451
Rwanda	-.601	-.058	.409	.678	.351
Poorest Quartile	-.319	-.02	.199	.256	.227
World	-.101	-.01	.038	.044	.043

Notes: Table shows model simulations of the change in agriculture share of GDP driven by the effects of climate change.

Figure 10: Projected Impact of Climate Change on Agricultural GDP Share



Notes: Map shows the model simulations of the change in the agriculture share of GDP driven by climate change.

large relative declines in agricultural productivity. In contrast, the trade effect dominates in Rwanda, where the domestic production share of agricultural expenditures falls from 85% to 54%. Other countries, such as Canada, Denmark, Kenya, and Mozambique see an increase in agricultural specialization because of increased exports driven by improvements in relative agricultural productivity compared to their trading partners.

Figure 10 shows the worldwide change in agriculture's share of GDP. On average, the global agriculture share of GDP rises from 3.8% to 4.3% because agricultural productivity falls in more places than it rises, raising ω_{ak}, and net exports for the world are zero. The dominance of the 'food problem' is particularly relevant in poor countries disproportionately suffering from extreme heat. The population-weighted average change in agriculture's share of GDP in the poorest quartile of the world is +2.8 percentage points. Thus, simulations suggest that climate change will push more people in poor countries into farming as the productivity of those farms declines dramatically.

7.4 Aggregate Productivity and Willingness-to-Pay

The estimated sectoral productivity effects combined with the changes in sectoral specialization map directly into changes in aggregate productivity. Table 7 shows the change in real GDP for each counterfactual in select countries, deflating nominal income at the country level using the Tornqvist price index from Equation 24.

The results make clear that projected reallocation *exacerbates* the impact of climate change on aggregate productivity in most countries, as well as globally on average. Global GDP declines 1.9% in the counterfactual that holds sectoral shares fixed, but declines 2.1% when allowing for reallocation. GDP in the poor-

Table 7: Counterfactual Percent Change in GDP - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.002	0	.001
Canada	-.022	-.007	-.018	-.018	-.016
China	-.072	-.036	-.043	-.045	-.045
Denmark	.109	.006	0	0	.005
Ethiopia	-.313	-.102	-.163	-.218	-.217
India	-.381	0	-.074	-.131	-.127
Kenya	-.054	-.044	-.037	-.038	-.034
Mozambique	-.217	-.104	-.14	-.192	-.199
Rwanda	-.601	-.058	-.334	-.557	-.508
Poorest Quartile	-.319	-.02	-.083	-.132	-.126
World	-.101	-.01	-.019	-.023	-.021

Notes: Table shows model simulations of the change in GDP driven by the effects of climate change.

est quartile of countries falls by 8.3% in the no reallocation counterfactual, and falls 12.6% with reallocation. This happens for two reasons. First, as discussed in Section 7.3, the 'food problem' pushes up the labor share of agriculture in many countries while agricultural productivity declines dramatically. Second, as Dingel, Meng and Hsiang (2019) have shown, the spatial correlation of the productivity impacts heighten their importance. Since food prices in Rwanda are a function of agricultural productivity in Rwanda and its closest trading partners, the losses to Rwanda intensify when accounting for shocks that hit their neighbors.

How can reallocation that worsens aggregate productivity and measured GDP be consistent with optimizing behavior? In Table 8, I calculate the willingness-to-pay (WTP) to avoid climate damages under each counterfactual as the equivalent variation loss in income at the baseline equilibrium set of wages and prices, as shown in Equation 23. The results show that the full reallocation counterfactual reduces the welfare consequences of climate change, as captured by willingness-to-pay, even while increasing the impact on GDP. The WTP under the no reallocation counterfactual is particularly dramatic because it forces agents to deviate from optimal consumer behavior. This highlights that the no reallocation counterfactual is actually an unrealistic scenario. In the presence of declining income and large projected increases in food prices, keeping fixed the expenditure share on food would require declines in the quantity of food consumed that are strongly inconsistent with the observed low substitutability between food and non-food. To summarize the intuition, people are willing to sacrifice income (GDP) to reallocate expenditures toward food when food prices rise because they need food to survive.

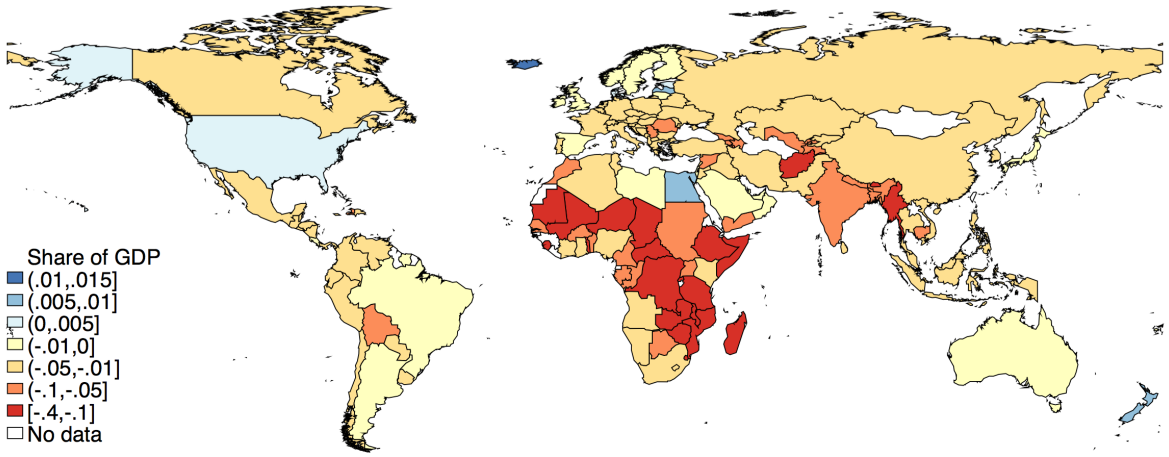
Figures 11 and 12 show the global distribution of willingness-to-pay to avoid

Table 8: Equivalent Variation Change in Welfare (Share of GDP) - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.008	-.002	0
Canada	-.022	-.007	-.018	-.016	-.014
China	-.072	-.036	-.057	-.04	-.04
Denmark	.109	.006	.003	0	.005
Ethiopia	-.313	-.102	-.364	-.171	-.169
India	-.381	0	-.311	-.085	-.082
Kenya	-.054	-.044	-.052	-.035	-.031
Mozambique	-.217	-.104	-.279	-.143	-.147
Rwanda	-.601	-.058	-.725	-.434	-.387
Poorest Quartile	-.319	-.02	-.277	-.092	-.088
World	-.101	-.01	-.04	-.018	-.017

Notes: Table shows model simulations of the willingness-to-pay to avoid the productivity effects of climate change using the non-homothetic CES measure of utility.

Figure 11: Equivalent Variation Change in Welfare from Climate Change

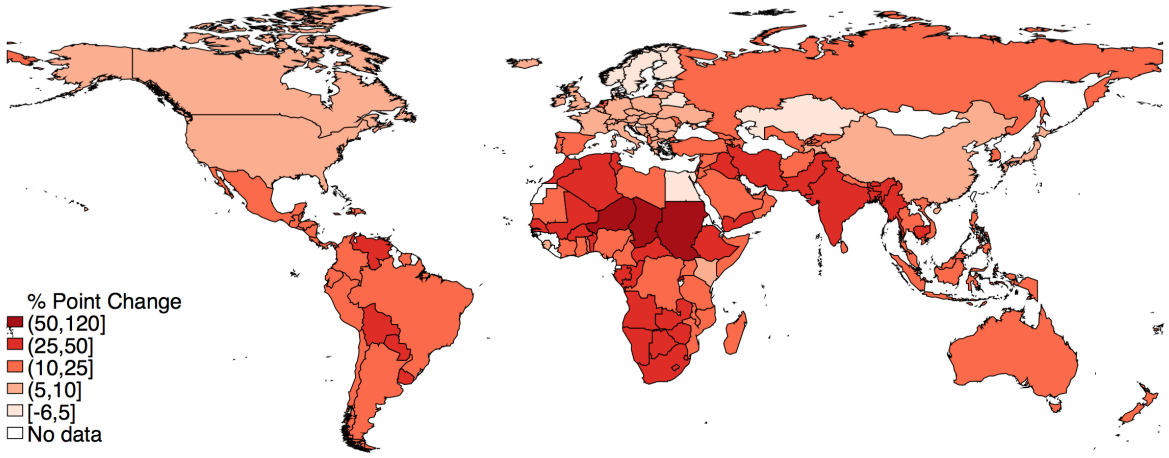


Notes: Map shows model simulations of the willingness-to-pay to avoid the effects of climate change as a share of GDP.

climate change, and the change in food prices, P_{ak}, which comprise a key driver of the welfare losses. Food prices rise in 156 of the 158 countries, and rise by at least 25% in 41 countries containing over 32% of the world's population.[43] Climate

[43]The large changes in food prices also imply that the incidence of these losses may fall on urban consumers more than on farmers suffering lost productivity. I investigate the distributional

Figure 12: Projected Percentage Change in Food Prices



Notes: Map shows model simulations of the change in the aggregate food price index, P_{ak}, driven by climate change.

change does net damage as measured by WTP in 150 countries, and causes welfare losses exceeding 8% of GDP in 32 countries covering 27% of the world's population. Because the losses are concentrated in poor countries, global willingness-to-pay is only 1.7% of GDP. However, the population-weighted average global losses are 4.7% of GDP, and the population-weighted average for countries in the bottom quartile of income is 8.8% of GDP. The interpretation of this number is that climate change will cost the average person in the poorest quartile of the world nearly 9% of their income. Note that these results account neither for the costs of firm-level adaptation investments nor for the benefits of anticipated economic growth, both of which will be examined in Section 7.6.

7.5 Low Trade Cost Counterfactual

The analysis of sectoral reallocation and aggregate productivity in Sections 7.3 and 7.4 demonstrates that openness to trade mitigates the harm from climate change by counteracting 'the food problem.' To further investigate the importance of trade policy, I run an additional counterfactual exercise in which I replace all bilateral trade costs, τ_{jkn}, for manufacturing and agriculture with a 100% tariff ($\tau = 2$), representing approximately the 90th percentile of trade openness.[44] I choose this number rather than 0% to acknowledge the fact that some level of shipping costs, regulatory discrepancies, and language barriers are inherent to cross-country trade, so no amount of policy change could make trade perfectly costless. A 100% tariff-

consequences of climate change within countries further in Appendix F.3.

[44]The 10th percentile of bilateral tariff-equivalent trade costs in my estimates is 122% in agriculture and 85% in manufacturing.

equivalent trade cost is approximately equal to the cost I estimate for shipping food from Belgium to Australia; representing an ambitious, yet realistically feasible, change in global trade policy.

To disentangle the benefits of trade for climate change adaptation from the more general gains from trade, I rescale each country's vector of sectoral productivity parameters, Z_{jk}, such that I continue to match the baseline levels of GDP per capita in the initial equilibrium. Note, however, that without the estimated high barriers to trade in developing countries the model can no longer match the observed global pattern of the agriculture share of GDP. In this hypothetical world of increased openness, developing countries import substantially more food from richer countries with high relative productivity in agriculture even in the absence of climate change.

Table 9 shows the WTP to avoid climate change under different trade cost scenarios for a select subset of countries especially vulnerable to climate change. Two things about these results are worth noting. First, as shown in Table 9, reducing trade barriers dramatically reduces the costs of climate change in the hardest-hit countries. Overall, the WTP for the average person in the lowest quartile of global income is only 2.9%, relative to 8.8% in the estimated trade cost case.

Second, the effects of openness to trade vary substantially across countries. For 40 countries representing 15.1% of the global population, WTP to avoid climate change as a share of GDP is *higher* in the low trade cost scenario.[45] The intuition

Table 9: Equivalent Variation Change in Welfare (Share of GDP)
Alternative Trade Cost Cases

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Rwanda	-.434	-.387	-.086
Central African Republic	-.428	-.356	-.037
Chad	-.25	-.226	-.032
Malawi	-.225	-.225	-.119
Zimbabwe	-.223	-.212	-.074
Zambia	-.208	-.199	-.001
Ethiopia	-.171	-.169	-.091
Sierra Leone	-.13	-.164	-.105
India	-.085	-.082	-.013
Poorest Quartile	-.092	-.088	-.029
World	-.018	-.017	-.013

Notes: Table shows model simulations of the willingness-to-pay to avoid the effects of climate change under different scenarios - autarky, estimated global barriers to trade, and an alternative scenario setting all bilateral trade costs to approximately the 90th percentile level of trade openness.

[45] To be clear, these countries still experience overall gains from trade. But once those general gains are netted out, they suffer larger climate change damages in this scenario.

for this result is as follows. When trade barriers are high and local consumption depends mostly on local production, the effects of deteriorating productivity are also concentrated locally. Conversely, more trade makes the world more interdependent and dilutes the effects of a local shock across many countries. If consumption in Austria is more linked to production in Zimbabwe, then Austrian consumers suffer more from shocks that hit Zimbabwe. Conversely, Zimbabwean consumers insulate themselves from the local shock by consuming a more diversified global portfolio of products.

Overall, trade reduces the aggregate global willingness-to-pay to avoid climate change by 7.4% relative to autarky under existing global trade policy, and by 30.7% under the specified alternative assumption of freer trade. This pattern holds much more starkly in poor countries. For the average person in the poorest quartile of the world, trade reduces WTP by 4.5% relative to autarky under existing policy, but by 68.2% under freer trade. I discuss possible policy mechanisms to realize these gains in Section 9.

7.6 Future Income Baseline

The results in Sections 7.1 to 7.5 use projections for future temperature change, but hold the baseline global economy fixed at the present day equilibrium. In this section, I endeavor to account for the effects of anticipated economic growth by allowing global income levels to evolve according to an example set of projections from the Shared Socioeconomic Pathway (Scenario Three) developed by Cuaresma (2017) of the International Institute for Applied Systems Analysis.[46]

Allowing for economic growth to take place has two important effects on the aggregate consequences of climate change. First, the agriculture share of GDP declines as countries grow richer due to nonhomothetic preferences for food, reducing the aggregate consequences of agriculture-specific productivity shocks. I capture this effect in the model by applying projected income growth to 2080 as sector-neutral increases in the baseline values of Z_{jk}. Second, the results from Section 4 imply that sensitivity to temperature for manufacturing and services firms declines markedly as countries become richer. I capture this by re-evaluating the sensitivity to temperature shown in Figures 5a and 5b at 2080 levels of log GDP per capita. Appendix Figures A-25 and A-26 show that the effects of temperature on non-agricultural productivity accounting for adaptation are substantially muted, even in this relatively low growth scenario that projects only slightly more than a doubling of global income between 2015 and 2080.

Table 10 shows the impact of expected economic growth on the agriculture share of GDP and expected willingness-to-pay. The willingness-to-pay numbers in Columns 5 and 6 of Table 10 also incorporate the firm-level adaptation costs calculated in Appendix D, thus accounting more comprehensively for the costs as well as benefits of adaptation. This particular future scenario includes little to

[46]Use of the Shared Socioeconomic Pathways in future projections of climate change damages follows from the work of Carleton et al. (2018) and many others.

Table 10: Equivalent Variation Change in Welfare (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfac-tual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Central African Republic	1.47	.299	.287	.094	-.436	-.316
Rwanda	1.14	.409	.39	.322	-.394	-.366
Zimbabwe	4.17	.302	.122	.14	-.248	-.111
Malawi	2.84	.436	.309	.357	-.244	-.167
Zambia	1.52	.36	.28	.318	-.233	-.181
Chad	1.13	.257	.213	.243	-.226	-.221
Sierra Leone	1.49	.139	.177	.173	-.204	-.146
Ethiopia	1.23	.359	.333	.376	-.19	-.182
India	3.24	.161	.087	.106	-.082	-.045
Poorest Quartile	3.05	.199	.126	.144	-.10	-.062
World	2.2	.038	.025	.028	-.027	-.015

Notes: Table shows model simulations of the effects of projected economic growth on the agriculture share of GDP and the willingness-to-pay to avoid climate change in select countries. Example economic growth projections come from Cuaresma (2017).

no projected growth for many currently poor countries, allowing for contrast with those that grow faster. This comparison shows the importance of economic growth in reducing climate change damages. Table 10 shows that Zimbabwe and Malawi get substantially richer in this projection and their agriculture share of GDP and welfare losses decline markedly. In contrast, climate change continues to be very harmful to countries that grow slowly, such as Rwanda and Chad.

The results in Table 10 show that the aggregate global WTP for climate change is 2.7% of GDP at current global income levels and 1.5% at future projected incomes. The average WTP for a person in the bottom quartile of the world is 10.0% from the present baseline and 6.2% from the future baseline. To summarize the importance of the distributional consequences of climate change, I follow Jones and Klenow (2016) to calculate the willingness-to-pay of a Rawlsian social planner taking the certainty equivalent of being any person in the world with random probability.[47] The Rawlsian welfare losses from climate change are 6.2% of global GDP from the present income baseline and 3.6% of global GDP from the future baseline, more than twice as high as the aggregate willingness-to-pay calculated by summing across agents.

[47] Following Jones and Klenow (2016) I use log utility in this calculation.

7.7 Model Robustness

I consider robustness to three alternative model assumptions in Appendix F. In Appendix F.1 I represent subsistence requirements for food using generalized Stone-Geary preferences instead of the nonhomothetic CES specification in the baseline model. In Appendix F.2 I use lognormal, rather than Frechet, distributions for sector-country productivities across varieties. The numerical solution method described in Appendix E.1 allows for solving the model for an arbitrary range of productivity distributions and demand systems within and across sectors.[48] In Appendix F.3 I relax the assumption that all workers within a country have the same wage by laying out a version of the model with heterogeneous workers by skill type. This extension allows for capturing an additional feature of the data; agricultural workers generally have lower wages than non-agricultural workers. Appendix Table A-4 shows that the main counterfactual simulation results are very similar under the first two alternative modeling assumptions. The third extension with heterogeneous workers is not amenable to quantification, but I demonstrate the qualitative robustness of the main results and use the extension to explore additional dimensions of the implications of climate change for comparative advantage across sectors and the distributional impact on low and high skill workers.

8 Supporting Empirical Evidence

In this section, I present country-level panel regression evidence consistent with the model counterfactuals. In particular, the results in Section 7 suggest that the 'food problem' outweighs the trade response, on average, in driving sectoral reallocation due to climate change. This finding is supported by the simulated method of moments inference that underlies my parameter estimates, is consistent with both cross-sectional and historical patterns of sectoral specialization in the world, and is further bolstered by existing empirical evidence that aims to isolate the causal effect of agricultural productivity on structural transformation. In particular, Gollin, Hansen and Wingender (2018) proxy for improvements in agricultural productivity using variation in the development, diffusion, and climatic suitability for high-yielding crop varieties and Bustos, Caprettini and Ponticelli (2016) study the introduction of genetically engineered soybean seeds in Brazil. Both papers find that rising agricultural productivity drove labor out of agriculture and into industry.

More recent work by Fiszbein and Johnson (2020) provides evidence that the relationship between agricultural productivity and structural change varies with trade openness as predicted by the model in Section 7. They use a similar high-yielding crop variety instrument to show that agricultural productivity growth reduces agriculture's employment share in more closed economies, but raises it in a

[48] By avoiding the need for analytical expressions for key model components, this method offers flexibility similar to the procedure of Adao, Costinot and Donaldson (2017) while preserving representation of all components of the competitive equilibrium with no restrictions on the counterfactual exercises that can be considered.

subset of countries sufficiently open to trade. Since only a minority of countries meet their threshold of openness, their results further support the conclusion that the 'food problem' dominates comparative advantage in driving reallocation when agricultural productivity improves. Here, I present evidence relevant to the converse more representative of climate change - that exogenous declines in agricultural productivity increase the agriculture share of GDP and labor on average.

Table 11 summarizes the data sources used in this part of the analysis.[49] Following Schlenker and Roberts (2009), I use "growing degree days" (GDD) between 0°C and 29°C and "killing degree days" (KDD) above 29°C as temperature transformations representing positive and negative shocks to agricultural productivity respectively. I aggregate GDD and KDD to the country level for each year weighting by each pixel's share of cropland.[50]

Table 11: Country-Level Panel Data

Variable	Data Source
Temperature	Berkeley Earth Surface Temperature Dataset
Ag Share of GDP	World Bank
Ag Share of Labor Force	International Labour Organization
Food Share of Imports	UN Comtrade
GDP	World Bank

Notes: Data covers 164 countries from 1960-2012 with varying coverage by country and dataset. Economic data from all sources above are retrieved from the World Bank Databank.

I estimate the following panel regression with observations at the country-year level for four separate outcome variables - log GDP, food share of imports, agricultural share of GDP, and agricultural share of labor:

$$Y_{it} = \beta_1 GDD_{it} + \beta_2 KDD_{it} + \delta_i + \kappa_t + \epsilon_{it} \tag{27}$$

The regression exploits idiosyncratic variation in weather controlling for country fixed effects, δ_i, and year fixed effects, κ_t, to estimate the plausibly causal effect of shocks to agricultural productivity. I weight observations by their share of the

[49] I use BEST temperature data with a 1°global grid in this specification because aggregating GMFD temperature data from a 0.25°grid for every country worldwide exceeds my available computational resources.

[50] Following standard procedure in estimating temperature effects on agricultural productivity, degree days are calculated by fitting a sinuisoidal curve through daily minimum and maximum temperature, and then integrating the proportion of each day above a certain threshold.

global agricultural labor force to recover expected reallocation for the average farm worker in the world.

The results in Table 12 are broadly consistent with the model simulations in Section 7. The composition of imports shifts toward food in response to negative agricultural productivity shocks (KDD), and away from food in response to positive shocks (GDD), but the magnitudes of these changes are small. Consistent with an important role for 'the food problem,' the agriculture share of GDP and labor rise with KDD and fall with GDD, with magnitudes roughly similar to those in the model. In the regression, the agriculture share of GDP rises by slightly under 1 percentage point for an agriculture-biased shock that reduces GDP by 12%. Similarly, in the model simulations, the agriculture share of GDP rises by 2.1 percentage points in countries suffering large declines in agricultural productivity (<10 percentage points) from an average agricultural productivity fall of 29.5%.

Table 12: Country-Level Panel Regression

	(1)	(2)	(3)	(4)
	log(GDP)	Food Share of Imports	Ag Share of GDP	Ag Labor Share
KDD X 100	-0.121	0.00258	0.00875	0.00991
	(-2.31)	(0.64)	(1.08)	(1.55)
GDD X 100	0.0505	-0.00429	-0.00140	-0.00138
	(1.64)	(-2.45)	(-1.54)	(-0.38)
Observations	3602	2916	3171	3715
Country FE	X	X	X	X
Year FE	X	X	X	X
Ag Labor Weights	X	X	X	X

Notes: t-statistics in parentheses. Reported Driscoll and Kraay (1998) standard errors are robust to heteroskedasticity, spatial correlation, and autocorrelation of up to 5 lags. Results come from estimating Equation 27 with crop-area weighted growing and killing degree days. Data covers 164 countries from 1960-2012 with varying coverage by country and outcome variable. Economic data from all sources above are retrieved from the World Bank Databank.

The results from the country-level regressions are imprecise and insufficient in isolation to make full general equilibrium projections or welfare calculations relating to sectoral reallocation in response to climate change.[51] Taken together with the analysis in Sections 6 and 7 and the existing body of evidence, however,

[51] I show results for the unweighted regressions in Appendix Table A-1. I gain precision in the unweighted specification because the agriculture labor share weights are missing for a nontrivial share of the observations, but have a less interesting interpretation of the coefficients as effects on the average country in the world rather than on the average unit of agricultural labor.

these results reinforce the important role of the 'food problem' in mediating the aggregate consequences of climate-driven agricultural productivity shocks.

9 Discussion

This paper has three sets of implications relevant to policy on climate change and development. First, the results inform cost-benefit analysis on policies to reduce greenhouse gas emissions and mitigate climate change. These results are not a comprehensive evaluation of the costs of climate change - I omit international migration, uncertainty, health effects, and non-temperature effects such as storms and sea-level rise, among other topics, from the analysis. I do, however, address an existing challenge in the literature by estimating the welfare consequences of the global change in aggregate productivity in a framework that accounts for reallocation of production between agriculture and non-agriculture.

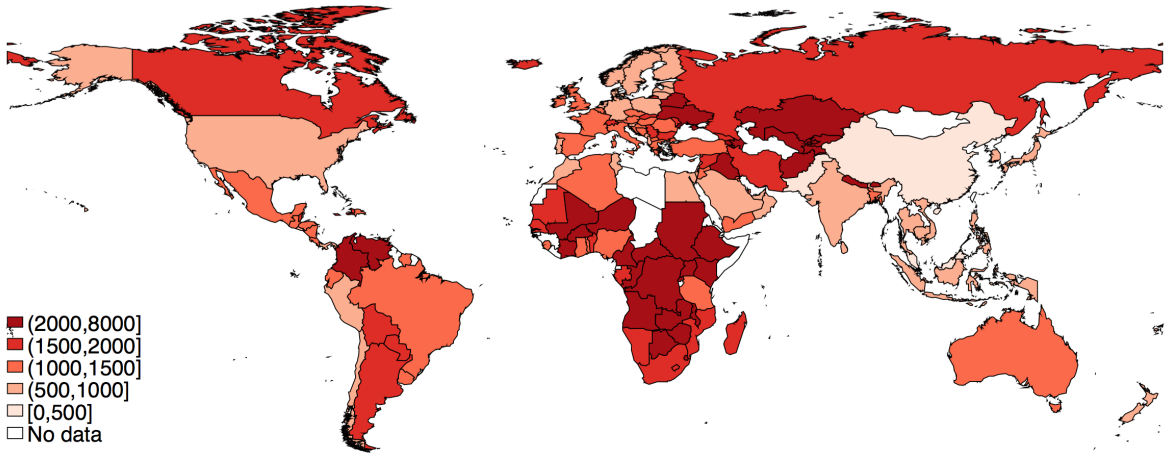
Second, the results inform decisions about the best way to channel efforts to adapt directly to the consequences of climate change. If it were true that agricultural activity is likely to shift substantially away from hot developing countries, optimal investments in adaptation might focus on retraining farm workers to transition to non-agricultural occupations. Instead, the finding that climate change is more likely to increase specialization in agriculture in hot countries underscores the urgent need to reduce the temperature-sensitivity of production through technology, irrigation, heat-resistant crop varieties, or other means. The agricultural productivity consequences projected by Cline (2007) will take place gradually and worsen far into the future, and need not be invariant to efforts to reduce them.

Third, and perhaps most importantly, the results speak to the importance of reducing barriers to trade in developing countries as a mechanism for climate change adaptation. The results in Section 7.4 show that increasing trade openness could dramatically reduce exposure to climate damages in the poorest countries in the world. Reducing tariffs would be one place to start, but tariffs account for a relatively small proportion of estimated trade costs. As Tombe (2015) documents at length, red tape barriers appear to be a far more important deterrent in many places. Figures 13a and 13b show data from the World Bank Ease of Doing Business Indicators on fees and delays associated with importing a container.

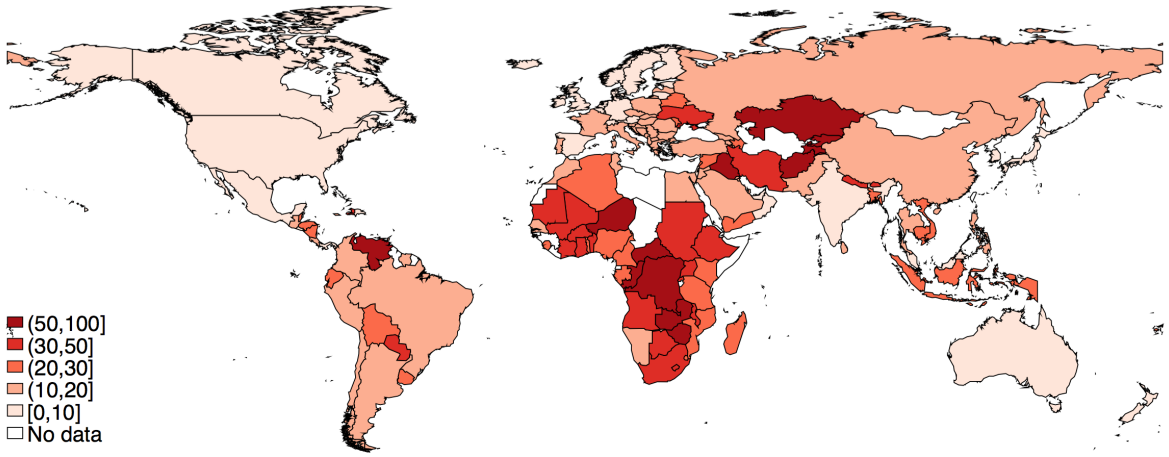
The average country in Sub-Saharan Africa requires 9 documents and over $2700 in fees for customs clearance, document processing, customs brokerage, terminal handling, and inland transport to import a 20-foot container of goods, exclusive of tariffs and unofficial payments. Importing a shipment to Sub-Saharan Africa also requires waiting an average of 37 days upon arrival at the border for compliance with customs clearance, inspection procedures, and document preparation, likely a prohibitive length of time for many food imports. Such patterns help clarify the statistic presented in this paper that the domestic production share of agricultural consumption ranges around 90% in much of this part of the world. Given that these types of trade barriers do not involve international negotiations or physi-

Figure 13: Non-Tariff Barriers to Trade

(a) Direct Costs to Import a 20-Foot Long Container (USD)



(b) Days to Import a Container



Notes: Panel A shows the direct cost to import one container of goods. Costs include documents, administrative fees for customs clearance, terminal handling charges, and inland transport, but not tariffs or taxes. Panel B shows the average number of days required to import a container. Delays include customs clearance, government inspection procedures, and documentary compliance requirements. Data comes from the World Bank Ease of Doing Business Index.

cal constraints to shipping over long distances, they could represent a relatively tractable target for reforms that could make a substantial impact on climate change adaptation.

10 Conclusion

Standard intuition suggests that reallocation should improve outcomes. Falling productivity raises prices and encourages substitution to other products. But this logic does not hold for broad categories of necessary consumption, such as food.

If a fall in productivity causes the price of corn to rise sharply, people can adapt by eating more rice. But when people become poorer and the relative price of food rises, they cannot compensate by substituting away from food.

This paper investigates the importance of subsistence requirements for food for the general equilibrium and welfare consequences of climate change. I show that climate change predominantly shifts comparative advantage in agriculture away from the equator as the effects of extreme temperatures on non-agricultural productivity are generally smaller than those in agriculture. On average, however, the effect of declining agricultural productivity moves specialization toward, rather than away from, agriculture because of the special properties of consumer preferences for food. Countries experiencing large climate change losses in agriculture that are more open to trade suffer less because they are more able to increase imports of food and shift production toward other sectors. Overall, reducing barriers to trade could decrease the losses from climate change by more than half for the poorest quartile of the world's population.

I conclude with several suggestions for future research. First, while my work is informative for cost-benefit analysis of climate change mitigation, additional effort is required to integrate these general equilibrium effects directly into calculating the social cost of carbon. Second, while my analysis shows that reducing barriers to trade is a necessary condition to induce sectoral reallocation that curtails the costs of climate change, I cannot conclude that it is sufficient. A low trade cost counterfactual in which specialization in agriculture shifts away from the equator still relies on assumptions about diminishing returns to producing tradable manufactured goods in developing countries, as well as on the availability of complementary inputs such as soil quality necessary to expand agricultural production in colder countries. A final topic concerns the political economy of trade policy regarding food. Policymakers often prioritize "food security" as a stated aim, implying a preference for domestic food production secure from interference by foreign countries. To the extent that this goal conflicts with adaptation to climate change in light of large declines in agricultural productivity in certain regions, it may be worth examining this tradeoff more closely, both in practice and in perception.

References

Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson, "Nonparametric counterfactual predictions in neoclassical models of international trade," *American Economic Review*, 2017, *107* (3), 633–89.

Alfaro, Laura and Maggie X Chen, "Selection and market reallocation: Productivity gains from multinational production," *American Economic Journal: Economic Policy*, 2018, *10* (2), 1–38.

Baqaee, David Rezza and Emmanuel Farhi, "The macroeconomic impact of microeconomic shocks: beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Baumol, William J and William G Bowen, "Performing Arts: The Economic Dilemma (New York: The Twentieth Century Fund)," 1966.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *The Review of Economic Studies*, 2016, *83* (1), 87–117.

Burnside, Craig and Martin Eichenbaum, "Factor Hoarding and the Propagation of Business Cycle Shocks," *American Economic Review*, 1996.

_ , _ , and Sergio Rebelo, "Labor hoarding and the business cycle," *Journal of Political Economy*, 1993, *101* (2), 245–273.

Burstein, Ariel and Jonathan Vogel, "International trade, technology, and the skill premium," *Journal of Political Economy*, 2017, *125* (5), 1356–1412.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural productivity and structural transformation: Evidence from Brazil," *American Economic Review*, 2016, *106* (6), 1320–65.

Carleton, Tamma A, Amir Jina, Michael T Delgado, Michael Greenstone, Trevor Houser, Solomon M Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan B Nath et al., "Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits," Technical Report, National Bureau of Economic Research 2020.

Carleton, Tamma, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E Kopp, Kelly McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Justin Simcock, Arvid Viaene, Jiacan Yuan, and Alice Zhang, "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," 2018.

Chen, Wei, Xilu Chen, Chang-Tai Hsieh, and Zheng Michael Song, "A Forensic Examination of China's National Accounts," 2019.

Cline, William R, *Global warming and agriculture: End-of-century estimates by country*, Peterson Institute, 2007.

Colmer, Jonathan, "Temperature, labor reallocation, and industrial production: Evidence from India," *American Economic Journal: Applied Economics*, 2018.

Comin, Diego A, Danial Lashkari, and Martí Mestieri, "Structural change with long-run income and price effects," *National Bureau of Economic Research Working Paper*, 2015.

Costinot, Arnaud, Dave Donaldson, and Cory Smith, "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world," *Journal of Political Economy*, 2016, *124* (1), 205–248.

Cuaresma, Jesús Crespo, "Income projections for climate change research: A framework based on human capital dynamics," *Global Environmental Change*, 2017, *42*, 226–236.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced trade," *American Economic Review*, 2007, *97* (2), 351–355.

Deryugina, Tatyana and Solomon M Hsiang, "Does the environment still matter? Daily temperature and income in the United States," *National Bureau of Economic Research Working Paper*, 2014.

Deschenes, Olivier and Michael Greenstone, "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather," *American Economic Review*, 2007, *97* (1), 354–385.

Deschênes, Olivier and Michael Greenstone, "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US," *American Economic Journal: Applied Economics*, 2011, *3* (4), 152–85.

Desmet, Klaus and Esteban Rossi-Hansberg, "On the spatial economic impact of global warming," *Journal of Urban Economics*, 2015, *88*, 16–37.

Dingel, Jonathan I, Kyle C Meng, and Solomon M Hsiang, "Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate," *National Bureau of Economic Research Working Paper*, 2019.

Driscoll, John C and Aart C Kraay, "Consistent covariance matrix estimation with spatially dependent panel data," *Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Eaton, Jonathan and Samuel Kortum, "Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Fiszbein, Martin and Will Johnson, "Agricultural Productivity, International Trade, and Structural Change," 2020.

Gollin, Douglas, Casper Worm Hansen, and Asger Wingender, "Two blades of grass: The impact of the green revolution," *National Bureau of Economic Research Working Paper*, 2018.

_ , Stephen L Parente, and Richard Rogerson, "The food problem and the evolution of international income levels," *Journal of Monetary Economics*, 2007, *54* (4), 1230–1255.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez, "Capital allocation and productivity in South Europe," *The Quarterly Journal of Economics*, 2017, *132* (4), 1915–1967.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect inference," *Journal of applied econometrics*, 1993, *8* (S1), S85–S118.

Hicks, Joan Hamory, Marieke Kleemans, Nicholas Y Li, and Edward Miguel, "Reevaluating agricultural productivity gaps with longitudinal microdata," *National Bureau of Economic Research Working Paper*, 2017.

Hulten, Charles R, "Growth accounting with intermediate inputs," *The Review of Economic Studies*, 1978, *45* (3), 511–518.

Jeffrey, Stephen, Leon Rotstayn, MA Collier, Stacey Dravitzki, Carlo Hamalainen, Chris Moeseneder, KK Wong, and Jozef Syktus, "Australia's CMIP5 submission using the CSIRO-Mk3. 6 model," *Aust. Meteor. Oceanogr. J*, 2013, *63*, 1–13.

Jones, Charles I and Peter J Klenow, "Beyond GDP? Welfare across countries and time," *American Economic Review*, 2016, *106* (9), 2426–57.

Lagakos, David and Michael E Waugh, "Selection, agriculture, and cross-country productivity differences," *American Economic Review*, 2013, *103* (2), 948–80.

Liu, Maggie, Yogita Shamdasani, and Vis Taraz, "Climate Change, Structural Transformation, and Infrastructure: Evidence from India," 2020.

Lyu, Changjiang, Kemin Wang, Frank Zhang, and Xin Zhang, "GDP management to meet or beat growth targets," *Journal of Accounting and Economics*, 2018, *66* (1), 318–338.

Matsuyama, Kiminori, "Agricultural productivity, comparative advantage, and economic growth," *Journal of Economic Theory*, 1992, *58* (2), 317–334.

Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw, "The impact of global warming on agriculture: a Ricardian analysis," *The American Economic Review*, 1994, pp. 753–771.

Nordhaus, William D, "An optimal transition path for controlling greenhouse gases," *Science*, 1992, *258* (5086), 1315–1319.

Schlenker, Wolfram and David B Lobell, "Robust negative impacts of climate change on African agriculture," *Environmental Research Letters*, 2010, *5* (1), 014010.

_ and Michael J Roberts, "Nonlinear temperature effects indicate severe damages to US crop yields under climate change," *Proceedings of the National Academy of sciences*, 2009, *106* (37), 15594–15598.

Seppanen, Olli, William J Fisk, and QH Lei, "Effect of temperature on task performance in office environment," 2006.

Somanathan, E, Rohini Somanathan, Anant Sudarshan, Meenu Tewari et al., "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing," 2015.

Teignier, Marc, "The role of trade in structural transformation," *Journal of Development Economics*, 2018, *130*, 45–65.

Tombe, Trevor, "The missing food problem: Trade, agriculture, and international productivity differences," *American Economic Journal: Macroeconomics*, 2015, *7* (3), 226–58.

Törnqvist, Leo, "The Bank of Finland's consumption price index," 1936.

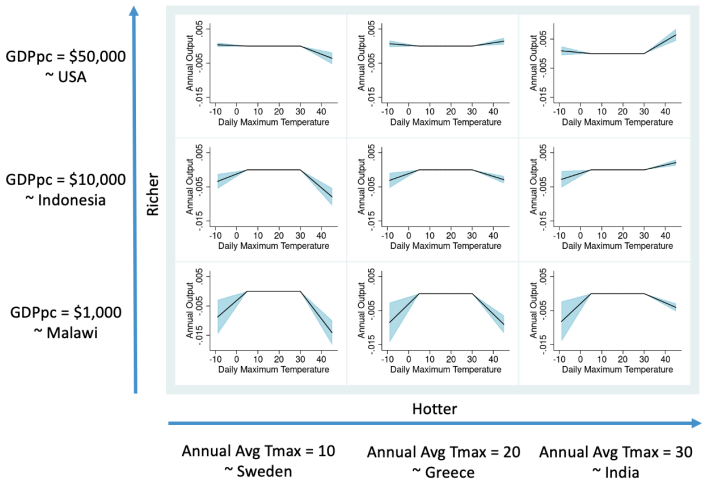
Uy, Timothy, Kei-Mu Yi, and Jing Zhang, "Structural change in an open economy," *Journal of Monetary Economics*, 2013, *60* (6), 667–682.

Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang, "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," *Journal of Environmental Economics and Management*, 2018, *88*, 1–17.

Zivin, Joshua Graff and Matthew Neidell, "Temperature and the allocation of time: Implications for climate change," *Journal of Labor Economics*, 2014, *32* (1), 1–26.

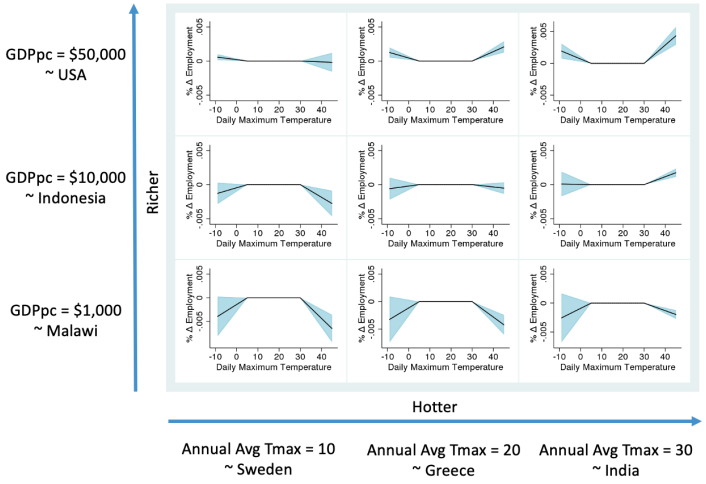
Appendix A: Additional Regression Results

Figure A-1: Predicted Heterogeneous Response of Annual Manufacturing Revenue to Daily Maximum Temperature



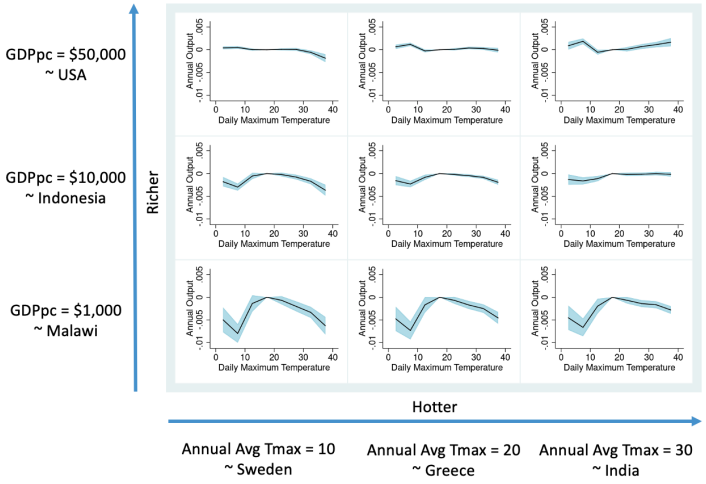
Notes: Figure shows the predicted effect of temperature on the log of manufacturing revenues at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 3 of Table 2. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-2: Predicted Heterogeneous Response of Annual Manufacturing Employment to Daily Maximum Temperature



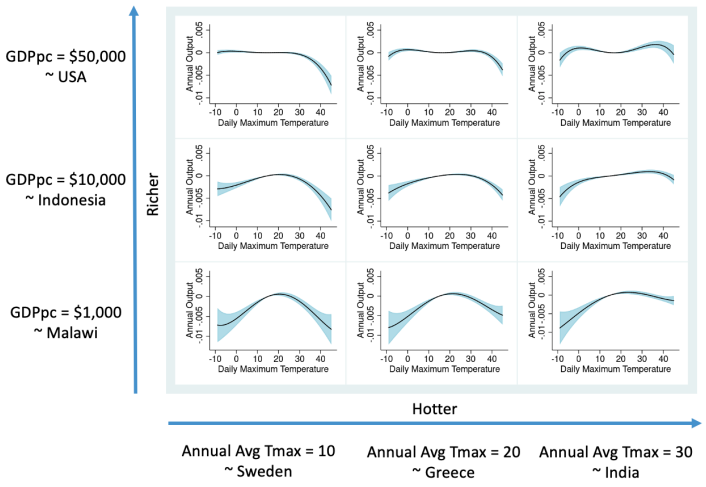
Notes: Figure shows the predicted effect of temperature on the log of manufacturing employment at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 4 of Table 2. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-3: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Bins of Daily Maximum Temperature



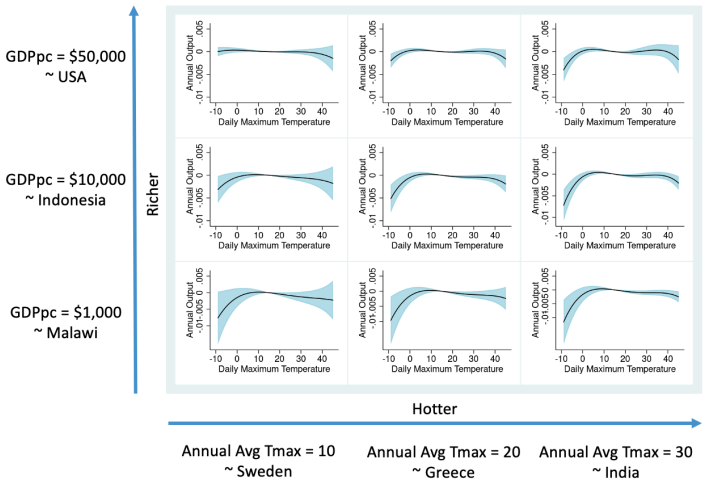
Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using bins of daily maximum temperature in the specification from Equation 8. Days are grouped into 5°C bins. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-4: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Polynomial of Degree 4 of Daily Maximum Temperature



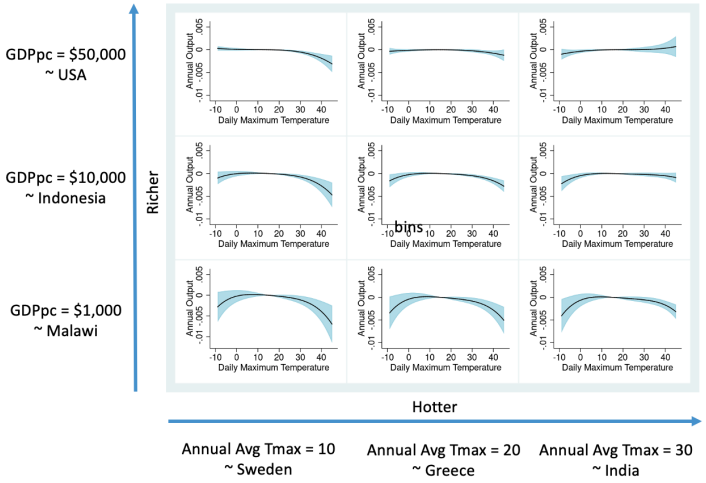
Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using a polynomial of degree four in daily average temperature in the specification from Equation 8. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-5: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE



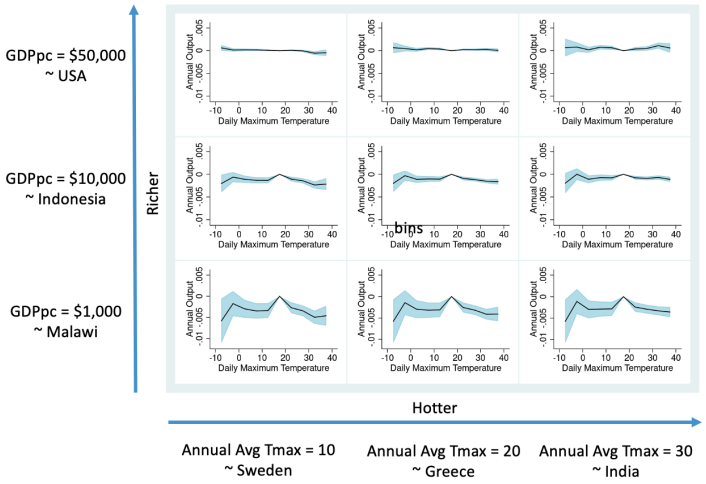
Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-6: Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE



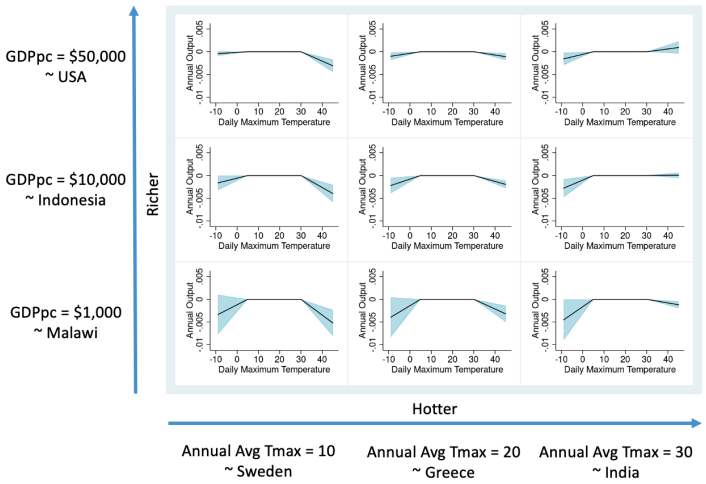
Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-7: Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE



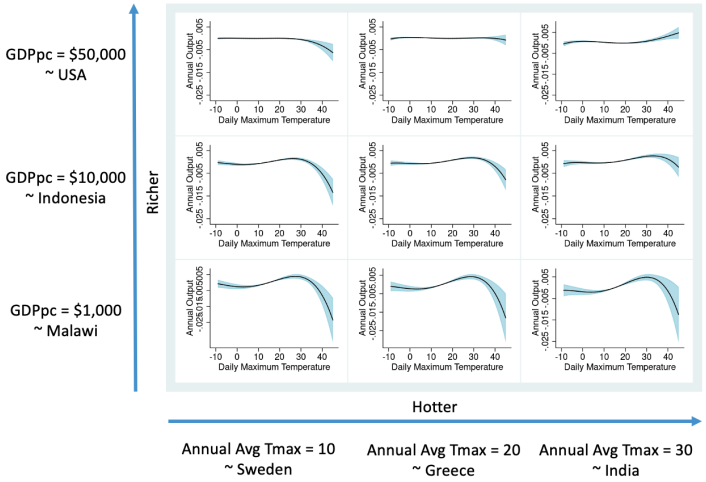
Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and bins of daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-8: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - Controls for Capital



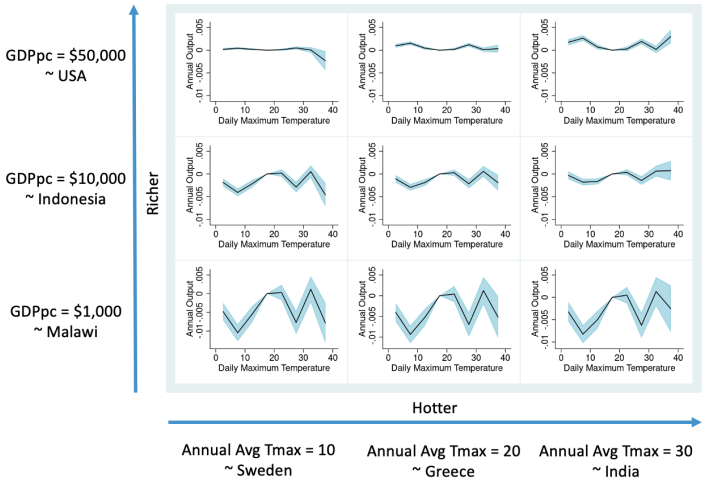
Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with controls for capital. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-9: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Polynomial of Degree 4 of Daily Maximum Temperature



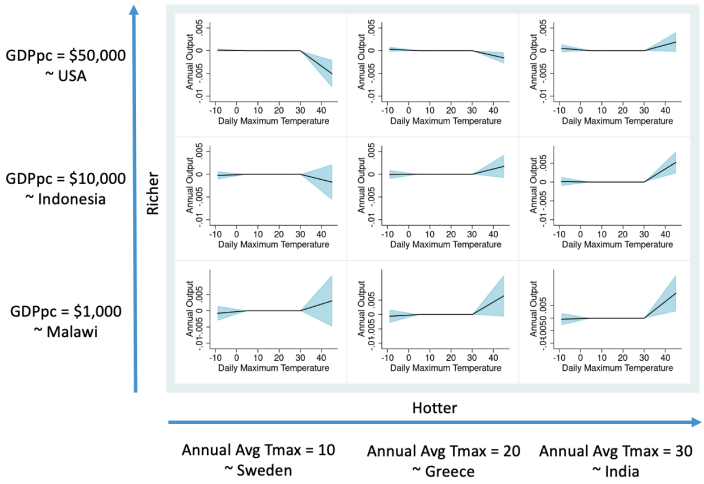
Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-10: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Bins of Daily Maximum Temperature



Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with bins of daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-11: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature



Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

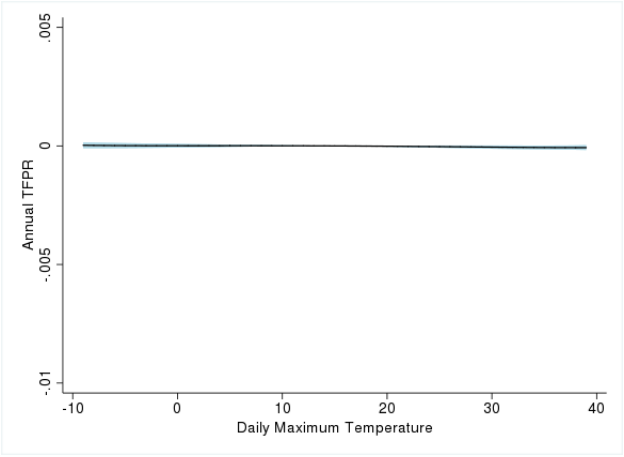
Table A-1: Country-Level Panel Regression

	(1)	(2)	(3)	(4)
	log(GDP)	Food Share of Imports	Ag Share of GDP	Ag Labor Share
KDD X 100	-0.0223	0.00638	0.0165	0.00483
	(-0.55)	(1.80)	(3.92)	(3.14)
GDD X 100	0.00251	-0.00191	-0.00165	-0.00113
	(0.44)	(-2.87)	(-1.53)	(-1.74)
Observations	7561	5775	5522	3718
Country FE	X	X	X	X
Year FE	X	X	X	X
Ag Labor Weights				

Notes: t-statistics in parentheses. Reported Driscoll and Kraay (1998) standard errors are robust to heteroskedasticity, spatial correlation, and autocorrelation of up to 5 lags. Results come from estimating Equation 27 with crop-area weighted growing and killing degree days. Data covers 164 countries from 1960-2012 with varying coverage by country and outcome variable. Economic data from all sources above are retrieved from the World Bank Databank.

Appendix B: U.S. Results

Figure A-12: Estimated Response of U.S. Annual Manufacturing TFPR to Daily Maximum Temperature



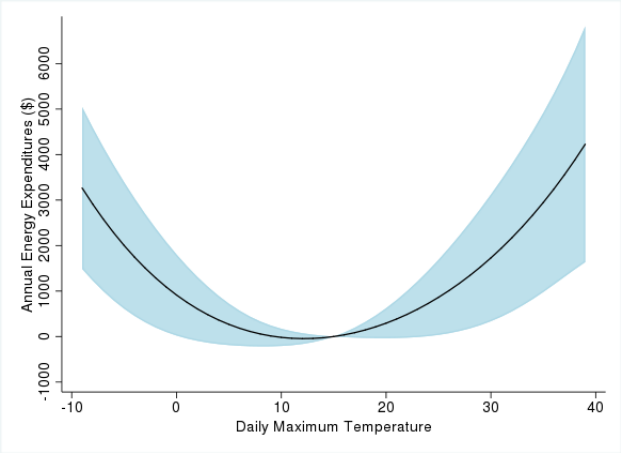
Notes: Figure shows the estimated effect of temperature on manufacturing TFPR using the specification from Equation 7 with a polynomial of degree four in daily maximum temperature. 95% confidence interval is shown in blue. Outcome data comes from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD.

Table A-2: U.S. Productivity Results

	Revenue/Worker	Revenue	Employment	TFPR	Revenue/Worker	Revenue/Worker
TMax-30	-0.0000109	0.0000220	0.0000330	0.00000134	-0.0000422	0.0000110
	(-2.21)	(2.01)	(3.49)	(0.33)	(-2.97)	(0.46)
5-TMax	0.0000365	0.0000338	-0.00000269	-0.00000685	-0.0000226	0.000154
	(5.65)	(2.65)	(-0.26)	(-1.30)	(-1.71)	(3.56)
Observations	2852000	2852000	2852000	2852000	2852000	2852000
Firm FE	X	X	X	X	X	X
Country X Year FE	X	X	X	X	X	X
State X Year FE					X	
Sales Weighting						X

Notes: t-statistics in parentheses. Dependent variables all in logs. Standard errors two-way clustered at the firm and county-by-year level. Estimates use the regression model from Equation 7 with outcome variable data from 1976-2014 from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau and temperature data from GMFD.

Figure A-13: Estimated Response of U.S. Annual Manufacturing Plant-Level Energy Expenditures to Daily Maximum Temperature



Notes: Figure shows the estimated effect of temperature on manufacturing energy expenditures using the specification from Equation 7 with a polynomial of degree four in daily maximum temperature. 95% confidence interval is shown in blue. Energy expenditures are the sum of cost of fuels and electricity expenditures in the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD.

Table A-3: U.S. Energy Results

	log(Energy Expenditure)	Energy Expenditures	log(Energy Expenditures)	Energy Expenditures
TMax-30	0.0000822	0.0000890	251.1	6056
	(6.03)	(3.24)	(4.45)	(1.32)
5-TMax	0.0000108	0.00000184	490.8	13840
	(0.78)	(0.04)	(3.57)	(1.69)
Observations	2852000	2852000	2852000	2852000
Firm FE	X	X	X	X
Country X Year FE	X	X	X	X
Sales Weighting		X		X

Notes: t-statistics in parentheses. Standard errors two-way clustered at the firm and county-by-year level. Estimates use the regression model from Equation 7 with outcome variable data from 1976-2014 from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau and temperature data from GMFD. Dependent variable is the sum of electricity expenditures and cost of fuels, in logs or levels.

Appendix C: China Results

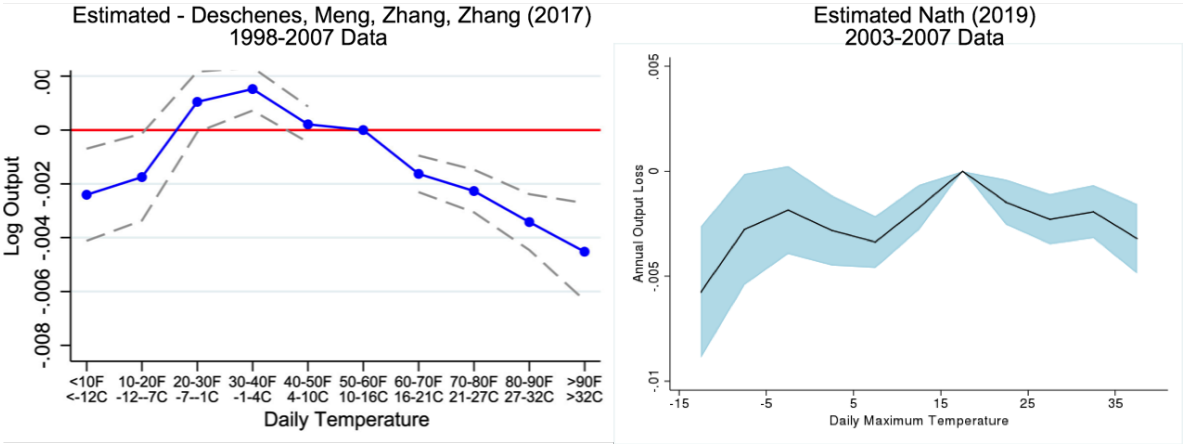
This section explains the data quality issues that lead me to estimate the results in Section 4.1 excluding data from China. At a high level, I find evidence consistent with the conclusions of Chen, Chen, Hsieh and Song (2019) that Chinese micro-data after 2007 are unreliable due to systematic manipulation by local officials. The details are as follows.

To start with, Zhang, Deschenes, Meng and Zhang (2018) analyze data from China for the years 1998-2007 and find that both cold and hot temperatures harm output and productivity, consistent with my findings. Using the overlapping subset of years from my data, which goes from 2003-2012, I am able to replicate their findings fairly closely, as shown in Appendix Figure A-14. Notably, I am also able to use my main results from the rest of my global data in Figure 3 to closely predict the response of output to temperature in China based on their income level and average climate. My prediction and the estimates from Zhang, Deschenes, Meng and Zhang (2018) are shown in Figure A-15. While I slightly overpredict sensitivity to cold and underpredict sensitivity to heat, my results are broadly consistent with their findings, lending external validity to my work.

However, when I estimate the response to temperature in my full sample of Chinese firms from 2003-2012, I produce the highly anomalous results shown in Figure A-16. This estimate using my full sample of Chinese data implies that extreme temperatures sharply and statistically significantly *increase* output, a finding inconsistent with my results from any other country in the world. Notably, this anomalous result begins to appear by including later years starting with 2008 in the regression, the same year Chen, Chen, Hsieh and Song (2019) start to find discrepancies in the data. They state that "local statistics increasingly misrepresent the true numbers after 2008" and "the micro-data of the ASIF [have] overstated aggregate output."

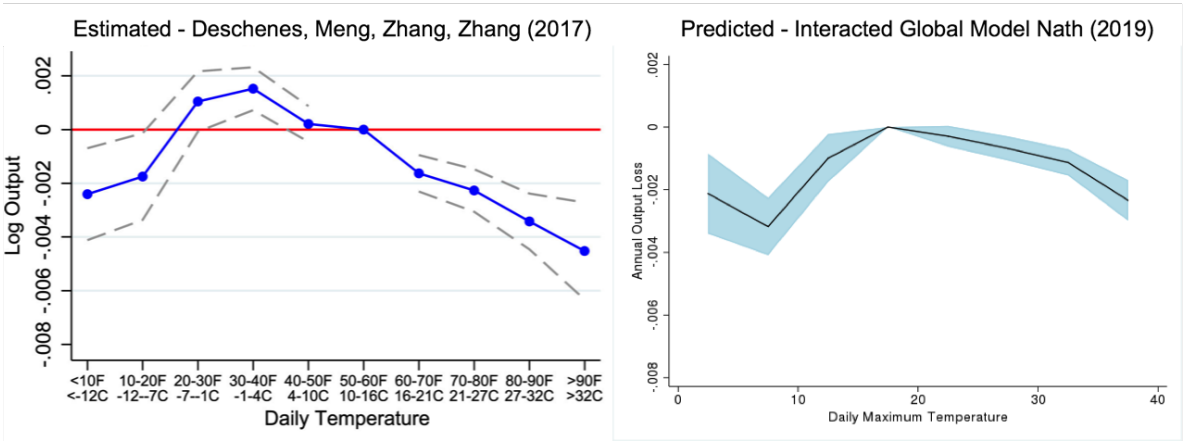
A somewhat puzzling fact is that my results suggest that this documented manipulation of data in China is systematically correlated with temperature. One plausible hypothesis is that Chinese provincial officials inflate reported manufacturing output to meet GDP targets in response to declines in other sectors more susceptible to temperature, such as agriculture. These targets have historically played a central role in the evaluation and promotion of government officials, and Lyu, Wang, Zhang and Zhang (2018) demonstrate that reported provincial GDP just barely hits target thresholds with implausible frequency. I cannot provide further evidence on the particular sources and methods of manipulation, but given the widespread external documentation of problems with this subset of the Chinese firm data and my very short panel that would remain when excluding these years in China, I exclude this dataset entirely from my main analysis. Still, I view the consistency of both my replication and predictions with the results of Zhang, Deschenes, Meng and Zhang (2018) as validating my central analysis.

Figure A-14: China Replication - Overlapping Years



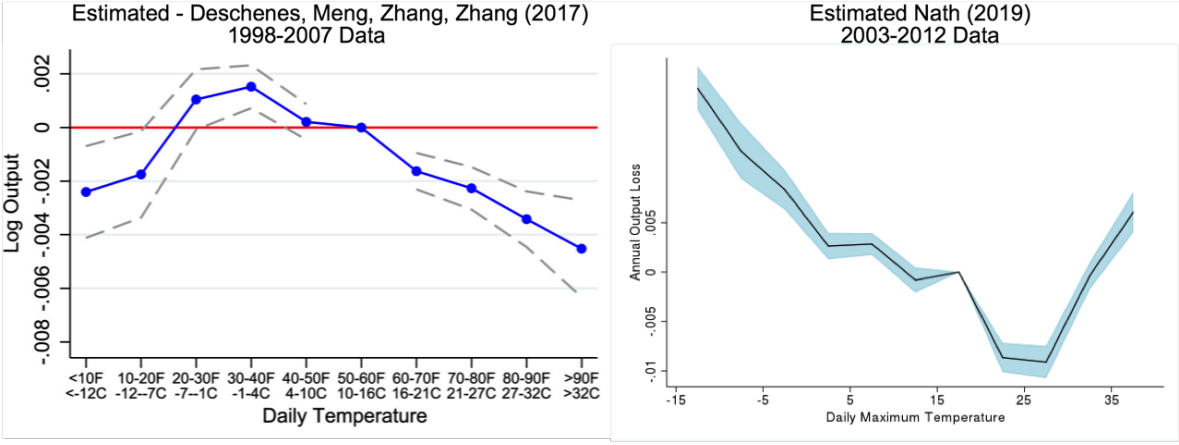
Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result using data from the same source for 2003-2007 - the overlapping years of my data coverage. Temperature data is from GMFD.

Figure A-15: China Manufacturing Temperature Sensitivity - Estimated and Predicted



Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows the predicted effect of temperature in China from evaluating my global interacted specification from Column 2 of Table 2 at China's income and average long-run temperature from 1998-2007. I do not use any data from China in my estimation or prediction, but replicate the pattern closely.

Figure A-16: China Replication - Different Years



Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result using data from the same dataset for 2003-2012 - the years of my data coverage. Temperature data is from GMFD.

Appendix D: Adaptation Benefits and Costs

In this section I explain how I use revealed preference methods developed by Carleton et al. (2018) to infer the costs firms incur from reducing the sensitivity of their production to extreme temperatures. To build intuition start by considering a simple example of otherwise identical firms in two cities, Seattle and Houston. Houston is hotter than Seattle, but Seattle heats up over the course of the century such that its exposure to CDD in 2100 is that of Houston in 2020. Let β represent lost annual revenues from exposure to a cooling degree day, a function of the adaptation investments the firm chooses to make. The annual costs of extreme heat to a firm in Seattle are given by $CDD_{Seattle} * \beta_{Seattle}$. Since Seattle suffers little exposure to extreme heat, its firms choose a lower (more negative) β than firms in Houston, as I find in my empirical estimates. If Seattle firms had chosen the Houston β associated with greater expected exposure to heat, the marginal benefits they would obtain are as follows:

$$MB = CDD_{Seattle} * (\beta_{Houston} - \beta_{Seattle})$$

Given that Seattle firms do not choose $\beta_{Houston}$, we know that the marginal costs of this incremental reduction in temperature sensitivity must exceed the marginal benefits. By repeating this logic for the firm's estimated temperature sensitivity for every year of warming from $Seattle_{2020}$ to $Seattle_{2100}$, we can construct the full marginal cost curve for the Seattle firm's projected change in chosen β from 2020

to 2100:

$$TC = \sum_{t=2020}^{2099} MC_t = \sum_{t=2020}^{2099} CDD_t * (\beta_{t+1} - \beta_t) \tag{28}$$

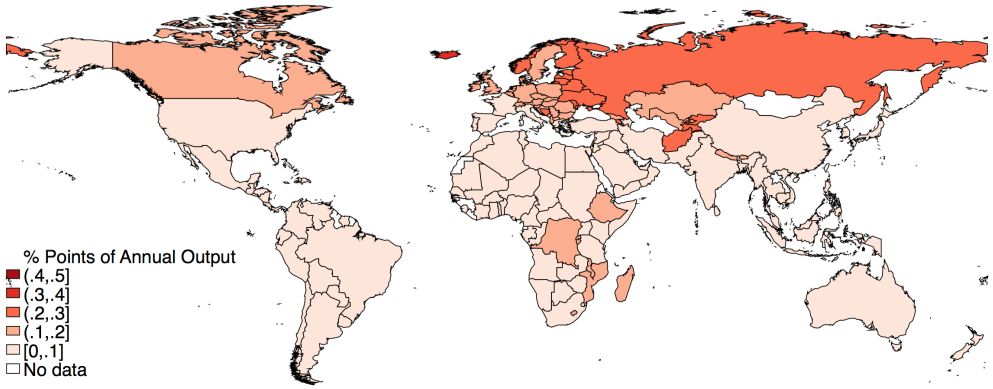
Note that the continuous version of Equation 28 also follows straight from the firm's first-order condition in the framework in Section 3.1. The firm's lost revenues from extreme heat are $CDD * \beta$ so the marginal benefit the firm receives from a reduction in β is given by CDD. Since the firm's optimal choice of β equates marginal benefit to marginal cost, we have marginal cost $c_\beta = CDD$ for the full range of CDDs.

The total benefits of future adaptation for firms in Seattle are given by the change in damages from choosing their optimal level of adaptation for expected heat exposure in 2100 rather than remaining at the adaptation level they choose in 2020:

$$TB = CDD_{2100} * (\beta_{2100} - \beta_{2020}) \tag{29}$$

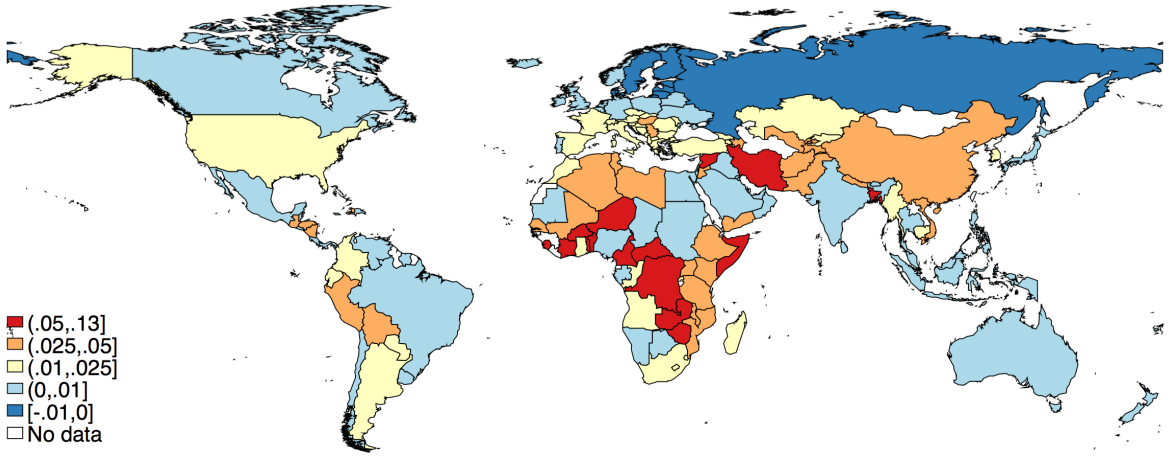
Because CDDs are increasing as countries become hotter, the benefits of adaptation in Equation 29 exceed the costs in Equation 28. Figure A-17 shows predicted manufacturing sensitivity to a hot day at end-of-century temperatures, which is substantially muted relative to the sensitivities at current temperatures shown in Figure 5a. Figure A-18 show the costs of achieving this reduced sensitivity, as calculated using Equation 28, and Figure A-19 show the net benefits of firms adapting to changes in expected exposure to extreme heat.

Figure A-17: Predicted Effect of a 40°C Day on Annual Manufacturing Revenue per Worker At 2080 Average Temperatures



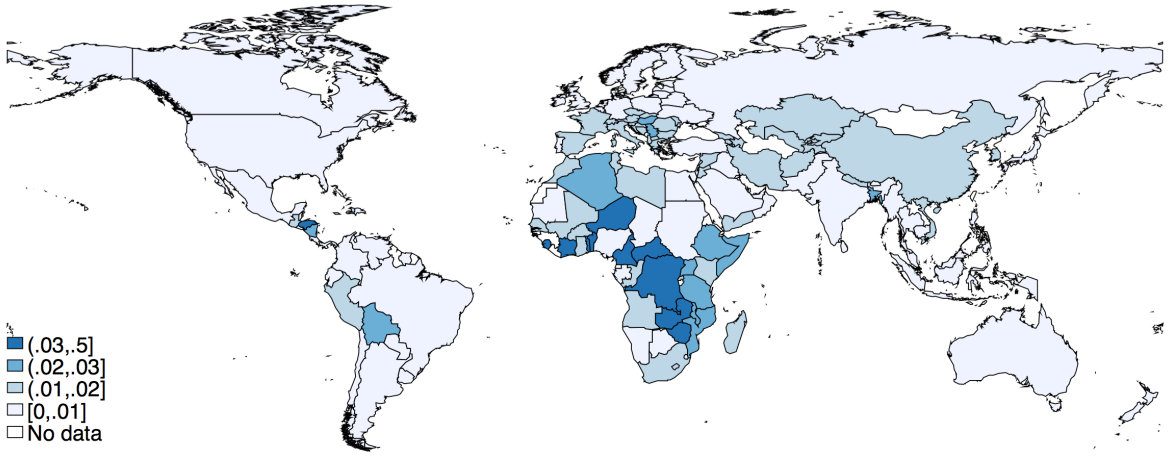
Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's level of income and end-of-century long-run average temperature. Temperature sensitivities are lower in this figure than in Figure 5a because my results predict that firms will adapt to hot temperatures as the world warms.

Figure A-18: Firm-Level Adaptation Costs
(Share of Manufacturing Output)



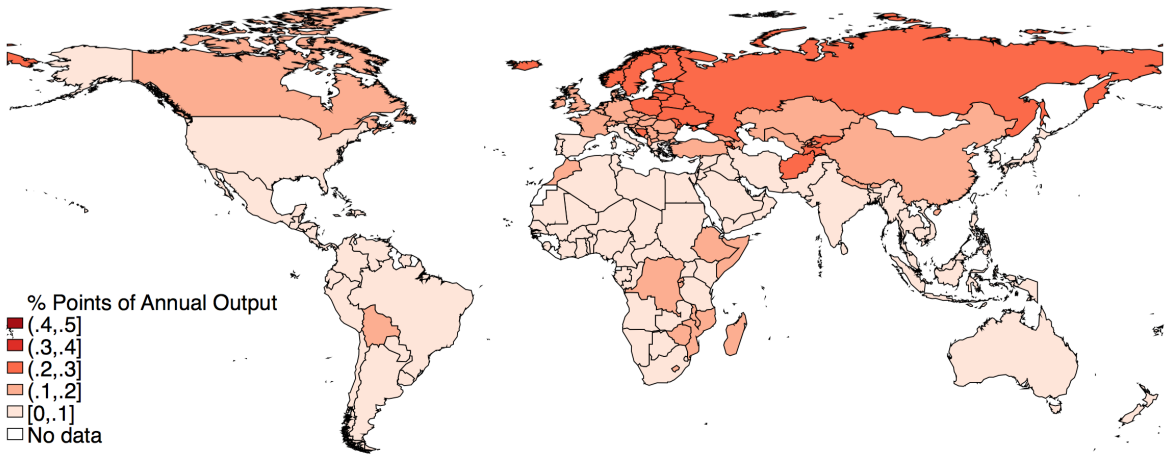
Notes: Map shows my calculations of the costs firms pay to achieve the lower temperature sensitivity shown in Appendix Figure A-17 compared to Figure 5a. I infer these costs using a revealed preference approach developed by Carleton et al. (2018) that infers adaptation costs from the foregone benefits firms would have attained by reducing their heat sensitivity. The procedure is detailed in Appendix D.

Figure A-19: Firm-Level Adaptation Net Benefits
(Share of Manufacturing Output)



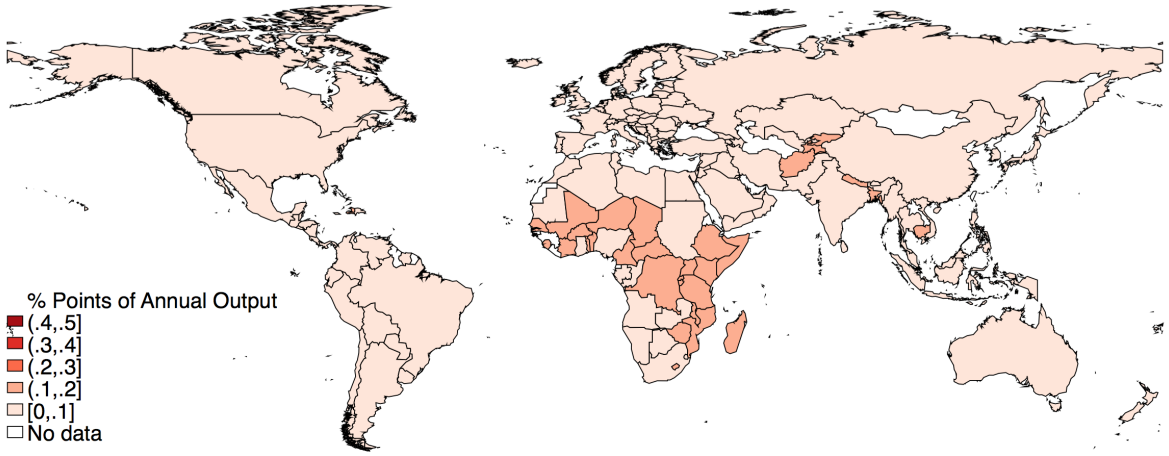
Notes: Map shows my calculations of the net benefits firms receive by investing to reduce their heat sensitivity as the climate warms. The benefits come from reducing heat sensitivity to the level shown in Appendix Figure A-17 compared to the original level in Figure 5a. The inferred costs are shown in Appendix Figure A-18. The procedure to calculate these costs and benefits is detailed in Appendix D.

Figure A-20: Predicted Effect of a 40°C Day on Annual Services Revenue per Worker



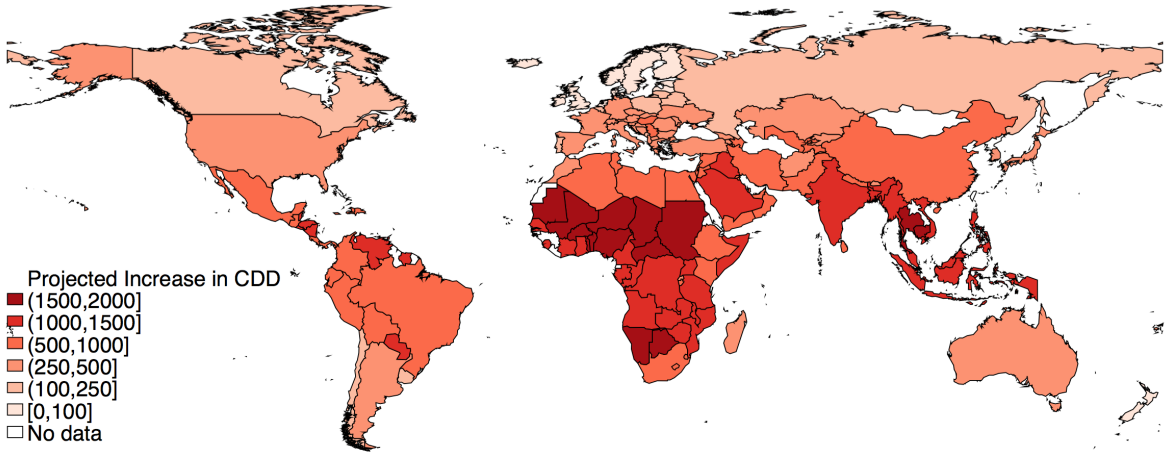
Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-21: Predicted Effect of a -5°C Day on Annual Services Revenue per Worker



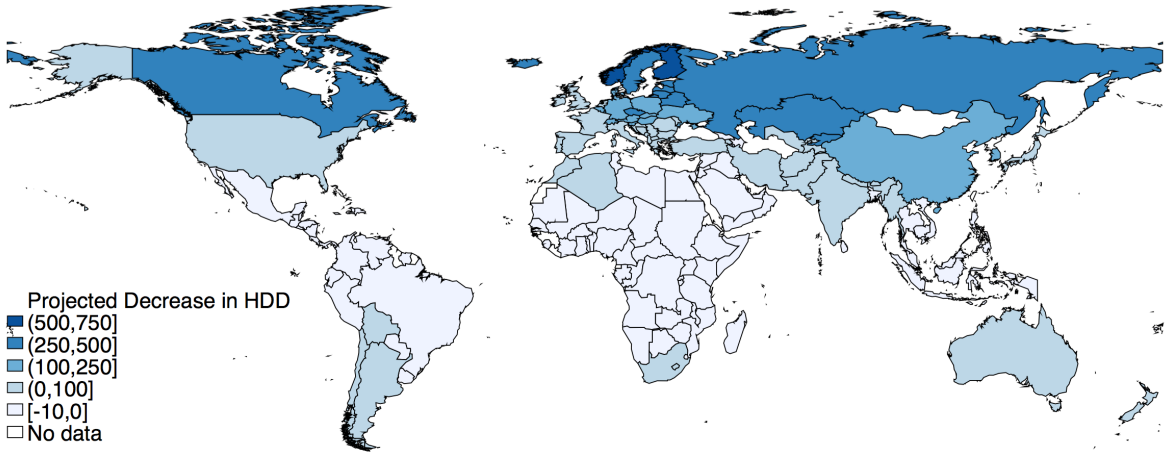
Notes: Map shows the predicted annual percentage point loss in revenue per worker from a -5°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-22: Projected Change in Exposure to Extreme Heat



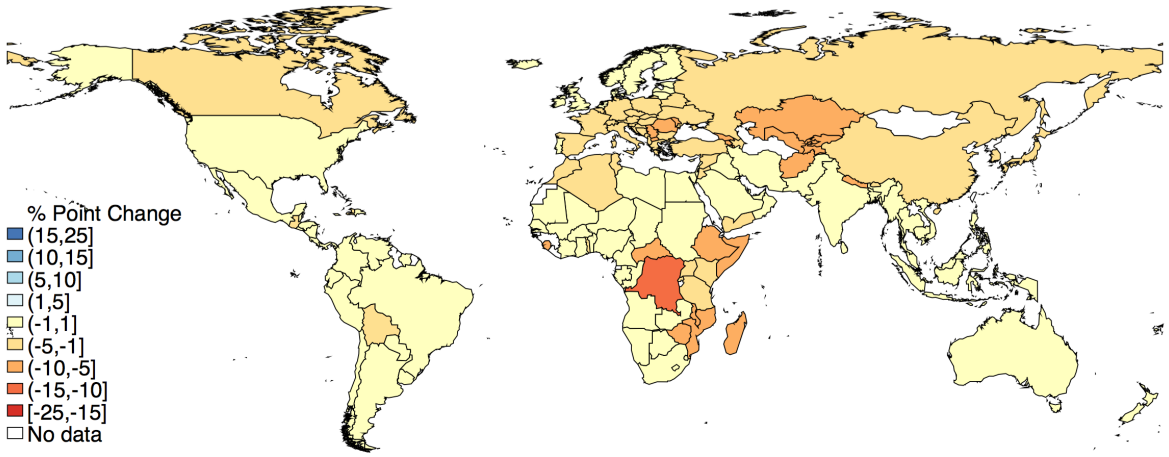
Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme heat as measured by the change in cooling degree days above 30°C from 2015 to 2080-2099.

Figure A-23: Projected Change in Exposure to Extreme Cold



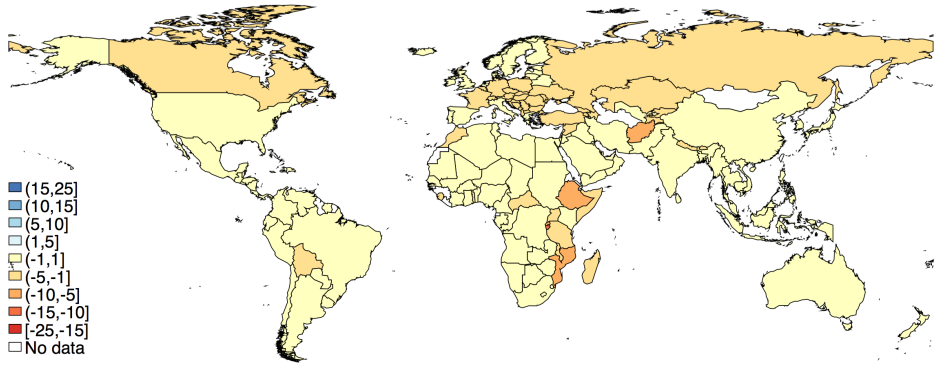
Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme cold as measured by the change in heating degree days below 5°C from 2015 to 2080-2099.

Figure A-24: Projected Impact of Climate Change on Services Productivity



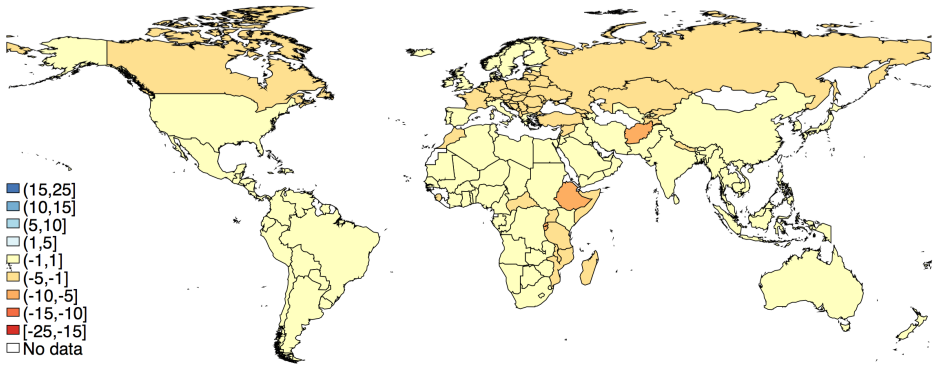
Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 5 of Table 2 at each country's income and end-of-century long-run average temperature.

Figure A-25: Projected Impact of Climate Change on Manufacturing Productivity Accounting for Economic Growth and Adaptation



Notes: Map shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 2 of Table 2 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017). These estimates that account for economic growth show reduced losses relative to those in Figure 8a because my empirical results suggest that firms in richer countries have reduced exposure to extreme temperatures.

Figure A-26: Projected Impact of Climate Change on Services Productivity Accounting for Economic Growth and Adaptation



Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 5 of Table 2 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017). These estimates that account for economic growth show reduced losses relative to those in Appendix Figure A-24 because my empirical estimates suggest that firms in richer countries have reduced exposure to extreme temperatures.

Appendix E: Model Estimation

Appendix E.1: Solution Algorithm & Simulated Method of Moments

I solve the model presented in Section 5 numerically by directly simulating each component. I draw productivities, z_{ijk}, from Frechet distributions for 20,000 varieties for each sector for each country. I assign production of each variety in each country to the world's lowest cost producer based on wages, trade costs, and the productivity draw z_{ijk}. This allows me to calculate prices, revenues, and bilateral imports and exports for each sector and country. I start by guessing a vector of wages across all countries and then iterate until the equilibrium conditions hold and national income equals national expenditures for each country.

I use a combination of calibration and estimation to set the model parameters. I set the trade elasticities to the values estimated by Tombe (2015); $\theta_a = 4.06$, and $\theta_m = 4.63$. I calibrate the relative levels of Z_{jk} to match relative value-added per worker in agriculture, manufacturing, and services, and the overall level of $\{Z_{ak}, Z_{mk}, Z_{sk}\}$ to match country-level nominal GDP.[52] I estimate the consumption parameters to minimize the sum of squared distance from sectoral share data, and choose bilateral trade costs to match the data on bilateral trade flows by sector.

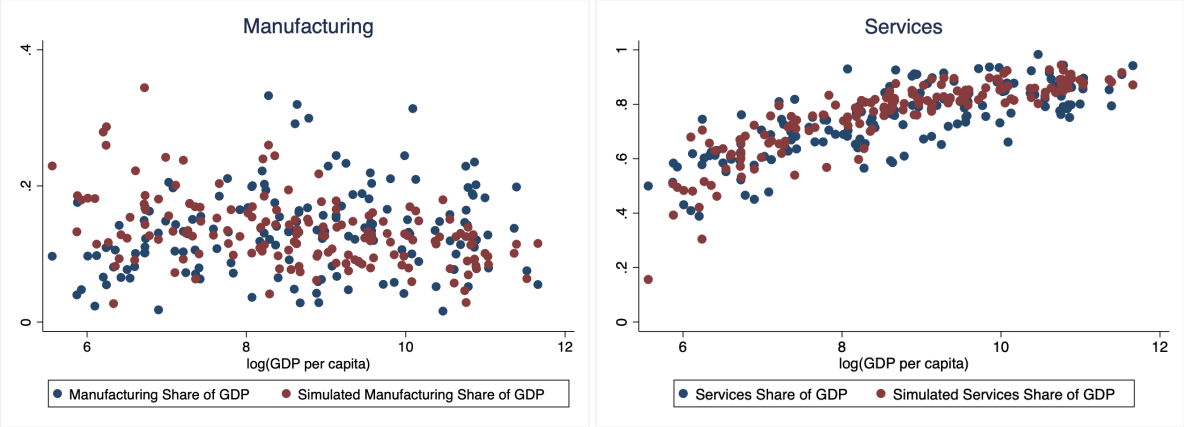
For the trade moments, I obtain data from UN Comtrade and classify HS 1988/92

[52] Since trade flows are in nominal terms, I match nominal GDP in the model for consistency. The nonhomothetic price index deflates nominal income to a measure of welfare.

codes 1-24 as agriculture and 28-97 as manufacturing to best approximate food and non-food imports. Since trade data is reported in gross output terms but GDP is in value-added, I follow Tombe (2015) and deflate the trade data by country-sector-level value-added to output ratios obtained from the United Nations Statistical Division. Following recommendations from UN Comtrade documentation, I use importer-reported trade data where possible, but default to exporter-reported data for smaller developing countries with large discrepancies between importer and exporter reported data.

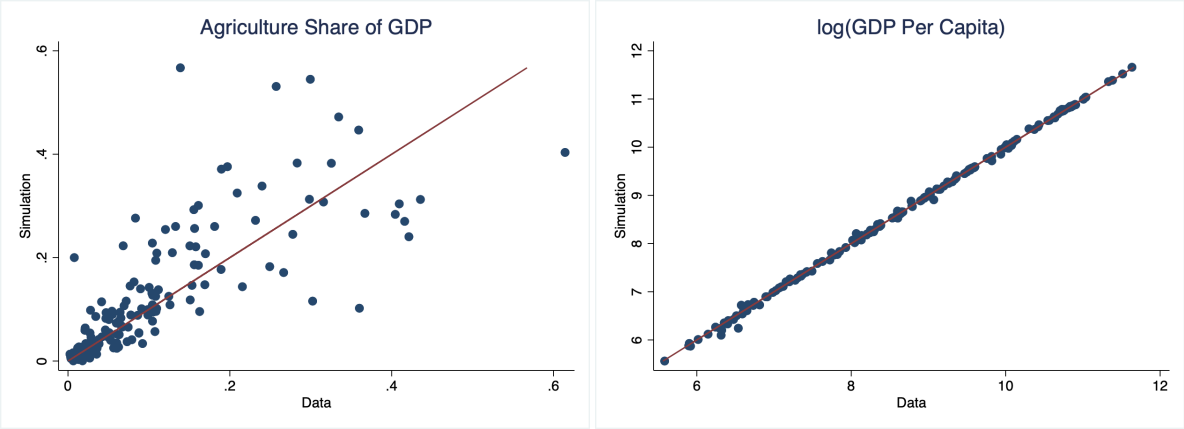
Appendix E.2: Additional Model Fit Figures

Figure A-27: Sectoral GDP Shares - Data vs. Simulation



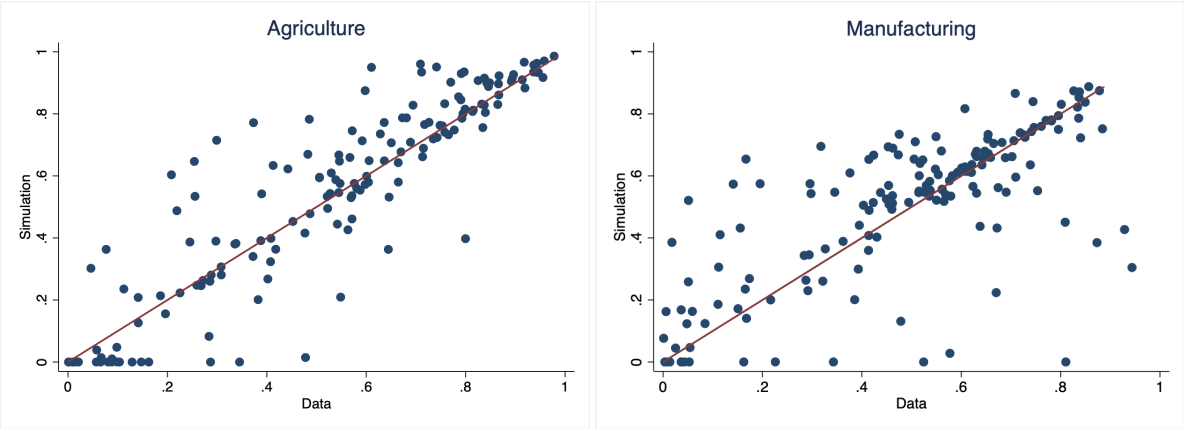
Notes: Left graph shows the fit of simulated manufacturing share of GDP in the model to data from the World Bank. Right graph shows the same comparison for services share of GDP.

Figure A-28: GDP Per Capita and Agriculture Share of GDP - Data vs. Simulation



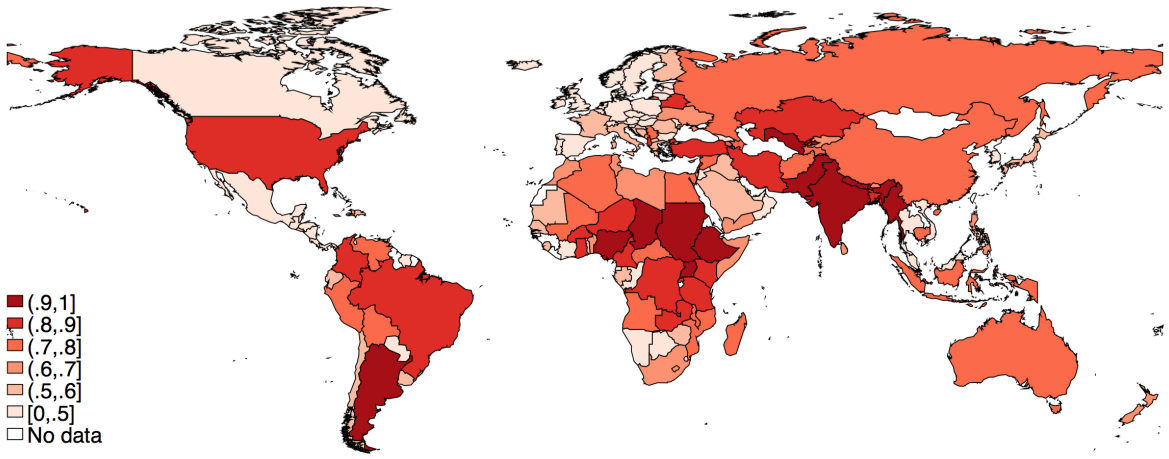
Notes: Left graph shows another view of the fit of simulated agriculture share of GDP in the model to data from the World Bank also shown in Figure 6. Right graph shows the same comparison for GDP per capita. A perfect fit would have all data points be on the 45°line such that the simulated and actual values are equal. The simulation explains over 60% of the variation in the agriculture share of GDP and over 99% of the variation in per capita income.

Figure A-29: Domestic Production Share of Expenditures - Data vs. Simulation



Notes: Graph shows the fit of simulated domestic production share of agricultural (left) and manufacturing (right) consumption in the model to data from Comtrade. As shown in Section 5.2, openness to food imports is a crucial parameter governing the response of labor reallocation to climate change. The simulation explains over 80% of the variation in the data for this moment.

Figure A-30: Domestic Production Share of Expenditures in Agriculture - Model Simulation



Notes: Figure shows that the share of expenditures on domestically produced goods in agriculture is very high in many developing countries with high barriers to trade. Table 5 shows that these simulated values track closely to the data.

Appendix E.3: Cline Productivity Estimates

The analysis in Cline (2007) uses micro-data from 18 countries in Africa, North and South America, and Asia representing over 35% of the world's agricultural production to estimate Ricardian cross-sectional regressions of agricultural output (in dollars) from grains, fruits, vegetables, and livestock as a function of temperature, precipitation, and irrigation. Because we expect farmers to optimize crop choice and land use decisions in response to local long-run climate conditions, I interpret the estimated effects of temperature and precipitation from these cross-sectional regressions as net of adaptation through choice of crops and livestock. Projections using the empirical estimates are averaged with projections from leading crop models from agronomy, which also account for adaptation through crop-switching and adjusted farming techniques. The crop model projections in Cline (2007) account for reallocation across crop types within country, shifting planting dates, and increased irrigation and fertilizer use. None of the estimates in the analysis account for any response of international trade.

Appendix F: Model Robustness

In this section, I evaluate the robustness of the counterfactual model simulations presented in Section 7 to three sets of different assumptions - an alternative specification for nonhomothetic consumer preferences, an alternative functional form to represent sector-country level productivity distributions, and a version of the model that allows for heterogeneous workers in each country.

Appendix F.1: Stone-Geary Preferences

I test that the model predictions are robust to the way nonhomothetic consumer preferences are specified by estimating a version of the model in which the representative agent in country k has the following generalized Stone-Geary preferences over the sectoral final goods in agriculture, manufacturing, and services:[53]

$$U(C_{ka}, C_{km}, C_{ks}) = \left(\omega_a^{\frac{1}{\sigma}} (C_{ka} - \overline{C_{ka}})^{\frac{\sigma-1}{\sigma}} + \omega_m^{\frac{1}{\sigma}} (C_{km} - \overline{C_{km}})^{\frac{\sigma-1}{\sigma}} + \omega_s^{\frac{1}{\sigma}} (C_{ks} - \overline{C_{ks}})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{30}$$

This specification is ubiquitous in the literature on structural transformation and has the advantage of intuitively capturing subsistence requirements for food by specifying a level of consumption below which people cannot survive. However, the model fit to the data is much weaker with Stone-Geary preferences than with the primary nonhomothetic CES specification, as shown in Figure A-31.

Table A-4 shows that the results in this version of the model are very similar to the baseline specification. For the poorest quartile of the global population, climate change increases agriculture's share of the labor force by 2.8 percentage points, reduces GDP by 10.7 percentage points, reduces welfare (as captured by willingness-to-pay) by 7 percentage points, and raises food prices by 37%. These results are very similar to the results in the baseline specification.

Appendix F.2: Lognormal Productivity Distributions

I estimate a version of the model with lognormal, rather than Frechet, sector-country productivity distributions to test robustness to functional form. In this specification the productivity draw, z_{ijk}, received by each intermediate goods producer is drawn from a lognormal distribution with sector-specific variance parameter φ_j and sector-country specific mean parameter Z_{jk}:

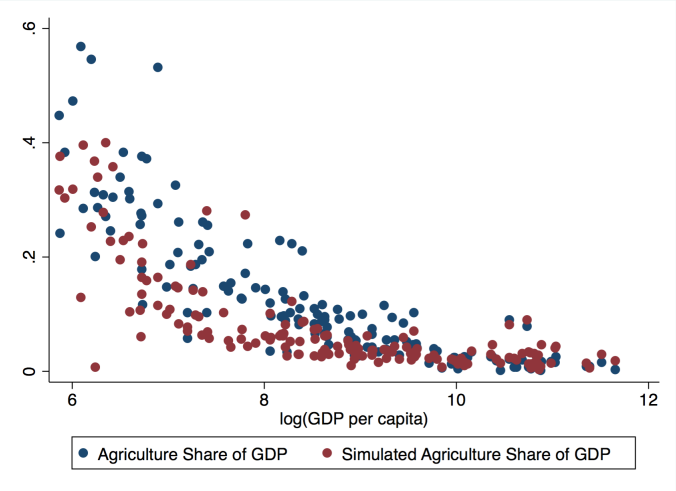
$$z_{ijk} \sim F_{jk} \text{ where } F_{jk}(z_i) = \Phi\left(\frac{(ln(x) - Z)}{\varphi}\right)$$

I estimate φ_j to match the standard deviation of the productivity distributions in the Frechet case, which yields estimates of φ_a = 0.337 and $\varphi_m = 0.398$. I estimate Z_{jk} to match both the ratio of value-added per worker across sectors and the overall level of value-added per worker in each country, as in the Frechet case.

Table A-4 shows that the results with lognormal productivity distributions are very similar to the baseline specification, with slightly larger declines in GDP and

[53] The consumption parameter estimates from applying the simulated method of moments procedure to this version of the model are $\sigma = 0.89$, $\omega_a = 0.020$, $\omega_m = 0.141$, $\omega_s = 0.839$, $\overline{C_a} = 75.5$.

Figure A-31: Agriculture Share of GDP - Data vs. Simulation
Stone-Geary Specification



Notes: Graph shows the fit of simulated agriculture share of GDP to data from the World Bank with an alternative model specification using Stone-Geary preferences over sectoral consumption. The best fit with Stone-Geary preferences has an R^2 of only 0.43 and dramatically underpredicts the agriculture share in middle-income countries especially. In contrast, the chosen nonhomothetic CES preferences from Comin, Lashkari and Mestieri (2015) explain over 60% of the variation.

Table A-4: Climate Change Counterfactual Results Summary
Alternative Model Assumptions

Country	Δ Ag Labor Share	Δ GDP	Willingness to Pay	Δ Food Prices
Baseline				
World	.005	-.021	-.017	.223
Poorest Quartile	.028	-.126	-.088	.377
Lognormal Productivity				
World	.005	-.023	-.018	.209
Poorest Quartile	.022	-.131	-.09	.338
Stone-Geary Preferences				
World	.003	-.018	-.015	.219
Poorest Quartile	.028	-.107	-.07	.371

welfare, and slightly smaller changes in food prices and agricultural labor shares in low-income countries.

Appendix F.3: Heterogeneous Workers

The baseline model makes the assumption that each country contains a population of representative agents that each receive the same wage. In practice, we observe that wages differ substantially across sectors. Agricultural workers have lower wages than non-agricultural workers in most parts of the world, and especially so in poor countries.

In this section, I specify a version of the model with heterogeneous worker skill levels and explore how this extension affects the primary comparative statics of interest in the paper.[54] While an alternative model specification with adjustment costs that impede moving across sectors could also replicate the pattern in the macro data, recent empirical evidence points to worker heterogeneity as the central force underlying sectoral wage differences. In particular, Hicks, Kleemans, Li and Miguel (2017) find that workers experience only small gains in wages by moving from agriculture to non-agriculture when controlling for individual-level fixed effects. This suggests that low wages in agriculture stem from the different characteristics of the people working in that sector, rather than from barriers that prevent them from realizing large productivity and wage gains from a potential move into non-agricultural sectors.

In the version of the model with worker heterogeneity I start by assuming that each country has a fixed endowment of high-skill and low-skill workers, $\overline{L_H}$ and $\overline{L_L}$. Intermediate goods producers in each of the three sectors employ workers of both types and have sector-specific CRS production functions with varying skill-intensity (for simplicity I assume that manufacturing and services have the same skill-intensity):

$$\begin{aligned} Y_{ia} &= z_{ia} l_{Hia}^{\beta} l_{Lia}^{1-\beta} \\ Y_{im} &= z_{im} l_{Him}^{\alpha} l_{Lim}^{1-\alpha} \\ Y_{is} &= z_{is} l_{His}^{\alpha} l_{Lis}^{1-\alpha} \\ \alpha &> \beta \end{aligned} \tag{31}$$

Manufacturing and services are more high-skill intensive than agriculture, as reflected by the high-skill labor production elasticities $\alpha > \beta$. Solving the firm's problem gives the following optimal ratio of high-skill and low-skill workers employed in each sector as a function of the production elasticities and relative wages:

[54] Note that I do not estimate this version of the model because the simplified framework with two types of workers does not straightforwardly correspond to the data. In addition, it is not obvious how to estimate separate production elasticities by sector and skill-level, even if high-skill and low-skill workers were well-defined in practice.

$$\frac{L_{Hm}}{L_{Lm}} = \frac{\alpha}{1-\alpha}\left(\frac{w_L}{w_H}\right)$$
$$\frac{L_{Ha}}{L_{La}} = \frac{\beta}{1-\beta}\left(\frac{w_L}{w_H}\right)$$

With $\alpha > \beta$, these conditions imply that manufacturing and services firms will employ a higher share of high-skill workers than agricultural firms for any set of relative wages. The relative wage will adjust to satisfy both these conditions as well as the labor market clearing conditions in both sectors - total employment by skill type across the three sectors must add up to the country-level endowment of each skill type - such that wages respond both to productivity and to the relative scarcity of each type of worker.

This version of the model leaves several predictions of the baseline specification unchanged, and makes two distinct predictions worth highlighting. The predictions of the baseline model that carry through in this extension concern the basic dynamics of sectoral reallocation in response to a productivity shock. As in the baseline model, a decline in agricultural productivity (Z_a falls) will raise the marginal cost of production for firms in agriculture, forcing them to raise prices in a competitive market. The variety-level increases in p_a will raise the corresponding aggregate price index for the final good in agriculture, P_a. Consumer preferences remain as in the baseline specification, so Equation 25 governing the expenditure share in agriculture will continue to dictate that ω_{ak} rises in response to the rise in P_a and the decline in real wages associated with the productivity shock. As in the baseline model, Equation 26 shows that agriculture's share of GDP will rise with the expenditure share if the response of net exports to the change in comparative advantage is not sufficiently large. Thus, the competing forces of subsistence food requirements and international trade that govern the primary sectoral reallocation comparative statics are qualitatively robust to the extension with worker heterogeneity.

The model extension adds two dimensions of richness to our understanding of sectoral reallocation following a productivity shock in agriculture: more information about the distributional consequences of climate change and a more nuanced representation of comparative advantage. First, incorporating heterogeneous workers into the model allows me to examine the distributional consequences of climate change within, in addition to across, countries. On this point, the model predicts that the relative wage of low-skill workers to high-skill workers rises with the revenue share of agriculture.[55] Since agriculture is the less skill-intensive sector, the 'food problem' actually works to partially insulate farmers from the welfare

[55]The outline of the proof of this statement is as follows. In a perfectly competitive market with low-skill and high-skill workers as the only inputs to production, each sector's revenues are split between their workers according to their Cobb-Douglas production elasticities. So total income for

costs of declining agricultural productivity. Intuitively, inelastic demand for the sectoral output good causes a strong response of the output price that raises the relative wages of the low-skill workers disproportionately employed in that sector. So while the relationship between greater openness to international trade, sectoral reallocation, and aggregate productivity remains similar in the case of heterogeneous workers, the extended model suggests that the adaptation gains from trade will likely be smaller for agricultural and other low-skill workers if trade moves domestic production away from that sector and dampens the increase in its output price.

The second insight of the model with heterogeneous workers is that comparative advantage depends not only on the relative aggregate productivities in each sector, but also on the relative scarcity of high-skill and low-skill workers. Burstein and Vogel (2017) use a very similar model to specify a generalized definition of comparative advantage that incorporates both these Ricardian and Heckscher-Ohlin forces. In this framework, comparative advantage evolves endogenously with sectoral reallocation as relative wages shift with labor demand. Movement into (away from) agriculture raises (lowers) the relative wage of low-skill workers and shifts comparative advantage further toward (away from) manufacturing. For the primary climate change counterfactuals of interest in the paper, this additional channel would have the effect of attenuating the degree of sectoral reallocation in both directions. If the 'food problem' shifts production toward agriculture when its productivity falls, the resulting increase in the relative wage of low-skill workers pushes comparative advantage further toward manufacturing and endogenously strengthens the importance of the trade response pulling labor away from agriculture. Similarly, in the case of relatively free trade, production moving away from agriculture would reduce the relative wage of low-skill workers and endogenously dampen the movement of comparative advantage away from agriculture.

Overall, extending the model to represent workers of heterogeneous skill type leaves the fundamental predictions about climate change and sectoral reallocation unchanged, but sheds additional light on the forces underlying comparative advantage and the distributional consequences of climate change.

---

each category is given by:

$$w_L\overline{L_L} = (1-\beta)R_a + (1-\alpha)R_m + (1-\alpha)R_s$$
$$w_H\overline{L_H} = \beta R_a + \alpha R_m + \alpha R_s$$

Consider a 1% increase in the revenue share of agriculture, r_a, and a 1% decline in the revenue share of manufacturing, r_m. The change in low-skill share of total income is given by $\alpha-\beta$ and the change in the high-skill share of total income is given by $\beta-\alpha$. With $\alpha > \beta$ the low-skill share of total income rises. Since the total number of low-skill and high-skill workers is fixed, $\frac{w_L}{w_H}$ also rises.