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

Ishan Nath *

Princeton University

May 4, 2022

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 through adapting to global warming by shifting labor toward manufacturing and increasing imports of food. To investigate the likelihood that such gains will be realized, I embed the 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 sectors to meet subsistence consumer needs. The productivity effects of global warming reduce welfare by 5-8% for the poorest quartile of the world with trade barriers held at current levels, but by over 40% 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, Will Rafey, Levi Crews, 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, the Federal Reserve Bank of San Francisco Virtual Seminar on Climate Economics, 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

1 Introduction

Existing research suggests that climate change will cause major changes in agricultural productivity across the world during the 21st century. Figure 1a shows the average impact projection from a range of estimates in this literature, which suggests that global warming will reduce agricultural productivity by 30-60% in hot, largely agrarian, regions such as Sub-Saharan Africa and South Asia while having neutral to positive effects in colder parts of the world.[1] These forecasts suggest large potential gains from shifting the geography of agricultural production. If tropical regions reallocate production toward non-agricultural sectors while agricultural specialization moves toward temperate climates, the damage caused by climate change might be substantially reduced. 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 1a will push agriculture away from the equator. First, these estimates show changes in the absolute advantage of agriculture, whereas international trade responds to comparative advantage across sectors. 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 that existing evidence suggests temperature also affects non-agricultural productivity, the change in comparative advantage is not immediately clear.

Second, comparative advantage does not exclusively, or even primarily, determine sectoral specialization. Figure 1b shows that poor countries have high agricultural GDP shares despite a much lower ratio of relative value-added per worker in agriculture compared to non-agriculture than that of rich countries. 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 its low 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 with a quantitative macroeconomic model to assess the potential for sectoral reallocation

[1] The agricultural productivity impacts displayed in Figure 1a take the average of estimates produced by Hultgren et al. (2021), Cline (2007), Iglesias and Rosenzweig (2010), and Costinot, Donaldson and Smith (2016). Sections 7.1, 7.6, and Appendix E.3 contain further details on each of these sources, which produce very similar estimates despite using a wide range of methods and data. More broadly, the global projections shown in Figure 1a are also consistent with a large body of papers that produce more local estimates of the impacts of climate change on agriculture, such as Mendelsohn, Nordhaus and Shaw (1994), Deschenes and Greenstone (2007), Schlenker and Roberts (2009), and Schlenker and Lobell (2010), among many others.

to contribute to climate change adaptation. First, to project changes in agricultural comparative advantage, I provide the first global micro estimates of the impact of rising temperatures on labor 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 make projections that account for firm-level adaptation by allowing these treatment effects to vary with income and expectations of temperature.

Second, I construct a global model of sectoral specialization and trade that explains existing patterns of agricultural 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 expenditures 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 these consumer needs for food, but in practice this channel is weak in developing countries. This paper calculates that 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 expenditure shares necessary for people with low incomes to meet subsistence needs to eat.

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.[2] To quantify the relative strengths of these mechanisms, I calibrate 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 simulate the effects of climate change on sectoral specialization, trade, prices, GDP, and welfare.

The paper has three key findings. First, I find that extreme temperatures can have substantial

[2]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 quartiles of the world, implying an increase in the shadow value of migration but also fewer resources for people in poor countries to pay the costs of migration. I leave it to other papers, such as recent work by Cruz Alvarez and Rossi-Hansberg (2021), to quantify the response along this margin.

effects on non-agricultural productivity in some contexts, but with strong evidence consistent with successful adaptation by firms in rich countries and those that experience given extremes frequently. For the least adapted firms 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.[3] Effects are about half as large in middle-income countries, and smaller still for extreme temperatures common to a given location. 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.[4]

I combine these estimates of predicted temperature sensitivities with global climate model predictions of future temperatures to project the country-level effects of climate change on manufacturing and services productivity. Overall, rising temperatures project to reduce manufacturing productivity by 1.7% for the world on average with more substantial effects (5-14%) in some poor countries. While these effects are non-trivial, they are generally much smaller than previously estimated effects in agriculture. Thus, the change in the global comparative advantage of agriculture is qualitatively similar to its change in absolute advantage, implying large potential gains if hot parts of the world are able to reallocate production away from farming.

The second key finding is that model simulations suggest that 'the food problem' dominates comparative advantage in driving the response to climate change, limiting the projected gains from sectoral reallocation. Climate change raises the agriculture share of GDP by 2.7 percentage points in the poorest quartile of the world even as agricultural productivity in those countries 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. Thus, trade reduces climate damages by only 1.2% for the poorest quartile, and 1.7% for the world overall, relative to a scenario that assumes countries start in autarky, largely because those countries most susceptible to climate change are also least open to trade. Overall, the productivity effects of extreme temperatures caused by climate change reduce welfare by 2.8% for the world and 7.6% for the poorest quartile when measured from the present day baseline, and by 1.5% and 4.7%, respectively, from a future baseline that accounts for economic growth reducing temperature sensitivity and agricultural specialization over time.[5]

The third key finding is that counterfactual policy simulations suggest that greater openness to trade could dramatically reduce climate damages in poor countries. To quantify the importance

[3] I find similar effects for manufacturing and services firms, though I lack data coverage for services firms in poor countries where the effects of temperature are most detectable.

[4] 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 the welfare analysis. See Appendix D for details.

[5] Note that this paper does not conduct a comprehensive welfare analysis of climate change. Other important effects, such as hurricanes, sea level rise, health costs, and catastrophic tail risks, are beyond the scope of the analysis.

of tradability for climate change adaptation, I consider a separate counterfactual in which I assign all countries to about the 90th percentile level of trade openness, consistent with poor countries achieving future levels of openness comparable to the present day OECD. In this hypothetical scenario, welfare costs are 14% lower for the world overall, and 42% lower for the poorest quartile than in the main projection with trade barriers calibrated to match observed trade flows. Greater tradability reduces climate damages by making specialization more responsive to shifting comparative advantage, thereby reducing the climate-induced rise in food price indices globally, and especially so in hot parts of the world where agricultural productivity suffers domestically. Thus, the results suggest that understanding the causes of low levels of trade in developing countries could be an important topic for research on climate change adaptation.

This paper relates to several literatures on climate change and macroeconomic development, including a small number of closely related papers. For example, Costinot, Donaldson and Smith (2016) project substantial climate change adaptation gains from reallocation across types of crops, though their framework does not allow for analyzing reallocation across aggregate sectors. In their model, income effects are precluded by the quasi-linear utility specification, the expenditure share on agriculture is fixed in each country, non-agricultural production is not subject to climate damages, and trade in all non-agricultural sectors is costless.[6] The other most similar work consists of two papers by Desmet and Rossi-Hansberg (2015) and Conte, Desmet, Nagy and Rossi-Hansberg (2020), which examine migration and trade in a dynamic model and project meaningful adaptation gains from changes in the global distribution of sectoral specialization. This paper addresses a similar question, but is the first to focus specifically on incorporating the non-homothetic preferences that underlie 'the food problem' and drive sectoral specialization in poor countries.[7] Critically, this paper's results break with the full range of previous work by showing that global warming is likely to *intensify* agricultural specialization in the hottest parts of the world under current levels of trade openness, limiting the gains from reallocation in the most vulnerable places.

This paper's empirical work on temperature and productivity builds on country-level estimates produced by Somanathan, Somanathan, Sudarshan and Tewari (2021) 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 (2021). The model counterfactuals relate to empirical work by Colmer (2021) and

[6]Another paper by Gouel and Laborde (2021) builds on Costinot, Donaldson and Smith (2016) in several ways, such as including livestock in the analysis, but also does not include model components necessary to analyze cross-sector reallocation.

[7]Subsequent work by Cruz Alvarez (2022) builds on this paper by incorporating non-homothetic preferences into a dynamic model that also allows for analyzing migration at both the subnational and international levels.

Liu, Shamdasani and Taraz (2020) that examines the local relationship between temperature and sectoral reallocation in Indian districts. Their 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 this paper's model predictions about tradability and the food problem. Finally, some of the 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 contributes to the advancing 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 contributes to a nascent body of work - including Balboni (2019), Barrage (2020), Conte (2021), Cruz Alvarez and Rossi-Hansberg (2021), Fried (2021), and Rudik, Lyn, Tan and Ortiz-Bobea (2021) - that combines detailed micro-data with quantitative macroeconomic models to study the economics of climate change adaptation. This unifying approach aims to capture the strengths of each prior branch of literature: empirical estimates that map directly into model parameters, allowing the researcher to evaluate climate impacts, 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 temperature and non-agricultural productivity. Section 5 lays out the model. Section 6 explains the model calibration 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, years, and data sources included in the dataset. 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. 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. 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. More details on the data construction are contained in Appendix A.1.

The sample covers 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. 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 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.[8] 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, the data contains information about hot, cold, rich, and poor countries, and the degree to which the effects of temperature might differ across these contexts.

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

[8] 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 Appendix C.

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. 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.

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 section of the paper to enable the model simulations that follow. 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.

3.1 Conceptual Framework for Firm-Level Temperature Exposure

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 can cause physical fatigue, impair cognitive function, or increase the disutility of labor such that workers reduce effort or minutes spent working.[9]

[9] 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.

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

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

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.[10] 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:

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

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. Each type of adaptation investment is available at a constant marginal cost, given by c_h and c_c, respectively. 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:

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

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.[11] Thus, the firm's optimal adaptation investment 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.

[10] 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.

[11] Optimal adaptation investment is decreasing in worker effort because there are concave returns to effort.

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 the 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, the 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):

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

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 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. (2020), 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:

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

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.[12]

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. 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 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 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.[13]

[12] 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.

[13] Empirical estimation of adaptation in the climate impacts literature broadly relies heavily on cross-sectional

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 this 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 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.[14]

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

variation because of the inherent difficulty in finding quasi-experimental variation in long-run conditions.

[14]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 national or global product market prices.

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 2 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 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.

Though the results are robust to a range of specification choices, it is worth noting some limitations of this empirical approach. First, panel data on production in the informal sector is not widely available, and thus not included in the analysis. Second, the panel approach to establishing identification necessarily sacrifices the ability to measure any permanent component of temperature effects that is absorbed by the firm fixed effects. While I capture adaptation by observing how marginal effects differ by local climate, the assumption that effects accumulate additively in the long run is difficult to verify.[15] Finally, aggregating from firm-level effects to industry level effects in Section 7.1 requires the assumption of homogeneous effects across firms within a sector for each country. If temperature affects different manufacturing firms in India differently (a level of granularity beyond which I am able to systematically measure in the global data), then entry, exit, and reallocation within sectors could cause the aggregate effect to differ from the firm-level measure.

Overall, the net effect of these caveats could either increase or decrease the effect of future temperature change on non-agricultural productivity. If the true effects are larger than projected, the welfare consequences projected in Section 7 will be larger than projected and the potential

[15] Intuitively, consider a regression that measures the effects of a drought. If farmers adapt by irrigating their crops from a finite pool of groundwater, the effects of repeated drought exposure that depletes their stock would not aggregate linearly from the effect of a single drought. This dimension of permanent changes has generally not been accounted for in the empirical climate impacts literature, and would not appear in this paper's projections to the extent that similar mechanisms exist for temperature and non-agricultural productivity.

gains from reallocating production from agriculture to non-agriculture in hot places will be smaller, whereas the reverse would be true if the effects are smaller than projected. In light of these limitations, it is perhaps worth reiterating that this paper finds that the effects of rising temperatures on global non-agricultural productivity are approximately an order of magnitude smaller than the consensus of estimated effects in agriculture. This stark differential suggests that the qualitative findings about the potential gains from sectoral reallocation in Section 7 would be robust to even substantial adjustments to the projections made here.

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.[16] 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 2). Figure 3a 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 using global data in Figure 2, I find precisely estimated null effects of temperature on output-per-worker in the U.S.[17] The U.S. data also includes information on other inputs lacking from the global sample, allowing me to observe some of the adaptation costs incurred by U.S. firms. Appendix Figure 3b shows that the average U.S. plant increases expenditures on electricity and other fuels by several thousand dollars for each extremely hot and cold day, presumably for cooling and heating expenses.[18] 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 calibrate the model. Figure 4a shows the predicted effects of a day with maximum temperature of 40°C on annual manufacturing revenue per worker and Figure 4b shows the effect of a -5°C day. Consistent with intuition about adaptation and the results displayed in Figure 2, poor countries and those which experience given temperatures

[16]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).

[17]The result displayed in Figure 3a 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.

[18]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.

less frequently are more susceptible to extreme realizations.[19]

Projecting the impacts of climate change also requires accounting for adaptation by adjusting the temperature sensitivities shown in Figures 4a and 4b 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 temperatures in Appendix Figure A-16.[20] 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 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 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. (2020). Appendix D covers the details of this calculation.

The model simulations in Section 7 also require 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.[21] 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.[22] 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, so I make projections under the assumption that the income gradient of temperature sensitivity in manufacturing is similar to that of services. Appendix Figures A-19 and A-20 show predicted current global sensitivity to hot and cold days in

[19]Note that following Carleton et al. (2020), 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.

[20]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 7.5 I also allow for economic growth to make countries richer in the future, further reducing their temperature sensitivity.

[21]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.

[22]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.

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. The next section lays out the model used to simulate how sectoral reallocation and trade respond to the estimated changes in comparative advantage.

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, prices, GDP, 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 sectoral reallocation and its corresponding effect on climate change adaptation.

5.1 Model Ingredients

Consumption

Following the demand system specified in Comin, Lashkari and Mestieri (2021), 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.[23]

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},

[23]Nonhomothetic CES preferences improve model fit substantially compared to 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.

in sector j in country k:

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

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

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

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

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

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:[24]

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

Intermediate goods producers for each variety in each country receive a productivity draw, z_{ijk}, drawn from a Fréchet distribution with sector-specific shape parameter θ_j and sector-country specific scale parameter 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

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

country, k, and importer country, n. For a given country pair, the iceberg trade costs are allowed to differ both across sectors and with the direction of shipment. 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, firms in country k price domestically produced goods at $\frac{w_k}{z_{ijk}}$. When selling to foreign country n and incurring the cost of trade, the intermediate goods producer in country k sets the price of the exported good at $\frac{\tau_{jkn} w_k}{z_{ijk}}$. The final goods producer in each country sources each variety from the lowest-priced intermediate goods producer across all countries, such that the price of intermediate good i in sector j in country k is as follows:

$$p_{ijk} = \min_n \left\{ \frac{\tau_{jnk} w_n}{z_{ijn}} \right\} \tag{17}$$

The sectoral final goods prices are given by the CES price index of all intermediate varieties used in that sector:

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

Finally, the final goods producer's demand function for variety i 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.

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

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.

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

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:[25]

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

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 GDP

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,

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

(ω^0_{jk} and ω^1_{jk}), to construct an aggregate price index with which to deflate nominal income:

$$P^T_k = \prod P_{jk}^{(\omega^0_{jk}+\omega^1_{jk})/2} \longrightarrow GDP_k = \frac{w_k L_k}{P^T_k} \tag{24}$$

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 GDP 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 force labor toward low productivity agriculture rather than allowing it to reallocate to sectors less affected by climate change.

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 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.[26] 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 ($\frac{P_{ak}}{P_k}$).[27] 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

[26]This equation also appears in Comin, Lashkari and Mestieri (2021). 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.

[27]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.

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 protective 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.[28] 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.[29]

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

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 next calibrate the model parameters.

[28]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.

[29]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).

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 counterfactual simulations. In addition, I solve the model in levels rather than changes because the latter approach would preclude some of the key counterfactuals presented in Section 7.[30]

6.1 Parameter Estimates

I infer model parameters from the data using a combination of calibration and estimation. Table 3a 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, respectively, for 158 countries. Appendix E.1 contains details about data construction and the SMM procedure.

Table 3b 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.[31] 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 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.[32]

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

[30] Solving the model in changes following Dekle, Eaton and Kortum (2007) is not feasible for the key results presented in Section 7.5 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 in the baseline that accounts for economic growth but does not include 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.

[31] 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.

[32] 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 5b shows that the pattern of relative prices in the model is similar to that of the data.

a sector sharply increases with its relative price. This paper's estimate of σ to target the global cross-section of sectoral shares matches up well with that of Comin, Lashkari and Mestieri (2021), who use various historical panel datasets to estimate σ between 0.2 and 0.6.[33]

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 (2021). In particular, the estimates in Table 3b 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.[34] 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 A-4 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 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-29, most developing countries import little of their food. In the data, the average person in the poorest quartile of the world consumes 91% domestically produced food (89% in the simulation) compared to 45% in the richest quartile (52% 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 over 60% of the variation in the global agriculture share of GDP.[35] This is a relatively strong fit considering that only the seven free parameters in Table 3b were chosen to match 316 target moments consisting of GDP shares for agriculture, manufacturing, and services in 158 countries. As shown in Figure 5a, the nonhomothetic CES demand specification enables

[33] 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 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.

[34] 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 the sample.

[35] 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-30.

the simulation to closely mirror the smooth decline of agricultural GDP with log income per capita observed in the global cross section.

The model also reproduces the general pattern of high relative prices for agricultural consumption in poor countries - a moment I do not target in the calibration. In Figure 5b, 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 trade costs calibrated to rationalize observed trade flows tightly link domestic consumption to domestic production, causing many developing countries to specialize in agriculture despite its low productivity.[36] In the next section, I use the model to investigate the welfare consequences of the projected sectoral reallocation that occurs in response to climate change.

7 Model Counterfactuals

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

7.1 Estimated Productivity Impacts and Comparative Advantage

I start by projecting the impacts of climate change on country-sector level productivity. 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 future temperature distributions in 2080-2099.[37] I obtain future temperature predictions from the CSIRO-MK-3.6.0 model produced by Jeffrey et al. (2013).[38]

Figure 6a and Appendix Figure A-23 show the projected changes in manufacturing and services productivity, respectively. The results suggest that climate change will have nontrivial effects on

[36]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.

[37]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 7.5.

[38]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-21 and A-22. I multiply the country-level coefficients by the projected changes in hot and cold temperatures to get the impacts shown here.

non-agricultural productivity. Population-weighted global average manufacturing productivity falls by 1.7% in the projections, with small improvements of up to 3.2% in 11 richer, colder countries and declines of more than 5% (and up to 14.2%) in 28 poorer, hotter countries. The results for services are qualitatively very similar, though less central for the model simulations about comparative advantage and trade.

For agricultural productivity, a rich literature already exists on the projected impacts of global warming so it is not necessary to make new projections. Instead, I take the unweighted country-level average of estimates from four leading sources in the literature. The four sources used in this paper are Hultgren et al. (2021), which makes projections using panel estimates from a global dataset with subnational resolution following a similar procedure to that used in the non-agricultural estimates in this paper; Cline (2007), which uses Ricardian estimates from a separate collection of global micro-data; Iglesias and Rosenzweig (2010), which uses projections from leading crop models assembled by the International Consortium for Application of Systems Approaches to Agriculture (ICASA); and Costinot, Donaldson and Smith (2016), which use crop model estimates from the UN Food and Agriculture Organization's Global Agro-Ecological Zones database. For each source, I use projections for the same high-emissions scenario used in the non-agricultural estimates in this paper for consistency.

I take the average across these sources because each contains its own advantages and drawbacks, but Table 7 shows that the key results in this paper are very similar when using each agricultural impact projection on its own. The robustness of this paper's results to the choice of agricultural productivity estimates is not surprising given that these projections are remarkably quantitatively similar despite the wide range of methods and datasets they employ. The population-weighted global average decline in agricultural productivity ranges between 18% and 21% in Costinot, Donaldson and Smith (2016), Iglesias and Rosenzweig (2010), and Cline (2007), with a slightly smaller decline of 14.3% projected in Hultgren et al. (2021). Appendix E.3 contains further details on the methods employed by each of these sources and their relative strengths and weaknesses.

Figure 6b brings together the estimated manufacturing productivity effects with the agricultural productivity estimates from the literature to show the change in the relative productivity of the model's tradable sectors for every country. The pattern shows clearly that climate change shifts comparative advantage in agriculture toward colder countries far from the equator on average. While the negative effects on manufacturing productivity are concentrated in similar parts of the world as declines in agricultural productivity, they are generally smaller in magnitude. The projected population-weighted global average change in agriculture's productivity is -18.3%; more than ten times larger than in manufacturing. Every country in Africa, South Asia, and Latin America has larger estimated productivity losses in agriculture than manufacturing, implying large potential gains from reallocating production in these locations away from farming. The extent to which such gains from adaptation through sectoral reallocation are realized will depend on the degree to which specialization responds to Ricardian comparative advantage.

To evaluate the effects of climate change on trade, sectoral reallocation, and welfare, I embed these empirically estimated productivity changes into the model by adjusting each sector-country specific aggregate productivity Z_{jk} by its corresponding climate change effect. I then calculate equilibrium wages, prices, sectoral shares, and trade flows in counterfactual scenarios with and without climate change, and in intermediate counterfactuals that decompose the effects by allowing for only certain dimensions of reallocation. I start by considering a climate change counterfactual relative to a baseline scenario at present day incomes and productivities taken from the model calibration described in Section 6. This allows me to disentangle the mechanisms by which sectoral reallocation mediates the welfare effects of climate change. In Section 7.5, I consider the implications of accounting for economic growth that affects baseline productivities and incomes before the effects of climate change enter the model.

7.2 Trade Flows and Sectoral Specialization

Figure 7a shows the simulated effect of climate change on agricultural net exports. Consistent with the estimated change in comparative advantage, the predominant pattern is that hotter countries experiencing large declines in agricultural productivity move toward importing more food, while cooler countries with neutral or improving agricultural productivity move toward exporting more food. For instance, Norway and Canada increase agricultural net exports from 0.9% to 1.8% and 0.5% to 1.9% of GDP respectively. Conversely, most countries in 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 5 out of 158 countries decrease agricultural net exports by more than 10% of GDP. This feature of the projection follows from the low baseline trade shares in agriculture in the countries for whom climate change hits hardest. As discussed in Section 6, the import share of agricultural consumption in the poorest quartile of the world is less than 10%, implying large trade costs in the model calibration and a muted trade response to global warming.

The change in trade flows only partially determines sectoral reallocation. As shown in Section 5.2, 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:

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

The change in net exports shown in Figure 7a 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.

I decompose the competing effects of climate change on the agriculture share of GDP by running separate counterfactuals with and without trade adjustment. 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 corresponding change in expenditure shares. This provides the change in ω_{ak}, which would equal the change in agriculture's share of GDP under the assumption that each country started in autarky with production shares equal to consumption shares. In contrast, the full equilibrium counterfactual incorporating trade provides the net effect of both forces governing reallocation. Table 4a displays the baseline, autarky counterfactual, and trade-inclusive counterfactual agriculture shares of GDP for a selection of countries.

The results in Table 4a show that the consumption response and trade response both have substantial effects on specialization in agriculture, with significant heterogeneity across countries. In Ethiopia, India, and Zambia the 'food problem' effect dominates and the agriculture share of GDP rises in response to climate change despite large relative declines in agricultural productivity and a shift toward importing more of their food. In contrast, the trade effect dominates in Ghana and Niger, where the domestic production share of agricultural expenditures falls from 81% to 69% and 84% to 64% respectively, and climate change reduces agricultural specialization on net. This difference suggests that the latter two countries are more open to trade on the margin in the calibration or that their change in comparative advantage was especially extreme, though conclusions about any one country cannot be interpreted literally due to simulation error in the baseline model fit. Other countries, such as Canada, Norway, and Mozambique, see an increase in agricultural specialization because of increased exports driven by improvements in relative agricultural productivity compared to the world as a whole, and their trading partners in particular.

Figure 7b 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 +2.7 percentage points. Thus, simulations suggest that climate change will keep more people in poor countries working on farms even as their productivity declines dramatically.

7.3 Reallocation and Welfare

This section explores the welfare consequences of the sectoral reallocation that increases the labor share of agriculture in hotter countries. I start by showing the projected impact of climate change on real GDP in each counterfactual considered in Section 7.2 - no reallocation, autarky (expenditure share change only), and full reallocation. The effects of climate change, and of the endogenous reallocation response, on GDP differ meaningfully from the welfare effects, but correspond to what could plausibly be measured in data and help to illustrate the mechanisms of the model.

To calculate the impact of climate change on GDP, I deflate nominal income at the country level using the Tornqvist price index from Equation 24 with prices and sectoral shares that correspond to each counterfactual. In the no reallocation counterfactual, prices respond directly to the estimated changes in sectoral productivity, but sectoral shares stay fixed. In the autarky counterfactual, prices change and sectoral shares respond as dictated by the change in consumption shares. Finally, in the full reallocation counterfactual, prices and sectoral shares both reflect their levels in equilibrium with climate change that incorporates the full response of consumption shares and trade.

Table 4b shows the impact of climate change on real GDP in each counterfactual for a selection of countries. The results make clear that projected reallocation actually *exacerbates* the decline in aggregate output in most countries, as well as globally on average. Global GDP declines 1.8% in the naïve counterfactual that holds sectoral shares fixed, but by 2.1% when allowing for the full effects of reallocation. This effect is especially stark in poor countries. GDP in the poorest quartile of countries falls by 6.3% in the no reallocation counterfactual, and by 9.4% with reallocation.

The projected impact of climate change on GDP is larger when accounting for the endogenous change in sectoral shares for two reasons. First, as shown in Section 7.2, the 'food problem' pushes up the labor share of agriculture in many countries where agricultural productivity declines dramatically. 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. 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 the impact of climate change on GDP be consistent with optimizing behavior? In Table 4c, I calculate the equivalent variation percent change in welfare using Equation 23 under each counterfactual. The results show that the full reallocation counterfactual reduces the *welfare* consequences of climate change dramatically relative to the no reallocation scenario, even while exacerbating the decline in GDP. Agents have an extremely high WTP to avoid climate change in the no reallocation counterfactual because this naïve hypothetical forces them to deviate from their optimal consumption bundles, and thus does not represent an equilibrium. In particular, keeping fixed the expenditure share on food in the presence of declining incomes and large projected increases in food prices would require sharp declines in the quantity of food consumed. This hypothetical with very low food consumption has a severe corresponding effect on welfare. Thus, to summarize the intuition of the results in Tables 4b and 4c, people are willing to sacrifice income (GDP) to reallocate expenditures toward food when food prices rise because they need food to survive.

While reallocating expenditures toward food as prices rise significantly tempers the welfare effects of climate change relative to the extreme example of fixed consumption shares, accounting for the impact of trade does little to affect the results. Comparing Column 2 and Column 3 of Table 4c shows that the welfare effects of climate change are closely approximated by the counterfactual that assumes all countries start in autarky. The willingness-to-pay to avoid climate change is only 1.7% lower for the world as a whole, and 1.2% lower for the poorest quartile, in the full reallocation counterfactual than in autarky. The few exceptions to this broader pattern are those countries for which the trade response was particularly strong in the simulation. For instance, in Niger, trade reduces the welfare costs by over 23% as the relatively large increase in agricultural net imports (12.5 percentage points of GDP) allows production to move substantially away from agriculture.

Figure 8 summarizes the welfare effects of climate change in the full reallocation counterfactual. Panel (a) shows the global distribution of willingness-to-pay to avoid climate change, and Panel (b) displays the change in food prices, P_{ak}, which comprise a key component of the welfare effects. Food prices rise in all countries, and rise by at least 20% in 81 countries containing about half of the world's population.[39] Climate change does net damage as measured by WTP in 152 out of 158 countries, with the worst losses concentrated in poor countries. The total willingness-to-pay across countries amounts to 1.7% of global GDP as rich countries accounting for a disproportionate share of global income are hit relatively less hard, but the average person in the poorest quartile of the world suffers losses totaling 6.4% of income. In 20 African countries, the WTP to avoid warming exceeds 10% of 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 incorporated in the more detailed welfare analysis in Section 7.5.

[39]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 consequences of climate change within countries further in Appendix F.2.

7.4 Low Trade Cost Counterfactual

The analysis of sectoral reallocation and welfare in Sections 7.2 and 7.3 shows that openness to trade can reduce the harm from climate change by counteracting 'the food problem' and allowing specialization in hotter places to move away from the hardest hit sector. To further investigate the importance of trade costs, 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 in the calibration.[40] 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.[41] A 100% tariff-equivalent trade cost is approximately equal to the calibrated cost for shipping food from Belgium to Australia; representing an ambitious, yet realistically feasible, change in global trade costs.

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 5 shows the welfare effects of climate change in the low trade cost scenario for a select subset of especially vulnerable countries. Two things about these results are worth noting. First, the results imply that reducing trade barriers could dramatically reduce 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 3.7% of income in the low trade cost case, relative to 6.4% in the estimated trade cost case. Global food prices rise by less in this scenario, and especially so in poor countries, as agricultural specialization closely follows comparative advantage and thus moves away from hotter regions. Table 5b shows that the average person in the poorest quartile of the world experiences a 20% increase in food prices in the low trade cost case, in contrast to the 27% increase they face in the previous counterfactual with estimated trade costs.

Second, the effects of openness to trade vary substantially across countries. For 35 countries representing 13% of the global population, WTP to avoid climate change as a share of GDP is *higher* in the low trade cost scenario.[42] The intuition for this result is as follows. When trade barriers

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

[41] I show robustness results for a variety of trade cost scenarios, including frictionless trade, in Section 7.6.

[42] 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.

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, the global willingness to pay to avoid climate change is 14% lower under the specified alternative scenario with freer trade than in the results from Section 7.3 with existing levels of trade barriers. This pattern holds much more starkly in poor countries. For the average person in the poorest quartile of the world, the welfare costs fall by 42% in the low trade cost counterfactual. I discuss possible policy mechanisms to realize these gains in Section 9.

7.5 Accounting for Baseline Economic Growth

This paper's welfare analysis constitutes only a partial assessment of the costs of climate change. Given that the focus is on the potential benefits of sectoral reallocation and whether they are likely to be realized, I omit other important topics such as international migration (see e.g. Missirian and Schlenker (2017) and Cruz Alvarez and Rossi-Hansberg (2021)), health effects (Heutel, Miller and Molitor, 2017), hurricanes (Bakkensen and Barrage, 2018), sea-level rise (Desmet, Kopp, Kulp, Nagy, Oppenheimer, Rossi-Hansberg and Strauss, 2018), and uncertainty in sensitivity of global temperatures to emissions (Cai and Lontzek, 2019), all of which will likely play an important role in the welfare consequences of climate change. Nevertheless, in this section, I endeavor to improve the estimates of the welfare effects of sectoral productivity changes by accounting for the effects of baseline economic growth.

The results in Sections 7.1 to 7.4 use projections for future temperature change, but hold the baseline global economy fixed at the present day equilibrium in order to isolate the mechanisms driving reallocation. In this section, I allow global income levels in the no climate change baseline 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, which have been used widely in research on the economics of climate change.[43]

Allowing for economic growth to take place has two important effects on the welfare 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-

[43] Use of the Shared Socioeconomic Pathways in future projections of climate change damages follows from the work of Carleton et al. (2020), Newell, Prest and Sexton (2021), Kalkuhl and Wenz (2020), and many others. I use Scenario Three here because its projections have tracked most closely to realized growth in the early decades of the century, but I show results for other growth scenarios in Section 7.6.

neutral increases in the baseline values of Z_{jk}.[44] 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 4a and 4b at 2080 levels of log GDP per capita. Appendix Figures A-24 and A-25 show that the effects of temperature on non-agricultural productivity accounting for this 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 6 shows the impact of expected economic growth on the agriculture share of GDP and the willingness-to-pay to avoid climate change in a selection of the most vulnerable countries. The willingness-to-pay numbers in Columns 2 and 3 of Table 6b also incorporate the firm-level adaptation costs calculated in Appendix D, thus accounting more comprehensively for the costs as well as benefits of within-sector adaptation. This particular future scenario includes little to 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 6a shows that Zimbabwe and the Democratic Republic of Congo 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 Ethiopia.

The results in Table 6b show that the aggregate global WTP for climate change is 2.8% 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 7.6% from the present baseline and 4.7% from the future baseline.[45] 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.[46] The Rawlsian welfare losses from climate change are 5.4% of global GDP from the present income baseline and 3.1% of global GDP from the future baseline, approximately twice as high as the aggregate willingness-to-pay calculated by summing across agents.

7.6 Model Simulation Robustness

In Table 7, I examine the robustness of the model simulations to a variety of alternative assumptions. The top panel shows alternative versions of the baseline results from Sections 7.2 and 7.3. First, I simulate a version of the model that uses generalized Stone-Geary preferences to represent prefer-

[44]Święcki (2017) shows that productivity growth in recent decades has been very similar in agriculture and manufacturing, though slower in services.

[45]The numbers from the present-income baseline here differ from those presented in Section 7.3 because they also account for the firm-level costs of adaptations to reduce heat sensitivity.

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

ences for food with a sharp minimum subsistence requirement rather than the smoothly declining budget share implied by the nonhomothetic CES specification. Appendix F.1 covers the calibration details for this version of the model, which produces results very similar to the baseline across all dimensions.

The second, third, and fourth sets of alternative results use individual sets of agricultural productivity estimates from Hultgren et al. (2021), Iglesias and Rosenzweig (2010), and Cline (2007), rather than the average across the range of sources. The advantage of the Hultgren et al. (2021) estimates is that they use an empirical approach following Carleton et al. (2020) and the non-agricultural productivity projections in this paper. In particular, that paper uses panel regressions to establish plausible causality and strives to account carefully for within-crop farmer adaptation. The advantage of the Iglesias and Rosenzweig (2010) estimates is that they use crop models widely cited in the climate change literature (e.g. Parry et al. (2004)), though recent work by Adamopoulos and Restuccia (2018) has shown that such crop models do not account for existing cross-country differences in agricultural productivity because they do not model technology. In addition, a key disadvantage of both Hultgren et al. (2021) and Iglesias and Rosenzweig (2010) is that they account only for staple grain crops and do not project crop-switching. In contrast, the estimates from Cline (2007) also include impacts on vegetables, fruits, and livestock, and allow for reallocation between all product categories within agriculture. The disadvantage of Cline (2007), however, is that its cross-sectional regressions require stronger assumptions to be interpreted as causal.[47] This paper's main projections use an average across sources to avoid allowing the potential drawbacks of any individual source to unduly influence the results.

The results in Table 7 show that the model simulations in this paper are similar across all sets of agricultural estimates. The key results regarding the impact of warming on the agricultural labor share, global GDP, welfare, and food prices have very similar magnitudes across estimates, with the worst impacts consistently concentrated in poor countries. The only noticeable difference between the estimates is that the welfare effects of warming are modestly more concentrated in poor countries when using Cline (2007) and modestly less so when using Hultgren et al. (2021), perhaps because richer countries with more temperate climates benefit most from crop-switching (which is accounted for only in Cline (2007)). Given the wide range of methods and datasets used in these sets of agricultural impact estimates, the results in Table 7 provide support for the robustness of the main findings of this paper. In addition, further support for the quantitative findings is provided by a range of other studies that estimate climate change impacts for only a given region (e.g. Schlenker and Lobell (2010) for Africa) and find similar results to the globally comprehensive projections that

[47] Estimates from Costinot, Donaldson and Smith (2016) are also used in the country-level averages shown in Figures 1a and 6b, but cannot be used for a robustness check in Table 7 because they were produced for only 50 countries that are not representative of the 158 countries used in the model simulations in this paper.

can be used directly in this paper's model simulations. It is worth noting, however, that a drawback common across the range of existing work on agriculture and climate change is that impacts are measured on a per-acre basis rather than a per-worker basis, as is the case in the model in this paper. Future empirical work evaluating whether the effects of climate on labor productivity differ from those on land productivity would provide an additional dimension of robustness to these results.

In Appendix F.2, I consider a version of the model with multiple factors of production, such as land or differentially educated labor. I do not calibrate this version of the model, but show that the model's key comparative statics regarding the competing effects of the food problem and trade remain qualitatively unchanged. I also show that this version of the model provides two additional insights about the main results. First, modeling heterogeneous workers allows for analyzing the within-country distributional effects of various counterfactuals. In particular, I show that greater openness to trade in places where agricultural productivity declines would reduce the relative wage of the type of workers intensively used in agriculture. Intuitively, trade barriers benefit those employed in unproductive agricultural sectors by keeping domestic food prices high. Thus, the climate change adaptation gains from increased openness are likely to accrue disproportionately to workers employed in non-agricultural sectors who primarily consume, rather than produce, food, though both types of workers may gain in absolute terms.

Second, with multiple inputs, comparative advantage in a given sector endogenously weakens as specialization shifts toward that sector because of the corresponding rise in the relative price of the input more intensively used in that sector. The baseline model already incorporates diminishing returns to specialization through the Fréchet distribution representing the continuum of producers in a given sector-country pairing. As specialization in a sector increases, lower quality producers become active, implicitly capturing effects such as agricultural producers moving toward lower quality land or manufacturing firms hiring less trained labor. Explicitly modeling these multiple types of production inputs introduces an additional dimension of diminishing returns into the model through the input price channel. Thus, the richer version of the model introduces an additional barrier to adaptation through sectoral reallocation, suggesting that the simulations with the baseline model likely overestimate the gains from global reallocation. This strengthens the main finding that shifting trade patterns contribute little to climate change adaptation with estimated trade barriers, and suggests caution in interpreting the magnitude of the adaptation gains from trade in the alternative simulations with reduced trade costs.

While I do not explicitly consider a model extension with mobility frictions between sectors for reasons explained in Appendix F.2, it is worth noting that this addition would also have a similar effect on the counterfactuals. If frictions prevent workers from reallocating from agriculture to non-agriculture, as some evidence in the literature suggests, this would further impede adaptation through sectoral reallocation and reinforce the main counterfactual in which the gains from this channel are limited. However, the presence of such domestic labor market frictions would also suggest further caution in interpreting the magnitude of adaptation gains that can be achieved from lowering trade barriers in the alternative counterfactuals.

The second panel of Table 7 shows how the results vary across a range of assumptions about trade costs.[48] Reducing all trade barriers by half has only a modest effect on the welfare costs of climate change, reducing losses by 2% overall and by 6% for the poorest quartile, perhaps because the trade costs estimated to rationalize the minimal existing trade flows in poor countries are so prohibitively high to begin with. In contrast, global damages fall by 14% and poorest quartile damages by 43% in the counterfactual with the more dramatic reduction in trade barriers considered in Section 7.4 where all bilateral trade costs move to approximate current rich country levels. For comparison, I also show results for a frictionless trade benchmark version of the model, though it is unlikely to be achievable in practice. With no trade costs, all countries face the same change in food prices and the welfare costs fall substantially further, though the losses remain concentrated in poor countries that experience larger declines in sectoral productivities and real wages.

The third panel of Table 7 reproduces the results from Section 7.5 under a variety of assumptions about baseline economic growth. Unsurprisingly, the results show that the world is less susceptible to climate change in the higher growth scenarios and that poor countries are less susceptible in the scenarios with greater convergence in which they grow faster.[49] More importantly for the research question in this paper, the results show that climate change raises agriculture's share of GDP in poor countries in all the growth scenarios. While the 'food problem' becomes less binding overall when economic growth is faster, climate change accentuates its importance from each of the future income baselines. Thus, the results can be interpreted as suggesting that global warming will slow the pace of structural transformation from any given baseline.

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

[48]Note that in each trade cost scenario, I rescale the country-sector productivity parameters such that baseline income is unchanged.

[49]Note that the results in this paper do not incorporate any endogenous feedback between economic growth and emissions, so these results pertain only to the contribution of economic growth to climate change adaptation when taking the degree of climate change as given.

(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 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 8a summarizes the data sources used in this part of the analysis.[50] 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.[51]

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

The results in Table 8b 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

[50] 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.

[51] 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.

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 the largest declines in agricultural productivity (>10 percentage points, with an average of 29.5%).

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.[52] Taken together with the analysis in Sections 6 and 7 and the existing body of evidence, however, 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, but do address an existing challenge in the literature by estimating the welfare consequences of global productivity changes 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 agricultural production through technology, irrigation, heat-resistant crop varieties, or other means. The agricultural productivity consequences projected in the climate change literature 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 9a and 9b

[52]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.

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 physical 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 can substantially improve outcomes. Falling productivity raises prices and encourages substitution to other products. But this logic weakens when applied to 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. This implies large potential gains from hot countries reallocating production away from farming. On average, however, model simulations suggest that these countries will largely move specialization toward, rather than away from, agriculture due to high trade costs in poor countries and the special properties of consumer preferences for food. Countries facing severe climate change damages in agriculture that are more open to trade suffer less because they are more able to increase imports of food and shift production toward less affected sectors. Overall, moving poor country trade barriers to the existing global frontier of openness 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 this 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 this analysis shows that reducing barriers to trade is a necessary condition to induce sectoral reallocation that substantially curtails the costs of climate change, I cannot conclude that it is sufficient. The precise magnitude of gains in the low trade cost counterfactual relies on out-of-sample assumptions about the degree of diminishing returns to specialization from producing more manufactured goods in hotter countries or expanding agriculture in colder countries. The model presented in this paper captures these diminishing returns through the Fréchet shape parameter, but omits a detailed representation of mechanisms such as soil quality or worker skill that could provide additional insight into the feasibility of dramatic sectoral shifts in a world with more trade. 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

Adamopoulos, Tasso and Diego Restuccia, "Geography and agricultural productivity: Cross-country evidence from micro plot-level data," Technical Report, National Bureau of Economic Research 2018.

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.

Allen, Treb, Costas Arkolakis, and Xiangliang Li, "On the equilibrium properties of network models with heterogeneous agents," Technical Report, National Bureau of Economic Research 2020.

Alvarez, Jose Luis Cruz, "Global Warming and Labor Market Reallocation," Technical Report 2022.

_ and Esteban Rossi-Hansberg, "The economic geography of global warming," Technical Report, National Bureau of Economic Research 2021.

Bakkensen, Laura and Lint Barrage, "Climate shocks, cyclones, and economic growth: bridging the micro-macro gap," Technical Report, National Bureau of Economic Research 2018.

Balboni, Clare Alexandra, "In harm's way? infrastructure investments and the persistence of coastal cities." PhD dissertation, The London School of Economics and Political Science (LSE) 2019.

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

Barrage, Lint, "The Fiscal Costs of Climate Change," in "AEA Papers and Proceedings," Vol. 110 2020, pp. 107–12.

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

Beach, Robert H, Timothy B Sulser, Allison Crimmins, Nicola Cenacchi, Jefferson Cole, Naomi K Fukagawa, Daniel Mason-D'Croz, Samuel Myers, Marcus C Sarofim, Matthew Smith et al., "Combining the effects of increased atmospheric carbon dioxide on protein, iron, and zinc availability and projected climate change on global diets: a modelling study," *The Lancet Planetary Health*, 2019, *3* (7), e307–e317.

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.

Cai, Yongyang and Thomas S Lontzek, "The social cost of carbon with economic and climate risks," *Journal of Political Economy*, 2019, *127* (6), 2684–2734.

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.

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*, 2021, *13* (4), 101–24.

Comin, Diego, Danial Lashkari, and Martí Mestieri, "Structural change with long-run income and price effects," *Econometrica*, 2021, *89* (1), 311–374.

Conte, Bruno, "Climate change and migration: the case of africa," *Working Paper*, 2021.

__ , Klaus Desmet, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg, "Local sectoral specialization in a warming world," Technical Report, National Bureau of Economic Research 2020.

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.

__ , Robert E Kopp, Scott A Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H Strauss, "Evaluating the economic cost of coastal flooding," Technical Report, National Bureau of Economic Research 2018.

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.

Fried, Stephie, "Seawalls and stilts: A quantitative macro study of climate adaptation," Technical Report, Federal Reserve Bank of San Francisco 2021.

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.

Gouel, Christophe and David Laborde, "The crucial role of domestic and international market-mediated adaptation to climate change," *Journal of Environmental Economics and Management*, 2021, *106*, 102408.

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

Heutel, Garth, Nolan H Miller, and David Molitor, "Adaptation and the mortality effects of temperature across US climate regions," *The Review of Economics and Statistics*, 2017, pp. 1–33.

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.

Hultgren, Andrew, Tamma Carleton, Michael Delgado, Diana Gergel, Michael Greenstone, Trevor Houser, Solomon Hsiang, Amir Jina, Robert Kopp, Steven Malevich, Kelly McCusker, Terin Mayer, Ishan Nath, James Rising, Ashwin Rode, and Jiacan Yuan, "Estimating the Global Impact of Climate Change on Staple Crops," Technical Report, Working Paper 2021.

Iglesias, Ana and Cynthia Rosenzweig, "Effects of Climate Change on Global Food Production under Special Report on Emissions Scenarios (SRES) Emissions and Socioeconomic Scenarios: Data from a Crop Modeling Study.," Technical Report, Columbia University 2010.

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.

Kalkuhl, Matthias and Leonie Wenz, "The impact of climate conditions on economic production. Evidence from a global panel of regions," *Journal of Environmental Economics and Management*, 2020, *103*, 102360.

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.

Missirian, Anouch and Wolfram Schlenker, "Asylum applications respond to temperature fluctuations," *Science*, 2017, *358* (6370), 1610–1614.

Myers, Samuel S, Antonella Zanobetti, Itai Kloog, Peter Huybers, Andrew DB Leakey, Arnold J Bloom, Eli Carlisle, Lee H Dietterich, Glenn Fitzgerald, Toshihiro Hasegawa et al., "Increasing CO 2 threatens human nutrition," *Nature*, 2014, *510* (7503), 139–142.

Newell, Richard G, Brian C Prest, and Steven E Sexton, "The GDP-temperature relationship: implications for climate change damages," *Journal of Environmental Economics and Management*, 2021, *108*, 102445.

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

Parry, Martin, Cynthia Rosenzweig, Ana Iglesias, Günther Fischer, and Matthew Livermore, "Climate change and world food security: a new assessment," *Global environmental change*, 1999, *9*, S51–S67.

Parry, Martin L, Cynthia Rosenzweig, Ana Iglesias, Matthew Livermore, and Gunther Fischer, "Effects of climate change on global food production under SRES emissions and socio-economic scenarios," *Global environmental change*, 2004, *14* (1), 53–67.

Rudik, Ivan, Gary Lyn, Weiliang Tan, and Ariel Ortiz-Bobea, "Heterogeneity and Market Adaptation to Climate Change in Dynamic-Spatial Equilibrium," 2021.

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.

Smith, Matthew R and Samuel S Myers, "Impact of anthropogenic CO 2 emissions on global human nutrition," *Nature Climate Change*, 2018, *8* (9), 834–839.

Somanathan, Eswaran, Rohini Somanathan, Anant Sudarshan, and Meenu Tewari, "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing," *Journal of Political Economy*, 2021, *129* (6), 1797–1827.

Świecki, Tomasz, "Determinants of structural change," *Review of Economic Dynamics*, 2017, *24*, 95–131.

Taylor, Charles A and Wolfram Schlenker, "Environmental Drivers of Agricultural Productivity Growth: CO2 Fertilization of US Field Crops," Technical Report, National Bureau of Economic Research 2021.

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.

Zhu, Chunwu, Kazuhiko Kobayashi, Irakli Loladze, Jianguo Zhu, Qian Jiang, Xi Xu, Gang Liu, Saman Seneweera, Kristie L Ebi, Adam Drewnowski et al., "Carbon dioxide (CO2) levels this century will alter the protein, micronutrients, and vitamin content of rice grains with potential health consequences for the poorest rice-dependent countries," *Science advances*, 2018, *4* (5), eaaq1012.

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.

Tables & Figures

Figure 1: Motivating Evidence on Climate Change and Agricultural Specialization

(a) Projected Impact of Climate Change on
Agricultural Productivity, 2080-2099

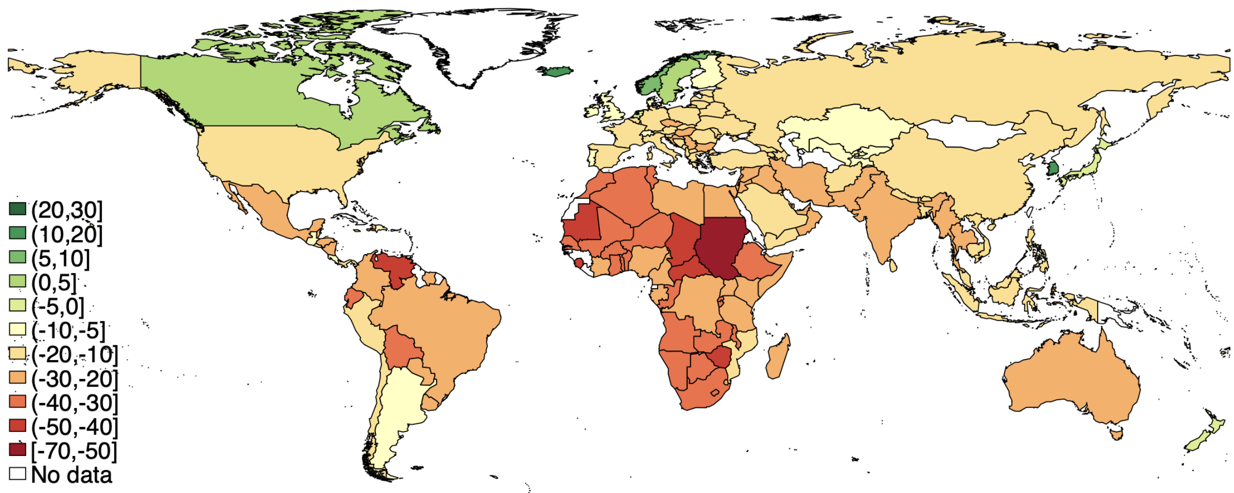

(b) Comparative Advantage and Specialization in Agriculture

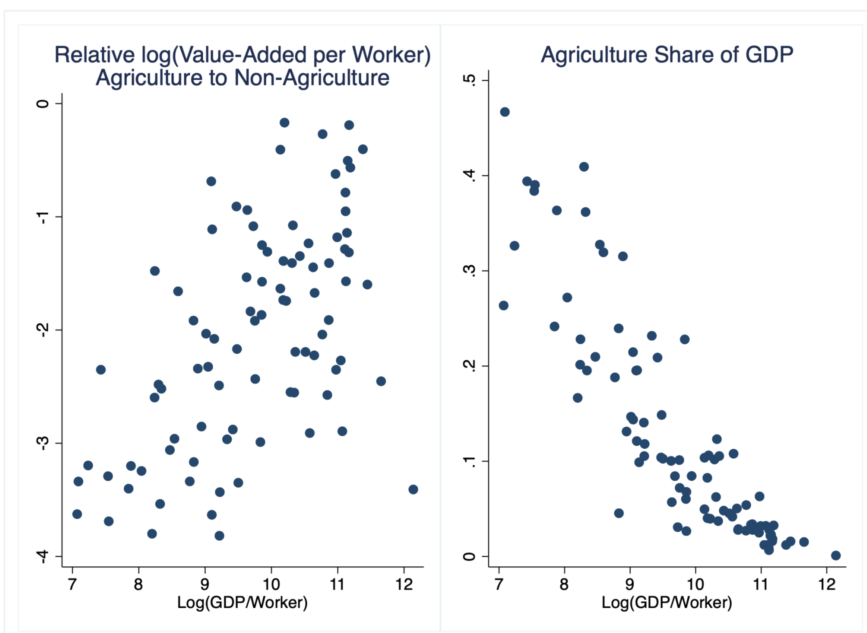

Notes: The map in Panel (a) shows the projected impact of climate change on agricultural productivity averaged across estimates from Hultgren et al. (2021), Cline (2007), Iglesias and Rosenzweig (2010), and Costinot, Donaldson and Smith (2016). See Appendix E.3 for more information on the methods and data used in each of these sources. The graphs in Panel (b) show data from Tombe (2015) indicating that poorer countries specialize heavily in agriculture despite having low agricultural relative to non-agricultural productivity, compared with richer countries. Data on relative value-added per worker in the left graph adjusts for prices for the global cross-section in 2005.

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). Survey datasets include manufacturing firms, and Amadeus data includes both manufacturing and services firms. Data coverage extends from the 3rd to the 99th percentile of the global distribution of per-capita income in 2014, and the 1st to the 90th percentile of long-run average temperature by country. This allows for estimating how the effects of extreme temperatures vary across rich and poor countries and hot and cold countries.

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 estimating 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. Figure 2 evaluates the interaction terms from the results in Column 2 to show the magnitude of temperature effects across places at varying hypothetical levels of income and temperature.

Figure 2: Estimated Global Response of Annual Manufacturing Revenue per Worker to Daily Maximum Temperature

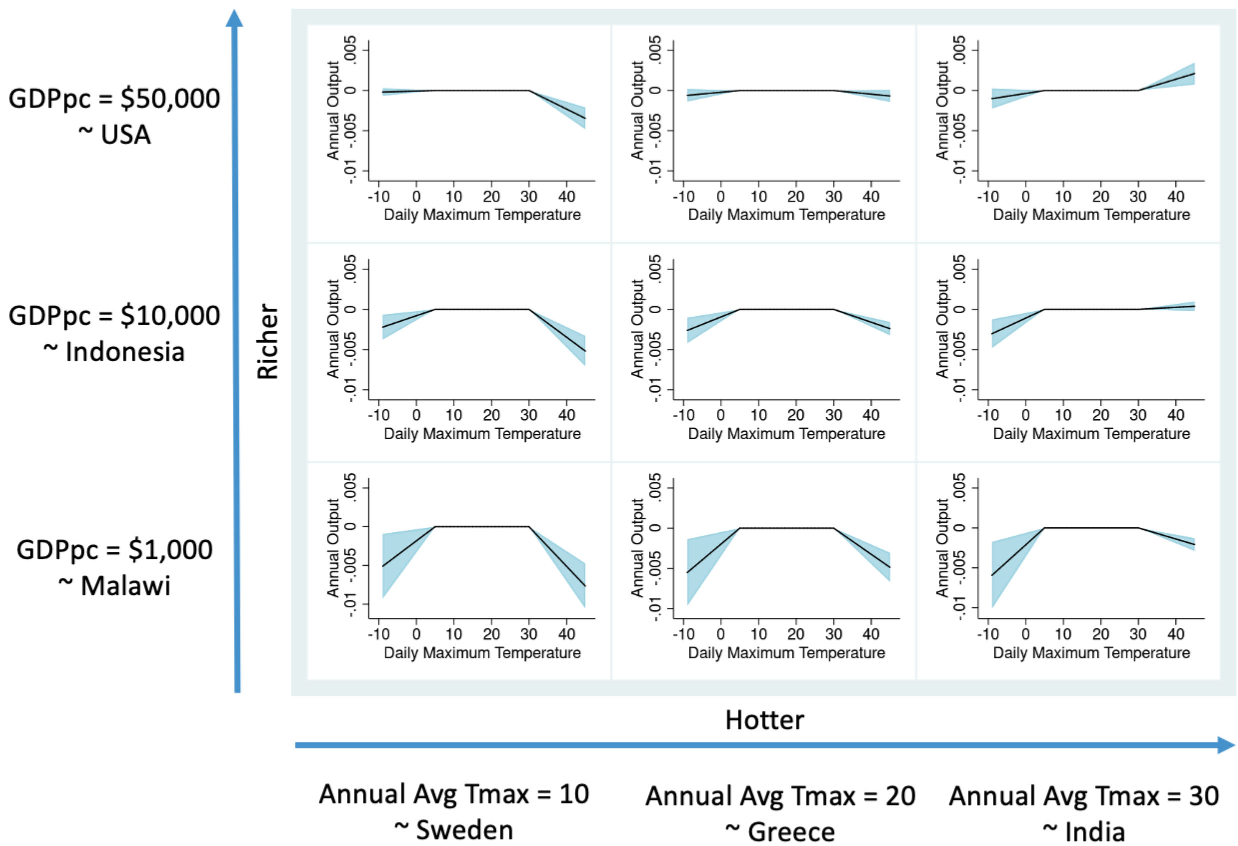

Notes: Graphs show the predicted effect of exposure to daily maximum temperature on the log of firm-level revenue per worker at varying levels of income and long-run average temperature by evaluating the interacted panel regression from Column 2 of Table 2. The specification includes firm and country-by-year fixed effects. 95% confidence intervals are shown in blue, and standard errors are two-way clustered at the firm and county-by-year level. Plots can be interpreted as the effect of moving a single day in the year from the moderate temperature range to the given temperature shown on the x-axis. For instance, for a hypothetical firm in a poor country with a temperate climate shown in the bottom middle cell, the results imply that a single day of exposure to extreme heat or extreme cold can reduce annual revenue per worker by about 0.4%.

Figure 3: Temperature Effects on U.S. Manufacturing

(a) Estimated Response of Annual Plant-Level Revenue per Worker to Daily Maximum Temperature

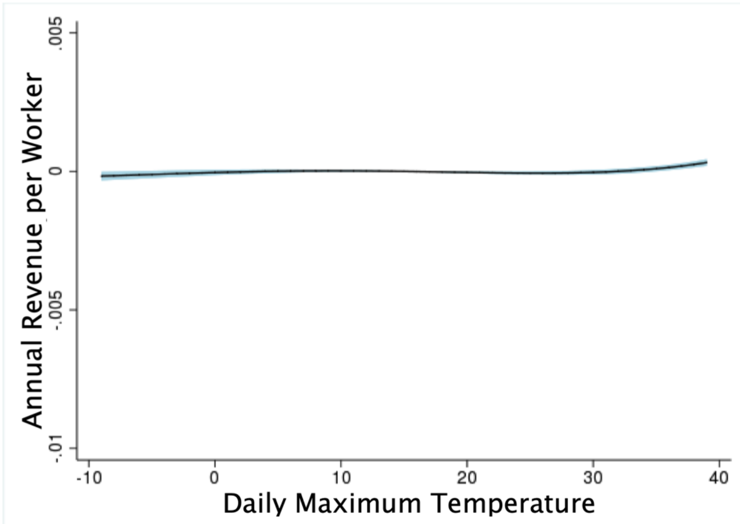

(b) Estimated Response of Annual Plant-Level Energy Expenditures to Daily Maximum Temperature

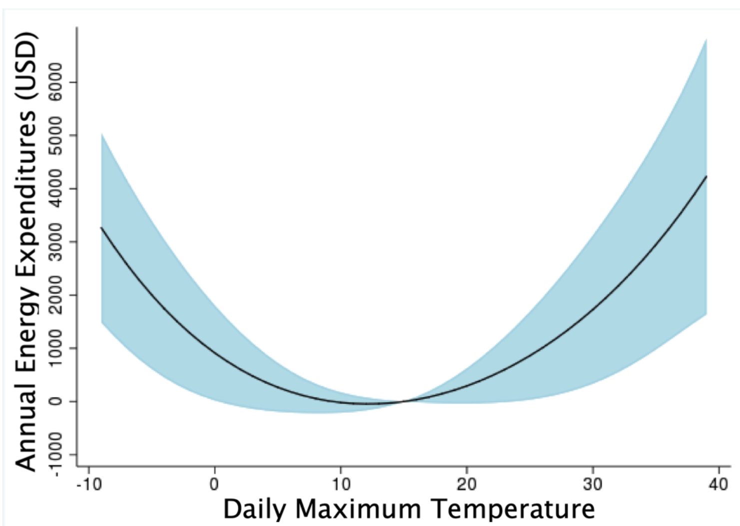

Notes: Panel (a) shows the response of annual revenue per worker to daily maximum temperature for U.S. manufacturing plants estimated using the panel regression specification in Equation 7 with a polynomial of degree four. Panel (b) shows the same specification with plant-level energy expenditures as the dependent variable. Energy expenditures are the sum of cost of fuels and electricity expenditures. Both regressions include plant and year fixed effects. The 95% confidence interval is shown in blue, and standard errors are two-way clustered at the firm and county-by-year level. 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.

Figure 4: Predicted Effect of Extreme Temperatures on Annual Manufacturing Revenue per Worker

(a) 40°C Day

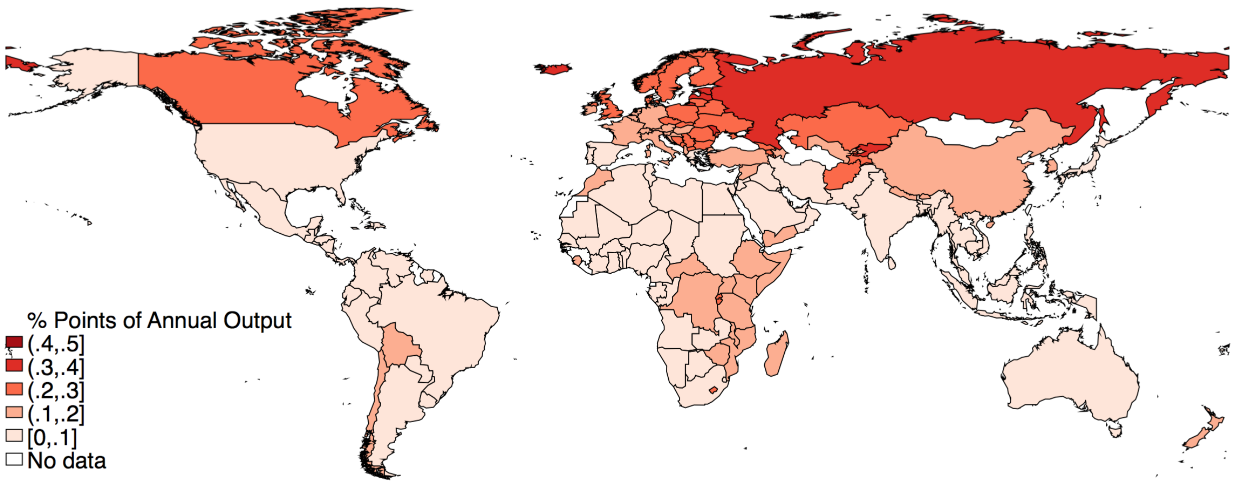

(b) -5°C Day

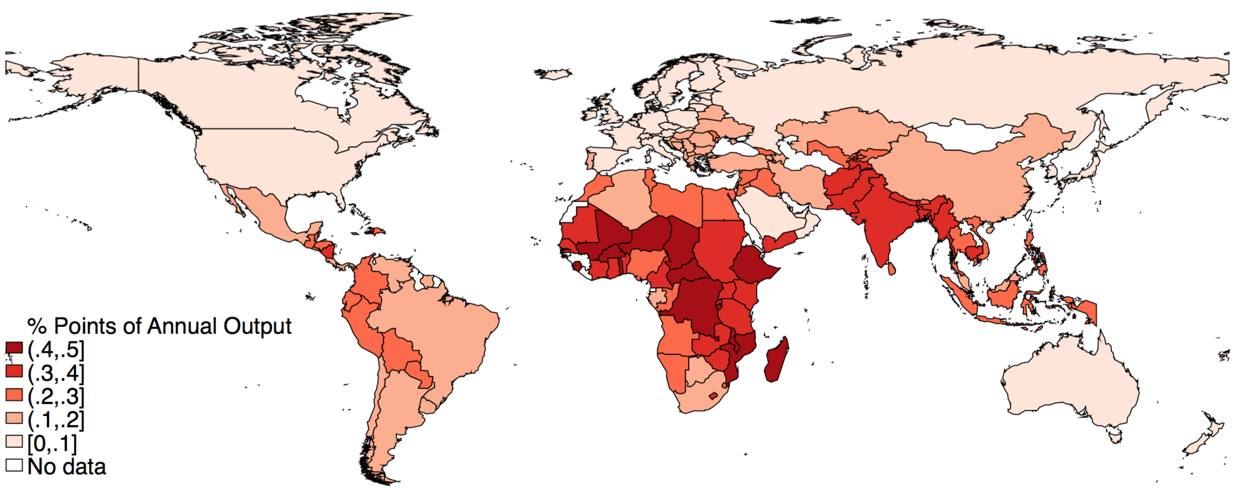

Notes: Maps show predicted annual percentage point loss in revenue per worker from a 40°C day and -5°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's GDP per capita and long-run average temperature. These estimates come from estimating the panel regression specification in Equation 8, which includes firm and country-by-year fixed effects and interacts the effects of temperature with local GDP per capita and long-run average temperature. Table 1 displays the firm-level panel data used in the estimation. Temperature data, both for the estimation and for the projected effects in these maps, comes from GMFD.

Table 3: Model Calibration Summary

(a) Model Parameters and Target Moments

Parameters	Data Moment	Data Source
σ	Sectoral GDP Shares	World Bank
Ω_a, Ω_m, Ω_s	Sectoral GDP Shares	World Bank
ϵ_a, ϵ_m, ϵ_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

(b) 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: Panel (a) shows the data sources for moments targeted in the simulated method of moments procedure presented in Section 5. Data is for the global cross-section in 2011, accessed from the World Bank Databank. Panel (b) shows the estimated values of key consumer preference parameters, which track closely with the estimates presented in Comin, Lashkari and Mestieri (2021). Standard errors in parentheses are 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.

Figure 5: Model Fit Summary

(a) Targeted Moment: Agriculture Share of GDP - Data vs. Simulation

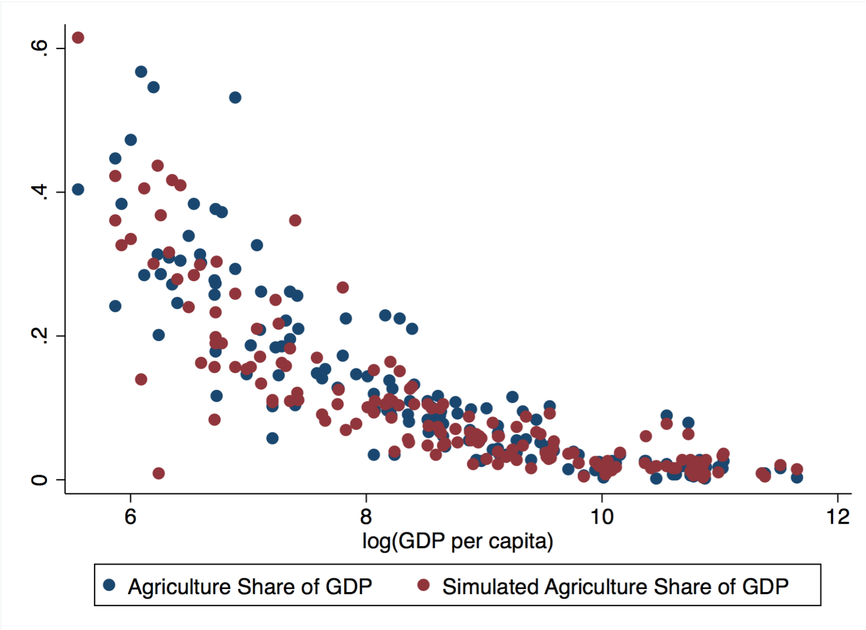

(b) Nontargeted Moment: Relative Price of Food - Data vs. Simulation

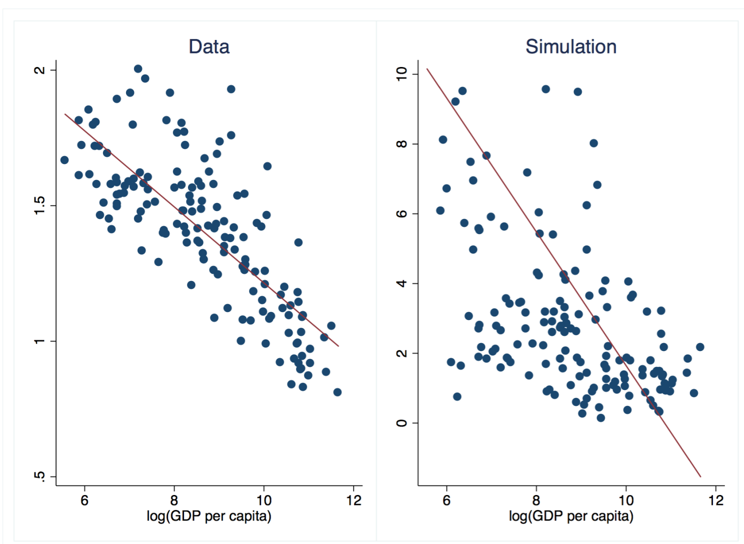

Notes: Panel (a) shows the model's fit to a targeted moment: the agriculture share of GDP across countries. The simulation explains over 60% of the variation in the data, and reproduces the smooth pattern of non-homotheticity observed in the empirical relationship between agricultural shares and income. Panel (b) shows the model's fit to a nontargeted moment: the relative price of food versus non-food. The left graph 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.

Figure 6: Projected Impact of Climate Change on Productivity

(a) Manufacturing

(b) Agriculture Relative to Manufacturing

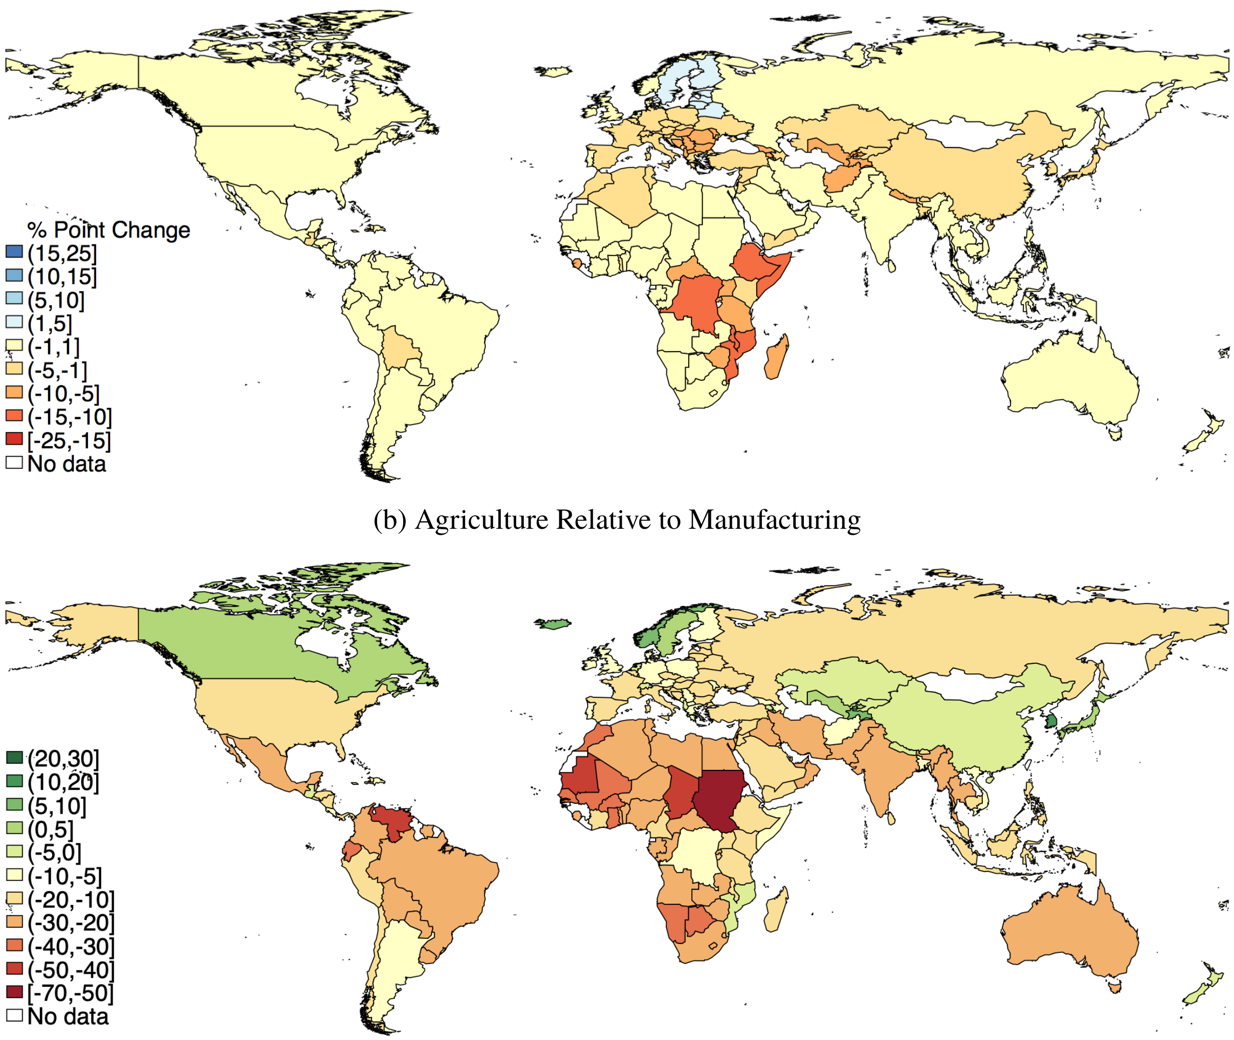

Notes: Panel (a) shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated 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 average change in agricultural productivity from four sources in the literature minus my estimate of the change in manufacturing productivity, shown above, in percentage points. Agricultural productivity impacts come from an average of estimates from Hultgren et al. (2021), Cline (2007), Iglesias and Rosenzweig (2010), and Costinot, Donaldson and Smith (2016), each of which is described in more detail in Appendix E.3. The pattern in Panel (b) shows that hotter parts of the world are likely to suffer much larger declines in agricultural productivity than manufacturing productivity, implying potential gains from reallocation if these places were able to move production away from farming.

Figure 7: Simulated Effects of Climate Change on Trade and Sectoral Reallocation

(a) Change in Agricultural Net Exports Relative to No Climate Change

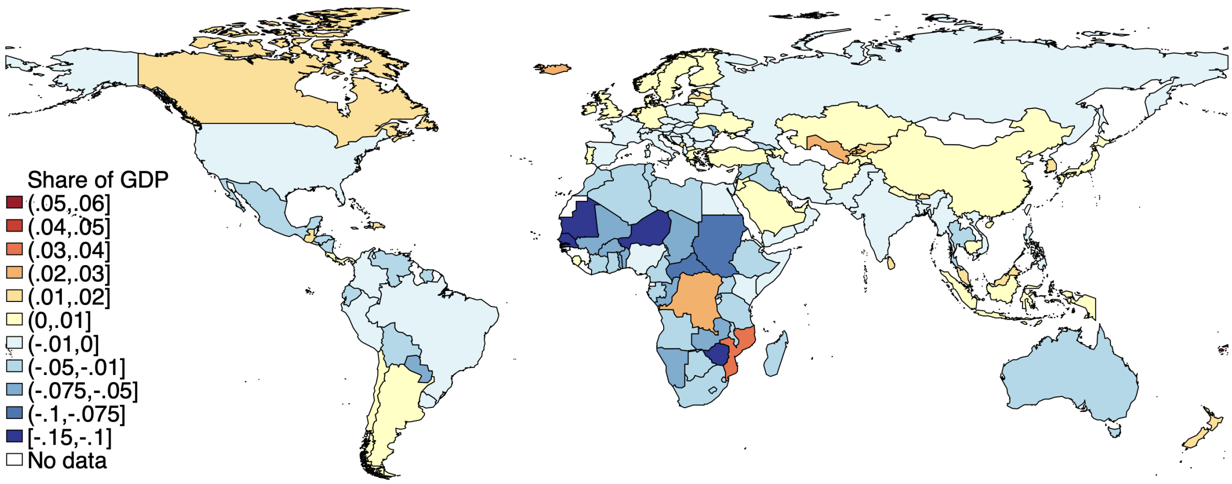

(b) Change in Agricultural GDP Share Relative to No Climate Change

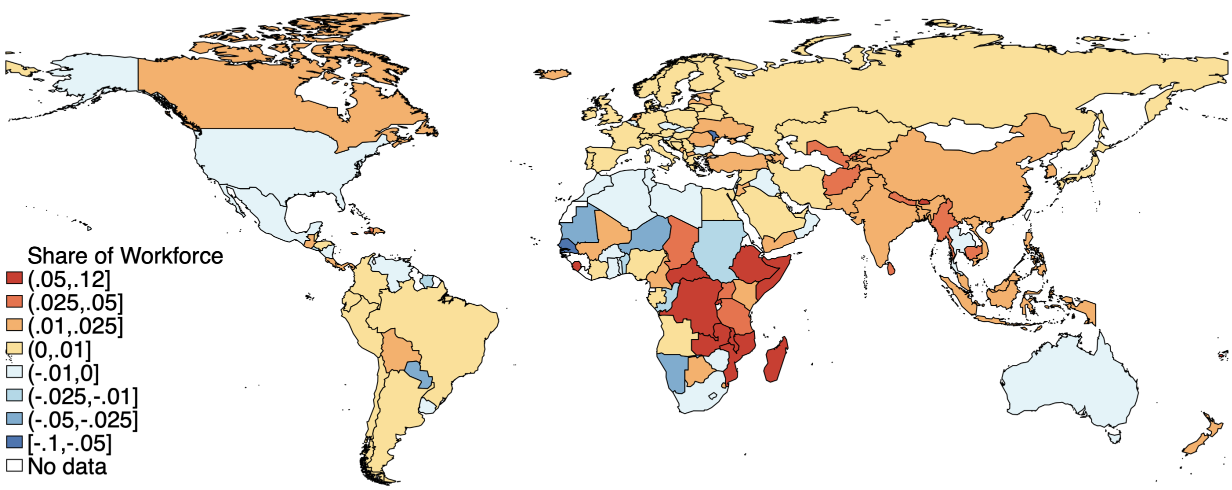

Notes: Panel (a) shows model simulations of the impact of climate change on agricultural net exports as a share of GDP driven by the effects of climate change on sector-level productivity and comparative advantage shown in Figure 6b. On average, hotter countries move modestly toward importing more food and colder countries move modestly toward exporting more food. Panel (b) shows model simulations of the impact of climate change on agricultural GDP share. As shown in Equation 26, this reallocation results from the net effect of the change in trade flows and the change in agriculture's expenditure share as incomes and prices change. The map shows that many poor countries increase their production and labor shares in agriculture even as agricultural productivity in these places falls sharply because the increase in food imports is not sufficient to meet domestic demand for food.

Table 4: Impact of Climate Change With and Without Reallocation

(a) Simulated Agriculture Share of GDP by Scenario

Country	No Reallocation	Autarky	Full Reallocation	Relative Agricultural vs. Manufacturing Productivity Change
Ethiopia	.359	.445	.419	-.234
Ghana	.181	.227	.177	-.365
India	.161	.191	.185	-.216
Mozambique	.367	.418	.453	-.091
Niger	.325	.438	.292	-.353
Norway	.02	.019	.029	.062
Zambia	.36	.481	.414	-.364
Poorest Quartile	.199	.238	.226	-.206
World	.038	.043	.043	-.166

(b) Impact of Climate Change on GDP by Scenario

Country	No Reallocation	Autarky	Full Reallocation	Relative Agricultural vs. Manufacturing Productivity Change
Ethiopia	-.172	-.233	-.232	-.234
Ghana	-.054	-.074	-.071	-.365
India	-.038	-.064	-.064	-.216
Mozambique	-.13	-.175	-.186	-.091
Niger	-.147	-.241	-.194	-.353
Norway	-.004	-.004	-.003	.062
Zambia	-.159	-.297	-.287	-.364
Poorest Quartile	-.063	-.094	-.093	-.206
World	-.018	-.021	-.021	-.166

(c) Impact of Climate Change on Welfare (Equivalent Variation as a Share of Income) by Scenario

Country	No Reallocation	Autarky	Full Reallocation	Relative Agricultural vs. Manufacturing Productivity Change
Ethiopia	-.391	-.18	-.18	-.234
Ghana	-.244	-.059	-.059	-.365
India	-.146	-.039	-.04	-.216
Mozambique	-.252	-.131	-.138	-.091
Niger	-.417	-.169	-.13	-.353
Norway	-.002	-.004	-.002	.062
Zambia	-.44	-.184	-.179	-.364
Poorest Quartile	-.184	-.065	-.064	-.206
World	-.038	-.017	-.017	-.166

Notes: Table shows model simulations of sectoral reallocation, welfare, and measured GDP for a selection of countries in counterfactuals that allow for no reallocation, a change in expenditure shares only ("Autarky"), and a change in expenditure shares and trade flows ("Full Reallocation"). Rows marked "Poorest Quartile" show population-weighted outcomes for the poorest quartile. Rows marked "World" show global totals as a share of GDP (equivalently a GDP-weighted average). The right-most column shows the relative impact of climate change on agricultural versus manufacturing productivity from Figure 6b to benchmark the change in comparative advantage in each country.

Figure 8: Global Welfare Impact of Climate Change Productivity Effects

(a) Impact of Climate Change on Welfare (Equivalent Variation)

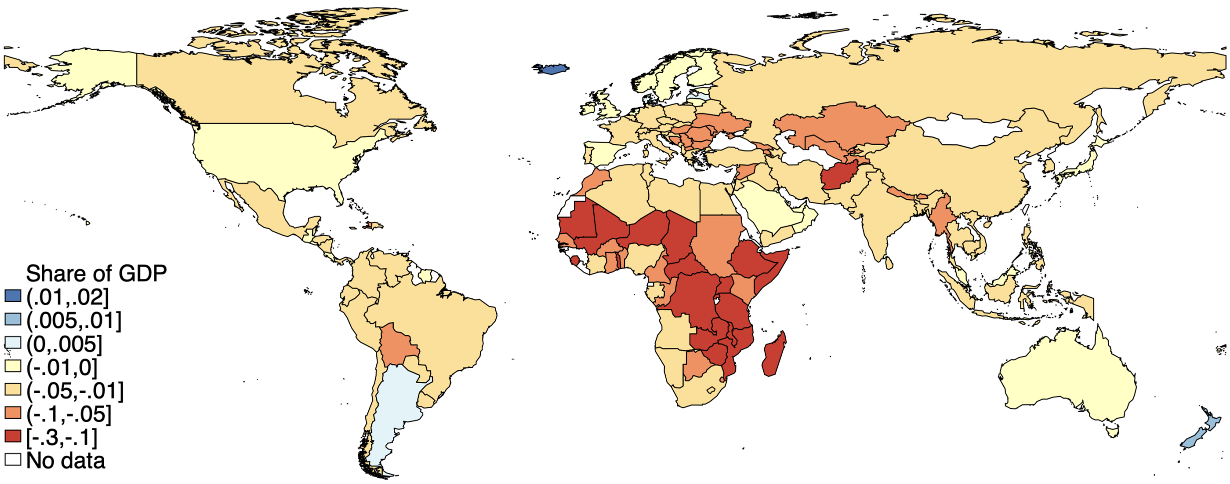

(b) Projected Impact of Climate Change on Domestic Food Price Index (Percentage Points)

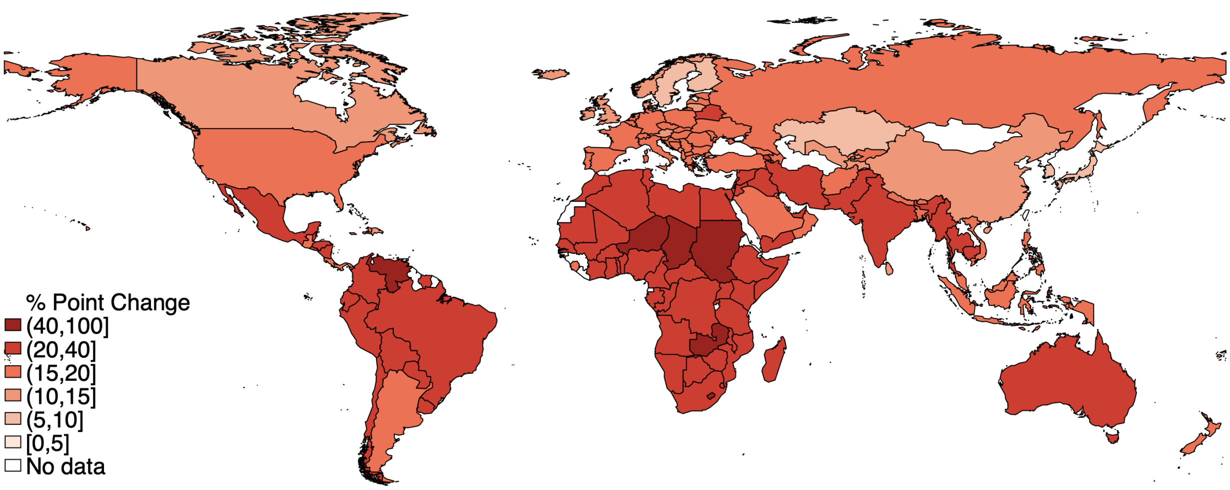

Notes: Panel (a) shows the equivalent variation welfare impacts of the productivity effects of climate change in each country in the model simulation. Panel (b) shows the corresponding impacts on the aggregate country-level food price index, P_{ak}, a key driver of the welfare effects in the simulation. Note that the effects here do not account for a variety of other pathways for climate change impacts, such as health effects, sea-level rise, or risk-aversion to low probability catastrophic outcomes. The results shown in these maps are from the simulation that uses current per-capita income as the baseline and estimated levels of trade costs. The results in Table 6 and Table 7 show how the welfare effects differ when the baseline is adjusted to allow for various projections of future economic growth, and the results in Table 5 and Table 7 show how the effects vary for different hypothetical levels of trade costs.

Table 5: Impact of Climate Change under Alternative Trade Cost Scenarios
Selected Countries

(a) Equivalent Variation

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Central African Republic	-.271	-.271	-.050
Malawi	-.263	-.263	-.137
Zimbabwe	-.246	-.227	-.084
Rwanda	-.205	-.204	-.101
Sierra Leone	-.162	-.194	-.123
Democratic Republic of Congo	-.196	-.194	-.138
Ethiopia	-.18	-.18	-.108
Zambia	-.184	-.179	-.008
Somalia	-.169	-.169	-.123
Poorest Quartile	-.065	-.064	-.037
World	-.017	-.017	-.015

(b) Impact on Food Prices

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Central African Republic	.784	.294	.193
Malawi	.579	.298	.193
Zimbabwe	.713	.379	.193
Rwanda	.535	.317	.193
Sierra Leone	.692	.170	.193
Democratic Republic of Congo	.361	.213	.193
Ethiopia	.506	.363	.195
Zambia	.573	.431	.193
Somalia	.351	.253	.193
Poorest Quartile	.300	.268	.197
World	.239	.218	.192

Notes: Panel (a) shows model simulations of the equivalent variation welfare effects of climate change under three different scenarios: autarky, in which only expenditure shares, but not trade flows, are allowed to adjust; the equilibrium counterfactual allowing full adjustment under existing trade barriers calibrated to match the data on trade flows for each importer-exporter-sector pairing, and an alternative "low trade cost" scenario that sets all bilateral trade costs to a 100% tariff-equivalent ($\tau = 2$), which is approximately the 90th percentile level of trade openness observed in the calibration. Panel (b) shows the corresponding change in the aggregate food price index, P_{ak} - a key component of the welfare effects - that rises by over 100% in some countries. The table focuses on a subset of countries with the greatest vulnerability to climate change. Table 7 contains summary results for a variety of additional trade cost scenarios.

Table 6: Impact of Climate Change from Alternative Baseline Incomes
Selected Countries

(a) Baseline and Counterfactual Agriculture Shares of GDP

Country	Projected GDP Per-Capita 2080 / Present in SSP 3	Agriculture GDP Share Present Baseline	Agriculture GDP Share 2080 Baseline	Agriculture GDP Share 2080 Climate Change Counterfactual
Central African Republic	1.47	.299	.287	.226
Malawi	2.84	.436	.309	.359
Zimbabwe	4.17	.302	.122	.125
Sierra Leone	1.49	.139	.177	.159
Democratic Republic of Congo	10.19	.421	.151	.177
Somalia	2.27	.334	.201	.220
Zambia	1.52	.360	.280	.325
Rwanda	1.14	.409	.390	.513
Ethiopia	1.23	.359	.333	.384
Poorest Quartile	3.05	.199	.126	.145
World	2.20	.038	.025	.028

(b) Equivalent Variation from Present and Future Income Baselines

Country	Projected GDP Per-Capita 2080 / Present in SSP 3	Equivalent Variation from Present Baseline	Equivalent Variation from 2080 Baseline
Central African Republic	1.47	-.333	-.249
Malawi	2.84	-.283	-.188
Zimbabwe	4.17	-.265	-.117
Sierra Leone	1.49	-.234	-.157
Democratic Republic of Congo	10.19	-.223	-.050
Somalia	2.27	-.221	-.171
Zambia	1.52	-.212	-.164
Rwanda	1.14	-.209	-.200
Ethiopia	1.23	-.200	-.185
Poorest Quartile	3.05	-.076	-.047
World	2.20	-.028	-.015

Notes: This table shows how allowing for baseline economic growth affects the impacts of climate change. The simulation presented here adjusts baseline income by taking the country-level growth projections from Shared Socioeconomic Pathway Scenario Three developed by Cuaresma (2017) and applying these as sector-neutral productivity growth in the model. The impacts of climate change are then calculated by comparing the future baseline equilibrium to a counterfactual in which sector-country productivities are affected by the climate change impacts projected in Figure 6b. Panel (a) shows that baseline agriculture share of GDP falls as countries grow richer in the future due to nonhomothetic preferences, but that climate change still increases agricultural specialization on average in poor countries relative to the no climate change baseline. Panel (b) shows that the willingness to pay to avoid climate change is lower when allowing for baseline economic growth, though it remains substantial in poor countries. Note that the willingness-to-pay presented here includes the costs of firm-level adaptation, calculated as described in Appendix D. For results for other economic growth scenarios, see Table 7.

Table 7: Impacts of Climate Change - Model Simulation Robustness

Model Scenario	Δ Ag Labor Share	Δ GDP	Equivalent Variation	Δ Food Prices
Robustness of Main Results				
Baseline Results				
World	.005	-.021	-.017	.218
Poorest Quartile	.027	-.093	-.064	.268
Stone-Geary Preferences				
World	.003	-.018	-.016	.217
Poorest Quartile	.031	-.085	-.057	.266
Hultgren et al. Agriculture Estimates				
World	.005	-.02	-.017	.203
Poorest Quartile	.021	-.058	-.043	.167
Iglesias-Rosenzweig Agriculture Estimates				
World	.005	-.021	-.017	.216
Poorest Quartile	.029	-.091	-.063	.263
Cline Agriculture Estimates				
World	.005	-.021	-.017	.223
Poorest Quartile	.028	-.126	-.088	.377
Alternative Low Trade Cost Counterfactual Cases				
Reduce All Trade Costs by Half				
World	.004	-.02	-.017	.189
Poorest Quartile	.014	-.085	-.06	.235
Reduce All Trade Costs to $\tau = 2$ (90th Percentile Openness)				
World	.003	-.017	-.015	.191
Poorest Quartile	-.003	-.045	-.037	.197
Reduce All Trade Costs to $\tau = 1$ (Frictionless Trade)				
World	.002	-.015	-.013	.199
Poorest Quartile	-.002	-.028	-.027	.199
Alternative Economic Growth Scenarios				
SSP 1 - Moderate Growth, High Convergence				
World	.003	-.01	-.011	.184
Poorest Quartile	.009	-.029	-.017	.275
SSP 2 - Moderate Growth, Moderate Convergence				
World	.003	-.01	-.012	.182
Poorest Quartile	.012	-.036	-.023	.27
SSP 3 - Low Growth, Low Convergence				
World	.003	-.012	-.016	.178
Poorest Quartile	.019	-.062	-.047	.255
SSP 4 - Moderate Growth, Low Convergence				
World	.003	-.011	-.014	.179
Poorest Quartile	.016	-.052	-.036	.261
SSP 5 - High Growth, High Convergence				
World	.002	-.009	-.008	.185
Poorest Quartile	.007	-.023	-.012	.28

Notes: Table summarizes results under a variety of assumptions and scenarios. Top panel shows robustness to a different consumer preference specification and alternative sources for the impact of climate change on agricultural productivity. Middle panel shows how results vary with assumptions about trade costs. Bottom panel shows how results vary with assumptions about baseline economic growth. For more details on each of these, see Section 7.6.

Table 8: Country-Level Panel Regressions on Sectoral Reallocation

(a) Country-Level Panel Data

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

(b) Country-Level Panel Regression Results

	(1) log(GDP)	(2) Food Share of Imports	(3) Ag Share of GDP	(4) 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 degree days (GDD) and killing degree days (KDD). GDD measure 24 hour increases of 1°C between 0 and 29°C, typically helpful for crops, and KDD contain the corresponding measure for increases above 29°C, typically harmful for crops. Data summarized in Panel (a) 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.

Figure 9: Non-Tariff Barriers to Trade

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

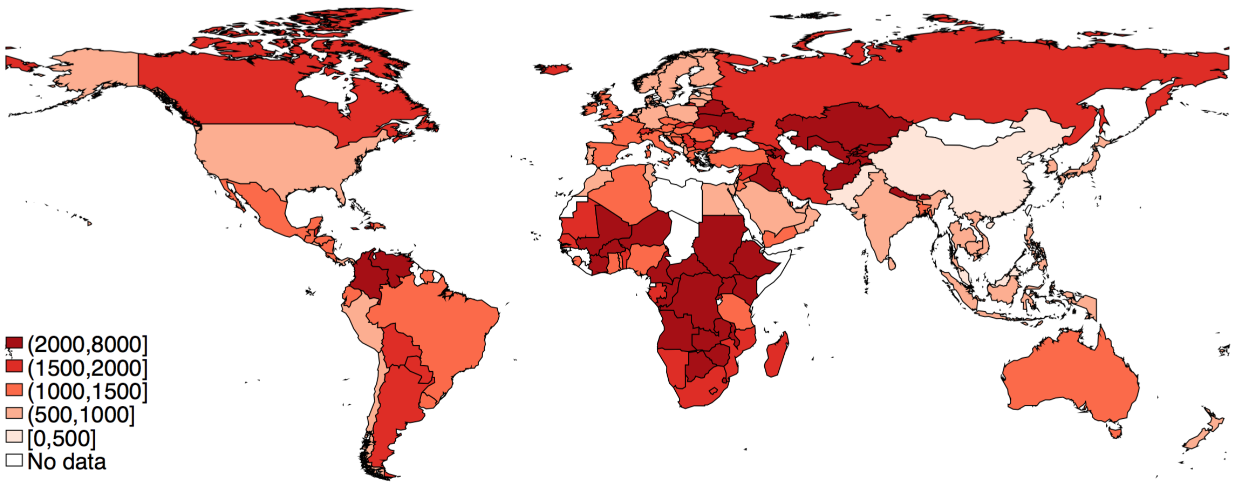

(b) Days to Import a Container

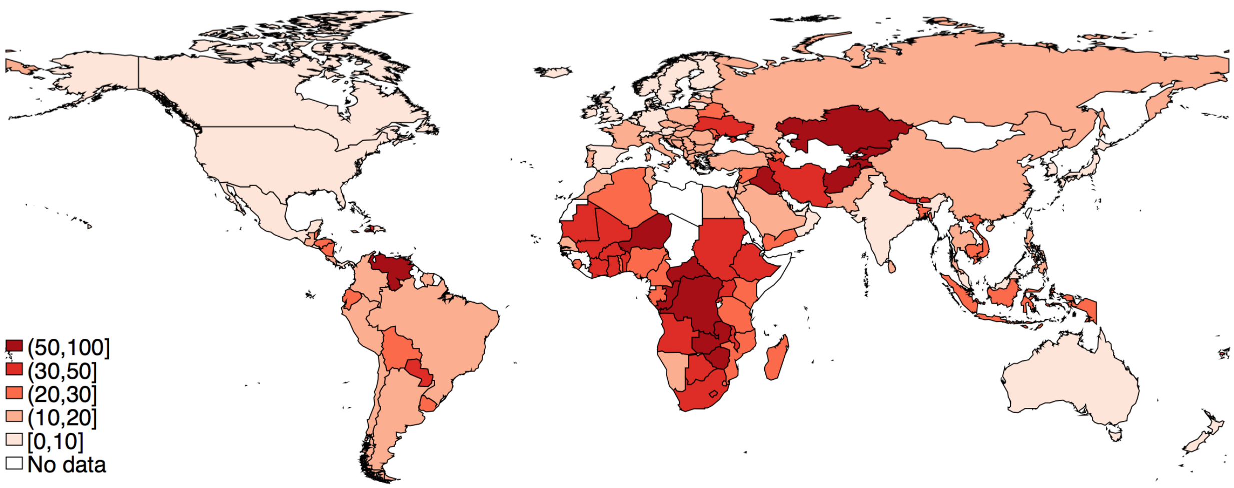

Notes: This figure shows possible underlying causes of the high barriers to trade calibrated in the model to match the low levels of trade flows in developing countries. Panel (a) shows the direct cost to import one container of goods in US dollars. Costs include documents, administrative fees for customs 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 for both panels comes from the World Bank Ease of Doing Business Index.

Appendix A: Further Information On Empirical Estimation

Appendix A.1: Data Construction

Amadeus Data Collection: 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.

BVD collects data from many countries around the world in their Amadeus and Orbis series, but I restrict my analysis to those European countries that have mandatory reporting requirements and thus contain comprehensive nationally representative samples according to Bloom, Draca and Van Reenen (2016). Denmark, Ireland, and Portugal are additional countries with mandatory reporting requirements that were unavailable to me due to data licensing restrictions (Denmark) and missing or outdated geographic identifiers (Ireland and Portugal).

I drop a small proportion of firms marked mining, construction, utilities, and agriculture, though results are very similar when including these firms in the pooled sample.

Merge Details: I merge firm-level data to climate data at the county level. Government surveys provide county information for each firm directly. The Amadeus data provides city name and zip code, which I match to the county level using the GeoPostcodes dataset from GeoData Limited. 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. For some countries, the administrative unit to which I aggregate is more comparable to a town than a county.

Appendix A.2: Additional Regression Results

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

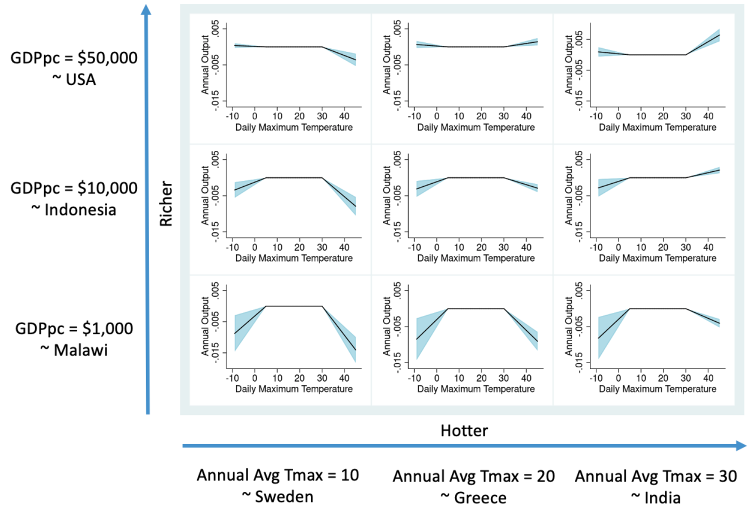

Notes: Figure shows the predicted effect of temperature on the log of manufacturing revenues at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 3 of Table 2. 95% confidence intervals are shown in blue. Outcome variables come from data sources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

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

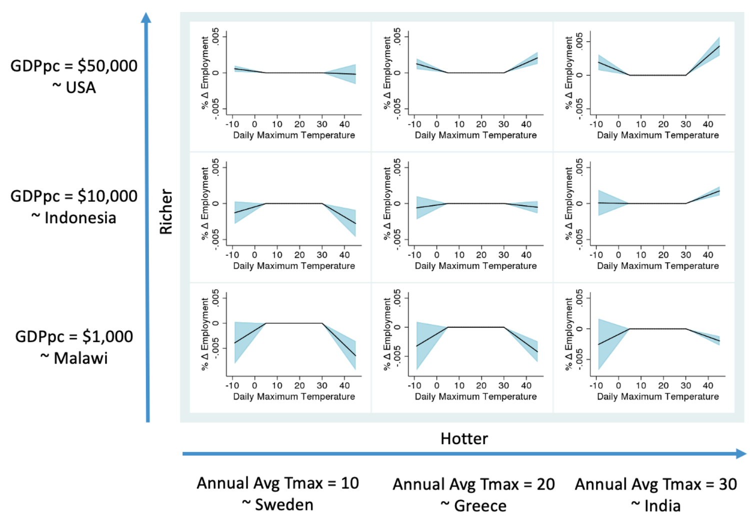

Notes: Figure shows the predicted effect of temperature on the log of manufacturing employment at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 4 of Table 2. 95% confidence intervals are shown in blue. Outcome variables come from data sources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

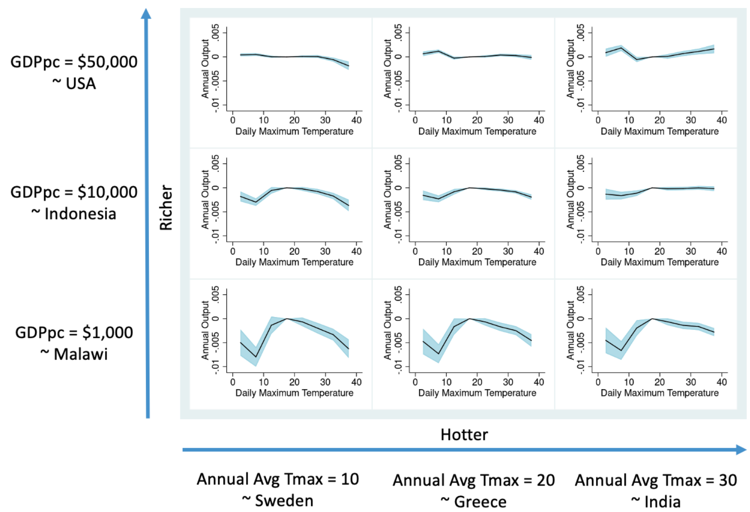

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using bins of daily maximum temperature in the specification from Equation 8. Days are grouped into 5°C bins. 95% confidence intervals are shown in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

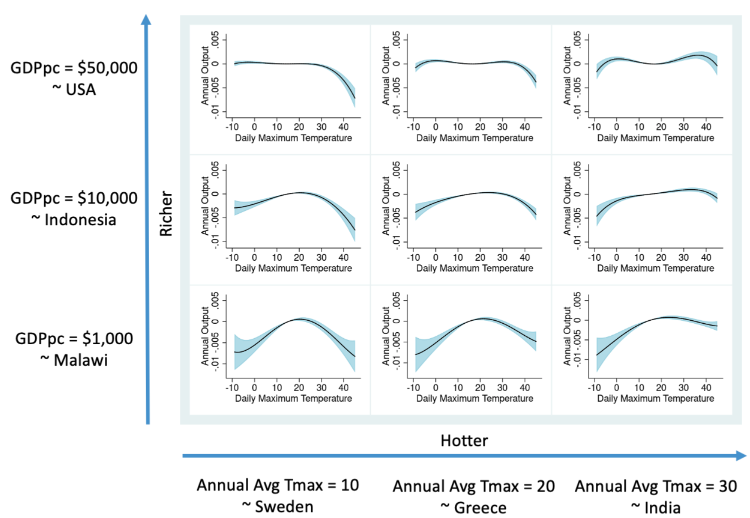

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using a polynomial of degree four in daily average temperature in the specification from Equation 8. 95% confidence intervals are shown in blue. Outcome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

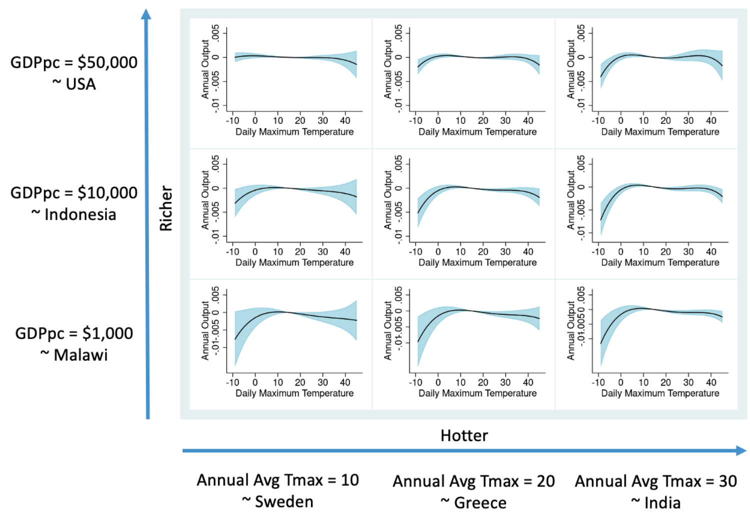

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. 95% confidence 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

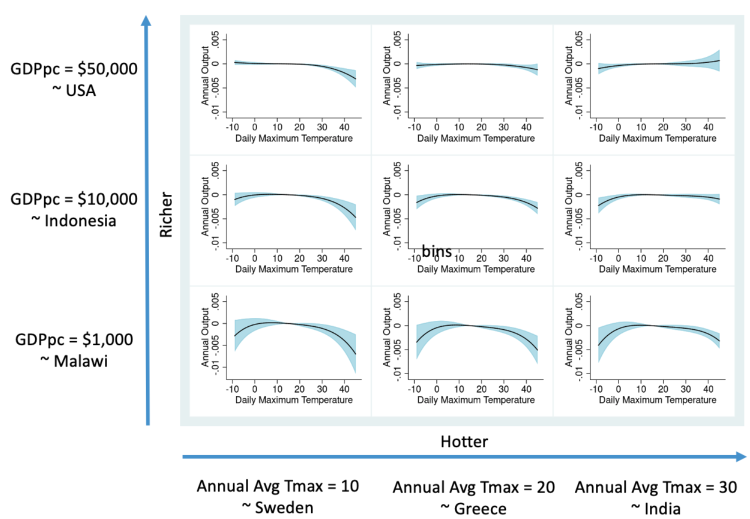

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in 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

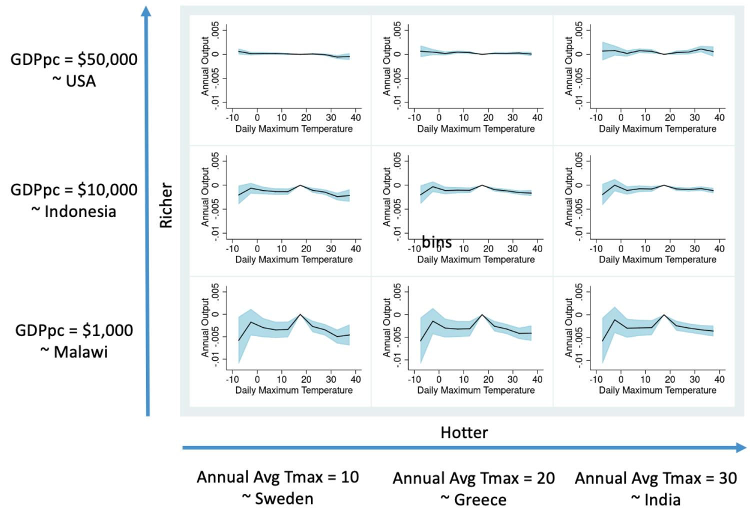

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and bins of daily maximum temperature. 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

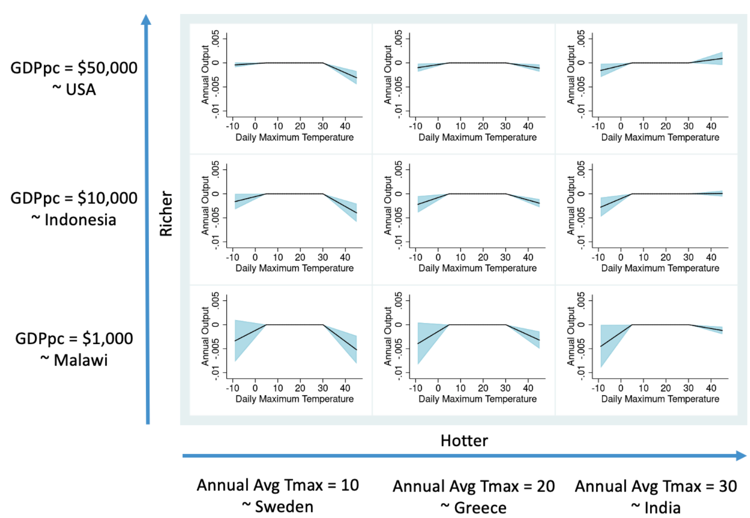

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with controls for capital. 95% confidence intervals are shown in blue. Outcome variables come from data sources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

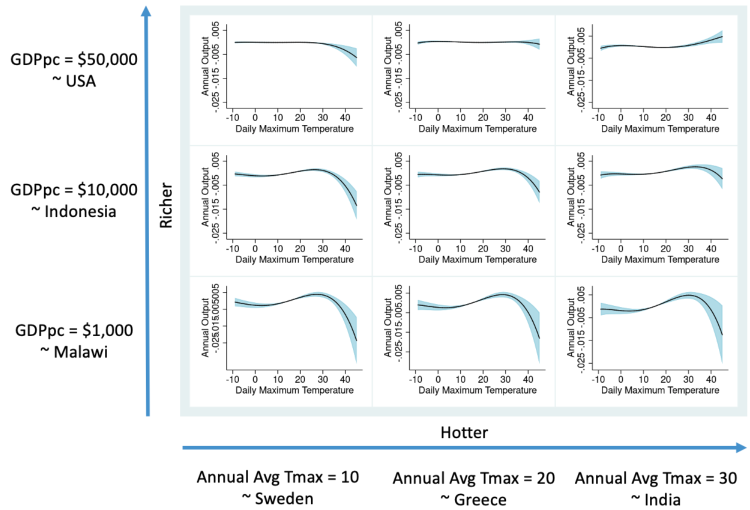

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome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

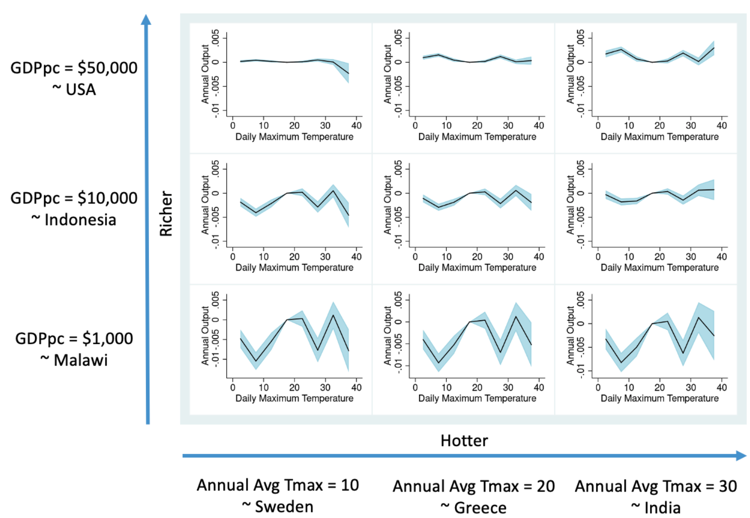

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with bins of daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data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

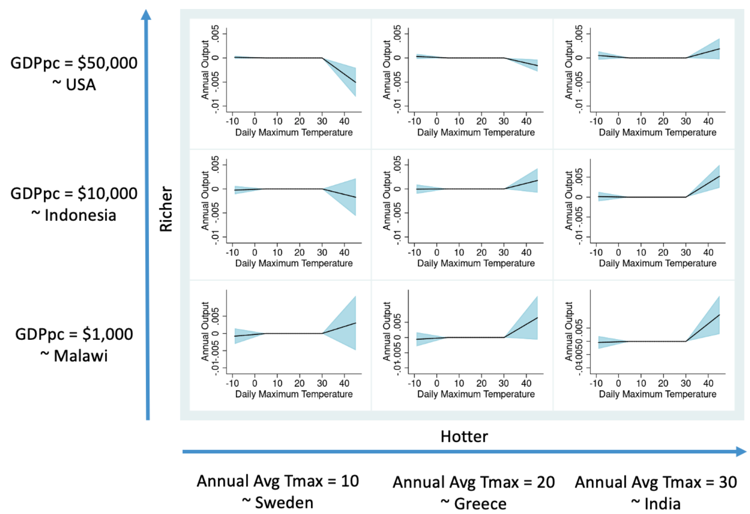

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Table 1 and temperature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

Table A-1: Country-Level Panel Regression

	(1) log(GDP)	(2) Food Share of Imports	(3) Ag Share of GDP	(4) 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: This regression is unweighted version of the one contained in Table 8b. 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 degree days (GDD) and killing degree days (KDD). GDD measure 24 hour increases of 1°C between 0 and 29°C, typically helpful for crops, and KDD contain the corresponding measure for increases above 29°C, typically harmful for crops. 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

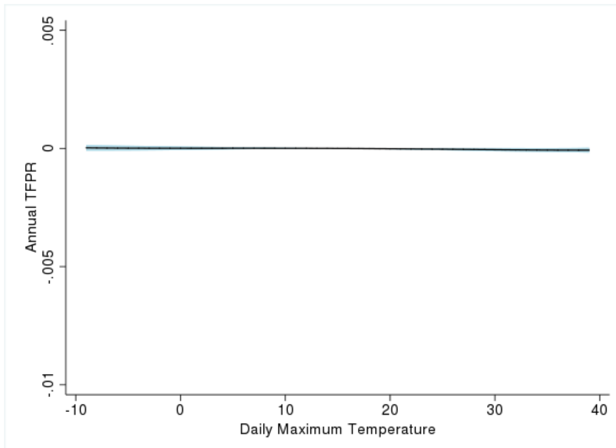

Notes: Figure shows the estimated effect of temperature on manufacturing TFPR using the specification from Equation 7 with a polynomial of degree four in daily maximum temperature. 95% confidence interval is shown in blue. Outcome data comes from the Annual Survey of Manufacturers and Census of Manufacturers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

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 the broader findings in this paper. 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-13. Notably, I am also able to use the main results from the rest of my global data in Figure 2 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-14. While I slightly overpredict sensitivity to cold and underpredict sensitivity to heat, these results are broadly consistent with their findings, lending external validity to this paper's findings. 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-15. 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."

Figure A-13: China Replication - Overlapping Years

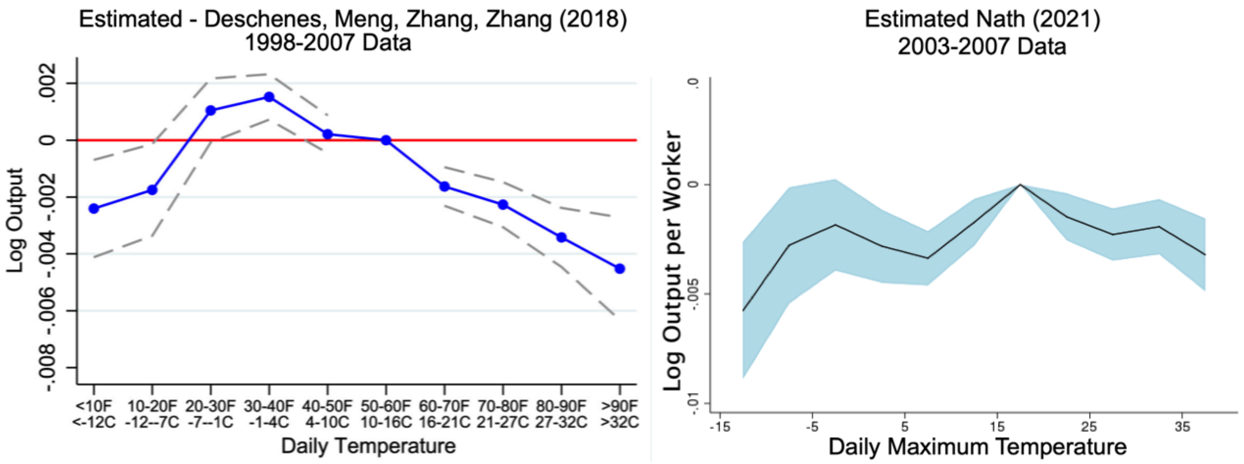

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result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

Figure A-14: China Manufacturing Temperature Sensitivity - Estimated and Predicted

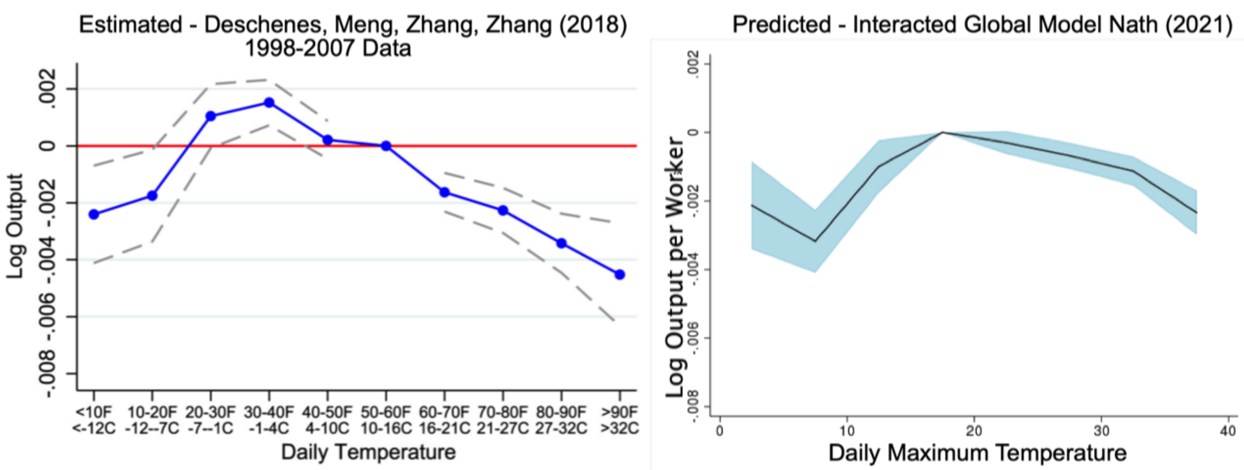

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows the predicted effect of temperature 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-15: China Replication - Different Years

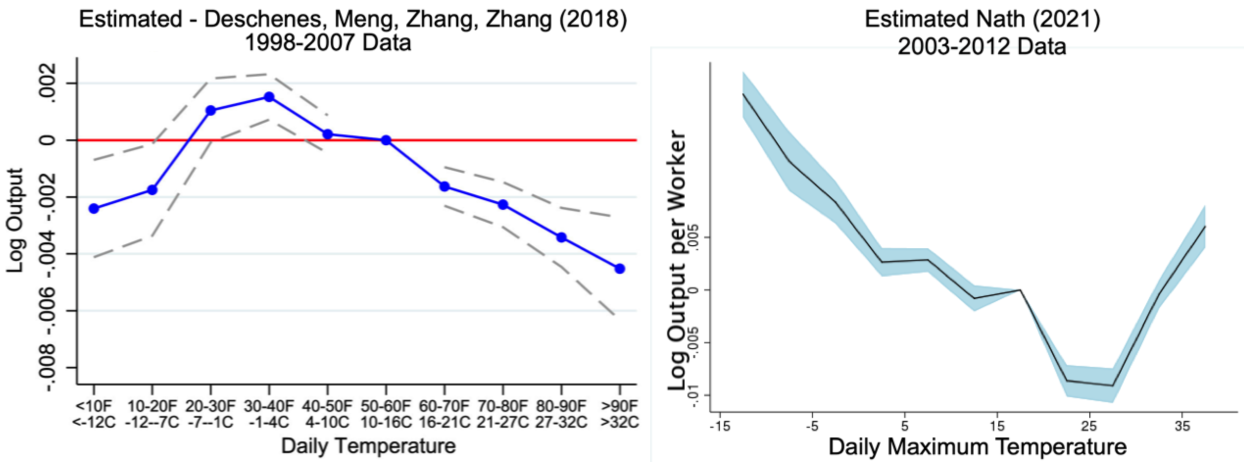

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result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

A somewhat puzzling fact is that these 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 the 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 the central analysis in this paper.

Appendix D: Adaptation Benefits and Costs

In this section I explain how I use revealed preference methods developed by Carleton et al. (2020) to infer the costs firms incur from reducing the sensitivity of production to extreme temperatures as their expectations adjust to global warming. 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 the 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.

This approach to calculating adaptation costs is subject to several assumptions, among them that all firms across the world face a common cost function for adaptation technologies that is invariant to local conditions, and that firms optimize their adaptation decisions on the margin in the long run. See Carleton et al. (2020) for a more detailed description of the assumptions under which Equation 28 recovers a valid estimate of adaptation costs.

In addition to the costs, we can also calculate 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-16 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 4a. Figure A-17 show the costs of achieving this reduced sensitivity, as calculated using Equation 28, and Figure A-18 show the net benefits of firms adapting to changes in expected exposure to extreme heat.

Figure A-16: Predicted Effect of a 40°C Day on
Annual Manufacturing Revenue per Worker
At 2080 Average Temperatures

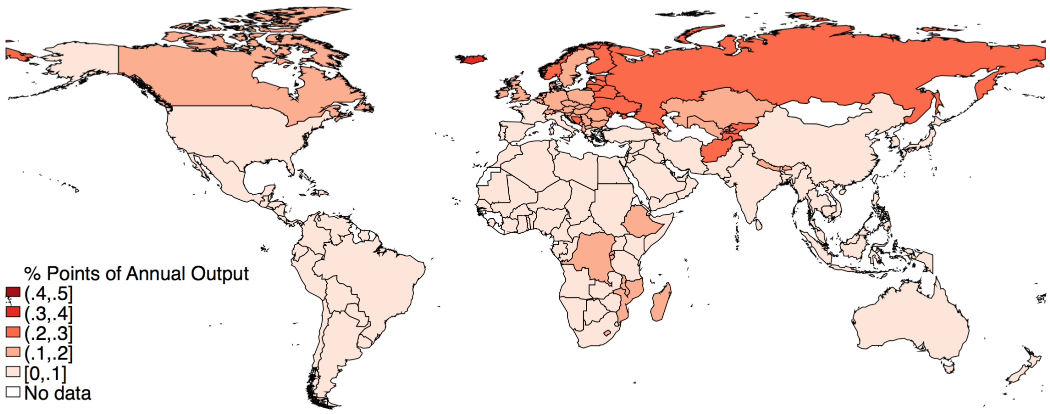

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's level of income and end-of-century long-run average temperature. Temperature sensitivities are lower in this figure than in Figure 4a because the results predict that firms will adapt to hot temperatures as the world warms.

Figure A-17: Firm-Level Adaptation Costs
(Share of Manufacturing Output)

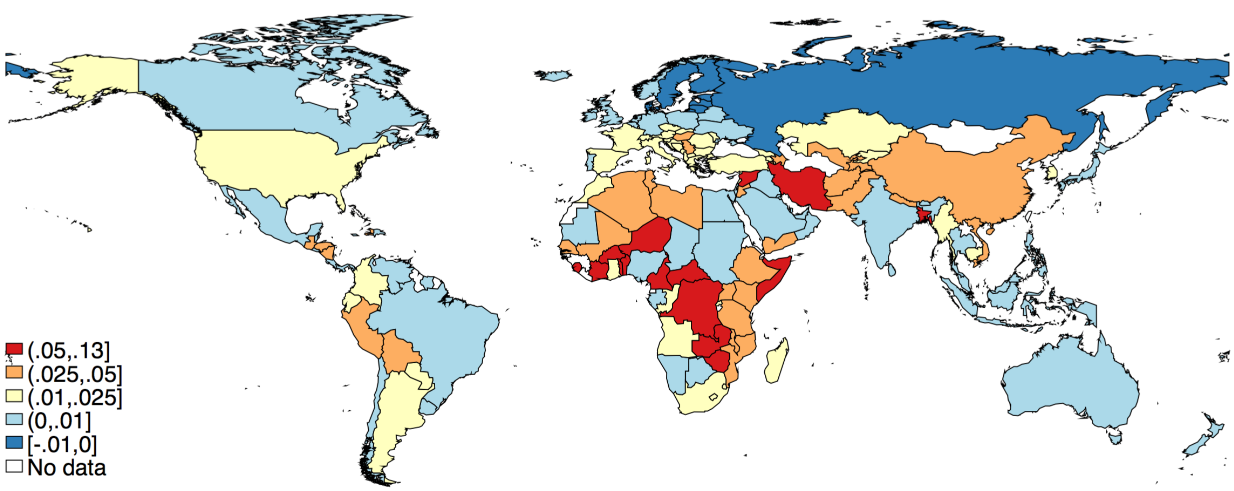

Notes: Map shows the calculations of the costs firms pay to achieve the lower temperature sensitivity shown in Appendix Figure A-16 compared to Figure 4a. I infer these costs using a revealed preference approach developed by Carleton et al. (2020) that infers adaptation costs from the foregone benefits firms would have attained by reducing their heat sensitivity. The procedure is detailed in Appendix D.

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

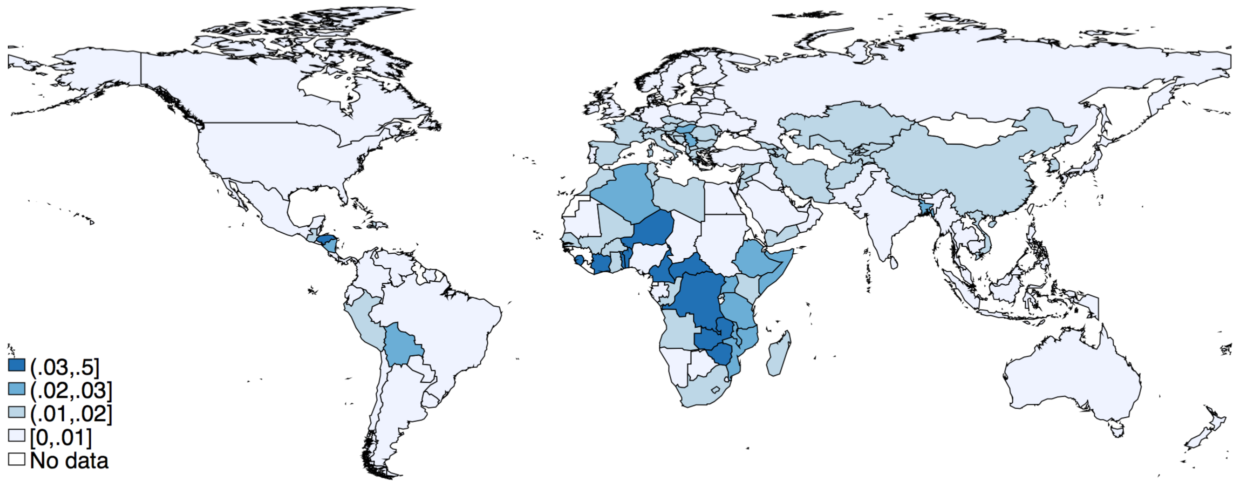

Notes: Map shows the calculations of the net benefits firms receive by investing to reduce their heat sensitivity as the climate warms. The benefits come from reducing heat sensitivity to the level shown in Appendix Figure A-16 compared to the original level in Figure 4a. The inferred costs are shown in Appendix Figure A-17. The procedure to calculate these costs and benefits is detailed in Appendix D.

Figure A-19: Predicted Effect of a 40°C Day on Annual Services Revenue per Worker

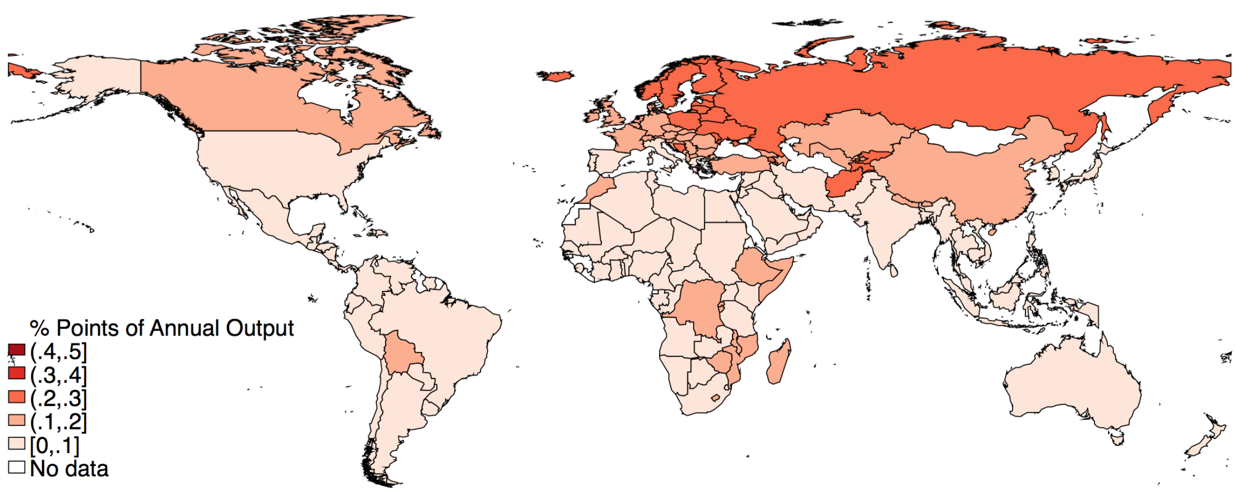

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-20: Predicted Effect of a -5°C Day on Annual Services Revenue per Worker

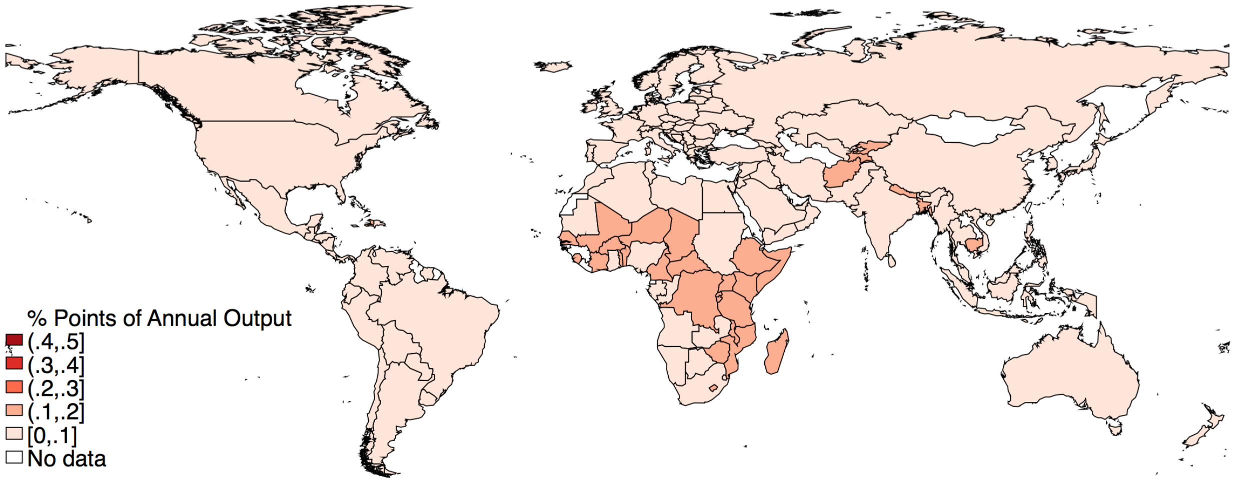

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a -5°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-21: Projected Change in Exposure to Extreme Heat

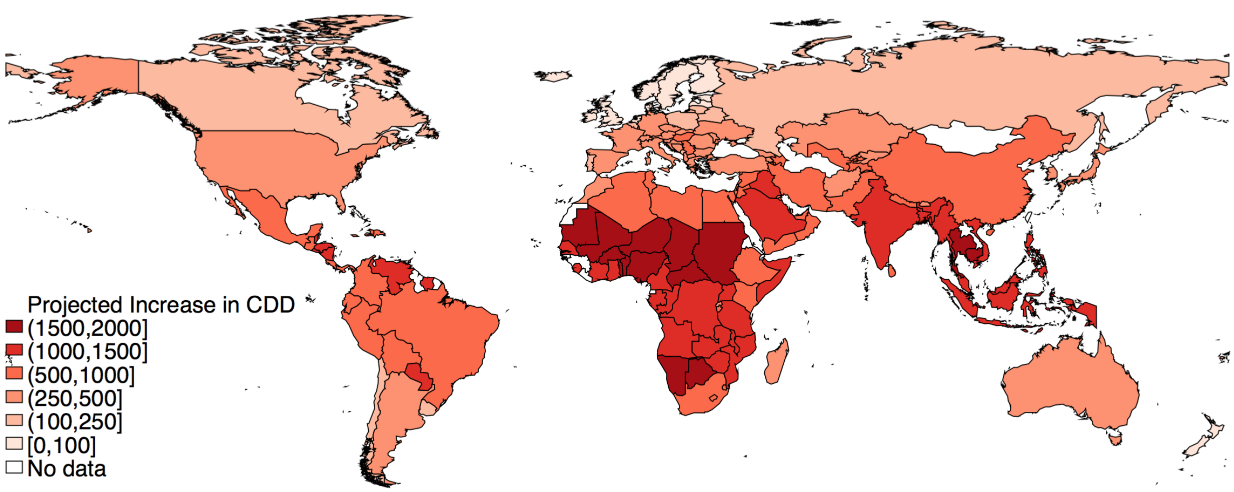

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme heat as measured by the change in cooling degree days above 30°C from 2015 to 2080-2099.

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

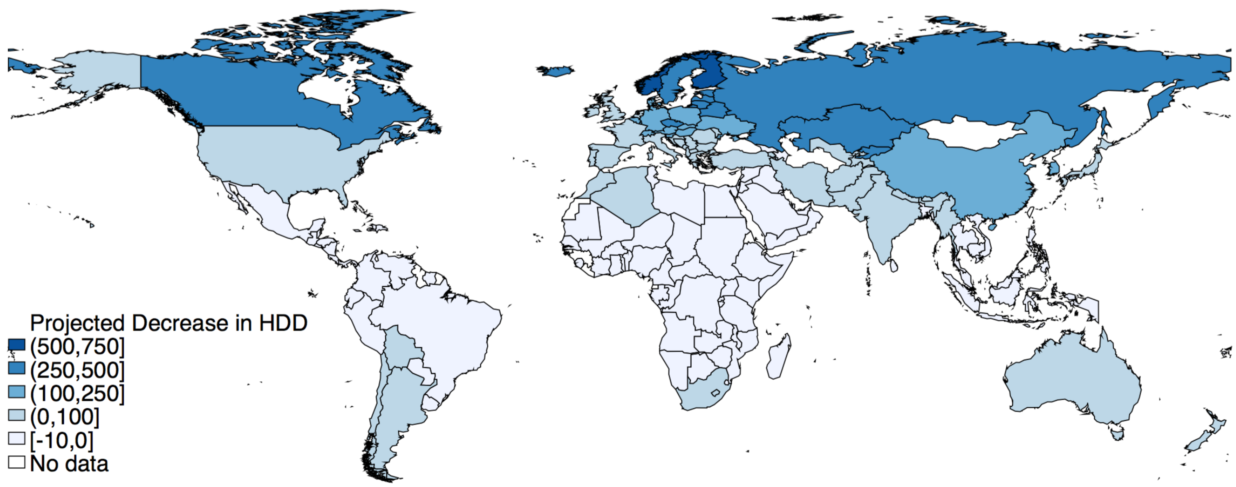

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme cold as measured by the change in heating degree days below 5°C from 2015 to 2080-2099.

Figure A-23: Projected Impact of Climate Change on Services Productivity

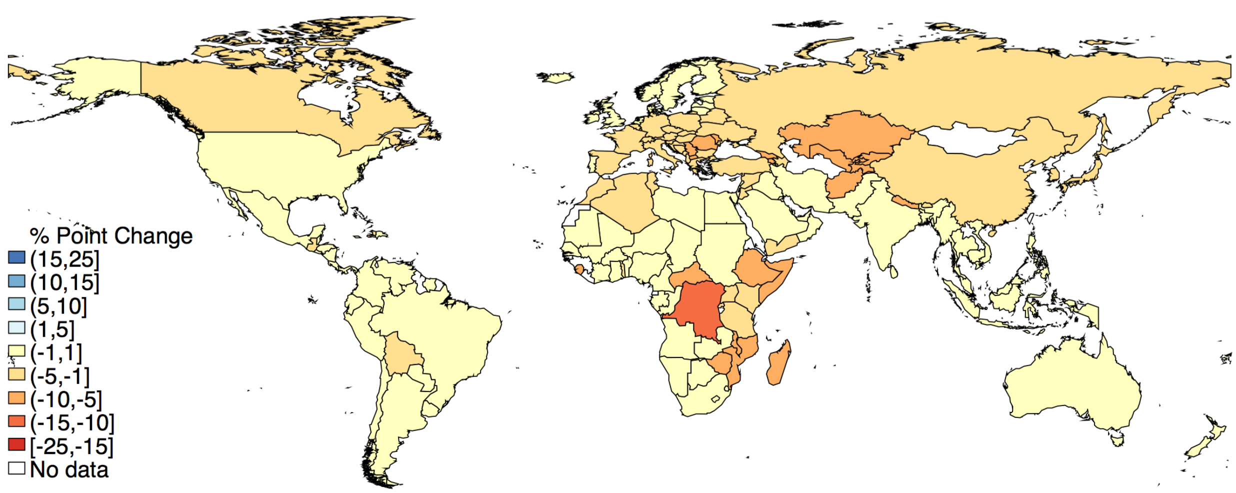

Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating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

Figure A-24: Projected Impact of Climate Change on Manufacturing Productivity Accounting for Economic Growth and Adaptation

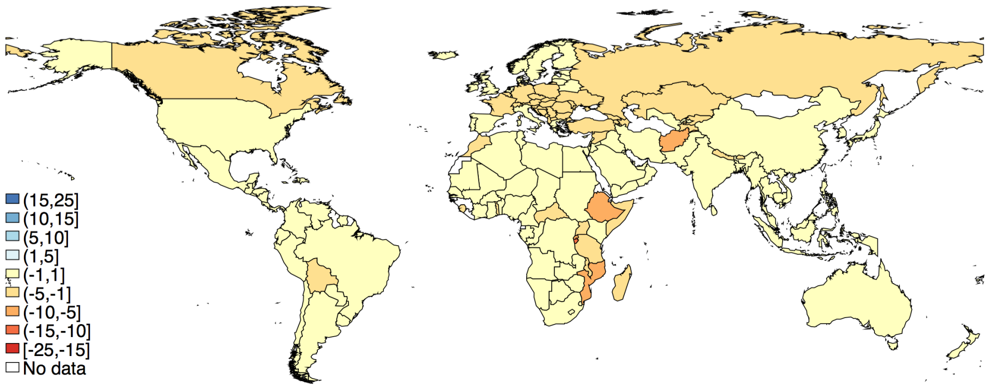

Notes: Map shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by 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 6a because the empirical results suggest that firms in richer countries have reduced exposure to extreme temperatures.

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

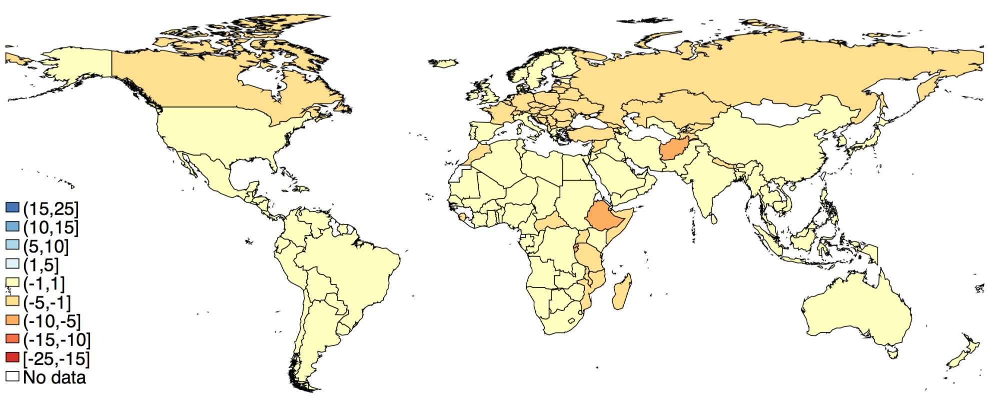

Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating 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-23 because the empirical estimates suggest that firms in richer countries have reduced exposure to extreme temperatures.

Appendix E: Model Calibration Details

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

I solve the model presented in Section 5 numerically as follows. First, I guess a vector of wages. Given the model parameters, this implies a set of sector-by-country price indices, P_{jk}, and bilateral sectoral trade shares, π_{jnk}, following the Eaton and Kortum (2002) structure within each sector in the model.[53] The sectoral price indices, together with the guess of wages, imply sectoral expenditure shares, ω_{jk}, consumption quantities, C_{jk}, and average cost indices, P_k, following the non-homothetic CES preference specification across sectors (Equations 11 and 12). Given these objects, I can check whether Equation 20 holds and expenditures equal incomes (trade balances) in each country. This is the final step necessary for the set of moments to constitute an equilibrium. If the condition fails, I repeat the procedure with a new set of wage guesses until Equation 20 holds (to within approximately 0.1%). See Allen, Arkolakis and Li (2020) for analysis of the uniqueness of equilibria in this class of models.

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.[54] 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 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.

[53] With the Fréchet productivity distributions, the equation for the sectoral price indices is as follows:

$$P_{jk} = \Gamma\left(1 + \frac{1-\eta}{\theta_j}\right)^{\frac{1}{1-\eta}} \left[\sum_{n \in N} Z_{jn}(\tau_{jkn} w_n)^{-\theta_j}\right]^{-1/\theta_j}$$

And the equation for the sectoral bilateral trade shares is given by:

$$\pi_{jnk} = \frac{Z_{jn}(\tau_{jnk} w_n)^{-\theta_j}}{\sum_{m=1}^{N} Z_{jm}(\tau_{jmk} w_m)^{-\theta_j}}$$

[54] 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.

Appendix E.2: Additional Model Fit Details

Table A-4: 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}			1.009
(Ag Domestic Production Share)			(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-26, A-27, and A-28.

Figure A-26: Sectoral GDP Shares - Data vs. Simulation

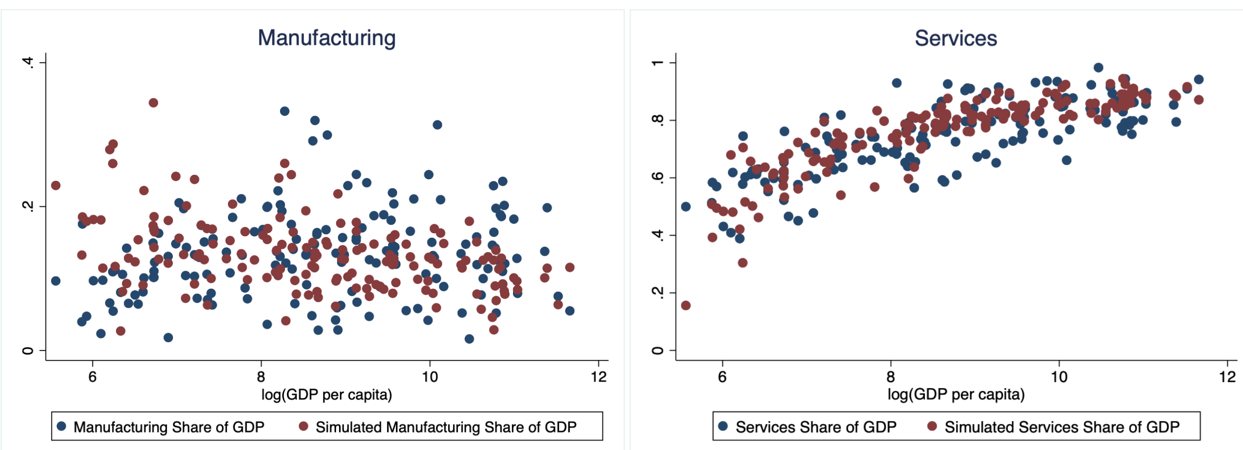

Notes: Left graph shows the fit of simulated manufacturing share of GDP in the model to data from the World Bank. Right graph shows the same comparison for services share of GDP.

Figure A-27: GDP Per Capita and Agriculture Share of GDP - Data vs. Simulation

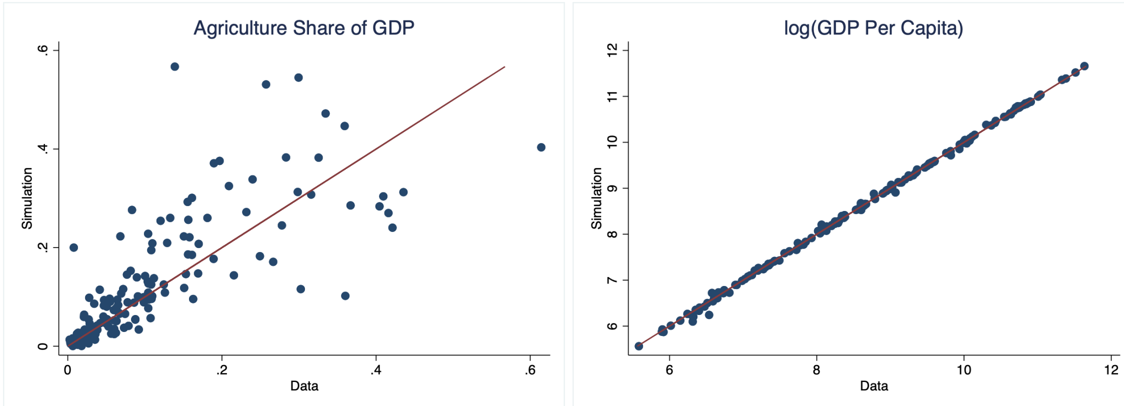

Notes: Left graph shows another view of the fit of simulated agriculture share of GDP in the model to data from the World Bank also shown in Figure 5a. Right graph shows the same comparison for GDP per capita. A perfect fit would have all data points be on the 45°line such that the simulated and actual 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-28: Domestic Production Share of Expenditures - Data vs. Simulation

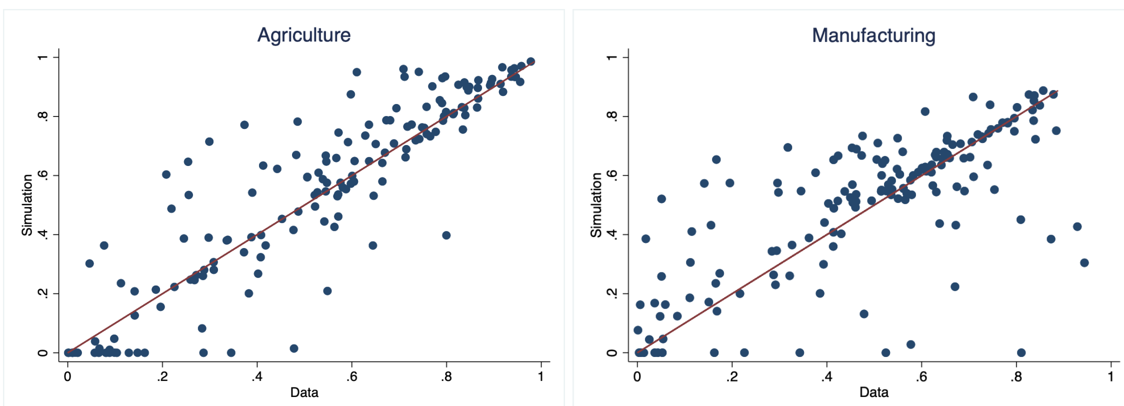

Notes: Graph shows the fit of simulated domestic production share of agricultural (left) and manufacturing (right) consumption in the model to data from Comtrade. As shown in Section 5.2, openness to food imports is a crucial parameter governing the response of labor reallocation to climate change. The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

Figure A-29: Domestic Production Share of Expenditures in Agriculture - Model Simulation

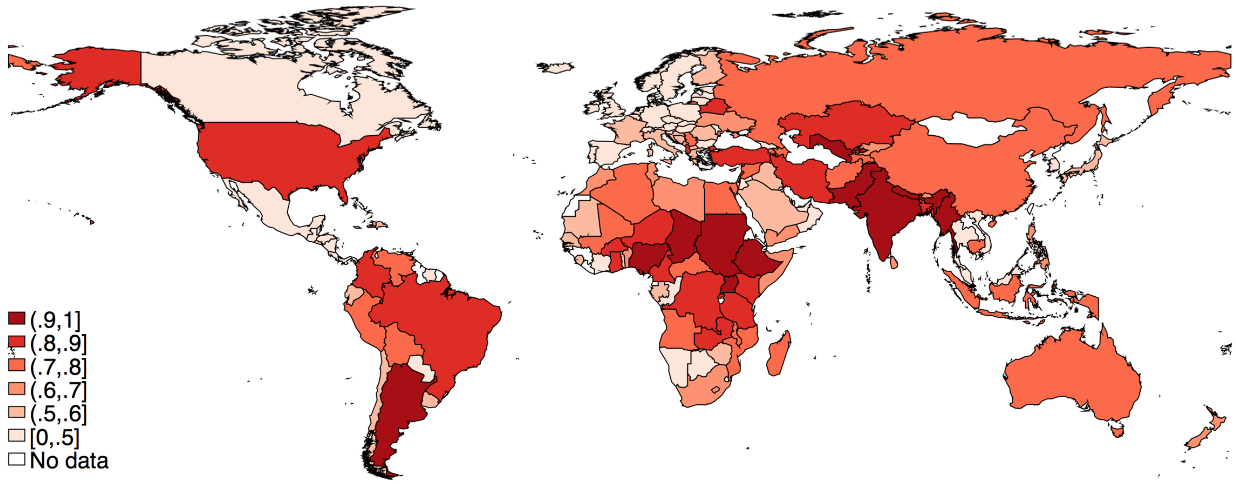

Notes: Figure shows that the share of expenditures on domestically produced goods in agriculture is very high in many developing countries with high barriers to trade. Table A-4 shows that these simulated values track closely to the data.

Appendix E.3: Agricultural Productivity Estimates

In this section, I briefly summarize the methods from the four sources of agricultural productivity estimates used in this paper: the quasi-experimental panel regression approach taken in Hultgren et al. (2021), the Ricardian approach taken in Cline (2007), the crop modeling approach taken in Iglesias and Rosenzweig (2010), and the separate set of crop model estimates used in Costinot, Donaldson and Smith (2016). In the main estimates in this paper, I take the unweighted country-level average impact projection across these four sources. Note that Table 7 shows that this paper's central results and qualitative conclusions are robust to the particular set of agricultural productivity estimates chosen.

The analysis in Hultgren et al. (2021) closely follows the empirical approach employed by Carleton et al. (2020) and used in Section 3 of this paper. The paper uses panel data from 12,658 sub-national administrative units across 55 countries for six major global staple crops (maize, soybeans, wheat, rice, cassava, and sorghum) in its empirical implementation. The analysis employs a cross-validation approach to select the moments from the temperature and precipitation distributions that matter most for each crop, and accounts for crop-level adaptation to future climate conditions using similar methods to this paper. I use results from the CSIRO-MK-3.6.0 climate model for consistency with the manufacturing projections made in Section 7.1, and take the average change in yield for each country weighting each crop by its present day share of national output.

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.

The analysis in Iglesias and Rosenzweig (2010) uses the IBSNAT-ICASA crop model to project global changes in wheat, rice, maize, and soybeans (I take a weighted average of the crop-level productivity change, using production shares as weights). This model contains a bottom-up representation of the physiological processes of crop growth, with functions that capture the effects of solar radiation, temperature, precipitation, soil characteristics, and management practices such as irrigation and the application of fertilizer. The model was parameterized using experimental evidence on crop growth from 124 sites across a wide range of local environments. The effects of climate change are simulated directly within the model for all regions throughout the world. See Parry et al. (1999) and Parry et al. (2004) for further details.

Finally, the estimates in Costinot, Donaldson and Smith (2016) use the UN Food and Agriculture Organization's Global Agro-Ecological Zones (GAEZ) projections for 10 crops: bananas, soybeans, cotton, sugarcane, maize, tomatoes, oil palm, wheat, rice, and white potatoes. The GAEZ agronomic model uses information on soil types, elevation, land gradient, rainfall, temperature, humidity, wind speed, and sun exposure to project the yield of each crop for each parcel of land under average historical baseline conditions and those projected in the future with climate change. The dataset contains separate projections for a range of climate scenarios and assumptions about complementary inputs, such as irrigation, fertilizer, and machinery. The main estimates in Costinot, Donaldson and Smith (2016) use the "high input" "rain-fed" set of projections from the Hadley CM3 A1F1 model scenario. National agricultural productivity estimates average across parcels for each crop and weight each crop by its share of national output.

Note that across all sets of agricultural productivity results, I do not account for any benefits of carbon fertilization. While recent work by Taylor and Schlenker (2021) has found that the fertilization effect from rising CO_2 concentrations has a substantial positive impact on crop yields, a range of scientific evidence suggests that it will have a substantial negative impact on crop nutrient content (see Beach et al. (2019), Zhu et al. (2018) Smith and Myers (2018), and Myers et al. (2014) for examples). In particular, field experiments and laboratory evidence show that rising CO_2

concentrations reduce the content of protein, zinc, iron, and vitamins B1, B2, B5, and B9 across a range of crops. Modeling suggests that the decline in nutrient density will more than offset the gains from the CO_2 effect on yields, such that the direct effect of rising CO_2 concentrations will cause a net decline in nutrient productivity. Given that more than twice as many people globally suffer from malnutrition (insufficient access to nutrients) as undernutrition (insufficient access to calories), my assessment is thus that the agricultural productivity estimates without carbon fertilization are most relevant to the subsistence food consumption mechanism central to the model presented in this paper.

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:[55]

$$U(C_{ak}, C_{mk}, C_{sk}) = \Bigg(\omega_a^{\frac{1}{\sigma}} (C_{ak} - \overline{C_{ak}})^{\frac{\sigma-1}{\sigma}} + \omega_m^{\frac{1}{\sigma}} (C_{mk} - \overline{C_{mk}})^{\frac{\sigma-1}{\sigma}} + \omega_s^{\frac{1}{\sigma}} (C_{sk} - \overline{C_{sk}})^{\frac{\sigma-1}{\sigma}} \Bigg)^{\frac{\sigma}{\sigma-1}} \tag{30}$$

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, particularly for middle-income countries, as shown in Figure A-30.

Table 7 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

[55]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-30: Agriculture Share of GDP - Data vs. Simulation
Stone-Geary Specification

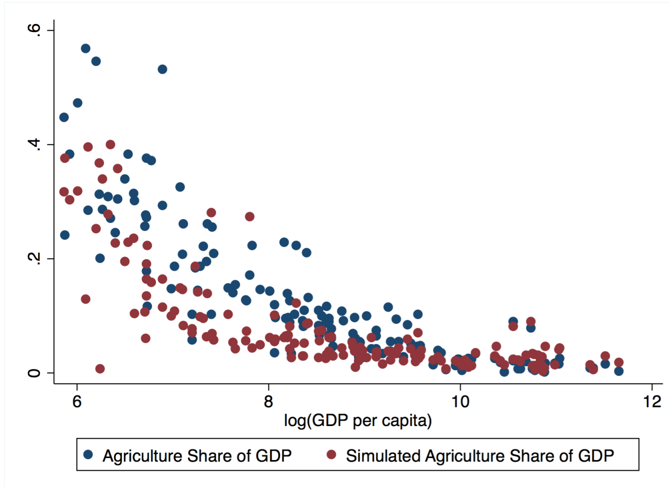

Notes: Graph shows the fit of simulated agriculture share of GDP to data from the World Bank with an alternative model specification using Stone-Geary preferences over sectoral consumption. The best fit with Stone-Geary preferences has an R^2 of only 0.43 and dramatically underpredicts the agriculture share in middle-income countries especially. In contrast, the chosen nonhomothetic CES preferences from Comin, Lashkari and Mestieri (2021) explain over 60% of the variation.

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: Multiple Factors of Production

The baseline model makes the assumption that production in each sector scales linearly with labor as the only input. In this section, I consider an extension of the model that allows for multiple factors of production, such as differentially educated labor, or land. I show that this version of the model maintains the key comparative statics regarding the competing effects of the food problem and trade, and also allows for additional qualitative insights about within-country distributional effects and the forces underlying comparative advantage. In this section, I detail a specification with heterogeneous workers in the production function, but the qualitative takeaways would be similar for a version with labor and land as inputs.

Considering a version of the model with heterogeneous workers is motivated by the fact that real world wages can differ substantially across sectors, whereas 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 agricultural workers have lower wages than non-agricultural workers in most parts of the world, and especially so in poor countries. 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-education and low-education 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 education-intensity (for simplicity I assume that manufacturing and services have the same education-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-education intensive than agriculture, as reflected by the high-education labor production elasticities $\alpha > \beta$. Solving the firm's problem gives the following optimal ratio of high-education and low-education workers employed in each sector as a function of the production elasticities and relative wages:

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

With $\alpha > \beta$, these conditions imply that manufacturing and services firms will employ a higher share of high-education 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 education type across the three sectors must add up to the country-level endowment of each education 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-education workers to high-education workers rises with the revenue share of agriculture.[56] Since agriculture is the less education-intensive sector, the 'food problem' actually works to partially insulate farmers from the welfare 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-education workers disproportionately employed in that sector. So while the relationship between greater openness to international trade, sectoral reallocation, and welfare 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-education workers if trade moves domestic production away from that sector and dampens the increase in its output price.

[56]The outline of the proof of this statement is as follows. In a perfectly competitive market with low-education and high-education 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 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-education share of total income is given by $\alpha - \beta$ and the change in the high-education share of total income is given by $\beta - \alpha$. With $\alpha > \beta$ the low-education share of total income rises. Since the total number of low-education and high-education workers is fixed, $\frac{w_L}{w_H}$ also rises.

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-education and low-education 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-education 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-education 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-education workers and endogenously dampen the movement of comparative advantage away from agriculture. Thus, relative to the baseline model, this extension presents an additional barrier that diminishes the potential for shifting trade flows to contribute to climate change adaptation.

Overall, extending the model to represent multiple factors of production leaves the fundamental predictions about climate change and sectoral reallocation unchanged, but allows for additional components of comparative advantage and sheds additional light on the distributional consequences of climate change.