Global income distributions under climate change

W. Matthew Alampay Davis*

July 22, 2025
link to latest version

Abstract

Understanding how climate change impacts global economic inequality is critical to the design of equitable and politically viable climate policy. Yet where evidence of this relationship is available, analysis is generally limited to coarse comparisons of country-level aggregates, remaining agnostic to disparities across individuals within those countries. This paper advances this literature using newly available distributional income data to incorporate within-country dimensions of this 'climate inequality'. First, I document new evidence that temperature shocks persistently exacerbate inequalities within countries by disproportionately reducing income growth among the lowest income-earners, especially in warm-climate economies. Second, I use these empirical results to compare the historically observed distribution of global income to its counterfactual distribution had global warming been stabilized at 1980 levels. I estimate that systematic warming between 1980-2016 increased global income inequality by 4.0% [1.6,6.6], with absolute increases in between-country and within-country inequality contributing equally to this total effect. These contributions in turn correspond to proportional increases in inequality of 2.6% [0.0, 5.6] within countries and 8.7% [4.9, 13.3] between countries. To my knowledge, these distributional results constitute a novel contribution to the climate impacts literature and altogether offer the most comprehensive evidence yet of the globally regressive economic impact of climate change.

*wm.alampaydavis@gmail.com

This work was previously circulated as my job market paper titled "Climate inequality"
I acknowledge funding support from Columbia's Program for Economic Research.

"Note that the Bengal famine might have been a compensation test victory because a lot of people gained a lot in 1943. The new destitutes mostly died, but all we are checking is whether they *could* have been compensated. I mean, what kind of 'improvement' is that? There is no progress trying to do welfare economics without taking on the real problems of inequality and poverty. This could not be done, then or now."

Amartya Sen (Sen et al. 2020)

1 Introduction

Addressing anthropogenic climate catastrophe and narrowing economic inequalities between and within countries are widely recognized as urgent and interrelated global priorities (United Nations 2015; Masson-Delmotte et al. 2018; Chancel, Piketty, et al. 2022; Pande 2023). Indeed, a robust empirical literature has consistently found that climate change disproportionately burdens near-tropical, typically less developed economies on average, widening disparities between the 180 or so countries (Ricke et al. 2018; Diffenbaugh and Burke 2019). Yet, alarmingly little evidence exists relating climate change to inequality among the billions of *people* resident within these national borders[2]. If these within-country impacts are also substantial and regressive, as most environmental economists believe to be likely (Howard and Sylvan 2021), then our present understanding of climate change is agnostic to its most important social implications: global mass immiseration and social stratification.

Addressing this ignorance is thus critical to the design of a sufficient global response to the climate crisis (Masson-Delmotte et al. 2018). For example, economists have long advocated for mitigation policy centered around pricing instruments like carbon taxes or quotas which adjust the private costs of marginal emissions to reflect their total social costs (Weitzman 1974). Classically, a pricing schedule which achieves this internalization is said to be efficient in the limited Kaldor-Hicks sense that enough aggregate benefit is obtained as a result of the policy to offset the aggregate losses from unmitigated emissions, regardless of how those benefits are distributed. In recent years, demand for more intentional equity considerations have led to institutional departures from the pure utilitarianism of traditional cost-benefit analysis, allowing the least well off to be assigned greater welfare weighting (Management and Budget 2023; Prest et al. 2024). But even with these refinements, mitigation policy is only efficient in the welfare-preserving sense to the extent that those disproportionately burdened by the marginal emission and the incidence of the pricing instrument itself are sufficiently compensated (Sen et al. 2020). Thus, any efficient and equitable climate policy requires actionable evidence identifying the distribution of climate impacts.

Consideration of within-country inequality impacts is likely to be pivotal for overcoming resistance to ambitious climate policy proposals. The historical unpopularity of carbon pricing mechanisms are well-documented and often explicitly motivated by distributional concerns, as dramatically demonstrated by the 2018 Gilets Jaunes protests opposing impending fuel tax increases. A growing body of social science research corroborates a more general positive association between public support for national climate policies and their perceived progressivity (Maestre-Andrés et

[2]Within-country income inequality now accounts for almost 70% of total global income inequality as measured by the entropy-based Theil index, an increase from just over 40% in 1980. (Chancel, Piketty, et al. 2022)

al. 2019) or their bundling with pro-poor social and structural policies such as a jobs guarantee or affordable housing reform (Bergquist et al. 2020). Dechezleprêtre et al. (2025) document experimental evidence that these equity considerations are much more salient than aggregate climate impacts and Furceri et al. (2023) suggest that successfully implemented carbon pricing policies are only politically costly in periods of heightened inequality.

Despite these first-order policy implications, the literature relating climate change to economic inequality *within* countries remains sparse mainly due to the historical unavailability of reliable and comparable distributional data (Atkinson 2001). The far more active literature on between-country climate inequality is made possible by the regular measurement of country-level aggregates such as gross domestic production which is rigorously standardized by the System of National Accounts specifically to facilitate international and intertemporal comparisons (United Nations 2010). By one account, the greater academic, commercial, and political interest in aggregate production has crowded out investment in research infrastructure to support the construction of distributional welfare measures (Jorgenson 2018). While secondary inequality datasets providing estimates of country-year inequality date at least as far back as Deininger and Squire (1996), the value of these products for causal analysis and cross-country comparison is undermined by substantial inconsistencies across input sources in spatiotemporal coverage, how inequality is operationalized, and the opacity of their construction (Atkinson and Brandolini 2001; Jenkins 2015). Nonetheless, Cevik and Jalles (2023) and Gilli et al. (2024) are early efforts using these data products to empirically study temperature impacts on the distribution of economic outcomes for multiple countries. Other works such as Hsiang et al. (2017) , Marx (2024) , and Dasgupta et al. (2023) make important progress by instead using rich but idiosyncratic distributional data from individual countries to document evidence of regressive climate incidence within the United States, France, and South Africa respectively. Warmer developing countries where aggregate climate impacts are known to be greatest generally lack the capacity to produce similarly fine-grained data.

I am able to overcome these myriad empirical limitations by taking advantage of newly available income data from the 2024 update to the World Inequality Database. As described in greater detail in Section 2, this new generation of distributional data represents a major advance in the measurement of global inequality because its construction, modeled after the System of National Accounts, systematically integrates information from various primary sources to capture national income distributions more completely and transparently. The final dataset used throughout this paper takes the form of a well-balanced annual panel of country-quantile pre-tax income since 1980 and critically spans an expanded set of countries that now extends to most developing countries. The present analysis of within-country climate impacts thus benefits from the unprecedented granularity and near-global extent of this data, overcoming the steep tradeoff between subnational resolution and global generalizability that has constrained existing approaches. As a result, I am able to document new evidence that identified temperature shocks impose large, regressive, and persistent effects on income distributions, especially in warmer geographies. I then incorporate these empirical results into a counterfactual analysis which allows me to estimate the regressive effect of climate change on global income inequality decomposed into its between-country and within-country components for the first time.

My empirical approach begins with a brief discussion of data in Section 2.3, where I describe how I construct temperature shocks to address recent critiques that traditional constructions fail to satisfy necessary conditions for validity as an explanatory variable. Section 3 describes how I use the method of local projections to estimate "cumulative temperature multipliers" which isolate the dynamic cumulative effects of a unit temperature shock on income while accounting for autocorrelation in the shock itself. In practice, this entails simply estimating a series of single-equation regressions for every income quantile and projection period of interest then adapting recent econo-

metric innovations to make inferences about the persistence of this effect (described in more detail in **??**.). I argue this intuitive semi-parametric approach is particularly well-suited to overcoming important limitations in the existing climate impacts literature. Empirical results in Section 3.2 depict how the expected impacts of transitory 1°C temperature shocks vary across income groups, how these income gradients in turn vary across countries, and how these effects accumulate and dissipate over time. More specifically, I report evidence that a pulse of temperature has substantial and persistent impacts on distributions of national income which are particularly pronounced and regressive in warmer countries. For the bottom decile in a representative 26° country, the average effect of a unit temperature shock is a 7% reduction in cumulative effect over the ensuing three years; this compares to a smaller loss of 2.2% for the median adult in the same country and a 3.0% gain for the bottom decile in a representative $10^c irc$ country.

Section 4 summarizes an analysis comparing the observed evolution of the global distribution of income between and within countries to a counterfactual in which global warming was held fixed at 1980 levels. This entails applying bootstrapped estimates of the cumulative temperature multipliers to historical income and weather data to construct a range of plausible and compounding outcomes absent systematic warming for each year from 1980 to 2016. My results imply that systematic warming over these 37 years has unambiguously exacerbated total global inequality, both between and within countries. This effect mainly stems from foregone income growth for the poorest 20% of global income-earners. For the lowest positive income-earning global percentile in our data, the analysis implies that incomes today would be 29% higher than observed had the climate been stable over this period. The evidence is also consistent with the standard finding in the literature on aggregate impacts that climate change is extremely likely to have reduced total global income over this period. I find that climate change over this period increased total global inequality by 4.0% [1.6, 6.6] as measured by the mean logarithmic deviation index with increases to within-country and between-country inequality contributing essentially equally to this total effect. As within-country inequality now comprises a dominant share of total inequality, these represent proportional increases of 2.6% [0.0, 5.6] and 8.7% [4.9, 13.3] respectively. To the best of my knowledge, these estimates represent a novel contribution to the vast climate impacts literature. Section 5 concludes.

1.1 Related literature

Spurred by methodological advances, accessibility of computing power, and political urgency, the past decade has seen a proliferation of economic research dedicated to estimating the social impacts of climate change. To contextualize the relative sparsity of within-country distributional impact studies in climate economics, it is instructive to familiarize the reader with the subfield's taxonomy of empirical designs into complementary "top-down" and "bottom-up" approaches each premised on an underlying utilitarian cost-benefit logic.

Top-down approaches restrict attention to a widely available social outcome of interest. This outcome is taken as a comprehensive enough proxy for social welfare that climatically attributable impacts to the stock of this outcome may represent the socially relevant climate "damage". The pioneering works in this category include Dell et al. (2012) and Burke, Hsiang, et al. (2015), which relate annual observations of country-level temperature to GDP growth. Empirical designs in this category commonly take the form fixed effects regression models run on country-year panels. The advantage of the former approach is in its convenience and interpretability: the wide availability of cross-country economic aggregates such as those provided by national accounts substantially simplifies the task of calculating global climate impacts.

Alternatively, bottom-up approaches consider impacts to distinct "sectors" of an economy;

mortality (Carleton, Jina, et al. 2022), energy consumption (Rode et al. 2021), and labor (Graff Zivin and Neidell 2014) are common areas of focus, for example. Results using this approach can be presented as partial costs or combined with other sectoral damages in an integrated model which accounts for their interdependencies. The benefit of this approach is in its ability to account for disaggregated and non-market impacts to which single-dimensional aggregates are largely indifferent.

The simplicity of the top-down approach and comprehesiveness of the bottom-up approach pose a critical tradeoff between geographic disaggregation and global coverage. As perhaps the most prominent existing work directly concerned with characterizing global 'climate inequality', Diffenbaugh and Burke (2019) estimates that the ratio between the top and bottom deciles of the global income distribution was 25% larger in 2010 than it would have been had anthropogenic climate forcings been constrained to 1960 levels, all else held equal. However, this quantity is calculated by assigning all individuals their country's mean GDP and crudely defining global deciles accordingly, an unavoidable restriction when limited to 160 or so distinct economic units. The authors remark that "documenting the impact of global warming on within-country inequality remains an important challenge."

Hsiang et al. (2017) is a particularly impressive example of a study which makes the much less common opposite compromise in this tradeoff. Integrating fine-scaled agricultural, crime, energy, mortality, and labor data, the authors determine that climate risks in the continental United States are disproportionately borne by counties in the US South and Midwest, regions generally already poorer than their coastal counterparts; climate change then is likely to imply worsening inequality within the United States. But to make more generalizable claims about the implications of global warming for global inequality, analyses must extend to settings which lack comparably high-quality and granular data. Elsewhere, Marx (2024) uses rich tax data from France to infer that years with marginally more days above 30°C are associated with wider income disparities between the richest and poorest cantons (administrative units described as roughly equivalent to 10 municipalities). Dasgupta et al. (2023) uses data from perhaps the most unequal country in the world to project a 3-6 Gini point increase in South Africa's income distribution by 2100 under a moderate warming scenario.

The distributional focus of this paper complements a much more active area of research measuring unequal impacts across countries. Like Diffenbaugh and Burke (2019), a subset of these studies explore the implications of this type of inequality directly, for example by highlighting the inverse relationship between country-level contributions to global greenhouse gas emissions and exposure to their social consequences (Chancel and Piketty 2015; Ricke et al. 2018; Kotz et al. 2024). There is by now general agreement that non-linearities in the response of aggregate variables to marginal warming concentrate damages and risk in relatively warm and poor countries which have contributed proportionally little to anthropogenic warming (Howard and Sylvan 2021). However, estimated magnitudes of these impacts are notoriously divergent, largely because of unsettled disagreement over the persistence of these non-linear effects (Newell et al. 2021). We draw from recent work by Nath et al. (2024) and Bilal and Känzig (2024) which use less restrictive time series methods to circumvent these statistical bugbears.

Similar non-linearities emerge at the micro level across a range of outcomes as varied as labor supply (Graff Zivin and Neidell 2014), agricultural yields (Schlenker and Roberts 2009), and standardized test performance (Park et al. 2020). Carleton and Hsiang (2016) overviews the proliferation of multidisciplinary studies in this category whose results describe patterns of setting-specific inequality suggestive of a broader micro-level environmental inequality similar to those observed in the aggregate. The present study bridges this micro evidence with the aforementioned macro evidence on inequalities between countries, estimating distributional responses consistent

with the general features of both categories of impacts. We also document evidence of negative impacts to an elite minority, which I have not seen examined in existing work.

The literature on climate impacts on economic inequality *within* countries has remained comparatively sparse in large part due to an absence of adequate research infrastructure. Estimation of distributional equivalents of the aggregate damage functions standard in the literature requires data comparable in credibility and coverage to domestic production data whose construction across countries follows rigorous standards set forth by the international System of National Accounts (United Nations 2010). Historically, distributional measurement of welfare has lacked and arguably been crowded out by the wider interest and investment in measuring aggregate production (Jorgenson 2018). While secondary inequality data products compiling country-year inequality variables date at least as far back as Deininger and Squire (1996), the validity of these data products for cross-country causal analysis is limited by inconsistencies across primary sources in construction and availability (Atkinson and Brandolini 2001). A recent empirical evaluation of the most widely used of these datasets concluded they are "not sufficiently credible" (Jenkins 2015). In the absence of more credible global inequality measures, Cevik and Jalles (2023) and Gilli et al. (2024) estimate the model of Burke, Hsiang, et al. (2015) replacing their GDP growth outcome variable with the inequality variables provided in these early-generation products.

The research programs of Thomas Piketty, Emmanuel Saez, and their co-authors since the turn of the century are widely credited with revitalizing public and academic interest in the measurement of within-country inequality. Beginning with work tracing the long-run evolution of income and wealth inequality in France (Piketty 2003) and the United States (Piketty and Saez 2003), their mobilization of extensive tax returns data has enabled credible measurement of top incomes over long time horizons, stimulating new programs of research and equipping the general public with new political language around which to mobilize (Jones 2015). This paper takes advantage of the cross-country research apparatus that has rapidly developed and been made accessible in recent years with the intention of maturing the economic analysis of climate change in a similar manner. Palagi et al. (2022) is the only other climate-economic study I could locate which makes use of the same " accounts" (The World Inequality Lab 2024) used in this study. The authors find that extreme precipitation negatively impacts low-income individuals much more severely in countries more dependent on agriculture.

Finally, I locate this paper in the context of a growing demand for distributional considerations in economic analyses of climate change and the design of climate policy. For example, Chancel, Bothe, et al. (2023) estimates that within-country inequality of carbon emissions have recently exceeded the equivalent "carbon inequality" between countries and advocates for targeted carbon taxation programs based on individual footprints to more efficiently internalize the climate change externality. The modern SCC framework accommodates equity considerations by concavifying the underlying utility function so that the welfare weights assigned to individuals are inversely related to their relative levels of consumption; a logarithmic specification is commonly chosen, for example. Variations of this more progressive utilitarianism have been institutionalized to some degree in a handful of countries including recently in the United States (Management and Budget 2023). Theoretically, the way equity weighting is implemented should result in substantially higher SCC estimates depending on the extent to which the distributive effect of climate change is modeled as regressive. This paper provides the first comprehensive estimate of the global regressivity of climate change at subnational resolution.

2 Data

2.1 Distributional income data

We make use of the recently published grouped income series from the World Inequality Database (The World Inequality Lab 2024). Its predecessor, the World Top Income Database launched in 2011, drew from comprehensive tax records from over 30 mostly Western countries to capture the distribution of income among their top income-earners with unprecedented precision. A critical innovation used to develop the new product is the systematic integration of this tax data with national accounts and household survey data in order to impute the full income distribution within those countries (Piketty, Saez, and Zucman 2019) in a manner patterned after the System of National Accounts (United Nations 2010). The newest generation of WID data extends coverage within those countries—now claiming to capture "100% of income" in the United States (Piketty, Saez, and Zucman 2018)—and to countries with relatively limited available data. By 2019, historical data for over 120 countries had been incorporated into the database. As of this writing, their equal-split adult pre-tax income series was available for 215 countries and territories collectively representing almost the entire global population with distributional data imputed every year since 1980. All income estimates are converted to 2023 US dollars using the provided PPP-adjusted exchange rates. Accompanying metadata scores the availability and reliability of each category of input source by country, unsurprisingly indicating that the input mix of these sources varies substantially by country with less developed countries being more reliant on household surveys and wealthier countries on tax records.

Notably for the percentile-level panels, the income data exhibits substantial bottom-coding with almost every country-year reporting exactly zero income for at least the bottom four percentiles of the pre-tax income distribution. Even if an accurate measure of adult unemployment, this evokes well-known deficiencies of income as an identifier and measure of poverty compared to alternatives like consumption expenditures (Meyer and Sullivan 2003). For example, while the unemployed subpopulation of a given economy may well be composed of its most impoverished, it is likely to also include unemployed individuals of all levels of precarity and social status and does not account for security provided by communal and social safety nets. It also positively biases estimated changes in welfare among the poor since zero incomes mechanically cannot decrease.

Given these inherent disadvantages of income measures as a proxy for relevant welfare, we intentionally do not make precise claims about the poverty implications of our findings. Instead, we consider the incomes of the bottom 10 percentiles as a group in all analyses and consider estimated impacts for this quantile to be a conservative bound on true impacts on the very poorest because it implicitly assumes perfect equality among its constituent population (Cowell 2011).

2.2 Historical weather data

We merge all income data with an annual country-level panel of population-weighted near-surface air temperature data derived from the Global Meteorological Forcing Dataset for Land Surface Modeling (GMFD), which corrects model-based biases in NCEP/NCAR reanalysis product using observational meteorological station data (Sheffield et al. 2006). The GMFD reports historical weather data for the entire global surface at 0.25-degree resolution[3] for every three hours from the start of January 1, 1948 to the end of December 31, 2015. We use the product which

[3]At this resolution, cells near the equator are approximately 27 km^2 (17 mi^2).

reports daily averages of near-surface air temperature.

While other temperature series of similarly high resolution can provide a larger estimation sample by beginning before 1948 or updating more recently than 2015, results reported in the broad climate impacts literature have been generally consistent across data products. We favor the GMFD product as it has the unique feature of being the only reference product used to bias-correct and downscale the CMIP6 projection data series produced by the NASA Earth Exchange Global Daily Downscaled Projections; this allows the results of the present analysis to seamlessly integrate with future work exploring its implications under different future warming scenarios or attribution studies which attempt to decompose estimated climate change impacts into their anthropogenic and natural contributions, as introduced in **??**.

2.3 Temperature as an explanatory variable

It has been customary in the climate impacts literature to use functions of observed weather variables such as temperature or precipitation as explanatory variables in causal analyses. Historically, their use in panel settings has been justified by an intuition that annual fluctuations in these geophysical conditions are plausibly exogenous with respect to most economic activity. This intuition has been increasingly challenged by works such as Nath et al. (2024), which argues that because levels of temperature do not satisfy conditions of an identified shock[4] and exhibit substantial autocorrelation, they are ill-suited for dynamic causal inference.

Addressing these concerns, I define temperature shocks $\widehat{\tau}_{it}$ as deviations in temperature levels relative to a climate 'state' variable $\overline{T}_{it}$. This state variable is in turn defined as the M-period moving average of local temperature $\overline{T}_{it}$:

$$\overline{T}_{it} := \frac{1}{M} \sum_{m=1}^{M} T_{i,t-m} - 18 \tag{1a}$$

$$\widehat{\tau}_{it} := T_{it} - \overline{T}_{it} \tag{1b}$$

This definition of the state variable $\overline{T}_{it}$ can be contrasted against the time-invariant state variable used in Nath et al. (2024) where the climate state variable is constant for each country. I favor this moving-average construction because the median country in my sample observes, on average, a 0.26°C (IQR: 0.18-0.36) increase in this state variable every 15 years as a result of global warming trends. Additionally, this definition of the state variable has an appealing coherence with the convention in climatology of using long-run moving averages of weather realizations to define the slow-moving state of the climate at a fixed period. Most commonly, these "climate normals" are defined as 30-year averages to accommodate natural variation which would occur even in a stable climate[5] and so we set $M = 30$ in our main analysis; Kahn et al. (2021) is another economics study which uses the same construction.

[4] As outlined in the handbook chapter of Ramey (2016), a valid shock in a time series setting must i) measure unanticipated movements in relevant exogenous variables, ii) be exogenous with respect to current and lagged endogenous variables, and iii) be uncorrelated with other exogenous shocks included in the model. Temperature specified in levels do not satisfy the first condition and the third condition is violated in commonly used non-linear specifications where temperature enters as a polynomial.

[5] The most economically relevant example of this natural variability is the El Niño Southern Oscillation, an ocean-warming phenomenon which can vary average global surface temperatures by as much as 0.4°C across its three phase cycles which may span 2-7 years. Other minor oscillations can span multiple decades.

We subtract 18 in Equation (1a) so that the main effects we attribute to the temperature shock τ correspond to an intermediate representative climate of 18°C, roughly the 40th percentile of average temperatures in our country-year data.

Figure 1: Country-level time series of identified temperature shocks

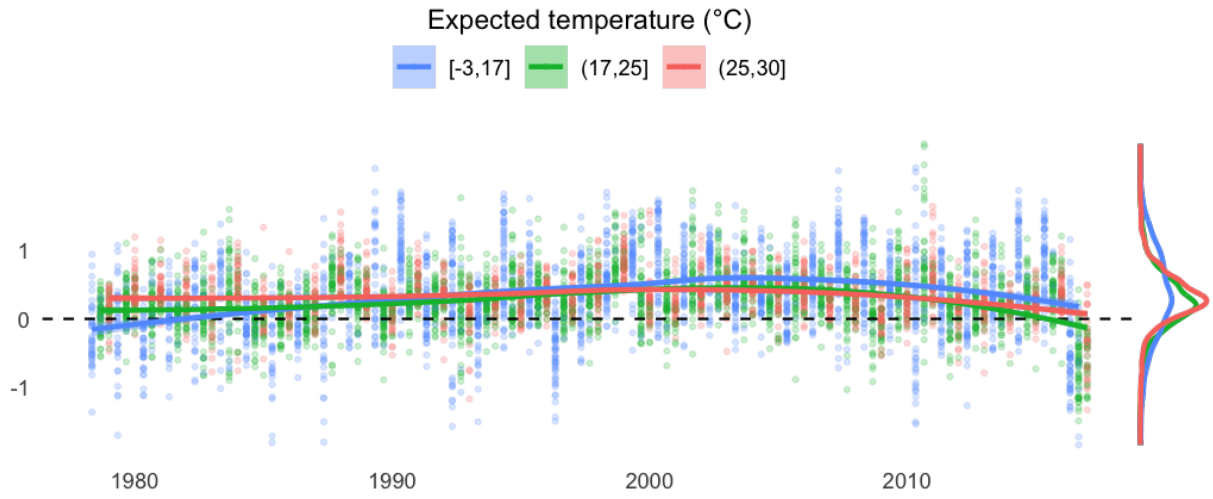

Points correspond to shocks observed in individual country-years. The three colors correspond to an equal-sized binning of country-years according to their moving-average temperature. Flexible local regression curves describe time trends for each bin; segments with more positive values correspond to periods of greater systematic warming. Marginal distribution plots on the right show that cooler country-years observe greater variance in shocks and more positive values.

Figure 1 depicts the implied history of temperature shocks $\widehat{\tau}_{it}$ constructed in this manner with individual points corresponding to distinct country-year observations. These observations are divided into three equally sized groups according to their corresponding climate states $\overline{T}_{it}$; points and local regression curves are colored accordingly. The marginal distribution plots on the right demonstrate that colder seasonal countries experience more positive and more volatile temperature shocks than their warmer counterparts, a reflection of both greater natural weather variability in these settings and the fact that absolute anthropogenic warming is greater in regions closer to the Earth's poles[6].

The positive values of the curves in Figure 1 over almost the entire observation period captures the idea that systematic warming is characterized by the increased frequency and greater magnitude of positive (hot) shocks relative to negative (cold) shocks. Intervals where the curves are closer to 0 thus reflect relative balance in the incidence of hot and cold shocks, corresponding to periods of relative climate stability and slower warming.

[6]Figure A5 provides a more direct illustration of this phenomenon for different global climate models

2.4 Modeling the persistence of social impacts: levels vs. growth effects

Comparisons across columns in Figure 2 illustrate the conceptual distinction between a shock which imposes transitory level effects (first column) and permanent growth effects (third column); the middle column describes a middle ground where an effect is persistent but eventually vanishing. Effects are measured in logged levels (top panels) and first logged differences (bottom), a convenient approximation of a growth rate. Filled lines depict the evolution of the outcome variable in the absence (black) and presence (red) of a unit shock which occurs in period 0. Filled areas in orange measure the cumulative loss (or gain) in the outcome which may be attributed to the shock.

Figure 2: Levels and growth under different persistence structures

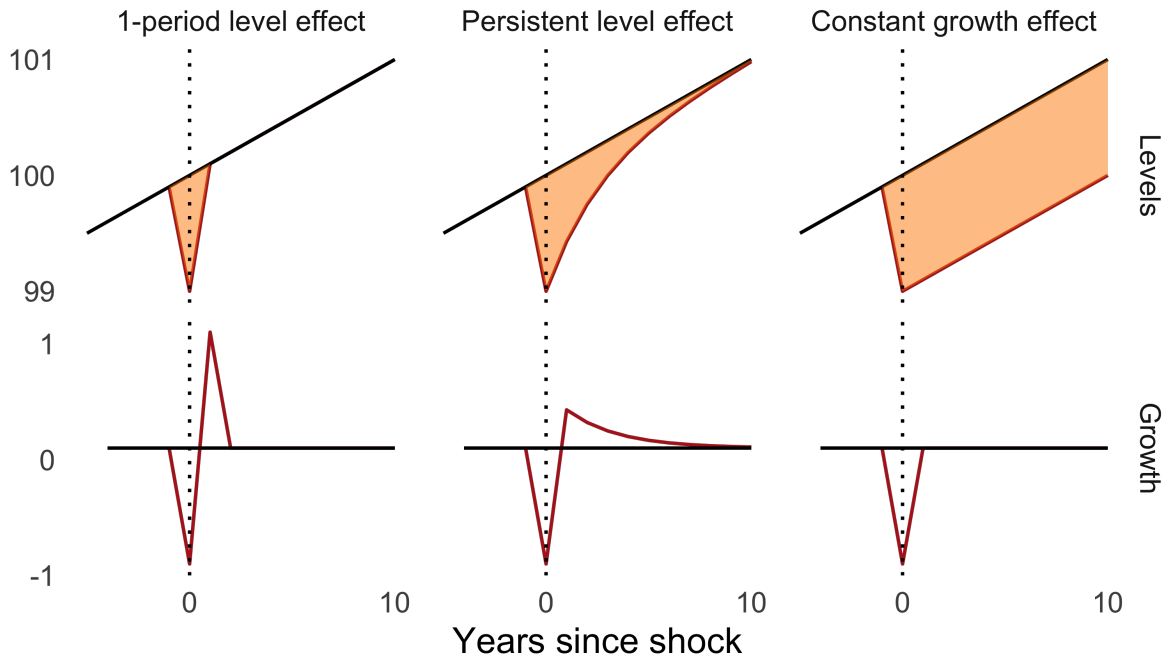

Illustrative diagrams depicting annual levels (top row) and growth (bottom row) of a social variable of interest. Dotted vertical lines mark the timing of a transitory temperature shock. Red lines represent trajectories resulting from a transitory temperature shock while black filled lines represent their counterfactual trajectory had they continued their pre-shock linear trend. Filled areas in orange correspond to the implied cumulative loss in the outcome variable ("damages") attributable to the shock. Columns correspond to different assumptions about how persistent the impact of a temperature shock is on the social outcome.

Under a one-period level effect, the level of, say, GDP falls by an estimated level $\widehat{\beta}_0$ in the year of the shock but fully recovers to its pre-shock trend (black) in the next period. In this scenario, the production lost in the year of the shock is exactly offset by an above-trend rate of growth which exactly offsets the negative growth effect in the year of the shock. The shock has no additional effect on future production and the area in orange represents the one-off impact from the shock.

Under the constant growth effect depicted in the rightmost panels, the same shock incurs the same effect on the levels but this effect is permanent so that GDP in all future periods is reduced by the same amount relative to the no-shock counterfactual. The difference between the two tra-

jectories thus compound over time as observed by the orange area which grows linearly with the projection horizon. The corresponding growth effect depicted in the bottom panel is entirely transitory with no compensatory positive growth relative to trend.

The middle column describes an intermediate case where the economy requires several periods to return to its pre-shock trend. Cumulative differences attributable to the one-off shock continue to rise over time but to a diminishing degree until eventually compensatory positive growth exactly offsets the initial negative growth, i.e., the integral of the growth curve equals zero in the long run.

Now consider the following two models, which we may refer to as the levels and growth models respectively.

$$\log Y_{it} = \beta^L X_{it} + v_{it} \tag{2a}$$

$$\Delta \log Y_{it} = \beta^G X_{it} + u_{it} \tag{2b}$$

For simplicity, imagine the outcome variable Y_{it} represents a conventional economic aggregate like GDP and X_{it} a perfectly identified temperature shock. In the levels model, β^L represents an estimate of the change in log GDP attributable to a unit shock. Graphically, this impact is represented by the vertical distance between the black and red lines in the top-left plot. Since this is an entirely static model, there is no mechansim for X_{it} to affect future values of Y_{it} and so the model assumes a perfect return to trend in all ensuing periods. In the growth plot directly below, this is represented by a $\beta^L - 0 = \beta^L$ effect in period 0, an exactly offsetting $0 - \beta^L = -\beta^L$ effect on growth in period 1, and no effect on growth in all ensuing periods. In models specified in levels, level effects which persist for p periods can be accommodated by including p additional lags of X in the regression. A permanent growth effect can only be accommodated with the inclusion of infinite lags.

In the growth model, $\beta^G \neq 0$ represents a permanent change in the difference between consecutive values of Y_{it}. This is represented by the gap between the black pre-shock trend line and the observed GDP series in the top-right plot staying constant at β^G for all periods following the shock and is never offset. Level effects which persist for p periods can be accommodated by including p lags of X and finding that the sum of the $p + 1$ coefficients is exactly 0 so that the initial negative growth effect is eventually offset by sufficient compensatory growth.

The differing restrictions imposed by these stylized regression models is a simplification of perhaps the most influential unsettled question in climate economics: the persistence of climate impacts. For example, two studies projecting the proportional loss in cumulative global production by 2100 arrive at point estimates of -20% (Burke, Hsiang, et al. 2015) and -3.4% (Casey et al. 2023) under the same worst-case warming scenario. Methodologically, the two studies use similar methods and data except for their choice of outcome variable.

Intermediate persistence structures where level effects are diminishingly persistent can be accommodated by either model by including sufficient lags of the explanatory variables. Of course, identifying variation is reduced with the addition of additional explanatory variables, a problem exacerbated in specifications where shocks enter non-linearly and settings where shocks exhibit significant autocorrelation, both general features of modern climate models using absolute temperature as their explanatory variables. In addition, models specified in levels are notoriously vulnerable to unit root concerns, further undermining credible identification.

Indeed, the appendix to Burke, Hsiang, et al. (2015) concludes that the identifying variation which remains after lag augmentation to test for persistence "cannot reject the hypothesis that this effect is a true growth effect nor the hypothesis that it is a temporary level effect", rendering these

extremely influential assumptions effectively unfalsifiable. Absent substantially improved data, overcoming this issue requires richer data than provided by national accounts or an alternative methodological approach relatively robust to misspecification.

3 Empirical analysis

3.1 Methods: cumulative temperature multipliers

To characterize the dynamic effects of an identified temperature shock on national income distributions, we adapt the method of local projections first introduced in Jordà (2005). As Jordà (2023) notes, this semi-parametric method of estimating impulse response functions has found growing appeal outside of applied macroeconomics in recent years as recent econometric literature has demonstrated its flexibility in incorporating non-linearities (Cloyne et al. 2023), its coherence with the potential outcomes framework (Dube et al. 2025), and its robustness to misspecification and non-stationarity (Montiel Olea and Plagborg-Møller 2021; Montiel Olea, Plagborg-Møller, et al. 2024; Piger and Stockwell 2025). Berg et al. (2024) , Bilal and Känzig (2024) , and Nath et al. (2024) are the first works I am aware of to apply this method to estimate dynamic economic responses to climate shocks, demonstrating how these features can help address the identification issues described in Section 2.

Our application of this method to our inequality setting entails the following steps:

1. Estimating the state-dependent response of income to a unit temperature shock
2. Estimating the state-dependent response of temperature shocks to a unit temperature shock
3. Using both to estimate a state-dependent multiplier to isolate the dynamic income response net of the shock's effect on future shocks
4. Statistically evaluating the significance and persistence of this multiplier
5. Repeating the above steps for each income quantile of interest

We briefly describe this process below and provide additional exposition in **??** using an application to country-level GDP data.

3.1.1 Estimating state-dependent impulse responses

We use country-year panel datasets to estimate dynamic responses to temperature for an outcome of interest represented by the variable y_{it}. For example, y may represent the average logarithm of income of the bottom 10% of individuals in country i in year t and we may repeat the estimation process for each income decile and use the results to infer a total effect on national income distributions.

The method of local projections entails estimating $H+1$ iterations of the following single-equation model, corresponding to different projection horizons $h \in \{0, 1, ..., H\}$ periods after an identified temperature shock:

$$\begin{aligned} \Delta_h y_{i,t+h} &:= y_{i,t+h} - y_{i,t-1} \\ &= \beta_{1h}\widehat{\tau}_{it} + \beta_{2h}\widehat{\tau}_{it} \cdot \overline{T}_{it} + \lambda_h \overline{T}_{it} + \boldsymbol{\gamma}_h^\top \mathbf{Z}_{it} + \mu_i + \eta_t + u_{i,t+h} \end{aligned} \tag{3}$$

In our main specification, the matrix $\mathbf{Z}_{it}$ includes lagged controls $\{\widehat{\tau}_{i,t-j}, \widehat{\tau}_{i,t-j} \cdot \overline{T}_{it}, \Delta y_{i,t-j}\}_{j=1}^{p}$. The inclusion of p lags of the temperature shock $\widehat{\tau}$ is included to address potential autocorrelation in the explanatory variable. Interactions of these lagged shocks with the state variable capture potential state-dependencies in these autocovariance structures motivated partially by the finding in Figure 1 that shocks in cooler climates tend to be more positive on average and feature greater year-to-year variability. Lags of the outcome are included as a guard against trend non-stationarity. Finally, the sets μ_i and η_t of fixed effects are included to absorb time-invariant unit-specific effects and common year-specific effects. In all, this general model differs from the state-dependent model of Nath et al. (2024) only in how temperature shocks are constructed and in permitting the state variable to vary over time as described in Section 2.3.

For each projection horizon h, the main regression coefficients of interest attribute to a unit temperature shock a main effect β_{1h} representing a change in the "long difference" between y observed h periods after the shock and y observed the period before the shock. Equivalently, this can be interpreted as the change in y h periods after the shock relative to a counterfactual absent the shock. In most of our results, y represents the logarithm of income for some quantile of the national income distribution; in these cases, coefficients are interpreted as percentage-point impacts on income.

An interaction term β_{2h} allows the effect at each period to scale with an observation's time-varying expected temperature $\overline{T}_{it}$, defined in Section 2.3. This 'state dependence' is motivated by the standard result in the climate impacts literature that the marginal effect of a given temperature shock on aggregate economic outcomes is increasingly negative for warmer climates.

The inclusion of a main effect λ_h associated with the time-varying state variable allows for the possibility that expected impacts may depend on changes in expected temperatures, net of the impact of climate change on the distribution of temperature shocks. For example, if a hot country stabilizes at a climate that is 1°C warmer than some reference period, the productive potential of the country is altered even in the absence of temperature shocks because the new climate will be more or less conducive to different economic activities. Neglecting this likelihood amounts to assuming perfect climate adaptation in the long run and that climate change is only socially impactful through its positively biasing the distribution of annual temperature shocks.

It is worthwhile to take stock of the assumptions underlying the sequence of regressions specified in Equation (3). First, because coefficients for each horizon are estimated through distinct regressions using differently constructed variables from different subsamples of the data, there are in principle no restrictions imposed by the model on the shape of the resulting impulse response function for a given value of the state variable; it is in this sense that the method is partially non-parametric.

This flexibility is particularly appealing in the context of estimating responses to weather shocks. As summarized in Section 2.4, fully parametric methods require the modeler to assume either of the extreme persistence patterns distinguished in Figure 2 with limited ability to accommodate intermediate persistence dynamics. By instead allowing the modeler to remain mostly agnostic to the data generating process, this method circumvents a notorious source of extreme estimation instability and can accommodate any persistence structure in principle without restriction[7]. Graphically, the sequence of coefficient estimates $\{\beta_{1h}\}_{h=0}^{H}$ in Equation (3) maps directly to the differences between the black and red time series in the first "levels" row; equivalently the sequence defined as the first difference in this sequence of coefficients correspond to the difference between the red and black time series depicted in the second "growth" row. Thus, a one-period level effect would be represented by a sequence where $\beta_{1h} = 0$ for all horizons $h > 0$ while a

constant growth effect would imply a sequence converging to a non-zero constant[8].

The single-equation local projections model is asymptotically equivalent regardless of whether the outcome is specified as h-period long differences as in Equation (3) or as h-period projections in levels. However, recent work has shown that the former has particularly appealing finite-sample properties which make it especially robust to misspecification, largely eliminating well-known biases associated with highly persistent and even non-stationary processes (Piger and Stockwell 2025; Jordà and Taylor 2025). This is particularly valuable in a climate impacts literature where GDP regressions are commonplace despite their vulnerability to unit root concerns (Nelson and Plosser 1982; Campbell and Mankiw 1987).

Still, the model is not entirely free of restrictions. Most critically, the model specified in Equation (3) imposes within-horizon linearity assumptions on the coefficients of interest. Since the shock enters additively into Equation (3), responses to the shock are assumed to scale proportionally so that a 1°C shock has exactly twice the impact of a 0.5° shock and the exact negative of the effect of a -1°C shock. By the same principle, the model also imposes linearity on the interaction effect, imposing symmetry in the state dependence such that holding all other factors fixed, the expected impact of a shock is assumed to be $k\beta_{2h}$ percentage points more positive in a climate some quantity k degrees warmer than baseline and $-k\beta_{2h}$ more negative in a climate k degrees cooler than baseline. While appreciating that unpredictable non-linearities abound in the interactions between social and climate systems (Dietz et al. 2021), I consider these simplifications much more preferable to generalized alternative models.

The gains to unbiasedness from minimally parametric models naturally come with costs in estimation efficiency; IRFs estimated by local projection regressions are fundamentally noisier than their parametric and structural counterparts such as vector autoregressions[9]. For this reason, we expect resulting estimates to be less precise than more restrictive methods, particularly for increasingly distant projection horizons since the available sample size in a fully balanced panel reduces linearly in h.

Collecting estimates of the coefficients of interest $\{\widehat{\beta}_{1h}, \widehat{\beta}_{2h}\}_{h=0}^{H}$ from Equation (3) and adopting the notation of Jordà (2023), our set of local projection estimates imply a family $\mathcal{R}_{\tau \to y}(h; \overline{T})$ of state-dependent impulse responses as a function of the state variable $\overline{T}$:

$$\mathcal{R}_{\tau \to y}(h; \overline{T}) := \widehat{\beta}_{1h} + \widehat{\beta}_{2h} \cdot \overline{T} \tag{4}$$

3.1.2 Estimating state-dependent cumulative multipliers

To make progress accounting for the persistence of temperature shocks themselves (see: Section 2.3), Figure 3 depicts impulse responses which result from using temperature shocks themselves as an outcome variable y as constructed by two commonly used datasets for gridded historically observed weather. We plot separate impulse responses for three representative states corresponding to the 20th and 80th percentiles of country-year expected temperatures in our sample

[8]Residuals across the projection horizon are generally autocorrelated but do not affect consistency of point estimates.

[8]Of course, under a conservative null hypothesis of non-persistence at all horizons, a permanent growth effect cannot be affirmatively evidenced with finite data and the modeler can only produce evidence of a lower bound on this persistence.

[9]Different compromises in this bias-variance tradeoff can be obtained by imposing additional restrictive parametric assumptions such as the functional approximation approach of Barnichon and Matthes (2018) which can substantially diminish estimation error.

(10°C colored blue and 26°C colored red) and their average (18°C, roughly the 40th percentile and colored green). This coloring scheme and these reference temperatures will be used throughout the rest of this paper. [10]

Figure 3: State-dependent response of temperature shocks to a 1°C temperature shock

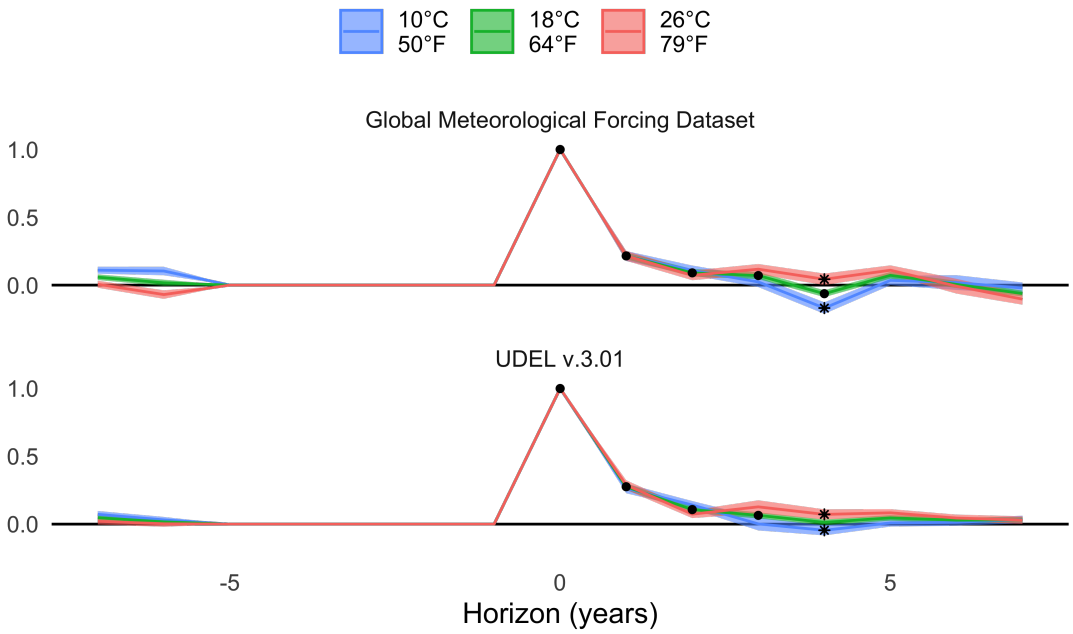

State-dependent impulse responses of identified temperature shocks estimated using two commonly used temperature datasets. Dots in black denote significant evidence of autocorrelation and asterisks denote significant state dependencies in this autocorrelation. Fo both products, persistence is large in magnitude for the first 1-2 periods. Negative horizons are included as a check for significant pre-shock trends.

Points along the IRFs depicted in black correspond to point estimates which are deemed significantly persistent by a joint hypothesis test with $\alpha = 0.05$. The underlying test is based on the significance-band test recently introduced in Inoue et al. (2025) and adapted for our state-dependent setting as described in Section A2.2. In short, we have adopted a convention where filled circles correspond to significant main effects and asterisks correspond to significant interaction effects under a joint null hypothesis of non-persistence at each horizon. The presence of a significant effect up to horizon h is interpreted as evidence of a lower bound of h-period persistence. Thus, the IRFs presented in Figure 3 constitute significant evidence that temperature shocks positively influence successive temperature shocks for at least four periods although the magnitude of this persistence is mostly diminished within 1-2 years.

[10] Note that for expositional convenience, we have transformed our state variable to center on the intermediate reference temperature of 18°C so that the responses in green correspond to the sequence of estimated main effects $\widehat{\beta}_{1h}$ while the responses in blue and red incorporate equivalent but opposite interaction effects corresponding to ± 8°C deviations from this reference state

As emphasized by Nath et al., we will want to directly account for this persistence when attributing social impacts to a unit temperature shock. An analogy can be drawn to the estimation of multipliers in empirical macroeconomics: when estimating the impulse response of, for example, GDP to a fiscal policy shock, one must account for the tendency of stimulus programs to be implemented in stages or to be followed by additional, if typically smaller, rounds of stimulus. Otherwise, the effect sizes implied by simple impulse response functions misattribute cumulative changes in the outcome only to the initial unit shock; the calculation of multipliers scale these effects at each horizon by the corresponding accumulation of shocks. As such, we will want to present our results in the form of these cumulative temperature multipliers rather than as raw impulse response functions.

In **??**, I describe how I modify methods for efficiently estimating cumulative temperature multipliers in order to implement this transformation while retaining the ability to statistically test for persistence. Ultimately, the process of estimating cumulative temperature multipliers for some outcome y amounts to estimating a cumulative variant of the original single-equation model in Equation (3). That is, for every discrete income quantile q of interest and projection horizon $h \in \{0, 1, ..., H\}$, we estimate the following state-dependent model:

$$\Delta_h^c y_{i,t+h}^q = m_{1,h}^q \widehat{\tau}_{i,t+h}^c + m_{2,h}^q \widehat{\tau}_{i,t+h}^c \overline{T}_{it} + \lambda_h \overline{T}_{it} + \mathbf{Z}_{it}^q \boldsymbol{\gamma}_h^q + \mu_i + \eta_t + \varepsilon_{i,t+h}^q \tag{5}$$

Here, variables denoted by a superscript c are cumulative versions of the original variables constructed using the one-step process summarized by Algorithm A1 in **??**. The quantile-specific matrix of controls $\mathbf{Z}_{it}^q$ include lags $\ell \in \{1, ..., L\}$ of the shocks, the interaction of the shock with the state variable, the outcome $\Delta y_{i,t-\ell}^q$, and country-level aggregate growth $\Delta y_{i,t-\ell}$.

State-dependent cumulative temperature multiplier estimates $\{\widehat{m}_{1,h} + \widehat{m}_{2,h} \overline{T}_{it}\}_{h=0}^H$ describe the effect of an identified temperature shock on the cumulative stock of outcome y^q after h periods as a function of the expected temperature.

3.2 Empirical results: evidence for a temperature effect on within-country inequality

Panel (a) of Figure 4 summarizes the results of estimating separate sets of state-dependent cumulative temperature multipliers for all 10 within-country income deciles. Here, color-coding follows the same intuitive scheme as previous plots with blue and red corresponding to representative expected temperatures of 10 and 26°C respectively[11]. The points in asterisks indicate significant interactions which impose differentiated effects based on expected climate. These effects persist for between 1-3 years, concentrating negative impacts in hot countries while keeping cooler countries relatively insulated, consistent with the aggregate results standard in the climate literature and the dynamic GDP results reported in **??**. Estimation error is larger than when estimating GDP impacts in Figure A2, unsurprising given the difference in temporal coverage and data quality.

Importantly, we are now able to compare impacts across quantiles within countries. Focusing on hot-climate countries, the estimates reported in Panel (a) imply an incidence pattern wherein negative effects of a transitory temperature shock to hot countries are found at all quantiles but are particularly dire at the bottom. For 26°C countries, the 90% confidence intervals for the impact of a temperature shock on cumulative income a year after the shock is [-10.0,-2.9] percent for the

[11] We omit the green 18°C main effect to avoid over-busy plots, but point estimates are implied as the average of the two depicted curves due to the linearity of the interaction effect.

Figure 4: Temperature multipliers of national income by quantile

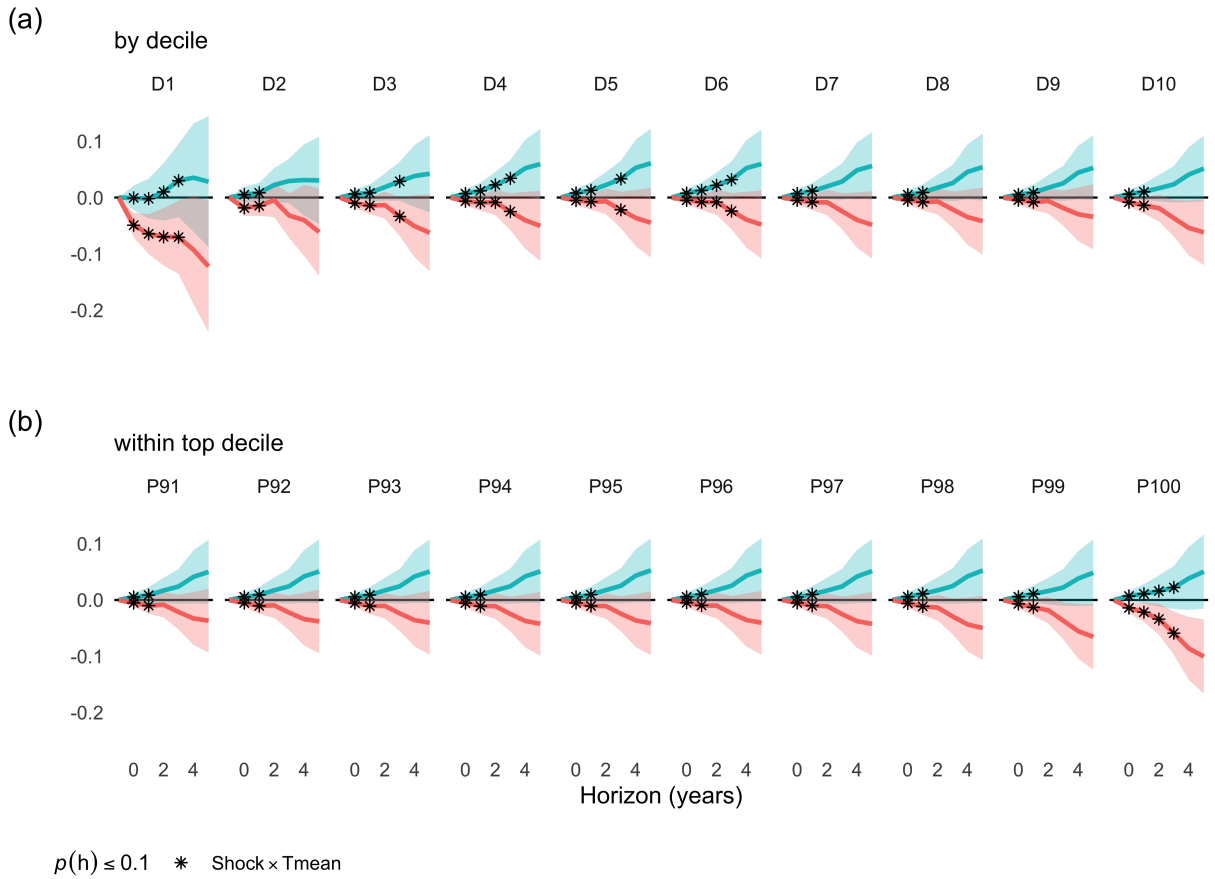

Estimates of the cumulative temperature multiplier for different quantiles of national income for two representative 'states' each: a 10°C country in blue and a 26° country in red. Asterisks denote statistically significant state-dependencies. The top row (a) depicts results for each decile of the national income distribution while the bottom row (b) depicts results decomposing for each percentile within the top decile (i.e., D10 in the top row).

bottom decile compared to [-3.2, 0.3] percent for the second decile. This generally improves as one moves up the income distribution, stabilizing at approximately [-2.4,0.6] percent by the fourth decile through to the 99th percentile. Interestingly, the estimated impacts then discontinuously drops to [-3.9, -0.5] for the richest percentile. Panel (b) shows this discontinuity within the top decile and also indicates this effect is significantly persistent for two additional years. In contrast, cumulative growth incidence is much flatter across the income distribution with point estimates suggesting positive impacts are somewhat larger for the middle classes between the third and sixth deciles.

Figure 5 and the bottom panel of Figure 5 capture these discontinuities at the extremes more clearly by plotting point estimates of percentile-level impacts over time for our three representative climates. These figures imply that negative expected impacts to hot countries are indeed concentrated on both the bottom 10% and the top 1% of the distribution[12]. To the best of my knowledge, this particular vulnerability of the highest incomes in hot and generally poor countries to transitory environmental shocks has not been previously observed or theorized although it is consistent with recent work outside environmental economics demonstrating this group's disproportionate sensitivity to aggregate conditions even compared to other members of the top income decile (Roine et al. 2009; Alvaredo et al. 2018). In ongoing work, I investigate the extent to which this may be driven by weather-induced capital depreciation, capital shares in agricultural income, and the relative unavailability of weather derivatives in these settings.

This first set of results describe inequalities in the growth incidence of impacts across a national income distribution as a function of the local climate. The disproportionate impacts to the bottom decile strongly imply a regressive effect: In the aftermath of an isolated temperature shock, a representative member of the the bottom decile in a 26°C country can expect to earn 7% less in total over the next three years than they would had the shock never occurred. This compares to losses of 2.2% for the median adult in their country and a *gain* of 3.0% for their counterpart in a country with a 10°C climate. At the same time, a representative member of the 100th percentile in the same 26°C country, represented by the thinnest sliver at the top of the rightmost distributional contour plot in (b), can expect to earn 5.9% less over those three years than they would otherwise compared to 3.6% for the 99th percentile.

[12] As mentioned in Section 2, the income data for the bottom 10% is not well suited for decomposition into percentile effects but one would expect an ideal decomposition would reveal a similarly steep within-decile gradient in which negative impacts are driven by the poorest percentiles within the bottom 10%.

Figure 5: Temperature multipliers of national income by quantile, point estimates

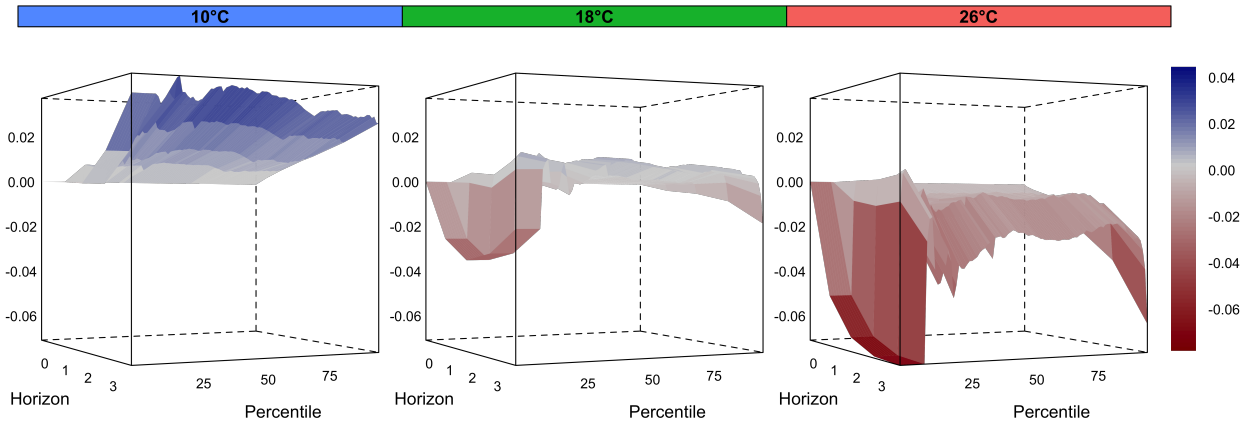

(a) Temperature multipliers of income. 3D surfaces represent estimates of the temperature multiplier of income for the bottom decile and percentiles 11-100 of a national income distribution. Curves in the multiplier-horizon plane represent dynamics for different income quantiles while curves in the multiplier-percentile plane are akin to h-period growth incidence curves attributable resulting from a pulse of temperature. Columns correspond to different representative climates.

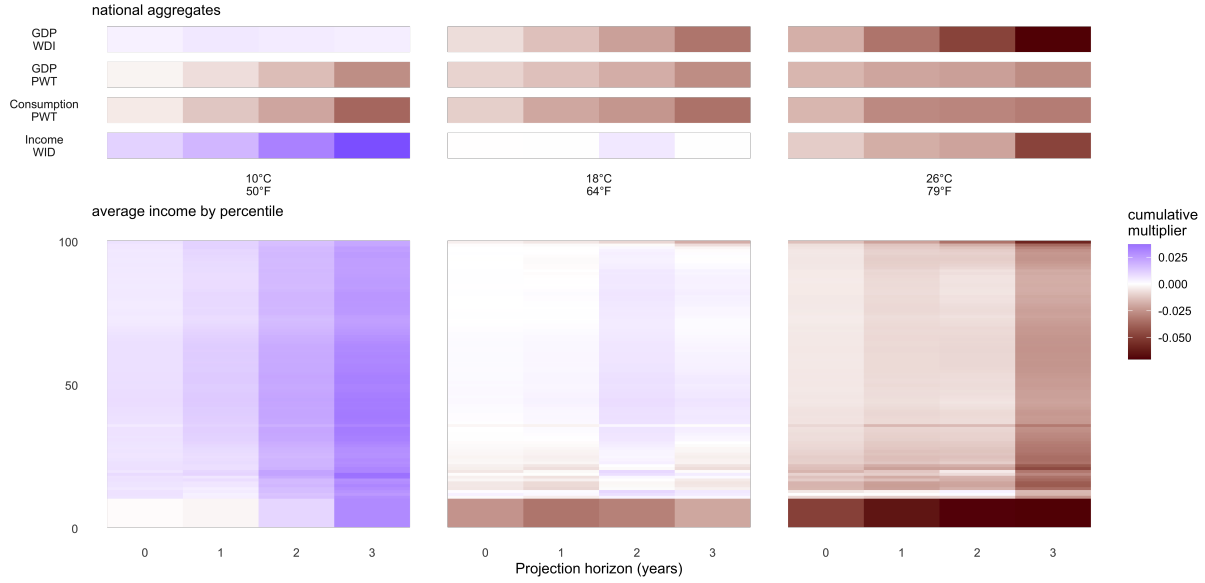

(b) Aggregate vs. distributional multipliers. Planes at the bottom flatten the surfaces in panel (a) by using color gradients to represent positive (blue) and negative (brown) multipliers. Strips at the top are the equivalent multipliers estimated for country-level economic aggregates.

4 Counterfactual analysis

4.1 From identified temperature shocks to systematic climate change

The results presented in Section 3.2 are expressed as responses to a 1°C shock, which notably is much larger in magnitude than shocks we would expect to observe in a typical year in any country. Referring back to the time series of shocks illustrated in Figure 1, the median magnitude of a shock varies by the rate of climate change which will in turn vary across geographies and time. For the warmest 33% of country-years in our data, the median shock was approximately 0.19°C in absolute value in the 1970s and then increased to between 0.30-0.37° in the ensuing decades. Thus, hot-country effects can be understood as responses to a shock 3-5 times larger than is typical. For the coldest 33% of country-years, the median magnitude of a shock was approximately 0.35°C in the 1970s and 1980s, 0.51-0.53°C in the 1990s and 2000s, and 0.59°C in the 2010s. Their effect sizes then can be interpreted as the response to a shock just 2-3 times larger in magnitude than the median.

So while those results consistently demonstrated much larger distributional responses to the same shock for residents in warmer geographies, this disparity is at least partially offset by the relatively mild incidence of shocks in these settings. The counterfactual analysis we document in this section is valuable as a means of interpreting these reduced-form results practically, accounting for the spatial incidence of temperature shocks as well as other factors such as the absolute distribution of global incomes across country-quantiles and the spatial distribution of the global population across different geographies as observed in the historical record.

In the absence of a structural climate-economy model, it is standard practice to extrapolate reduced-form effects estimated using short-term weather variation to infer long-term impacts of sustained climate change. Broadly, this involves applying the estimated impact model to a counterfactual climate scenario to produce a counterfactual trajectory for the social outcome of interest. The difference between the social outcome of interest under this counterfactual relative to its trajectory under a reference scenario (e.g., for historical attribution studies like the present analysis, the actual historically observed trajectory of this outcome) are then attributed to the differences in the climate under the two scenarios.

In our case, we are estimating the effect of changes to the global climate system between 1980 and 2016 on the global distribution of income. Using the observed history of country-year temperatures, we calculate the implied history of country-year termperature shocks as defined in Equations (1a) and (1b). For the bottom national income decile and each of the 11th to 100th national income percentiles, we collect $B = 500$ bootstrap estimates of the state-dependent cumulative temperature multipliers of income according to Equation (5). Integrating these percentile-bootstrap specific estimates to the observed histories of temperature shocks implies a counterfactual history of country-percentile-year income absent any temperature shocks between 1980 and 2016. The differences between this simulated history of global income distributions and the observed history of global income distributions are then attributed to changes in the climate over this 37-year period.

This process (and its analogous implementation for projecting future climate change impacts) offers several significant advantages over prior approaches to counterfactual attribution. First, since estimates of cumulative damages result from local projection regressions, dynamic impacts are identified directly from medium-term long differences in income and disciplined by a statistical test of their persistence rather than imposed by modeling choices as described in Section 2.4.

Second, models which use levels of temperature as their treatment variable have been challenged on the grounds that these counterfactual exercises require them to extrapolate effects of

treatments increasingly beyond the range of observed data. This is much less of a concern in our setting where the treatment variable is an identified temperature shock which is dynamically anchored by a time-varying state variable. As a result, the treatment variable remains largely within the range of the historical record even under rapid warming because the state variable gradually adjusts to recent experience. In this sense, our model accommodates a notion of adaptation in the sense of Hsiang (2016) wherein the "direct" effects of exposure to a temperature shock can be partially offset by Bayesian "belief effects"; the same temperature level is less of a 'shock' the more experience one has with similar temperatures. In our case, a one-time discrete shift from one climate to a new climate would be fully internalized after 30 years.

More rapid adaptation can be accommodated by lowering the value of M defining the state variable. Alternatively, Bilal and Känzig (2024) and Nath et al. (2024) define temperature shocks as the residuals from autoregressive distributed lag models (ARDLs) of temperature levels. In so doing, they impose a model of much more immediate and complete adaptation so that if a 26°C country were to be exposed to, for example, 40°C weather five years in a row, the fifth year will impose no economic cost because it would register as a 0°C shock. This difference in assumed adaptability is why their time series of historically observed temperature shocks is largely homoskedastic around a constant mean of approximately zero while ours is positively trended in periods of warming (as seen in Figure 1).

Our more conservative construction offers a number of attractive advantages. First, it is arguably more consistent with the empirical literature on climate adaptation; a recent review of this work finds "limited evidence" of observable adaptation moderating the direct impacts of weather shocks in the context of climate change (Burke and Emerick 2016). Secondly, it lends itself to a simple, intuitive, and consistent approach to estimating damages from climate change: the empirical analysis isolates an identified temperature shock to estimate its social impacts and then the counterfactual analysis represents climate change as a positive trended bias in the distribution of these shocks and uses the empirical result to calculate damages.

This intuition is not available to approaches which define temperature shocks as residuals from an ARDL. Because the time series of temperature shocks reverts to a mean of zero within the periodicity of its maximal lag, the distribution of temperature shocks are nearly homoskedastic and symmetrical about 0 even amid climate change. As a result, the same intuitive approach to counterfactual analysis would imply near-zero cumulative social costs from climate change.[13]

4.2 Results

In constructing these counterfactual series, each bootstrap b assigns each country-quantile-year to a global percentile corresponding to their population-weighted income ranking across the global income distribution for that year. Finally, we calculate the long-term growth between 1981 and 2016 implied by each series for each bootstrap series. The resulting long-term growth incidence curve is depicted in Figure 6.

These results may be interpreted as impacts of observed climate change insofar as it the influence of climate change is sufficiently well-captured by asymmetries in the incidence of temperature shocks and changes to expected temperatures over the 36-year period of analysis. **??** presents

[13]To arrive at non-zero climate damages, Bilal and Känzig (2024) and Nath et al. (2024) instead seem to use the conventional "delta method" appropriate for parametric models which use temperature levels as their treatment. This entails first linearizing the difference in temperature levels before and after the relevant period of warming; for example, under the most extreme warming scenario between 2010 and 2100, the population-weighted mean temperature increase over this period is 4.3°C. These differences are then seemingly interpreted as a representative long-term 'shock' to be multiplied by the coefficient on the original temperature shock.

additional results decomposing these effects into their anthropogenic and natural contributions by applying methods from climate attribution science to output from state-of-the-art climate models, arriving at largely the same conclusions.

Here, the historically observed level of growth over this period is represented by a black filled line, reproducing a version of the "elephant curve" famously documented in Lakner and Milanovic (2016). The blue band depict the 90% confidence interval for the counterfactual removing the influence of temperature shocks on income. Regions where this band is above the observed incidence curve imply deprivations of income attributable to the incidence of temperature shocks over this period. Point estimates imply probable net harm for around 65% of the global sample.

Despite proportionally milder absolute warming in the poorest countries (see Figure A5), this negative incidence still concentrates on the world's poorest 20%, representing 1.2 billion people in our 2016 sample and mostly located in warm developing countries. Importantly, as in Alvaredo et al. (2018), calculations underpinning this figure omit the bottom 10% of each country's income distribution because of the inclusion of unemployed adults and compositional instability effects which arise from assigning deciles of very large populations to percentiles representing smaller populations. For the poorest percentile in our data, we find that incomes absent these disproportionate temperature shocks would be 29% [18, 41] higher than they are today. As these effects exclude precisely the subpopulation we have already found to be the most economically vulnerable to temperature shocks, this growth incidence curve should be interpreted as a likely very conservative estimate of deprivations to the global poor.

Elsewhere, estimates also imply likely harm from warming for the 52nd-97th percentiles. That the top 3% are relatively unaffected despite the negative effects for the top 1% in Section 3 is because these groups are dominated by extreme incomes in relatively cold countries. Effects on the global middle class are relatively tempered with slightly positive point estimates estimated for essentially the entire second quartile, also consistent with the within-country incidence curves which found the middle class experienced the mildest economic impacts across all climates.

Despite these different-signed effects, the counterfactual growth incidence band is noticeably flatter than the observed elephant curve, implying that historical climate change has been globally regressive in addition to acting as a drag on global economic growth. Figure 6 plots the time series for the global equivalents of the within-country inequality measures introduced earlier. In our within-country analysis, we identified regressive transfers from the bottom decile and progressive transfers from the top percentile. The historical effects represented in these graphs summarize how these distributional patterns aggregate globally according to index-specific social welfare functions[14]. All three measures are increasing in inequality and for each, redistributive effects are found to be large and regressive: proportionally, the global Theil index increases by 0.9% [-1.8, 2.8], the Gini index by 0.5 points [-0.2, 1.0], and the mean logarithmic deviation index by 3.9% [1.6, 6.4].

---

[14]One interpretation of the Theil measure is as a geometric measure of inequality. For any given income distribution, an individual with $100,000 giving $10 to someone with $10,000 would improve the Theil inequality index by the exact same amount as someone with $1,000 giving the same $10 to someone with $100 because the ratio of incomes between the donor and recipient are the same in both cases (Cowell 2011).

Conversely, a $100 income earner regressively transferring 10% of their income to a $1,000 income earner would offset the progressivity of a $100,000 income-earner transferring a ten-thousandth of their income to a $10,000 income earner. If that is still not intuitive, then you are in good company (Sen 1973).

[14]One interpretation of the mean logarithmic deviation measure is as the average disutility of a given income distribution compared to an equal distribution of the same total income, assuming logarithmic individual utility. Of the three measures, the Theil places the most welfare weight on redistribution among the top quantiles while the MLD places

Figure 6: Global income inequality (1980-2016), observed vs. counterfactual

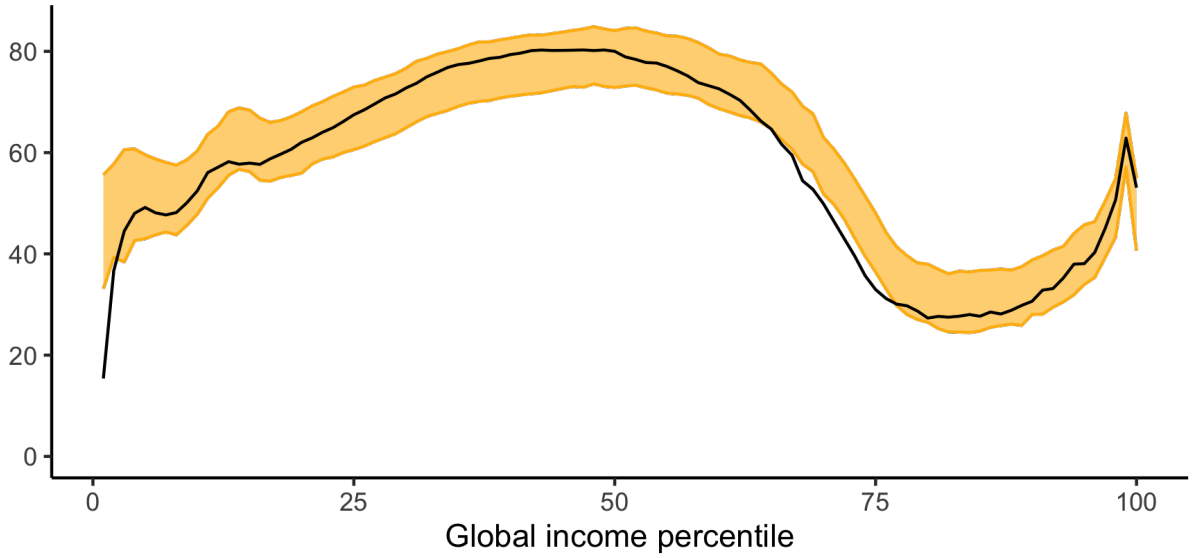

(a) Change in income levels (%), 1980 vs. 2016. Black line depicts observed increases in income levels by global income percentile among positive income-earners. Bands in gold depict the 95% confidence interval of the counterfactual increase had warming stayed fixed at 1980 levels.

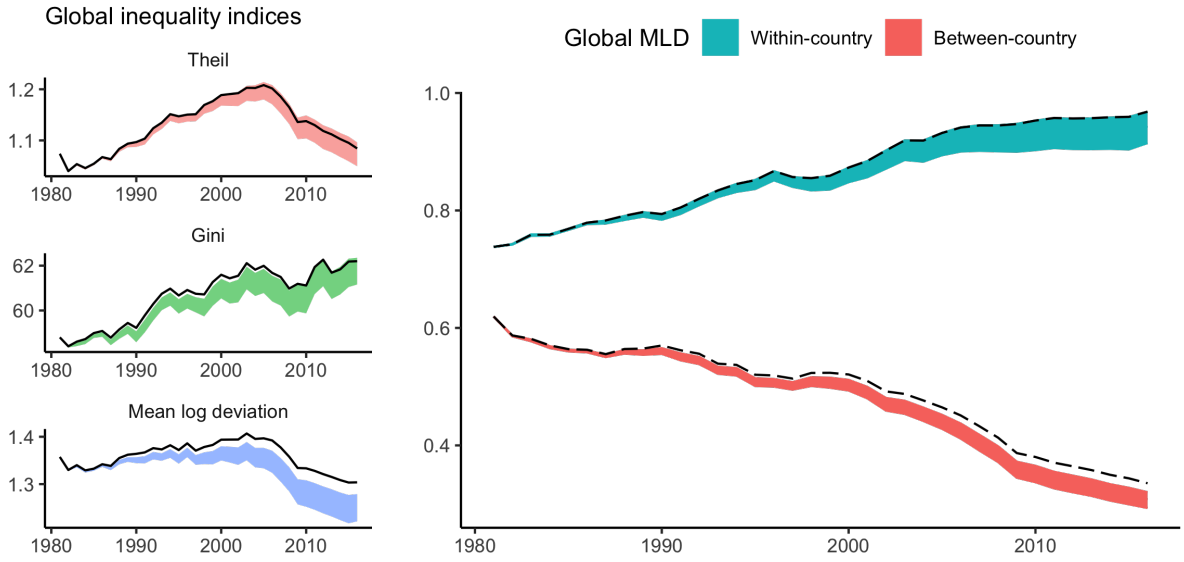

(b) Global inequality, total and decomposed. On left, black lines depict the observed evolution of global inequality from 1980 to 2016 by three common inequality indices. Colored bands represent the 95% confidence interval for their evolution under the same counterfactual. On right, we decompose the mean log deviation series into its between-country and within-country contributions.

The right panel in Figure 6 decomposes the effect on the global MLD index into its within-country and between-country dimensions. We find that climate change has unambiguously exacerbated both varieties of inequality, increasing between-country inequality by 8.7% [4.9, 13.3] and within-country inequality by 2.6% [0.0, 5.6]. To the best of my knowledge, this represents the first empirical estimate of the latter quantity in the vast empirical literature on climate impacts.

5 Conclusion

This paper represents an important advance in our understanding of the social welfare implications of climate change. Taking advantage of newly available income data, I first present new evidence that temperature shocks impact economic inequality between *people*, building on existing distributional analysis which has largely been restricted to coarse comparisons between countries. I show that temperature shocks have historically widened income inequalities within countries by quantifying the disproportionate and persistent burden faced by the lowest income quantiles, particularly in hotter countries.

A counterfactual analysis suggests that systematic global warming between 1980 and 2016 substantially increased global inequality, an effect equally driven by its thoroughly studied effect widening between-country disparities just as much as by this newly documented effect exacerbating within-country inequality. For the lowest percentile of income earners globally, these effects are most pronounced, amounting to cumulative income losses on the order of 29% and concentrated in populous developing countries.

Incorporating these distributional implications is essential not just for understanding climate change as a social phenomenon, but also for informing the efficient and equitable design of climate policy. These results challenge the adequacy of prevailing utilitarian cost-benefit approaches to even non-environmental impact evaluations which have been criticized for embedding a "profound indifference to inequality" (Sen et al. 2020). In addition, distributional impacts are likely key to identifying and overcoming barriers to the political viability of sufficiently ambitious climate policy. As Chancel, Piketty, et al. (2022) write, "Policy reforms which do not properly factor in the degree of inequality in a country, and the winners and losers of these reforms, are unlikely to be publicly supported and are likely to fail."

Yet, the limitations of the present analysis invite further research across the interdisciplinary climate community emphasizing these disaggregated distributional impacts. For one, even perfectly accurate income data is known to be a flawed measure of welfare among the impoverished. This is partially reflected in our data where bottom-coding groups all unemployed peoples together, regardless of class and social security, requiring my analysis to consider the bottom deciles within these countries together. There is likely to be incredible material inequality within this group that income data is poorly suited to distinguish. Related considerations have made much less widely available consumption expenditures the preferred measure of welfare for the global poor in development economics and innovations to overcome this data gap are necessary to better capture the poverty implications of climate change (Jean et al. 2016).

the most welfare weight on redistribution among the bottom quantiles.

REFERENCES

Alvaredo, Facundo et al. (May 2018). "The Elephant Curve of Global Inequality and Growth". *American Economic Review: Papers and Proceedings* 108, pp. 103–108.

Atkinson, Anthony B. (June 2001). "The Strange Disappearance of Welfare Economics". *Kyklos* 54.(2), pp. 193–206.

Atkinson, Anthony B. and Andrea Brandolini (Sept. 2001). "Promise and Pitfalls in the Use of "Secondary" Data-Sets: Income Inequality in OECD Countries as a Case Study". *Journal of Economic Literature* 39.(3).

Barnichon, Regis and Christian Matthes (Nov. 2018). "Functional Approximation of Impulse Responses". *Journal of Monetary Economics* 99, pp. 41–55.

Berg, Kimberly A., Chadwick C. Curtis, and Nelson C. Mark (Oct. 2024). "GDP and temperature: Evidence on cross-country response heterogeneity". *European Economic Review* 169.

Bergquist, Parrish, Matto Mildenberger, and Leah C Stokes (May 2020). "Combining climate, economic, and social policy builds public support for climate action in the US". *Environmental Research Letters* 15.(5).

Bilal, Adrien and Diego R. Känzig (Nov. 2024). *The Macroeconomic Impact of Climate Change: Global vs. Local Temperature*. NBER Working Paper 32450. National Bureau of Economic Research.

Burke, Marshall and Kyle Emerick (2016). "Adaptation to Climate Change: Evidence from US Agriculture". *American Economic Journal: Economic Policy* 8.(3).

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

Campbell, John Y and N Gregory Mankiw (Nov. 1987). "Are output fluctuations transitory?" *Quarterly Journal of Economics* 102.(4), pp. 857–880.

Carleton, Tamma and Solomon M. Hsiang (Sept. 2016). "Social and economic impacts of climate". *Science* 353.(6304).

Carleton, Tamma, Amir Jina, et al. (Nov. 2022). "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". *Quarterly Journal of Economics* 137.(4), pp. 2037–2105.

Casey, Gregory, Stephie Fried, and Ethan Goode (June 2023). "Projecting the Impact of Rising Temperatures: The Role of Macroeconomic Dynamics". *IMF Economic Review* 71, pp. 688–718.

Cevik, Serhan and João Tovar Jalles (Mar. 2023). "For whom the bell tolls: Climate change and income inequality". *Energy Policy* 127.

Chancel, Lucas, Philippe Bothe, and Tancrède Voiturez (Jan. 2023). "Climate Inequality Report 2023: Fair taxes for a sustainable future in the Global South". *Environmental Research Letters* 18.

Chancel, Lucas and Thomas Piketty (Nov. 2015). "Carbon and inequality: from Kyoto to Paris—Trends in the global inequality of carbon emissions". *World Inequality Lab*.

Chancel, Lucas, Thomas Piketty, et al. (2022). *World Inequality Report*. Paris: World Inequality Lab.

Cloyne, James, Òscar Jordà, and Alan M. Taylor (Feb. 2023). *State-Dependent Local Projections: Understanding Impulse Response Heterogeneity*. NBER Working Paper 30971. National Bureau of Economic Research.

Cowell, Frank (Mar. 2011). *Measuring Inequality*. Oxford University Press.

Dasgupta, Shouro, Johannes Emmerling, and Soheil Shayegh (Nov. 2023). "Inequality and growth impacts of climate change—insights from South Africa". *Environmental Research Letters* 18.

Dechezleprêtre, Antoine et al. (Apr. 2025). "Fighting Climate Change: International Attitudes toward Climate Policies". *American Economic Review* 115.(4), pp. 1258–1300.

Deininger, Klaus and Lyn Squire (Sept. 1996). "A New Data Set Measuring Income Inequality". *World Bank Economic Review* 10.(3).

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (July 2012). "Temperature Shocks and Economic Growth: Evidence from the Last Half Century". *American Economic Journal: Macroeconomics* 4.(3), pp. 66–95.

Dietz, Simon et al. (Aug. 2021). "Economic impacts of tipping points in the climate system". *Proceedings of the National Academy of Sciences* 118.(34).

Diffenbaugh, Noah S. and Marshall Burke (Apr. 2019). "Global warming has increased global economic inequality". *Proceedings of the National Academy of Sciences* 116.(20).

Dube, Arindrajit et al. (May 2025). *A Local Projections Approach to Difference-in-Differences Event Studies*. NBER Working Paper 31184. National Bureau of Economic Research.

Furceri, Davide, Michael Ganslmeier, and Jonathan Ostry (July 2023). "Are climate change policies politically costly?" *Energy Policy* 178.

Gilli, Martino et al. (Sept. 2024). "Climate change impacts on the within-country income distributions". *Journal of Environmental Economics and Management* 127.

Graff Zivin, Joshua and Matthew Neidell (2014). "Temperature and the Allocation of Time: Implications for Climate Change". *Journal of Labor Economics* 32.(1).

Howard, Peter and Derek Sylvan (Mar. 2021). "Gauging Economic Consensus on Climate Change". *Institute for Policy Integrity, New York University School of Law*.

Hsiang, Solomon M. (Oct. 2016). "Climate Econometrics". *Annual Review of Resource Economics* 8, pp. 43–75.

Hsiang, Solomon M. et al. (June 2017). "Estimating economic damage from climate change in the United States". *Science* 356.(6345), pp. 1362–1369.

Inoue, Atsushi, Òscar Jordà, and Guido M. Kuersteiner (Feb. 2025). "Inference for local projections". *Econometrics Journal*.

Jean, Neal et al. (Aug. 2016). "Combining satellite imagery and machine learning to predict poverty". *Science* 353.(6301), pp. 790–794.

Jenkins, Stephen P. (Aug. 2015). "World income inequality databases: an assessment of WIID and SWIID". *Journal of Economic Inequality* 13, pp. 629–671.

Jones, Charles I. (Nov. 2015). "Pareto and Piketty: The Macroeconomics of Top Income and Wealth Inequality". *Journal of Economic Perspectives* 29.(1), pp. 29–46.

Jordà, Òscar (Mar. 2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review* 95.(1), pp. 161–182.

— (Sept. 2023). "Local projections for applied economics". *Annual Review of Economics* 15, pp. 607–663.

Jordà, Òscar and Alan M. Taylor (Mar. 2025). "Local Projections". *Journal of Economic Literature* 63.(1), pp. 59–110.

Jorgenson, Dale W. (Sept. 2018). "Production and Welfare: Progress in Economic Measurement". *Journal of Economic Literature* 56, pp. 867–919.

Kahn, Matthew E. et al. (Dec. 2021). "Long-term macroeconomic effects of climate change: A cross-country analysis". *Energy Economics* 104.

Kotz, Maximilian, Anders Levermann, and Leonie Wenz (Apr. 2024). "The economic commitment of climate change". *Nature* 628, pp. 551–557.

Lakner, Christoph and Branko Milanovic (Aug. 2016). "Global income distribution: From the fall of the Berlin Wall to the Great Recession". *World Bank Economic Review* 30.(2), pp. 203–232.

Maestre-Andrés, Sara, Stefan Drews, and Jeroen van den Bergh (June 2019). "Perceived fairness and public acceptability of carbon pricing: a review of the literature". *Climate Policy* 19.(9), pp. 1186–1204.

Management, Office of and Budget (Nov. 2023). *Circular No. A-4: Regulatory Analysis*.

Marx, Nicolas Longuet (Apr. 2024). *Investigating the Impact of Temperature Shocks on Income Disparity*. CEEP Working Paper 32.

Masson-Delmotte, V. et al., eds. (2018). *Global Warming of 1.5řC*. IPCC Special Report SR15. Geneva, Switzerland: Intergovernmental Panel on Climate Change (IPCC).

Meyer, Bruce D. and James X. Sullivan (Nov. 2003). "Measuring the well-being of the poor using income and consumption". *Journal of Human Resources* 38, pp. 1180–1220.

Montiel Olea, José Luis and Mikkel Plagborg-Møller (July 2021). "Local projection inference is simpler and more robust than you think". *Econometrica* 89.(4), pp. 1789–1823.

Montiel Olea, José Luis, Mikkel Plagborg-Møller, et al. (Aug. 2024). *Double Robustness of Local Projections and Some Unpleasant VARithmetic*. Version 2. arXiv preprint. eprint: `arXiv:2405.09509`.

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (July 2024). *How Much Will Global Warming Cool Global Growth?* NBER Working Paper 32761. National Bureau of Economic Research.

Nelson, Charles R and Charles I Plosser (1982). "Trends and random walks in macroeconomic time series: Some evidence and implications". *Journal of Monetary Economics* 10.(2), pp. 139–162.

Newell, Richard G., Brian C. Prest, and Steven E. Sexton (2021). "The GDP-Temperature relationship: Implications for climate change damages". *Journal of Environmental Economics and Management*.

Palagi, Elisa et al. (Oct. 2022). "Climate change and the nonlinear impact of precipitation anomalies on income inequality". *Proceedings of the National Academy of Sciences* 119.

Pande, Rohini (Aug. 2023). "The climate crisis is a crisis of inequality". *Science* 381.(6661).

Park, R. Jisung et al. (May 2020). "Heat and Learning". *American Economic Journal: Economic Policy* 12.

Piger, Jeremy and Thomas Stockwell (May 2025). *Differences from Differencing: Should Local Projections with Observed Shocks be Estimated in Levels or Differences?* Tech. rep. SSRN Working Paper.

Piketty, Thomas (Oct. 2003). "Income Inequality in France, 1901–1998". *Journal of Political Economy* 111.(5), pp. 1004–1042.

Piketty, Thomas and Emmanuel Saez (Feb. 2003). "Income Inequality in the United States, 1913–1998". *Quarterly Journal of Economics* 118.(1).

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (May 2018). "Distributional National Accounts: Methods and Estimates for the United States". *Quarterly Journal of Economics* 133.(2), pp. 553–609.

— (May 2019). "Simplified Distributional National Accounts". *AEA Papers and Proceedings* 109, pp. 289–295.

Prest, Brian C. et al. (Aug. 2024). "Equity weighting increases the social cost of carbon". *Science* 385.(6710), pp. 715–717.

Ramey, Valerie A. (Nov. 2016). "Macroeconomic Shocks and Their Propagation". *Handbook of Macroeconomics* 2, pp. 76–162.

Ricke, Katharine et al. (Sept. 2018). "Country-level social cost of carbon". *Nature Climate Change* 8, pp. 895–900.

Rode, Ashwin et al. (2021). "Estimating a Social Cost of Carbon for Global Energy Consumption". *Nature* 598, pp. 308–314.

Roine, Jesper, Jonas Vlachos, and Daniel Waldenström (Aug. 2009). "The long-run determinants of inequality: What can we learn from top income data?" *Journal of Public Economics* 93.(8), pp. 974–988.

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

Sen, Amartya (Dec. 1973). *On Economic Inequality*. Oxford, UK: Oxford University Press.

Sen, Amartya, Angus Deaton, and Tim Besley (Aug. 2020). "Economics with a Moral Compass? Welfare Economics: Past, Present, and Future". *Annual Review of Economics* 12, pp. 1–21.

Sheffield, Justin, Gopi Goteti, and Eric F. Wood (July 2006). "Development of a 50-Year High-Resolution Global Dataset of Meteorological Forcings for Land Surface Modeling". *Journal of Climate* 19.(13). Dataset, pp. 3088–3111.

The World Inequality Lab (Feb. 2024). *Distributional National Accounts Guidelines: Methods and Concepts used in the World Inequality Database*. Dataset.

United Nations (July 2010). *System of National Accounts 2008*. Studies in Methods: Series F. New York: United Nations.

— (Sept. 2015). *Transforming our world: the 2030 Agenda for Sustainable Development*. Resolution adopted by the General Assembly on 2015-09-25.

Weitzman, Martin L. (Oct. 1974). "Prices vs. Quantities". *Review of Economic Studies* 41.(4), pp. 477–491.

A1 Data

A1.1 Construction of inequality indices

Data from the World Inequality Database can be extracted at the country-percentile level with variables reporting averages, minimums, maximums, and shares of national income. As grouped data, its use for the construction of inequality indices still requires the modeler to make assumptions about the distribution of household incomes within these disaggregated groups. This becomes increasingly influential for the construction of entropy-based inequality measures such as the mean logarithmic deviation or the Theil index for higher levels of inequality across groups.

For a given inequality measure J, the lower bound J_L is calculated by assuming perfect equality within all groups such that all individuals are assigned their group mean income. The upper bound J_U is calculated by assigning a proportion of the individuals within a group the minimum income of the group and the remainder the maximum income of the group.

Figure A1: Sensitivity of inequality indices to within-quantile inequality

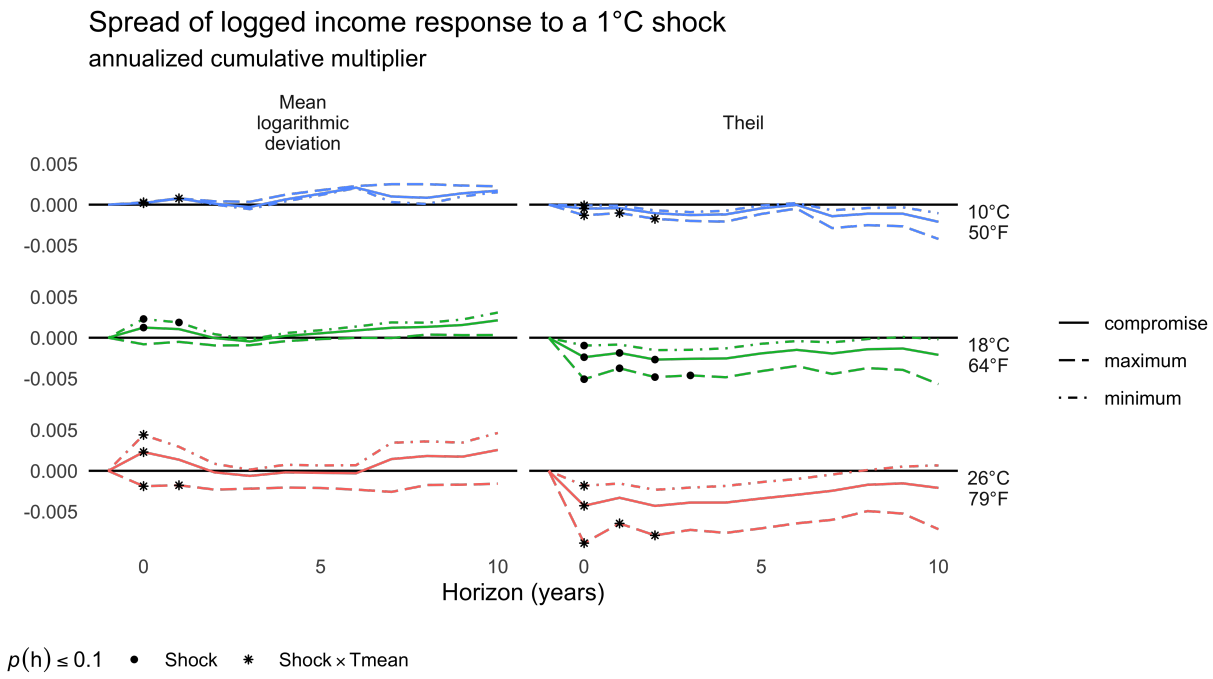

Cowell 2011 describes multiple 'compromise' assumptions between the two extremes, finding that they each approximate an average where the lower bound is given twice the weight of the upper bound. Figure A1 depicts IRFs of the two inequality measures under each of these three assumptions. For our purposes, we use the compromise construction whenever a choice is necessary.

A2 Methods

Here, we provide additional exposition of our application of the method of local projections for inferring a dynamic social response to the temperature shock described in Section 2.3. The supplementary analysis here uses country-level GDP from the Penn World Tables (Feenstra et al. 2015) as the outcome variable of interest. The results documented here may also be compared to other estimates in the prolific literature on global climate impacts estimated by top-down methods, which most commonly use similar country-level aggregates as outcomes.

A2.1 State-dependent impulse response of GDP

We first estimate Equation (3) on a country-year panel dataset using the h-period long difference in the logarithm of GDP as our outcome variable $\Delta_h y_{i,t+h}$. Figure A2 illustrates the estimated impulse response of GDP up to a projection horizon of 6 years after a unit temperature shock. Colors correspond to three representative temperatures depicted in different colors. The plots in the second row depict the same functions annualized by simplying dividing estimates at each projection horizon h by $h + 1$.

The contemporaneous estimates in the IRFs in the first row of Figure A2 depict a pattern now standard in the aggregate impacts literature featuring near-null temperature effects in cold climates, mild effects in mid-temperature climates, and the largest negative effects in hot climates. The evolution of these impacts over time is less commonly depicted and here produce a pattern qualitatively similar to those reported by Nath et al. wherein per capita GDP appears to stay persistently depressed over the entire projection horizon. The second row depicts the same estimates but annualized by scaling down estimates by their corresponding horizon $h + 1$ to more clearly depict convergence behavior relative to the contemporaneous effect of the shock.

Global inequality effects of climate change thus are not just dependent on the shape of the impulse response responses across quantiles but also the distribution of warming over space and time as well as the actual distributions of incomes across quantiles. The simulation exercises to be described in Section 4 are valuable because they account for four relevant climate inequalities: inequality of temperature effects across an income distribution, inequality in the incidence of temperature effects across different climates, inequality in the global distribution of income, and spatial inequality in the incidence of absolute warming.

A2.2 Inferring persistence from local projection estimates

The colored error bands depicted in Figure A2 reflect pointwise estimation error. This is particularly important to include in this context since the volume of regressions underlying a single LP-IRF precludes convenient summary in a single regression table.[15] One may be tempted to infer from the fact that the hot-country bands in red do not intersect the horizontal axis over the entire projection horizon that these estimates constitute significant evidence to reject a joint null hypothesis of level effects which persist for fewer than 15 years.

However, the uncertainty depicted by error bands corresponds only to individual hypothesis tests. Since estimates and residuals across an IRF are serially correlated—for example, the green bands straddle the horizontal axis but all estimates are negative— the relevant test of significance is a joint hypothesis test not represented by pointwise error. The significance of the main effect of

[15] For example, the results depicted in Figure A2 derive from 16 different regressions. Later results derive from as many as 91. All graphically depicted errors and significance tests presented in this paper correspond to a significance level of $\alpha = 0.1$.

Figure A2: State-dependent dynamic responses to a 1°C shock

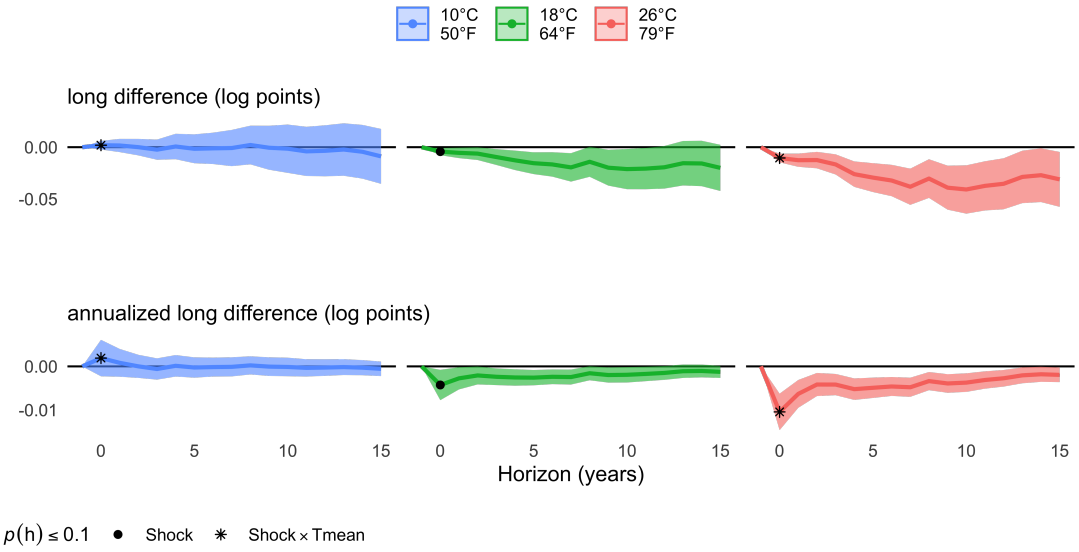

(a) State-dependent impulse response of GDP. The plots in the top row correspond to estimates of impulse responses of GDP estimated by the $H+1$ local projections regressions specified in Equation (3) for three representative temperatures depicted in different colors. The plots in the bottom row depict the same estimates and error bands annualized by dividing values at each horizon h by $h+1$

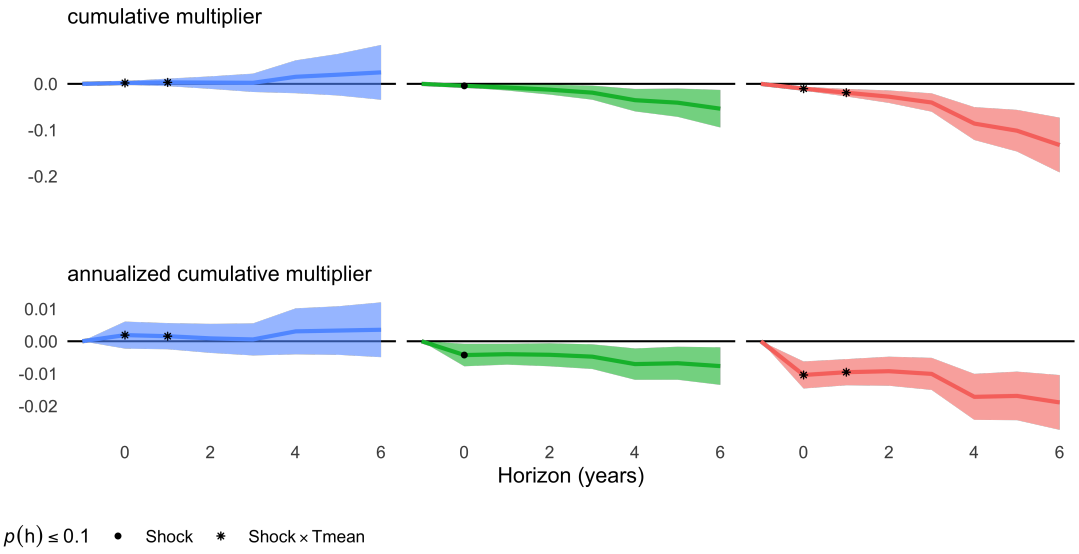

(b) State-dependent temperature multiplier of GDP. The cumulative multiplier is interpreted as the change in *cumulative* GDP over h periods attributable to a unit shock in $\hat{\tau}$, net of its impact on future shocks $\hat{\tau}$. Annualized multipliers represent the same estimates but divided by the projection horizon.

the impulse response $\mathcal{R}_{\tau \to y}(h)$ amounts to testing the null hypothesis that all main effects up to h are 0:

$$H_0(h) : \beta_{1,0} = \beta_{1,1} = ... = \beta_{1,h} = 0$$

Since the number of coefficients being considered in the joint hypothesis increases in h, the coverage of the joint hypothesis test must be adjusted for the inclusion of each additional horizon to compensate for the problem of multiple hypothesis testing. Inverting this joint null hypothesis, Inoue et al. (2025) derives convenient estimators for the implied "significance band" defined to satisfy the following condition:

$$\mathbb{P}\left[\bigcap_{h=0}^{H}\left\{\zeta_{\frac{\alpha}{2(H+1)}}\frac{\sigma_h}{\sqrt{T-h}} < \widehat{\beta}_h < \zeta_{1-\frac{\alpha}{2(H+1)}}\frac{\sigma_h}{\sqrt{T-h}}\right\}\right] \geq 1-\alpha$$

The scaling of the significance level α by $2(H+1)$ is an implementation of the Bonferroni correction. In practice, rejecting the null for the main effect corresponding to horizon $h+1$ is sufficient to reject the null of non-persistence up to horizon h. Graphically, these bands are represented by a region straddling the horizontal axis such that estimates of β_{1h} contained within it are insufficiently large enough to reject the hypothesis of a null effect that the impulse response persists up to period h, analogous to the bands drawn to measure persistence in time series correlograms. The width of these bands increases with the projection horizon because the available estimation sample for a near-balanced panel is reduced by one period per unit for each additional projection.

These significance bands are specific to the main effect and are represented by the gray regions in the middle panels of Figure A3. Since ours is a state-dependent model, I derive equivalent significance bands specific to the interaction effects. These are defined relative to the point estimates of the main effects and are depicted by the gray regions in the panels on the left and right. Points outside these regions are considered sufficient to reject the null hypothesis that interaction effects do not persist up to h periods. By the linearity assumption, the null hypothesis of zero interaction effects for a 10° climate is rejected if and only if it is rejected for a 26° climate.

For comparison, Figure A3 depicts a placebo test where I construct significance bands equivalent to those described above to test the persistence in the effect of temperature on another widely available and potentially non-stationary growth rate: that of population. As I am not aware of any theory or evidence to suggest temperature shocks would meaningfully impact country-level population growth[16], we should expect the tests to return precise null results. Indeed, we find with all point estimates resulting from two separate population datasets easily contained by the corresponding significance bands. Notably, if we were to misuse the error bands as indicators of statistical significance, we would misinterpret the estimated IRFs as evidence to reject the null for large subintervals of the projection horizon for mid-temperature and hot climates. Instead, the entirety of these error bands are easily enveloped by the significance bands.

Since including significance bands, color-coded error bands, and point estimates for multiple IRFs can clutter space, I adopt a convention of omitting the significance bands entirely and instead coloring black the point estimates which are located outside the implicit error bands. In all IRFs beginning with those depicted in Figure A2, rounded points correspond to significant main effects

[16]While there is an established literature on climate impacts on mortality rates (Barreca and Shimshack 2012; Barreca, Clay, et al. 2016; Carleton, Jina, et al. 2022), absolute magnitudes are miniscule relative to total populations

Figure A3: Impulse responses to a 1°C shock, including significance bands

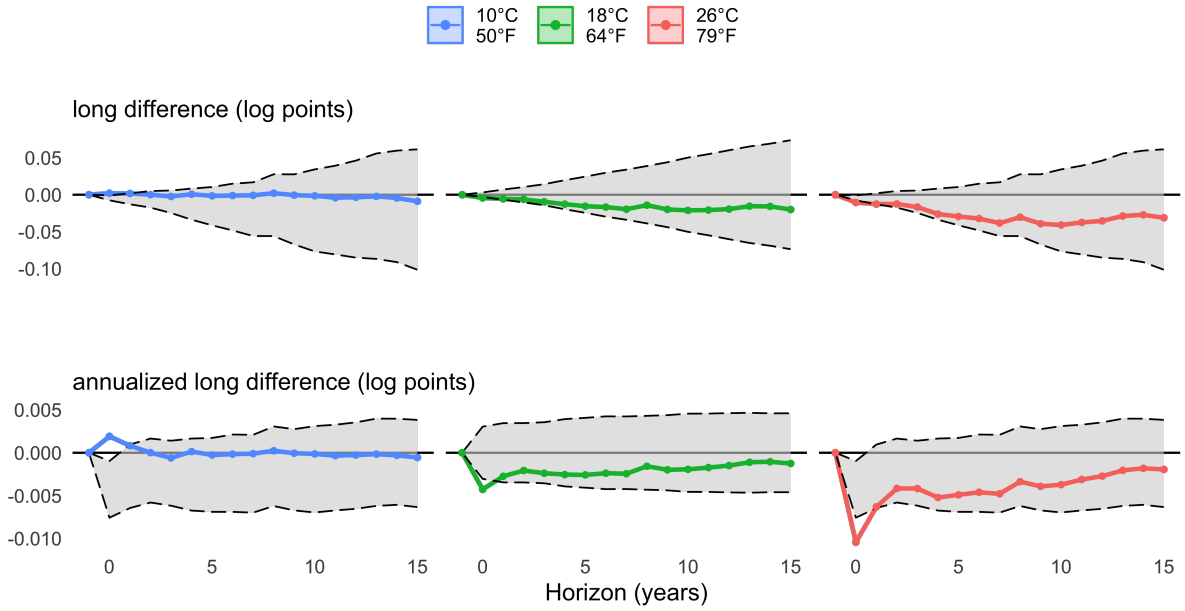

(a) GDP impulse response. Gray significance bands in the middle panels correspond to the threshold magnitudes for significant main effects while those in the left and right panels are centered on the point estimates of these main effects (green filed lines) and correspond to the threshold for significant state interaction effects.

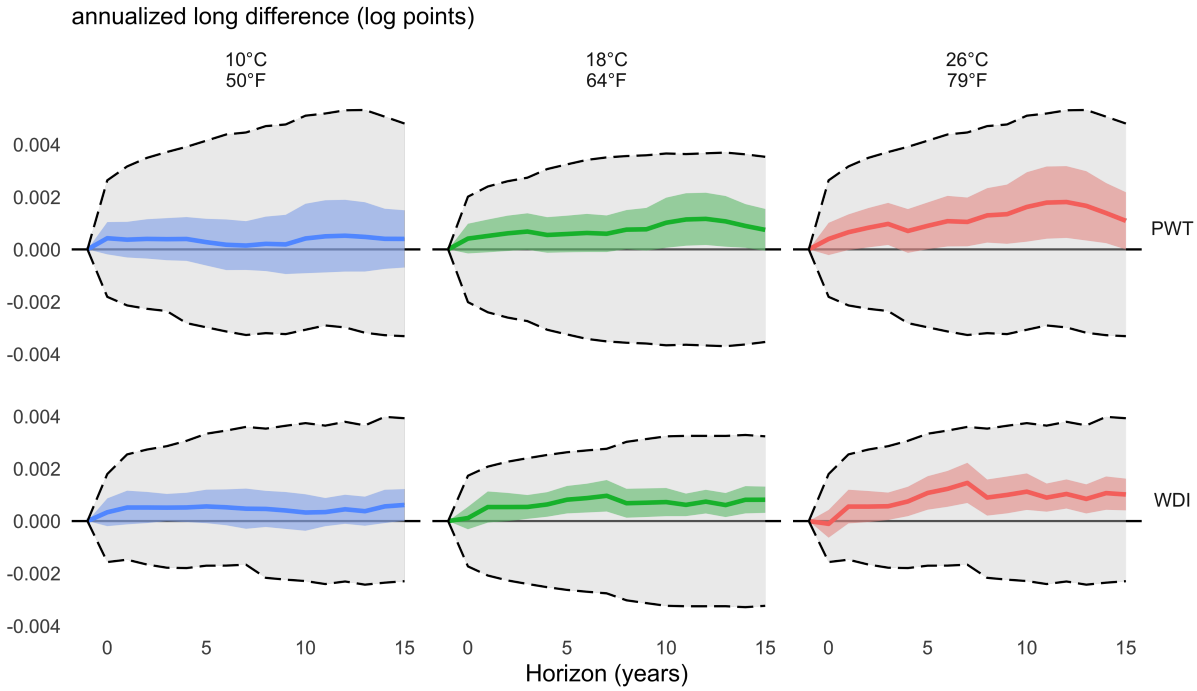

(b) Placebo test: national population response. Equivalent to the bottom row of panel (a) but for national population, an outcome variable we do not expect temperature to affect and with conventional error bands also included. Rows differ by source of population data.

while asterisks correspond to significant interaction effects. By this test, these results cannot reject the null hypothesis of one-period level effects.

A2.3 The state-dependent temperature multiplier

Another reason to prefer temperature shocks to levels of temperature as treatment variables is that models which use the latter neglect to sufficiently account for autocorrelation in temperature itself when estimating dynamic effects (Nath et al. 2024; Bilal and Känzig 2024). An analogy can be drawn to the estimation of multipliers in empirical macroeconomics: when estimating the impulse response of, for example, GDP to a fiscal policy shock, one must account for the tendency of stimulus programs to be implemented in stages or to be followed by additional stimulus. Otherwise, the effect sizes implied by simple impulse response functions misattribute cumulative changes in the outcome only to the initial unit shock; multipliers scale these effects at each horizon by the corresponding accumulation of shocks.

Figure 3 depicts the impulse response of temperature shocks to a pulse of itself for our three representative climates. Each exhibits very similar persistence patterns across two independent temperature dataset. An application of the persistence test covered in Section A2.2 implies autocorrelations which persist for up to 10 years though only the first 1-2 periods are qualitatively large (approximately 21% and 9% respectively with insignificant heterogeneity across climates).[17] Cumulative multipliers are estimated in order to scale down the estimated impulse response of GDP $\mathcal{R}_{\tau\to y}(h)$ by the impulse response of the shock to itself $\mathcal{R}_{\tau\to\tau}(h)$ to account for this dynamic treatment schedule.

Cumulative responses are traditionally computed by first constructing *cumulative* versions of the two IRFs: $\mathcal{R}^c_{\tau\to y}(h) := \sum_{j=0}^{h} \mathcal{R}_{\tau\to y}(j)$ for the outcome of interest y and $\mathcal{R}^c_{\tau\to y}(h) := \sum_{j=0}^{h} \mathcal{R}_{\tau\to\tau}(j)$ for the shock $\widehat{\tau}$. Then the cumulative temperature multiplier $m_{\tau\to y}(h)$ would be calculated as the ratio

$$m_{\tau\to y}(h) := \frac{\mathcal{R}^c_{\tau\to y}(h)}{\mathcal{R}^c_{\tau\to\tau}(h)}$$

However, this approach has a couple of inconvenient downsides. First, as noted by Jordà and Taylor (2025), computation of standard errors of a ratio of two random variables is complicated and non-standard. Second, it is computationally inefficient in the sense of requiring $2(h+1)$ (or $4(h+1)$ in our state-dependent case) separate local projections to obtain the corresponding multiplier for each projection horizon h.

Instead, we follow the recommendation to use the one-step method described in Ramey (2016) and applied in Ramey and Zubairy (2018). We adapt the method for our state-dependent setting through the procedure summarized in Algorithm A1.

We can use the resulting coefficients to trace a cumulative multiplier function $\mathcal{R}^c_{\tau\to y}(h;\overline{T})$ analogous to Equation (4) but which accounts for the effect of persistence in the shock. Since estimation of the coefficients follows directly from a system of local projection regressions, standard inference is preserved and the significance testing methods outlined in Section A2.2 conveniently remain available. Figure A2 demonstrates an application of the method to estimate temperature multipliers of GDP. The estimated multiplier over horizon h is interpreted as the difference in cumulative

[17] This contrasts with the finding of NRK who measure shock persistence using levels of temperature to infer that temperature shocks remain 10-20% higher for at least nine years following a unit shock.

GDP over the h periods following an isolated 1°C shock relative to a counterfactual where the shock had never occurred.

Estimated cumulative multipliers are found to be substantial at larger horizons but the significance band test cannot reject a null of non-persistence beyond one period. 90% confidence intervals for the effect of an identified temperature shock on cumulative GDP are [-0.49, 1.1] percent for 10°C climates and [-2.7,-1.1] percent for 26°C climates.

Algorithm A1: One-step state-dependent cumulative multiplier

for $h = 0$ **to** H **do**

Estimate the IRF $\mathcal{R}_{\tau \to \tau}(h)$ as defined in Equation (4), collecting state-dependent shock coefficients $\widehat{\alpha}_{1,h}, \widehat{\alpha}_{2,h}$;

Define $\Delta_h^c y_{i,t+h} := \sum_{j=0}^{h} \Delta_j y_{i,t+j}$;

Define $\widehat{\tau}_{i,t+h}^c := \widehat{\tau}_{it} \left[\sum_{j=0}^{h} \widehat{\alpha}_{1,j} + \widehat{\alpha}_{2,j} \overline{T}_{it} \right]$;

Estimate local projection:

$\Delta_h^c y_{i,t+h} = m_{1,h} \widehat{\tau}_{i,t+h}^c + m_{2,h} \widehat{\tau}_{i,t+h}^c \overline{T}_{it} + \lambda_h \overline{T}_{it} + \mathbf{Z}_{it} \boldsymbol{\gamma}_h + \mu_i + \eta_t + \varepsilon_{i,t+h}$;

end

Output: Multiplier coefficients $\widehat{m}_{1,h}, \widehat{m}_{2,h}, \widehat{\lambda}_h$ for $h \in \{0, ...H\}$

A3 Supplementary results

Figure A4: Temperature multipliers estimated from household surveys

(a) Impacts on income by decile

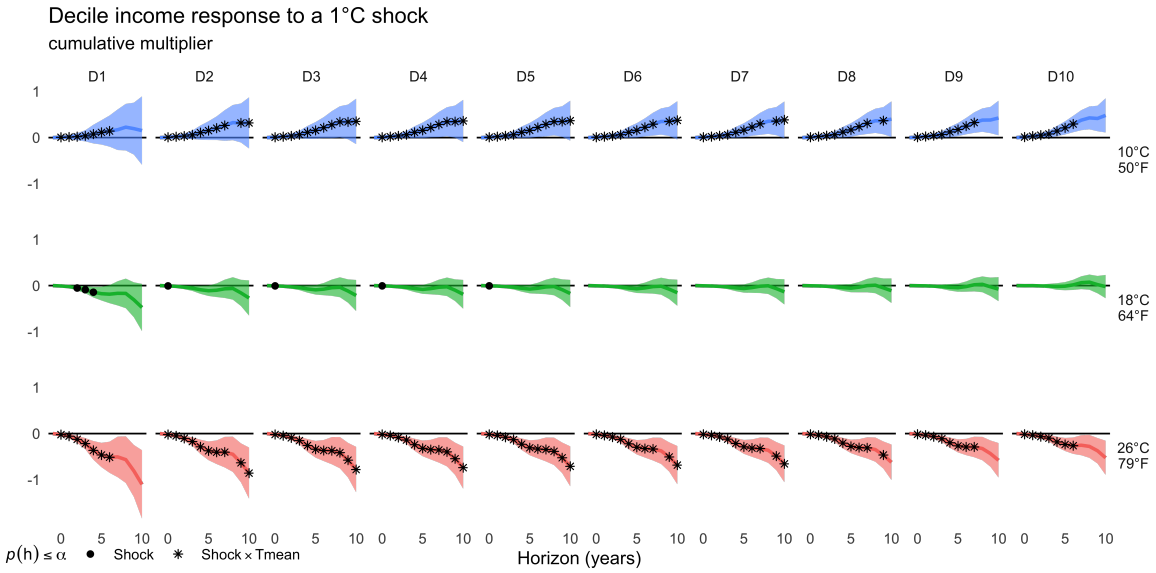

(b) Impacts on consumption by decile

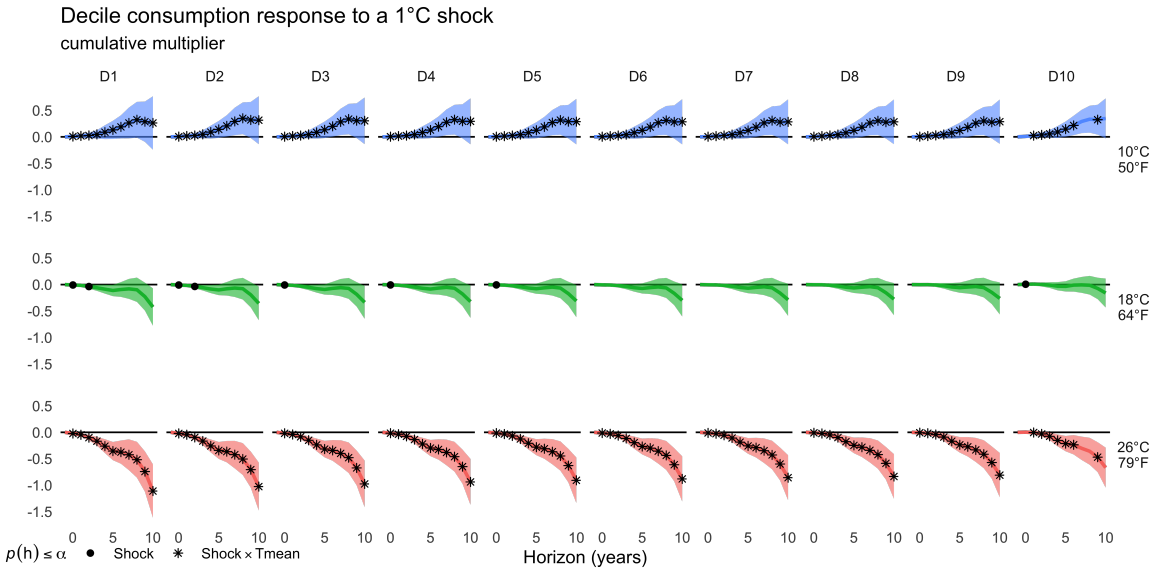

Data from Lahoti et al. (2016)

A4 Anthropogenic attribution

To attribute a global inequality effect to historical anthropogenic climate change, we implement a counterfactual analysis which decomposes observed temperature effects into their natural and anthropogenic contributions, adapting methods commonly used in climate attribution science.

A4.1 Climate simulation data

Climate simulation data comes from the contributions of climate modeling groups to the World Climate Research Program's Coupled Model Intercomparison Project (CMIP). The CMIP, initiated in 1995, is a framework to coordinate, standardize, and disseminate the results of state-of-the-art simulations of the global climate for the benefit of the international climate research community. Data used in this study comes from CMIP6, the sixth and newest[18] generation of models (Eyring et al. 2016). Data is provided as a collection of globally gridded time series summarizing the joint evolution of several hundred climate variables under various pre-specified calibrations. We will primarily be interested in the models' simulations of near-surface air temperature.

All participating models are "coupled", meaning that they explicitly account for interactions and feedback between distinct components of the global climate system such as the atmosphere, the cryosphere, land surface, and the ocean. Because the manner in which this integration is achieved is idiosyncratic to each model, computationally expensive, and highly sensitive to external calibrations, coordinating the set of "experiments" that participating models run is necessary for comparability and interpretability. For example, one of the four indexes which define an experiment is the realization index which corresponds to the set of geophysical initial conditions at the beginning of the simulation period. By holding the initial conditions fixed, differences in results across models within the same experiment can be more readily attributed to the distinguishing features of the model without being confounded by differences in implementation. Similarly, differences in the same climate model's results across different perturbations to the initial conditions provide a measure of the model's internal variability. The CMIP thus provides a systematic way of decomposing total variation across models and experiments both internal and external to the model. Coordination also enables identification of systematic and idiosyncratic biases in model designs that then inform the next generation of models.

Among the other three indexes which define an experiment, the forcing index is central to our counterfactual analysis. Forcings are factors which affect the net transfer of radiative energy in the climate system, essentially the energy the Earth receives from the sun minus that which it expels back into outer space. Anthropogenic forcings include, for example, industrial greenhouse gas emissions, deforestation, and the transformation of land for arable agriculture. Natural forcings include aerosol expulsions from major volcanic activity and variation in solar irradiance such as those associated with the 11-year solar cycle.[19] Comparing simulations which only differ in their forcings allow one to attribute differences in outcomes to the differences in forcings; we describe next how comparisons of this kind form the basis of both our historical and projection counterfactual analyses.

[18] After a delayed rollout, outputs from most CMIP6 models were made accessible by 2022

[19] Natural forcings are distinct from sources of natural variability such as the ENSO cycle because they are external to coupled climate models.

Figure A5: Model-implied spatial distributions of global warming, 1980-2014

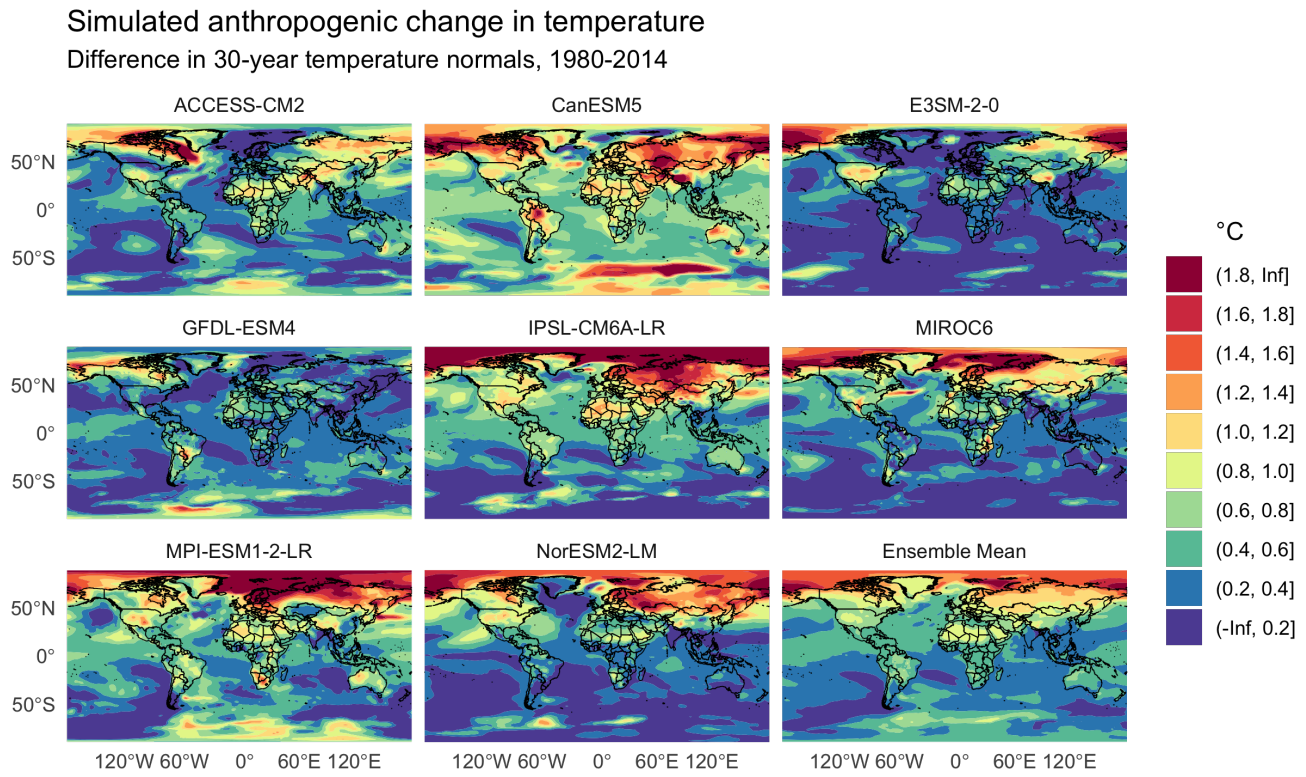

Simulated spatial distribution of anthropogenic warming for eight randomly selected climate models. The bottom-right plots the average distribution for all 13 CMIP6 models which ran both historical and historical-natural simulations.

A4.2 Construction of historical counterfactuals

Our historical counterfactual analysis uses simulations from a category of experiments called "historical runs". For the CMIP6 generation of models, historical simulations begin at a pre-industrial baseline period in 1850 and run through to 2014. Among the forcing scenarios used in these historical experiments is one simply labeled "historical", which attempts to represent the impacts of influential forcings actually imposed on the climate system over this period, such as the deforestation of the Amazon rainforest and the eruptions of Krakatoa in 1883 and Mt. Pinatubo in 1991. Results from this historical run are required of all participating CMIP6 models since these "hindcasts" can be directly compared to observational data in order to evaluate model accuracy. Models collectively perform remarkably well at reproducing the historical record (Fan et al. 2020; Zelinka et al. 2020; Yazdandoost et al. 2021). For this reason, studies which use CMIP data for the purpose of counterfactual analysis are advised to include the full available "multi-model ensemble" for experiments of interest (Tebaldi and Knutti 2007; Knutti et al. 2010).

A subset of these models also simulate a set of optional 'historical-natural' experiments which are defined identically except with anthropogenic forcings held fixed at pre-industrial levels. The differences between the two forcings by the same model holding all other variables fixed is then interpreted as the model's simulation of anthropogenic contributions to climate change. Comparisons between historical and historical-natural experiments are commonly used in the subfield of attribution science concerned with quantitatively measuring anthropogenic contributions to the intensity, frequency, or probability of weather events or trends.

For each climate model m and country i, we use an observed series T and the two model-specific historical series $\mathcal{T}^{hist}$ and $\mathcal{T}^{nat}$ to construct a counterfactual series $\widetilde{T}$ where anthropogenic forcings have been held fixed since the first year of the simulation period. The iterative process is summarized in Algorithm A2 and illustrated in Figure A6 using data from the Philippines as an example. In panel (a), differences between model-specific simulations of national temperature when including all forcings (orange) and when excluding anthropogenic forcings (green horizontal axis). Local regression fits are depicted in red. Light-blue shaded regions correspond to 30-year reference periods used to define temperature normals (dark blue dots) for the start and end of the simulation period 1980-2014.

In panel (b), The linearized difference between normals depicted as dark blue lines in (a) are subtracted from the observed temperature series (black) starting in 1980 to produce counterfactual temperature series (green). Each of these is interpreted as the temperature history that would have been realized had anthropogenic forcings been held fixed at 1980 levels according to a specific model of the climate system.

In panel (c), linearizing the difference between the climate normals defined at either end of the simulation period and subtracting the result from an observed temperature series implements the standard "delta method" for constructing bias-corrected climate counterfactuals. Counterfactual economic series for a random selection of five climate models are depicted in green and are constructed by iteratively applying the decile-level dynamic income responses depicted in Figure 4 to the observed and counterfactual temperatures depicted in (b). Their difference is interpreted as the change in income attributable to anthropogenic forcings since 1980. Green bands correspond to 90% confidence intervals generated from 500 bootstrap estimates of the response functions.

Figure A6: Constructing counterfactuals holding anthropogenic forcings fixed

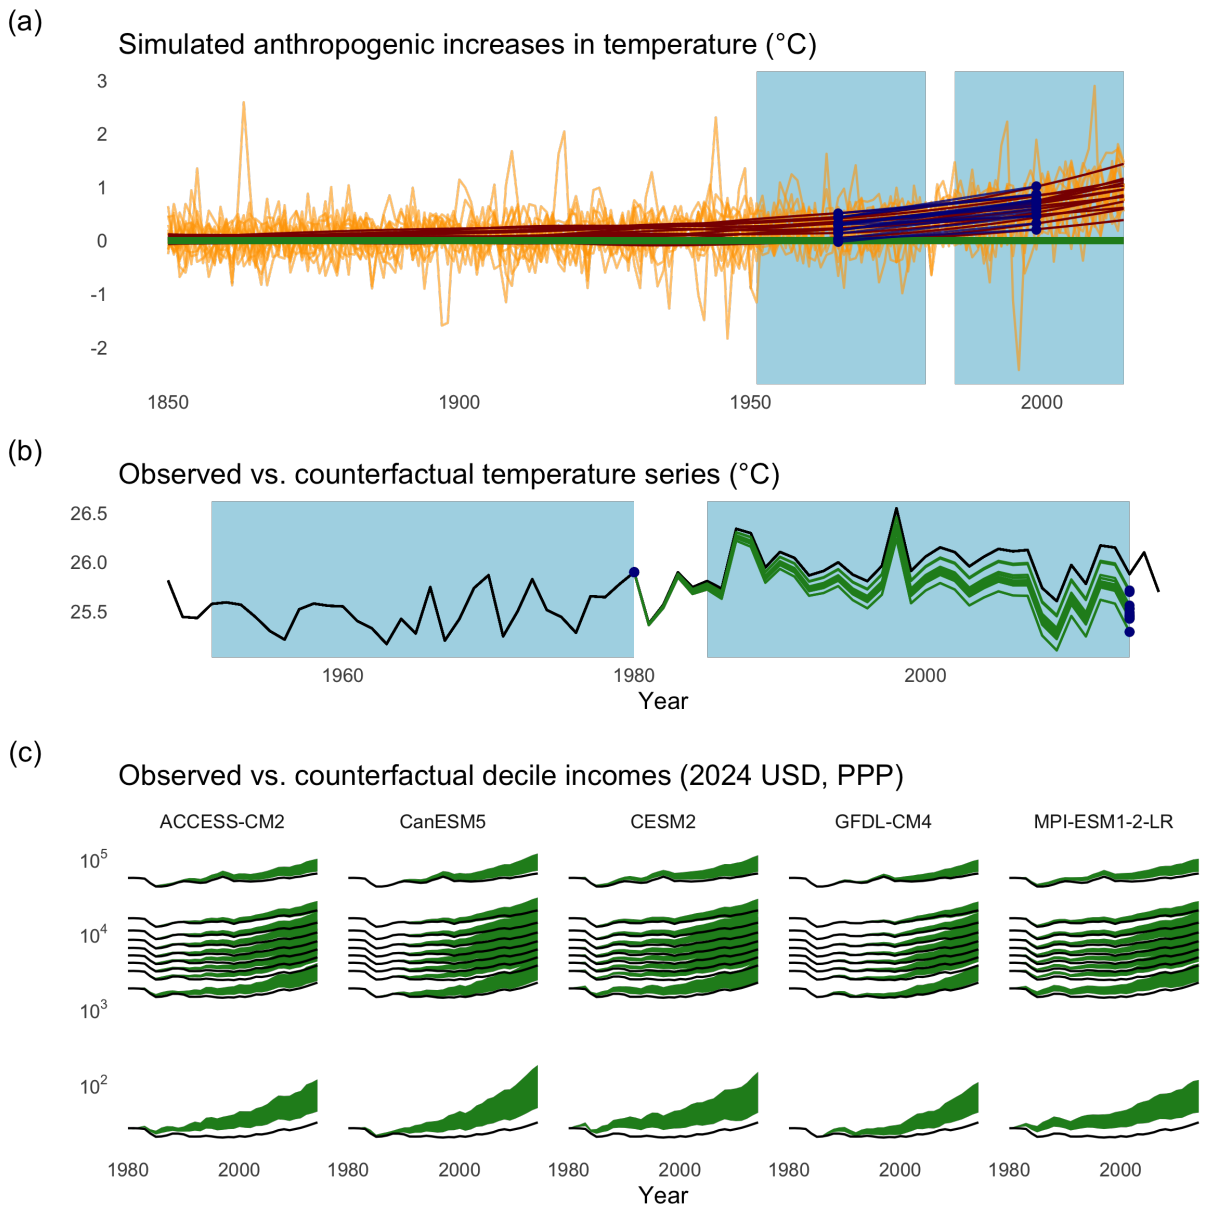

Demonstration of the delta method for constructing bias-corrected counterfactuals using data for the Philippines as an example.

Algorithm A2: Constructing temperature counterfactuals holding anthropogenic forcings fixed at 1980 levels

Input: $\{\mathcal{T}^{hist}_{m,i,t}\}$,$\{\mathcal{T}^{nat}_{m,i,t}\}$,$\{T_{it}\}$

for $t^* \in \{1980, 2014\}$ **do**

$\overline{\mathcal{T}}^{hist}_{m,i,t^*} \leftarrow \frac{1}{30}\sum_{j=1}^{30} \mathcal{T}^{hist}_{m,i,t^*-j}$;

$\overline{\mathcal{T}}^{nat}_{m,i,t^*} \leftarrow \frac{1}{30}\sum_{j=1}^{30} \mathcal{T}^{nat}_{m,i,t^*-j}$;

$\overline{\delta}_{m,i,t^*} \leftarrow \overline{\mathcal{T}}^{hist}_{m,i,t^*} - \overline{\mathcal{T}}^{nat}_{m,i,t^*}$;

end

for $t = 1980$ **to** 2014 **do**

$\overline{\delta}_{m,i,t} \leftarrow (t - t_0)\frac{\overline{\delta}_{m,i,2014} - \overline{\delta}_{m,i,1980}}{2014-1980}$;

$\widetilde{T}^{nat}_{m,i,t} \leftarrow T_{it} - \overline{\delta}_{m,i,t}$;

end

Construct states $\{\overline{T}^{nat}_{m,i,t}\}$ by applying Equation (1a) to $\{\widetilde{T}^{nat}_{m,i,t}\}$ then subtracting by $\widetilde{T}^{nat}_{m,i,1980}$;

Construct shocks $\{\widehat{\tau}^{nat}_{m,i,t}\}$ by applying Equation (1a) and Equation (1b) to $\{\mathcal{T}^{nat}_{m,i,t}\}$;

Output: $\{\widetilde{T}^{nat}_{m,i,t}\}$,$\{\overline{T}^{nat}_{m,i,t}\}$,$\{\widehat{\tau}^{nat}_{m,i,t}\}$

Algorithm A3: Constructing economic counterfactuals holding anthropogenic forcings fixed at 1980 levels

Input: $\{\widetilde{T}^{nat}_{m,i,t}\}$,$\{\overline{T}^{nat}_{m,i,t}\}$,$\{\widehat{\tau}^{nat}_{m,i,t}\}$, B bootstrap samples indexed b

Initializing series:

for $b = 1$ **to** B **do**

Specify desired simulation interval: $(t_0, t_1) \leftarrow (1980, 2014)$;

Estimate $\mathcal{R}^{c}_{\tau \to Y}(h; \overline{T}, b)$ from Equation (5);

Specify maximum projection horizon H using persistence test from Section A2.2;

Define IRFs $\mathcal{R}_{\tau \to Y}(h; \overline{T}, b) := \mathcal{R}^{c}_{\tau \to Y}(h; \overline{T}, b) - \mathcal{R}^{c}_{\tau \to Y}(h-1; \overline{T}, b)$;

Collect IRF coefficients;

$\{\widehat{\beta}_{1,h,b}, \widehat{\beta}_{2,h,b}, \widehat{\lambda}_{h,b}\}_{h=0}^{H}$;

Define $\delta^{\tau}_{m,b,i,t} := \widehat{\tau}^{nat}_{m,i,t} - \widehat{\tau}_{it}$;

Define $\delta^{\tau T}_{m,b,i,t} := \widehat{\tau}^{nat}_{m,i,t} \cdot \overline{T}^{nat} - \widehat{\tau}_{it} \cdot \overline{T}_{it}$;

Define $\delta^{T}_{m,b,i,t} := \overline{T}^{nat}_{m,i,t} - \overline{T}_{it}$;

Initialize $\delta^{Y}_{m,b,i,t} \leftarrow 0$;

Initialize $\widetilde{Y}^{nat}_{m,b,i,t_0} := Y_{i,t_0}$

end

Iteratively constructing series:

for $t = t_0 + 1$ **to** t_1 **do**

for $b = 1$ **to** B **do**

for $h = 0$ **to** H **do**

$\delta^{Y}_{m,b,i,t+h} \leftarrow \delta^{Y}_{m,b,i,t+h} + \widehat{\beta}_{1,h,b} \cdot \delta^{\tau}_{m,b,i,t} + \widehat{\beta}_{2,h,b} \cdot \delta^{\tau T}_{m,b,i,t} + \widehat{\lambda}_{h,b} \cdot \delta^{T}_{m,b,i,t}$;

end

$\widetilde{Y}^{nat}_{m,b,i,t} \leftarrow \left(1 + \log(\frac{Y_{it}}{Y_{i,t-1}}) + \delta^{Y}_{m,b,i,t}\right) \widetilde{Y}^{nat}_{m,b,i,t-1}$

end

end

Output: $\widetilde{N}^{nat}_{m,b,i,t}$

Figure A7: Anthropogenic growth incidence relative to 1980 baseline by climate model

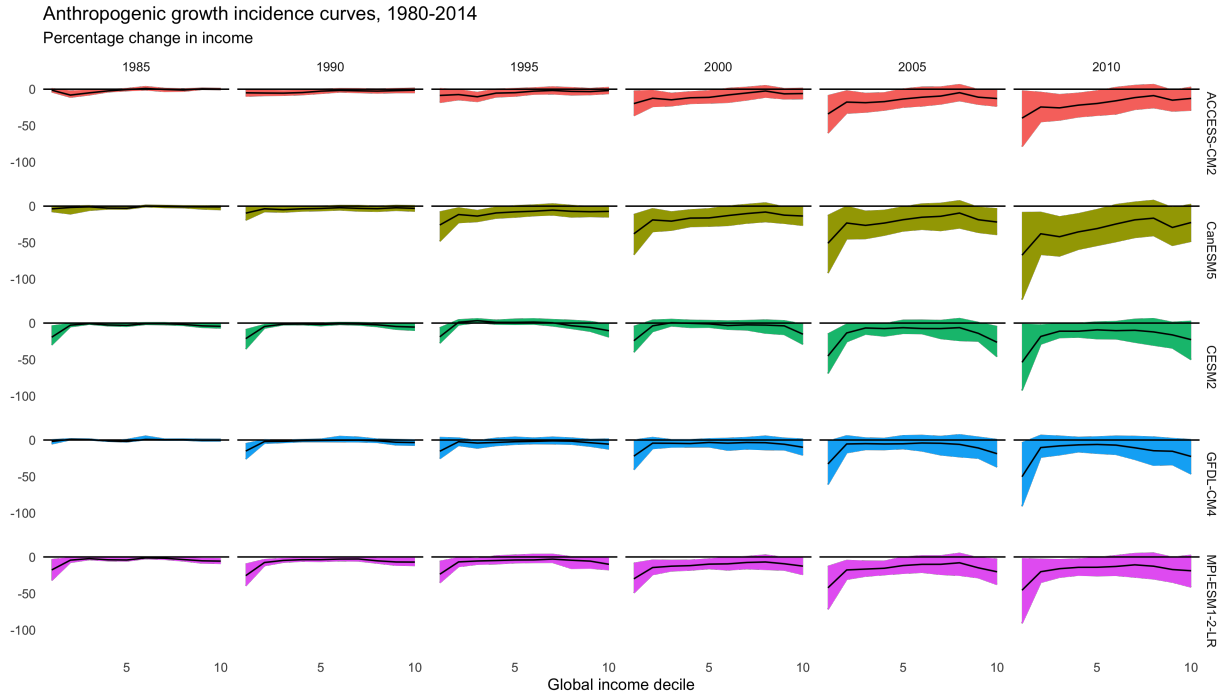

REFERENCES FOR APPENDIX

Barreca, Alan, Karen Clay, et al. (Feb. 2016). "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century". *Journal of Political Economy* 124.(1), pp. 105–159.

Barreca, Alan I. and Jay P. Shimshack (Oct. 2012). "Absolute Humidity, Temperature, and Influenza Mortality: 30 Years of County-Level Evidence from the United States". *American Journal of Epidemiology* 176.(7), pp. 114–122.

Bilal, Adrien and Diego R. Känzig (Nov. 2024). *The Macroeconomic Impact of Climate Change: Global vs. Local Temperature*. NBER Working Paper 32450. National Bureau of Economic Research.

Carleton, Tamma, Amir Jina, et al. (Nov. 2022). "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". *Quarterly Journal of Economics* 137.(4), pp. 2037–2105.

Cowell, Frank (Mar. 2011). *Measuring Inequality*. Oxford University Press.

Eyring, Veronika et al. (May 2016). "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization". *Geoscientific Model Development* 9.(5), pp. 1937–1958.

Fan, Xuewei et al. (Oct. 2020). "Global surface air temperatures in CMIP6: Historical performance and future changes". *Environmental Research Letters* 15.(10).

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (Oct. 2015). "The Next Generation of the Penn World Table". *American Economic Review* 105.(10). Dataset, pp. 3150–3182.

Inoue, Atsushi, Òscar Jordà, and Guido M. Kuersteiner (Feb. 2025). "Inference for local projections". *Econometrics Journal*.

Jordà, Òscar and Alan M. Taylor (Mar. 2025). "Local Projections". *Journal of Economic Literature* 63.(1), pp. 59–110.

Knutti, Reto et al. (Jan. 2010). *Meeting Report of the IPCC Expert Meeting on Assessing and Combining Multi Model Climate Projections*. Tech. rep. Geneva, Switzerland: Intergovernmental Panel on Climate Change.

Lahoti, Rahul, Arjun Jayadev, and Sanjay G. Reddy (July 2016). "The Global Consumption and Income Project (GCIP): An Overview". *Journal of Globalization and Development* 7.(1).

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (July 2024). *How Much Will Global Warming Cool Global Growth?* NBER Working Paper 32761. National Bureau of Economic Research.

Ramey, Valerie A. (Nov. 2016). "Macroeconomic Shocks and Their Propagation". *Handbook of Macroeconomics* 2, pp. 76–162.

Ramey, Valerie A. and Sarah Zubairy (Apr. 2018). "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data". *Journal of Political Economy* 126.(2).

Tebaldi, Claudia and Reto Knutti (Aug. 2007). "The use of the multi-model ensemble in probabilistic climate projections". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365.(1857).

Yazdandoost, Farhad et al. (Mar. 2021). "Evaluation of CMIP6 precipitation simulations across different climatic zones: Uncertainty and model intercomparison". *Atmospheric Research* 250.

Zelinka, Mark D. et al. (Jan. 2020). "Causes of Higher Climate Sensitivity in CMIP6 Models". *Geophysical Research Letters* 47.(1).