Available online at www.sciencedirect.com
ScienceDirect
European Economic Review 00 (2024) 1–24
European Economic Review

GDP and Temperature: Evidence on Cross-Country Response Heterogeneity ⋆

Kimberly A. Berg[a], Chadwick C. Curtis[b,*], Nelson C. Mark[c]

[a]Department of Economics, Farmer School of Business, Miami University
[b] Department of Economics, Robins School of Business, University of Richmond
[c]Department of Economics, University of Notre Dame; National Bureau of Economic Research, USA

Abstract

We estimate individual country real GDP per capita growth responses to country, global, and idiosyncratic temperature shocks. Negative growth responses to country and global temperature at longer horizons are found for all Group of Seven countries. Positive country (global) responses are found for approximately eight (seven) of the nine poorest countries at longer horizons. Both country and idiosyncratic temperature shocks have more negative than positive effects on growth across countries, but it is more evenly split for the global temperature shock. After controlling for average temperature, positive growth responses to global temperature shocks are more likely for countries that are poorer, have experienced slower growth, are more educated (higher high school attainment), and more open to trade.

Keywords: Climate, Temperature, Growth
JEL: E23, O13, Q54, Q56

1. Introduction

This paper studies how GDP responds to temperature change and how these responses vary across countries. We decompose a country's temperature into global (common) and idiosyncratic components, then use local projections (Jordà, 2005) to estimate the responses of real GDP per capita growth to each temperature component. The distribution of cross-country responses to country, global, and idiosyncratic temperature shocks displays heterogeneity in sign and in magnitude. The number of countries with negative growth responses to country and idiosyncratic temperature shocks exceed positive responses. The cross-country growth response to shocks to global temperature are more evenly split between positive and negative. In many cases, the responses of a given country to global and idiosyncratic temperature shocks go in opposite directions. These differences highlight the importance of including both sources of temperature fluctuations in order to isolate the broader effects of global temperature change.

⋆We thank Galina Hale, Claudia Ruiz-Ortega, Yener Altunbas, Scott Milliman, two anonymous referees, Editor Òscar Jordà, and the seminar participants at Indiana University, European Economic Association Meetings, IWH-CREQ-GW Macro Workshop, Annual Conference of Central Bank of Brazil, LACEA-LAMES conference, Virginia Association of Economists, James Madison University, Effects of Climate on the Business Cycle and the Economy conference, Asian Meetings of the Econometric Society, International Symposium on Regional Integration and Financial Development in Northeast Asia at Liaoning University, and EC^2 Conference at Aarhaus University for comments. Kimberly Berg thanks Miami University for summer research financial support. All errors are our own.

*Corresponding Author: Department of Economics, Robins School of Business, University of Richmond, 1 UR Dr., Richmond VA, USA
Email addresses: bergka@miamioh.edu (Kimberly A. Berg), ccurtis2@richmond.edu (Chadwick C. Curtis), nmark@nd.edu (Nelson C. Mark)

Much of the related literature, discussed below, employs panel regression techniques that imposes extensive homogeneity restrictions across countries. These studies either find uniformly negative effects of higher temperature on growth for all countries, or negative effects only for poor countries with inconclusive results for the rich. This paper departs from the literature along two main dimensions. First, in contrast to panel regressions, we use local projections for each individual county, which imposes few restrictions and allows for complete cross-country response heterogeneity. The second departure lies in the decomposition of country-level temperature into orthogonal global and idiosyncratic components, whereas much of the previous research regressed country i real GDP per capita growth on a measure of country i specific temperature. Our global and idiosyncratic components are similar in spirit to a permanent and transitory decomposition. The global component, reflecting global warming, is trending upward and shocks to it are permanent. The idiosyncratic component, on the other hand, is transitory by construction.

The idiosyncratic component is similar to the regressand in panel regressions with time-fixed effects. Global temperature is more systematic and less noisy than country temperature, since it is a cross-country average. If each country is a small open economy, not only would that country's own temperature matter for growth, but temperature should work indirectly through its effect on the rest-of-world, then circling back to the country in question via trade and financial linkages. These external, spill-over effects can be captured by global temperature shocks.

As climate is the distribution of weather variables, this decomposition is also useful because changes in global temperature is the shift in the mean of the temperature distribution. Global temperature, therefore, is conceptually close to climate change. We find impulse responses of rich country real GDP per capita growth to global temperature variation tend to be negative. At longer horizons, all of the Group of Seven (G-7) countries have negative responses, whereas many of the responses by poor countries tend to be positive.

To study the determinants of cross-country response variation, we regress the local projection impulse response coefficients on various country characteristics. This methodology draws on research strategies used in finance (e.g., Lustig and Richmond (2020) who regress the exchange rate's dollar-factor 'beta' on gravity variables). Note that there is no 'generated regressor' problem in this cross-sectional analysis because the estimated response coefficients are the dependent variable. We find that the country-level attributes have little explanatory power for the idiosyncratic responses. Country-level characteristics do a better job of explaining the impact from country and global temperature. In particular, at longer horizons, positive growth responses to global temperature shocks, which would appear to be anomalous, are more likely for countries that are poorer, have experienced slower growth, are more educated (higher high school attainment), and more open to trade.

We also run local projections using an alternative measure of temperature, the number of days in specific temperature bins (relatively cold to relatively hot days), which examines non-linear responses to temperature. When looking at greater movements in temperature – a country moving to an additional hot vs. cold day – we find the results qualitatively align with the specifications using average temperature. These also point to substantial heterogeneity in the direction of the responses. Many of the rich countries are harmed by having a higher frequency of hot days while many gain, including some of the poorest countries.

A central motivation for this project is to shed light on limited and conflicting conclusions in the literature regarding impact heterogeneity of temperature variation on real GDP per capita growth. Depending on the particular study, the empirical literature that employs panel regressions find either an inverse relationship between temperature and GDP for all countries, or an inverse relationship that holds only for poor countries. A path-breaking study in this literature is Dell et al. (2012), who use international data in estimation with country and time-fixed effects. An important motive for their panel regression approach was to use country fixed effects to control for omitted-variables bias that was present in an earlier generation of studies of cross-sectional regressions of time-averaged GDP on temperature.[1] Dell et al. (2012) reports that increased temperature lowers GDP per capita growth, but only for poor countries. Leta and Tol (2019) and Henseler and Schumacher (2019) report similar results for total factor productivity growth. Burke et al. (2015), on the other hand, find increased temperature to have a negative effect on GDP growth, but do not find differential impacts between rich and poor countries. In Nath et al. (2023), increases in temperature lower GDP per capita in hotter countries. Bansal and Ochoa (2011) find increasing global temperature lowers GDP growth of all

[1]The most prominent candidate for an omitted variable may be institutional quality, which is controlled for by the country fixed effect in panel regressions. Studies by Acemoglu et al. (2002), Easterly and Levine (2003), and Rodrik et al. (2004) argue institutions are main drivers of long-run growth outcomes.

countries with larger effects on low latitude countries.[2]

Informed by the extant literature, our prior beliefs were that the time-series variation would reveal a distribution of negative local projection coefficients with the far left tail populated primarily by poor, low latitude countries. It was surprising for us to estimate the direction of growth responses to be more evenly split between positive and negative and to find that many of the richer countries fall on the negative side of the distribution. Our empirical approach relaxes homogeneity restrictions which is prevalent in the existing literature. As a result, we unmask significant heterogeneity.

Some broader implications follow from this project. First, the pattern of cross-country response heterogeneity can supplement the ethical arguments presented by Stern (2008) to incentivize rich countries to invest in abatement strategies. The evidence that rich countries are directly economically damaged by warming should naturally incentive them to invest in climate mitigation.[3] Furthermore, if environmental policy is largely informed by observing past relationships – and we show that the sign of these responses are not uniform across countries – our results identify an additional reason why forming a global consensus on future abatement strategies is difficult. Our results can also inform refinements to damage function specifications in integrated assessment models (IAM) that compute welfare costs and evaluate the social cost of carbon. Since much of the empirical literature finds higher temperatures to be more economically damaging to poorer and hotter regions, regional IAMs, informed by such empirical damage estimates produce similar regional damage projections.[4] The geographical variation provided by our country-specific assessments to the knowledge base can provide more detailed specifications of IAM damage functions. However, our findings show that growth in many countries, and many poor countries, respond positively to historical global temperature change. We believe these findings are robust given the historical record but acknowledge that these historical relationships between growth and temperature might change in the future following several degrees (Celsius) of additional warming due to 'tipping points' or adaptation.[5]

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 discusses substantive ways our analysis departs from panel regressions. The local projection analysis is reported in Section 4. Section 4.3 undertakes a cross-sectional analysis to understand the country-level attributes that, in part, account for the results. Section 5 considers an alternative measure of temperature and Section 6 concludes.

2. GDP, Country Temperature and Its Two-Factor Decomposition

Our empirical analysis begins with an analysis of how real GDP per capita growth responds to temperature variation. In this section we first describe the sources and construction of our economic and temperature data. Section 2.2 then describes the decomposition of country-level temperature into global and idiosyncratic temperature components.

2.1. Data

Real GDP per capita is from the World Bank's *World Development Indicators*. These data are valued in constant 2010 United States dollars and have a maximal span from 1960-2017. The main empirical analysis uses only those 137 countries that have at least 30 consecutive years of observations.[6] In the analysis of Section 4.3, we also use the World Bank's, *World Development Indicators* to represent country characteristics.

Our temperature observations are population-weighted by year and country. The source is *Terrestrial Precipitation: 1900-2017 Gridded Monthly Time Series (V 5.01)* Matsuura and Willmott (2018). This is a monthly dataset

[2]The panel regression approach to study the economic effects of climate was introduced by Deschênes and Greenstone (2007), who estimated the effect of temperature on agricultural profits in the United States. Also, focusing on the United States, Colacito et al. (2019) reports higher summer temperatures are damaging to output growth in southern states and the negative impacts are by geography, not income. Hsiang et al. (2017), who examines growth in county-level income, similarly finds income is negatively impacted by temperature in the south and southwest, and increases in the north.

[3]In the absence of a global coordinated effort, Stern (2008) appeals to two ethical considerations to get the rich, industrialized countries to shoulder disproportionate costs of future abatement. First, industrialized countries are responsible for most of the current stock of greenhouse gasses and have gotten rich by generating those emissions. Second, poor countries are just beginning to overcome poverty through rapid growth and should not be forced to slow.

[4]DICE, FUND, and PAGE are prominent IAMs that serve as the main policy models employed by the U.S. Environmental Protection Agency. Regional IAMS have been developed by Hassler and Krusell (2012), Nordhaus and Yang (1996), Tol (2019), and Ricke et al. (2018), amongst others.

[5]See Barreca et al. (2016), Kim et al. (2022), and Gandhi et al. (2022) for examples of adaptation to weather related phenomenon.

[6]The full list of countries and the available sample time period for each country are listed in Online Appendix A.

estimated from weather station records and interpolated to a 0.5-degree by 0.5-degree latitude/longitude grid. We aggregate the monthly data to annual observations by node. We overlay the temperature data with population data in 2000 from the *Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11* (Center for International Earth Science Information Network, 2018). The data provides population counts at a 2.5 minute by 2.5 minute latitude/longitude grid. We use the population weights to obtain population-weighted temperatures by country and year, which is the standard approach in the literature (Kahn et al. 2019 and Dell et al. 2012).[7]

The temperature data is gridded and is interpolated among several ground stations. If we had consistent temperature measurement, temperature would be plausibly exogenous to any individual country's GDP. However, it has been pointed out that potential endogeneity arises if the underlying ground station temperature availability is dependent upon real GDP per capita growth (see Schultz and Mankin (2019) on the relation of civil conflict and discontinuity of weather station temperature readings). To address potential endogeneity concerns, we show in Online Appendix B that real GDP per capita growth is uncorrelated with weather station availability, thus mitigating these concerns. Going forward, we assume temperature is exogenous.

2.2. *Temperature*

We analyze how GDP growth responds to changes (shocks) to population-weighted country temperature $\tau_{j,t}$, and also how it reacts to changes to a common global component G_t, and associated idiosyncratic component $I_{j,t}$, obtained as a two-factor decomposition of country temperature. The global temperature (G_t) is the population-weighted cross-sectional average of $\tau_{j,t}$ across the N countries,

$$G_t = \sum_{j=1}^{N} \omega_j \tau_{j,t} \tag{1}$$

where ω_j is the population weight, using the year 2000 as the weight, on country j. The idiosyncratic temperature component is country temperature not explained by the global temperature and is given as the residual from a regression of country temperature ($\tau_{j,t}$) on global temperature (G_t),

$$I_{j,t} = \tau_{j,t} - \delta_j G_t - \alpha_j, \tag{2}$$

where α_j is the country intercept and δ_j is the slope coefficient (factor loading) on global temperature (G_t). By construction, the idiosyncratic component is stationary.

Figure 1. Population-Weighted Global Annual Temperature

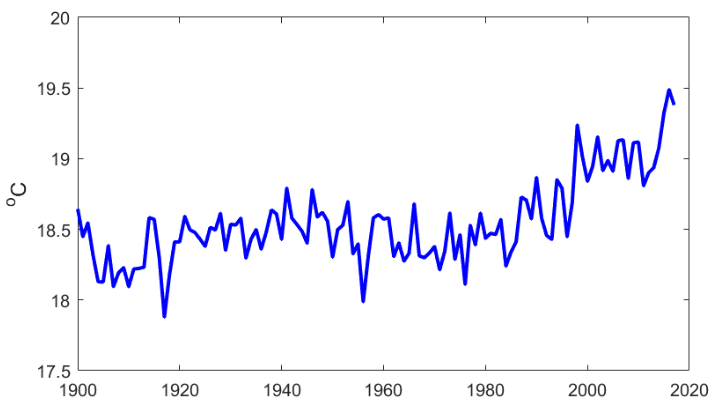

[7]We do not consider precipitation since earlier empirical work finds little or no effect of precipitation on income growth at the annual frequency.

Figure 1 plots annual global temperature from 1900-2017. It is reasonably stable from 1900 to 1980. After 1980, an upward trend is visually obvious, rising by about 1^{o}C over 40 years. We describe how trending global temperature is dealt with econometrically in Section 3.2 below.

Figure 2. Global Slope Coefficients (δ) from Equation (2)

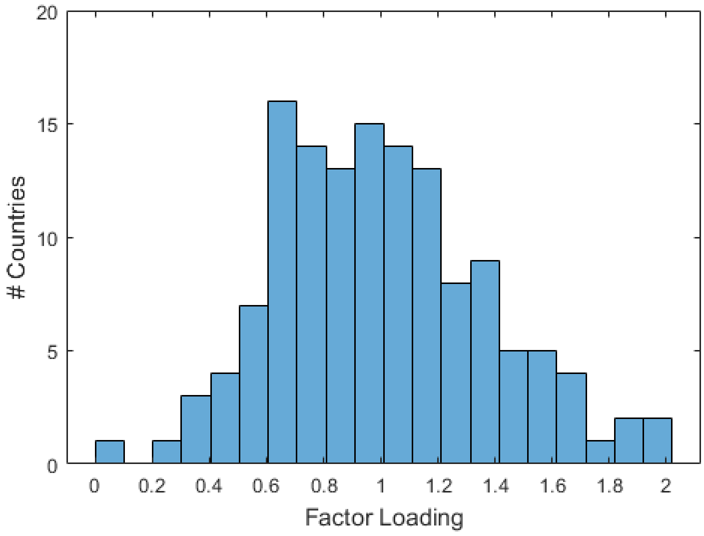

Figure 2 displays a histogram of the global temperature factor loadings (the δ_j) from the country-specific regressions in equation 2. The coefficients on all factor loadings are positive, meaning country-level temperatures all vary directly with global temperature. The distribution is centered around $\delta_j = 1$; by construction, country temperature varies one-to-one with global temperature, on average. However, the dispersion of the estimates highlights that some country's actual temperature are more (less) impacted by global temperature changes. There are four countries with factor loading coefficients below 0.4. These countries are Bolivia, Papua New Guinea, Cuba, and Bangladesh. Local temperatures in locations in close proximity to oceans have been generally less affected by global temperature change which applies to the latter three countries.

3. Departures from Panel Regression

The related literature widely adopts panel regression estimation procedures with time-fixed effects to investigate the relationship between temperature changes and real GDP per capita growth (hereafter, *growth*).[8] Two features of panel regression with time-fixed effects obscure the relationship between temperature and growth. The first is the manner in which time-fixed effects removes the global component from growth and temperature, thus resulting in a regression of coarsely constructed idiosyncratic growth on idiosyncratic temperature variation. The effects of actual country-level and global temperature variation on growth are not estimated. The second feature is the extensive homogeneity restrictions imposed on the slope coefficient of interest. While an objective of panel regression is to exploit cross-sectional variation to shrink standard errors, the imposition of extensive homogeneity should be imposed only when such restrictions are not rejected by the data. Section 3.2 shows they are rejected.

[8] Kahn et al. (2019) is an exception, who estimate panel autoregressive-distributed lag models. They also find negative GDP growth impacts of temperature, but no differences between rich and poor.

3.1. Time Fixed Effects

To illustrate the two points raised above, let $y_{j,t}$ be 100 times log real GDP per capita of country $j = 1, \ldots, N$ in time t, so that $\Delta y_{j,t}$ is the growth rate in percent. Without loss of generality, we abstract from time-invariant country-fixed effects. Consider the panel regression of growth, $\Delta y_{j,t} = y_{j,t} - y_{j,t-1}$, on the country's annual temperature, $\tau_{j,t}$, with time-fixed effects, θ_t,

$$\Delta y_{j,t} = \theta_t + \beta \tau_{j,t} + \epsilon_{j,t}. \tag{3}$$

Taking the cross-sectional average of equation (3) gives

$$\frac{1}{N}\sum_{j=1}^{N} \Delta y_{j,t} = \theta_t + \beta \frac{1}{N}\sum_{j=1}^{N} \tau_{j,t} + \frac{1}{N}\sum_{j=1}^{N} \epsilon_{j,t}. \tag{4}$$

Subtracting equation (4) from equation (3) eliminates the time-fixed effect giving,

$$\Delta y_{j,t} - \frac{1}{N}\sum_{j=1}^{N} \Delta y_{j,t} = \beta\left(\tau_{j,t} - \frac{1}{N}\sum_{j=1}^{N} \tau_{j,t}\right) + \left(\epsilon_{j,t} - \frac{1}{N}\sum_{j=1}^{N} \epsilon_{j,t}\right). \tag{5}$$

The variables in equation (5) are not growth and temperature, but are deviations of growth and temperature from their global averages. They are coarsely constructed idiosyncratic components of growth and temperature. Estimating the panel regression with time-fixed effects, equation (3), is equivalent to running stacked least squares on equation (5). The coefficient of interest β, does not measure the growth response to variations in the country's temperature. It measures the relative (to the world) growth response to relative (to the world) variations in temperature. If the panel estimate of β is negative, we can infer a country's growth is lower than average when it's temperature is hotter than average, but we cannot infer that global warming lowers growth.

3.2. Rejecting Extensive Homogeneity Restrictions with Local Projection at Horizon Zero

The literature has allowed modest amounts of heterogeneity on the slope for broad classes of countries (e.g., above and below median income) with dummy variable interactions. If one's interest is to study individual country responses, constrained (pooled) estimation should not proceed if the homogeneity restrictions are rejected. As a precursor to our main empirical work, we test, and reject, the extensive (i.e., across large numbers of countries) homogeneity restrictions that might typically be imposed in panel regressions.

Consider the regression of real GDP per capita growth on country temperature (the local projection at horizon 0),

$$\Delta y_{j,t} = \beta_j^{\tau} \tau_{j,t} + x_{j,t}' \gamma_j + \epsilon_{j,t} \tag{6}$$

where $x_{j,t}' \gamma_j$ denotes inclusion of control variables such as lagged growth and potentially lagged temperature. These controls are redundant, given our maintained assumption that temperature is plausibly exogenous to individual country GDP, but including them has the potential benefit of reducing the residual variance and the coefficient standard errors. The timing of the variables conforms to those used in Dell et al. (2012).

Next, consider regressing per capita GDP growth on global (G_t) and idiosyncratic ($I_{j,t}$) temperature factors.

$$\Delta y_{j,t} = \beta_j^{G} G_t + \beta_j^{I} I_{j,t} + x_{j,t}' \gamma_j + \epsilon_{j,t}, \tag{7}$$

where $x_{j,t}' \gamma_j$ denotes lags of growth, global temperature, and idiosyncratic temperature as controls. Since G_t and $I_{j,t}$ are an orthogonal decomposition of $\tau_{j,t}$, the global and idiosyncratic temperatures together contain the same information as the country temperature. The decomposition is useful, in the sense that thinking of climate as the multivariate probability distribution of weather variables, where changes in G_t represent variations in the mean of the global temperature distribution. This idea more closely represents the concept of climate change than country temperature alone (Hsiang, 2016). The decomposition also allows us to examine the similarity or differences in a country's GDP response to global and it's own country-specific temperature change.

Although country and global temperature are most likely to be nonstationary, they enter the regressions in levels. This specification gives a direct relationship between temperature and growth and is justified by West (1988), who

Table 1. Tests of the Homogeneity Restrictions

	All Countries		Poor Countries	
	χ^2_1	p-val	χ^2_1	p-val
Country	208.927	0.000	132.719	0.000
Global	188.399	0.002	131.565	0.000
Idiosyncratic	187.846	0.002	115.420	0.005

Notes: The Wald test statistic of the hypothesis of equality of slope coefficients on temperature variables is χ^2_1 under the null. Poor countries are those whose average real GDP per capita across the sample is below the median.

established asymptotic normality of the least squares estimator of β when $\tau_{j,t}$ (or G_t) is nonstationary. West's analysis was extended by Park and Phillips (1988), who established asymptotic normality of the least squares estimator of β when both current and lagged values of the nonstationary independent variable are included under exogeneity, which we assume.

We test the homogeneity restrictions for country temperatures by estimating equation (6) jointly as a system for $j = 1, .., N$, and performing a Wald test of the hypothesis of equal slope coefficients, $H_0 : \beta^\tau_1 = \beta^\tau_2 = \cdots = \beta^\tau_N$. The test statistic is distributed as χ^2_1 under the null. We estimate the system by weighted least squares.[9] Likewise, we test the homogeneity restrictions for global and idiosyncratic temperatures by estimating equation (7) in the same fashion and again perform Wald tests of the hypothesis of equal slope coefficients.

The left columns of Table 1 shows the results using all countries in the sample. The homogeneity restrictions are rejected by the data for the country as well as the global and idiosyncratic estimates, as the p-values of the test statistics are statistically significant. For robustness, we also estimate the system using unweighted least squares which also rejects the homogeneity restrictions for both temperature components at the 5 percent level of significance (not reported in text).

Next, we report that the split between positive and negative betas is not simply a split between rich and poor countries. In the right columns of Table 1, the test is applied only to poor countries–those whose average real GDP per capita across the sample is below the median. Here as well, the test of the homogeneity restrictions across poor countries is rejected for all temperature measures.

Overall, since the homogeneity restrictions are rejected at small p-values, we conclude that extensive pooling is not appropriate given this heterogeneity, even among poor countries.

4. Impulse Responses by Local Projections

This section first discusses our local projection specification and follows with the presentation of the main empirical results. We then undertake a cross-sectional analysis using country characteristics to help understand the results.

4.1. Local Projection Specification

Our local projections are the sequence of regressions at horizons $h \in \{0, ..., 5\}$ years, estimated separately for each country with at least 30 annual per capita GDP observations $j \in \{1, ..., 137\}$. Using country temperatures, they are,

$$y_{j,t+h} - y_{j,t-1} = \beta^\tau_{j,h}\tau_{j,t} + x'_{j,t}\gamma_{j,h} + \epsilon_{j,t+h} \tag{8}$$

where $y_{j,t}$ is 100 times log real GDP per capita of country j at time t, $\tau_{j,t}$ is the country's temperature in degrees Celsius, and $x_{j,t}$ is the vector of (potentially) lagged growth and temperature as controls. For the decomposition of temperature into global and idiosyncratic components, the local projections are

$$y_{j,t+h} - y_{j,t-1} = \beta^G_{j,h}G_t + \beta^I_{j,h}I_{j,t} + x'_{j,t}\gamma_{j,h} + \epsilon_{j,t+h}, \tag{9}$$

[9] If $\Delta\mathbf{Y}$ is the stacked vector of growth rate observations $\Delta y_{j,t}$ and $\mathbf{X}$ is the stacked matrix of independent variables, then the weighted least squares estimator is $\beta = (\mathbf{X}'\mathbf{\Omega}\mathbf{X})^{-1}(\mathbf{X}'\mathbf{\Omega}\mathbf{Y})$ where $\mathbf{\Omega}$ is a block diagonal weighting matrix where the weights are the inverse of the standard deviation of the residuals.

where G_t is global temperature at time t and $I_{j,t}$ is idiosyncratic country j temperature at time t.[10]

Lags of the control variables are determined by the Akaike's Information Criterion (AIC) for each country j and horizon h.[11] While the use of AIC may err on the side of overparameterization, we proceed given its potential of reducing residual variance and coefficient standard errors. Here we again rely on the results of West (1988) and Park and Phillips (1988) to establish our use of including country temperature ($\tau_{j,t}$) and global temperature (G_t) in levels for our estimations.

The sample length for our countries ranges from 30 to 57 annual observations. As shown by Jordà (2005) and Plagborg-Møller and Wolf (2021), the local projection coefficients are asymptotically equivalent to the impulse response function from a vector autoregression. In our analysis, we scale the country, global, and idiosyncratic local projection beta coefficients at each horizon by the cross-sectional standard deviation of the estimates. The interpretation of these standardized local projection betas are in units of standard deviation from a 1^oC increase in temperature.

Since impulse responses from vector autoregressions are colloquially referred to as responses to 'shocks,' we similarly refer to the local projection estimates as growth responses to temperature 'shocks' even though the regressor is a temperature variable (and not a 'shock' *per se*). At horizons $h > 0$, the overlapping dependent variable observations induce serial correlation in the error terms, which we address with Newey and West (1987) standard errors.

4.2. Local Projections with Country, Global, and Idiosyncratic Temperature

This section reports results for the responses of growth to country ($\tau_{j,t}$), global (G_t), and idiosyncratic ($I_{j,t}$) temperature shocks. We report summaries of the results rather than showing all of the impulse response figures. The full set of impulse responses to country, global, and idiosyncratic temperature shocks are shown in Online Appendix D.[12]

Figure 3 displays the histograms of the country (Panel A), global (Panel B), and idiosyncratic (Panel C) temperature standardized local projection betas at horizons 0 and 5. For the different temperature components and horizons, the local projection estimates are approximately centered around zero and extensive heterogeneity is observed.

Table 2. Country, Global, and Idiosyncratic Temperature Standardized Local Projection Summary

	A. Country Temperature						B. Global Temperature						C. Idiosyncratic Temperature					
	Standardized Local Projection Betas						Standardized Local Projection Betas						Standardized Local Projection Betas					
Horizon	0	1	2	3	4	5	0	1	2	3	4	5	0	1	2	3	4	5
# neg	76	75	73	73	78	76	69	55	65	67	72	65	76	78	69	75	77	73
# pos	61	62	64	64	59	61	68	82	72	70	65	72	61	59	68	62	60	64
# sig neg	6	9	15	19	23	24	11	8	6	8	12	21	11	12	15	16	17	16
# sig pos	5	10	14	15	20	22	4	11	9	11	14	24	6	7	7	5	5	5

Notes: This table shows the count of country (Panel A), global (Panel B), and idiosyncratic (Panel C) temperature standardized local projection (estimates from equations (8) and (9)) betas that are negative (neg), positive (pos), and statistically significant at the 5 percent level (sig neg and sig pos). Specifications are determined by Akaike's Information Criterion (AIC).

Table 2 reports a summary of the local projection betas. Panel A shows results for the country temperature shock and Panels B and C show the results for the global and idiosyncratic temperature shocks, respectively. Comparing across horizons reveals similar numbers of positive and negative point estimates. Negative betas often outnumber positive betas for the country and idiosyncratic temperature shocks (Panels A and C), but the number of positive and negative betas for the global temperature shock is slightly more even (Panel B).

Figure 4 plots the standardized local projection betas at horizons 0 and 5 onto a world map. Results for country temperature shocks are in Panel A, global temperature shocks are in Panel B, and idiosyncratic temperature shocks are in Panel C. Negative responses are shown in red and positive responses in green. We split the size of the coefficients into quartiles and the deep green is the top quartile of the responses to temperature variation and the deep red is the lowest quartile containing the most negative responses. A map of statistical significance of the results is available in Online Appendix E.

[10] Burke et al. (2015) and papers that followed posit there may be non-linear effects of temperature on GDP growth. We address this in Section 5.

[11] The local projection specifications for each country are reported in Online Appendix C.

[12] Non-standardized local projection betas are reported in the impulse response figures.

Figure 3. Country, Global, and Idiosyncratic Temperature Standardized Local Projection Betas

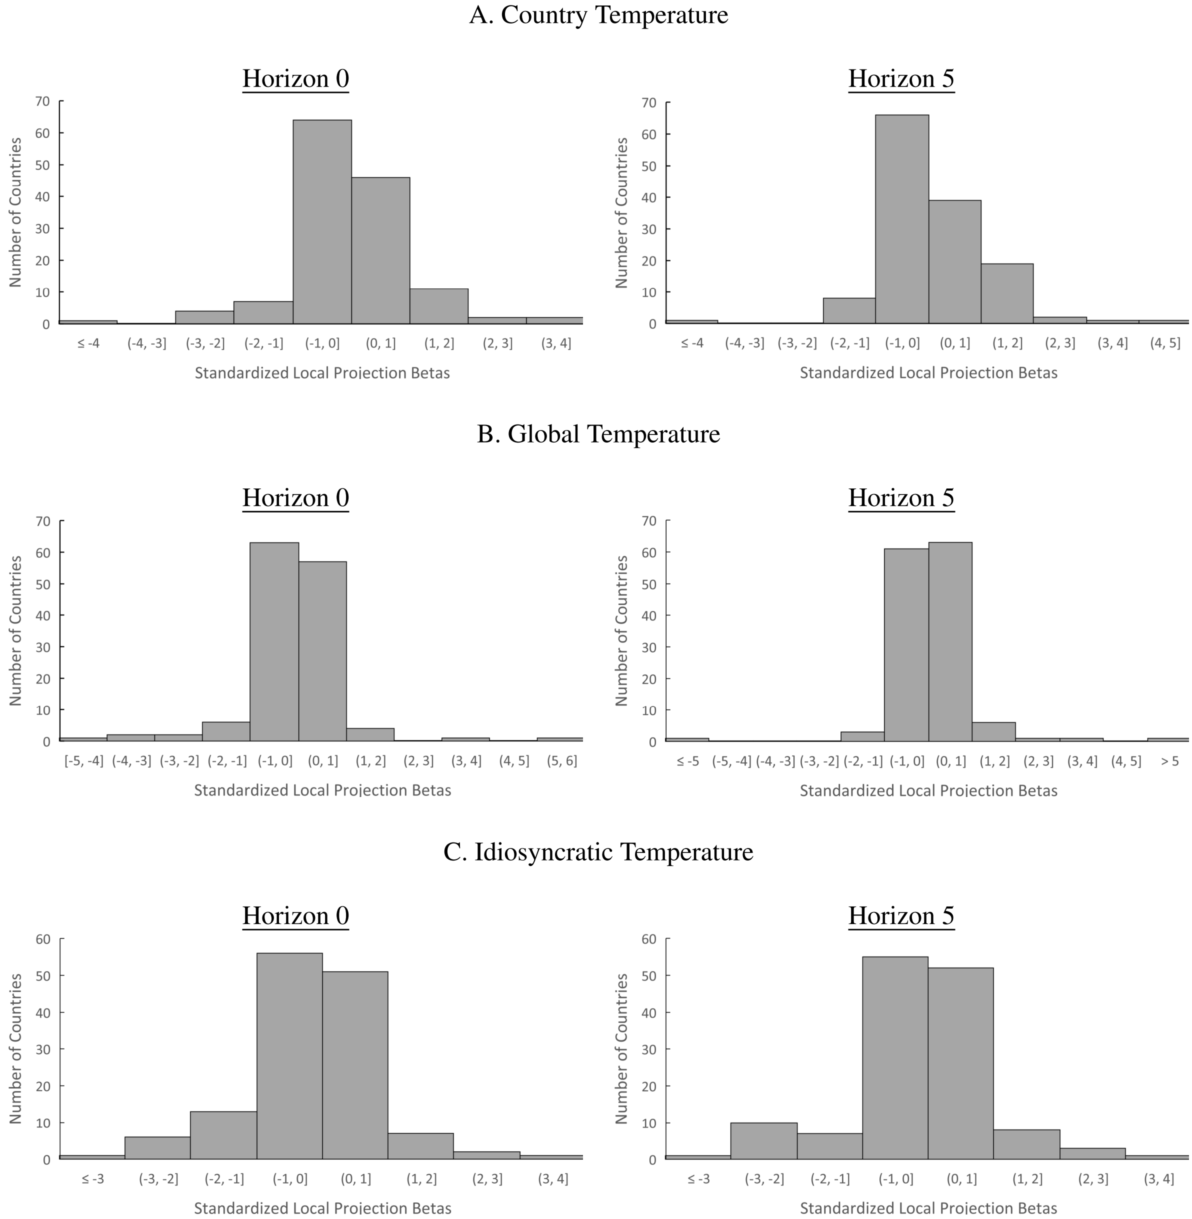

Notes: Distribution of country temperature ($\tau_{j,h}$), global temperature (G_t), and idiosyncratic temperature ($I_{j,t}$) standardized local projection betas from equations (8) and (9) for $j = 1, ..., 137$ and $h = 0$ and 5. Specifications are determined by Akaike's Information Criterion (AIC).

We first look at the responses to country temperature shocks. Visually, there is heterogeneity in the direction of the responses across countries and horizons. At horizon 0, large, negative responses are found for relatively high (e.g. Italy and South Korea), middle (e.g. Brazil and Indonesia), and low (e.g. Republic of the Congo and Mali) income countries. At horizon 5, almost all of the rich countries have negative responses. Regionally, from the far corner of continental Southeast Asia across the island nations through to New Zealand all have negative responses, as do some of the larger countries in the Americas such as Brazil and Mexico. Surprisingly, some of the poorest countries experience significantly positive growth responses to positive country temperature shocks, particularly throughout

Figure 4. Country, Global, and Idiosyncratic Temperature Standardized Local Projection Betas

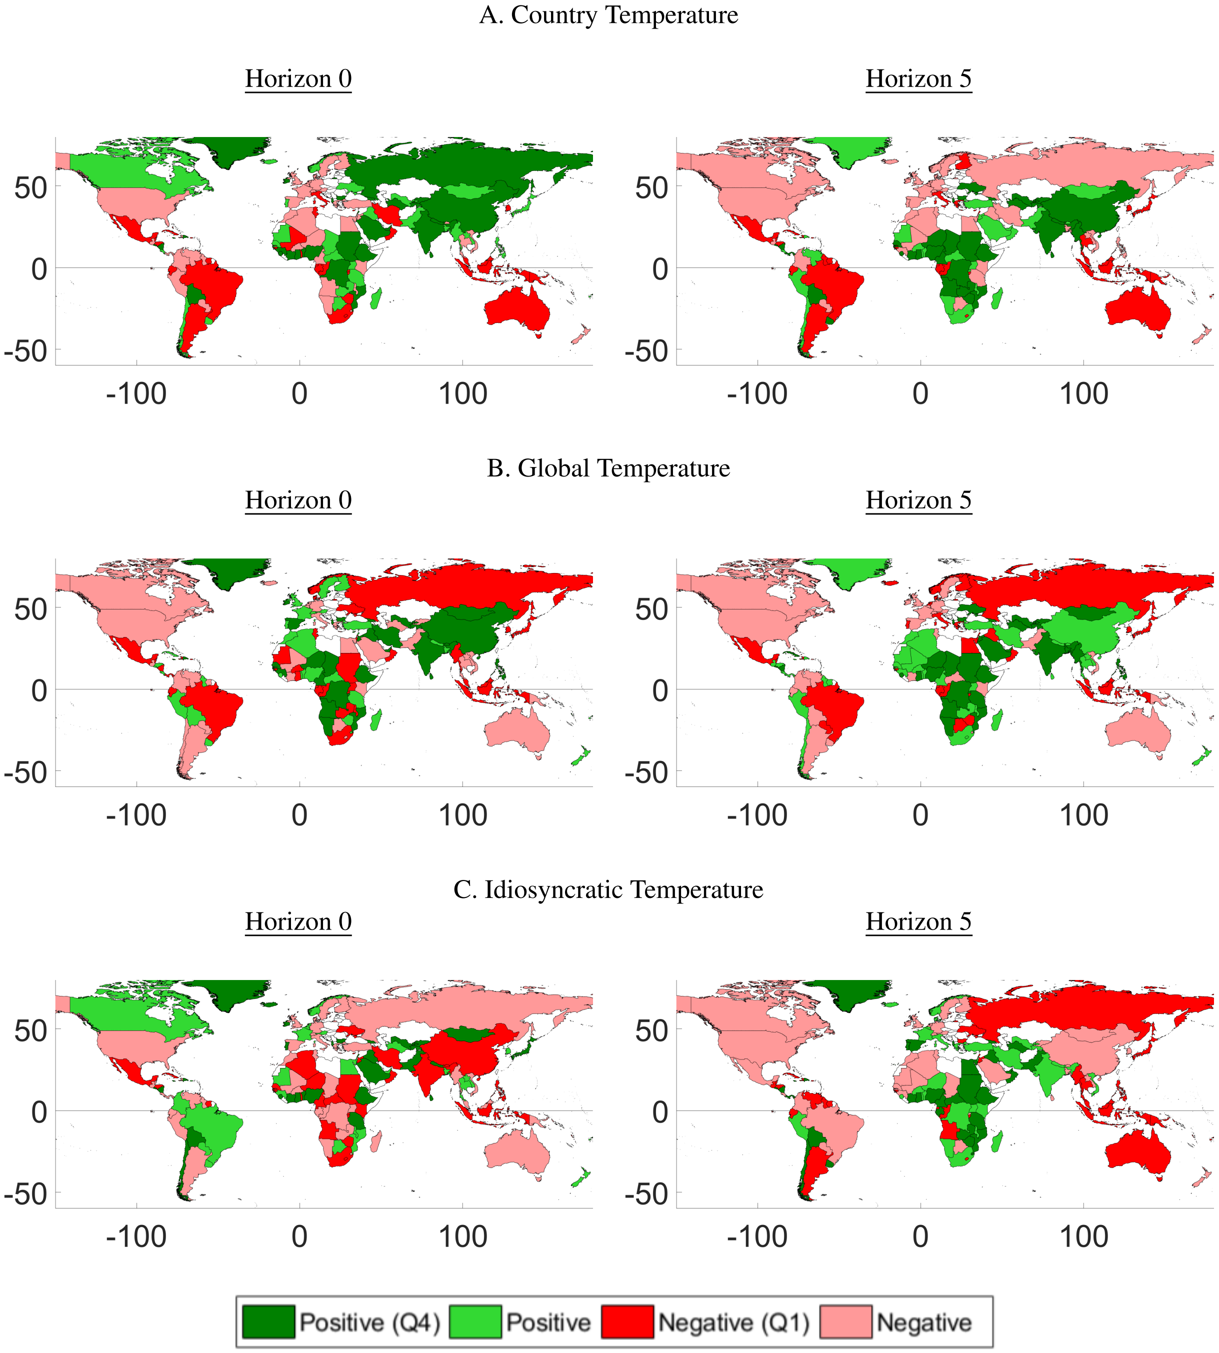

Notes: Country (Panel A), global (Panel B), and idiosyncratic (Panel C) temperature standardized local projection betas are from equations (8) and (9) for $h = 0$ and $h = 5$. Specifications are determined by Akaike's Information Criterion (AIC). Green are positive coefficients and red are for the negative coefficients. Deep green are the fourth quartile of beta estimates and deep red are the first quartile of beta estimates.

much of Sub-Saharan Africa. Large Asian countries such as China and India likewise have relatively large, positive responses.

Moving next to the global responses to global temperature shocks, at horizon 0 there are again large, negative responses for higher and lower income countries. However, some countries such as Zambia, Uganda, and Ghana exhibited positive responses to the country, but negative to the global temperature shock. As in the country temperature shock, at horizon 5 almost all of the rich countries have negative responses to the global temperature component. Many less developed and developing countries in Sub-Saharan Africa and South Asia have positive responses to positive global temperature shocks. Some of the larger oil producing countries in OPEC have large positive responses by horizon 5 such as Angola, Nigeria, Saudi Arabia, and UAE. Algeria's and Iran's are positive, but less large, while Venezuela's and Iraq's are negative.

Next, we look at the response to idiosyncratic temperature shocks. Here, at horizons 0 and 5, negative responses outnumber positive ones and there is less of a pattern amongst rich countries. At horizon 5, the responses of rich countries such as the United States, Canada, Germany, and Japan are all negative. At longer horizons, the coefficient signs are negative for many countries in Southeast Asia and Oceania, and several countries in Central and South America.

Interestingly, a visual comparison between the global and idiosyncratic shocks within each horizon shows the direction of a country's response are sometimes at odds with each other. Prominent examples include China and India at horizon 0, as well as for oil states such as Angola and Iran. The same sized coefficients on global and idiosyncratic temperature should not be interpreted to mean the two temperature components are equally important. Global temperature is trending up while idiosyncratic temperature is, by construction, stationary around zero. This makes the global temperature shocks quantitatively more important. For example, the country temperature shock results contain information from both temperature components. At horizon 5, of the countries whose signs differ between the idiosyncratic and global shocks, the county temperature shock coefficients agree with the global component in around 60 percent of the cases. This is consistent with Byrne and Vitenu-Sackey (2024), who likewise decompose country temperature into global and idiosyncratic components for 30 countries in an alternative method to ours, and find the global component is quantitatively more important on GDP growth.

Figures 5 and 6 display the local projection impulse responses to country, global, and idiosyncratic temperature shocks for a set of rich (Figure 5) and poor (Figure 6) countries. The rich are represented by the G-7 countries plus Australia and China and the poor are the nine poorest countries in our sample, based on average real GDP per capita over the sample.[13] Amongst the rich, real GDP per capita initially declines for many of the countries following an increase in country, global, and idiosyncratic temperature. Amongst the poorest countries, the direction of the responses of real GDP per capita from positive country, global, and idiosyncratic temperature shocks are more country specific.

Next, we put the country-level responses in context of the world-wide response. Table 3 reports the GDP-weighted responses across all countries for each horizon and temperature shock. The responses are the not standardized beta coefficients, so the interpretation is the annual growth response to a 1^{o}C increase in each temperature component. To follow the population weights used to construct the weighted temperatures, the GDP weights are from the year 2000. First, notice across all temperature shocks and horizons, the world-wide responses are negative. Although we do not estimate a pooled regression, these results are consistent with the literature that finds overall growth in pooled samples are harmed by increases in temperature. Typical estimates show that a local temperature increase by 1^{o}C reduces GDP by less than 1 percent on impact (Dell et al., 2012), which is in line with our idiosyncratic response at horizon 0. In the medium run, this grows to 1-3 percent (Dell et al. (2012); Burke et al. (2015)), as is the case for our results as well. In a related paper to ours, Bilal and Känzig (2024) also decompose temperature into local and global components in an alternative way to ours and find an increase by 1^{o}C to local temperature decreases GDP by 1.5 percent after 5 years, and an increase in global temperature can reduce GDP by nearly 10 percent after 5 years. The magnitude of our global response at horizon 5 is smaller at -4.3 percent, but we likewise find the global effect to be larger in magnitude than the idiosyncratic component.

Our results thus indicate that both observations are consistent: 1.) On average, world-wide growth may be harmed by increases in temperature and 2.) many countries respond positively to temperature shocks in the historical data and, surprisingly, this applies to a significant set of the least developed countries.

[13] China is grouped with the rich countries, not on the basis of per capita GDP but because it is the world's second largest economy.

Figure 5. Impulse Responses of Growth to Country (Red Small Dashed), Global (Black Solid), and Idiosyncratic (Blue Dashed) Temperature Shocks–G-7 Plus Australia and China

A. Country Temperature

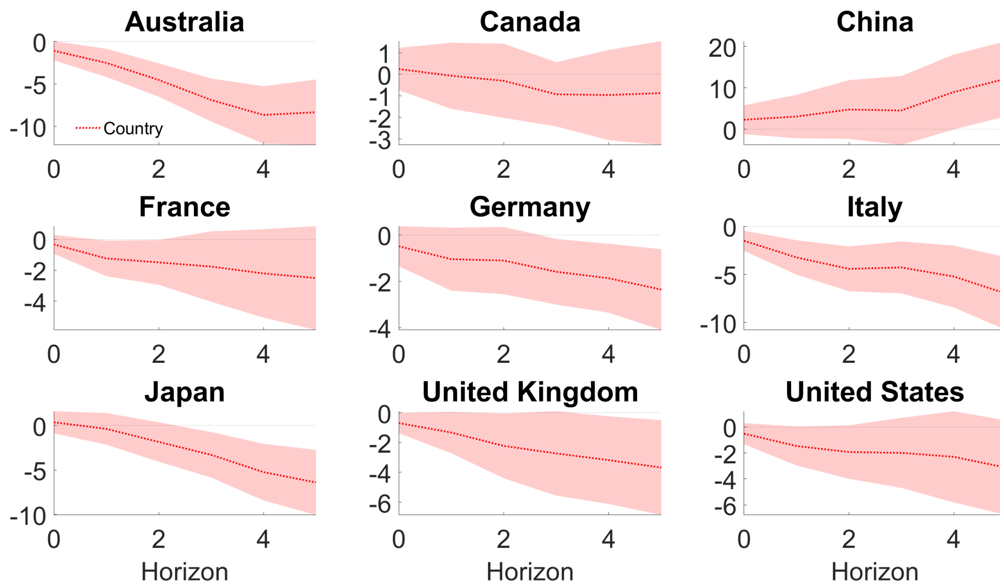

B. Global and Idiosyncratic Temperature

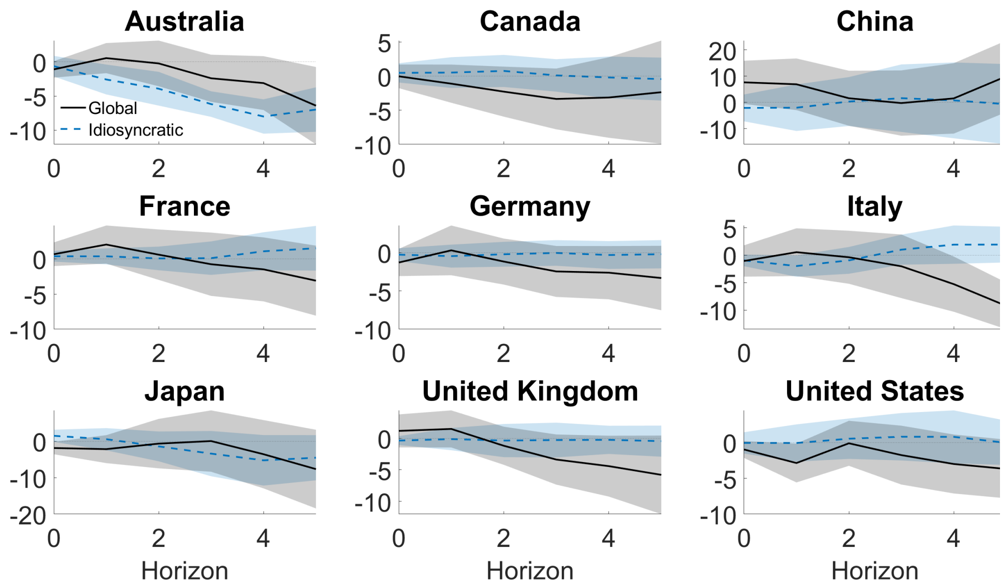

Notes: Local projection betas are not standardized. Shaded areas are plus and minus 1.96 standard error bands. Specifications are determined by Akaike's Information Criterion (AIC).

To summarize this section, the local projection results establish three main findings. First, there is substantial heterogeneity in the responses across countries, irrespective of the source of temperature fluctuations. Second, the variation in growth responses from the different sources of temperature fluctuations highlights the importance of sep-

Figure 6. Impulse Responses of Growth to Country (Red Small Dashed), Global (Black Solid), and Idiosyncratic (Blue Dashed) Temperature Shocks–Nine Poorest Countries

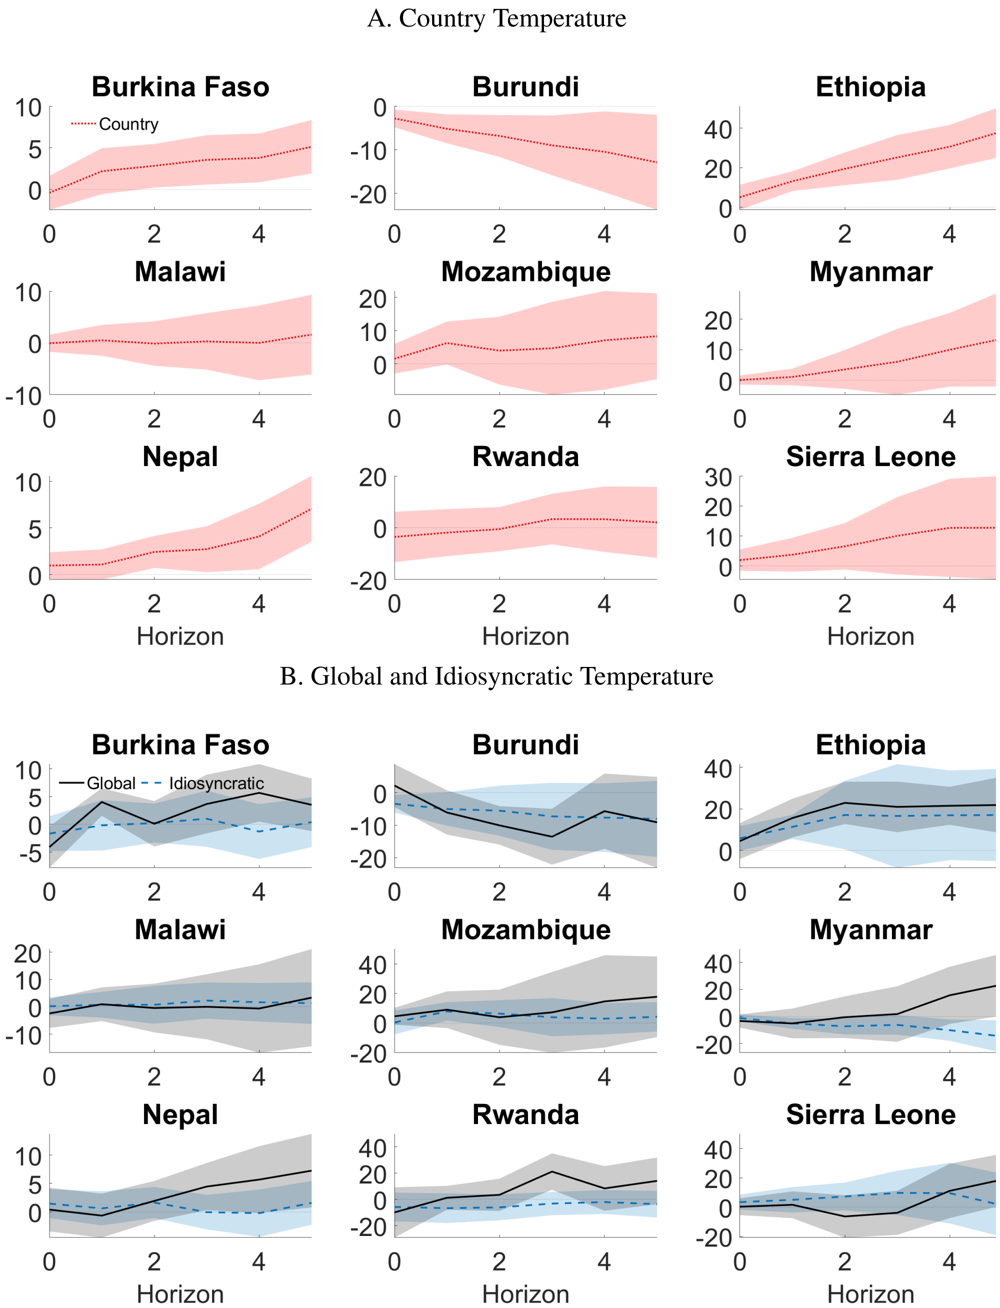

Notes: Local projection betas are not standardized. Shaded areas are plus and minus 1.96 standard error bands. Specifications are determined by Akaike's Information Criterion (AIC).

arately considering idiosyncratic, country-specific from global temperature change. While the signs of the growth responses from global and idiosyncratic temperatures sometimes coincide, there are instances where they go in opposite directions. Finally, we show some of the most developed countries evidently face substantial economic damages from global temperature change. Our aggregated results are consistent with previous findings, but this masks variation in the direction and size of growth responses due to temperature change.

Table 3. World-Wide Responses Weighted by Country GDP

	Horizon					
	0	1	2	3	4	5
Country	-0.238	-0.627	-0.987	-1.544	-1.840	-2.166
Global	-0.687	-1.125	-1.013	-1.996	-3.220	-4.286
Idiosyncratic	-0.271	-0.511	-0.565	-0.701	-1.186	-1.550

Notes: This table shows the GDP-weighted (using 2000 as a reference year) not standardized responses across all countries for the country, global, and idiosyncratic temperature shocks at horizons $h = 0, ..., 5$.

4.3. *Cross-Sectional Response Heterogeneity and Country Characteristics*

What explains the response heterogeneity across countries? This section investigates how country characteristics, including geographic, economic, demographic, and political factors can explain the variation in responses. The analysis is based on a cross-sectional regression of the country/global/idiosyncratic standardized local projection betas on these country characteristics.[14] Although the betas are estimated, there is no 'second stage' or generated regressors problem because the estimated response coefficients are the dependent variable in the regressions. If X_j is the vector of country $j's$ characteristics and the constant, for each country, global, and idiosyncratic standardized impulse response estimate ($\hat{\beta}_{j,h}$) we run the cross-sectional regression

$$\hat{\beta}_{j,h} = X_j'\gamma_h + u_{j,h}, \tag{10}$$

at horizons $h = 0, \ldots, 5$.

The variables we use are based on the following considerations. In light of panel studies finding response differences between rich and poor countries, we include average log real GDP per capita (log(GDPPC)). Extant research would lead one to expect log income to enter with a positive coefficient. We also consider a country's long-term growth rate (L.T. Growth), which is constructed as the growth rate of real GDP per capita over the full sample of observations. The country's average openness (Openness), which are exports plus imports as a share of GDP, captures the degree of economic connectedness to the rest of the world. We also include the average share of agriculture and industrial employment (Ag.&Indu.Empl) since labor productivity in these sectors have been seen as a direct channel through which temperature affects the economy. Agricultural workers, especially in poorer countries, are directly exposed to temperature as are the crops themselves, and Deryugina and Hsiang (2014), Deschênes and Greenstone (2007), Nelson et al. (2014), and Dietz and Lanz (2019) report empirical damage estimates to agriculture from high temperatures. Labor productivity has been found to suffer from higher temperatures for un-airconditioned factories (e.g., Somanathan et al. 2021). Average high-school attainment (High School) gives a coarse measure of human capital accumulated and Democracy examines the potential role of political responses to temperature. We include the interquartile range (IQR) of country temperature as experience with wider annual temperature variation may influence a country's response to temperature change.

Except for temperature, the data are from the World Bank's, *World Development Indicators* and are the country's time series average over the available sample span. Democracy is the World Bank's Index of Democratization. Average temperature is included primarily as a control variable.

Table 4 shows the correlations amongst these variables. The inverse relationship between temperature and income in the cross-section (correlation -0.531) has been well studied (Dell et al., 2009). In what follows, we present specifications with the entire list of variables together as regressors to mitigate potential omitted variables bias.

[14]Recently, Lustig and Richmond (2020) employed the same methodology to regress exchange rate betas on gravity variables.

Table 4. Correlation Matrix of Explanatory Variables

	L.T. Growth	Openness	High School	Democracy	Ag.&Indu. Empl.	Temperature	IQR
log(GDPPC)	0.145	0.257	0.711	0.740	-0.886	-0.531	0.376
L.T. Growth		-0.033	0.097	0.145	-0.168	-0.166	0.216
Openness			0.199	0.112	-0.276	-0.107	0.072
High School				0.630	-0.635	-0.684	0.588
Democracy					-0.649	-0.624	0.304
Ag.&Indu.Empl.						0.420	-0.278
Temperature							-0.740

Table 5 shows cross-sectional regressions of the local projection betas at horizons 0, 3, and 5, on country characteristics. The local projection betas are standardized and the bold indicates statistical significance at the 5 percent level and * indicated statistical significance at the 10 percent level.

Looking at the results for country temperatures at horizon 5, higher long term growth is statistically associated with lower country temperature shock responses while high school attainment and agricultural and industrial employment are associated with higher responses. To put these coefficient estimates in perspective, a 10 percentage point increase in high school attainment - equivalent of moving from the high school attainment in Mongolia to Denmark - is associated with a 1.59 standard deviation increase in β_5^τ. Across all horizons and temperature components, we do not find a definitive relationship between average temperature and growth responses. This is in contrast to Cruz and Rossi-Hansberg (2024) who find productivity gains for colder regions due to warmer temperatures and productivity damages to hotter regions.

The global temperature results are qualitatively similar to the country temperature results. At longer horizons, richer, and faster growing countries are more likely to have a lower response to higher global temperature. Countries that do more trade and that are more educated tend to have more favorable growth responses to higher global temperature. A 10 percentage point increase in openness is associated with an increase in the global response at horizon 5 by 0.49 standard deviations. In terms of the global temperature betas, the effect of temperature on GDP need not be restricted to temperature within its borders. While some part of a global temperature shock may represent the direct effect of country temperature on GDP, a good portion may also be the effect on the rest-of-world (ROW) economy and subsequent indirect effects on individual countries through trade and finance linkages, which we proxy with openness. To suggest a potential mechanism, the positive coefficient on openness and negative coefficient on per capita GDP is consistent with the following: suppose the global temperature shock has the effect of an uncertainty shock and has a larger effect on colder and richer countries. This affects these countries like a negative aggregate demand shock which improves the terms of trade for poorer and hotter countries. Countries that are more open to trade are able to benefit from this.[15] Indeed, Lee et al. (2022) find, at the four year horizon, exchange rates tend to appreciate for hotter, open countries from temperature shocks.[16]

Lastly, the idiosyncratic temperature responses are largely unsystematic in that these country characteristics are generally not significant. The exceptions are high school attainment at horizon 0 and agricultural and industrial share in horizon 5, which do align with the results from country temperature. The reason why we do not find much significance in the country characteristics may be related to the findings in Bilal and Känzig (2024) and Byrne and Vitenu-Sackey (2024). They also consider growth effects from local and global temperature effects and find the global temperature is quantitatively more impactful on growth. In this regard, if GDP changes are driven primarily by the global component, then we may expect the country attributes to interact more with this primary driver rather than the less important idiosyncratic component.

[15] Berg and Mark (2022) show how an uncertainty shock causes terms-of trade deterioration in the country experiencing the shock.

[16] This is not the same for climate disasters. Hale (2022) shows safe country currencies appreciate relative to risky country currencies following a climate disaster shock.

Table 5. Cross-Sectional Analysis of Temperature Standardized Local Projection Betas

	Country			Global			Idiosyncratic		
Horizon	0	3	5	0	3	5	0	3	5
log(GDPPC)	-0.194	-0.188	-0.121	-0.045	-0.234	-0.252	-0.183	0.144	0.210
	(-1.548)	(-1.616)	(-1.074)	(-0.329)	**(-2.000)**	**(-2.402)**	(-1.291)	(1.029)	(1.536)
L.T. Growth	-0.206	-0.282	-0.246	0.026	-0.229	-0.303	0.016	0.022	0.012
	(-2.065)	**(-3.045)**	**(-2.744)**	(0.240)	**(-2.450)**	**(-3.621)**	(0.137)	(0.196)	(0.115)
Openness	0.224	0.158	0.120	0.446	0.493	0.486	-0.009	-0.189	-0.139
	(-0.972)	(-0.738)	(-0.577)	(1.766)*	**(2.286)**	**(2.519)**	(-0.033)	(-0.732)	(-0.554)
High School	1.138	1.565	1.591	0.531	1.742	1.875	1.005	0.714	0.754
	(2.579)	**(3.818)**	**(4.016)**	(1.098)	**(4.220)**	**(5.079)**	**(2.008)**	(1.446)	(1.568)
Democracy	0.725	-0.898	-1.570	-0.244	-1.122	-1.629	1.540	-1.469	-1.442
	(-0.654)	(-0.872)	(-1.578)	(-0.200)	(-1.083)	(-1.757)*	(1.225)	(-1.184)	(-1.195)
Ag.&Indu.Empl	0.928	1.146	1.650	0.568	0.713	0.898	-0.271	1.112	2.051
	(1.080)	(1.437)	**(2.139)**	(0.603)	(0.888)	(1.250)	(-0.278)	(1.157)	**(2.193)**
Temperature	0.020	0.012	0.009	0.026	0.026	0.009	0.026	0.010	0.003
	(1.035)	(0.655)	(0.527)	(1.194)	(1.425)	(0.521)	(1.175)	(0.448)	(0.158)
IQR	0.050	0.030	0.023	0.044	0.053	0.034	0.036	0.004	-0.003
	(2.337)	(1.483)	(1.200)	(1.861)*	**(2.616)**	(1.905)*	(1.462)	(0.166)	(-0.125)
R-Square	0.226	0.327	0.359	0.094	0.366	0.461	0.100	0.049	0.075
Observations	121	121	121	121	121	121	121	121	121

Notes: Local projection betas are standardized. T-ratios in parentheses. Significance at the 5% level indicated by bold face and at the 10% level by '*'.

5. An Alternative Measure of Temperature

The previous sections studied GDP per capita growth responses to variations in annual average temperature. In this section, we employ an alternative measure of temperature to explore potential non-linearities in the responses within countries.

5.1. Temperature Measurement

Our alternative measure of temperature uses daily temperature observations to calculate the number of relatively cold, normal, and hot days in a country and year. For our data, we use Berkeley Earth[17] daily temperatures on a $1^o \times 1^o$ Latitude-Longitude Grid. We then reconstruct our population-weighted country-by-country temperature by day in the same manner we constructed our annual temperatures.

Beginning with the country-level temperatures, for each country we calculate the temperature quartiles using the full sample of observations from 1960 to 2017. For each year, we count the number of days that the country experiences each of three temperature classifications–the number of days when temperature is in the fourth quartile, the interquartile range, and the first quartile. We refer to temperatures in the fourth quartile as 'hot', those in the interquartile range as 'normal' and those in the first quartile as 'cold'. The separation of temperature into bins is similar to that used in studies such as Somanathan et al. (2021), Zivin and Neidell (2014), Barreca et al. (2016), Park et al. (2020), and Barreca et al. (2015).

Analogously, we employ the same method to calculate the global temperature quartiles, except in this case, we classify days into global temperature bins. In contrast to the country-level temperature, the number of days in each bin and year are identical across countries. In this analysis, we omit the idiosyncratic component as the interpretation of a discrete 'hot' or 'cold' idiosyncratic day is not clear.

To provide context for variation in temperature across countries, Table 6 shows the interquartile range for global temperature and selected countries. In our sample, the United Arab Emirates is the hottest country where a quarter of the days are above 33.34^{o}C and a quarter below 21.82^{o}C. Greenland is the coldest country, but since very few people live there, we also list Iceland, where temperatures above 8.01^{o}C are considered hot and below -1.18^{o}C are cold. Average global temperatures above 23.51^{o}C are considered hot and below 14.19^{o}C are cold.

Figure 7 shows days in each of the global temperature bins over the sample. Due to global warming, the number of hot days has trended up and normal and cold days have trended down.

[17] *Berkeley Earth*, Global Daily Land Average Temperature, $1^o \times 1^o$ Latitude-Longitude Grid. Accessed at *berkeleyearth.org* on February 14, 2024.

Table 6. Temperature Quartiles for Global and Selected Countries

	Celsius	
	Q4	Q1
Global	23.51	14.19
Greenland	3.24	-8.75
Iceland	8.01	-1.18
United Arab Emirates	33.34	21.82

5.2. *Empirical Specification*

We again perform local projections as the sequence of regressions at annual horizons $h \in \{0, ..., 5\}$ for each country j. Let $D^H_{j,t}$ be the number of days of hot temperature in country j in year t, $D^N_{j,t}$ be days of normal temperature, and $D^C_{j,t}$ be days of cold temperature. The local projection can be specified as

$$y_{j,t+h} - y_t = \beta^H_{j,h} D^H_{j,t} + \beta^N_{j,h} D^N_{j,t} + \beta^C_{j,h} D^C_{j,t} + x'_{j,t}\gamma_{j,h} + \epsilon_{j,t+h},$$

where $x_{j,t}$ is the vector representing the constant and lags of GDP growth which are determined by the Akaike's Information Criterion (AIC) for each country j and horizon h. If we ignore leap years, $D^C_{j,t} = 365 - D^H_{j,t} - D^N_{j,t}$, so an equivalent specification is to let cold days be the omitted category which results in

$$y_{j,t+h} - y_t = \left(\beta^H_{j,h} - \beta^C_{j,h}\right) D^H_{j,t} + \left(\beta^N_{j,h} - \beta^C_{j,h}\right) D^N_{j,t} + x'_{j,t}\gamma_{j,h} + \epsilon_{j,t+h}, \tag{11}$$

where the slope coefficient on the number of hot days $\left(\beta^H_{j,h} - \beta^C_{j,h}\right)$ is the effect on growth from substituting a hot day for a cold day, and the coefficient on the number of normal days $\left(\beta^N_{j,h} - \beta^C_{j,h}\right)$ is the effect on growth from substituting a normal day for a cold day. Our estimation is performed on equation (11) for binned country temperature and an analogous specification holds for binned global temperature. These bins allow us to identify potential non-linear relationships between output and temperature.

Figure 7. Days in Hot, Normal, and Cold Global Temperature Bins

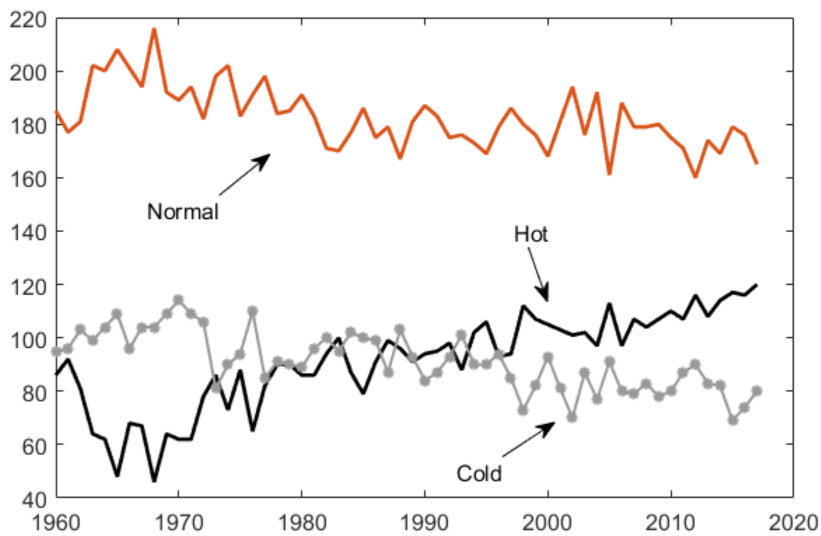

5.3. *Empirical Results*

We begin by analyzing the growth effects of an additional hot day. In comparison to the results using average temperature which reported responses from a 1^{o}C change in temperature, changes from a cold to a hot day capture larger changes in temperature. Figure 8 plots the results of substituting a hot day for a cold day on the global map. Panel A shows the results for country temperature and Panel B shows the responses for global temperature. Negative responses are shown in red and positive responses in green. We split the size of the coefficients into quartiles and the deep green is the fourth quartile of the responses to temperature variation and the deep red is the first quartile, which contains the most negative responses. Maps of statistical significance of the results are available in Online Appendix F.

We first look at the responses to country temperature. Similar to the previous results using average temperature, there is again substantial heterogeneity in the directional response of receiving an additional hot day. By horizon 5, many countries in Asia have relatively large positive responses, as do many in Sub-Saharan Africa. Most of the rich world, including Western Europe, Japan, South Korea, and Australia have large negative responses with the exceptions of Canada and New Zealand.

Figure 8. Substituting a Hot Day for a Cold Day

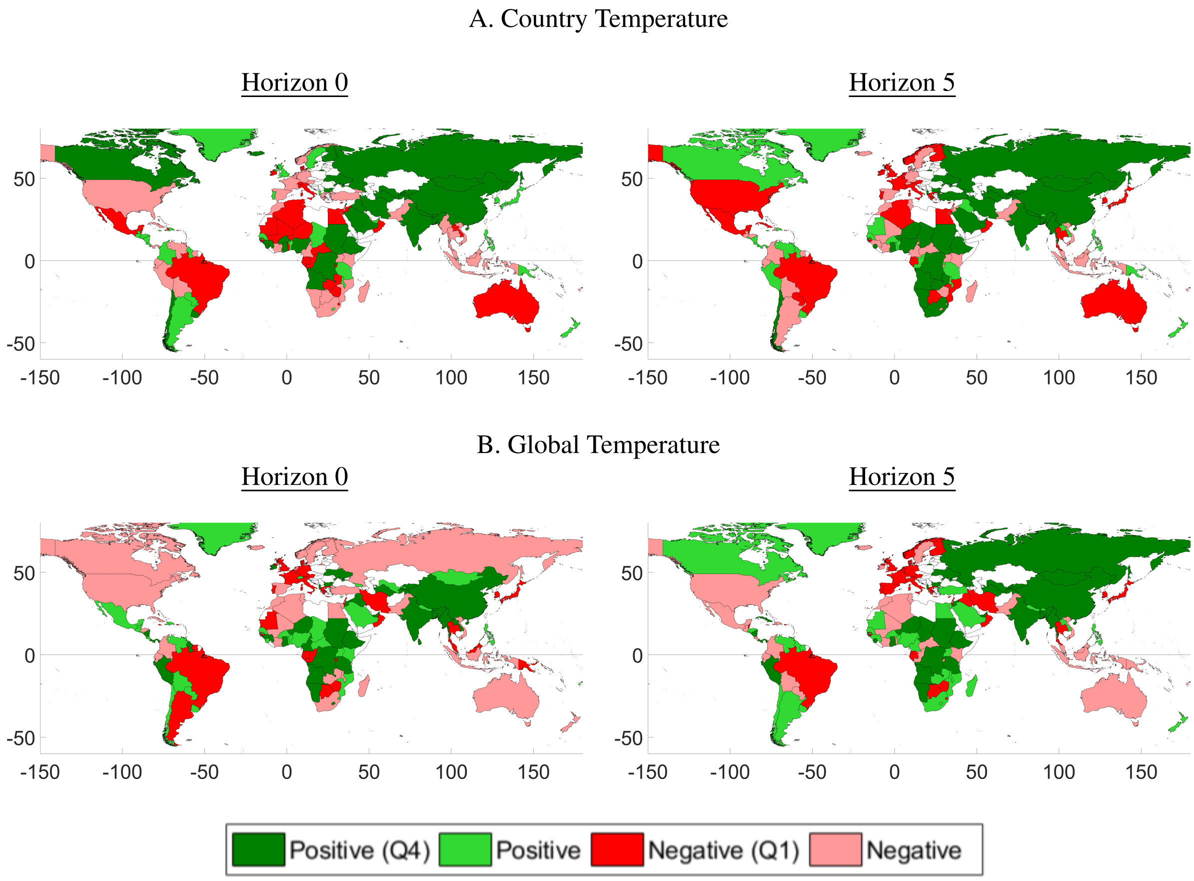

Notes: Figure shows differences in coefficient estimates of an additional day in the hottest compared to the coldest temperature bin, $\beta^H_{j,h} - \beta^C_{j,h}$, for country (Panel A) and global (Panel B) components estimated in equation (11). The comparisons are at horizons $h = 0$ and $h = 5$. Green are positive differences and red are negative differences. Deep green are the fourth quartile of differences and deep red are the first quartile of differences.

Moving next to the global temperatures in Panel B, we see similar, but not identical, patterns. In comparison to the country temperature at horizon 0, some differences in signs emerge such as Canada, Mexico, Iran, and Russia. Again, by horizon 5 much of the rich world still exhibits negative responses, but the relative magnitudes for the U.S. and Australia are less pronounced. Also, similar patterns for moving from a cold to a hot day for the country and global components exists for Asia and much of Africa.

With these results, we next ask how the responses from moving from a relatively cold to a hot day compare to our country betas using average temperatures (the results in Figure 4). Figure 9 shows the world map of this comparison. The countries in green are those whose signs on the responses are both positive from moving from a relatively cold to a hot day and betas using average temperatures. The countries in red are those whose signs on the responses are both negative from moving from a relatively cold to a hot day and betas using average temperatures. The countries in black are the ones whose signs are in disagreement.

Figure 9. Response Comparison of Results: Substituting a Hot Day for a Cold Day Responses vs. Average Temperature Responses

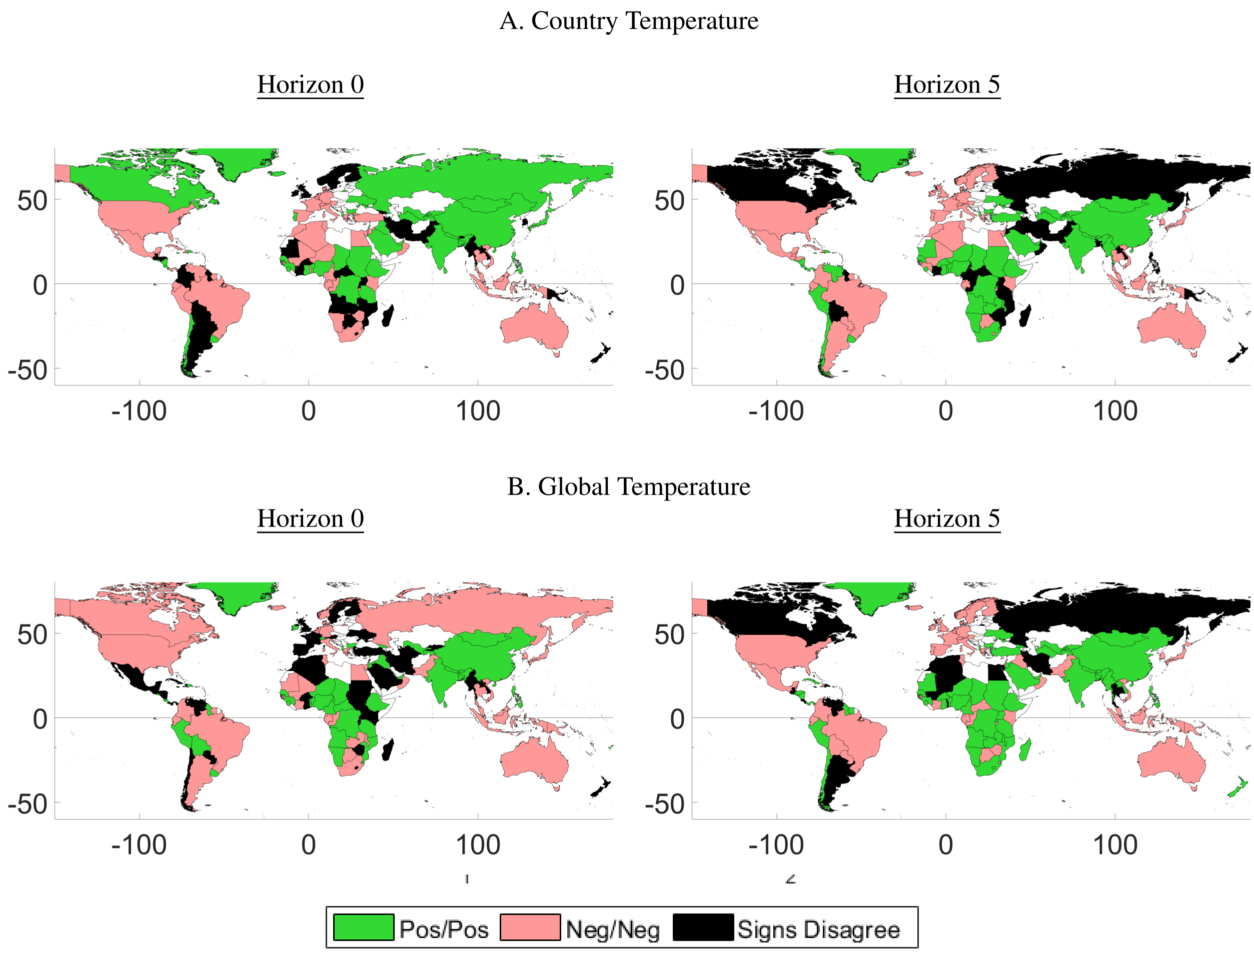

Notes: Figure shows differences in coefficient estimate signs of moving from a relatively cold to a hot day compared to country betas using average temperatures. The comparisons are at horizons $h = 0$ and $h = 5$. Green are sign agreements when both coefficients are positive, red are sign agreements when both coefficients are negative, and black are countries where there is a disagreement in sign.

We start with the responses to country temperature in Panel A. In comparison to the country betas from average temperature, the signs of the coefficients are similar to the results here. At horizon 0, 74 percent of the signs are in agreement. Notable differences with disagreeing signs include the United Kingdom and New Zealand (both positive for an additional hot day and negative for country temperature). At horizon 5, the signs are in agreement in 82 percent of the countries. For example, Iran and Pakistan have disagreeing signs (both positive for an additional hot day and negative for country temperature). The two largest countries by area, Russia and Canada, also stand out. They both had negative responses using average temperatures, but have positive responses when moving to an additional hot from a cold day.

For the global temperature in Panel B, in comparison to the global betas from average temperatures, the signs of the coefficients are similar to the results here. At horizon 0, 66 percent of the signs are in agreement. Notable differences with disagreeing signs include Mexico (positive for an additional hot day and negative for global temperature) and the United Kingdom and France (both negative for an additional hot day and positive for global temperature). At horizon 5, the signs are in agreement in 88 percent of the cases. Country examples with disagreeing signs include Egypt (positive for an additional hot day and negative for global temperature) and Iran (negative for an additional hot day and positive for global temperature), along with Russia and Canada. Even with this alternative temperature measure, at horizon 5, growth in many rich countries are harmed by moving from an additional cold to hot day. Likewise, growth for many relatively poorer countries in South Asia and Sub-Saharan Africa increase by moving from an additional cold to hot day, similar to the findings using average temperatures.

We show the results of moving from a cold to normal day in Online Appendix F. These results, in conjunction with the estimates of moving from a cold to a hot day, capture non-linear relationships between changes to output and temperature. To highlight the distribution of the responses, Figures 10 and 11 plot the coefficient for selected countries (relatively rich and poor). These show responses for both country and global temperatures at horizon 5 estimated from equation (11). Each interval is the percent change in GDP per capita at horizon 5 from an additional day moving from the coldest bin to a normal and hot day. Note that the cold day is the omitted category so it is always centered around 0. The shaded areas represent 95 percent confidence intervals. As we previously classified countries, Figure 10 shows the estimates for a set of rich countries – the G-7 countries plus Australia and China – and Figure 11 are results from the nine poorest countries based on average real GDP per capita over the sample.

Turning first to the country temperature responses for the set of rich countries in Panel A of Figure 10, for most countries (with the exceptions of Canada and China) there is an inverse relationship between temperature and output with the magnitude of the negative responses growing with higher temperatures. In fact, 6 of the 9 countries show statistically significant declines in GDP by horizon 5 by moving to an additional hot day. For these countries, damages are growing with rising temperatures as hot days show greater losses than having a normal temperature day. In the U.S., moving from a cold to a hot day reduces GDP per capita by 0.2 percent after 5 years whereas in China it increases GDP per capita by 0.6 percent over the same horizon.

The global temperature responses in Panel B show a similar pattern as the country responses. Here again, output is declining for most countries with the largest changes (in magnitude) for the larger movements to hot days compared to moving to a normal day.

Finally, moving to the set of poor countries in Figure 11, Panel A shows the country temperature response distributions. Recall from the maps that many of the relatively poor countries responded positively to increases in temperature. Here 7 of the 9 countries have positive responses from moving to an additional hot day by horizon 5, and 5 showing statistically significant positive responses. There is a positive relationship between temperature and output with the magnitude of the positive responses growing with higher temperatures for countries such as Burkino Faso, Ethiopia, Nepal, and Rwanda. For the other countries, this relationship is generally non-monotonic. At the high end, moving from a cold to a hot day in Myanmar increases GDP per capita by 0.8 percent after 5 years. At the other end in this sample of countries, moving from a cold to a hot day in Burundi reduces GDP per capita by 0.1 percent. As with the set of relatively rich countries, the patterns for the global temperatures are similar.

In summary, using this alternative temperature measure reinforces our finding that there is is substantial heterogeneity in the directions of the responses. Additionally, we likewise find that many of the rich countries are harmed by more high temperature days and many gain, including some of the poorest countries. By mapping out the distribution of the responses, we show that the signs of the responses of relative large movements from a cold to a hot day match, in most cases, the responses from average temperature movements.

Figure 10. Responses–G-7 Plus Australia and China at Horizon 5

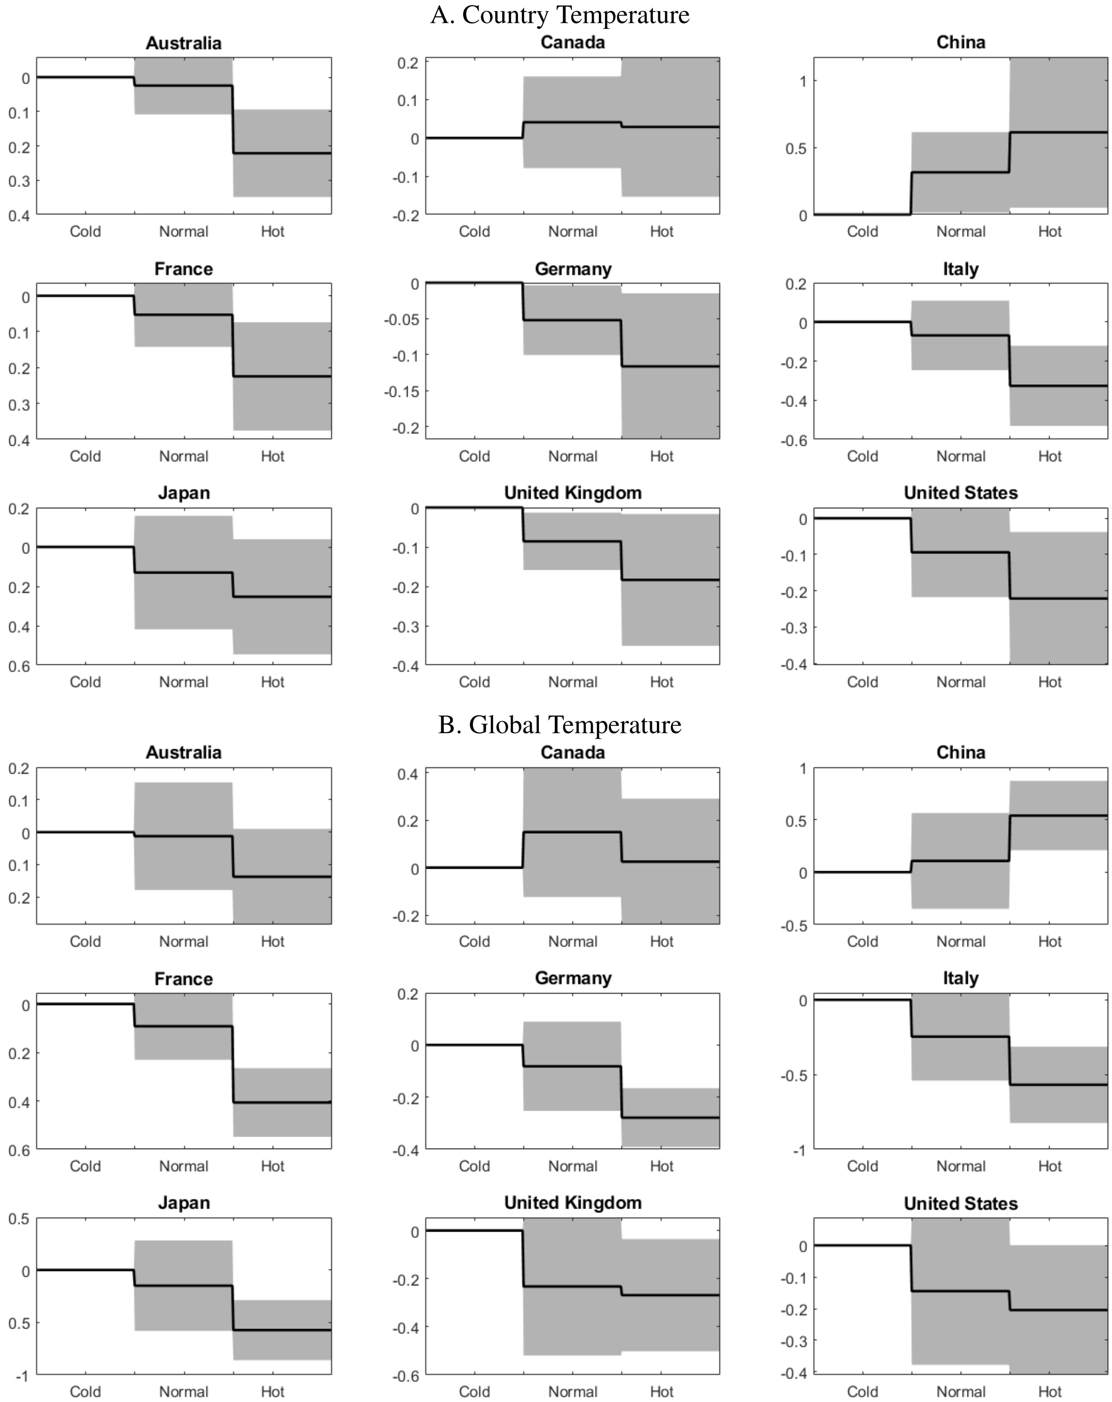

Notes: Estimates from equation (11) at horizon $h = 5$. Local projection betas are not standardized. Shaded areas are plus and minus 1.96 standard error bands. Specifications are determined by Akaike's Information Criterion (AIC).

Figure 11. Responses–Nine Poorest Countries at Horizon 5

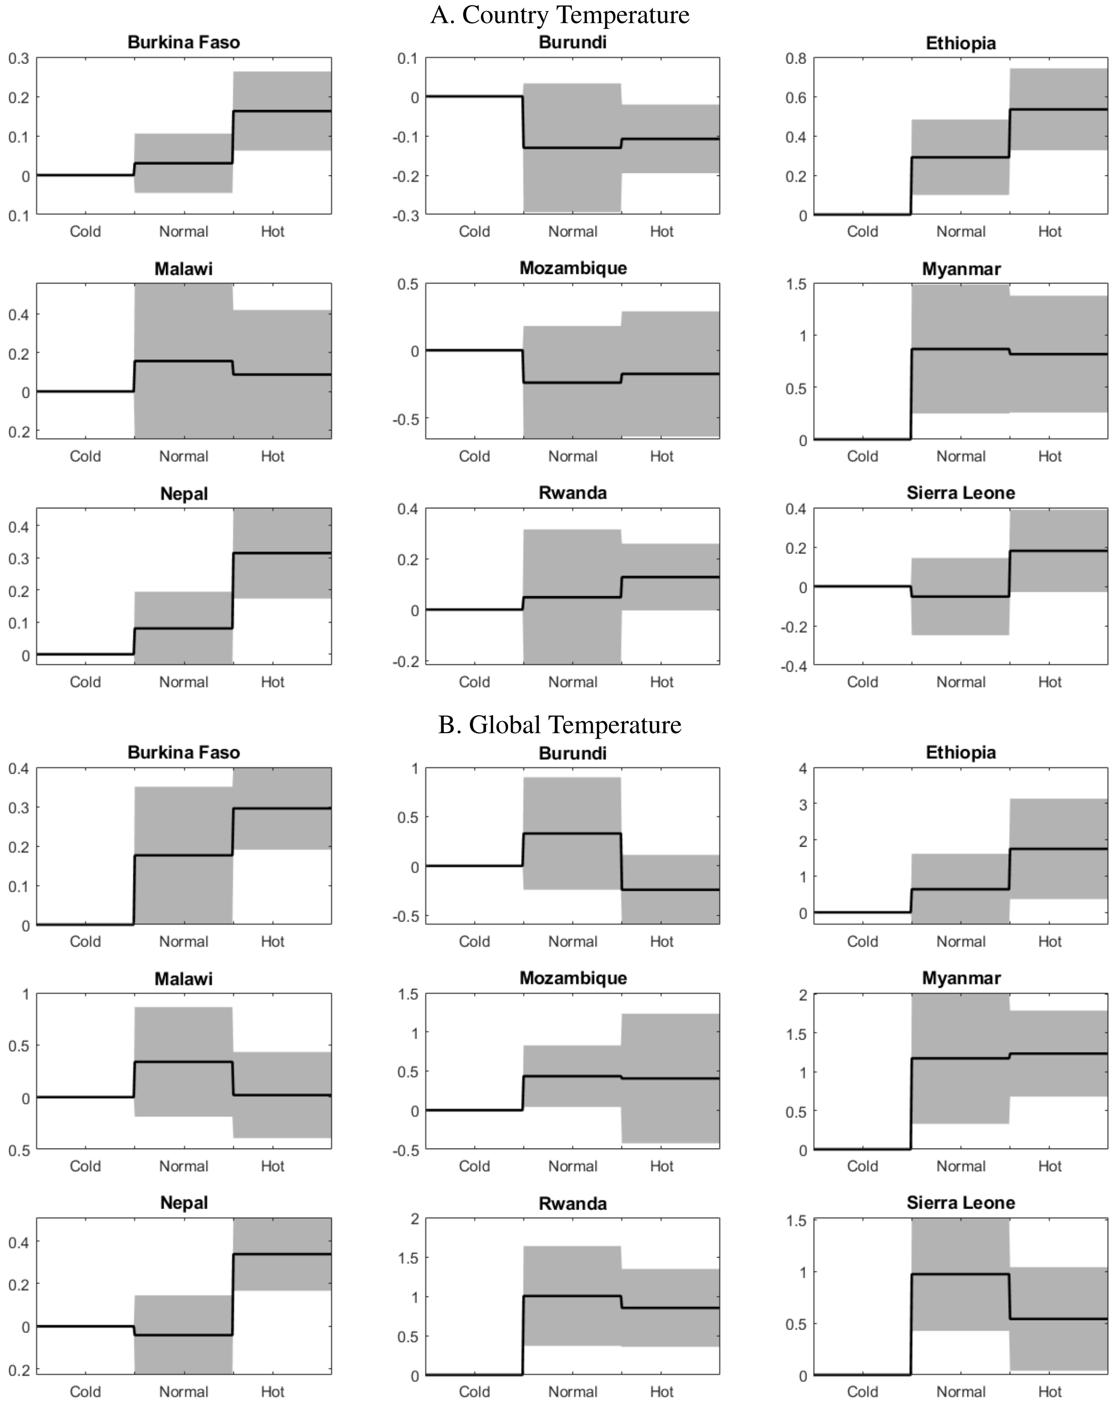

Notes: Estimates from equation (11) at horizon $h = 5$. Local projection betas are not standardized. Shaded areas are plus and minus 1.96 standard error bands. Specifications are determined by Akaike's Information Criterion (AIC).

6. Conclusion

This paper reexamines the relationship between rising temperature and real GDP per capita growth, but from a country-specific time series perspective using local projections (Jordà, 2005). We examine the growth responses from country, global, and idiosyncratic temperature variation. We find substantial heterogeneity across countries in the impulse responses of real GDP per capita growth to shocks to our temperature components–more than was previously reported in the literature. Qualitatively consistent with the previous literature though, when aggregating across countries to find a weighted world response, we find that responses to temperature increases are negative across all temperature shocks and horizons. Additionally, we find more countries have negative than positive impulse responses of real GDP per capita growth to increases in country and idiosyncratic temperature. On the other hand, it is more evenly split for the global temperature shocks. Richer countries, in particular, such as the United States, tend to experience negative impulse responses of real GDP per capita growth to increases in global temperature.

We find that growth responses are positive for many countries, including some of the poorest ones, from global temperature variation. These results come from the historical data. However, we do not assess the stability of these relationships moving forward. It would be highly speculative to think that the historical relationship between temperature and growth will continue in the future if global temperature rises 2^o-4^o Celsius. For this reason, we did not use our results to assess future damage.

Our analysis also investigates the country-level characteristics that might explain variation in the estimated growth responses to temperature change. These country characteristics did not explain growth responses to idiosyncratic temperature change. Responses to idiosyncratic temperature shocks are largely unsystematic. Variation in response to country and global temperature shocks are more systematically related to several country-level characteristics. For example, a country's GDP is more likely to respond positively to a global temperature shock if it is poorer, has grown less rapidly, is more open to trade, and more educated.

Our results may be helpful in framing climate change policy. As an ethical matter, Stern (2008) argues that rich countries should pay more for greenhouse gas abatement than developing countries, since the industrialized world has been responsible for emitting most of the current stock of greenhouse gasses. Beyond these ethical considerations, our findings that global temperature increases have resulted in significant economic damages to rich countries suggests that they have a self-interest in investing in abatement policies. If environmental policy is informed by historical relationships – and we show that direction of the growth responses are not uniform across countries – our results also suggest another challenge in forming a global consensus on future abatement strategies.

References

Acemoglu, D., Johnson, S., Robinson, J.A., 2002. Reversal of fortune: Geography and institutions in the making of the modern world income distribution. The Quarterly journal of economics 117, 1231–1294.

Bansal, R., Ochoa, M., 2011. Temperature, aggregate risk, and expected returns. NBER wp. 17575 .

Barreca, A., Clay, K., Deschênes, O., Greenstone, M., Shapiro, J.S., 2015. Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004. American Economic Review 105, 247–251.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., Shapiro, J.S., 2016. Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. Journal of Political Economy 124, 105–159.

Berg, K.A., Mark, N.C., 2022. Uncertainty, long-run, and monetary policy risks in a two-country macro model. mimeo, University of Notre Dame .

Bilal, A., Känzig, D.R., 2024. The macroeconomic impact of climate change: Global vs. local temperature. National Bureau of Economic Research Working Paper 32450.

Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. Nature 527, 235–239.

Byrne, J.P., Vitenu-Sackey, P.A., 2024. The macroeconomic impact of global and country-specific climate risk. Environmental and Resource Economics , 1–28.

Colacito, R., Hoffmann, B., Pham, T., 2019. Temperature and growth: A panel analysis of the united states. Journal of Money, Credit, and Banking 51.

Cruz, J.L., Rossi-Hansberg, E., 2024. The economic geography of global warming. Review of Economic Studies 91, 899–939.

Dell, M., Jones, B.F., Olken, B.A., 2009. Temperature and income: Reconciling new cross-sectional and panel estimates. American Economic Review 99, 198–204.

Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature shocks and economic growth: Evidence from the last half century. American Economic Journal: Macroeconomics 4, 66–95.

Deryugina, T., Hsiang, S.M., 2014. Does the environmenbt still matter? daily temperature and income in the united states. NBER wp. 20740 .

Deschênes, O., Greenstone, M., 2007. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. American Economic Review 97, 354–385.

Dietz, S., Lanz, B., 2019. Can a growing world be fed when the climate is changing? University of Neuchatel Institute of Economic Research IRENE WP 19-09 .

Easterly, W., Levine, R., 2003. Tropics, germs, and crops: how endowments influence economic development. Journal of monetary economics 50, 3–39.

Gandhi, S., Kahn, M.E., Kochhar, R., Lall, S., Tandel, V., 2022. Adapting to flood risk: Evidence from a panel of global cities. National Bureau of Economic Research Working Paper 30137.

Hale, G., 2022. Climate risks and exchange rates. Working Paper .

Hassler, J., Krusell, P., 2012. Economics and climate change: Integrated assessment in a multi-region world. Journal of the European Economic Association 10, 974–1000.

Henseler, M., Schumacher, I., 2019. The impact of weather on economic growth and its production factors. Climatic Change 154, 417–433.

Hsiang, S., 2016. Climate econometrics. Annual Review of Resource Economics 8, 43–75.

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larson, K., , Houser, T., 2017. Estimating economic damage from climate change in the united states. Science 356, 1362–1369.

Center for International Earth Science Information Network, C., 2018. Gridded population of the world, version 4 (gpwv4): Population count, revision 11. Columbia University .

Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95, 161–182.

Kahn, M.E., Mohaddes, K., Ng, R.N.C., Pesaran, M.H., Raissi, M., Yang, J.C., 2019. Long-term macroeconomic effects of climate change: A cross-country analysis. mimeo, Cambridge University .

Kim, H.S., Matthes, C., Phan, T., 2022. Severe weather and the macroeconomy. Working Paper .

Lee, S.O., Mark, N.C., Nauerz, J., Rawls, J., Wei, Z., 2022. Global temperature shocks and real exchange rates. Journal of Climate Finance 1.

Leta, M., Tol, R.S.J., 2019. Weather, climate and total factor productivity. Environmental and Resource Economics 73, 283–305.

Lustig, H., Richmond, R.J., 2020. Gravity in the exchange rate factor structure. Review of Financial Studies 33, 3492–3540.

Matsuura, K., Willmott, C., 2018. Terrestrial air temperature: 1900–2017 gridded monthly time series.

Nath, I.B., Ramey, V.A., Klenow, P.J., 2023. How much will global warming cool global growth? Working paper .

Nelson, G.C., van der Mensbrugghe, D., Ahammad, H., Blanc, E., Calvin, K., Hasegawa, T., Havlik, P., Heyhoe, E., Kyle, P., Lotze-Campen, H., von Lampe, M., d'Croz, D.M., van Meijl, H., Müller, C., Reilly, J., Robertson, R., Sands, R.D., Tabeau, C.S.A., Takahashi, K., Valin, H., Willenbockel, D., 2014. Agriculture and climate change in global scenerios: Why don't the models agree. Agricultural Economics 45, 85–101.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Nordhaus, W., Yang, Z., 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. American Economic Review 86, 741–765.

Park, J.Y., Phillips, P.C., 1988. Statistical inference in regressions with integrated processes: Part 1. Econometric Theory 4, 468–497.

Park, R.J., Goodman, J., Hurwitz, M., Smith, J., 2020. Heat and learning. American Economic Journal: Economic Policy 12, 306–339.

Plagborg-Møller, M., Wolf, C.K., 2021. Local projections and vars estimate the same impulse responses. Econometrica 89, 955–980.

Ricke, K., Drouet, L., Caldeira, K., Tavoni, M., 2018. Country-level social cost of carbon. Nature Climate Change 8, 894–900.

Rodrik, D., Subramanian, A., Trebbi, F., 2004. Institutions rule: the primacy of institutions over geography and integration in economic development. Journal of economic growth 9, 131–165.

Schultz, K.A., Mankin, J.S., 2019. Is temperature exogenous? the impact of civil conflict on the instrumental climate record in sub-saharan africa. American Journal of Political Science 63, 723–739.

Somanathan, E., Somanathan, R., Sudarshan, A., Tewari, M., 2021. The impact of temperature on productivity and labor supply: Evidence from indian manufacturing. Journal of Political Economy 129, 1797–1827.

Stern, N., 2008. The economics of climate change. American Economic Review: Papers & Proceedings 98, 1–37.

Tol, R.S., 2019. A social cost of carbon for (almost) every country. Energy Economics 83, 555–566.

West, K.D., 1988. Asymptotic normality, when regressors have a unit root. Econometrica , 1397–1417.

Zivin, J.G., Neidell, M., 2014. Temperature and the allocation of time: Implications for climate change. Journal of Labor Economics 32, 1–26.