Global Warming, Local Effects*

EMIR MALIKOV[1] JINGFANG ZHANG[2] RUIQING MIAO[3]

[1]University of Nevada, Las Vegas
[2]Alcorn State University
[3]Auburn University

June 21, 2026

Abstract

Empirical valuations of the macroeconomic damages from climate change typically impose a globally common temperature-output mapping, letting locations differ only through additive fixed effects. This restricts local heterogeneity to be purely intercept-shifting and hence climate-neutral (non-moderating), yet the factors that generate persistent differences in economic activity across locations are themselves plausible moderators of temperature sensitivity. When such nonneutral heterogeneity is correlated with within-location temperature variability, the globally constant slope may not recover a well-defined average marginal effect but a composite contaminated by heterogeneity bias. We instead adopt a geographically varying coefficient model that lets the temperature-economy relationship vary smoothly across space, delivering spatially contiguous, location-specific economic effects of warming that condition on local context beyond the usual baseline temperature. Estimated on a global municipality-level 1990–2022 panel of nearly 41,000 administrative units constructed from recent advances in machine-learning-based spatial downscaling and extrapolation of subnational income data, the model recovers a pronounced geographic bifurcation invisible to standard specifications: statistically negative effects of warming concentrate in the tropics and subtropics (41.4% of locations) and positive effects in high-latitude and high-altitude regions (9.8%). The globally constant specification eliminates this bifurcation and, in end-of-century projections, overstates aggregate global losses by a factor of two to three.

Keywords: Climate, Economic Growth, Geographic Heterogeneity, Temperature

JEL Classification: O44, Q54

*E-mail: emir.malikov@unlv.edu (Malikov), jzhang@alcorn.edu (Zhang), miaorong@auburn.edu (Miao).

1 Introduction

Quantifying the economic effects of climate change is among the most consequential empirical problems of today. Prominent estimates of aggregate damage from the projected rise in global temperatures by the end of century range from 1–3% of world GDP under calibrated level-effect models (Barrage and Nordhaus, 2024) to over 20% under empirical specifications that allow temperature to affect growth rates persistently (Burke et al., 2015; Bilal and Känzig, 2026). This order-of-magnitude disagreement reflects, among other things, different assumptions about the persistence and functional form of the temperature response, the source of identifying variation and, as we argue in this paper, the treatment of geographic heterogeneity in the temperature-economy relationship.[1] The dominant empirical approach in this literature relates per-capita GDP growth to temperature in a panel fixed-effects regression that exploits within-location variation for identification (e.g., Dell et al., 2012; Hsiang, 2010; Burke et al., 2015; Kalkuhl and Wenz, 2020; Kahn et al., 2021; Nath et al., 2024). A common and largely unexamined restriction embedded in this framework is that the temperature-economy relationship is governed by a single set of "globally" constant coefficients that apply identically to all locations in the sample. Under this restriction, unobserved cross-sectional heterogeneity enters only through additively separable unit fixed effects that shift intercepts but are assumed to be climate-neutral in the sense that they do not moderate the effect of temperature on economic activity. This paper shows that relaxing this restriction fundamentally changes the estimated magnitude and geography of climate damages.

This neutrality—or global homogeneity—assumption is particularly difficult to defend in a global climate-economy panel, where the factors that generate persistent differences in economic activity across locations (e.g., sectoral composition, institutional quality, adaptive capacity) are precisely the factors that would plausibly moderate temperature sensitivity. When such nonneutral local heterogeneity is correlated with the within-location variance of temperature, a globally constant coefficient model does not recover a well-defined average marginal effect but rather a composite that conflates heterogeneous temperature-economy mappings with the distribution of unobserved confounding moderators, producing what amounts to heterogeneity bias of the "random coefficients correlated with regressors" type (Chamberlain, 1982; Wooldridge, 2005). The resulting estimates may systematically mischaracterize both the magnitude and the geographic distribution of climate damages, overstating aggregate losses by averaging away the positive effects of warming in colder regions while inflating negative effects elsewhere.

In this paper, we employ a geographically varying coefficient framework for estimating the local economic effects of temperature that allows both the marginal effect of warming and its dynamic propagation to differ across space as smooth functions of geographic coordinates. Rather than estimating a single global temperature-economy relationship, we recover location-specific parameters for each of 40,997 admin-2-level geographic units (municipalities) worldwide over 1990–2022 using local kernel smoothing with adaptive nearest-neighbor bandwidth selection. The specification em-

[1] Also see Bilal and Stock (2025) for a recent review of the literature.

beds a distributed lag of temperature interacted with a location's long-run climatic baseline, enabling the data to distinguish between transitory level effects and persistent growth effects of temperature at each location separately. By formalizing the coefficient surfaces as unknown smooth functions of longitude and latitude, we impose only that nearby locations have similar temperature sensitivities, a restriction grounded in Tobler's first law of geography that "everything is related to everything else, but near things are more related than distant things," thereby encoding the inherent spatial structure of economic and climate data without restricting the relationship to be globally uniform.

The main advantages of our approach are therefore as follows. First, by allowing slopes to vary with geography we directly address heterogeneity bias, because the location-varying coefficients absorb the moderating role of unobserved local factors that would otherwise contaminate a globally constant estimate, thereby restoring the economic plausibility that the temperature-economy mapping is conditioned on local context.[2] Second, the approach delivers granular, location-specific measurements of climate sensitivity and damage projections rather than a single global average, providing the spatial resolution that adaptation planning requires. Third, the geographically varying specification nests the conventional globally constant model as a special case (obtained when the coefficient surfaces are flat), so the standard framework is directly testable against a flexible alternative. Fourth, by exploiting a comprehensive global admin-2-level panel constructed from recent advances in machine-learning-based spatial downscaling and extrapolation of subnational income data (Kummu et al., 2025), we bring the geographic granularity of the data into closer alignment with the spatial resolution of the econometric model, an alignment that country-level or even admin-1-level (provinces) datasets cannot provide.

Our empirical findings reveal a pronounced geographic bifurcation in the economic effects of temperature that is invisible to the globally constant coefficient specification. Accounting for local context via geographically varying parameters, 41.4% of admin-2 locations experience a statistically negative cumulative marginal effect of warming while 9.8% experience a statistically positive effect, with the remainder statistically indistinguishable from zero. Statistically negative effects concentrate in the tropics and subtropics, while positive effects cluster in colder high-latitude and high-altitude regions. This spatial pattern reflects heterogeneity in the coefficient surfaces themselves and not merely differences in temperature baselines. As a matter of fact, the conventional globally constant coefficient model, which also conditions marginal effects on local average temperature, produces a near-uniformly negative damage map that eliminates the bifurcation entirely. When translated into end-of-century income projections, even under the severe SSP5-8.5 scenario (high-emission fossil-fuel-intensive pathway with little additional mitigation) of the CanESM5 General Circulation Model that produces comparatively warmer projections, our geographically varying specification projects aggregate global losses of only about 25% (95% CI: –49%, –2%), whereas the globally homogeneous specification projects income losses of approximately 73% (95% CI: –96%, –50%), an overstatement by roughly a factor of three driven primarily by its inability to detect positive temperature effects in

[2] Albeit some of these local moderating factors may be endogenous with respect to economic activity, a geographic unit's location which we effectively use as their proxy is however immune to these concerns given its intrinsic strict exogeneity.

drier, colder and continental regions. This divergence persists across alternative warming scenarios and climate model projections.

Our paper relates most directly to the empirical climate-macroeconomy literature built on panel fixed-effects identification using variation in weather or climate innovations (Dell et al., 2012; Burke et al., 2015; Burke and Tanutama, 2019; Kahn et al., 2021; Nath et al., 2024).[3] These studies differ in their treatment of dynamics, nonlinearity and identifying variation, but their main specifications remain globally parameterized: non-temperature-dependent heterogeneity in coefficients, when considered, is introduced through interactions with observed states, coarse group splits or other forms of sample splitting. Dell et al. (2012) exploit within-country temperature fluctuations and allow the linear temperature effect to differ between poor and rich countries, finding large negative growth effects only in the former. Imposing a quadratic response function, Burke et al. (2015) also examine rich–poor heterogeneity but do not reject a common response function, which Burke et al. (2018) subsequently adopt to quantify the global economic damages avoidable under the UN mitigation targets. Burke and Tanutama (2019) study heterogeneity through income-quintile and country-specific sample subsetting. Kahn et al. (2021) model deviations of temperature from historical norms and test whether effects differ across rich–poor and hot–cold countries, concluding that they do not. Nath et al. (2024) estimate growth effects of temperature shocks and supplement the pooled model with country-by-country estimates, concluding that the latter are very similar to the state-dependent pooled model. Our spatial approach is distinct from these forms of parameter heterogeneity analysis. We allow the temperature-economy mapping itself to vary smoothly over geography and show that the parameter constancy restriction embedded in globally parameterized models can be an important source of heterogeneity bias in aggregate damage estimates.

In emphasizing the consequences of a particular modeling restriction, our paper also adds to an emerging strand of the literature concerned with how specification choices influence estimated economic impacts of climate change. Beyond Newell et al.'s (2021) demonstration that damage projections are sensitive to the choice among largely observationally equivalent GDP-temperature models, Cui et al. (2024) show that conventional model selection criteria can fail to identify the specification best suited for climate impact projections because, when all candidate response functions are imperfect approximations, a model that fits the outcome-weather relationship well under the historical climate need not continue to do so under the projected one.[4] We contribute to this line of inquiry by isolating a complementary specification margin that has received little scrutiny in the global climate-macroeconomy context, namely, the presumed spatial invariance of the temperature

[3] Bilal and Känzig (2026) are a standout in departing from the popular local weather identification strategy by exploiting global temperature variation. Although they examine parameter heterogeneity across regional and income-based groups, their framework continues to assume that those parameters are common across units within a group.

[4] A related specification margin concerns the very set of climate indicators entering the damage function. Waidelich et al. (2024) project damages from shifts not only in annual temperatures but also in temperature variability, rainfall patterns and extreme precipitation, finding the additional damages from the latter indicators to be comparatively small, although estimating the temperature dose-response function with these indicators included does add to projected global losses and widens the downside tail of the damage distribution. The quantitative weight of any given specification margin is thus, ultimately, an empirical matter.

response parameters.

A related strand of research examines geographic heterogeneity in climate effects on agriculture. Keane and Neal (2020) allow for cross-sectional—albeit, not "spatial" per se—and temporal slope heterogeneity in U.S. crop yield regressions, and Zhang et al. (2026) generalize the long-differences approach of Burke and Emerick (2016) by letting climate-agriculture coefficients vary over space explicitly, documenting a geographic divergence in projected crop yield changes in the U.S. due to warming and showing that spatially invariant specifications overestimate aggregate damages. In a related agricultural setting, Goh et al. (2025) develop a Bayesian hierarchical model with spatiotemporal Gaussian process priors that permits the effects of weather on U.S. corn yields to vary over space, documenting economically meaningful heterogeneity in these sensitivities too. Their approach shares with ours the premise that climate sensitivity parameters are spatially structured. We extend this logic from a single agricultural sector and a single country to the economy-wide temperature-growth relationship in a global panel. The subnational data dimension of our analysis is shared with Burke and Tanutama (2019), who assemble district-level GDP data across 37 countries and more than 11,000 districts, with Colacito et al. (2019), who document significant effects of seasonal temperatures on U.S. state-level output growth across a wide cross section of industries, with Hsiang et al. (2017), who value market and nonmarket climate damages in the U.S. at the county level by combining sector-specific dose-response functions with probabilistic climate projections to find losses concentrated in poorer and hotter southern counties, and with Kalkuhl and Wenz (2020), who use a panel of roughly 1,500 subnational regions from 77 countries to find that incomes respond to both annual temperature shocks as well as long-run temperature levels. These studies underscore the value of subnational economic data but maintain a globally parameterized framework. Our contribution is to combine subnational geographic granularity with a spatially flexible econometric specification that allows the temperature-economy relationship itself to vary across locations.

The rest of the paper unfolds as follows. Section 2 develops the geographically varying coefficient model, describes the local kernel estimation procedure and discusses the econometric motivation for allowing slopes to vary by geography. Section 3 describes the data. Section 4 presents the empirical results, including the estimated historical temperature-economy relationship and end-of-century income projections under alternative warming scenarios. Section 5 concludes.

2 Empirical Approach

To evaluate—and project—economic effects of a warming climate, we estimate the historical relationship between temperature and economic output across the globe by relating per-capita economic performance to local temperatures. To accomplish that, we employ a spatially varying model that allows the temperature-economy relationship to vary flexibly across space, enabling us to estimate it "locally" for each spatial unit.

Concretely, we assemble a global panel of economic output and climate variables spanning 1990–2022 for geographic units at the administrative level 2 (admin-2) across essentially the entire world.

The sample contains 40,997 unique locations. Each location i is indexed by a fixed geographic coordinate $\ell_i = (\text{lon}_i, \text{lat}_i)$, taken as the longitude-latitude pair of the admin-2 centroid. Let y_{it} denote real GDP per capita in location i and year t. Our outcome of interest is local annual per-capita GDP growth, $\Delta \ln y_{it} = \ln y_{it} - \ln y_{i,t-1}$. The main climate regressor of interest is the yearly average of the daily maximum temperature in location i and year t, T_{it}. Related to it also is the "long run" mean for each location i, $\overline{T}_i = t_{\max}^{-1} \sum_t T_{it}$, which characterizes local historical temperature profile, facilitating the differentiation between generally colder versus generally warmer locations. As is conventional, we also control for total annual precipitation P_{it}. Thus, letting climate sensitivity parameters not be constant but vary smoothly with ℓ_i, we have the following semiparametric temperature-economy relationship equation:

$$\Delta \ln y_{it} = \sum_{j=0}^{J} \left[\beta_j(\ell_i) + \gamma_j(\ell_i) \overline{T}_i \right] T_{it-j} + \delta(\ell_i) P_{it} + \alpha(\ell_i) t + \mu_i + \epsilon_{it} \quad \forall i = 1, \ldots, n; \; t = 1, \ldots, t_{\max}, \tag{2.1}$$

which includes the standard location fixed effects μ_i that absorb time-invariant confounders in local growth potential, thus helping control for persistent cross-space differences therein, as well as controls for local time trends in growth via $\alpha(\ell_i)t$. The key feature of (2.1) is that the nonparametric functions $\{\beta_j(\cdot), \gamma_j(\cdot); \; j = 0, \ldots, J\}, \delta(\cdot), \alpha(\cdot)$ are unknown smooth coefficient surfaces defined over geographic space; we revisit this shortly.

Specification (2.1) allows for dynamic temperature effects on economic activity and also permits those effects to depend on the temperature itself. Consider the dynamics first. Because the outcome variable in our analysis is the growth rate, a static model with only contemporaneous temperature terms ($J = 0$) would impose a pure *growth effect* mechanics onto the temperature-economy relation. However, temperature may instead (or also) have *level effects* that wash out as the economy bounces back over time, producing only short-run impacts on growth with subsequent reversal. As Dell et al. (2012) show using a stylized production framework, in a dynamic model relating economic growth rate to temperature a pure level effect shows up as a contemporaneous impact on growth $\beta_0(\ell_i) + \gamma_0(\ell_i)\overline{T}_i$ that is later offset by an approximately equal and opposite lagged effect $\sum_{j=1}^{J} \big(\beta_j(\ell_i) + \gamma_j(\ell_i)\overline{T}_i\big)$—so the temporary growth effect reverses as the income level effect reverts—netting a zero total effect when summed over periods $\sum_{j=0}^{J} \big(\beta_j(\ell_i) + \gamma_j(\ell_i)\overline{T}_i\big)$. On the other hand, a genuine growth effect does not reverse in the same way, in which case dynamic effects do not sum to zero, and the temperature shock has a lasting effect on the growth rate of economy, permanently affecting incomes in the long run. Since both temperature and GDP growth are persistent, omitting lags from the analysis can "create the appearance of growth effects" even if the true effect is primarily transitory (see Nath et al., 2024). These arguments therefore motivate our dynamic specification in (2.1) which includes a distributed lag of temperature. This prevents the model from hard-wiring a pure-growth interpretation as the data can deliver either reversal dynamics consistent with level effects or persistence consistent with growth effects.

Of note, the dynamic effects of temperature on local economy measurable from model (2.1) are conditional on local climatic baseline captured by $\overline{T}_i$. A standard approach to introduce such nonlinearity in temperature is through a quadratic polynomial, wherefrom the marginal effect of temperature varies with its contemporaneous realization (e.g., Burke et al., 2015). Instead, following Nath et al. (2024) we have temperature effects depend on the location's long-run temperature environment via interaction of $(T_{it}, T_{it-1}, \ldots, T_{it-J})$ with location-specific mean $\overline{T}_i$. This intuitively reinterprets nonlinearity in the temperature-economy relationship as state dependence, whereby given period's temperature changes may have heterogeneous effects on economic growth depending on the location's mean historical temperature. In this context, $\overline{T}_i$ provides a transparent baseline state variable that is predetermined with respect to year-to-year temperature fluctuations and is naturally interpretable as "climatic normalcy"[5] in a given geographic unit.

Most notably, in letting all parameters in (2.1) vary across space by formalizing them as unknown smooth functions of ℓ_i our specification generalizes the standard constant-coefficient "global" model of the temperature-economy relationship in an important respect. Specifically, it permits both the marginal effect of temperature and its dynamic propagation to differ across space all while respecting naturally occurring spatial congruence in economic and climate data, whereby geographic units are permitted to have heterogeneous temperature sensitivity parameters albeit not in an arbitrary, discontinuous fashion but in a disciplined manner that encodes the inherent "nearby places are more similar" feature of spatial data into the analysis by treating the estimated temperature-economy relationship as spatially continuous. Further, by explicitly incorporating geographic information in our model, not only do we obtain location-specific parameters but, through that, we also control for local heterogeneity that need not be additively separable and may *moderate* temperature effects on economic activity. The latter contrasts with the conventional constant-coefficient framework, which restricts that local heterogeneity—captured by unit fixed effects—is purely additive and occurs only through intercepts.

Indeed, an intercept shifter μ_i is meant to absorb any time-invariant determinants of the ith location's average income growth (including local geography, long-run institutional quality, historical development patterns and other differences across locations) that do not interact with climate. Controlling for such climate-neutral growth determinants via fixed effects is imperative but it may be insufficient for credible identification because it imposes a strong separability restriction, whereby unobserved local heterogeneity is allowed to shift the intercept but is assumed to be climate-neutral in the sense that it does not influence the marginal effect of temperature on economy. This neutrality restriction is particularly difficult to defend in a global climate-economy panel. Many of the reasons why locations differ economically are precisely the reasons why their economies would likely respond differently to temperature shocks. In other words, the same spatial factors that generate persistent differences in the rate of economic growth are plausible moderators of climate sensitivity. When those moderators are assumed away, a constant-coefficient model effectively attempts to average heterogeneous marginal effects into a single global coefficient. The question is whether the resulting, globally constant parameter meaningfully corresponds to any representative location. It

[5] A long-run temperature average also serves as a climatic baseline in Bento et al. (2023), who enter it additively alongside a weather deviation to separate long- from short-run responses. Our $\overline{T}_i$ instead enters multiplicatively, through interactions, because we target state dependence.

is unlikely so as it averages location-specific temperature effects with the weights that are induced by the joint distribution of temperature realizations and the unobserved moderating heterogeneity. That is, so long as nonneutral local heterogeneity is correlated with within-location temperature (the most plausible scenario), the global slope is generally not interpretable as even a well-defined average causal response.

The above is the central limitation of the fixed-effects-only approach that treats spatial heterogeneity as additively separable. The econometric consequence is, essentially, heterogeneity bias of the familiar "random coefficients correlated with regressors" type studied by Chamberlain (1982) and Wooldridge (2005).[6] To see precisely how this bias manifests, suppose the data generating process (DGP) admits location-specific temperature effects, $\Delta \ln y_{it} = b_i T_{it} + \mu_i + e_{it}$, but the estimated is a constant-slope specification, $\Delta \ln y_{it} = \beta T_{it} + \mu_i + u_{it}$ with $u_{it} = (b_i - \beta) T_{it} + e_{it}$, in which purportedly $\beta = E[b_i]$. The standard within (FE) estimator of β, however, has a probability limit that need not be $E[b_i]$ but, more plausibly, is a variance-weighted functional of the slope distribution.

Result 1 (Heterogeneity bias of the globally constant estimator) *Suppose the DGP is* $\Delta \ln y_{it} = b_i T_{it} + \mu_i + e_{it}$*, with* $\{(b_i, \mu_i, \{T_{it}, e_{it}\}_{t=1}^{t_{\max}})\}_{i=1}^{n}$ *i.i.d. across i,* $E[e_{it} | T_i, b_i, \mu_i] = 0$ *with* $T_i = (T_{i1}, \ldots, T_{it_{\max}})'$*, and* $E[s_{T,i}^2] > 0$*, where* $s_{T,i}^2 \equiv t_{\max}^{-1} \sum_t (T_{it} - \overline{T}_i)^2$ *is the within-i sample variance of temperature. Let* $\widehat{\beta}_{\mathrm{FE}}$ *denote the pooled OLS estimator on within-transformed data obtained under the constant-slope model* $\Delta \ln y_{it} = \beta T_{it} + \mu_i + u_{it}$*. As* $n \to \infty$ *with* $t_{\max}$ *fixed,*

$$\widehat{\beta}_{\mathrm{FE}} \xrightarrow{p} \frac{E[b_i s_{T,i}^2]}{E[s_{T,i}^2]} = E[b_i] + \frac{\mathrm{Cov}[b_i, s_{T,i}^2]}{E[s_{T,i}^2]},$$

so that $\widehat{\beta}_{\mathrm{FE}}$ *is a variance-weighted average of the location-specific slopes and is consistent for* $E[b_i]$ *if and only if* $\mathrm{Cov}[b_i, s_{T,i}^2] = 0$.

The result follows directly from the within-transformed OLS projection and is by no means novel. It is still useful because it clarifies the estimand implicitly recovered by the constant-coefficient FE estimator, whereby it does not generally target the unweighted average temperature effect $E[b_i]$ but rather a reweighting of the distribution of these effects by the cross-location distribution of within-location temperature variability, unless $\mathrm{Cov}[b_i, s_{T,i}^2] = 0$. While the latter is necessary as well as sufficient condition for a globally constant coefficient model to identify the average temperature sensitivity, it is difficult to defend in a global climate-economy panel. The structural moderators of b_i (sectoral composition, adaptive capacity, institutional quality) and the physical determinants of $s_{T,i}^2$ (continentality, latitude, climatic regime) are both spatially clustered, and there is no *a priori* economic or physical reason for the two clusterings to align such as to deliver a zero covariance.

Result 1 also implies that, had the analysis been univariate, the sign of the bias could be read off the sign of $\mathrm{Cov}[b_i, s_{T,i}^2]$, giving it direct empirical content. But in practice, the temperature-

[6] See also Pesaran and Smith (1995) for the dynamic heterogeneous panel analog and Graham and Powell (2012) for the average-partial-effect identification angle.

economy specification (2.1) is multivariate. Letting $\boldsymbol{x}_{it}$ collect all regressors and $\boldsymbol{b}_i$ be the corresponding vector of location-specific coefficients, the FE probability limit trivially generalizes to $\widehat{\boldsymbol{\beta}}_{\mathrm{FE}} \xrightarrow{p} \left(E[\boldsymbol{s}_i]\right)^{-1} E[\boldsymbol{s}_i \boldsymbol{b}_i] = E[\boldsymbol{b}_i] + \left(E[\boldsymbol{s}_i]\right)^{-1} E[\boldsymbol{s}_i (\boldsymbol{b}_i - E[\boldsymbol{b}_i])]$ with $\boldsymbol{s}_i \equiv t_{\max}^{-1} \sum_t (\boldsymbol{x}_{it} - \overline{\boldsymbol{x}}_i)(\boldsymbol{x}_{it} - \overline{\boldsymbol{x}}_i)'$, which is a matrix-weighted average of the heterogeneous slope vectors with weights given by the within-i second-moment matrix. Evidently, unlike in the univariate case, the matrix-weighted average of $\boldsymbol{b}_i$ does not admit a transparent sign rule. The bias in a multivariate constant-coefficient specification depends on the joint distribution of $\boldsymbol{b}_i$ with the entire within-i second-moment matrix $\boldsymbol{s}_i$ of regressors that include contemporaneous and lagged temperature, their interactions with the local baseline state, precipitation and the local time trend. Signing this object would require strong assumptions on the joint cross-location moments of $(\boldsymbol{b}_i, \boldsymbol{s}_i)$ that would be difficult to defend. We therefore treat the bias as an empirical object and read its direction off our estimates of historical temperature-economy relationship. The estimates in Table 1 provide the operative comparison, where we find that the constant-coefficient model returns a more negative mean temperature effect (by about 80% larger in the magnitude) than that obtained from a geographically varying model with heterogeneous coefficients, suggesting negative bias.

Allowing slopes to vary by geography is therefore not merely a "flexibility" add-on that delivers location-specific measurements of the economic effect of temperature. It is a way of restoring the economic plausibility that the temperature-economy mapping is conditioned on local context, while mitigating heterogeneity bias by modeling moderation explicitly. Since unobserved local moderators of climate sensitivity are spatially structured, treating coefficients as smooth functions of geographic coordinates provides a practical reduced-form way to condition on those moderators without observing them. The conceptualization here is not that each location has an unrelated coefficient. Rather, it is that the regressors and their moderators—and therefore the marginal effects—evolve gradually over space, generating natural spatial clustering. Imposing that nearby locations have similar parameters is a disciplined way to allow heterogeneity while avoiding the instability of estimating a separate slope for each location.

To summarize, for each admin-2 location our model delivers a local impulse-response-type mapping from temperature variations to per-capita income growth, with the mapping allowed to vary smoothly across geography. This is directly useful for damage accounting as it provides granular estimates of costs or benefits of warming that account for local climate-nonneutral contexts.

2.1 Local kernel estimation

We estimate (2.1) in within-transformed form. Because all of the coefficient surfaces in (2.1) are functions of the time-invariant location ℓ_i alone, subtracting the location-specific time mean from both sides of (2.1) eliminates the additive fixed effect μ_i without altering the coefficient structure. Letting $\widetilde{x}_{it} \equiv x_{it} - \overline{x}_i$ denote the within-transformation of x_{it}, with $\overline{x}_i$ its time mean at location i, the estimating equation reads

$$\widetilde{\Delta \ln y}_{it} = \sum_{j=0}^{J} \left[\beta_j(\ell_i) + \gamma_j(\ell_i)\overline{T}_i \right] \widetilde{T}_{it-j} + \delta(\ell_i)\widetilde{P}_{it} + \alpha(\ell_i)\widetilde{t} + \widetilde{\epsilon}_{it}, \tag{2.2}$$

in which $\overline{T}_i$ enters as an observed location-specific constant rescaling the demeaned temperature regressors. Additive time-invariant heterogeneity μ_i is thereby controlled for "nonparametrically" prior to smoothing.

Stack the demeaned regressors at observation (i,t) into $\widetilde{\boldsymbol{x}}_{it} = \left(\widetilde{T}_{it}, \overline{T}_i\widetilde{T}_{it}, \ldots, \widetilde{T}_{it-J}, \overline{T}_i\widetilde{T}_{it-J}, \widetilde{P}_{it}, \widetilde{t}\right)'$ and collect the parameters at a target location ℓ into $\boldsymbol{\theta}(\ell) = \left(\beta_0(\ell), \gamma_0(\ell), \ldots, \beta_J(\ell), \gamma_J(\ell), \delta(\ell), \alpha(\ell)\right)'$. Under the maintained smoothness of $\boldsymbol{\theta}(\cdot)$ in ℓ, the local-constant approximation $\boldsymbol{\theta}(\ell_i) \approx \boldsymbol{\theta}(\ell)$ holds for ℓ_i in a neighborhood of ℓ, and $\boldsymbol{\theta}(\ell)$ is recovered by the locally weighted least squares solution

$$\widehat{\boldsymbol{\theta}}(\ell) = \left[\sum_{i,t} \omega_i(\ell) \widetilde{\boldsymbol{x}}_{it} \widetilde{\boldsymbol{x}}'_{it} \right]^{-1} \sum_{i,t} \omega_i(\ell) \widetilde{\boldsymbol{x}}_{it} \widetilde{\Delta \ln y}_{it}, \tag{2.3}$$

where $\omega_i(\ell)$ is a non-negative kernel weight that decays with the geographic distance $\|\ell_i - \ell\|$, weighing each location-year observation based on how close the geographic coordinates of the ith location, $\ell_i = (\text{lon}_i, \text{lat}_i)$, are to those of the location ℓ. The estimator, thus, pools information from geographically proximate locations under the smoothness restriction, with all coefficient surfaces $\boldsymbol{\theta}(\ell)$ recovered at each ℓ from a correspondingly weighted local regression on the within-transformed data.

The speed with which kernel weights decay with the geographic distance between each ℓ_i and the target location ℓ is controlled by the bandwidth, the choice of which governs the effective sample size of each local regression. A fixed-radius bandwidth is ill-suited to a global admin-2-level panel because the spatial density of admin-2 units varies sharply across the globe, with oceans, deserts and sparsely populated regions inducing pronounced gaps in data coverage. A single global radius would over-smooth in densely subdivided areas while leaving coastal and island locations with too few admissible neighbors, since fixed geographic disks anchored at such locations contain no observations on the maritime side. To accommodate this irregularity, we employ an adaptive K-nearest-neighbor bandwidth. For each target location ℓ, the bandwidth is set to the Euclidean distance from ℓ to its Kth nearest sample coordinate:

$$r_K(\ell) = \left\| \ell_{(K)} - \ell \right\|, \tag{2.4}$$

where $\ell_{(K)}$ denotes the Kth nearest neighbor of ℓ among $\{\ell_i\}_{i=1}^n$. The bandwidth $r_K(\ell)$ is just the Kth order statistic on the distances $\|\ell_i - \ell\|$, and it is specific to a target location ℓ. The kernel weight entering (2.3) is then

$$\omega_i(\ell) = \mathcal{K}\left(\frac{\|\ell_i - \ell\|}{r_K(\ell)} \right) \quad \forall i = 1, \ldots, n, \tag{2.5}$$

with $\mathcal{K}(\cdot)$ a second-order Gaussian kernel. By construction, every local regression draws on a constant "effective" sample of K land-based locations regardless of coastal or island geography, and the bandwidth $r_K(\ell)$ adapts automatically to the local density of admin-2 units. Of note, although the location coordinates ℓ_i are bivariate, $r_K(\ell)$ modulates a univariate quantity, namely, the geographic distance $\|\ell_i - \ell\|$. This is because, different from the standard kernel fitting (using fixed bandwidth) in which data observations would have been weighted using the product of univariate kernels corresponding to each one of the two elements in $\ell_i - \ell$, the adaptive kernel fitting weights data using a norm thereof $\|\ell_i - \ell\|$. For this reason, when employing adaptive nearest-neighbor methods, smoothing variables are typically rescaled prior to the analysis to ensure comparability because the nearest-neighbor ordering is not scale-invariant in a multivariate case. In our application, however, we do not need to rescale coordinates ℓ_i (i.e., latitude and longitude) because they are already measured on the same cardinal scale, and the resulting Euclidean distance retains a direct physical interpretation. The optimal smoothing parameter K is selected by leave-one-location-out cross-validation.

3 Data

The subnational data on average income come from Kummu et al. (2025). It is a comprehensive, publicly available gridded per-capita GDP dataset (adjusted for Purchasing Power Parity) downscaled to the admin-2 level (about 43 thousand units) for 1990–2022.[7] This dataset is, to our knowledge, the only publicly available source providing subnational GDP per capita at admin-2 granularity with global coverage and annual frequency over a panel of this length. The admin-2 resolution is essential for our geographically varying coefficient approach, since country-level or even admin-1-level data would be too coarse to identify the local, yet spatially smooth, heterogeneity in temperature sensitivity that is central to our analysis. The dataset is constructed via spatial downscaling of reported subnational GDP data from official sources such as the OECD, Eurostat and national statistics agencies across 89 countries and 2,708 provincial units, using boosted ensemble machine learning algorithms trained on auxiliary covariates, including urbanization levels, travel time to nearest city and income inequality, with cross-validated $R^2 = 0.79$ and test-set $R^2 = 0.80$ against reported subnational data (Kummu et al., 2025). This type of downscaling and extrapolation methodology has been becoming increasingly standard in the climate economics literature for subnational and regional analyses. We acknowledge that any downscaled income measure is a modeled estimate and thus noisier than direct national accounts data, but this is the cost of the spatial granularity our analysis requires.

Furthermore, although classical measurement error in the outcome would not bias the slope-function estimators of interest but only inflate their variance, the more pertinent concern is that the downscaling imputation error in the log-GDP per capita $\ln y_{it}$ may not be classical to begin with. Specifically, the auxiliary predictors facilitating Kummu et al.'s (2025) downscaling are themselves spatially smooth and plausibly correlated with temperature in levels, so the measurement error may inherit both features. This concern is, however, largely neutralized by the identification structure of our specification. Because $\beta_j(\ell_i)$ and $\gamma_j(\ell_i)$ are pinned down by within-location, year-to-year fluctua-

[7] The dataset is publicly available at the following repository: https://doi.org/10.5281/zenodo.10976733.

tions in temperature, the relevant question is not whether the downscaling error has these features in levels but whether they survive into the residual variation that actually identifies our slope functions. Of note, the downscaling predictors are essentially time-invariant or very slow-moving at the annual frequency over a 33-year horizon, so the downscaling error is to inherit the same low-frequency profile. With the outcome in our analysis itself constructed as a first difference of the downscaled log-income series, the persistent component of the error is already attenuated before any estimation. The within transformation by μ_i and the detrending by $\alpha(\ell_i)t$ then eliminate what remains of the slow-moving component, leaving as a potential contaminant only its high-frequency error, which by construction has no counterpart in the cross-sectional, slow-moving downscaling predictors. For this surviving high-frequency component to bias our slope estimates, it would have to co-vary with weather-like, idiosyncratic deviations of T_{it-j} from the local trend within a location, a structure that has no obvious source in the downscaling. What ultimately survives is, at most, a residual noise component that inflates the variance of our estimates and renders the resulting inference conservative.

Panel (a) of Figure 1 plots average historical annual growth rate of per-capita GDP in our data. Upon merging these with weather data, we have a global panel of $n = 40,997$ admin-2 units across 2,722 admin-1 provinces. As can be seen from the figure, the global coverage is virtually complete.[8] Next, we describe the climate data used in the analysis, beginning with the historical weather variables and then turning to the projected climate data.

Daily maximum temperature and monthly total precipitation over 1990–2022 are obtained at a resolution of 6 arc-minutes from the ERA5-Land dataset developed by the European Centre for Medium-Range Weather Forecasts. We use daily *maximum* temperature rather than daily mean temperature because daily maxima better capture heat extremes that constitute the primary channel through which temperature affects economic activity, including nonlinear declines in labor productivity, total factor productivity, agricultural yields, human health and mortality (e.g., Lobell and Field, 2007; Graff Zivin and Neidell, 2014; Barreca et al., 2016; Zhang et al., 2018a; Callahan and Mankin, 2022; Zhang et al., 2024). We first obtain the annual average of daily maximum temperature and annual total precipitation of each grid cell within each country, and then aggregate the grid-level weather variables up to the admin-2 level as specified in the per-capita GDP dataset. The grid-level temperature and precipitation are weighted by the grid-level population count in 1989. Population-weighting ensures that the climate variables reflect the temperature conditions experienced by the economically active population rather than those of uninhabited terrain such as deserts, high-altitude ranges or ice-covered regions, where temperature variation is climatically informative but less relevant economically. We choose to use population data from the year 1989 because it precedes our sample period and, thus, is plausibly predetermined with respect to local income shocks, mitigating potential endogeneity concerns about economic growth influencing population distribution and trends. Historic grid-level population data are obtained from the Copernicus Land Change

[8]Admin-2 units that have not been merged are in pure island nations/territories, islands/overseas territories and a few coastal municipalities.

Lab at the Utrecht University.[9] Panel (b) of Figure 1 visualizes the geographic distribution of cross-year mean of average annual daily maximum temperatures, $\{\overline{T}_i\}_{i=1}^n$, in our historical data.

The projected future temperature data are obtained at the 15-arc-minute grid level from the Earth Exchange Global Daily Downscaled Projections (NEX-GDDP-CMIP6) dataset compiled by the U.S. National Aeronautics and Space Administration that harmonizes projection results from the Coupled Model Intercomparison Project Phase 6 (CMIP6). Specifically, we obtain end-of-century daily maximum temperatures in 2096–2100 projected by General Circulation Models (GCMs) CanESM5 and NorESM2-MM. These two GCMs are chosen to bracket a wide range of plausible warming futures: CanESM5 has among the highest equilibrium climate sensitivity (ECS) in the CMIP6 ensemble, producing comparatively warm projections, while NorESM2-MM has a more moderate ECS, yielding cooler projections (Hausfather, 2019). Together, the pair spans much of the inter-model uncertainty in projected warming magnitudes. For each one of the two GCMs, we obtain the projected climate data under two future warming scenarios: Shared Socioeconomic Pathways (SSPs) 2-4.5 and 5-8.5. Spanning a policy-relevant range of future radiative forcing, SSP2-4.5 represents a moderate-emission pathway consistent with intermediate mitigation effort, whereas SSP5-8.5 represents a high-emission, fossil-fuel-intensive pathway with little additional mitigation. The daily grid-level weather variables are first used to calculate the grid-level annual average maximum temperature for each grid cell, which are then aggregated up to the admin-2 level using the same population weighting as that used for the historical weather variables.

Finally, we use projected annual population data for 2096-2100 when computing future aggregate economic impacts of warming between the end of our historical sample and the end of century. These data (under SSPs 2 and 5) are obtained at the 2.5-arc-minute grid level from the Inter-Sectoral Impact Model Intercomparison Project.[10] We then aggregate these data to the admin-2 level.

4 Results

We estimate location-specific economic effects of temperature via a local-constant kernel smoother on a global panel of admin-2 geographic units. The *geographically varying* (GV) model (2.1) is estimated with a distributed lag order $J = 1$.[11] The optimal number of nearest neighbor locations selected through cross-validation is $K = 570$.[12] The corresponding bandwidths imply distances to the Kth nearest neighbor with an interquartile range between 1,887 and 2,555 km and an average of 2,343

[9] Weblink: `https://landuse.sites.uu.nl/datasets/` (accessed August 15, 2025).

[10] Weblink: https://data.isimip.org/search/tree/ISIMIP2b/SecondaryInputData/socioeconomic/pop/ssp2soc/ (accessed August 15, 2025).

[11] Adding further lags does not alter the results qualitatively while, as expected, producing noisier estimates.

[12] Given the large number of unique locations, K was cross-validated in two steps. In the first step, we grid-searched from 40 to 4,000 in increments of 40 to identify a provisional best K. In the second step, we searched within ±40 of this provisional value using increments of 2. To make cross-validation computationally feasible, we rounded latitude and longitude to zero decimal places, reducing the original 40,997 locations to 8,370 larger geographic units. This rounding assumes that locations within a one-degree grid cell share the same coefficients for the purpose of choosing K, but the bandwidths to the Kth nearest neighbor $r_K(\ell)$ used in estimation are based on the original, non-rounded geographic coordinates.

km.[13] Besides our semiparametric GV model, for benchmarking we also estimate a standard *globally constant* (GC) model that our specification in (2.1) nests as a special case, whereby all parameter functions are a priori assumed to be fixed coefficients across all locations: $\boldsymbol{\theta}(\ell_i) = \boldsymbol{\theta}\ \forall i$.[14]

For inference, we use wild residual block bootstrap with cluster resampling at the province (admin-1) level; there are 2,722 such blocks/regions. Two distinct notions of statistical significance appear throughout our analysis. Where the primary interest is in whether a given location's experienced economic effect or projected income change due to warming is detectably positive or negative, we employ one-sided tests at the 5% significance level. We color-code locations accordingly in figures: green indicates a statistically positive effect (the lower one-sided 95% bound exceeds zero), red a statistically negative effect (the upper one-sided 95% bound falls below zero) and gray a neither statistically positive nor negative effect. In tables summarizing our estimates and projections, we follow the more conventional practice and report two-sided 95% confidence intervals.

4.1 Testing for Global Homogeneity

Before turning to the local estimates, we establish that the data reject the globally constant coefficient restriction in favor of our geographically varying alternative. We do so via a formal joint test of global homogeneity (i.e., constancy of all slope parameters in the temperature-macroeconomy relationship across space) built on Ullah's (1985) nonparametric goodness-of-fit statistic, which discriminates between the restricted GC specification under the null and the unrestricted GV specification under the alternative. The statistic measures the proportional reduction in the residual sum of squares delivered by the flexible model, $T_n = (\text{RSS}_{GC} - \text{RSS}_{GV})/\text{RSS}_{GV}$, where RSS_{GC} and RSS_{GV} are the residual sums of squares from the GC and GV fitted specifications, respectively. As with a classical parametric "F-test," T_n converges to zero under the null of global homogeneity and is positive under the alternative, rendering the test one-sided. We approximate the null distribution of T_n by bootstrap that resamples residuals from the restricted GC specification at the admin-1 level, mirroring the inference procedure used throughout. The test firmly rejects global homogeneity, with the bootstrap $p < 0.01$, suggesting that the globally constant specification is statistically inadequate, which motivates the local estimates to which we now turn.

4.2 Historical Relationship

For each admin-2 geographic unit, our GV model delivers a cumulative marginal effect of a 1°C increase in temperature on per-capita income growth accounting for local context via location-specific parameters, evaluated at that location's historical average temperature $\overline{T}_i$. The cumulative effect combines the contemporaneous and lagged responses, $\left[\beta_0(\ell_i) + \gamma_0(\ell_i)\overline{T}_i\right] + \left[\beta_1(\ell_i) + \gamma_1(\ell_i)\overline{T}_i\right]$, and

[13] Geographic coordinates are measured in decimal degrees. For convenience, we convert to kilometers using the approximation $1° \approx 111$ km, which is exact at the equator.

[14] When forcing all parameters in (2.1) to be fixed, the resultant treatment of time trends becomes "global," viz., αt. To allow for more flexibility in capturing common time shocks in the GC model, we replace this globally linear trend with admin-1-level year fixed effects.

therefore summarizes the net annual growth effect of a temperature shock after accounting for any partial reversal one period later.

Figure 2 presents the cross-location distribution of these cumulative marginal effects evaluated at a sequence of counterfactual average-temperature values, $T^* = 5°, 10°, \ldots, 35°$C. Because the marginal effect at each location depends linearly on the temperature baseline through $\gamma_j(\ell_i)$, varying T^* traces out how the growth response to a uniform 1°C warming changes with the climatic environment in which it occurs. Two patterns are apparent. First, the counterfactual cross-location median and mean both decline monotonically with T^*: near-zero mean effects turn increasingly negative at $T^* \geq 10$°C, reaching approximately −0.9 percentage points (pp) at $T^* = 35$°C. Second, the distributions exhibit wide cross-location dispersion at every evaluation point, with the interquartile range spanning positive and negative territory throughout most of the T^* domain. This within-T^* dispersion reflects the geographic heterogeneity in coefficient surfaces $\beta_j(\ell_i)$ and $\gamma_j(\ell_i)$ that our model is designed to capture. Namely, locations sharing a similar climatic baseline can nonetheless have substantially different growth responses to warming owing to differences in the local context that moderates those responses. This pattern is broadly consistent with the concave temperature-productivity relationship documented at the country level by Burke et al. (2015), who find that economic output peaks at an annual average temperature of approximately 13°C (also see Nath et al., 2024). Our estimates reveal a qualitatively similar gradient at a much finer geographic resolution, with the additional observation that the cross-location dispersion around that gradient is empirically large and cannot be captured by a single globally constant parameter.

Figure 3 disaggregates cumulative marginal effect functions derived from our GV model by Köppen-Geiger climate class.[15] Each panel overlays the location-specific effect functions (thin light-blue lines) and their within-class median (dark-blue solid line) for one of the four major climate classes, alongside the fixed-coefficient benchmark from the standard GC model (black dashed line, identical across panels). The histogram above each panel displays the empirical distribution of $\overline{T}_i$ within that class, indicating the temperature range over which the estimates are most empirically relevant. Several observations emerge. In Tropical locations (panel a), cumulative marginal effects are overwhelmingly negative across the range of observed average temperatures, with the within-class median falling steeply and lying well below the GC benchmark. The Dry class (panel b) is more heterogeneous, with effect functions spanning a wide range and the median hovering near zero or mildly negative at the higher temperatures typical of arid regions. Temperate locations (panel c) display a gradient that is qualitatively similar to the full-sample pattern in Figure 2: mildly positive or near-zero effects at lower average temperatures turning negative at higher ones, with the median tracking closer to but still diverging from the fixed-coefficient line. Continental locations (panel d) have the most favorable profile, with the within-class median remaining positive or near zero over the temperature range that characterizes this class.

[15] This is a widely adopted scheme for categorizing global climates. The classification divides Earth's climate into five main climatic zone groups based on patterns of seasonal precipitation and temperature: A (Tropical), B (Arid), C (Temperate), D (Continental) and E (Polar) (Beck et al., 2023). We focus on A–D for obvious reasons. Spatial climate classification data are available at `https://www.gloh2o.org/koppen/`.

To account for the sampling variability, we now turn to the question of statistical significance and its geographic distribution. Figure 4 reports GV point estimates (scatter dots) of the contemporaneous, lagged and cumulative marginal effects for all locations, sorted by their corresponding one-sided 95% confidence bounds (solid line). In each row, locations with a statistically positive effect at the 5% level appear in green (with the lower bound above 0) and those with a statistically negative effect appear in red (with the upper bound below 0), with gray denoting locations where the null of zero effect cannot be rejected in the relevant one-sided direction. Accounting for local contexts, 31.7% of locations experience a statistically negative contemporaneous effect of temperature, 8.3% a statistically positive effect and the remaining 60.0% statistically no different from zero. The lagged effect is more muted: 20.8% statistically negative, 5.0% statistically positive and 74.1% statistically zero. Summing the two components, the cumulative marginal effect is statistically negative in 41.4% of locations and statistically positive in 9.8%, with 48.9% statistically zero.

The dynamics underlying these results warrant brief discussion. The cumulative marginal effect nets the contemporaneous and lagged responses. When a contemporaneous negative effect is followed by a partially offsetting positive lagged response (or vice versa), the cumulative effect is attenuated relative to the contemporaneous one, consistent with a mechanism in which some portion of the initial growth effect reverses in the subsequent period. As discussed in Section 2, such partial reversal is the empirical signature of what Dell et al. (2012) characterize as level effects on income that dissipate through subsequent growth recovery, as distinct from permanent growth effects whose dynamic components do not offset. That over half the locations retain a statistically nonzero cumulative effect indicates that, for a substantial share of the global sample, the temperature-income relationship is not purely transitory. This is consistent with the view that temperature has more than a fleeting influence on the trajectory of local incomes, at least in the parts of the world where the estimated effects are largest in magnitude. Overall, historical evidence indicates that warming significantly dampens economic growth in 41.4% of admin-2 geographic units. While a minority, the share of regions with statistically positive cumulative effects (9.8%) is also nontrivial and mostly concentrated in colder higher-latitude/altitude regions. This finding aligns with Burke et al.'s (2015) country-level quadratic specification that implies gains for cold countries located to the left of the per-capita income peak but whose fixed-coefficient setup does not permit these gains to vary with local context within a country.

Figure 5 maps the geographic distribution of estimated historical realizations of local cumulative marginal effects of warming, with the color scale ranging from green (positive) through gray (no different from zero) to red (negative), where significance is assessed via the one-sided 5% test described above. Panel (a) reports the estimates from the GV specification. A clear geographic gradient emerges in the map. Statistically positive effects are concentrated in colder high-latitude or high-elevation regions, most visibly across northern North America, Russian, Central Asia along with parts of the Himalayan/Tibetan region, with smaller pockets in northern and western Europe (predominantly along the Atlantic). Statistically negative effects dominate warmer tropics and subtropics, especially Latin America, Southeast and East Asia, Maritime Continent and most of Australia. Between these extremes, the locations that have been statistically insensitive to warming span most of the United States, western and central Europe, the Middle East, India, most of Africa and large parts of interior China. Thus, once the temperature-economy relationship is allowed to vary with local context, the implied temperature effects are spatially heterogeneous and broadly aligned with local climatic exposure: warming is beneficial in many cold locations, harmful in many already warm locations and statistically muted in much of the intermediate range.

Panel (b) of Figure 5 displays the corresponding estimates from the fixed-coefficient GC specification. The contrast is stark. With a single set of coefficients governing the entire globe, the estimated marginal effects are mechanically organized primarily by baseline temperature (see Figure 1(b) for comparison) rather than by locally varying response parameters. The map therefore assigns negative growth effects of temperature to most of the populated world outside the far northern high latitudes, roughly south of 55°N. Positive effects are confined to a narrow set of very cold regions, mainly parts of northern Canada and northern Siberia, while many colder but economically important regions become statistically insignificant. Relative to panel (a), the GC model compresses local heterogeneity into a global temperature-response schedule, producing a much more uniform pattern of damages outside the coldest locations. This comparison illustrates the empirical importance of allowing the temperature-economy relationship to vary with local context. Specifically, imposing global homogeneity substantially changes both the sign and significance of estimated marginal effects across large parts of the world, eliminating the geographic bifurcation in climate damages that is a defining feature of the local GV estimates.

It is worth emphasizing that the geographic bifurcation in temperature effects conditional on local context apparent in panel (a) of Figure 5 is not driven by cross-location differences in the temperature baseline $\overline{T}_i$ alone. The spatially homogeneous GC specification also conditions marginal effects on $\overline{T}_i$ through the interaction of T_{it} with $\overline{T}_i$, yet it produces a near-uniformly negative damage map (panel b) rather than a bifurcated one. What distinguishes the GV specification is the geographic (via ℓ_i) variation in the coefficients $\{\beta_j(\ell_i)\}_{j=0}^{1}$ and $\{\gamma_j(\ell_i)\}_{j=0}^{1}$ themselves. The GC model forces a common parameter vector $(\beta_0, \beta_1, \gamma_0, \gamma_1)$ that governs the temperature-economy mapping globally, so all cross-location variation in marginal effects is channeled through $\overline{T}_i$, and locations with the same temperature baseline therefore receive identical marginal effects regardless of differences in their economic structure, sectoral composition or adaptive capacity. When these unobserved local moderators of temperature sensitivity are themselves correlated with geographic location, as discussed in Section 2, a single global coefficient cannot represent them, and the resulting marginal-effect gradient in $\overline{T}_i$ is contaminated by the heterogeneity it averages over. The GV specification, by contrast, allows the coefficients to vary spatially and thereby absorbs the moderating role of local economic context, producing marginal effects that reflect not just where a location sits on the temperature scale but how its economy translates temperature into output. This distinction clarifies the economic content of the bifurcation.

In warm tropical and subtropical regions, the GV coefficients capture an economic context in which temperature damages are amplified relative to the global average: greater dependence on outdoor labor and rain-fed agriculture, less penetration of climate control and, more broadly, lower adaptive capacity (see Hallegatte et al., 2016). Labor productivity in these settings declines nonlinearly with heat exposure (Graff Zivin and Neidell, 2014; Zhang et al., 2018b; Somanathan et al., 2021), and agricultural productivity is disproportionately harmed by warming (Ortiz-Bobea et al., 2021). The local coefficients internalize these features, yielding marginal effects that are more negative than what a global average slope would assign to the same $\overline{T}_i$. In colder high-latitude and high-altitude regions, the local coefficients instead capture an economic context in which warming relaxes binding constraints. Specifically, shorter winters extend construction and agricultural seasons, reduce heating expenditures and lower cold-related morbidity; agricultural growing seasons lengthen, and previously marginal land becomes cultivable or more productive, particularly in northern Eurasia. These benefits are structural features of the local economy that the GV coefficients detect but a single global coefficient averages away. In the intermediate temperate range, the local coefficients reflect partially offsetting channels, yielding net effects that are small and statistically indistinguishable from zero. The within-T^* dispersion visible in Figure 2 (wide interquartile ranges at any given counterfactual temperature) corroborates this reading. Namely, even among locations sharing a common temperature baseline, substantial variation in marginal effects persists because the GV coefficients capture differences in the local economic context that moderate the temperature-economy relationship. The bifurcation is therefore an empirical consequence of allowing heterogeneous local contexts to shape the estimated temperature sensitivity, not an artifact of the state dependence in how temperature influences economy.

Table 1 quantifies this contrast using population-weighted summary statistics of the cumulative marginal effects. Under the GV specification accounting for local context, the population-weighted mean cumulative effect of a 1°C increase in temperature is -0.31 pp but statistically insignificant at the 5% level, albeit the first quartile is significant at -0.74 pp, indicating that the lower portion of the population-weighted distribution of marginal effects is detectably negative. Under the globally homogeneous GC specification, the corresponding mean is more material and statistically nonzero at -0.57 pp with, in fact, the entire interquartile range lying significantly below zero. The GC model therefore estimates the average damage of warming by roughly 80% more relative to the GV specification, not to a small degree, by eliminating most positive effects in the upper tail. This buttresses our argument that when location-specific heterogeneity in temperature sensitivity is correlated with the temperature realizations themselves, a globally-constant-coefficient model does not recover a well-defined average marginal effect but rather a composite that conflates the heterogeneous temperature-economy mapping with the distribution of the moderating local context.

4.3 Future Projections

We now translate the estimated historical temperature-economy relationship into projections of economic effects of future warming. Our projection methodology follows the framework of Burke et al. (2015) adapted to our state-dependent specification with a distributed lag. For each admin-2 location, future GDP per capita paths are constructed recursively with and without the additional growth effect implied by projected warming temperatures, and the percentage difference between the two forecast paths at a given future year provides the projected average income effect of warming for that location.

Concretely, for each admin-2 unit i, the projected average income path with warming is generated recursively as $y_{it} = y_{it-1}(1+\eta_i+\delta_{it})$, whereas the no-warming counterfactual is $y_{it} = y_{it-1}(1+\eta_i)$, initialized at observed 2022 per-capita GDP. Here, η_i is the location-specific average historical annual growth rate absent future climate change (see the map in Figure 1(a)), and δ_{it} is the additional—over the historical baseline—annual growth effect implied by projected warming temperatures. Assuming no change in precipitation to focus on the impact of temperature, we compute δ_{it} as

$$\begin{aligned}\delta_{it} = &\left[\beta_0(\ell_i)T^+_{it} + \gamma_0(\ell_i)T^+_{it}\overline{T}^+_i + \beta_1(\ell_i)T^+_{it-1} + \gamma_1(\ell_i)T^+_{it-1}\overline{T}^+_i\right] - \\ &\left[\beta_0(\ell_i)\overline{T}_i + \gamma_0(\ell_i)\overline{T}_i\overline{T}_i + \beta_1(\ell_i)\overline{T}_i + \gamma_1(\ell_i)\overline{T}_i\overline{T}_i\right],\end{aligned}$$

where the superscript "+" indicates projected future temperatures for years $t \geq 2023$. The first bracket is the cumulative growth rate effect predicted using projected future temperatures, for which we use historical values of the local state variable $\overline{T}_i$ for the future $\overline{T}^+_i$. The second bracket is the cumulative growth rate effect estimated using average historical temperature in each location, where all future T^+_{it} has been set equal to $\overline{T}_i$. In spirit with Burke et al. (2015), to compute δ_{it} for each year between 2023 and 2100, future temperatures are obtained by linearly interpolating between the observed 2022 realizations and the average projected daily maximum temperatures in 2096–2100, and projected temperatures exceeding the historical highest annual daily maximum of 39°C are capped at that threshold.[16]

We present results for the SSP5-8.5 scenario using climate projections from two GCMs that bracket a range of warming magnitudes: CanESM5 (warmer projections) and NorESM2-MM (cooler projections). Corresponding results for the more moderate SSP2-4.5 socioeconomic scenario are reported in Appendix A.

Figure 6 plots projected per-capita income paths with and without climate change under SSP5-8.5 for both GCMs. The colored lines represent income trajectories accounting for the warming-induced growth increment δ_{it}, while the gray lines represent the no-warming counterfactual in which each location grows at its historical average rate η_i. By mid-range horizon, divergence between the two sets of trajectories is already apparent, and by the end of century it is pronounced: many locations—predominantly in the (sub)tropics—have income paths that fall substantially below their no-warming counterfactuals, while a subset of locations—predominantly in higher latitudes/altitudes—track above or near their counterfactuals. The dispersion across locations is wider under CanESM5 (panel a) than under NorESM2-MM (panel b), reflecting the larger warming magnitudes projected by the former. These fan charts illustrate at the location level the same geographic bifurcation documented in the

[16] For $t \in [2022, 2100]$, the projected future temperature in a location i is thus given by $T^+_{it} = \min\{T_{i2022} + \frac{\Delta T^+_i}{78}(t-2022), 39\}$, where ΔT^+_i is the change between year-2022 temperature and the average of projected 2096–2100 temperatures. Resultant is a smooth, year-by-year trajectory from the historical observation in 2022 to the projected end-of-century climate.

historical estimates: the geographically varying temperature-economy relationship, when combined with heterogeneous warming trajectories, produces a distribution of future income effects that is far from uniformly negative.

Figure 7 aggregates these location-level effects into population-weighted measures of aggregate income change due to warming, separately for the full sample and for each Köppen-Geiger climate class. In each panel, the purple line reports the projected aggregate effect based on the GV specification, and the black line the corresponding projection from the GC model. Table 2 provides a quantitative summary of these projected aggregate income effects for both GCMs. For the full sample, both model specifications project aggregate income losses that grow over time, but the magnitudes differ markedly. By 2100, the GV model projects an aggregate global loss of approximately 25% under CanESM5, compared with roughly 73% under the GC model. The decomposition by climate class reveals the source of this divergence. In Tropical regions, both specifications project large losses. The critical difference appears in the other climate classes. For Dry regions, the GV model projects a large positive aggregate effect (approximately +98% under CanESM5 by 2100), whereas the GC model projects a loss of −84%. For Continental regions, the pattern is similar: +32% by the end of century under the GV specification vs. −51% under the GC fixed-coefficient specification . For Temperate regions, both project losses, but the GV model produces a substantially smaller estimate (−33% vs. −73%). Under NorESM2-MM, the same qualitative patterns hold with smaller magnitudes throughout. The SSP2-4.5 counterparts, reported in Figure A.2 in the appendix, exhibit the same structure with still more attenuated magnitudes. Across all scenario-GCM combinations, the GC specification that imposes global homogeneity onto temperature-economy relationship projects uniformly negative and substantially larger aggregate damages for every climate class that inflate the global aggregate, because it is incapable of detecting the positive effects of temperature that the GV model identifies in colder and drier regions. From Table 2 and its counterpart table A.1 for the SSP2-4.5 scenario in Appendix A, the overstatement of aggregate global damages relative to the geographically varying specification is by a factor of approximately two to three.

Figure 8 maps the geographic distribution of end-of-century (year 2100) projections of per-capita income effects of warming under SSP5-8.5 for both GCMs. The top row displays projections from the GV specification, the bottom row from the GC model. Based on the GV specification, we observe a clear geographic pattern, whereby projected average income gains (green) are concentrated in northern latitudes and high-elevation regions, spanning a broad Eurasian–interior Asian corridor running from western Russia and Siberia through Central Asia, the Caucasus, northern Iran and the Himalayan/Tibetan region toward Mongolia, northern China and the Russian Far East, with pockets in West and South Europe along the Atlantic, western Magreb, Iceland and Greenland, as well as New Zealand in the southern hemisphere. Whereas, projected losses (red) dominate the tropics and subtropics. The economic loss areas are much more spatially extensive than benefit areas. The intensity of the projected losses is greatest in Central America, northern South America and the Amazon/Andean tropics, coastal East Asia, the Maritime Continent and Australia, the Arabian peninsula and (sub)tropical Africa, especially parts of West/Central Africa and southern Africa.

In contrast but expectedly, the GC specification produces projections the maps of which are almost entirely red, with the intensity of losses more uniform across latitudes. The fixed-coefficient model projects large income losses even in high-latitude regions where the local-context-adapting GV model estimates positive effects, illustrating how the imposition of a globally constant temperature sensitivity can fundamentally mischaracterize the geographic distribution of climate effects.[17]

In sum, the projection exercise reinforces the central finding of the historical analysis, that is, failing to account for the geographic heterogeneity in temperature sensitivity, and in particular for the non-neutral role of local context in moderating that sensitivity, leads to substantially overstated global damage estimates and, more fundamentally, to a qualitative mischaracterization of the geography of climate impacts. The GV specification reveals a world in which the economic effects of warming are geographically bifurcated, with large projected losses in the (sub)tropics partially offset at the global level by projected gains in higher-latitude/altitude regions, rather than the uniformly negative landscape that the conventional GC model implies.

5 Conclusion

The globally constant coefficient restriction embedded in the standard fixed-effects framework is susceptible to heterogeneity bias in estimated climate damages because it forces a single temperature-economy mapping onto locations whose responses are moderated by unobserved local context. When those moderators are correlated with within-location temperature variation, the resulting global slope does not recover a well-defined average marginal effect but rather a composite that conflates heterogeneous temperature sensitivities with the distribution of confounding local factors.

We address this limitation by estimating a geographically varying coefficient model via local kernel smoothing on a global admin-2-level panel, with coefficient surfaces formalized as smooth functions of geographic coordinates. The specification imposes only that nearby locations have similar temperature sensitivities, a restriction grounded in the spatial regularity of economic and climate data that Tobler's first law of geography formalizes, thereby encoding the inherent spatial structure of the data without restricting the temperature-economy relationship to be globally uniform.

Our estimates reveal a pronounced geographic bifurcation in the economic effects of temperature that the globally constant specification is unable to detect. Statistically negative cumulative effects of warming concentrate in tropical and subtropical regions, where local economic structures amplify temperature damages relative to the global average, while statistically positive effects cluster in colder high-latitude and high-altitude locations, where warming relaxes binding climatic constraints on economic activity. Between these extremes, a substantial share of temperate locations exhibit effects that are statistically indistinguishable from zero, reflecting partially offsetting channels in which the local coefficients capture a more balanced economic context. The conventional globally constant specification eliminates this bifurcation entirely, producing a near-uniformly negative damage

[17] The corresponding SSP2-4.5 maps, reported in Figure A.3 in Appendix A, display the same geographic structure with moderately attenuated magnitudes.

map that overstates aggregate global damages by a factor of two to three across alternative warming scenarios and climate model projections. This overstatement arises primarily from the inability of a single global coefficient to detect the positive effects of temperature in drier, colder and continental regions that partially offset the large negative effects in the tropics at the aggregate level.

These findings suggest that accounting for geographic heterogeneity in the temperature-economy relationship (that goes beyond accommodating nonlinearities and/or state dependence) is not merely a refinement of existing damage estimates but changes the qualitative characterization of aggregate climate damages, with direct consequences for the calibration of integrated assessment models and the design of spatially targeted adaptation policy. Extensions to incorporate adaptation dynamics, sectoral decomposition or higher-order spatial interactions constitute a promising direction for future research.

References

Barrage, L. and Nordhaus, W. D. (2024). Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. *Proceedings of the National Academy of Sciences*, 121(13):e2312030121.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and 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(1):105–159.

Beck, H. E., McVicar, T. R., Vergopolan, N., Berg, A., Lutsko, N. J., Dufour, A., Zeng, Z., Jiang, X., van Dijk, A. I. J. M., and Miralles, D. G. (2023). High-resolution (1 km) köppen–geiger maps for 1901–2099 based on constrained cmip6 projections. *Scientific Data*, 10(724).

Bento, A. M., Miller, N., Mookerjee, M., and Severnini, E. (2023). A unifying approach to measuring climate change impacts and adaptation. *Journal of Environmental Economics and Management*, 121:102843.

Bilal, A. and Känzig, D. R. (2026). The macroeconomic impact of climate change: Global vs. local temperature. *The Quarterly Journal of Economics*, 141(2):889–944.

Bilal, A. and Stock, J. H. (2025). A guide to macroeconomics and climate change. NBER Working Paper No. 33567.

Burke, M., Davis, W. M., and Diffenbaugh, N. S. (2018). Large potential reduction in economic damages under UN mitigation targets. *Nature*, 557(7706):549–553.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

Burke, M. and Tanutama, V. (2019). Climatic constraints on aggregate economic output. NBER Working Paper No. 25779.

Callahan, C. W. and Mankin, J. S. (2022). Globally unequal effect of extreme heat on economic growth. *Science Advances*, 8(43):eadd3726.

Chamberlain, G. (1982). Multivariate regression models for panel data. *Journal of Econometrics*, 18(1):5–46.

Colacito, R., Hoffmann, B., and Phan, T. (2019). Temperature and growth: A panel analysis of the United States. *Journal of Money, Credit and Banking*, 51(2–3):313–368.

Cui, X., Gafarov, B., Ghanem, D., and Kuffner, T. (2024). On model selection criteria for climate change impact studies. *Journal of Econometrics*, 239(1):105511.

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

Goh, G., Yu, J., Kim, M., and Tack, J. (2025). Bayesian space–time varying coefficient modeling for climate econometrics: A spatial–temporal Gaussian process approach. *Journal of Econometrics*, page 106075.

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

Graham, B. S. and Powell, J. L. (2012). Identification and estimation of average partial effects in "irregular" correlated random coefficient panel data models. *Econometrica*, 80(5):2105–2152.

Hallegatte, S., Bangalore, M., Bonzanigo, L., Fay, M., Kane, T., Narloch, U., Rozenberg, J., Treguer, D., and Vogt-Schilb, A. (2016). *Shock waves: Managing the impacts of climate change on poverty*. World Bank Publications.

Hausfather, Z. (2019). Cmip6: The next generation of climate models explained. Accessed March 24, 2025.

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., and Houser, T. (2017). Estimating economic damage from climate change in the United States. *Science*, 356(6345):1362–1369.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.

Kahn, M. E., Mohaddes, K., Ng, R. N., Pesaran, M. H., Raissi, M., and Yang, J.-C. (2021). Long-term macroeconomic effects of climate change: A cross-country analysis. *Energy Economics*, 104:105624.

Kalkuhl, M. and Wenz, L. (2020). The impact of climate conditions on economic production: Evidence from a global panel of regions. *Journal of Environmental Economics and Management*, 103:102360.

Keane, M. and Neal, T. (2020). Climate change and U.S. agriculture: Accounting for multidimensional slope heterogeneity in panel data. *Quantitative Economics*, 11(4):1391–1429.

Kummu, M., Kosonen, M., and Masoumzadeh Sayyar, S. (2025). Downscaled gridded global dataset for gross domestic product (gdp) per capita ppp over 1990–2022. *Scientific Data*, 12(1):178.

Lobell, D. B. and Field, C. B. (2007). Global scale climate–crop yield relationships and the impacts of recent warming. *Environmental Research Letters*, 2:014002.

Nath, J. B., Ramey, V. A., and Klenow, P. J. (2024). How much will global warming cool global growth? *National Bureau of Economic Research Working Paper*, No. 32761.

Newell, R. G., Prest, B. C., and Sexton, S. E. (2021). The GDP-temperature relationship: Implications for climate change damages. *Journal of Environmental Economics and Management*, 108:102445.

Ortiz-Bobea, A., Ault, T. R., Carrillo, C. M., Chambers, R. G., and Lobell, D. B. (2021). Anthropogenic climate change has slowed global agricultural productivity growth. *Nature Climate Change*, 11:306–312.

Pesaran, M. H. and Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1):79–113.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

Ullah, A. (1985). Specification analysis of econometric models. *Journal of Quantitative Economics*, 1(2):187–209.

Waidelich, P., Batibeniz, F., Rising, J., Kikstra, J. S., and Seneviratne, S. I. (2024). Climate damage projections beyond annual temperature. *Nature Climate Change*, 14(6):592–599.

Wooldridge, J. M. (2005). Fixed-effects and related estimators for correlated random-coefficient and treatment-effect panel data models. *Review of Economics and Statistics*, 87(2):385–390.

Zhang, J., Malikov, E., and Miao, R. (2024). Distributional effects of the increasing heat incidence on labor productivity. *Journal of Environmental Economics and Management*, 125:102998.

Zhang, J., Malikov, E., Miao, R., and Ghosh, P. (2026). Geography of climate change adaptation in U.S. agriculture. *Journal of Regional Science*, 66(1):46–88.

Zhang, P., Deschenes, O., Meng, K., and Zhang, J. (2018a). Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17.

Zhang, P., Deschenes, O., Meng, K., and Zhang, J. (2018b). Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17.

Figures and Tables (in the order of appearance)

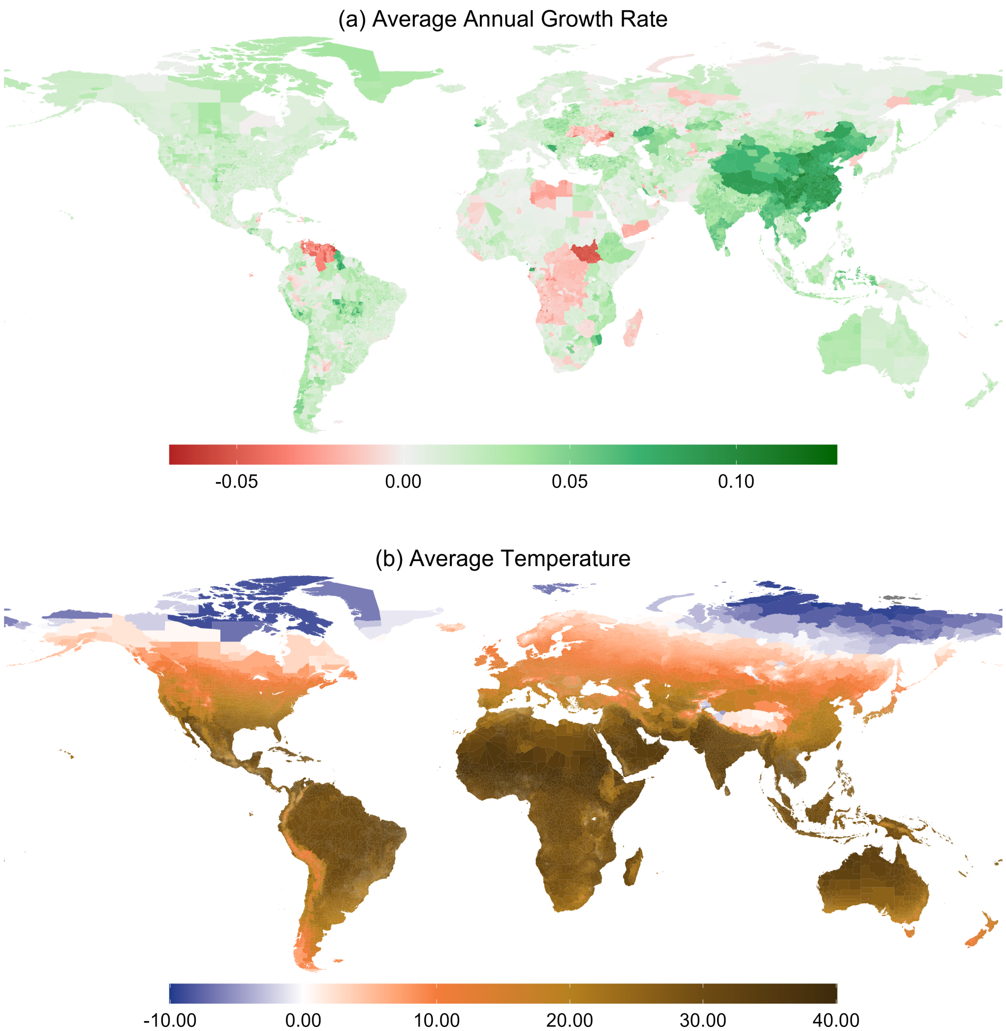

Figure 1. Historical Spatial Distribution of (a) Average Annual Growth Rate of Per-Capita Income and (b) and Average Temperature, 1990–2022

Notes: Average annual growth rate of per-capita income in location i is computed as $\left[y_{i2022}/y_{i1990}\right]^{1/32}-1$ using local yearly per-capita GDP y_{it}. Average temperature, $\overline{T}_i$, is computed for each location i as a cross-year mean of annual averages of daily maximum temperature.

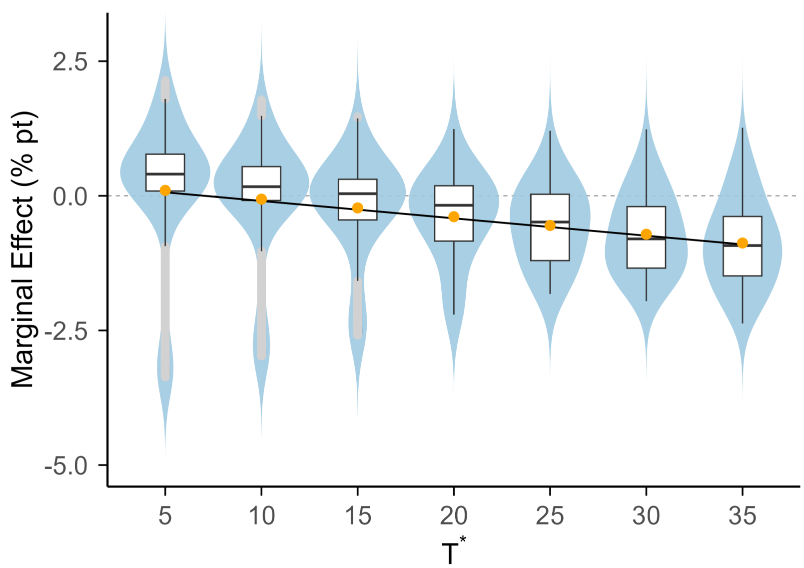

Figure 2. Cross-Location Distributions of Cumulative Marginal Effects of Warming Evaluated at Counterfactual Levels of Average Temperature

Notes: Violin boxplots show distributions of location-specific cumulative marginal effects of a 1°C increase in temperature, evaluated at counterfactual average-temperatures values, $[\beta_0(\ell_i)+\gamma_0(\ell_i)T^*]+[\beta_1(\ell_i)+\gamma_1(\ell_i)T^*]$ for $T^* = 5°, 10°, \ldots, 35°$C. Black horizontal lines inside boxplots mark cross-location medians; orange circles mark cross-location means. Estimates are in % points.

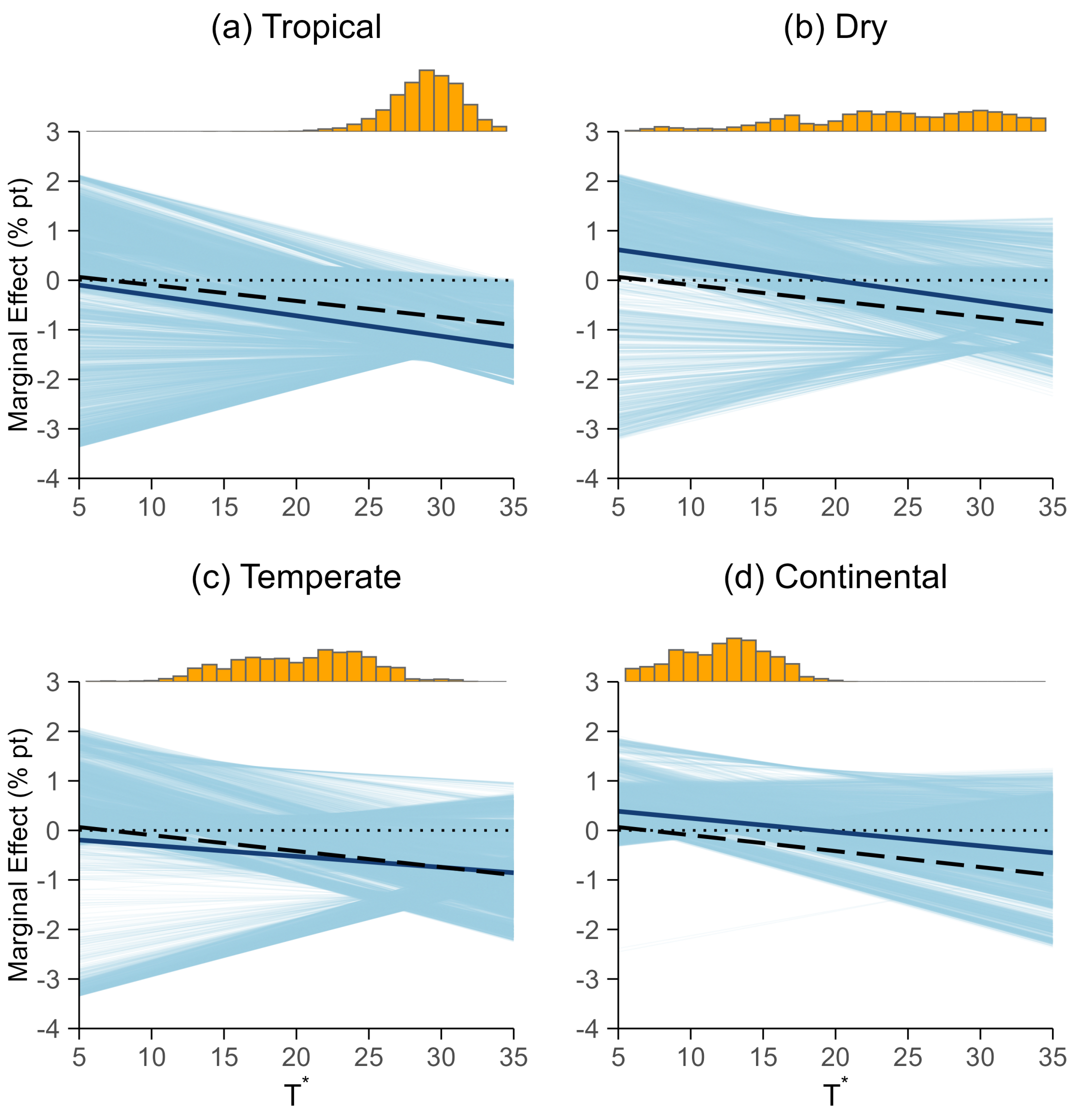

Figure 3. Cumulative Marginal Effect Functions of Warming by Climate Class

Notes: Each panel plots the cumulative marginal effect function of a 1°C increase in temperature, evaluated at counterfactual average-temperature values $T^* = 5°, 6°, \ldots, 35°$C, for locations belonging to one of the four major Köppen-Geiger climate classes: Tropical (a), Dry (b), Temperate (c) and Continental (d). For location i, the reported effect is $[\beta_0(\ell_i) + \gamma_0(\ell_i)T^*] + [\beta_1(\ell_i) + \gamma_1(\ell_i)T^*]$. The light-blue solid lines show location-specific estimates within each climate class, while the darker blue solid line shows the cross-location median marginal effect function. A black dashed line (identical across all panel) reports a fixed-coefficient benchmark. The histogram above each panel displays the within-class empirical distribution of average temperature. Estimates are in % points.

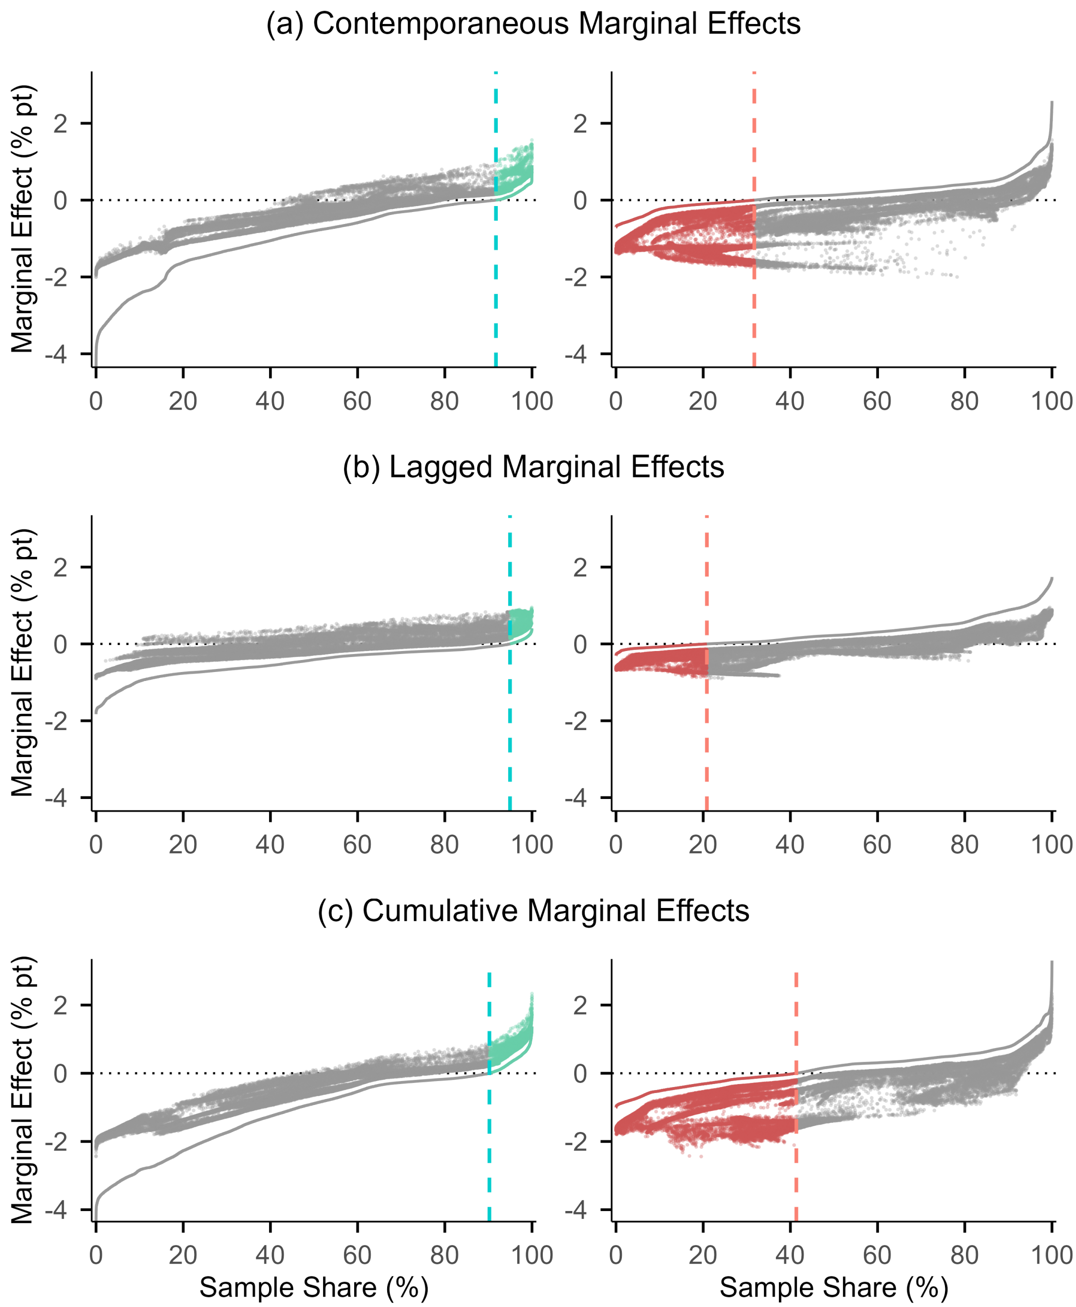

Figure 4. Local Point Estimates of (a) Contemporaneous, (b) Lagged and (c) Cumulative Marginal Effects of Warming, with the Corresponding One-Sided 95% Lower or Upper Bounds

Notes: The top, middle and bottom rows report, respectively, the contemporaneous marginal effect $\beta_0(\ell_i)+\gamma_0(\ell_i)\overline{T}_i$, the lagged marginal effect $\beta_1(\ell_i)+\gamma_1(\ell_i)\overline{T}_i$ and the corresponding cumulative marginal effect $\left[\beta_0(\ell_i)+\gamma_0(\ell_i)T^*\right]+\left[\beta_1(\ell_i)+\gamma_1(\ell_i)\overline{T}_i\right]$ of 1°C increase in temperature. In each row, the left (right) panel plots the corresponding point estimates–scatter dots—sorted by the corresponding lower (upper) one-sided 95% confidence bound, the solid line. Point estimates in green (red) indicate locations with a statistically positive (negative) effect at the 5% significance level. The horizontal axis reports the cumulative sample share after ordering locations by the relevant bound. Estimates are in % points.

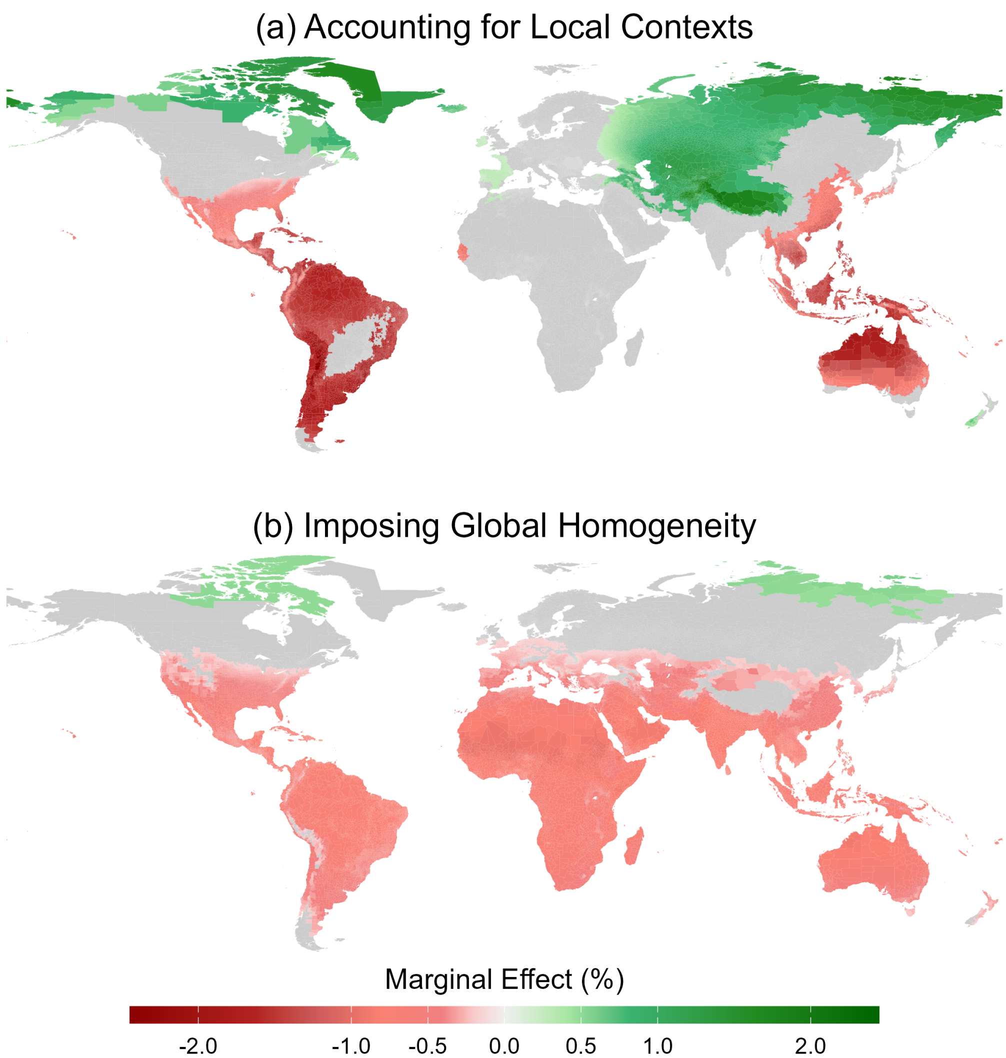

Figure 5. Geographic Heterogeneity in the Historical Local Cumulative Marginal Effects of Warming

Notes: Mapped are the estimated cumulative marginal effects of a 1°C increase in temperature, $\left[\beta_0(\ell_i)+\gamma_0(\ell_i)\overline{T}_i\right]+\left[\beta_1(\ell_i)+\gamma_1(\ell_i)\overline{T}_i\right]$, for each location i. Panel (a) uses estimates from our spatially varying coefficient model, thereby allowing the temperature-growth relationship to differ across locations as a function of local context. Panel (b) uses the corresponding globally-fixed-coefficient specification, which imposes a common temperature effect across locations. Green (red) indicates locations with a statistically positive (negative) effect at the 5% significance level; the remaining locations are in gray. Estimates are in % points.

Table 1. Summary of Historical Local Cumulative Marginal Effects of Warming

	Mean	1st Qu.	Median	3rd Qu.
Accounting for Local Contexts	–0.3128	–0.7435	–0.2456	0.0727
	(–0.715, 0.0895)	(–1.3697, –0.1174)	(–0.7256, 0.2344)	(–0.2276, 0.3731)
Imposing Global Homogeneity	–0.5670	–0.7524	–0.6295	–0.408
	(–0.8902, –0.2439)	(–1.1891, –0.3157)	(–0.9895, –0.2696)	(–0.6526, –0.1634)

Notes: Reported are population-weighted summary statistics of point estimates of historical location-specific cumulative marginal effects of a 1° increase in temperature, $\left[\beta_0(\ell_i)+\gamma_0(\ell_i)\overline{T}_i\right]+\left[\beta_1(\ell_i)+\gamma_1(\ell_i)\overline{T}_i\right]$. Estimates are in % points.

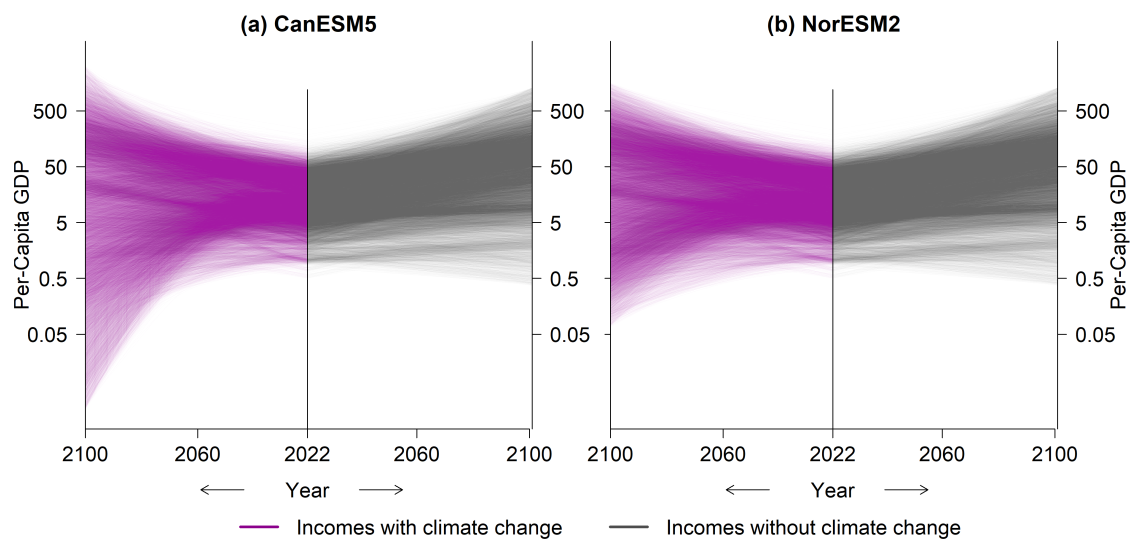

Figure 6. Projected Per-Capita Income Paths With and Without Climate Change

Notes: Plotted are projected paths of GDP per capita from 2022 to 2100 under two counterfactuals—with and without global warming—using end-of-century climate projections under the SSP5-8.5 scenario from CanESM5 (Panel a) and NorESM2-MM (Panel b) model. For each location i, the projected income path with warming is generated recursively as $y_{it} = y_{it-1}(1+\eta_i+\delta_{it})$, whereas the no-warming counterfactual is $y_{it} = y_{it-1}(1+\eta_i)$, initialized at observed 2022 per-capita GDP. Here, η_i is the location-specific average historical annual growth rate absent climate change, and δ_{it} is the additional—over the historical baseline—annual growth effect implied by projected warming temperatures. Future annual temperatures are constructed by linearly interpolating between the observed 2022 temperature and the mean projected temperature over 2096–2100. The colored lines show projected incomes accounting for climate change, while the gray lines show the corresponding no-warming counterfactuals. To facilitate readability, locations in the bottom and top 2.5% of year-2100 projections are omitted from the figure. The vertical axis is on a log scale.

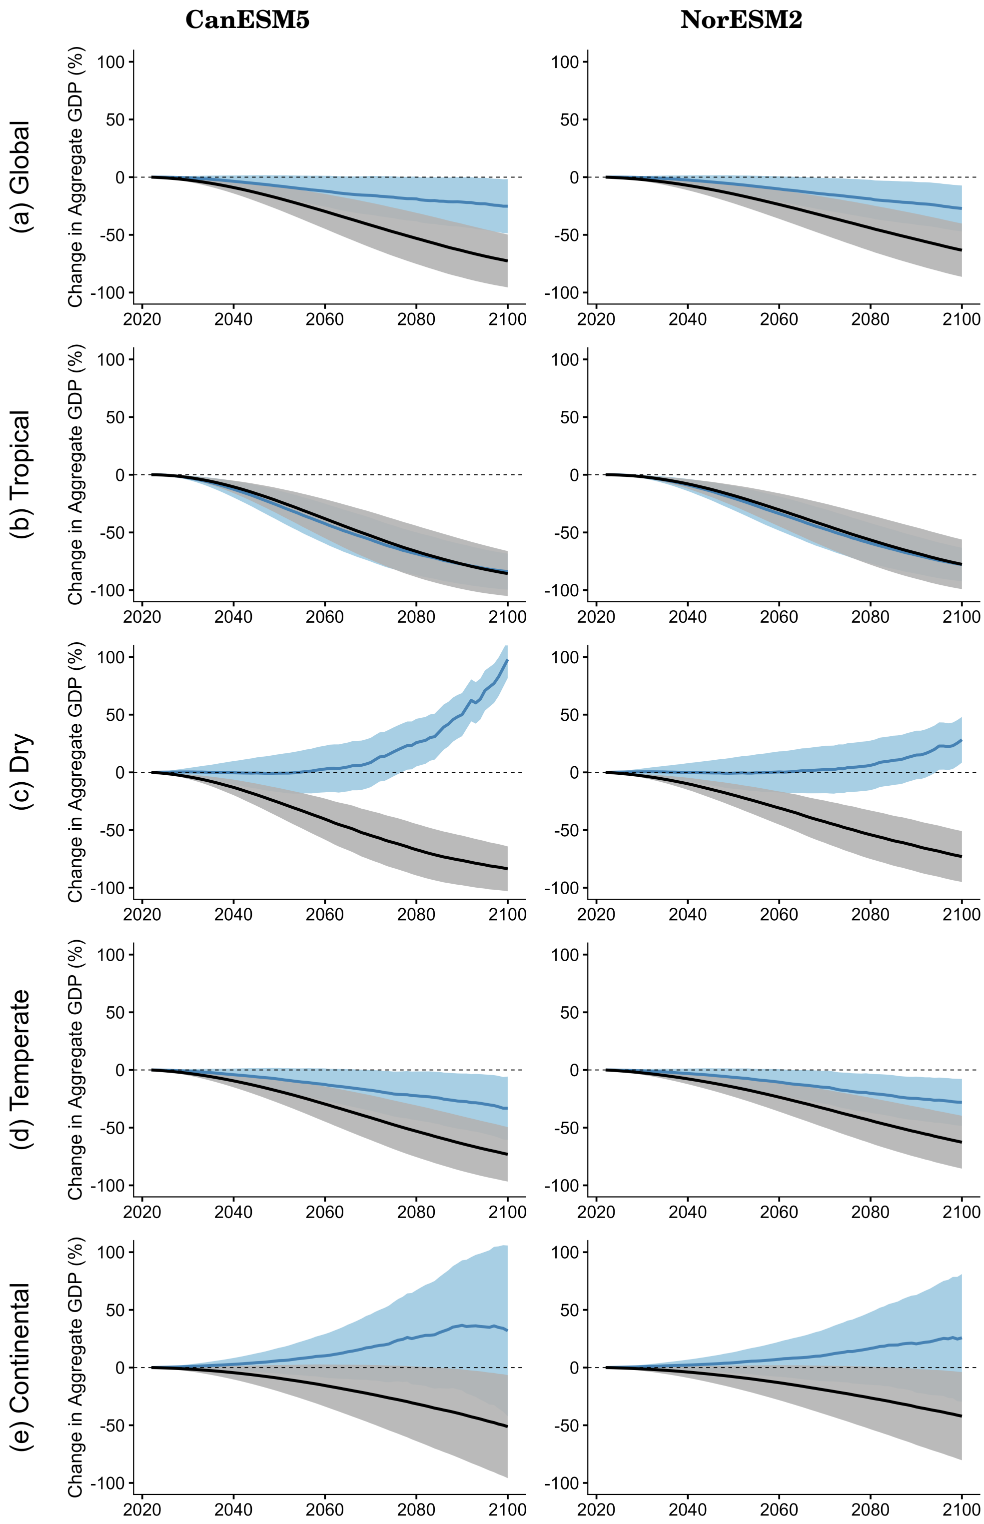

Figure 7. Projected Aggregate Income Effects of Warming by Climate Class

Notes: Plotted are projected % changes in aggregate income due to warming from 2022 to 2100 based on end-of-century climate projections under the SSP5-8.5 scenario from CanESM5 (left) and NorESM2-MM (right) model for the full global sample and separately by major Köppen-Geiger climate class. For any group G and year τ, the plotted effect is $[\sum_{i \in G} L_{i\tau} y_{i\tau}^{\text{no/warm}}]^{-1} \sum_{i \in G} L_{i\tau}(y_{i\tau}^{\text{w/warm}} - y_{i\tau}^{\text{no/warm}}) \times 100$, where $y_{i\tau}^{\text{w/warm}}$ and $y_{i\tau}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively, and $L_{i\tau}$ is projected population. The purple line reports the implied aggregate effect based on the spatially varying coefficient model, whereas the black line reports the corresponding effect from the globally fixed-coefficient specification. Shaded bands are two-sided 95% confidence intervals. To minimize outlier distortions, locations in the bottom and top 2.5% of income projections were omitted before aggregation.

Table 2. Projected Mid- and End-of-Century Aggregate Income Effects of Warming by Climate Class

	Global	Tropical	Dry	Temperate	Continental
			—CanESM5—		
2050 Impact Projection					
Accounting for Local Context	–7.75	–26.98	–0.68	–7.99	6.00
	(–17.08, 1.57)	(–40.12, –13.84)	(–17.09, 15.74)	(–17.66, 1.67)	(–5.15, 17.16)
Imposing Global Homogeneity	–18.47	–23.26	–26.25	–18.6	–9.31
	(–28.51, –8.42)	(–34.57, –11.95)	(–38.81, –13.69)	(–28.92, –8.29)	(–20.85, 2.23)
2100 Impact Projection					
Accounting for Local Context	–25.19	–83.97	97.91	–33.17	31.85
	(–48.68, –1.69)	(–99.51, –68.43)	(81.78, 114.04)	(–60.57, –5.77)	(–42.08, 105.77)
Imposing Global Homogeneity	–72.63	–85.63	–83.63	–73.12	–51.21
	(–95.54, –49.72)	(–105.11, –66.16)	(–103.06, –64.21)	(–96.72, –49.52)	(–95.85, –6.56)
			—NorESM2—		
2050 Impact Projection					
Accounting for Local Context	–5.91	–20.27	–0.65	–6.42	4.16
	(–13.36, 1.55)	(–30.69, –9.86)	(–14.18, 12.88)	(–14.1, 1.26)	(–5.01, 13.33)
Imposing Global Homogeneity	–14.58	–17.99	–19.74	–14.83	–7.95
	(–22.71, –6.45)	(–27.04, –8.94)	(–29.54, –9.94)	(–23.21, –6.45)	(–17.61, 1.72)
2100 Impact Projection					
Accounting for Local Context	–27.13	–77.79	28.27	–28.03	25.77
	(–47.03, –7.23)	(–92.23, –63.34)	(8.48, 48.06)	(–48.32, –7.74)	(–29.51, 81.05)
Imposing Global Homogeneity	–63.31	–77.65	–72.98	–62.63	–42.12
	(–86.4, –40.22)	(–99.13, –56.16)	(–95.03, –50.93)	(–85.53, –39.73)	(–80.34, –3.90)

Notes: Reported are projected % changes in aggregate income due to warming from 2022 to 2050 and to 2100 based on mid- and end-of-century climate projections from CanESM5 and NorESM2-MM models under the SSP5-8.5 scenario for the full global sample and separately by major Köppen-Geiger climate class. For any group G and year $\tau = \{2050, 2100\}$, the reported effect is $\left[\sum_{i\in G} L_{i\tau} y_{i\tau}^{\text{no/warm}}\right]^{-1} \sum_{i\in G} L_{i\tau}\left(y_{i\tau}^{\text{w/warm}} - y_{i\tau}^{\text{no/warm}}\right) \times 100$, where $y_{i\tau}^{\text{w/warm}}$ and $y_{i\tau}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively, and $L_{i\tau}$ is projected population. Two-sided 95% confidence intervals are in parentheses. To minimize outlier distortions, locations in the bottom and top 2.5% of income projections were omitted before aggregation.

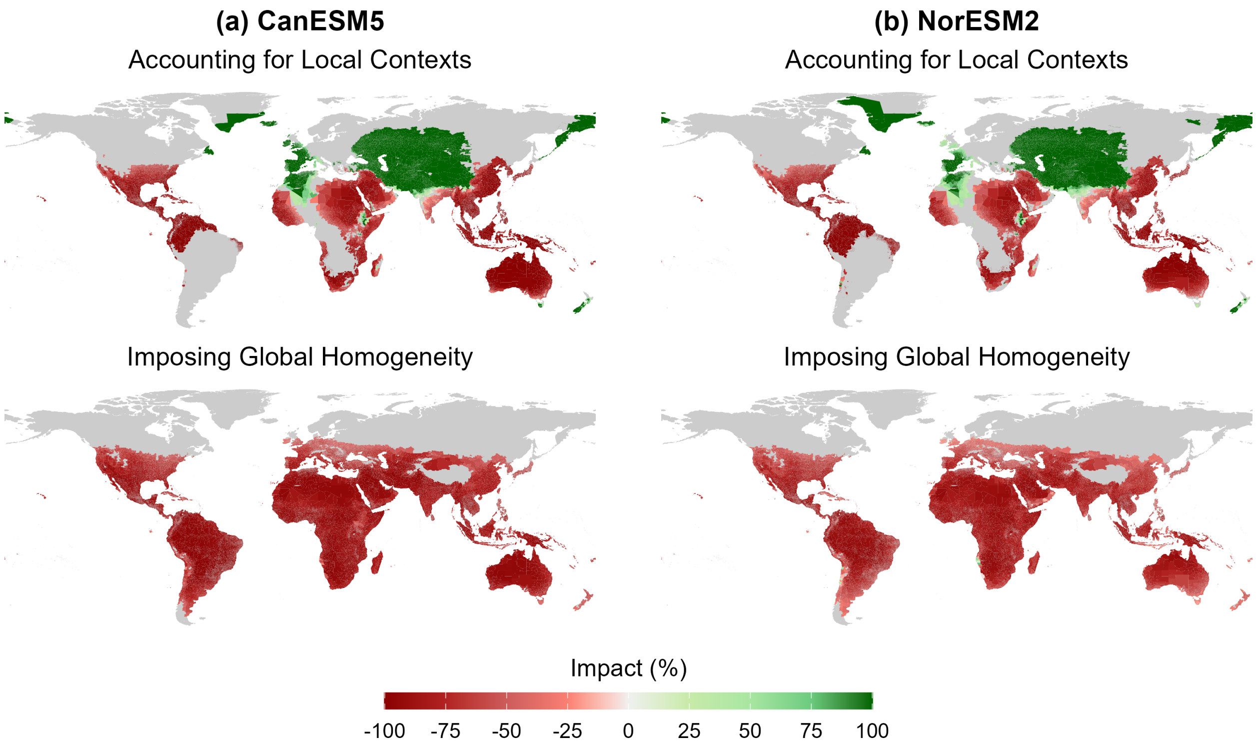

Figure 8. Geographic Distribution of End-of-Century Projections of Local Per-Capita Income Effects of Warming

Notes: Mapped are location-specific projections of per-capita income effects of warming by the year 2100: $(y_{i2100}^{\text{w/warm}} - y_{i2100}^{\text{no/warm}})/y_{i2100}^{\text{no/warm}} \times 100$, where $y_{i2100}^{\text{w/warm}}$ and $y_{i2100}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively. Panels (a) and (b) use end-of-century climate projections under SSP5-8.5 from the CanESM5 and NorESM2 models, respectively.The top-row maps plot estimates based on the spatially varying coefficient specification, which allows the temperature-growth relationship to differ across locations as a function of local context, whereas the bottom-row maps report the corresponding projections from the globally fixed-coefficient specification. Statistically—at the 5% level—positive (negative) projected effects are shown in green (red); the remaining locations are in gray. Estimates are in %.

Online Appendix

A Future Projections under the SSP2-4.5 Scenario

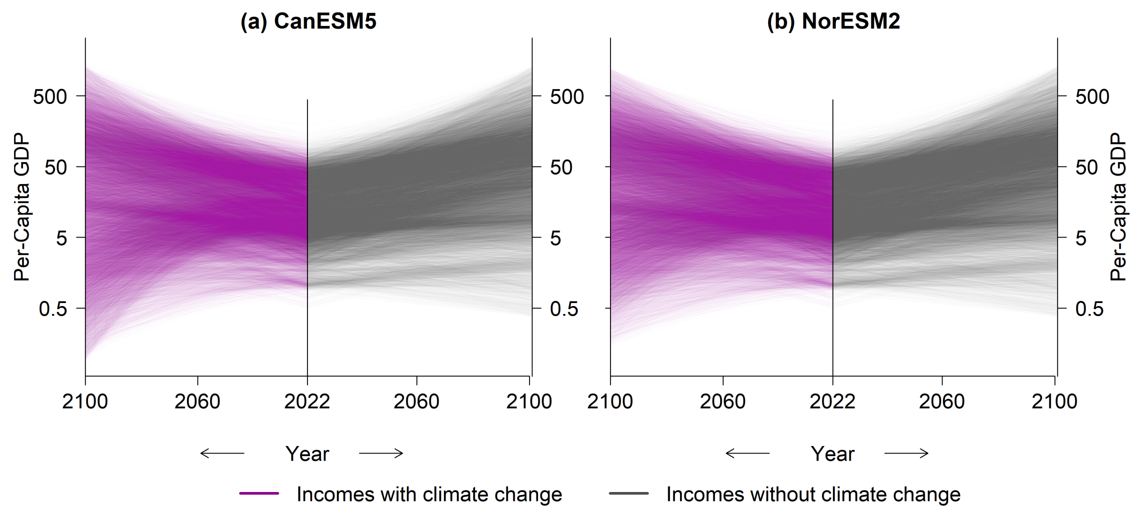

Figure A.1. Projected Per-Capita Income Paths under the SSP2-4.5 Scenario With and Without Climate Change

Notes: Plotted are projected paths of GDP per capita from 2022 to 2100 under two counterfactuals—with and without global warming—using end-of-century climate projections under the SSP2-4.5 scenario from CanESM5 (Panel a) and NorESM2-MM (Panel b) model. For each location i, the projected income path with warming is generated recursively as $y_{it} = y_{it-1}(1+\eta_i+\delta_{it})$, whereas the no-warming counterfactual is $y_{it} = y_{it-1}(1+\eta_i)$, initialized at observed 2022 per-capita GDP. Here, η_i is the location-specific average historical annual growth rate absent climate change, and δ_{it} is the additional—over the historical baseline—annual growth effect implied by projected warming temperatures. Future annual temperatures are constructed by linearly interpolating between the observed 2022 temperature and the mean projected temperature over 2096–2100. The colored lines show projected incomes accounting for climate change, while the gray lines show the corresponding no-warming counterfactuals. To facilitate readability, locations in the bottom and top 2.5% of year-2100 projections are omitted from the figure. The vertical axis is on a log scale.

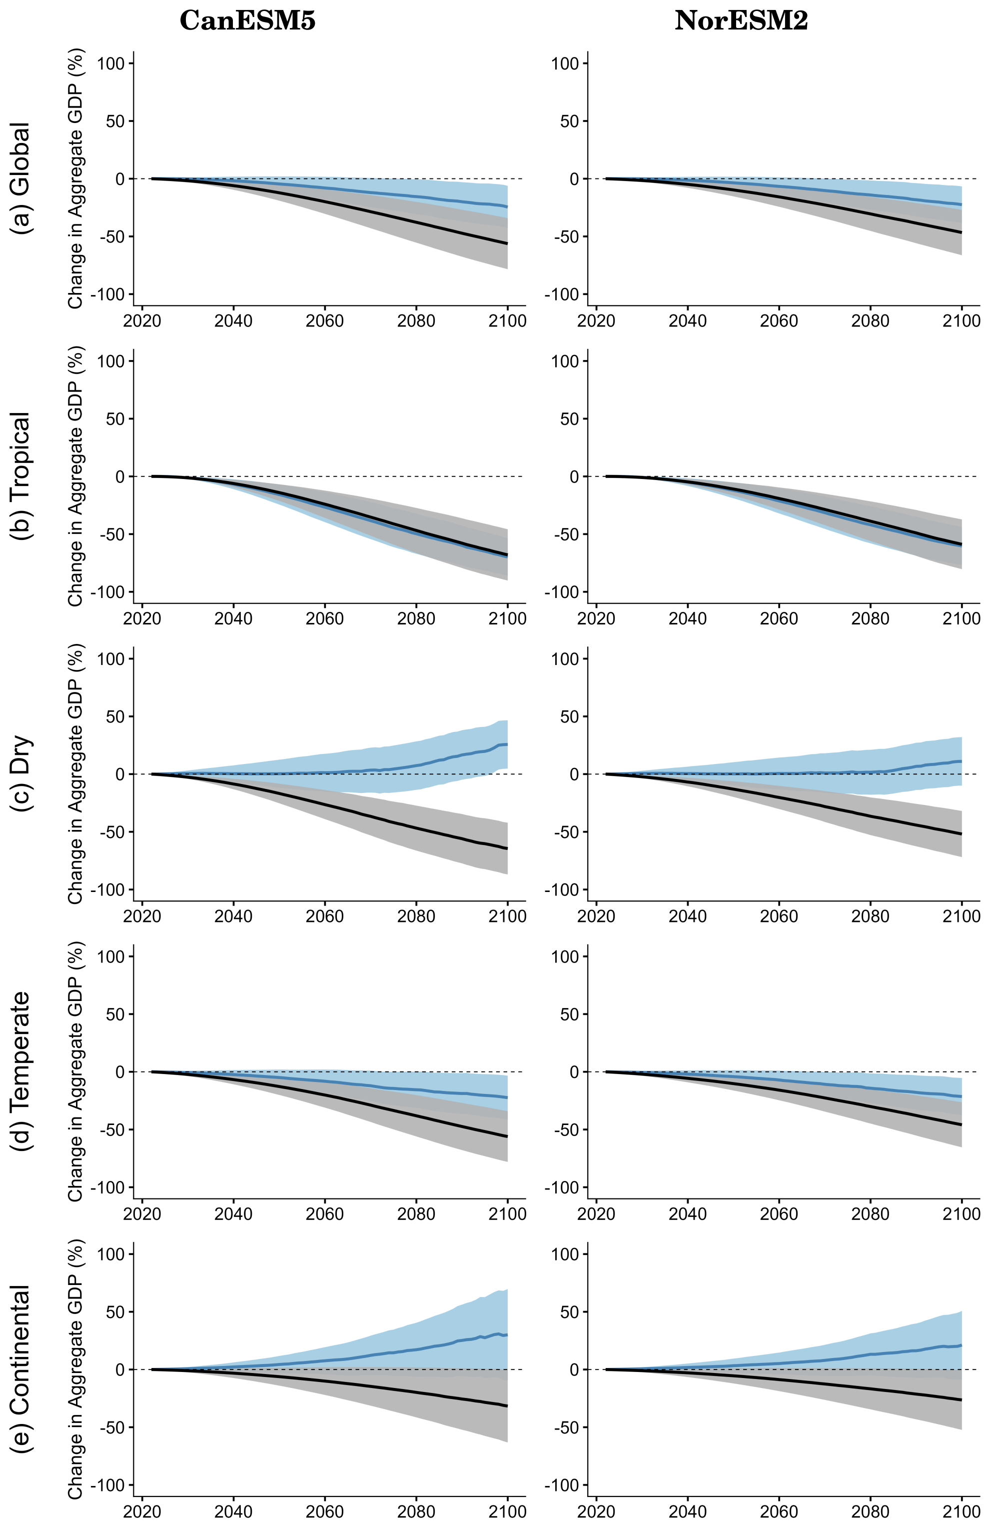

Figure A.2. Projected Aggregate Income Effects of Warming under the SSP2-4.5 Scenario by Climate Class

Notes: Plotted are projected % changes in aggregate income due to warming from 2022 to 2100 based on end-of-century climate projections under the SSP2-4.5 scenario from CanESM5 (left) and NorESM2-MM (right) model for the full global sample and separately by major Köppen-Geiger climate class. For any group G and year τ, the plotted effect is $\left[\sum_{i\in G} L_{i\tau} y_{i\tau}^{\text{no/warm}}\right]^{-1} \sum_{i\in G} L_{i\tau}\left(y_{i\tau}^{\text{w/warm}} - y_{i\tau}^{\text{no/warm}}\right) \times 100$, where $y_{i\tau}^{\text{w/warm}}$ and $y_{i\tau}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively, and $L_{i\tau}$ is projected population. The purple line reports the implied aggregate effect based on the spatially varying coefficient model, whereas the black line reports the corresponding effect from the globally fixed-coefficient specification. Shaded bands are two-sided 95% confidence intervals. To minimize outlier distortions, locations in the bottom and top 2.5% of income projections were omitted before aggregation.

Table A.1. Projected Mid- and End-of-Century Aggregate Income Effects of Warming by Climate Class under the SSP2-4.5 Scenario

	Global	Tropical	Dry	Temperate	Continental
			—CanESM5—		
2050 Impact Projection					
Accounting for Local Context	–4.54	–15.72	0.34	–4.98	4.36
	(–11.08, 2.01)	(–24.39, –7.05)	(–11.82, 12.51)	(–11.9, 1.95)	(–3.03, 11.74)
Imposing Global Homogeneity	–12.20	–14.03	–16.79	–12.76	–6.21
	(–19.06, –5.35)	(–21.22, –6.84)	(–25.27, –8.31)	(–20.07, –5.45)	(–13.99, 1.58)
2100 Impact Projection					
Accounting for Local Context	–24.56	–69.78	25.7	–22.4	30.11
	(–42.96, –6.15)	(–85.93, –53.64)	(4.78, 46.62)	(–41.49, –3.32)	(–9.46, 69.67)
Imposing Global Homogeneity	–56.28	–67.98	–64.5	–56.09	–31.68
	(–78.41, –34.15)	(–90.2, –45.77)	(–86.93, –42.08)	(–78.06, –34.11)	(–63.18, –0.17)
			—NorESM2—		
2050 Impact Projection					
Accounting for Local Context	–3.57	–12.07	0.25	–4.17	3.31
	(–8.84, 1.71)	(–18.96, –5.17)	(–9.83, 10.32)	(–9.73, 1.38)	(–2.85, 9.47)
Imposing Global Homogeneity	–9.74	–10.94	–13.03	–10.24	–5.45
	(–15.31, –4.16)	(–16.59, –5.28)	(–19.81, –6.25)	(–16.2, –4.29)	(–12.07, 1.17)
2100 Impact Projection					
Accounting for Local Context	–22.50	–60.33	11.04	–21.42	21.14
	(–38.35, –6.65)	(–76.75, –43.92)	(–9.95, 32.04)	(–37.46, –5.37)	(–8.5, 50.78)
Imposing Global Homogeneity	–46.67	–58.82	–51.91	–45.91	–26.37
	(–66.32, –27.02)	(–80.35, –37.28)	(–71.89, –31.94)	(–65.4, –26.41)	(–52.3, –0.44)

Notes: Reported are projected % changes in aggregate income due to warming from 2022 to 2050 and to 2100 based on mid- and end-of-century climate projections from CanESM5 and NorESM2-MM models under the SSP2-4.5 scenario for the full global sample and separately by major Köppen-Geiger climate class. For any group G and year $\tau = \{2050, 2100\}$, the reported effect is $\left[\sum_{i \in G} L_{i\tau} y_{i\tau}^{\text{no/warm}}\right]^{-1} \sum_{i \in G} L_{i\tau} \left(y_{i\tau}^{\text{w/warm}} - y_{i\tau}^{\text{no/warm}}\right) \times 100$, where $y_{i\tau}^{\text{w/warm}}$ and $y_{i\tau}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively, and $L_{i\tau}$ is projected population. Two-sided 95% confidence intervals are in parentheses. To minimize outlier distortions, locations in the bottom and top 2.5% of income projections were omitted before aggregation.

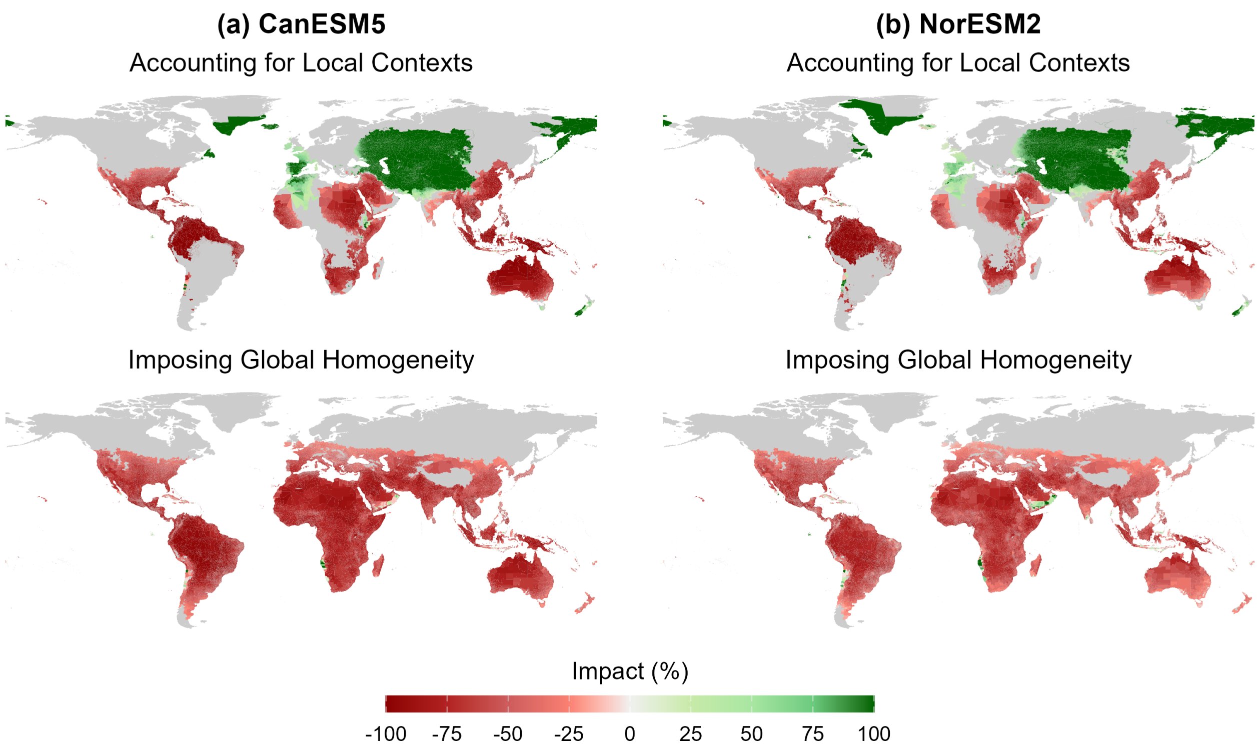

Figure A.3. Geographic Distribution of End-of-Century Projections of Local Per-Capita Income Effects of Warming under the SSP2-4.5 Scenario

Notes: Mapped are location-specific projections of per-capita income effects of warming by the year 2100: $(y_{i2100}^{\text{w/warm}} - y_{i2100}^{\text{no/warm}})/y_{i2100}^{\text{no/warm}} \times 100$, where $y_{i2100}^{\text{w/warm}}$ and $y_{i2100}^{\text{no/warm}}$ are projected GDP per capita with and without warming, respectively. Panels (a) and (b) use end-of-century climate projections under SSP2-4.5 from the CanESM5 and NorESM2 models, respectively.The top-row maps plot estimates based on the spatially varying coefficient specification, which allows the temperature-growth relationship to differ across locations as a function of local context, whereas the bottom-row maps report the corresponding projections from the globally fixed-coefficient specification. Statistically—at the 5% level—positive (negative) projected effects are shown in green (red); the remaining locations are in gray. Estimates are in %.