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

Ishan Nath *

Federal Reserve Bank of San Francisco

July 1, 2024

Abstract

This paper evaluates the potential for global reallocation between agricultural and non-agricultural production to contribute to climate change adaptation. Empirical estimates using a global sample of firms suggest that rising temperatures reduce productivity less in non-agriculture than agriculture, implying large potential gains if hot countries could increase food imports and shift labor toward manufacturing. However, model counterfactuals show that subsistence consumption needs and high trade barriers combine to create a "food problem" in which climate change instead intensifies agricultural specialization in especially vulnerable regions. Simulations suggest that reducing trade barriers can significantly reduce climate damages, especially in poor countries.

*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 numerous seminar and conference participants. 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. Any views expressed in this paper are those of the author and do not necessarily reflect those of the Federal Reserve System or its staff. 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 projection from a range of estimates in this literature, which suggests that global warming will reduce agricultural productivity by up to 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 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 reallocation 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 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 having a much lower ratio of value-added per worker in agriculture relative to non-agriculture compared with rich countries (Tombe, 2015). Projecting the future effects of climate change on sectoral reallocation requires accounting for the forces that drive poor countries to presently specialize in agriculture despite an apparent lack of existing comparative advantage.

This paper addresses these challenges by integrating local temperature treatment effect estimates with a quantitative macroeconomic model to assess the potential for sectoral reallocation to contribute to climate change adaptation. First, to project changes in comparative advantage between agriculture and other sectors, I provide the first global micro estimates of the impact of rising temperatures on labor productivity in manufacturing and services using data from a broad range of countries that cover over half the world's population and represent nearly the full range of temperature and income levels. Using methods developed by Carleton et al. (2022), I estimate plausibly causal treatment effects of extreme temper-

[1] The agricultural productivity impacts displayed in Figure 1a take the average of estimates produced by Cline (2007), Iglesias and Rosenzweig (2010), Costinot, Donaldson and Smith (2016), and Hultgren et al. (2021). Section 2.2 and Appendix E.4 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.

atures on output-per-worker, and make projections that account for firm-level adaptation by allowing these effects to vary with income and expectations of temperature.

Second, I embed the empirically estimated productivity effects in a global quantitative trade model that explains existing patterns of agricultural specialization as the result of two key forces: subsistence consumer preferences for food, and high barriers to trade. The model's consumer preferences incorporate two key features - non-homotheticity and low substitutability - that explain the high agricultural expenditure shares in low-income countries with high relative prices for food. In principle, imports could meet these domestic 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 expenditures shares necessary for people with low incomes to meet their subsistence need to eat, a phenomenon that Gollin, Parente and Rogerson (2007) labeled as the "food problem."

Thus, the model shows that the net effect of the two forces governing the sectoral reallocation response to global warming depends on openness to trade. If trade is relatively free, countries can dampen the effect of falling agricultural productivity by shifting production to other sectors; exporting more manufactured goods and importing more food. If trade is relatively closed, global warming can exacerbate the "food problem" and increase specialization in the most vulnerable sector in order to meet domestic demand for food. I calibrate the model to match global data on income levels, trade flows, and sectoral GDP shares, and simulate climate change counterfactuals with existing trade costs and in a range of alternative scenarios in which trade costs fall in the future.

The paper's results lay out the implications of the climate change "food problem" in several steps. First, I find that extreme temperatures have much smaller effects on non-agricultural productivity than the effects others have measured in agriculture, which implies large potential gains from sectoral reallocation. Exposure to extreme temperatures can have substantial effects on output per worker in some settings, with the least adapted firms experiencing annual losses of up to 0.4% from each extremely hot (40°C) or extremely cold (-5°C) day. However, I also find strong evidence consistent with adaptation by firms in more productive locations and those that experience a given extreme more frequently. In high-income countries, the effects of extreme days are virtually negligible, with some evidence of mild effects of hot days in cold places and cold days in hot places.[2] Combining the estimated temperature sensitivities with climate model projections suggests that future temperature changes will reduce global manufacturing productivity by about 1.7% on average, about an order of magnitude smaller than projected effects in agriculture. The hardest hit hot, poor countries project to suffer losses of

[2] 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. I also 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 in Appendix D.

up to 5-14%, though such magnitudes generally remain several times smaller than projected effects in agriculture. This suggests large potential gains if such locations were able to reallocate production away from farming.

Despite the potential gains from such reallocation, model counterfactuals suggest that global warming instead exacerbates the "food problem" and intensifies specialization in agriculture in hotter low-income countries where its productivity suffers most. While warming shifts net exports of agriculture toward colder regions and away from hotter ones, this response is modest in magnitude in the most vulnerable regions due to the high trade barriers the calibration infers from low existing trade shares. With this weak trade response, global warming raises agriculture's share of the labor force by 2.8 percentage points in the poorest quartile of countries in order for domestic production to meet demand for food in the face of the productivity shocks. Overall, accounting for trade reduces climate damages to welfare 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 global warming are also least open to trade.

Counterfactual model simulations demonstrate the paper's key policy implication, which is that greater openness to trade could dramatically reduce climate damages in low-income countries by alleviating the "food problem." Benchmark scenarios in which all countries move to the lowest levels of trade costs observed in the calibration reduce climate productivity damages by as much as half in low-income countries and by about a fifth for the world overall. In these hypotheticals with lower barriers to trade, specialization aligns more closely with shifting comparative advantage and the "food problem" no longer binds. By allowing agricultural production to move away from the most vulnerable regions, these scenarios limit the climate-induced rise in food price indices globally, and especially so in low-income countries. While a comprehensive breakdown of the underlying causes of low trade shares in poor countries remains unresolved in the literature, I show that reducing the proportion of the model's trade barriers explained by tariffs, trade agreements, and regulatory frictions can achieve most of the theoretical maximum welfare gains globally, and about a third in poor countries. The findings suggest that further research on how trade policy, regulatory frictions, and infrastructure investments shape the ability of low-income countries to engage in trade could be critical for identifying policies that promote climate change adaptation.

Before proceeding, it is worth stating clearly that this paper does not aim to provide a comprehensive assessment of the costs of climate change. The analysis omits a wide range of important factors including, but not limited to, international migration (see e.g. Missirian and Schlenker (2017) and Cruz Alvarez and Rossi-Hansberg (2021)), health effects (e.g. Heutel, Miller and Molitor (2017)), hurricanes (e.g. Bakkensen and Barrage (2018)), sea-level rise (e.g. Desmet, Kopp, Kulp, Nagy, Oppenheimer, Rossi-Hansberg and Strauss (2018)), and uncertainty (e.g. Cai and Lontzek (2019); Lemoine (2021)), all of which will likely play an important role in the welfare consequences of climate change.

This paper focuses specifically on the critical role of barriers to trade in preventing sectoral reallocation from alleviating the productivity effects of rising temperatures in low-income countries.

The paper builds on a small number of closely related papers. Costinot, Donaldson and Smith (2016) project substantial climate change adaptation gains from reallocation across crop types, but do not consider damages or trade in the non-agricultural sector or model a "food problem" in which trade barriers cause low-income countries to specialize in agriculture despite low relative productivity. They find that projected trade adjustments in the calibrated model have little impact on the welfare consequences of climate change, but do not investigate the central question in this paper about whether lowering observed trade barriers could contribute meaningfully to adaptation. A related paper by Gouel and Laborde (2021) extends the model from Costinot, Donaldson and Smith (2016) in several directions, including modeling a choice between crops and livestock, and finds that trade adjustment contributes substantially to climate change adaptation even with current levels of trade barriers. The other most similar work consists of papers by Desmet and Rossi-Hansberg (2015) and Conte, Desmet, Nagy and Rossi-Hansberg (2021), which examine migration and trade in dynamic models. These papers project meaningful global adaptation gains from agricultural specialization moving toward colder regions, but do not focus specifically on the combination of trade barriers and subsistence preferences that prevent low-income countries from realizing these benefits. This paper is the first to show that global warming is likely to intensify agricultural specialization in the hottest parts of the world with current levels of tradability, and that easing trade restrictions could reverse this "food problem."

More broadly, this paper's empirical work builds on country level estimates from 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 Uy, Yi and Zhang (2013), Tombe (2015), Teignier (2018), and a non-homothetic CES consumer preference specification from Comin, Lashkari and Mestieri (2021). Another related paper by Porteous (2019) focuses specifically on the welfare costs of high trade barriers in African agriculture. The model counterfactuals relate to empirical work by Colmer (2021) and Liu, Shamdasani and Taraz (2023) 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).

In the literature on climate change economics, this paper contributes to a nascent body of work that advances the frontier of methods by embedding credible empirical estimates into a general equilibrium model. Early work on climate impacts followed two primary tracks: macroeconomic models such as 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 adopted 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 follows the path of recent work such as Balboni (2019), Barrage (2020), Conte (2021), Cruz Alvarez and Rossi-Hansberg (2021), Fried (2021), Rudik, Lyn, Tan and Ortiz-Bobea (2021), and Bilal and Rossi-Hansberg (2023) that aims to unify these approaches by mapping empirical estimates from micro-data directly into the parameters of quantitative models, allowing the researcher to evaluate counterfactuals in a framework that captures equilibrium behavior and welfare.

The paper is structured as follows. Section 2 presents empirical evidence on extreme temperatures and non-agricultural productivity, along with stylized facts about climate change and agricultural specialization. Section 3 lays out the model, Section 4 describes the calibration, and Section 5 contains counterfactuals about climate change, sectoral reallocation, and trade policy. Section 6 concludes.

2 Empirical Evidence on Temperature and Productivity

This section examines the likely effects of global warming on sectoral productivity and comparative advantage. As documented in Figure 1a, a rich existing literature projects that warming will cause large damages to agricultural productivity, concentrated in tropical regions. Evaluating the potential for hotter places to adapt to these changes by reallocating economic activity to other sectors requires corresponding estimates of impacts in non-agricultural sectors, yet comparatively little research addresses this question. Existing studies using macro data find substantial effects of temperature on non-agricultural activity, but lack the statistical precision to compare magnitudes with effects estimated in agriculture (e.g. Jones and Olken, 2010; Dell, Jones and Olken, 2012; Burke, Hsiang and Miguel, 2015). A small number of studies using micro data also find substantial effects of temperature on non-agricultural productivity (e.g. Zhang, Deschenes, Meng and Zhang, 2018; Somanathan, Somanathan, Sudarshan and Tewari, 2021), but are limited in scope to a given location in a given time period, precluding general conclusions about the effects globally and in the future.

This paper uses a global sample of micro-data to project the effects of rising temperatures on non-agricultural productivity using methods from Carleton et al. (2022) that account for how effects vary across different contexts and allow for firms to adapt to their surroundings as the climate warms. The analysis finds that while temperature has substantial implications for non-agricultural productivity, the effects of global warming are likely to be about an order of magnitude larger in agriculture than non-agriculture, suggesting large potential gains for hot countries if they are able to shift away from existing patterns of specialization in agriculture.

2.1 Extreme Temperatures and Non-Agricultural Productivity

2.1.1 Empirical Approach

To project the effects of global warming on non-agricultural productivity, I start by assembling a broad global sample of micro-data in the manufacturing and services sectors across 17 countries. The data contains nationally representative firm-level panels obtained from government surveys in India, Colombia, Indonesia, China, and the United States, and from Bureau van Dijk's (BVD) Amadeus database in twelve European countries.[3] See Appendix Table A-1 for a full listing of the countries, years, and data sources, and Appendix A.1 for additional information on data construction. Note also that the main pooled specification in this paper excludes the U.S. because its data can be accessed only at a secure government facility as well as China for data quality reasons explained in Appendix C.

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. 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 covers 59% of the world's manufacturing output and 51% of the global population. The dataset also spans virtually the full range of climate and income levels in the global cross-section. According to the Penn World Tables, purchasing power parity 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.

For temperature exposure, I use data from Version 3 of the Global Meteorological Forcing Dataset (GMFD) produced at Princeton University. GMFD is a reanalysis dataset that reconstructs historical temperature at a 0.25°by 0.25°resolution using a combination of observational data and local climate models. Following Graff Zivin and Neidell (2014), I use daily maximum temperature

[3] Bureau van Dijk is a private company owned by Moody's Analytics that collects and distributes firm-level financial information collected from a combination of administrative sources and private surveys. Bloom, Draca and Van Reenen (2016) report that the data in many European countries contains nearly the full population of public and private firms as recorded in national business registers, and I restrict attention to those countries with mandatory filing requirements according to BVD documentation. See Gopinath, Kalemli-Özcan, Karabarbounis and Villegas-Sanchez (2017) and Alfaro and Chen (2018) for other examples of papers that use BVD data.

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 and 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, such as calculating degree-days or polynomials, at the pixel level, and average across pixels in a county weighting by population.

To estimate the effects of temperature on productivity, I start by noting that workers experience daily realizations of weather as emphasized by Deryugina and Hsiang (2014). 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:

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

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, similar to Schlenker and Roberts (2009), where output is allowed to vary linearly with daily maximum temperature above 30°C (Cooling Degree Days - CDD) and below 5°C (Heating Degree Days - HDD):

$$f(T) = \begin{cases} \beta_1(5 - T_{max}) & \text{if } T_{max} < 5 \\ 0 & \text{if } 0 \leq T_{max} \leq 30 \\ \beta_2(T_{max} - 30) & \text{if } T_{max} > 30 \end{cases} \tag{2}$$

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

The dependent variable in the analysis is revenue per worker. While the notion of productivity depicted by the model in Section 3 corresponds most closely to physical output per worker, this information is not available in my sample nor more generally in systematic firm or plant-level data with broad coverage, as documented by Syverson (2004). The estimates in this paper can be interpreted as the effect of temperature on physical productivity under the assumption that firms are price-takers in the output market and that the local shocks used for identification do not affect national or global product market prices. While these assumptions are difficult to test directly, the most relevant evidence comes from Somanathan, Somanathan, Sudarshan and Tewari (2021), who manually survey physical productivity in garment weaving, cloth sewing, and steel manufacturing plants in India. Their analysis finds that the effects of temperature on physical productivity in these industries closely mirrors the magnitude of the effects on revenue productivity measured in the national sample of Indian plants, which is also one of the datasets used in this paper. Thus, the best existing evidence suggests that the available measure of revenue productivity is a reasonable proxy for the ideal measure of productivity measured in physical units of output.

To isolate the causal impact of temperature on revenue per worker, I follow the standard approach of exploiting interannual variation in weather using the following panel regression:

$$\ln\left(\frac{Y_{it}}{L_{it}}\right) = \boldsymbol{\beta} F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \tag{3}$$

Firm fixed effects, δ_i, control for time-invariant features of productivity, and region-(country or state) by-year fixed effects, κ_{rt}, control for aggregate shocks. I cluster standard errors at the firm and county-by-year level to account for both serial and spatial correlation.

Equation 3 captures the average effect of temperature on productivity, but making global projections also requires measuring how this relationship varies across contexts and over time. Profit-maximizing firms have greater incentive to invest in adaptation when they are more productive or more exposed to extreme conditions, so we might expect the effects of extreme heat to be smaller in more productive and hotter locations, both in the global cross-section and over time as economic growth and global warming raise productivity and temperature. To capture this potential for firm-level adaptation, I follow Carleton et al. (2022) in allowing for heterogeneity by interacting the vector of temperature coefficients with measures of income and temperature, as follows:

$$\begin{aligned} ln\left(\frac{Y_{it}}{L_{it}}\right) &= \boldsymbol{\beta} F(T)_{it} + \boldsymbol{\gamma_1} ln(GDPpc)_{rt} \times F(T)_{it} \\ &+ \boldsymbol{\gamma_2} TMEAN_i \times F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \end{aligned} \tag{4}$$

The interaction variables in Equation 4 are country level annual GDP per capita from the Penn World Tables and long-run average daily maximum temperature from GMFD in the county containing firm i.[4] The coefficients on γ_1 and γ_2 measure the degree to which the effects of temperature on productivity differ in more productive and hotter locations.

[4] 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.

2.1.2 Effects of Temperature on Non-Agricultural Productivity

Table 1 contains the main results from estimating Equations 3 and 4. Column 1 displays the treatment effect of extreme temperatures for the average unit of output across 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 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 0.003% relative to a day in the moderate range of 5°C to 30°C. Column 2 shows the same regression but without weights, which can be interpreted as the effect on the average firm in the sample. This result is similar to the result from Column 1, perhaps because more of the data comes from rich countries where the effects are more muted.

The results from the interacted regression in Column 3 of Table 1 show that the effects of temperature depend heavily on context. Note that this specification is unweighted, following Carleton et al. (2022) and the recommendation of Solon, Haider and Wooldridge (2015) for settings where it is possible to model heterogeneous responses directly rather than averaging over them. The coefficients on CDD and HDD in Column 3 are large, negative, and precisely estimated, though the magnitudes cannot be interpreted directly as they correspond to a hypothetical location with zero income and average temperature. The coefficients on both interaction terms for log GDP per capita are large and positive, indicating that richer countries are insulated from the effects of both extreme heat and cold. Consistent with intuition about expectations driving adaptation, 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 conditions they experience more frequently.

To interpret the magnitude of the effects more intuitively, Figure 2 shows the predicted effects of temperature from Column 3 of Table 1 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 substantial effects on productivity in poor countries (bottom row), though the magnitudes depend critically on frequency of exposure. In a hypothetical low-income country with moderate long-run temperatures (bottom middle cell), a 40°C day reduces annual output per worker by about 0.4%, approximately equal to one full day of production from a typical working year. In hotter countries where firms experience these events more frequently, however, the effects are less than half as large.

Columns 4 and 5 of Table 1 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. Appendix Figures A-1 and A-2 show that these effects also primarily manifest only in poor countries. Finally, Column 6 shows the effects of temperature on a pooled sample of manufacturing and services firms. The effects are very similar to the sample of 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 in the pooled sample. 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.

I conduct a variety of robustness checks for the effects of extreme temperatures on non-agricultural productivity. Appendix A provides details of these specifications, which include results for alternative functional forms, lagged effects, additional control variables, and further interactions of covariates. In addition, I also test the external validity of the results using separate estimates with data from the United States. Predictions using the global interacted regression, which excludes the U.S. data for logistical reasons, suggest that temperature has a negligible effect on manufacturing revenue per worker in rich temperate countries (see the top middle cell of Figure 2, which corresponds closely to the average U.S. context). Consistent with this, Figure 3a shows a precisely estimated null effect of temperature on revenue per worker in U.S. manufacturing. The U.S. data also includes information on other inputs lacking from the global sample, allowing me to observe some of the adaptation costs firms incur. 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. 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, as shown in Appendix Figure A-14.

While the results in this section are robust to a wide range of specification choices, it is worth noting some limitations in the 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 regression approach 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.[5] Third, the dynamics of adaptation may evolve over time more broadly in ways that go beyond the data. For instance, if technological innovation reduces the costs of insulating production from extreme temperatures, the relationship between income, long-run average temperature, and temperature sensitivity could change in the future. Finally, aggregating from firm-level effects to industry level effects in Section 2.2 requires the assumption of homogeneous effects across firms within a sec-

[5]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.

tor for each country. If temperature affects different manufacturing firms in India differently (a level of granularity beyond which I can 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.

2.2 Projected Climate Impacts on Agriculture and Non-Agriculture

2.2.1 Projecting Impacts on Non-Agricultural Productivity

To project the sectoral productivity impact of global warming, I start by using the estimates from Equation 4 to predict the sensitivity of production to extreme temperatures 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. The sensitivities vary with each country's level of GDP per capita and long-run average temperature, based on the coefficients estimated in Column 3 of Table 1. Consistent with intuition about firm-level adaptation, extreme heat and cold have the largest effects in poorer locations and those which experience given temperatures less frequently.[6]

To account further for firms adapting to changing future conditions as the climate warms, Appendix Figure A-18 reevaluates the heat sensitivity at projected end-of-century temperatures in the climate scenario considered in Section 5. 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%). These effects are nearly an order of magnitude lower than similar estimates of extremely hot days on agricultural productivity, as the next section discusses further.

The adaptation benefits of adjusting to extreme heat come at a cost to firms, as shown in the effects of temperature on energy expenditures in U.S. manufacturing in Figure 3b. If it were costless to protect production from extreme heat, no firm would show effects of temperature on labor productivity. Instead, the results show that firms which experience given extremes infrequently are less adapted, implying that the costs they would incur to achieve a marginal reduction in temperature sensitivity exceed the benefits. In Appendix D, I use methods from Carleton et al. (2022) that leverage this intuition to infer a revealed preference measure of adaptation costs that I account for in the counterfactuals.

The model simulations in Section 5 also require projecting temperature sensitivity in services, which play an important role as the only nontradable sector in the model.[7] I make projections for services using the pooled sample of manufacturing and services firms due to the lack of services data

[6]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.

[7]While the primary reallocation of interest in this paper is between agriculture and non-agriculture, treating all non-agricultural production as potentially tradable would risk overstating the potential for policy to affect trade flows and contribute to climate change adaptation if some types of production, such as haircuts or health care services, are non-tradable by construction. Thus, I include a services sector with no trade in the model.

coverage in poor countries. 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 3 and 6 of Table 1.[8] 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-21 and A-22 show predicted current global sensitivity to hot and cold days in services using results from the pooled regression. I follow the same procedure to account for future adaptation benefits and costs as in manufacturing.

To project the impact of global warming, I combine the country-specific sectoral temperature sensitivities with projections of future temperature distributions at end-of-century. I obtain future temperature predictions for the climate scenario from Representative Concentration Pathway 8.5 from the CSIRO-MK-3.6.0 model produced by Jeffrey et al. (2013).[9] Figure 5a and Appendix Figure A-25 show the projected changes in manufacturing and services productivity, respectively. The results suggest that climate change will have meaningful, though moderate, effects on 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.

2.2.2 Comparison to Agricultural Productivity Impacts

I draw from a rich existing literature to show that the projected agricultural productivity effects of global warming are much larger than the non-agricultural impact estimates shown in the preceding section. Figure 1a shows 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 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. The sources

[8] A formal test shows that coefficients for manufacturing and services firms in the pooled regression have statistically indistinguishable responses to extreme heat.

[9] The estimates from the interacted model in Section 2.1 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 the last 20 years of the century relative to the first 20 years, which are shown in Appendix Figures A-23 and A-24. I multiply the country level coefficients by the projected changes in hot and cold temperatures to get the impacts shown here.

contain estimates for between four and ten crops each (and in the case of Cline (2007) accounts for revenues from all crop production and livestock), and I use projections throughout for the same high-emissions scenario used in the non-agricultural estimates. Appendix E.4 contains full details on the methods used in these papers, which each have their own advantages and drawbacks.

I use the average across these sources in the quantitative exercises, but the projections are remarkably similar to each other 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). Recall that this is in comparison to this paper's projection of a 1.7% global decline in manufacturing productivity. Thus, the projected effects of warming are about an order of magnitude larger in agriculture than in manufacturing. Figure 5b shows the relative sectoral effects by country. Every country in Africa, South Asia, and Latin America has larger estimated productivity losses in agriculture than manufacturing, with the magnitude of the difference measuring 20 to 50 percentage points in many places.

The agricultural impact estimates in the literature account for within-sector adaptation in similar ways to the non-agricultural estimates in this paper, and thus can be viewed as comparable. In particular, the empirical papers make projections that account for heterogeneous responses between locations with differing levels of development and exposure to extreme rainfall and temperatures, thus accounting for unobserved adaptation investments that may vary across contexts. Similarly, the sources that use crop models account for a range of adaptation mechanisms within agriculture that include the response of inputs such as irrigation, fertilizer, and machinery, reallocating planted acreage within country, and reoptimizing planting dates as the climate changes. Note that a drawback to many papers in this literature is that they do not account for adaptation by switching between crop types in response to changing climate conditions, just as the non-agricultural productivity projections in this paper do not incorporate potential reallocation between categories of manufacturing. The paper with projection methods most similar to this paper is Hultgren et al. (2021), which finds that the global average effect of a 40°C day reduces annual output by about 7% for soy, 5% for maize, and 3% for rice, in sharp contrast to the 0.1% global average effect of a 40°C day on annual manufacturing productivity estimated here. Thus, an extensive range of evidence using a variety of methods, as well as the most directly comparable empirical estimates using micro-data, all support the conclusion that the agricultural sector is likely to suffer far greater damages from warming than the manufacturing and services sectors.

2.3 Stylized Facts about Temperature, Productivity, and Sectoral Shares

The analysis in Sections 2.1 and 2.2 shows that rising global temperatures are likely to harm agricultural productivity far more than non-agricultural productivity, suggesting that the most exposed hot regions in the world might be able to adapt effectively if they were able to reallocate production away from agriculture. This section uses a combination of global data and existing literature to document a set of additional stylized facts about global patterns of agricultural productivity, specialization in agriculture, and the effects of temperature. Together, this information helps frame the subsequent model analysis of the potential for adaptation through sectoral reallocation.

Fact 1: Poor countries specialize in agriculture despite relatively low productivity.
Figure 1b (left panel) shows that agriculture's share of employment and output is much larger in the poorest countries in the world than in the richest countries. The average economy in the 10th percentile of GDP per capita in 2011 had agricultural employment of about 67%, compared with 3% in the 90th percentile richest countries. This pattern of specialization might cause one to conclude that poor countries hold comparative advantage in agriculture. However, Lagakos and Waugh (2013) calculate that the ratio of price-adjusted aggregate output per worker between 90th and 10th percentile countries is about 4 to 1 in manufacturing, and about 45 to 1 in agriculture. Thus, poor countries actually have much *lower* relative productivity in agriculture in comparison to other sectors, compared with rich countries. This suggests that existing global patterns of sectoral specialization could be driven by forces other than comparative advantage.

For the remaining stylized facts that follow, I produce country level panel regression estimates of the macroeconomic effects of exposure to extreme heat, which serve as motivating evidence for the global warming counterfactuals shown later in Section 5. Because the projections in Section 2.2 suggest that rising temperatures will have a disproportionate impact on agriculture, I use an agriculture-focused measure of exposure: "growing degree days" (GDD) between 0°C and 29°C and "killing degree days" (KDD) above 29°C, aggregated to the country level weighting each pixel by its share of cropland.[10] As shown by Schlenker and Roberts (2009), GDD and KDD represent positive and negative shocks to agricultural productivity, respectively. Using this measure of temperature, I estimate the following regression specification for several dependent variables in a historical panel with unbalanced coverage for 164 countries from 1960-2012.

$$Y_{it} = \beta_1 GDD_{it} + \beta_2 KDD_{it} + \delta_i + \kappa_t + \epsilon_{it} \tag{5}$$

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 the effect for the average farm worker in the world. Table 2 shows the effects of GDD and KDD on several dependent variables, and Appendix Table A-2 shows a version of the regression that interacts GDD and KDD with long-run average temperature and per-capita income. I summarize the key takeaways below.

[10] 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.

Fact 2: Extreme heat reduces GDP substantially in agricultural economies.
Column 1 of Table 2 shows that killing degree days have large effects on GDP. The table implies that 100 KDD would reduce GDP by about 12% when weighting countries by their share of global agricultural workers. Appendix Table A-2 shows that the effects are much larger in poorer and hotter regions. To put the magnitude of the effects in context, global warming projections from the emissions scenario used in Section 5 imply an increase in exposure up to several times larger than 100 KDD in hot, agrarian economies by the end of the century. This evidence, taken together with similar estimates by Dell, Jones and Olken (2012) and Burke, Hsiang and Miguel (2015), implies that future warming is likely to have a quantitatively meaningful effect on income levels in poor countries, underscoring the importance of the non-homotheticity mechanism considered in the model simulations to follow.

Fact 3: Extreme heat raises the food share of imports.
Column 2 of Table 2 shows that countries raise the food share of their imports in response to harmful temperature shocks (KDD) and reduce the food share of their imports in response to favorable temperature conditions (GDD). This suggests that extreme temperatures can have impacts on the sectoral composition of trade flows. Appendix Table A-2 shows that the food share of imports rises more in poorer and hotter regions in response to KDDs. It is worth noting, however, that the coefficients are imprecisely estimated and the magnitude of the effects is generally small. A typical low-income hot country would be expected to raise their food share of imports by only about 1% if faced with 100 additional KDDs. Overall, though, these results suggest that extreme temperatures have some modest effects on trade.

Fact 4: Extreme heat raises agriculture's share of production and the labor force.
Columns 3 and 4 of Table 2 show that exposure to killing degree days on cropland raises agriculture's share of GDP and the labor force, despite the harmful impact on agricultural productivity. Appendix Table A-2 shows that these effects are strongest in poorer and hotter countries. Though the coefficients are imprecisely estimated and reflect short-run rather than long-run effects, the magnitudes are consistent with projected future increases in extreme heat exposure raising agriculture's share of GDP by several percentage points in the hardest hit countries.

Together, the facts suggest that extreme heat causes substantial economic damage, elicits a mild trade response that raises the food import share, and moves production and labor into agricultural sectors that are relatively unproductive in poor countries on average. I interpret this as suggestive evidence of the impact of exposure to extreme climates on sectoral specialization, trade, and welfare, though the magnitudes cannot be compared directly to the model simulations that follow since the regressions measure short-run effects and do not account for international spillovers. In the next section, I present a model of sectoral specialization and trade that builds on existing work to rationalize the observed global pattern of agricultural specialization and productivity.

3 Model

This section lays out a static general equilibrium model of global production, consumption, and trade in agriculture, manufacturing, and services. I use the model to show the conditions under which reductions in agricultural productivity cause labor to reallocate away from the agricultural sector and to formalize the critical role of trade openness in governing this potential mechanism of adaptation. In the sections that follow, I use the model to quantify how temperature-driven changes in sectoral productivity affect sectoral specialization, trade flows, prices, and welfare under counterfactual policy scenarios.

3.1 Model Ingredients

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 for U_k:

$$\Omega_a^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_a}{\sigma}} C_{ak}^{\frac{\sigma-1}{\sigma}} + \Omega_m^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_m}{\sigma}} C_{mk}^{\frac{\sigma-1}{\sigma}} + \Omega_s^{\frac{1}{\sigma}} U_k^{\frac{\epsilon_s}{\sigma}} C_{sk}^{\frac{\sigma-1}{\sigma}} = 1 \tag{6}$$

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

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:

$$P_{ak}C_{ak} + P_{mk}C_{mk} + P_{sk}C_{sk} = w_k L_k \tag{7}$$

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

$$\omega_{jk} = \frac{P_{jk}C_{jk}}{w_k L_k} = \Omega_j \left(\frac{P_{jk}}{P_k}\right)^{1-\sigma} \left(\frac{w_k L_k}{P_k}\right)^{\epsilon_j-(1-\sigma)} \tag{8}$$

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

$$P_k = \left[\sum_{j\in\{a,m,s\}} \left(\Omega_j P_{jk}^{1-\sigma}\right)^{\frac{1-\sigma}{\epsilon_j}} \left(\omega_{jk}(w_k L_k)^{1-\sigma}\right)^{\frac{\epsilon_j-(1-\sigma)}{\epsilon_j}} \right]^{\frac{1}{1-\sigma}} \tag{9}$$

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

$$e(U_k|P_{ak}, P_{mk}, P_{sk}) = \left[\sum_{j \in \{a,m,s\}} \Omega_j U_k^{\epsilon_j} P_{jk}^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{10}$$

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.

$$Y_{jk} = \left(\int_0^1 \overline{y}_{ijk}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}} \tag{11}$$

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

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:

$$y_{ijk} = z_{ijk} * l_{ijk} \tag{12}$$

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

$$\text{and } Z_{jk} = f(\mu_{jk}, T_{jk}, E(T_{jk})) \tag{13}$$

The sector-country specific aggregate productivity parameters, Z_{jk}, connect the model to the empirical results in Section 2. 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 that shape baseline productivity, such as technology, institutions, and human capital. In the counterfactuals in Section 5, climate change enters the model by perturbing the vector of Z_{jk} with the empirically estimated productivity impacts that vary at the country-sector level.

Trade

The trade portion of the model follows Eaton and Kortum (2002). When selling to foreign countries, intermediate goods producers face an iceberg trade cost, τ_{jkn}, that varies by sector, j, exporter country, k, and importer country, n. 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:

$$p_{ijk} = \min_n \left\{ \frac{\tau_{jnk} w_n}{z_{ijn}} \right\} \tag{14}$$

The sectoral final goods prices are given by the CES price index of all intermediate varieties used in that sector, which can be expressed as follows with the Fréchet assumption:

$$P_{jk} = \left(\int_0^1 p_{ijk}^{1-\eta} di \right)^{\frac{1}{1-\eta}} = \Gamma \left(1 + \frac{1-\eta}{\theta} \right)^{\frac{1}{1-\eta}} \left[\sum_{n \in N} Z_{jn} (\tau_{jnk} w_n)^{-\theta} \right]^{-1/\theta} \tag{15}$$

Finally, the final goods producer's demand function for variety i is given by:

$$\overline{y}_{ijk} = \left(\frac{p_{ijk}}{P_{jk}} \right)^{-\eta} Y_{jk} \tag{16}$$

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.

Trade flows can be expressed as follows, where π_{jkn} represents the share of varieties in sector j in country n that are sourced from country k:

$$\pi_{jkn} = \frac{Z_{jk} (\tau_{jkn} w_k)^{-\theta_j}}{\sum_{m=1}^{N} Z_{jm} (\tau_{jmn} w_m)^{-\theta_j}} \tag{17}$$

This representation of trade incorporates Ricardian comparative advantage 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 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.

$$w_k L_k = \sum_{j \in \{a,m,s\}} \left(\pi_{jkk} P_{jk} C_{jk} + \sum_{n \neq k} \pi_{jkn} P_{jn} C_{jn} \right) \tag{18}$$

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:

$$L_k = L_{ak} + L_{mk} + L_{sk} \tag{19}$$

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 presence of trade the employment share equals the production share of revenues in each sector, incorporating net exports, which yields the following equation that also appears in the model of Uy, Yi and Zhang (2013):

$$l_{jk} = \pi_{jkk}\omega_{jk} + \sum_{n \neq k} \pi_{jkn}\omega_{jn} \frac{w_n L_n}{w_k L_k} \tag{20}$$

This condition illustrates the importance of both domestic consumer preferences and international trade in governing the allocation of labor across sectors. Intuitively, Equation 20 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 7, 8, 9, 11, 14, 15, and 16 hold) and trade balances (Equation 18 holds).

Willingness-To-Pay

I calculate the willingness-to-pay to avoid a given climate change impact as equivalent variation (EV) using the non-homothetic 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:

$$WTP_k = -EV_k = -\left[e(U_k^1; P_{ak}^0, P_{mk}^0, P_{sk}^0) - w_k^0\right] \tag{21}$$

3.2 Comparative Statics

I now use the model to show how non-homothetic preferences, low substitutability across sectors, and trade frictions combine to create a "food problem," which I define as follows:

Definition 1: A country, k, suffers from a "food problem" if an exogenous decrease in its agricultural productivity, $dZ_{ak} < 0$, raises its agricultural employment share, $dl_{ak} > 0$.

Proposition 1: *In a given small country, k, a decline in agricultural productivity ($dZ_{ak} < 0$) raises k's expenditure share in agriculture if preferences for food are non-homothetic and decreasing with real* *income ($\epsilon_a < 1-\sigma$, $\epsilon_a < \epsilon_m$, $\epsilon_a < \epsilon_s$), and sectoral goods are not substitutable ($\sigma < 1$). The fall in agricultural productivity also reduces k's production share of agricultural consumption in the domestic economy ($d\pi_{akk} < 0$) and all other countries ($d\pi_{akn} < 0 \;\forall\; n \neq k$). A "food problem" occurs in k if the first effect outweighs the second, such that the following condition holds:*[11]

$$dl_{ak} > 0 \qquad \text{if} \qquad \underbrace{\pi_{akk} d\omega_{ak}}_{\Delta \text{ Expenditure Share (+)}} + \underbrace{d\pi_{akk}\omega_{ak}}_{\Delta \text{ Domestic Production Share (-)}} + \sum_{n \neq k} \underbrace{d\pi_{akn}\omega_{an}}_{\Delta \text{ Gross Exports (-)}} > 0 \tag{22}$$

Corollary 1: *No small country k suffers from a "food problem" if trade is frictionless ($\tau_{jkn} = 1$ for all j, k, n).*

Corollary 2: *No small country k suffers from a "food problem" if preferences are homothetic CES ($\epsilon_j = 1-\sigma$ for all j) with $\sigma \geq 1$).*

The condition in Equation 22 follows directly from Equation 20, which shows that a sector's employment share depends on its expenditure share and net exports.[12] To see the intuition for how these objects respond to agricultural productivity, I start with the following equation for agriculture's expenditure share, expressed in logs:

$$\ln(\omega_{ak}) = \ln(\Omega_a) + \underbrace{(1-\sigma)\ln\left(\frac{P_{ak}}{P_k}\right)}_{\text{Substitution Effect}} + \underbrace{(\epsilon_a - (1-\sigma))\ln\left(\frac{w_k}{P_k}\right)}_{\text{Income Effect}} \tag{23}$$

Equation 23 shows that an agriculture-biased reduction in productivity - such as the projection for hot countries in Section 2.2 - has two effects on the expenditure share in agriculture.[13] First, falling 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 in the estimates presented in Section 4, then the reduction in real wage drives up the expenditure share on food, ω_{ak}. This is the effect of non-homotheticity. 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 ($\frac{P_{ak}}{P_k}$). If $\sigma < 1$, as is also the case in Section 4, then the rising relative price in agriculture raises its expenditure share. If food is not substitutable with other consumption, then its relative quantity falls less than the relative price rises, and the share of spending

[11] See Appendix E.1 for proofs of the proposition and both corollaries.

[12] Note that Equation 20 and Proposition 1 hold in the present model where labor is the only factor of production, there are no distortions affecting labor mobility across sectors, and revenue shares equal employment shares. In a more general model with both labor and capital represented, the direction of movement of agriculture's employment share depends also on the substitutability between labor and capital and the factor bias of the productivity shock. With a sufficient degree of substitutability, a labor-biased productivity decline in agriculture might reduce agricultural employment even if the agricultural revenue share of production rises. See Alvarez-Cuadrado, Van Long and Poschke (2017) for more details.

[13] This equation also appears in Comin, Lashkari and Mestieri (2021). They estimate that non-homotheticities (the income effect) account for about 75% of observed historical structural transformation, with changes in relative prices (the substitution effect) accounting for the rest.

on food goes up. 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.[14] This is similar to the logic that underlies Baumol's cost disease (Baumol and Bowen, 1966), a theory that endeavors to explain why low-substitutability service sectors with relatively low productivity growth tend to rise as a share of expenditures over time.

While these consumer preference forces contribute to a climate-driven "food problem," they do not constitute a sufficient condition to cause one on their own. Equation 22 shows that even if the expenditure share in agriculture rises, countries could, in principle, meet domestic demand by increasing agricultural imports. Equation 17 shows that a fall in domestic agricultural productivity, Z_{ak}, will reduce both π_{akk} and π_{akn}, exerting downward pressure on the employment share in agriculture as countries move toward raising food imports and reducing food exports. Thus, climate change causes a "food problem" only if this shift in trade flows is not large enough to outweigh the rising expenditure share. Corollaries 1 and 2 show that neither trade frictions nor non-homothetic and low-substitutability preferences alone can create a "food problem." Falling productivity raises employment in a sector only when consumers have limited capacity both to reduce consumption of that good and to substitute freely toward imports.

Proposition 1 shows the conditions under which the climate-driven "food problem" can exist, but not whether it is empirically relevant. The relative strength of the key mechanisms depends on both the preference parameters, ϵ_a and σ, and the matrix of bilateral trade frictions, τ_{jkn}, that govern the strength of the movements in the expenditure shares and trade shares. In the next section, I calibrate the model in order to quantify the magnitudes of these competing forces in the climate change counterfactuals.

4 Model Calibration

I solve the model numerically in levels, and use a simulated method of moments procedure along with estimates from the literature to calibrate the parameters.

4.1 Parameter Estimates

I use a combination of calibration and estimation to set the model parameters. I start by setting the trade elasticities to the values estimated by Tombe (2015): $\theta_a = 4.06$, and $\theta_m = 4.63$. I calibrate the baseline sectoral productivity parameters, Z_{jk}, to match the relative levels of sectoral value-added per worker and country level nominal GDP per capita from World Bank National Accounts data in 2011, which represents the baseline period. The bilateral sectoral trade cost parameters, τ_{jkn}, are

[14] These features of the model also explain why its predictions about the protective effects of reallocation differ from those of Costinot, Donaldson and Smith (2016). Their paper estimates an elasticity of substitution of 5.4 across crop varieties and 2.8 across crops, whereas this paper estimates an elasticity of 0.27 across sectors (see Section 4).

calibrated to match bilateral trade flows by sector from Comtrade, in which I classify HS 1988/92 codes 1-24 as agriculture and 28-97 as manufacturing. Appendix E.2 contains more details about the data construction.

The remaining parameters concern the non-homothetic CES preference specification. I estimate these using a simulated method of moments procedure that minimizes the sum of squared errors between simulated and empirical GDP shares across all countries, conditional on all other parameters. 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 with 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.[15]

Table 3b displays the estimated preference parameters. Two parameters in particular are critical in driving the "food problem" described in Section 3.2. First, the estimated cross-sector elasticity of substitution, $\sigma = 0.27$, means that raising the relative price of agriculture will increase its expenditure share through the substitution effect term in Equation 23. Second, the estimated agricultural non-homotheticity parameter, $\epsilon_a = 0.29$, implies that the consumption share of agriculture is strongly diminishing in real income. With the estimated parameters, Equation 23 shows that the income effect term is $\epsilon_a - (1 - \sigma) = -0.44$, implying that the expenditure share in agriculture will rise if a climate change shock causes real income to fall.

Note that the preference parameters estimated by targeting the global cross-section of sectoral shares match up closely with estimates from various historical panel regressions in Comin, Lashkari and Mestieri (2021). Their estimates for σ range from 0.2 to 0.6 across specifications, and for sectoral income elasticities range from 0.37 to 0.56 for agriculture, 0.83 to 1.03 for manufacturing, and 1.14 to 1.20 for services. The corresponding income elasticities from this paper's preference parameter estimates are 0.48 for agriculture, 0.98 for manufacturing, and 1.09 for services.[16]

4.2 Model Fit

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

[15] 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 6b shows that the pattern of relative prices in the model is similar to that of the data.

[16] The formula for the income elasticity in sector $j \in \{a, m, s\}$ 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.

tivity 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-31, 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).

The model explains about 96% of the variation in average agricultural GDP shares by decile of global income. Figure 6a shows that the simulations closely reproduce the average global pattern of agricultural specialization declining smoothly with log GDP per capita across countries. At the country level the model explains over 60% of the variation, which is substantially better than the 43% fit with Stone-Geary preferences.[17] The remaining variation not explained by the model represents idiosyncratic differences in agricultural specialization across countries, conditional on income levels. I choose not to introduce additional degrees of freedom to further refine the fit of agricultural specialization in order to keep the model parsimonious. This allows the counterfactuals to focus on the key forces described in Proposition 1 - non-homothetic consumer preferences, Ricardian comparative advantage, and barriers to trade. Other country-specific determinants of agricultural specialization - such as output subsidies and taxes, public procurement quotas, interventions in input markets for fertilizer and water, land market regulations, and public insurance schemes - affect agriculture's production shares and vary across countries in the baseline equilibrium, but are implicitly held constant in this paper's counterfactuals that focus on trade policy.[18] I show results in Section 5 by grouping of the global income and temperature distribution since the model closely matches patterns of sectoral specialization at this level of aggregation more so than for individual countries.

For further validation of the model calibration, I also show the model's fit to a key untargeted moment: the global pattern of high relative prices for agricultural consumption in poor countries. In Figure 6b, 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.

5 Model Counterfactuals

This section uses the model to quantify the role of the "food problem" described in Section 3.2 in mediating the welfare effects of global warming. I start with a counterfactual in the calibrated model with

[17] Non-homothetic CES preferences improve model fit substantially compared to generalized Stone-Geary preferences, particularly in middle income countries. I show robustness to using Stone-Geary preferences in Table 6 and Appendix F.1.

[18] See Anderson et al. (2008) for a discussion of the range of domestic policy interventions in agricultural markets.

existing levels of trade costs that shows how the productivity effects of climate change entrench specialization in agriculture in low-income countries where its productivity suffers most. The counterfactuals that follow show how lowering barriers to trade can reverse this "food problem" mechanism and instead allow for sectoral reallocation to contribute to climate change adaptation. In the final part of the section, I consider the implications of economic growth for sectoral specialization and climate change adaptation, as well as a number of robustness exercises.

5.1 Reallocation and Welfare With Existing Trade Costs

In the first counterfactual, I take the calibrated model from Section 4 and adjust each sector-country aggregate productivity parameter, Z_{jk}, by the projected climate change impact from Section 2.2. I then calculate equilibrium wages, prices, sectoral shares, and trade flows in counterfactual scenarios with and without climate change, and in intermediate counterfactuals that allow for decomposing the trade and consumption forces driving reallocation.

Recall from Equation 20 and Proposition 1 in Section 3.2 that sectoral reallocation depends on changes in trade flows and consumption shares across sectors. When poor, hot regions suffer large declines in agricultural productivity, they could in principle respond by importing more food from less vulnerable locations. Table 4a shows that the magnitude of this mechanism is modest in the model's simulations. In the baseline calibration, the poorest quartile of countries (most of which are also very hot) consume about 89% domestically produced food. When hit with the climate change shock, the domestic production share of consumption falls by only 1.8%. The model suggests that the barriers preventing these regions from importing much food in the baseline simulation also deter them from increasing imports substantially in response to climate change.

More broadly, the counterfactual changes in trade flows follow the direction implied by the global comparative advantage projections shown in Section 2.2, but the magnitude of the response is generally modest as a share of the economy. Colder countries with neutral to positive effects on agriculture generally increase net exports of food in the model simulations, though from a low base. For instance, Norway and Canada double and quadruple their net exports of food from 0.9% to 1.8% and 0.5% to 1.9% of GDP, respectively. Conversely, most hot countries in Sub-Saharan Africa and South Asia increase imports of food, with a few exceptions for whom the relative decline in agricultural productivity is small compared with their regional neighbors and trading partners. For instance, the Democratic Republic of Congo is an extremely poor country with most of its land located at high altitude, which gives it a relatively temperate climate for its region. Thus, its projected manufacturing effects are large relative to its agricultural effects in comparison with neighboring countries, and agricultural net exports rise in the climate change counterfactual. The change in net exports of agriculture is under 5% of GDP for the vast majority of countries, and under 2% of GDP on average in the hottest half of the world.

The limited response of global trade flows leaves scope for the consumption response to play a major role in governing sectoral reallocation. Tables 4b and 4c show the impact of climate change on the agricultural share of GDP and welfare across groups of countries, and decompose the relative strength of each mechanism by running separate counterfactuals with and without trade adjustment. The left column, labeled "No Reallocation," shows a counterfactual that allows for neither expenditure shares nor trade flows to adjust, such that agriculture's share of GDP is held fixed at its baseline value. The middle column, labeled "Autarky," considers a scenario that allows for only the expenditure share, ω_{ak}, to adjust, with price changes equal to the inverse of the projected change in productivity. Finally, the right column, labeled "Full Reallocation," shows the effects with the full adjustment of both expenditure shares and trade flows. Figure 7 shows the corresponding global map of the welfare effects of climate change in equilibrium, as well as a second map illustrating the key channel of rising food prices in vulnerable regions.

Table 4b demonstrates the critical role of the "food problem" in governing the response to temperature change. In the poorest quartile of the world's countries, the expenditure share in agriculture rises by four percentage points in the counterfactual that assumes autarky.[19] This follows from the income effect and substitution effect displayed in Equation 23, both of which play a large role in this setting with non-homothetic preferences and low substitutability across sectors. Climate change makes the poorest people poorer and makes food more expensive, and since people cannot substitute away from eating, they raise the share of their income spent on food. As discussed above, food imports in such places also increase, but the third column of Table 4b shows that the trade adjustment covers only about a quarter of the increase in the consumption share of food in low-income countries. While some of these countries have a stronger trade response in the simulation and shift specialization away from agriculture, on net warming raises agriculture's share of the labor force by about 14%, or 2.8 percentage points, in the lowest quartile of income. While the effect is strongest for low-income countries, the table also shows that agricultural specialization increases across all quartiles of the world's income and temperature distributions. This occurs both because of increased exports raising agricultural specialization in colder regions, and because the global expenditure share in agriculture rises from 3.8% to 4.3% as global agricultural productivity falls on average.

Table 4c shows the welfare effects of climate change in each counterfactual. The left column shows that welfare would fall especially severely in poorer and hotter countries if both expenditure shares and trade flows were held constant while prices and real wages adjust. This naïve hypothetical forces consumers to deviate from their optimal consumption bundles and thus does not represent an equilibrium, but helps illustrate that the adjustment in consumption shares is itself a critical mechanism of adaptation. While the food "problem" raises agricultural specialization and keeps

[19]Note that quartiles of the income and temperature distribution in Table 4 are based on the present day distribution of the global population.

more workers in the most vulnerable sector, the alternative in which people sharply reduce the quantity of food they consume when prices rise and incomes fall is dramatically worse.

In contrast, the difference in welfare effects between the autarky and full adjustment counterfactuals is almost negligible. 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 when accounting for trade. A few countries with an especially strong trade response constitute exceptions to this pattern in the simulation. For instance, trade reduces the welfare costs of warming by about 23% in Niger, which has one of the largest counterfactual increases in agricultural net imports (12.5 percentage points of GDP) and reallocates production substantially away from agriculture. Overall, however, the results suggest that trade contributes little to climate change adaptation because of the muted role it plays in the most severely affected countries. The next section investigates how potential changes in trade policy could amplify its importance in reducing the welfare impacts of climate change.

5.2 Trade Policy Counterfactuals

This section describes several counterfactual exercises with varying levels of trade barriers, τ_{jkn}, that demonstrate the potential role of trade and trade policy in alleviating the "food problem" and the welfare consequences of climate change. I start with a counterfactual that considers frictionless trade, $\tau_{jkn} = 1$ for all sectors and countries. This hypothetical is not achievable in practice but provides a measure of the theoretical maximum impact of this channel of adjustment. A second counterfactual provides a somewhat more realistic benchmark for the potential magnitude of gains by setting all bilateral trade costs for manufacturing and agriculture to approximately the 90th percentile of trade openness observed in the sample; a 100% tariff-equivalent value ($\tau = 2$). This scenario takes as given that even the most open countries face obstacles such as shipping costs, contracting frictions, and language barriers, and imagines a world in which trade costs in poor countries fall to approximately the level observed in rich countries today. The final two counterfactuals make more concrete policy evaluations that quantify the proportion of the hypothetical gains that can be achieved by specific instruments such as reducing tariffs, entering free trade agreements, and eliminating observable regulatory frictions.

To disentangle the benefits of trade for climate change adaptation from the more general gains from trade, I rescale the sectoral productivity parameters, Z_{jk}, in each counterfactual such that I continue to match the targeted levels of GDP per capita in the baseline equilibrium. It is worth noting, however, that without the estimated trade costs the model can no longer match the observed global distribution of agricultural specialization. In these hypothetical scenarios with increased openness, developing countries import substantially more food from richer countries even in the baseline equilibrium without climate change, consistent with the relatively low productivities of their agricultural sectors shown in Figure 1b. Thus, part of the mechanism through which trade facilitates adaptation occurs by reallocating production away from the most vulnerable sector even before accounting for the response to rising temperatures.

Table 5 shows the effects of climate change by trade cost counterfactual across groups of countries. Each panel shows a different trade cost scenario, and the columns show the impact of the productivity changes on the domestic production share of agricultural expenditures, agriculture's share of GDP, welfare, and food prices in hot, cold, rich, and poor economies. In the frictionless trade case, global food prices perfectly equalize across countries and the burden of warming is shared more evenly across space, though poor countries still suffer larger declines in welfare because their domestic productivity falls by more. Overall, eliminating all trade frictions reduces the welfare costs of warming by about 22% for the world, and by about 59% for the poorest quartile. In the more plausible benchmark scenario that moves all countries to the frontier of current global openness, the welfare costs fall by about 13% globally and by about 43% for the poorest quartile. In these scenarios that dramatically increase global tradability, the impact of global warming on food prices in poor countries is about a quarter lower, and the impact on welfare reduced by nearly half.

The effects of trade openness on climate change vulnerability vary tremendously across countries. While the poorest countries face lower climate damages when trade costs are lower, the willingness-to-pay to avoid climate change is actually *higher* in the low trade cost scenario with $\tau = 2$ than in the baseline counterfactual for 35 countries representing 13% of the global population. To be clear, these countries still experience overall gains from trade, but suffer larger climate change damages once those general gains are netted out. The intuition for this surprising result is as follows. When trade barriers are high and local consumption depends mostly on local production, the effects of deteriorating productivity are also concentrated locally. Conversely, more trade makes the world more interdependent and dilutes the effects of a local shock across many countries. If consumption in Austria is more linked to production in Zimbabwe, then Austrian consumers suffer more from shocks that hit Zimbabwe. Conversely, Zimbabwean consumers insulate themselves from the local shock by consuming a more diversified global portfolio of products. Thus, the global adaptation gains from trade openness represent the net impact of both increased and decreased vulnerability across countries.

The first two trade cost counterfactuals in Table 5 demonstrate the potential effects of tradability on alleviating the "food problem," but do not quantify the role of particular policies. To enable counterfactuals that evaluate the impact of specific policy instruments that facilitate trade, I run the following regression to decompose the share of the model calibrated trade costs, τ_{jkn}, that is explained by various policy and non-policy factors:

$$\begin{aligned}\ln(\tau_{jkn} - 1) = {} & \beta_{1j}\text{Tariff}_{jn} + \beta_{2j}\mathbb{1}(\text{Trade Agreement}_{kn}) + \beta_{3j}\text{Import Fees}_n \\ & + \beta_{4j}\text{Import Time}_n + \beta_{5j}\text{Distance}_{kn} + \beta_{6j}\mathbb{1}(\text{Contiguous}_{kn}) \\ & + \beta_{7j}\mathbb{1}(\text{CommLang}_{kn}) + \beta_{8j}\mathbb{1}(\text{Col}_{kn}) + \beta_{9j}\mathbb{1}(\text{CommCol}_{kn}) + \epsilon_{jkn}\end{aligned} \tag{24}$$

Equation 24 models the log of the calibrated sectoral bilateral trade costs as a function of four policy variables and five exogenous factors. The policy variables are sectoral import tariffs, an indicator for mutual participation in a trade agreement, total fees associated with regulatory clearance (exclusive of tariffs), and import processing times in days.[20] The exogenous factors are the average distance between countries using nighttime lights to weight locations within each country (Hinz, 2017), an indicator for contiguous borders, and indicators for common official or primary language, a colonial relationship post-1945, and a colonial relationship with a common country post-1945. Data for import tariffs, processing fees, and time delays come from the World Bank, and data for all other variables come from the CEPII Gravity Database (Conte, Cotterlaz and Mayer, 2022).

Appendix Table A-6 displays the results from estimating Equation 24 for agricultural and manufacturing trade costs. The estimates show strong partial correlations between the policy variables and the model's inferred trade costs, especially for agriculture. Participation in a common trade agreement is associated with a 36 log point decline in agricultural trade costs. Similarly, moving tariffs, regulatory import fees, and import processing delays from the 95th percentile to their minimum value would reduce agricultural trade costs by about 34, 50, and 71 log points, respectively. While these estimates rely on cross-sectional variation, the control variables for physical distance, contiguity, common language, and colonial relationships unsurprisingly explain a large amount of the variation in trade costs, and their coefficients have similar magnitudes in specifications that exclude the endogenous policy variables. The correlations in the residual variation are consistent with tariffs, trade agreements, and regulatory barriers playing a major role in shaping global agricultural trade.

In the bottom two panels of Table 5, I use the results from the model trade cost decomposition to run counterfactuals that consider specific policy scenarios.[21] In the first of these counterfactuals, all bilateral trading pairs enter trade agreements and eliminate tariffs. Specifically, I adjust $\ln(\tau - 1)$ for all trade costs in the model using the coefficients from Appendix Table A-6 and then exponentiate to get a new matrix of τ_{jkn}. I restrict all values of τ to be at least one in the counterfactual so that trade costs are never negative. The table shows that the effects of climate change on sectoral reallocation change substantially in this scenario. For the poorest quartile of countries, liberalizing trade policy reverses the sign of climate change's effect on the agriculture's share of GDP. Global warming raises the import share of agricultural consumption by 15 percentage points, as compared to 1.8 in the baseline with estimated trade costs, and reduces agriculture's share of labor by 1.9 percentage points instead of raising it by 2.8 percentage points. Thus, the model suggests that entering trade agreements and reducing tariffs would prevent the aspect of the "food problem" in which global warming strengthens specialization in the most vulnerable sector in the most vulnerable regions.

Globally, the welfare costs of warming fall by over 10% in this scenario with universal trade agree-

[20]I use the simple mean of tariffs across primary products for the agricultural sector and manufactured products for the manufacturing sector.

[21]Note that Costinot, Donaldson and Smith (2016) perform a similar decomposition in order to isolate the share of trade costs associated with physical distance and run a counterfactual that considers within-country trade costs.

ments and no tariffs, representing most of the gains that could be achieved in the benchmark scenario with $\tau = 2$. However, the poorest quartile achieves only about 7% of the reduction in climate damages experienced in the $\tau = 2$ scenario, despite alleviating the climate change aspect of the "food problem." This is because the non-climate related aspects of the "food problem" remain. Baseline specialization without climate damages is still concentrated in the vulnerable agricultural sector in poor countries in the scenario with trade agreements and tariff reductions. The average agricultural GDP share of the poorest quartile falls from 20.1% in the current policy baseline to 2.8% in the $\tau = 2$ scenario, but remains at 11.3% in the trade policy liberalization scenario. Intuitively, the simulations suggest that trade agreements and tariff reforms would prevent climate change from *exacerbating* the "food problem," but only incrementally alter the broader pattern of agricultural specialization in poor countries.

The final counterfactual in Table 5 considers the further step of reducing "red tape" barriers to importing. Tombe (2015) shows that regulatory frictions concentrate disproportionately in low-income countries, and argues that they represent the most important source of trade barriers in these locations. The average country in Sub-Saharan Africa requires 9 documents, over $2700 in fees, and a 37 day wait for customs clearance, document processing, and inspection procedures, exclusive of tariffs and unofficial payments. Processing delays have particularly severe consequences for agricultural goods. Hummels and Schaur (2013) estimate that each day of waiting adds a 3% tariff-equivalent cost for food products, over 50% more than for non-food. Appendix Figure A-32 shows that these regulatory obstacles are dramatically lower in other regions, consistent with the pattern in which low-income countries have an especially low import share of agricultural consumption (9%) in the baseline equilibrium.

The bottom panel of Table 5 shows that reforming observable regulatory frictions contributes more substantially to climate change adaptation, especially in low-income countries. The results show that the combination of trade agreements, tariff reductions, and reducing processing fees and time delays to the minimum value in the sample ($300 and 3 days) reduces the global welfare costs of warming by nearly 17%, achieving over three quarters of the potential adaptation gains from frictionless trade.[22] For the poorest quartile, however, this combination of policy reforms reduces the welfare costs of warming by only about a third of the adaptation gains achieved in the $\tau = 2$ benchmark (14% compared to 42%).

Overall, the results from Table 5 show that trade policy could play a critical role in facilitating adaptation to climate change. However, the counterfactuals also suggest that trade agreements, tariff reductions, and measurable red tape barriers can achieve only about one-third of the sizable hypothetical adaptation gains from moving poor countries to rich country levels of tradability. For the por-

[22]Note that the results in Appendix Table A-6 show a small negative association between regulatory barriers and manufacturing trade costs, for which such processing delays are less important, so this counterfactual only eliminates these barriers for agricultural trade costs.

tion of the model's trade costs not explained by the variables considered in Table A-6, it is worth noting that Adamopoulos (2011) finds that transportation costs are about 16 times higher in low-income countries than high-income countries. Atkin and Donaldson (2015) similarly find that within-country distance costs are four to five times higher for manufactured goods in Nigeria and Ethiopia than in the U.S., and Porteous (2019) further shows that the same holds for agriculture and that these costs are especially high in locations with unpaved roads. While it is not feasible to consider counterfactuals that improve road and port quality in a country level model, the benchmark free trade counterfactuals suggest that any reforms that improve tradability in poor countries, particularly in agriculture, could reduce their vulnerability to global warming.

5.3 Extensions and Robustness

This section considers the robustness of the main model results and an extension that quantifies the value of economic growth for climate change adaptation. The first panel of Table 6 shows the main counterfactual results under a variety of alternative assumptions. First, I simulate a version of the model that uses generalized Stone-Geary preferences to represent preferences for food with a sharp minimum subsistence requirement rather than the smoothly declining budget share implied by the non-homothetic CES specification. Appendix F.1 covers the calibration details for this version of the model, which produces results very similar to the baseline for agriculture's share of labor, the welfare costs globally and in poor countries, and the change in food prices. The second, third, and fourth sets of alternative results in the table use the individual productivity estimates from Hultgren et al. (2021), Iglesias and Rosenzweig (2010), and Cline (2007), rather than the average across the range of sources. Again, the results for the key dimensions of the "food problem" are broadly similar, though the welfare effects 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)). See Section 2.2 and Appendix E.4 for details about the strengths and weaknesses of each source of agricultural estimates.

In Appendix F.2, I consider a version of the model with multiple factors of production, such as land or heterogeneous labor.[23] I do not calibrate this version of the model, but show that the model's key comparative statics regarding the "food problem" remain qualitatively unchanged. I also show that this version of the model produces two additional insights about the results. First, with heterogeneous worker types, those that are more concentrated in agricultural production will gain less from trade openness. Intuitively, trade barriers benefit those employed in less productive agricultural sectors by keeping domestic food prices high. Second, in a model with multiple inputs, comparative advantage

[23] I do not explicitly consider a version of the model with intermediate inputs to production, but note that trade policy generally has greater implications for real income in such settings (Costinot and Rodríguez-Clare, 2014). Trade in intermediate inputs also plays an important role in the international transmission of technology and the evolution of agricultural productivity over time (Farrokhi and Pellegrina, 2023).

in a given sector endogenously weakens as specialization shifts towards that sector because of rising relative input prices. 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 lower trade costs.

The bottom panel of Table 6 shows how allowing the baseline productivities in the model to evolve with economic growth affects the results. In these counterfactuals, I rescale the productivity parameters, Z_{jk}, for each country to match the magnitude of economic growth projected by end-of-century in several Shared Socioeconomic Pathway scenarios developed by Cuaresma (2017).[24] Economic growth contributes to adaptation along two dimensions in the model. First, the agriculture share of GDP declines as countries grow richer due to non-homothetic preferences for food, reducing the consequences of agriculture-biased productivity shocks. Second, the results from Section 2 imply reduced temperature sensitivities within sector as countries grow richer, which I capture by re-evaluating the sensitivities at future projected levels of log GDP per capita. Thus, the table shows that the welfare effects of warming are lower for the world when accounting for the benefits of growth, and especially so in low-income countries. The costs of warming fall most in scenarios that project more growth and more convergence of global incomes, underscoring the power of growth for reducing the burden caused by global warming. Critically for the research question in this paper, however, growth does not eliminate the effects of the "food problem." The global warming counterfactuals raise agriculture's share of the labor force in poor countries across all scenarios. While economic growth accelerates the transition away from agriculture in poor countries, global warming slows it down.

6 Conclusion

This paper shows that non-agricultural production is substantially less vulnerable to global warming than agriculture, suggesting that countries in hot locations could potentially adapt effectively by reallocating production away from farming. However, model simulations show that warming is likely to instead strengthen the existing global pattern in which low-income countries specialize intensively in their relatively unproductive agricultural sectors. When rising temperatures reduce domestic agricultural productivity, subsistence consumer preferences for food and high barriers to trade combine to create a "food problem" that raises the share of workers allocated to the sector suffering the greatest productivity losses. Model simulations show that reducing barriers to trade can reverse the consequences of this climate-driven "food problem" in developing countries and substantially alleviate their exposure to the productivity consequences of warming.

The paper has three primary policy implications. First, the results inform a particular aspect

[24] I apply these changes as sector-neutral technological progress in Z_{jk}. Świecki (2017) shows that global productivity growth in recent decades has been similar in agriculture and manufacturing, though slower in services. The projections from Cuaresma (2017) have been used widely in research on the economics of climate change. See Carleton et al. (2022), Newell, Prest and Sexton (2021), and Kalkuhl and Wenz (2020) for examples.

of the literature on the economic costs of global warming. While this paper does not attempt to provide a holistic evaluation of these impacts since it omits critical factors such as growth effects, migration, and feedbacks between the economy and the climate, the results show the importance of accounting for the "food problem" in other work that measures climate damages and the social cost of carbon. Second, the results underscore the critical importance of technological innovation that can reduce the impacts of extreme weather on agriculture in hot, poor countries. If agricultural activity were likely to shift substantially away from these locations, investment in adaptation might be better focused on facilitating the transition to non-agriculture. Instead, the findings suggest that modernizing agricultural technology in developing countries and promoting innovations such as heat- and drought-resistant crop varieties could play a central role in successful adaptation.

The third key policy implication is that raising the level of trade integration in low-income countries could substantially reduce their exposure to global warming. This paper provides suggestive evidence that a meaningful portion of these gains can be achieved with policy instruments such as entering trade agreements, reducing tariffs, and reducing regulatory frictions that impede imports. Given the considerable magnitude of the potential adaptation gains from greater tradability, the results also suggest an important role for future research that provides more direct causal evidence on the relationship between specific policies and the low observed agricultural trade shares in low-income countries, and investigates the role of related factors like transportation infrastructure.

I conclude with a final suggestion for future research. While the results in this paper suggest that low-income countries in hot locations could adapt effectively to climate change by facilitating imports of food, in practice policymakers often show a stated and revealed preference for pursuing a notion of "food security" that instead prioritizes reducing food imports. If policymakers believe that promoting domestic production insulates them from global volatility driven by climate shocks and geopolitical risk, this could dissuade them from pursuing trade openness as a climate adaptation strategy. Thus, the results of this paper suggest that it would be useful to more closely examine the relationship between trade policy and food security, as well as the broader political economy of protectionism in agriculture in low-income economies.

References

Adamopoulos, Tasso, "Transportation costs, agricultural productivity, and cross-country income differences," *International Economic Review*, 2011, *52* (2), 489–521.

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-Cuadrado, Francisco, Ngo Van Long, and Markus Poschke, "Capital–labor substitution, structural change, and growth," *Theoretical Economics*, 2017, *12* (3), 1229–1266.

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

Anderson, Kym, Kym Anderson, Marianne Kurzweil, Will Martin, Damiano Sandri, and Ernesto Valenzuela, "Methodology for measuring distortions to agricultural incentives," Technical Report, World Bank 2008.

Atkin, David and Dave Donaldson, *Who's Getting Globalized?: The Size and Implications of Intra-national Trade Costs* number w21439, National Bureau of Economic Research Cambridge, MA, 2015.

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.

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.

Bilal, Adrien and Esteban Rossi-Hansberg, "Anticipating Climate Change Across the United States," Technical Report, National Bureau of Economic Research 2023.

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.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel, "Global non-linear effect of temperature on economic production," *Nature*, 2015, *527* (7577), 235–239.

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

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.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath et al., "Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits," *The Quarterly Journal of Economics*, 2022, *137* (4), 2037–2105.

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," *Journal of Economic Geography*, 2021, *21* (4), 493–530.

Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer, *The CEPII gravity database*, CEPII, 2022.

Costinot, Arnaud and Andrés Rodríguez-Clare, "Trade theory with numbers: Quantifying the consequences of globalization," in "Handbook of international economics," Vol. 4, Elsevier, 2014, pp. 197–261.

_ , 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.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken, "Temperature shocks and economic growth: Evidence from the last half century," *American Economic Journal: Macroeconomics*, 2012, *4* (3), 66–95.

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.

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.

Farrokhi, Farid and Heitor S Pellegrina, "Trade, technology, and agricultural productivity," *Journal of Political Economy*, 2023, *131* (9), 2509–2555.

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

Gollin, Douglas, 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.

Hinz, Julian, "The view from space: Theory-based time-varying distances in the gravity model," 2017.

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.

Hummels, David L and Georg Schaur, "Time as a trade barrier," *American Economic Review*, 2013, *103* (7), 2935–2959.

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, Benjamin F and Benjamin A Olken, "Climate shocks and exports," *American Economic Review*, 2010, *100* (2), 454–459.

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.

Lemoine, Derek, "The climate risk premium: how uncertainty affects the social cost of carbon," *Journal of the Association of Environmental and Resource Economists*, 2021, *8* (1), 27–57.

Liu, Maggie, Yogita Shamdasani, and Vis Taraz, "Climate change and labor reallocation: Evidence from six decades of the Indian Census," *American Economic Journal: Economic Policy*, 2023, *15* (2), 395–423.

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.

Porteous, Obie, "High trade costs and their consequences: An estimated dynamic model of African agricultural storage and trade," *American Economic Journal: Applied Economics*, 2019, *11* (4), 327–366.

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.

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.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge, "What are we weighting for?," *Journal of Human resources*, 2015, *50* (2), 301–316.

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.

Syverson, Chad, "Market structure and productivity: A concrete example," *Journal of political Economy*, 2004, *112* (6), 1181–1222.

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.

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 at End-of-Century

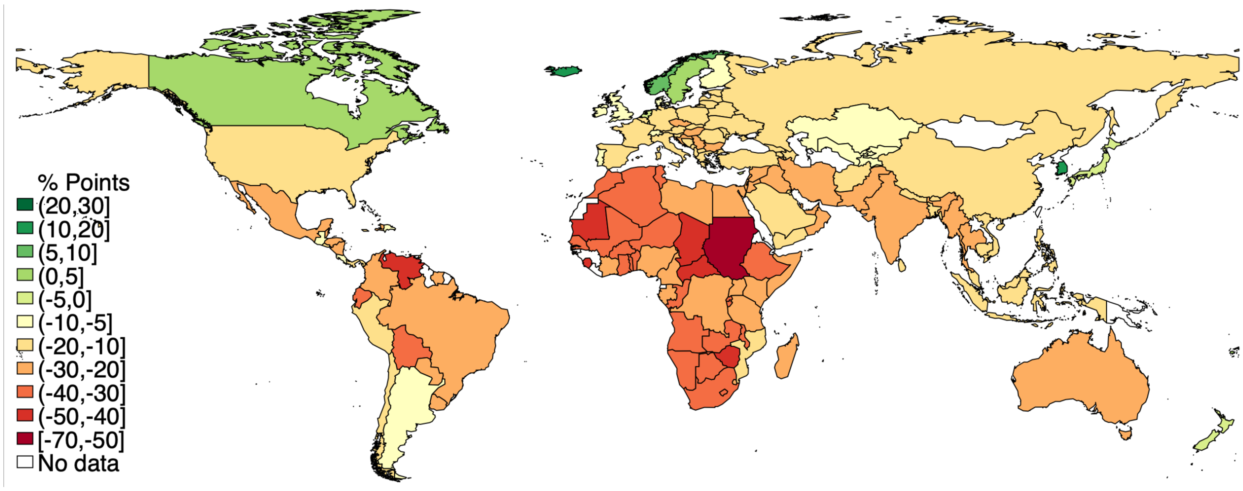

(b) Comparative Advantage and Specialization in Agriculture

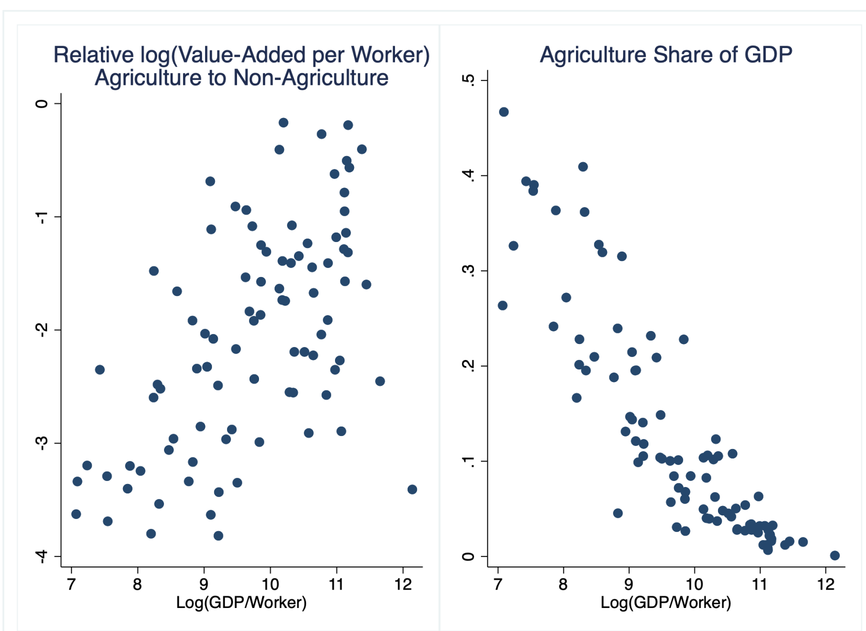

Notes: The map in Panel (a) shows the projected impact of climate change on agricultural productivity averaged across estimates from Hultgren et al. (2021), Cline (2007), Iglesias and Rosenzweig (2010), and Costinot, Donaldson and Smith (2016). See Appendix E.4 for more information on the methods and 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: Effects of Daily Temperature on Annual Revenue per Worker

	(1)	(2)	(3)	(4)	(5)	(6)
	Revenue/Worker	Revenue/Worker	Revenue/Worker	Revenue	Employment	Revenue/Worker
Cooling Degree Days Above 30 °C	-0.0000311	-0.0000403	-0.00119	-0.00250	-0.00131	-0.00100
	(-2.29)	(-3.36)	(-4.73)	(-6.80)	(-5.25)	(-4.03)
Heating Degree Days Below 5 °C	-0.0000315	-0.0000352	-0.000956	-0.00180	-0.000842	-0.000452
	(-2.15)	(-3.18)	(-2.15)	(-2.91)	(-1.92)	(-2.07)
Cooling Degree Days Above 30 °C X log(GDPpc)			0.0000715	0.000178	0.000107	0.0000595
			(4.07)	(6.79)	(6.06)	(3.65)
Cooling Degree Days Above 30 °C X Long-Run $\overline{T}$			0.0000186	0.0000334	0.0000148	0.0000160
			(4.85)	(6.24)	(3.93)	(3.96)
Heating Degree Days Below 5 °C X log(GDPpc)			0.0000898	0.000167	0.0000769	0.0000416
			(2.14)	(2.85)	(1.85)	(2.02)
Heating Degree Days Below 5 °C X Long-Run $\overline{T}$			-0.00000292	0.00000212	0.00000504	0.000000703
			(-1.54)	(0.93)	(2.85)	(0.59)
Number of Observations	4125776	4125776	4125776	4125776	4125776	17938084
Manufacturing	X	X	X	X	X	X
Services						X
Firm FE	X	X	X	X	X	X
Country X Year FE	X	X	X	X	X	X
Inverse Sample Size Weights	X					
GDP Weights	X					
Countries Included	15	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. Columns 1 and 2 show coefficients from estimating Equation 3 with and without weights, and Columns 3-6 show results from estimating Equation 4. Outcome variables come from the data sources listed in Appendix Table A-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. Figure 3a shows results for the United States and Appendix C shows results for China. Figure 2 evaluates the interaction terms from the results in Column 3 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

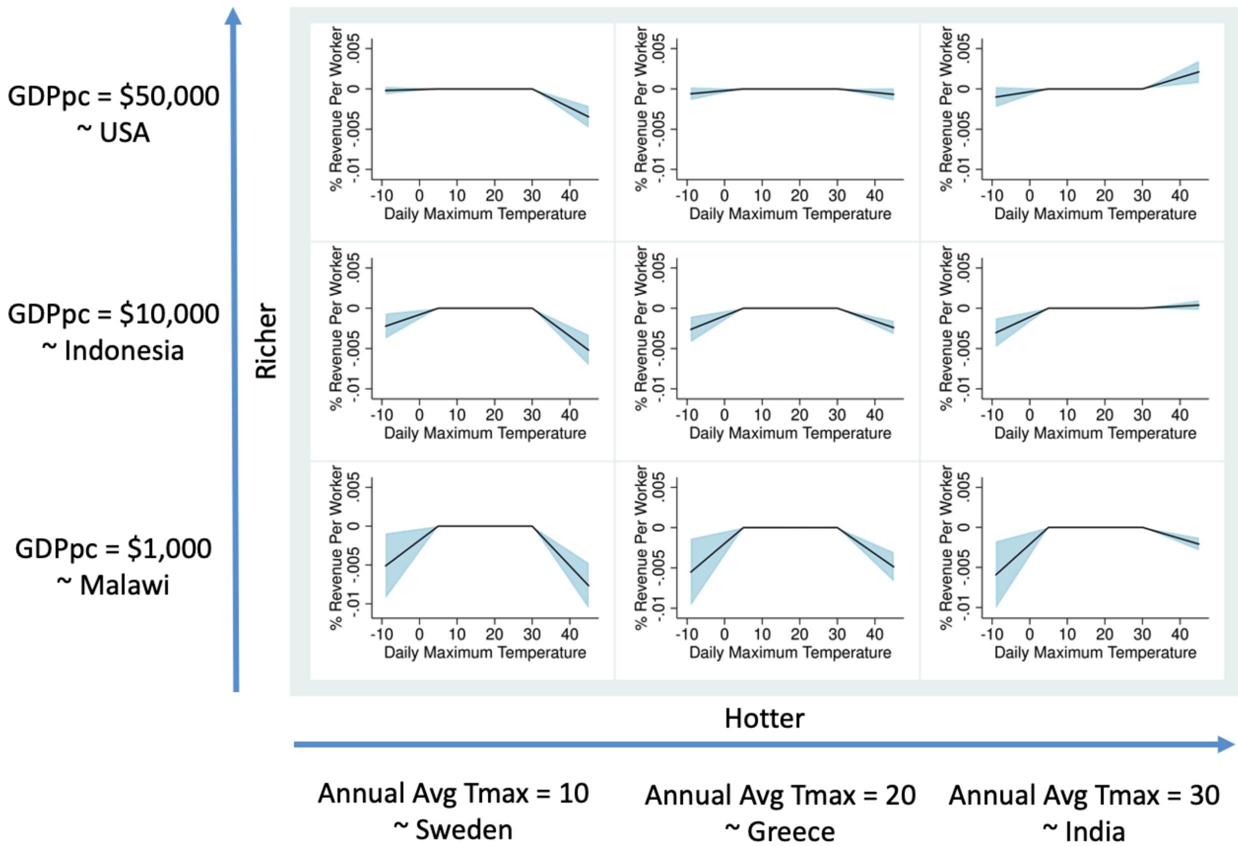

Notes: Graphs show the predicted effect of exposure to daily maximum temperature on the log of firm-level revenue per worker at varying levels of income and long-run average temperature by evaluating the interacted panel regression from Column 3 of Table 1. The specification includes firm and country-by-year 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

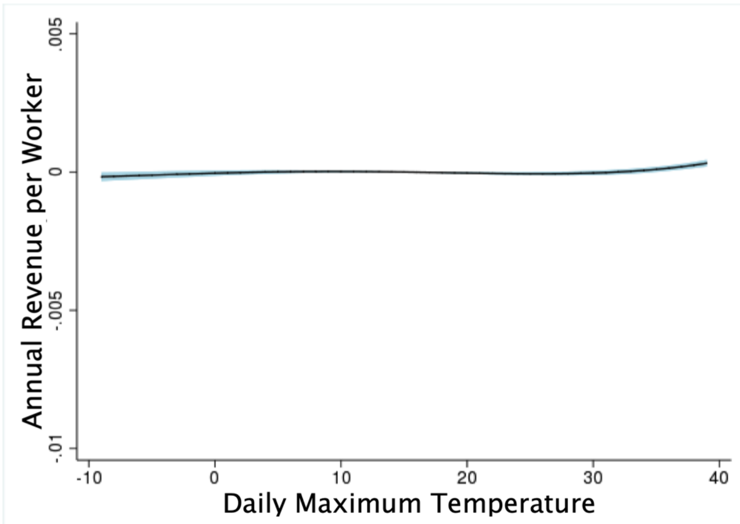

(b) Estimated Response of Annual Plant-Level Energy Expenditures to Daily Maximum Temperature

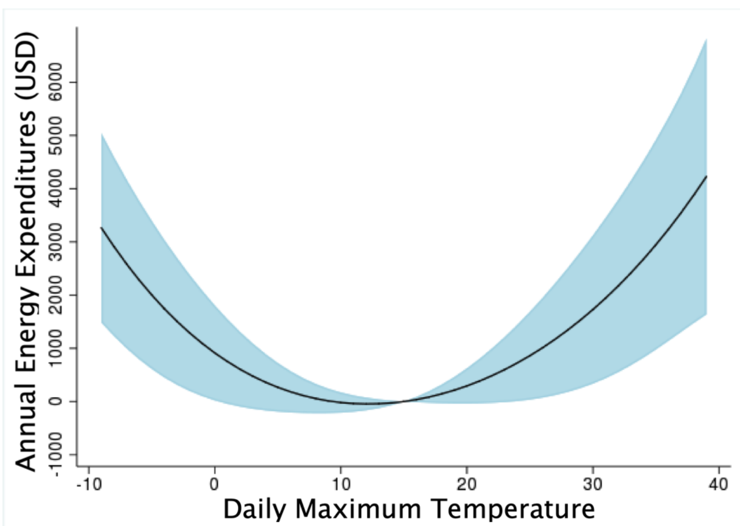

Notes: Panel (a) shows the response of annual revenue per worker to daily maximum temperature for U.S. manufacturing plants estimated using the panel regression specification in Equation 3 with a polynomial of degree four. Panel (b) shows the same specification with plant-level energy expenditures as 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

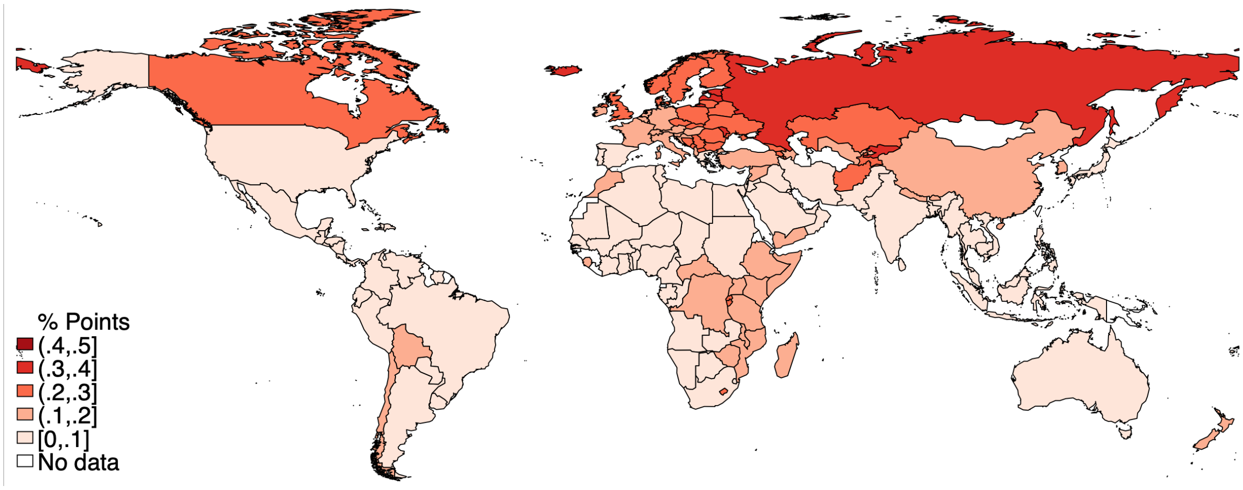

(b) -5°C Day

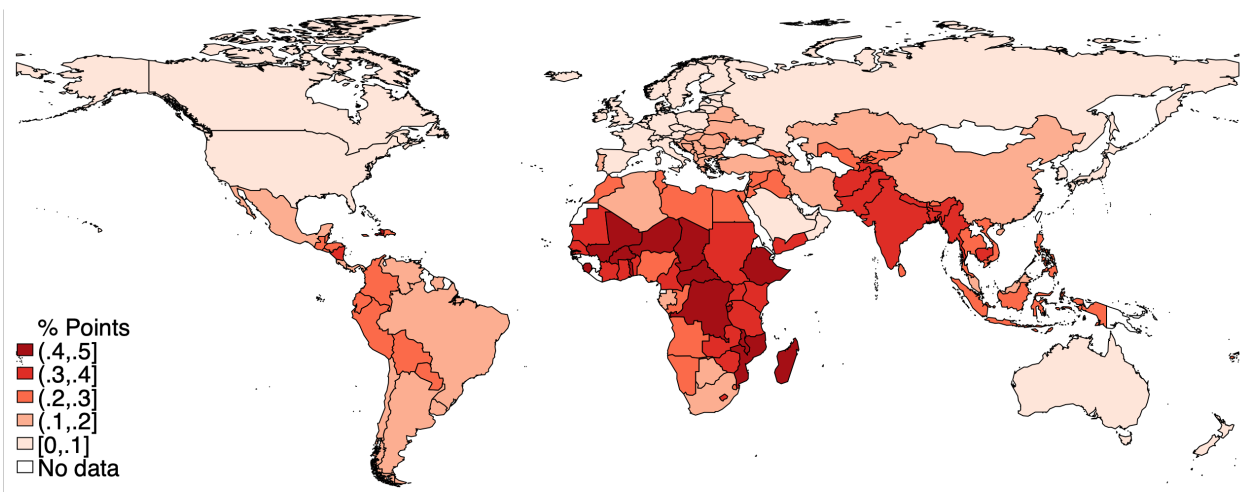

Notes: Maps show predicted annual percentage point loss in revenue per worker from a 40°C day and -5°C day obtained by evaluating the interaction regression in Column 3 of Table 1 at each country's GDP per capita and long-run average temperature. These estimates come from estimating the panel regression specification in Equation 4, 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. Appendix Table A-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.

Figure 5: Projected Impact of Climate Change on Productivity

(a) Manufacturing

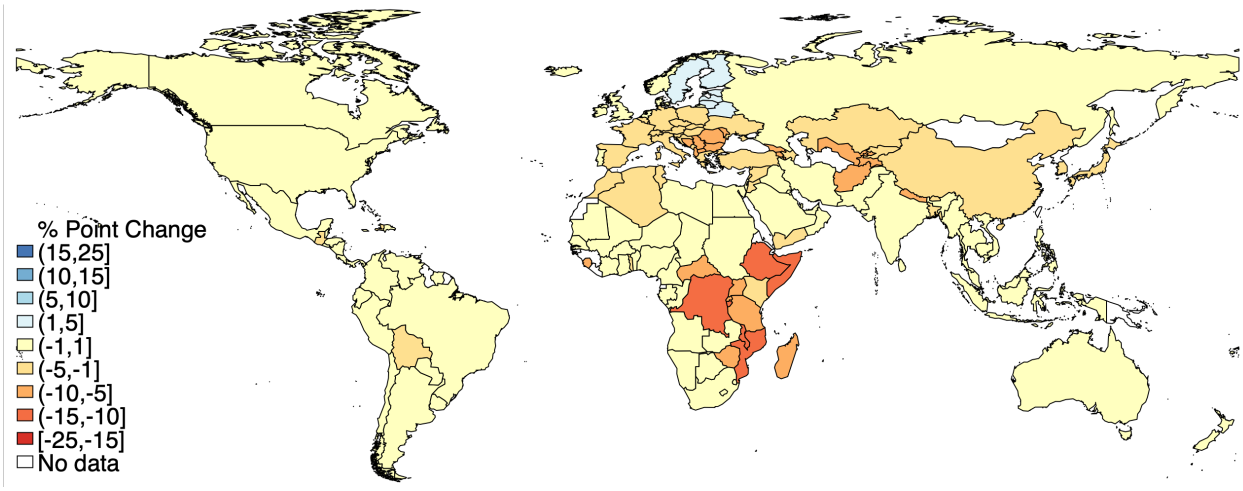

(b) Agriculture Relative to Manufacturing

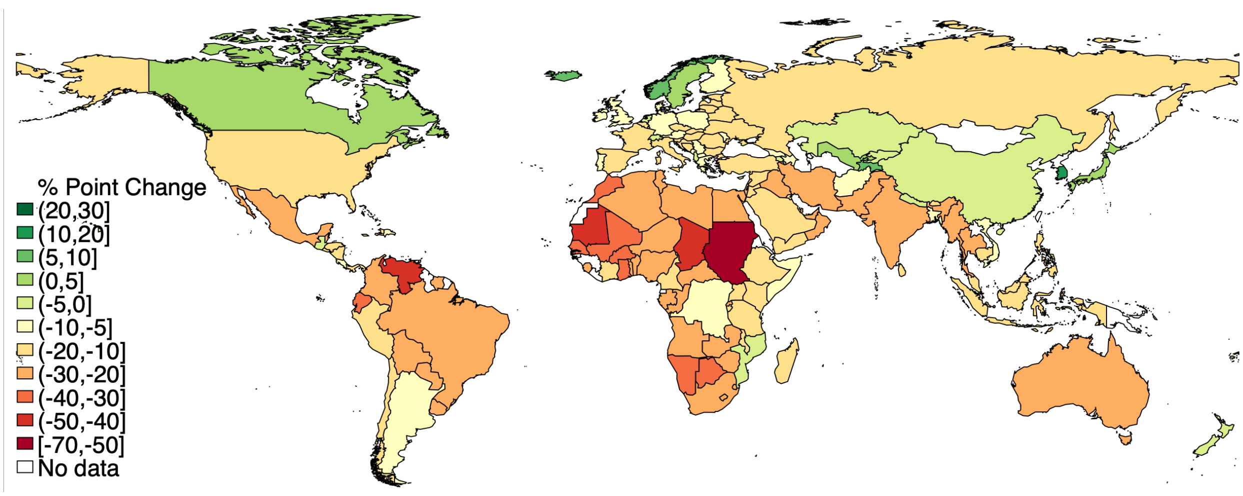

Notes: Panel (a) shows the projected impact of climate change on manufacturing productivity by end-of-century obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 3 of Table 1 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.4. 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.

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

Notes: t-statistics in parentheses. Reported Driscoll and Kraay (1998) standard errors are robust to heteroskedasticity, spatial correlation, and autocorrelation of up to 5 lags. Results come from estimating Equation 5 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. Countries in the regression are weighted by their share of the global agricultural labor force. Appendix Table A-2 contains results for a specification that interacts GDD and KDD with country level income and average temperature. 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.

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 3. 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 6: Model Fit Summary

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

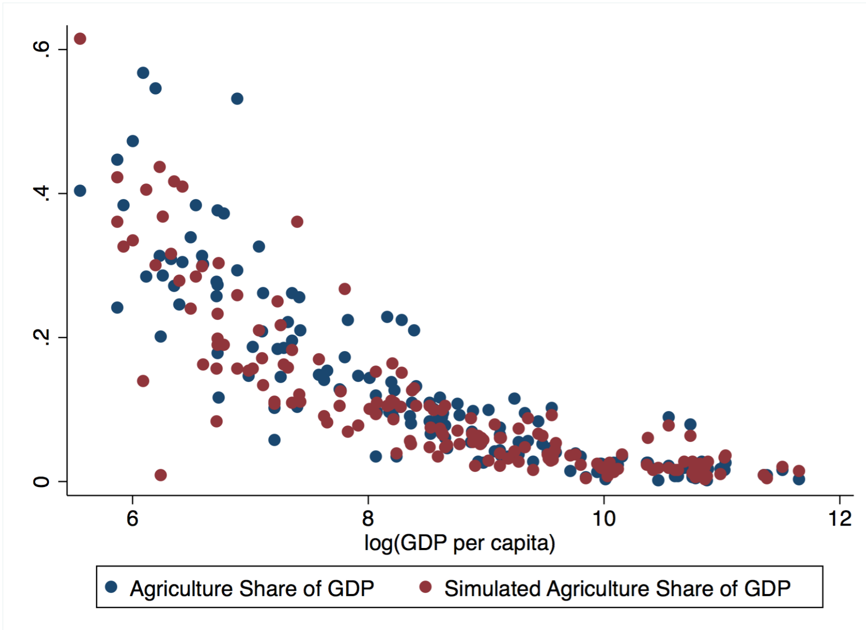

(b) Nontargeted Moment: Relative Price of Food - Data vs. Simulation

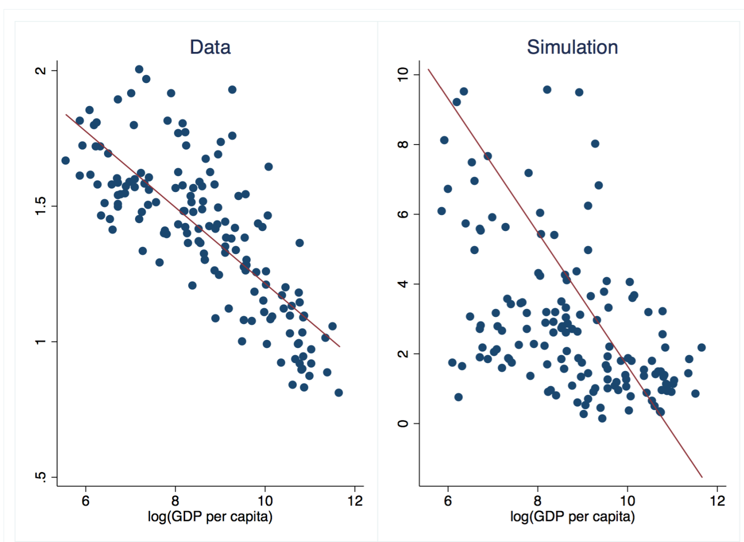

Notes: Panel (a) shows the model's fit to a targeted moment: the agriculture share of GDP across countries. The simulation explains over 60% of the variation in the data, and over 96% of the variation in average agricultural GDP share by decile of the global income distribution. Panel (b) shows the model's 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.

Table 4: Impact of Climate Change With and Without Reallocation

(a) Impact of Climate Change on Domestic Production Share of Agricultural Consumption

Country	Baseline	Full Reallocation
World	.742	.74
Poorest Quartile	.894	.876
2nd Quartile of Income	.71	.726
3rd Quartile of Income	.636	.575
Richest Quartile	.609	.616
Coldest Quartile	.648	.687
2nd Quartile of Temperature	.71	.643
3rd Quartile of Temperature	.876	.856
Hottest Quartile	.698	.668

(b) Impact of Climate Change on Agricultural Share of GDP

Country	No Reallocation	Autarky	Full Reallocation
World	.038	.043	.043
Poorest Quartile	.202	.242	.23
2nd Quartile of Income	.082	.095	.093
3rd Quartile of Income	.053	.066	.055
Richest Quartile	.03	.033	.034
Coldest Quartile	.051	.056	.06
2nd Quartile of Temperature	.114	.14	.126
3rd Quartile of Temperature	.163	.194	.186
Hottest Quartile	.125	.154	.136

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

Country	No Reallocation	Autarky	Full Reallocation
World	-.038	-.017	-.017
Poorest Quartile	-.189	-.067	-.066
2nd Quartile of Income	-.089	-.038	-.038
3rd Quartile of Income	-.082	-.022	-.021
Richest Quartile	-.029	-.014	-.013
Coldest Quartile	-.052	-.032	-.032
2nd Quartile of Temperature	-.135	-.05	-.049
3rd Quartile of Temperature	-.149	-.049	-.048
Hottest Quartile	-.142	-.04	-.039

Notes: Table shows model simulations of the impact of the productivity effects of climate change on the domestic production share of agricultural consumption, the agricultural share of GDP, and welfare for groups 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 labeled by quartile show population-weighted outcomes for each grouping of countries in the global income and temperature distribution in 2011. Rows marked "World" show global totals as a share of GDP (equivalently a GDP-weighted average).

Figure 7: Global Welfare Impact of Climate Change Productivity Effects

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

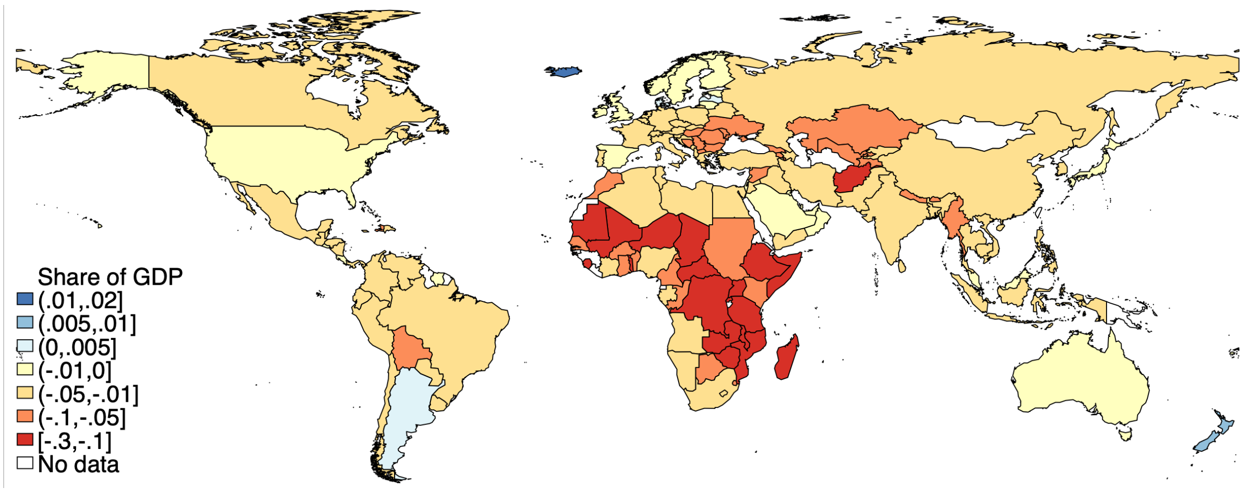

(b) Projected Impact of Climate Change on Domestic Food Price Index (Percentage Points)

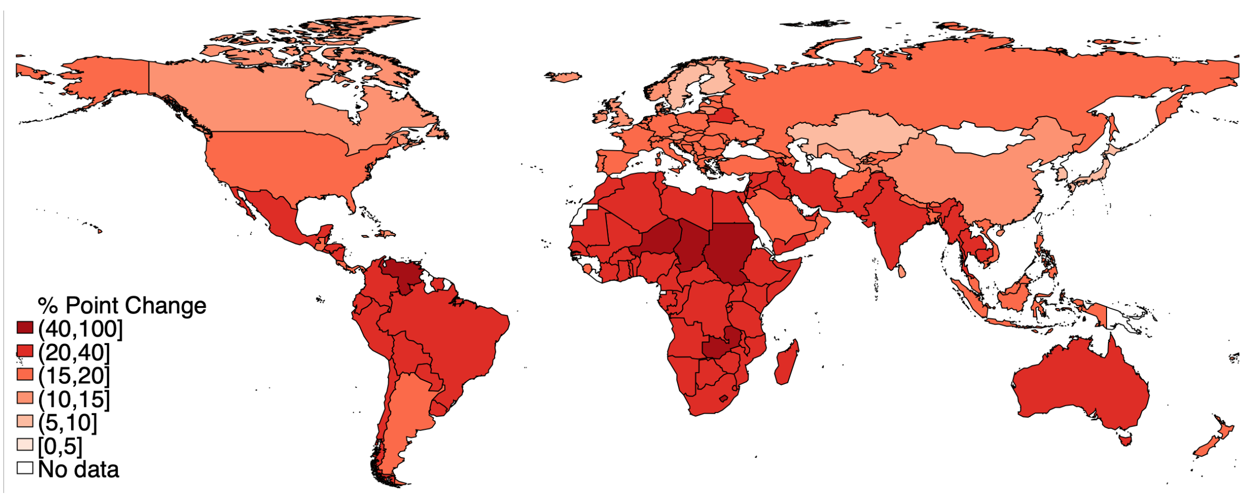

Notes: Panel (a) shows the equivalent variation welfare impacts of the productivity effects of climate change in each country in the model simulation. Panel (b) shows the corresponding impacts on the aggregate country level food price index, P_{ak}, a key driver of the welfare effects in the simulation.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

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

Country	Δ Ag Domestic Production Share	Δ Ag Share of GDP	Equivalent Variation	Δ Food Prices
Baseline - Estimated Trade Costs				
World	-.002	.005	-.017	.218
Poorest Quartile	-.018	.028	-.066	.273
Richest Quartile	.007	.004	-.013	.181
Coldest Quartile	.039	.009	-.032	.148
Hottest Quartile	-.03	.011	-.039	.27
Reduce All Trade Costs to $\tau = 1$ (Frictionless Trade)				
World	-.005	.002	-.013	.199
Poorest Quartile	-.001	-.002	-.027	.199
Richest Quartile	-.004	.002	-.011	.199
Coldest Quartile	-.001	.002	-.03	.199
Hottest Quartile	0	-.002	-.01	.199
Reduce All Trade Costs to $\tau = 2$ (90th Percentile Openness)				
World	.014	.003	-.015	.191
Poorest Quartile	-.015	-.003	-.038	.197
Richest Quartile	.011	.003	-.012	.192
Coldest Quartile	.026	.004	-.032	.186
Hottest Quartile	-.009	-.005	-.017	.196
Universal Free Trade Agreements & No Tariffs				
World	.001	.003	-.015	.114
Poorest Quartile	-.15	-.019	-.064	.206
Richest Quartile	0	.003	-.012	.113
Coldest Quartile	.036	.007	-.031	.102
Hottest Quartile	-.063	-.011	-.032	.174
Universal Free Trade Agreements, No Tariffs, & Reduced Regulatory Import Frictions				
World	.003	.003	-.014	.113
Poorest Quartile	-.148	-.024	-.057	.186
Richest Quartile	.003	.004	-.011	.112
Coldest Quartile	.032	.006	-.03	.108
Hottest Quartile	-.047	-.017	-.023	.143

Notes: Table shows model simulations of the impact of the productivity effects of climate change on the domestic production share of agricultural consumption, the agricultural share of GDP, welfare, and food prices for groupings of countries based on global quartiles of the world's income and temperature distributions in 2011. The top panel shows results with trade costs calibrated to match existing trade flows as in Section 4. The next two panels show results in two benchmark scenarios that replace trade costs for all bilateral pairs in both agriculture and manufacturing with either frictionless trade ($\tau = 1$, second panel) or with a uniform low level of trade costs ($\tau = 2$, third panel) equal to about the 90th percentile of openness in the calibration. The bottom two panels consider specific policy counterfactuals related to trade agreements, tariffs, and regulatory frictions. See Section 5.2 for more details on implementation.

Table 6: Impacts of Climate Change - Robustness and Economic Growth Extension

Model Scenario	Δ Ag GDP Share	Equivalent Variation	Δ Food Prices
Robustness of Main Results			
Baseline Results			
World	.005	-.017	.218
Poorest Quartile	.027	-.064	.268
Stone-Geary Preferences			
World	.003	-.016	.217
Poorest Quartile	.031	-.057	.266
Hultgren et al. Agriculture Estimates			
World	.005	-.017	.203
Poorest Quartile	.021	-.043	.167
Iglesias-Rosenzweig Agriculture Estimates			
World	.005	-.017	.216
Poorest Quartile	.029	-.063	.263
Cline Agriculture Estimates			
World	.005	-.017	.223
Poorest Quartile	.028	-.088	.377
Economic Growth Scenarios			
SSP 1 - Moderate Growth, High Convergence			
World	.003	-.011	.184
Poorest Quartile	.009	-.017	.275
SSP 2 - Moderate Growth, Moderate Convergence			
World	.003	-.012	.182
Poorest Quartile	.012	-.023	.27
SSP 3 - Low Growth, Low Convergence			
World	.003	-.016	.178
Poorest Quartile	.019	-.047	.255
SSP 4 - Moderate Growth, Low Convergence			
World	.003	-.014	.179
Poorest Quartile	.016	-.036	.261
SSP 5 - High Growth, High Convergence			
World	.002	-.008	.185
Poorest Quartile	.007	-.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. Bottom panel shows how results vary with assumptions about baseline economic growth. For more details on each of these, see Section 5.3.

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.

For defining the covariates used in Equation 4, I calculate long-term temperature as the 40-year average in the county containing firm i, and use country-level income since reliable comprehensive data on subnational income is not available.

Table A-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.

Appendix A.2: Additional Regression Results

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

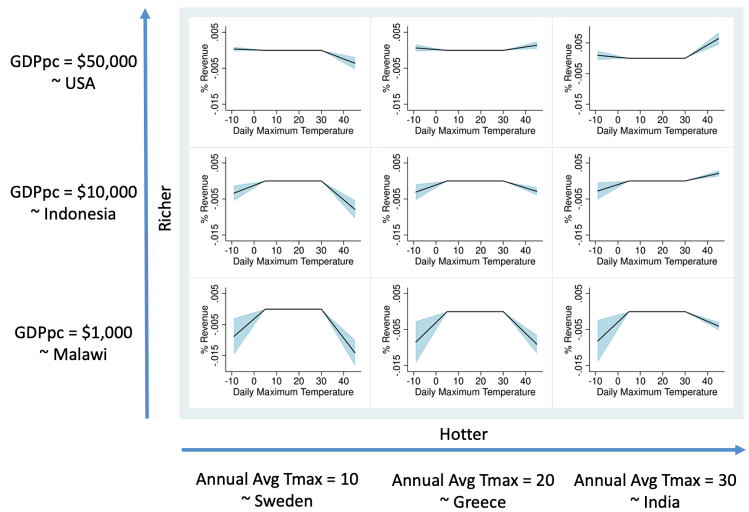

Notes: Figure shows the predicted effect of temperature on the log of manufacturing revenues at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 4 of Table 1. 95% confidence intervals are shown in blue. Outcome variables come from data sources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

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

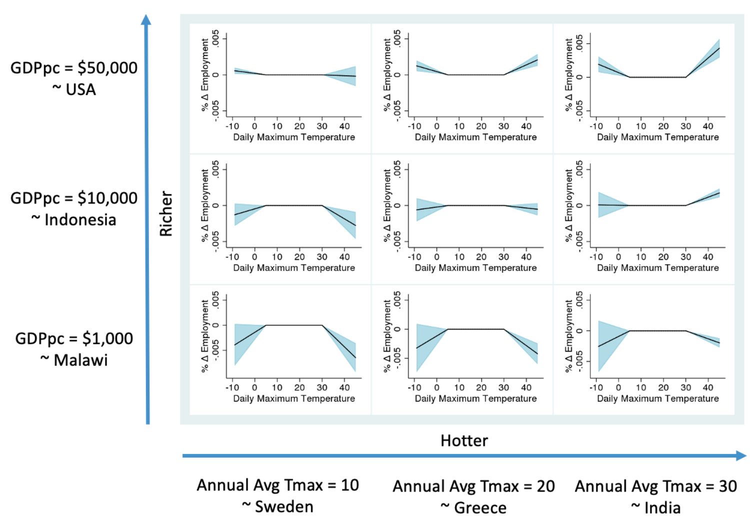

Notes: Figure shows the predicted effect of temperature on the log of manufacturing employment at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 5 of Table 1. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-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

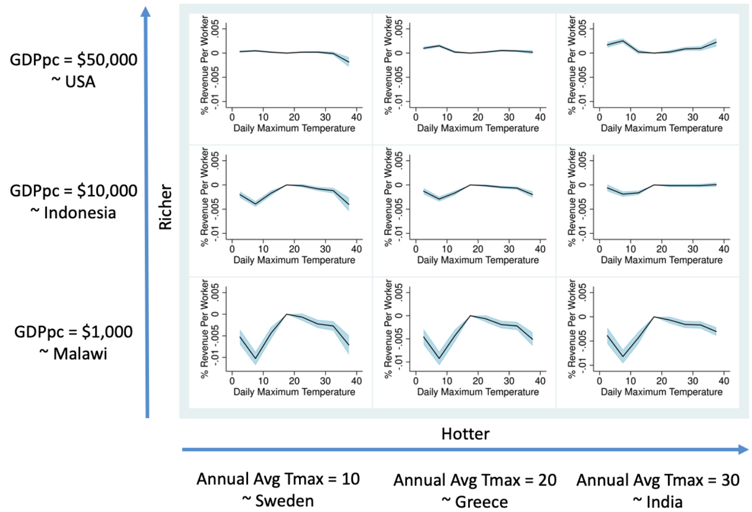

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using bins of daily maximum temperature in the specification from Equation 4. Days are grouped into 5°C bins. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-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

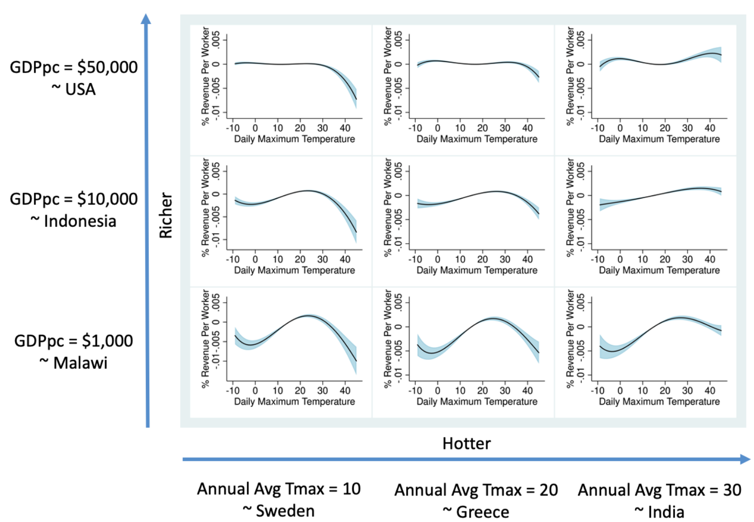

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using a polynomial of degree four in daily average temperature in the specification from Equation 4. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-5: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Polynomial of Degree 4 of Daily Maximum Temperature
Additional Interaction of GDP per capita and Long-Run Temperature

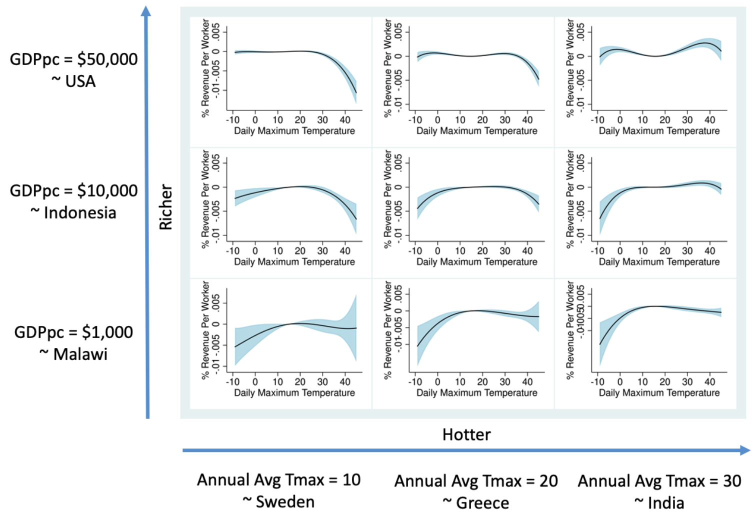

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using a polynomial of degree four in daily average temperature. Estimates are from the specification from Equation 4, with an additional term interacting the polynomial with both GDP per capita and long-run average temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

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

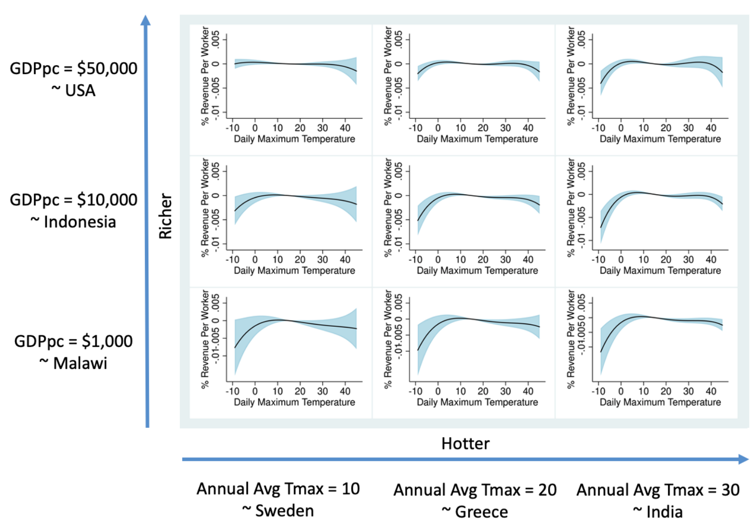

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-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

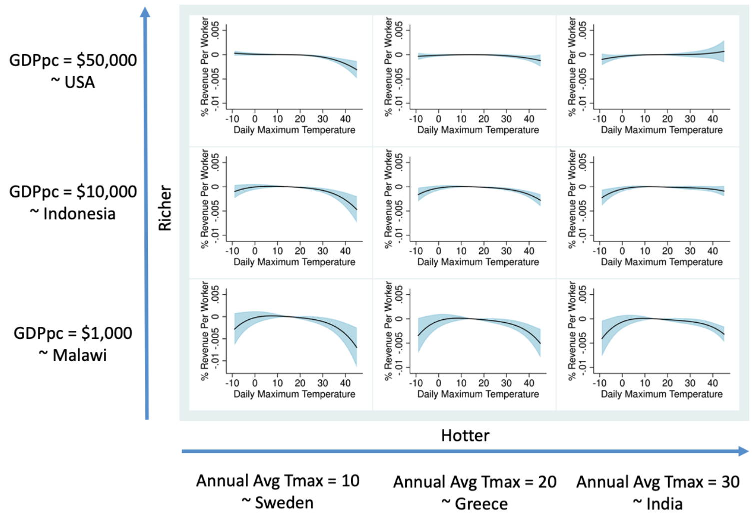

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 4 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-8: Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE

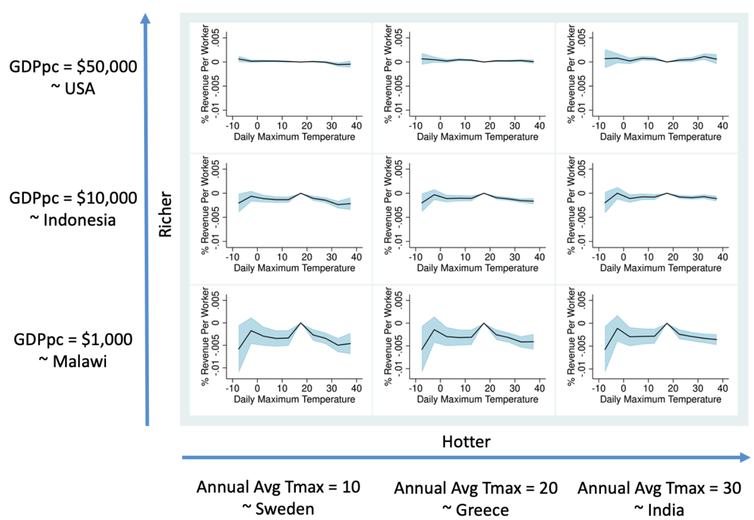

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 4 with state-by-year fixed effects and bins of daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-9: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - Controls for Capital

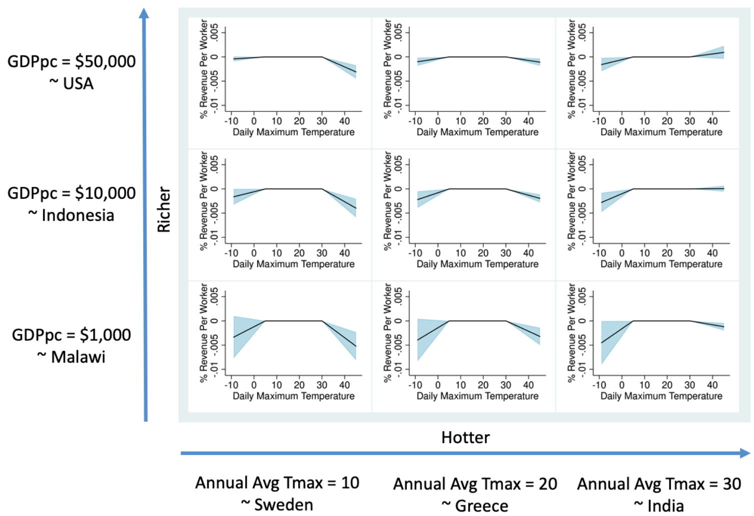

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4 with controls for capital. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-10: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - Controls for Precipitation and Lags of Temperature

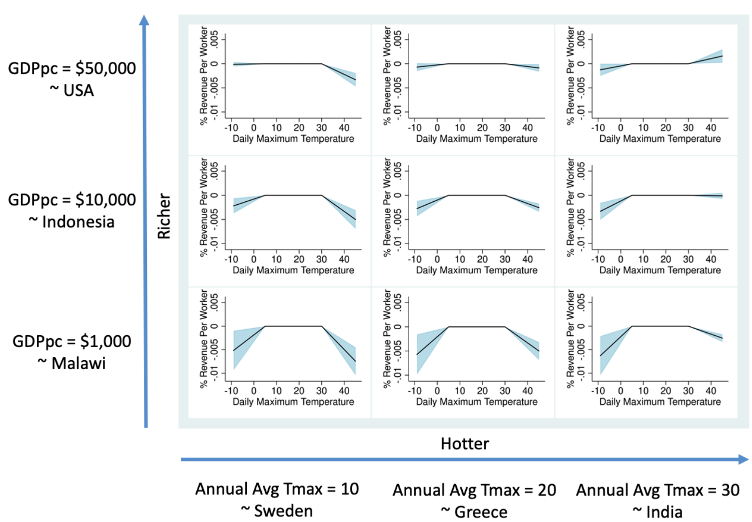

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4 with controls for a second degree polynomial in precipitation, and four lags of both CDD and HDD. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-11: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Polynomial of Degree 4 of Daily Maximum Temperature

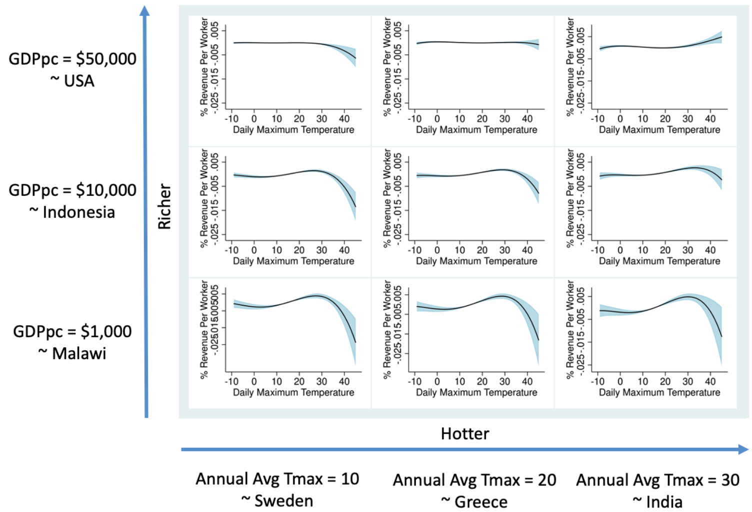

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4 with a polynomial of degree four in daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-12: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Bins of Daily Maximum Temperature

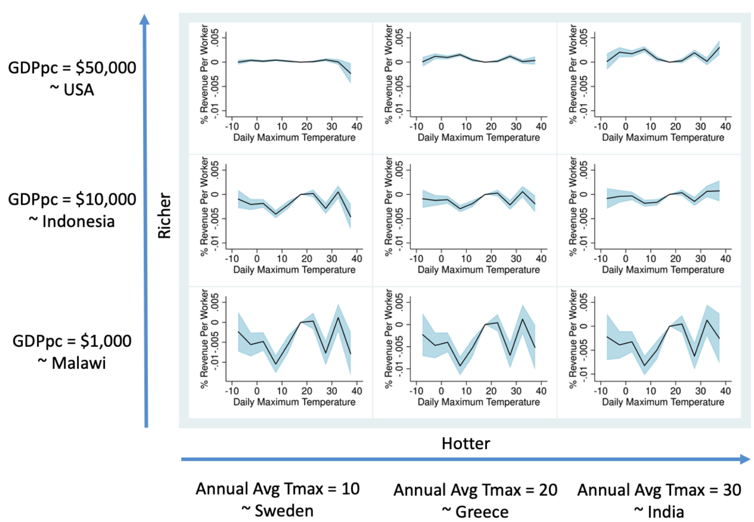

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4 with bins of daily maximum temperature. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Figure A-13: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature

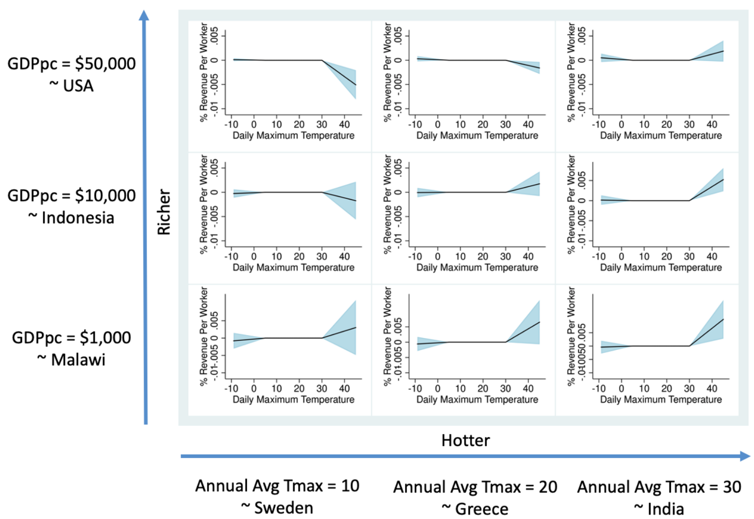

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 4. 95% confidence intervals are shown in blue. Outcome variables come from data sources listed in Appendix Table A-1 and temperature data is from GMFD.

Table A-2: Country Level Panel Regression

	(1) ln(GDP)	(2) Food Share	(3) Ag Share	(4) Ag Labor Share
KDD X 100	-0.367	-0.0157	0.0969	0.00442
	(-1.97)	(-1.23)	(1.88)	(0.51)
GDD X 100	-0.326	0.00945	-0.0622	0.000342
	(-4.21)	(2.10)	(-4.57)	(0.06)
GDD X 100 × ln($\overline{GDPpc}$)	0.0333	-0.00121	0.00570	0.000302
	(4.67)	(-2.82)	(5.06)	(0.62)
GDD X 100 × Long-Run $\overline{T}$	0.00208	0.00000724	0.000844	-0.000306
	(1.49)	(0.07)	(2.40)	(-1.75)
KDD X 100 × ln($\overline{GDPpc}$)	0.0674	-0.00292	-0.0125	-0.0000394
	(2.91)	(-2.09)	(-2.60)	(-0.06)
KDD X 100 × Long-Run $\overline{T}$	-0.00750	0.00169	0.000307	0.000177
	(-1.88)	(4.46)	(0.36)	(0.83)
Observations	7176	5514	5236	3586

Notes: t-statistics in parentheses. This regression is similar to the one contained in Table 2b, with the addition of interaction terms for the sample average of per-capita income and temperature in each country. 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 5 with crop-area weighted growing degree days (GDD) and killing degree days (KDD), and their interactions with average incomes and temperatures. 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-14: Estimated Response of U.S. Annual Manufacturing TFPR to Daily Maximum Temperature

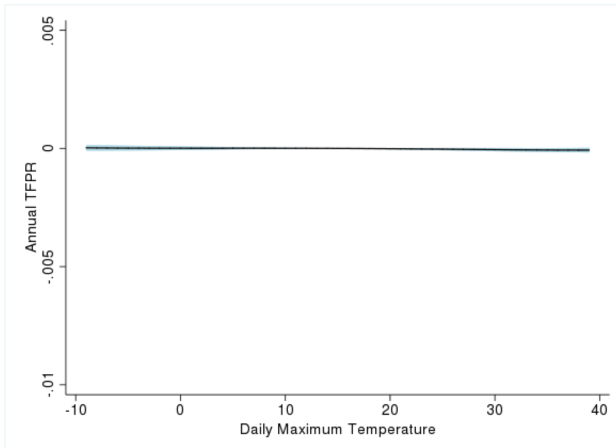

Notes: Figure shows the estimated effect of temperature on manufacturing TFPR using the specification from Equation 3 with a polynomial of degree four in daily maximum temperature. 95% confidence interval is shown in blue. Outcome data comes from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD.

Table A-3: 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 3 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-4: U.S. Energy Results

	ln(Energy Expenditure)	ln(Energy Expenditures)	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 3 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 2.1 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-15. 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-16. 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-17. This estimate using my full sample of Chinese data implies that extreme temperatures sharply and statistically significantly *increase* output, a finding inconsistent with the 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-15: China Replication - Overlapping Years

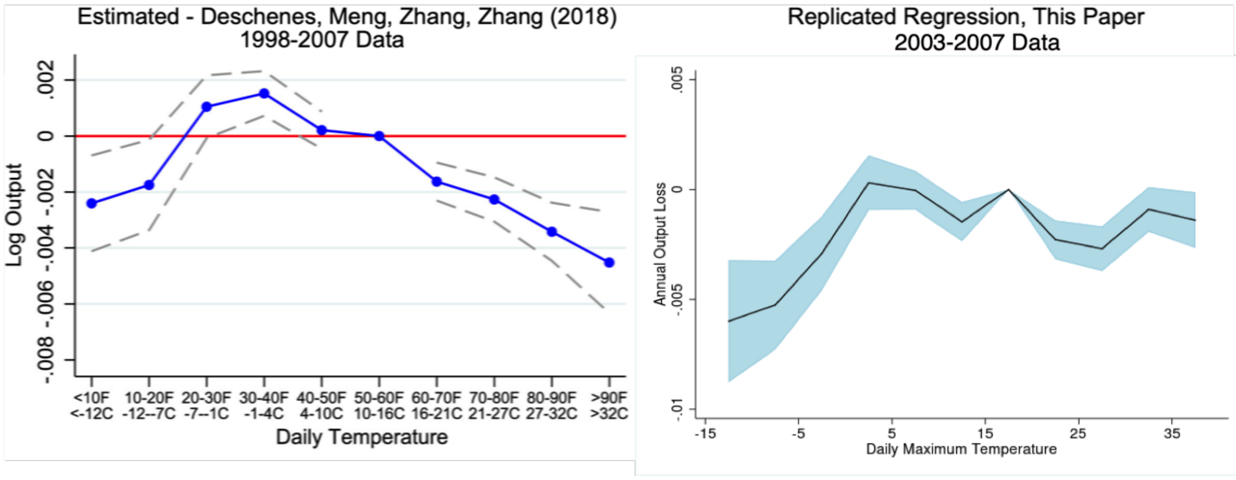

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result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

Figure A-16: China Manufacturing Temperature Sensitivity - Estimated and Predicted

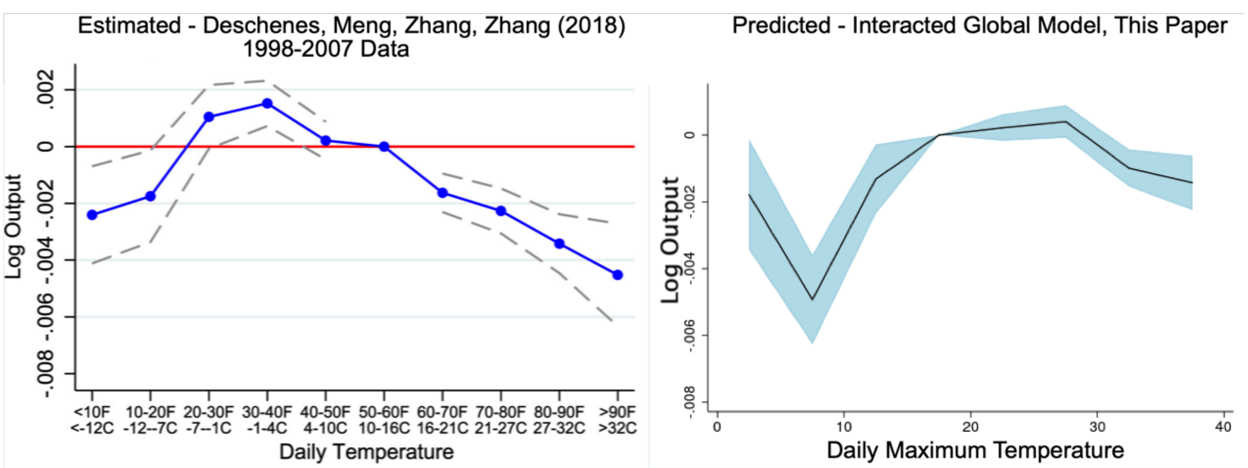

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows the predicted effect of temperature in China from evaluating my global interacted specification from Column 4 of Table 1 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-17: China Replication - Different Years

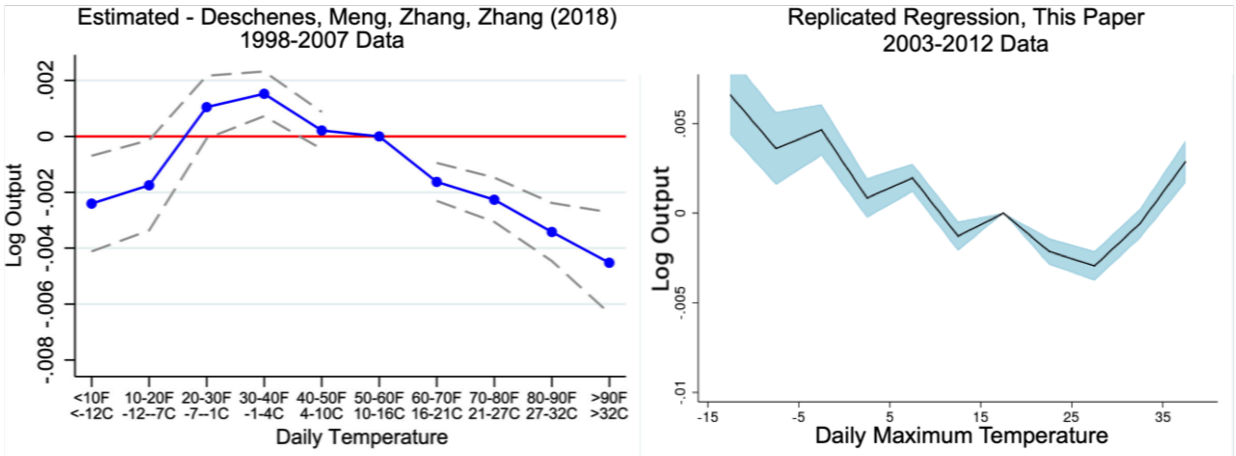

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result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

$$TC = \sum_{t=2020}^{2099} MC_t = \sum_{t=2020}^{2099} CDD_t * (\beta_{t+1} - \beta_t) \tag{25}$$

Note that the continuous version of Equation 25 also follows straight from the firm's first-order condition. 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 25 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, which 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:

$$TB = CDD_{2100} * (\beta_{2100} - \beta_{2020}) \tag{26}$$

Because CDDs are increasing as countries become hotter, the benefits of adaptation in Equation 26 exceed the costs in Equation 25. Figure A-18 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-19 show the costs of achieving this reduced sensitivity, as calculated using Equation 25, and Figure A-20 show the net benefits of firms adapting to changes in expected exposure to extreme heat.

Figure A-18: Predicted Effect of a 40°C Day on
Annual Manufacturing Revenue per Worker
At 2080 Average Temperatures

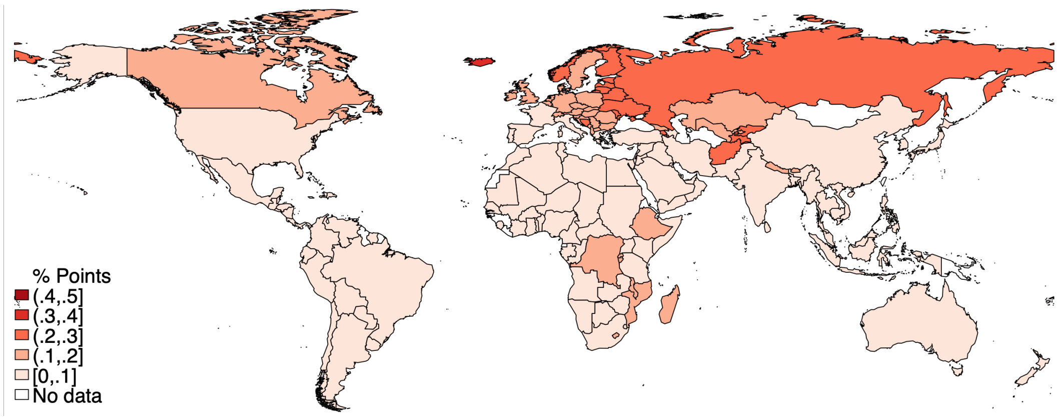

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression in Column 3 of Table 1 at each country's level of income and end-of-century long-run average temperature. Temperature sensitivities are lower in this figure than in Figure 4a because the results predict that firms will adapt to hot temperatures as the world warms.

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

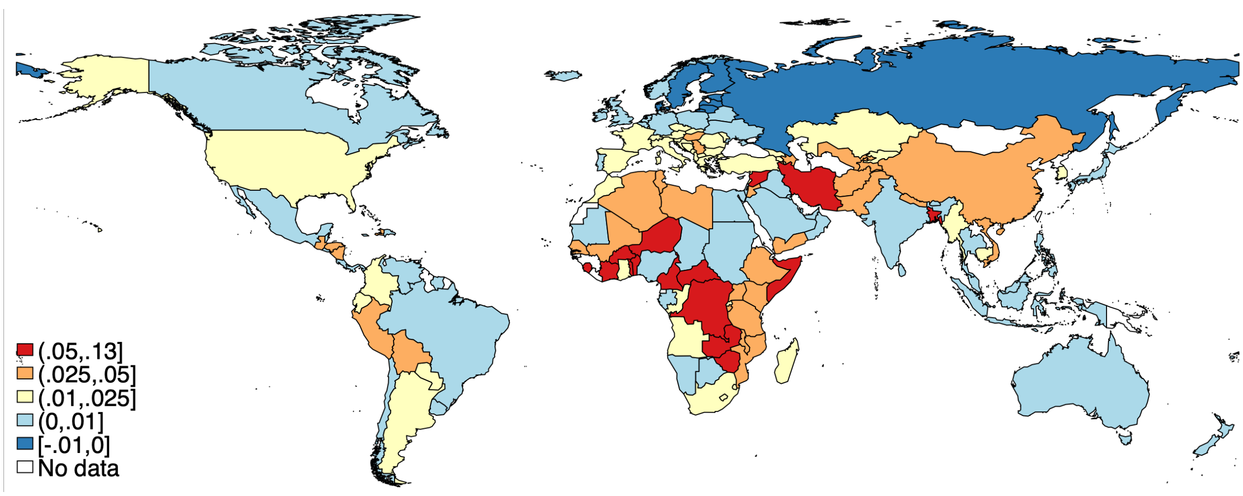

Notes: Map shows the calculations of the costs firms pay to achieve the lower temperature sensitivity shown in Appendix Figure A-18 compared to Figure 4a. I infer these costs using a revealed preference approach developed by Carleton et al. (2020) that infers adaptation costs from the foregone benefits firms would have attained by reducing their heat sensitivity. The procedure is detailed in Appendix D.

Figure A-20: Firm-Level Adaptation Net Benefits
(Share of Manufacturing Output)

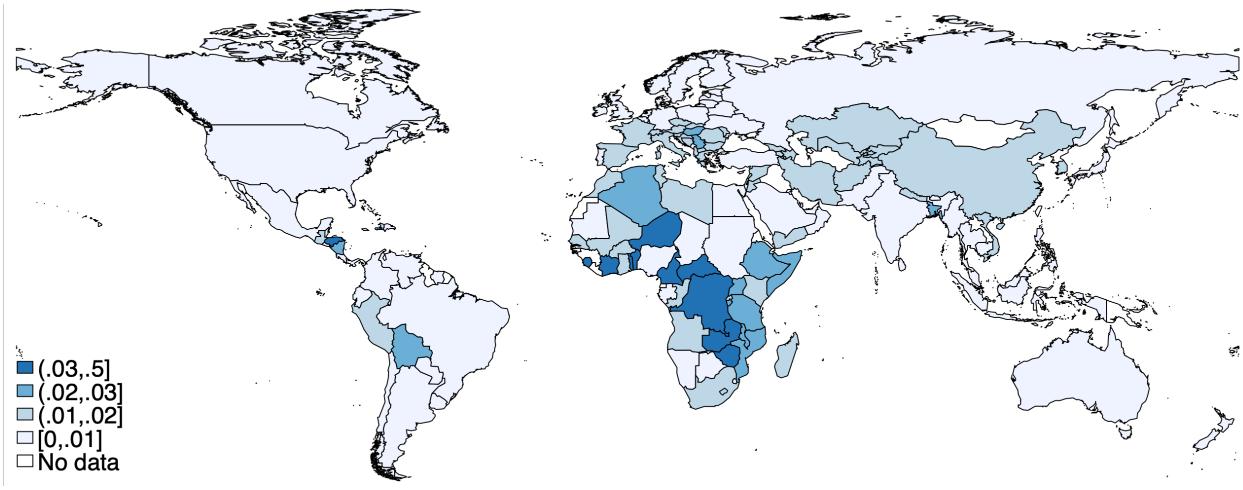

Notes: Map shows the calculations of the net benefits firms receive by investing to reduce their heat sensitivity as the climate warms. The benefits come from reducing heat sensitivity to the level shown in Appendix Figure A-18 compared to the original level in Figure 4a. The inferred costs are shown in Appendix Figure A-19. The procedure to calculate these costs and benefits is detailed in Appendix D.

Figure A-21: Predicted Effect of a 40°C Day on Annual Services
Revenue per Worker

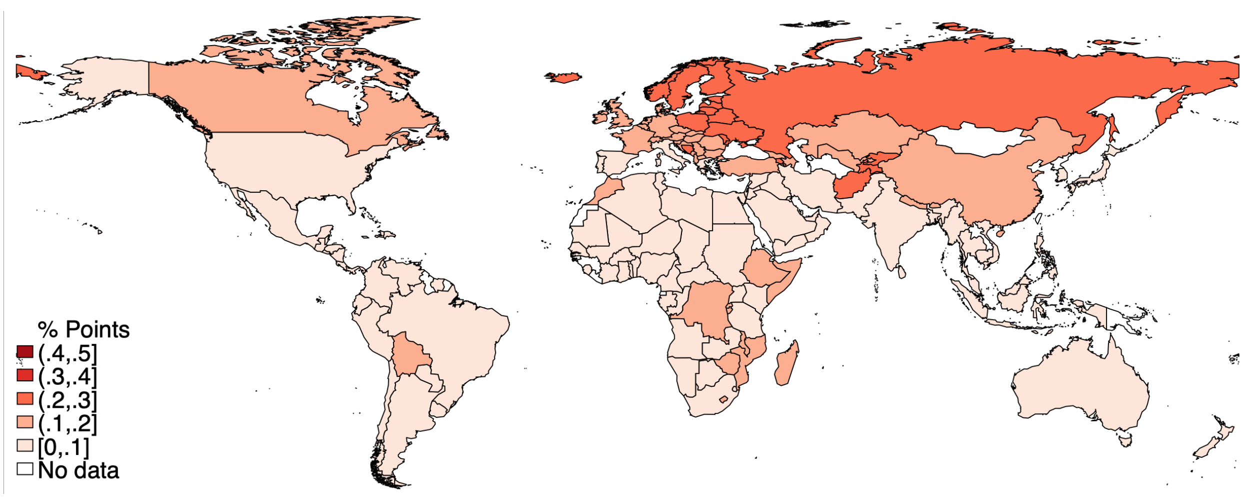

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 6 of Table 1 at each country's level of income and long-run average temperature.

Figure A-22: Predicted Effect of a -5°C Day on Annual Services Revenue per Worker

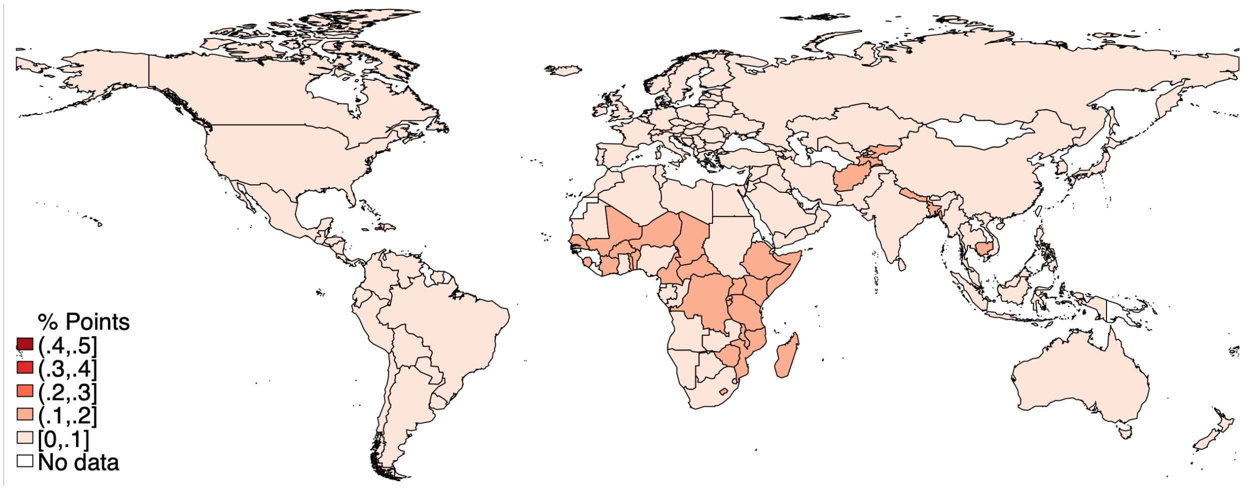

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a -5°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 1 at each country's level of income and long-run average temperature.

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

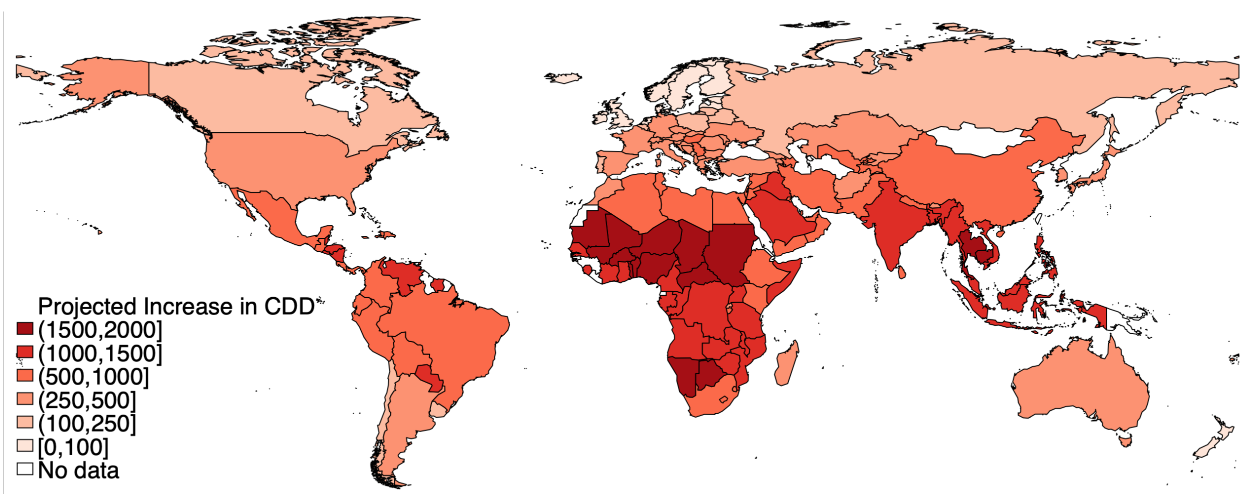

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme heat as measured by the change in average cooling degree days above 30°C from the first 20 years to the last 20 years of the century.

Figure A-24: Projected Change in Exposure to Extreme Cold

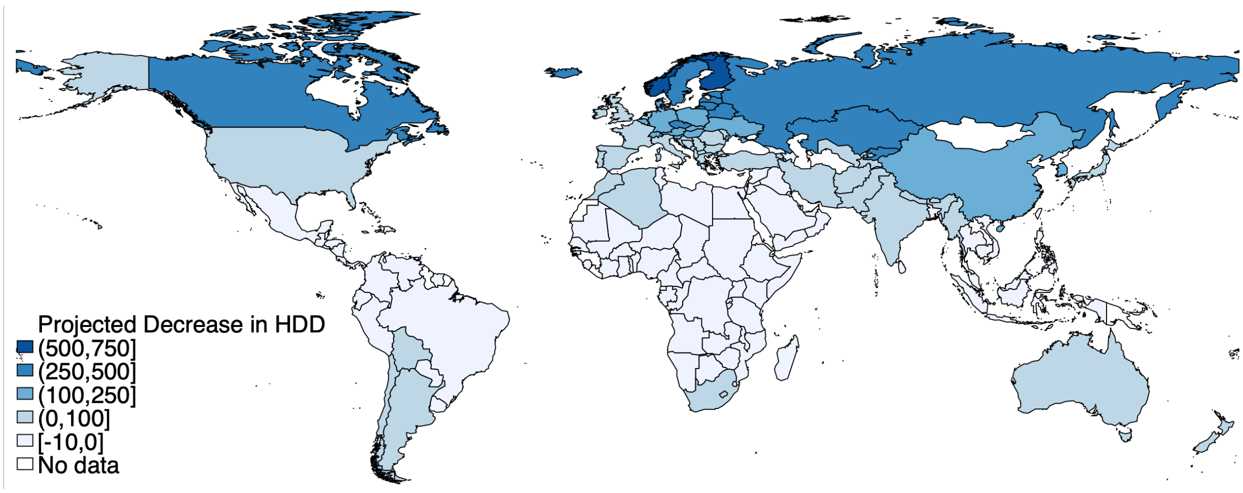

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of future exposure to extreme cold as measured by the change in average heating degree days below 5°C from the first 20 years to the last 20 years of the century.

Figure A-25: Projected Impact of Climate Change on Services Productivity

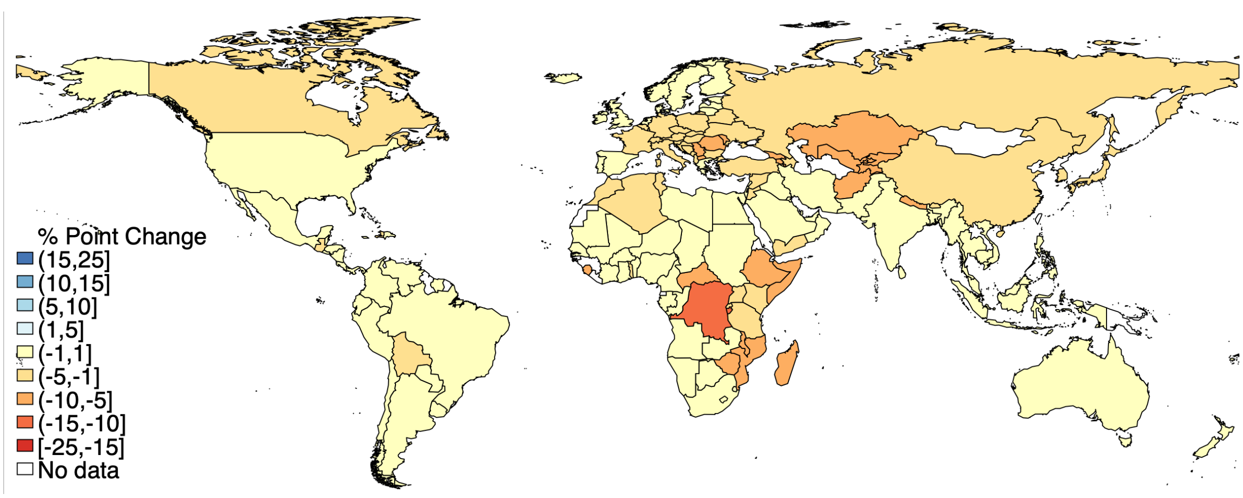

Notes: Map shows the projected impact of climate change on services productivity by end-of-century obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 6 of Table 1 at each country's income and end-of-century long-run average temperature.

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

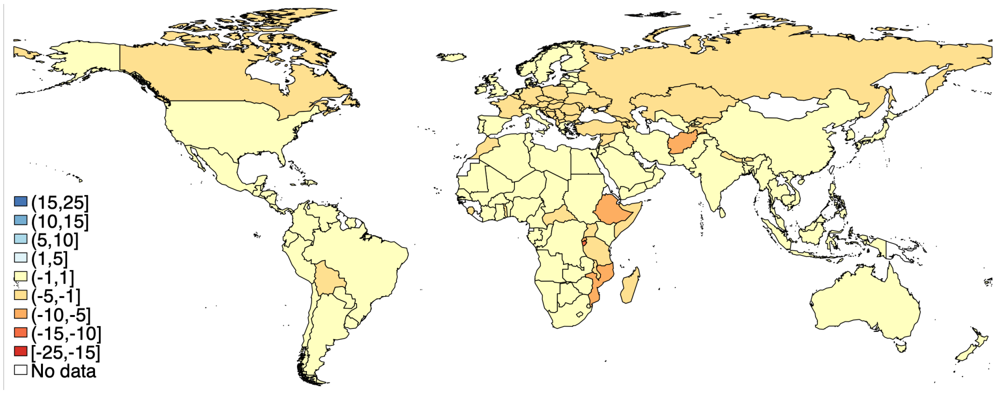

Notes: Map shows the projected impact of climate change on manufacturing productivity by end-of-century obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 3 of Table 1 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017) in Shared Socioeconomic Pathway 3. These estimates that account for economic growth show reduced losses relative to those in Figure 5a because the empirical results suggest that firms in richer countries have reduced exposure to extreme temperatures.

Figure A-27: Projected Impact of Climate Change on Services Productivity Accounting for Economic Growth and Adaptation

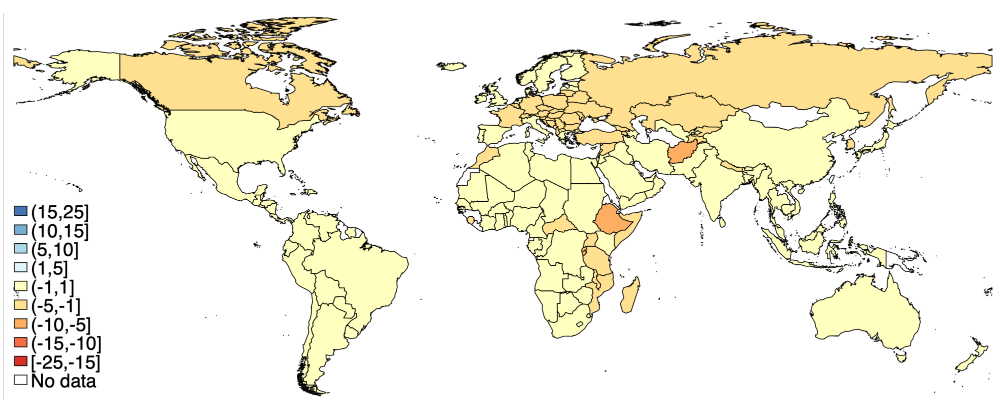

Notes: Map shows the projected impact of climate change on services productivity by end-of-century obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 6 of Table 1 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017) in Shared Socioeconomic Pathway 3. These estimates that account for economic growth show reduced losses relative to those in Appendix Figure A-25 because the empirical estimates suggest that firms in richer countries have reduced exposure to extreme temperatures.

Appendix E: Model Details

Appendix E.1: Proofs of Proposition and Corollaries

Proof of Proposition 1

Proof of the condition in Equation 22:

I start by showing that the condition in Equation 22 holds. Equation 20 provides the expression for the labor share in agriculture as a function of expenditure shares, wages, and country k's share of production in all countries:

$$l_{ak} = \pi_{akk}\omega_{jk} + \sum_{n \neq k} \pi_{akn}\omega_{jn}\frac{w_n L_n}{w_k L_k}$$

For simplicity, treat domestic labor as the numeraire so the domestic nominal wage, w_k, is normalized to 1. By the assumption that country k is "small," nominal wages in other countries, $w_n \ \forall \ n \neq k$, are unaffected by productivity changes in k. Thus, I treat the term $\frac{w_n L_n}{w_k L_k}$ as a constant since population in each country, L_n, is also held fixed in the model. Thus, totally differentiating the labor share equation yields the following:

$$dl_{ak} = \pi_{akk}d\omega_{ak} + d\pi_{akk}\omega_{ak} + \sum_{n \neq k} d\pi_{akn}\omega_{an}\frac{w_n L_n}{w_k L_k} + \sum_{n \neq k} \pi_{akn}d\omega_{an}\frac{w_n L_n}{w_k L_k}$$

The assumption that country k is "small," also gives us that sectoral price indices, P_{jn}, and the aggregate price index, P_n, in other countries do not respond to changes in productivity in country k. This is because country k accounts for a "small" share of consumption in other countries, n, so movements in exports from k to n are assumed to have a negligible effect on the price indices in n. By Equation 8, this means that movements in the expenditure share in agriculture, ω_{an}, can also be treated as "small" in countries other than k. This means that $d\omega_{an}$ is approximately 0 and the last term drops out. This gives us that:

$$dl_{ak} > 0 \qquad \text{if} \qquad \pi_{akk}d\omega_{ak} + d\pi_{akk}\omega_{ak} + \sum_{n \neq k} d\pi_{akn}\omega_{an}\frac{w_n L_n}{w_k L_k} > 0$$

which is the condition stated in Equation 22 of Proposition 1.

Proof that $d\pi_{akk}\omega_{ak} < 0$ when $dZ_{ak} < 0$:

Equation 17 provides an expression for π_{akk} as a function of Z_{ak}:

$$\pi_{akk} = \frac{Z_{ak}(w_k)^{-\theta_a}}{\sum_{m=1}^{N} Z_{am}(\tau_{amk} w_m)^{-\theta_a}}$$

Note that $\tau_{akk} = 1$ for domestic consumption, so that variable drops out from Equation 17. When considering the derivative $\frac{d\pi_{akk}}{dZ_{ak}}$, the small country assumption gives us that $\frac{dw_m}{dZ_{ak}} = 0$ for all m, and w_k has been normalized to 1. In addition, Proposition 1 considers a case where $dZ_{ak} < 0$ but productivity in other countries is not moving, so $\frac{dZ_{am}}{dZ_{ak}} = 0$. Thus, all the terms in this equation are constant except for Z_{ak}. Since the wages, productivities, and τ are all positive, this is sufficient to imply that $d\pi_{akn} < 0$ when $dZ_{ak} < 0$.

Proof that $\sum_{n \neq k} d\pi_{akn} \omega_{an} \frac{w_n L_n}{w_k L_k} < 0$ *when* $dZ_{ak} < 0$:
Once again, Equation 17 gives us the expression for π_{akn} as a function of Z_{ak} for all countries n that could import food from k:

$$\pi_{akn} = \frac{Z_{ak}(\tau_{akn} w_k)^{-\theta_a}}{\sum_{m=1}^{N} Z_{am}(\tau_{amk} w_m)^{-\theta_a}}$$

The equation is the same as the one above for π_{akk} with the exception of the import tariff, τ_{akn}, which is held constant for the change in Z_{ak}. Thus, again we have that all the terms in the equation for π_{jkn} are held constant for changes in Z_{ak}, which means that $dZ_{ak} < 0$ implies $d\pi_{akn} < 0$ since the productivities, wages, and trade costs are all positive values. Given that $\frac{d\omega_{an}}{dZ_{ak}} = 0$ and $\frac{dW_n}{dZ_a} = 0 \;\; \forall \;\; n \neq k$ follow from the small country assumption and that w_k is normalized to 1, this is sufficient for the third term in Equation 22 to be negative when $dZ_{ak} < 0$.

Proof that $\pi_{akk} d\omega_{ak} > 0$ *when* $dZ_{ak} < 0$:
The proposition states that $\pi_{akk} d\omega_{ak} > 0$ holds with the parameter conditions $\sigma < 1$, $\epsilon_a < 1 - \sigma$, $\epsilon_a < \epsilon_m$, and $\epsilon_a < \epsilon_s$. To see this, start with the equation for the expenditure share, ω_{ak}, expressed in logs:

$$\ln(\omega_{ak}) = \ln(\Omega_a) + (1-\sigma)\ln\left(\frac{P_{ak}}{P_k}\right) + (\epsilon_a - (1-\sigma))\ln\left(\frac{w_k}{P_k}\right)$$

The first term is a constant and is unaffected by $dZ_{ak} < 0$. Thus, for ω_{ak} to increase, it is sufficient to show that the second and third terms both increase when $dZ_{ak} < 0$, given the stated parameter conditions.

To start with, Equation 15 provides an expression for P_{ak} as a function of Z_{ak}. Given that we have normalized w_k to 1 and the small country assumption gives us that $\frac{dw_n}{dz_{ak}} = 0$ for all $n \neq k$, we have that $\frac{dP_{ak}}{dZ_{ak}} = (-1/\theta) \times Z_{jn}^{-1/\theta-1} \times$ constant. Thus, we have $\frac{dP_{ak}}{dZ_{ak}} < 0$ and P_{ak} rises with the fall in Z_{ak}. Equation 15 also shows that P_{mk} and P_{sk} are unaffected by the change in Z_{ak} since each sector's price index depends only on productivities and tariffs in that sector, and on nominal wages, which are unaffected by assumption. Given this, Equation 10 shows that utility, U_k, must fall to offset the rise in P_{ak}, given that nominal expenditures in country k are fixed at 1 and that P_{mk} and P_{sk} do not change. From the definition $P_k = \frac{w_k L_k}{U_k}$ and the assumptions that w_k and L_k do not change, this means that P_k must rise to offset the fall in U_k. Since w_k is assumed to be fixed, this means that the real wage $\frac{w_k}{P_k}$ falls when Z_{ak} falls. Given the requirement that $\epsilon_a < 1 - \sigma$, the final term is then the product of two negatives when $dZ_{ak} < 0$, and must be positive. So the decrease in agricultural productivity decreases the real wage, which raises the expenditure share in agriculture through the non-homotheticity channel.

Next, we must prove that the change in the second term above also raises the expenditure share. Recall from above that P_{mk} and P_{sk} remain fixed since Equation 15 shows that they depend only on Z_{mk}, τ_{mnk}, and w_n, and Z_{sk}, τ_{snk}, and w_n, respectively, all of which are held fixed in the counterfactual where Z_{ak} declines. In a standard homothetic setting, the aggregate price index is a weighted average of P_{ak}, P_{mk}, and P_{sk}, so this would be sufficient on its own to prove that P_k moves by less than P_{ak} such that the ratio $\frac{P_{ak}}{P_k}$ increases. In this setting the aggregate price index, P_k, is defined in Equation 9, which does not straightforwardly show that P_k rises less than one-for-one with P_{ak}, so the proof must take an alternative tactic.[25]

Thus, it is simpler to re-express Equation 23 in terms of relative expenditure shares for agriculture and manufacturing in terms of their relative prices:

$$\ln\left(\frac{\omega_{ak}}{\omega_{mk}}\right) = \ln\left(\frac{\Omega_a}{\Omega_m}\right) + (1-\sigma)\ln\left(\frac{P_{ak}}{P_{mk}}\right) + (\epsilon_a - \epsilon_m)\ln\left(\frac{w_k}{P_k}\right) \tag{27}$$

Recall it is shown above that P_{ak} increases when $dZ_{ak} < 0$ and P_{mk} does not change. Thus, the increase in the relative price of agriculture also increases its expenditure share relative to manufacturing, as well as relative to services, for which a corresponding version of the above equation holds. Since the parameter conditions also state that ϵ_m and ϵ_s are both greater than ϵ_a, the decline in real wage in the third term caused by P_{ak} increasing also increases the relative expenditure share in agriculture. Since the expenditure shares sum to 1, an increase in the relative expenditure share of agriculture compared with both other sectors requires its absolute expenditure share to increase. Thus, the in-

[25] Using the implicit function theorem and Equation 10, we can get the following expression for $\frac{dP_k}{dP_{ak}}$:

$$\frac{dP_k}{dP_{ak}} = \frac{(1-\sigma)\Omega_a(w_k L_k)^{\epsilon}_{a} P_k^{-\epsilon_a} P_{ak}^{-\sigma}}{\epsilon_a \Omega_a (w_k L_k)^{\epsilon_a} P_k^{-\epsilon_a - 1} P_{ak}^{1-\sigma} + \epsilon_m \Omega_m (w_k L_k)^{\epsilon_m} P_k^{-\epsilon_m - 1} P_{mk}^{1-\sigma} + \epsilon_s \Omega_s (w_k L_k)^{\epsilon_s} P_k^{-\epsilon_s - 1} P_{sk}^{1-\sigma}}$$

The parameter restrictions do not provide enough information to show that this condition must be less than one, such that P_k increases by less than P_{ak}.

crease in P_{ak} and the decline in $\frac{w_k}{P_k}$ caused by the decline in Z_{ak} both cause the expenditure share in agriculture, ω_{ak} to increase.

Proof of Corollary 1

This proof follows closely from above. With frictionless trade, producers in country k have no trade cost advantage from selling in country k relative to foreign producers. Thus, the "small country" assumption applies to country k's own prices in the same way as it applied to all other countries $n \neq k$ above. This means that dZ_{ak} is assumed to have a negligible impact on P_{ak} and w_k, in addition to other sectoral prices, and prices and wages in foreign countries. Since domestic sectoral prices and wages do not change, Equation 27 implies that sectoral expenditure shares do not change. Thus, with frictionless trade, $dZ_{ak} < 0$ implies $d\omega_{ak} = 0$. In contrast, the impact of $dZ_{ak} < 0$ on the trade shares from Equation 17 continues to hold as in the above proof of Proposition 1. Thus, with frictionless trade and the small country assumption the first term in Equation 22 is zero and the second two terms remain negative, meaning that the labor share in agriculture cannot increase and the "food problem" cannot hold in this context.

Proof of Corollary 2

In a closed economy all consumption is domestically produced so $\pi_{akk} = 1$, and there are no exports so $\pi_{akn} = 0 \ \forall \ n$. Thus, the second and third terms in Equation 22 must be zero - there is no trade response to a decline in agricultural productivity. With homothetic preferences we have $\epsilon_a = \epsilon_m = \epsilon_s$, which means that expenditure shares are independent of real income and the third term in Equation 23 is zero. Finally, with $\sigma > 1$, the effect in the above proof of Proposition 1 is reversed - an increase in the relative price of agriculture compared to that of other sectoral goods will reduce the expenditure share in agriculture. Thus, with homothetic preferences and $\sigma > 1$, the expenditure share in agriculture falls when $dZ_{ak} < 0$. Equation 20 shows that agriculture's share of labor equals the expenditure share, $l_{ak} = \omega_{ak}$, in a closed economy when $\pi_{akk} = 1$ and $\pi_{akn} = 0 \ \forall \ n$, which means that the declining expenditure share guarantees that the "food problem" cannot hold with these preferences even in a closed economy.

Appendix E.2: Solution Algorithm & Simulated Method of Moments

I solve the model presented in Section 3 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} from Equations 15 and 17. 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 Equations 8 and 9. Given these objects, I can check whether Equation 18 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 18 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.[26] 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. 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.

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.

Appendix E.3: Additional Model Fit Details

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

Table A-5: Summary of Model Fit

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

Notes: Table shows the results from regressing empirical moments in the data on their simulated counterparts. Data on nominal income levels and the agriculture share of GDP are from the World Bank. Data on the domestically produced share of expenditures in agriculture is constructed using Comtrade data. A coefficient of 1 with $R^2 = 1$ would constitute a perfect fit. The fit for other moments in the model is displayed in Appendix E.2 Figures A-28, A-29, and A-30.

Figure A-28: Sectoral GDP Shares - Data vs. Simulation

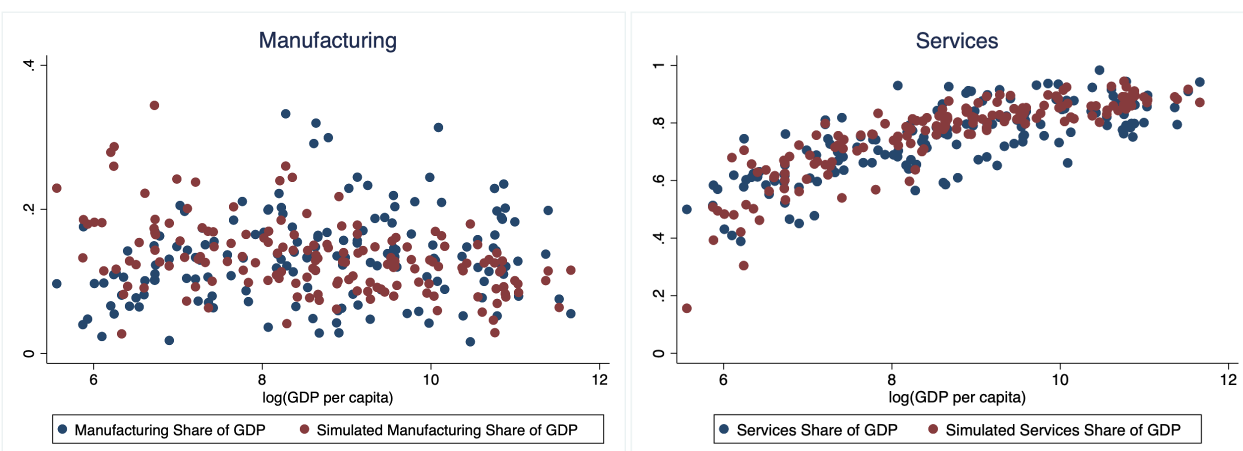

Notes: Left graph shows the fit of simulated manufacturing share of GDP in the model to data from the World Bank. Right graph shows the same comparison for services share of GDP.

Figure A-29: GDP Per Capita and Agriculture Share of GDP - Data vs. Simulation

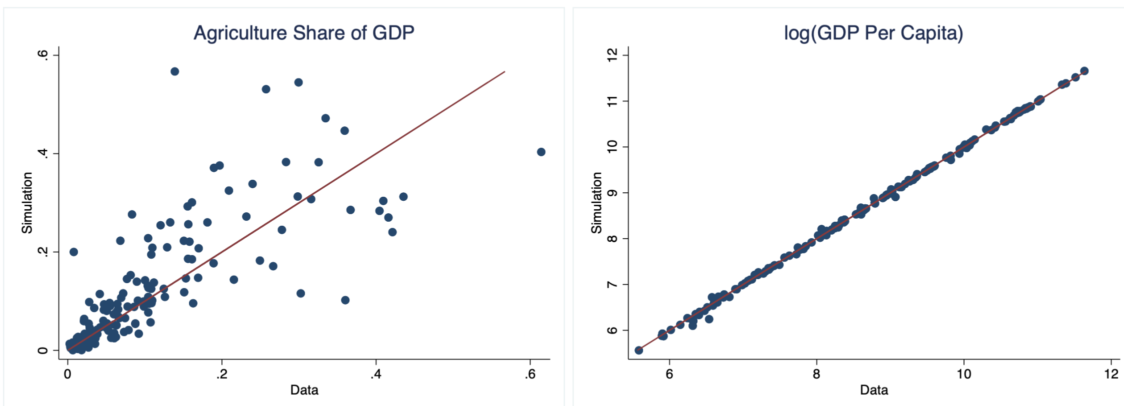

Notes: Left graph shows another view of the fit of simulated agriculture share of GDP in the model to data from the World Bank also shown in Figure 6a. Right graph shows the same comparison for GDP per capita. A perfect fit would have all data points be on the 45°line such that the simulated and actual 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-30: Domestic Production Share of Expenditures - Data vs. Simulation

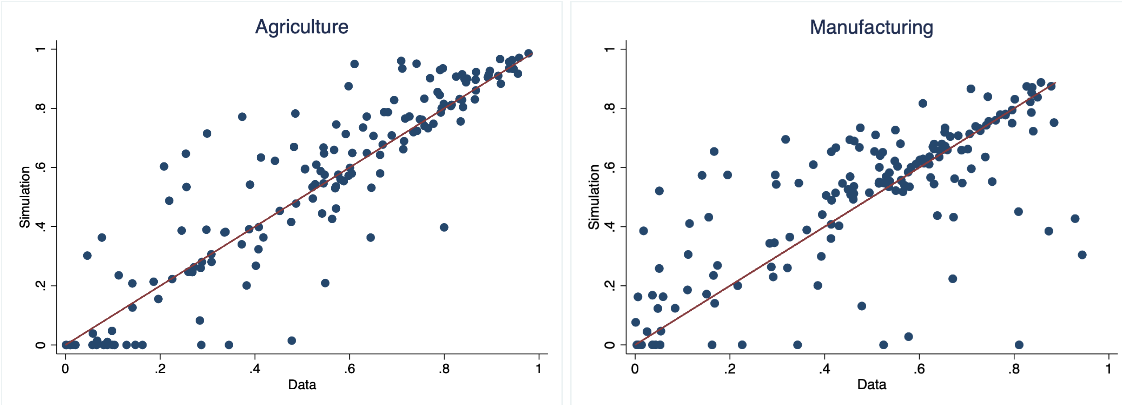

Notes: Graph shows the fit of simulated domestic production share of agricultural (left) and manufacturing (right) consumption in the model to data from Comtrade. As shown in Section 3.2, openness to food imports is a crucial parameter governing the response of labor reallocation to climate change. The simulation explains over 80% of the variation in the data for this moment.

Figure A-31: Domestic Production Share of Expenditures in Agriculture - Model Simulation

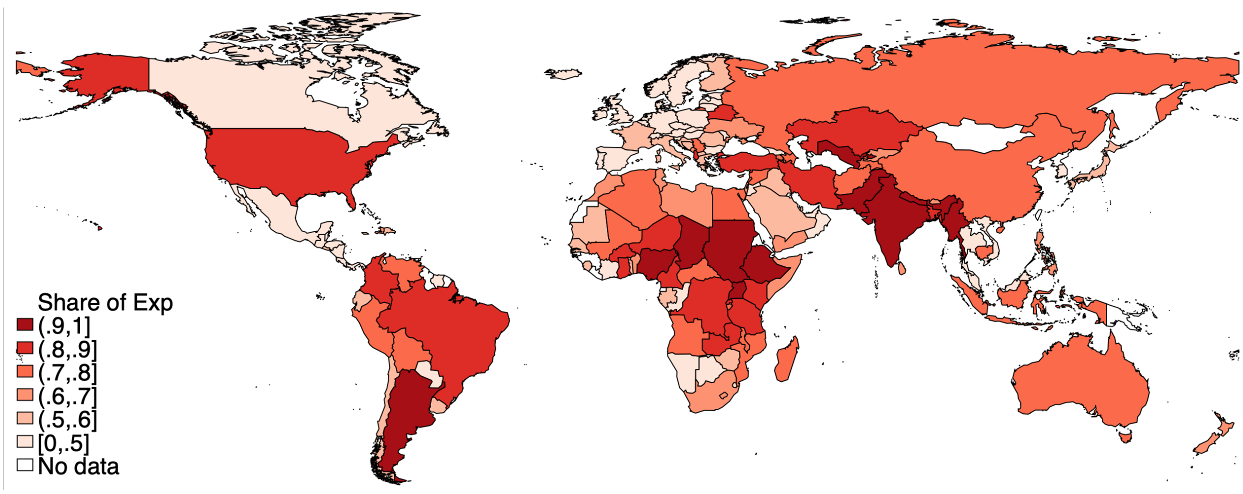

Notes: Figure shows that the share of expenditures on domestically produced goods in agriculture is very high in many developing countries with high barriers to trade. Table A-5 shows that these simulated values track closely to the data.

Appendix E.4: 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 6 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 2.1 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 2.2, 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 E.5: Decomposing Barriers to Trade

Figure A-32: Non-Tariff Barriers to Trade

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

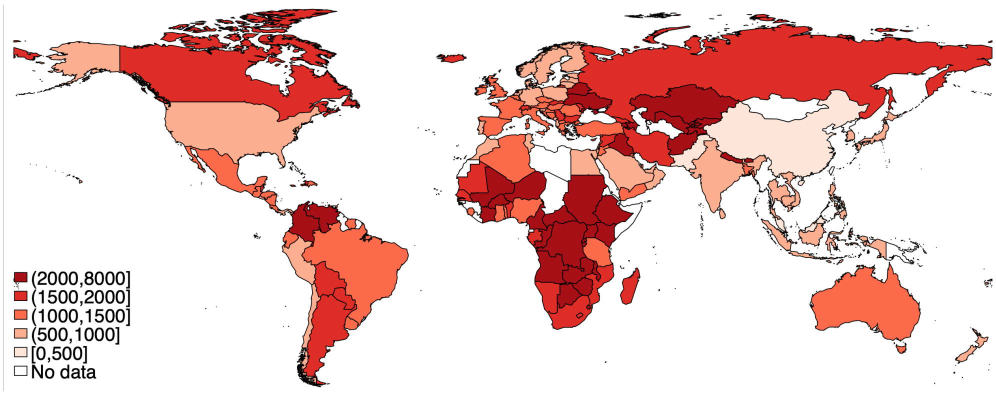

(b) Days to Import a Container

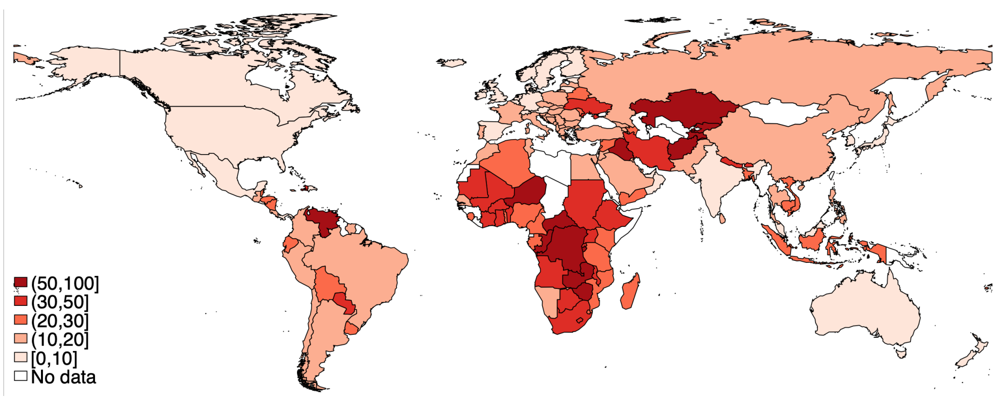

Notes: This figure shows possible underlying causes of the high barriers to trade calibrated in the model to match the low levels of trade flows in developing countries. Panel (a) shows the direct cost to import one container of goods in US dollars. Costs include documents, administrative fees for customs clearance, terminal handling charges, and inland transport, but not tariffs, taxes, or unofficial payments. 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.

Table A-6: Trade Cost Decomposition

	Log Ag Trade Cost		Log Manufacturing Trade Cost	
	(1)	(2)	(3)	(4)
Trade Agreement	-0.359		-0.618	
	(-14.33)		(-23.81)	
Mean Tariff - Primary Products	0.0158			
	(12.06)			
Mean Tariff - Manufactured Products			0.00167	
			(1.01)	
Cost (Fees) Per Container US$	0.000111		-0.0000774	
	(10.29)		(-6.70)	
Days to Import	0.0108		-0.00210	
	(12.46)		(-2.25)	
Contiguous	-0.666	-0.715	-0.456	-0.675
	(-11.00)	(-11.97)	(-7.36)	(-11.50)
Distance	0.0000141	0.0000201	0.00000433	0.0000237
	(6.88)	(10.24)	(2.02)	(12.08)
Common Language	-0.235	-0.182	-0.0390	-0.0608
	(-8.97)	(-6.92)	(-1.42)	(-2.32)
Common Colonizer Post-1945	0.0740	0.154	0.0495	0.0921
	(2.48)	(5.17)	(1.58)	(3.09)
Colonial Relationship Post-1945	-0.399	-0.493	-0.467	-0.483
	(-4.09)	(-4.91)	(-4.75)	(-4.95)
Constant	1.148	1.624	1.644	1.243
	(47.67)	(94.13)	(66.73)	(72.22)
Observations	21219	23585	21047	23376

Notes: t-statistics in parentheses. Dependent variable in each regression is the log of bilateral trade costs ($\ln \tau_{jkn} - 1$) for the agriculture and manufacturing sectors from the model in Section 3 using the calibration strategy described in Section 4. For the independent variables, data on tariffs, import fees, and days to import come from the World Bank. Import fees cover customs clearance, document processing, customs brokerage, and terminal handling, but are exclusive of tariffs. Import delays are associated with customs clearance, inspection procedures, and document preparation. Data on participation in trade agreements, contiguity, distance, common language, and colonial relationships come from the CEPII Gravity Database (Conte, Cotterlaz and Mayer, 2022). The physical distance variable uses the mean distance between countries, where locations within a country are weighted using nighttime lights (Hinz, 2017). See Section 5.2 for further details on the regression specification and related policy counterfactuals.

Appendix F: Model Robustness

In this section, I evaluate the robustness of the counterfactual model simulations presented in Section 5.1 to two sets of different assumptions - an alternative specification for non-homothetic consumer preferences 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 non-homothetic 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:[27]

$$U(C_{ak}, C_{mk}, C_{sk}) = \left(\omega_a^{\frac{1}{\sigma}}(C_{ak} - \overline{C}_{ak})^{\frac{\sigma-1}{\sigma}} + \omega_m^{\frac{1}{\sigma}}(C_{mk} - \overline{C}_{mk})^{\frac{\sigma-1}{\sigma}} + \omega_s^{\frac{1}{\sigma}}(C_{sk} - \overline{C}_{sk})^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \tag{28}$$

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 non-homothetic CES specification, particularly for middle-income countries, as shown in Figure A-33.

Table 6 shows that the results in this version of the model are very similar to the baseline specification. For the poorest quartile of the global population, climate change increases agriculture's share of the labor force by 2.8 percentage points, reduces GDP by 10.7 percentage points, reduces welfare (as captured by willingness-to-pay) by 7 percentage points, and raises food prices by 37%. These results are very similar to the results in the baseline specification.

Appendix F.2: 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,

[27]The consumption parameter estimates from applying the simulated method of moments procedure to this version of the model are $\sigma = 0.89$, $\omega_a = 0.020$, $\omega_m = 0.141$, $\omega_s = 0.839$, $\overline{C}_a = 75.5$. I set $\overline{C}_m$ and $\overline{C}_s$ to zero such that there is no subsistence consumption requirement in manufacturing or services.

Figure A-33: Agriculture Share of GDP - Data vs. Simulation
Stone-Geary Specification

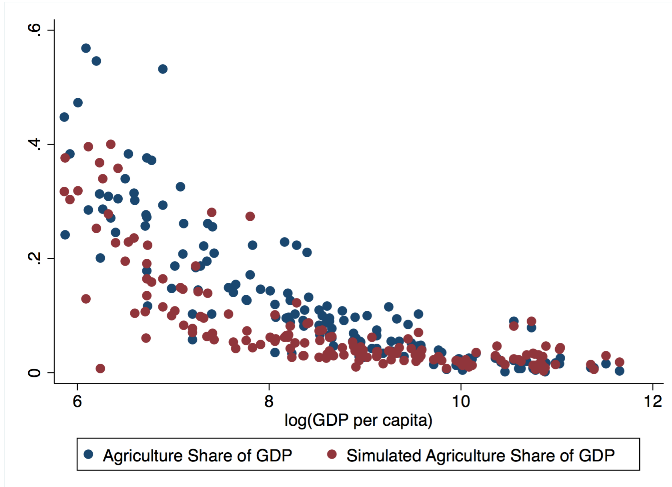

Notes: Graph shows the fit of simulated agriculture share of GDP to data from the World Bank with an alternative model specification using Stone-Geary preferences over sectoral consumption. The best fit with Stone-Geary preferences has an R^2 of only 0.43 and dramatically underpredicts the agriculture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

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):

$$\begin{aligned} Y_{ak} &= z_{ak} l_{Hak}^{\beta} l_{Lak}^{1-\beta} \\ Y_{mk} &= z_{mk} l_{Hmk}^{\alpha} l_{Lmk}^{1-\alpha} \\ Y_{sk} &= z_{sk} l_{Hsk}^{\alpha} l_{Lsk}^{1-\alpha} \\ \alpha &> \beta \end{aligned} \tag{29}$$

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 23 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 20 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.[28] 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.

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 to-

[28]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 j's total revenues, R_j, 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 agriculture's share of total revenues, r_a, and a 1% decline in manufacturing's share of total revenues, 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.

ward 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.