Climate inequality

Matthew Alampay Davis*

Job market paper

November 18, 2024
link to latest version

Abstract

This paper characterizes the relationship between anthropogenic climate change and global inequality, a subject intersecting two of the defining challenges of the 21st century but which remains virtually unstudied at subnational levels. To make overdue progress addressing this research gap, I improve upon frontier methods for estimating climate impacts to overcome outstanding methodological limitations and apply these refinements to fine-scaled data on globally representative income distributions. Next, I document new evidence that temperature shocks significantly and persistently impact distributions of income within countries, an effect driven by concentrations of harm onto the lowest income-earners in warm climates as well as a surprising vulnerability of the top 1% in these countries to environmental shocks. Integrating these inequality effects over observed distributions of income and the spatial incidence of global warming, I find that climate change between 1981 and 2016 regressively redistributed global shares in a reduced stock of income both between and within countries largely by depriving the world's poorest of economic opportunity that would otherwise have been available. Altogether, these results constitute the most comprehensive evidence yet of the regressive impact of climate change on global inequality.

*wm.alampaydavis@gmail.com

I thank Suresh Naidu, Jack Willis, David F. Hendry, and Jeffrey Shrader for their support and patience throughout the preparation of this manuscript. I thank Austeja Makareviciute and Aditya Jain for valuable research assistance. Feedback from my colleagues at Columbia's development seminar, sustainable development seminar, development colloquium, political economy colloquium, and student research breakfast helped improve this work. I acknowledge funding support from the Program for Economic Research.

1. Introduction

Mitigating anthropogenic climate change and constraining economic inequality are widely regarded as two of the defining social challenges of the 21st century. They are foundational to the 2030 Agenda adopted by all United Nations members, whose Sustainable Development Goals explicitly resolve to "take urgent action to combat climate change and its impacts" and "reduce inequality within and among countries" (United Nations, 2015). These are also understood to be intimately related objectives: the Intergovernmental Panel on Climate Change identifies "widening opportunities between and within countries" (Masson-Delmotte et al., eds, 2018) as first-order considerations in any sufficient policy response to climate change and the World Inequality Report remarks that "climate mitigation is largely a distributional issue, not only between countries but also within them" (Chancel et al., 2022).

These linkages are also appreciated by those who directly study the social impacts of climate change; the overwhelming majority of published environmental economists surveyed in 2021 indicated a belief that climate change is "likely" or "very likely" to exacerbate inequalities between (89%) and within (71%) countries (Howard and Sylvan, 2021). However, there remains virtually no evidence characterizing the economic incidence of climate change across the Earth's eight billion people beyond what is captured by aggregate differences between fewer than 200 national boundaries. Addressing this knowledge gap is an increasingly vital priority for understanding and equitably responding to the interconnected and cascading crises to come.

In fact, Chancel et al. (2022) estimates that within-country differences have comprised a rapidly growing majority portion of global inequality since the emergence of the modern international order in the second half of the 20th century. By the group-decomposable Theil measure, between-country differences now account for just over 30% of global income inequality, down from 57% as recently as 1980. Thus, an analysis of inequality which remains agnostic to material inequalities between persons risks increasing irrelevance exactly when it is most urgent. This paper aims to initiate and stimulate progress in this area by providing the most comprehensive characterization yet of the effects of climate change on global income inequality.

Contributing to this agenda requires overcoming several challenges. First, prevailing empirical methods commonly used to estimate climate impacts require refinement. Even in standard country-level analyses using the same sources of data, estimated impacts of climate change are known to diverge by orders of magnitude depending on the chosen specification for the "damage function" which causally relates variation in weather to variation in a

social outcome of interest (Pindyck, 2013). To avoid importing this imprecision into a nascent research setting, I extend cutting-edge impact models to incorporate features which address identification concerns flagged in recent literature. This methodological contribution is presented in Section 2 using an expositional application to conventional aggregate impacts for comparison to existing literature. In Section 2.2, I show how new tools from empirical macroeconomics (Inoue et al., 2024) may be used to conduct direct statistical tests of the persistence dynamics whose unfalsifiability has largely driven the extreme instability of climate damage estimates reported in the literature (Newell et al., 2021).

I proceed to apply this empirical framework to the most recent update to the World Inequality Database, which provides annual near-balanced percentile-level pre-tax income estimates for an expanded set of countries now representing almost the entire global population. In combination, this strictly improves the frontier of an analytical tradeoff between disaggregated analysis and global coverage which has historically characterized distributional analyses of climate change. As a result, we are able to document in Section 2.5 new empirical evidence demonstrating that the impacts of identified temperature shocks on within-country distributions of income are substantial and driven primarily by extreme tail effects in warm climates. More specifically, estimates imply substantial transfers of national income shares away from bottom quantiles and, perhaps surprisingly, the top 1% in warm climates. Since impacts to warmer countries are significantly negative across all quantiles, these inequality effects largely capture regressive redistribution of claims to diminished stocks of national income relative to the counterfactual absent a unit temperature shock.

In Section 3, we aggregate our analysis to produce the first reported estimates of the distributional impact of climate change on global income. This entails integrating the reduced-form within-country impacts over historically observed distributions of income and climate over the maximum observational period between 1981 and 2016. We find that climate-induced asymmetries in the incidence of temperature shocks over this 36-year period have unambiguously exacerbated global inequality, mainly through a high likelihood of depriving income growth for roughly the poorest 20% of global income earners. For the poorest percentile in our data, expected incomes absent climate impacts would be, conservatively, 29% higher than they are today. Evidence is also consistent with the standard finding in the established literature on aggregate impacts that climate change is very likely to have reduced total incomes over this period though net impacts are uncertain for the global middle class. In Section A1, I apply methods from climate attributional sci-

ence to state-of-the-art global climate models to quantify the contributions of anthropogenic forcings to these inequality effects.

Section 4 concludes by considering how these new findings complement and complicate existing aggregate-level evidence on global climate impacts. I close by considering the limitations of the present study design and propose promising lines of future research recognizing distributional welfare impacts as vital to the economic analysis of climate change and the design of equitable climate policy.

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 absence of 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 supporting a prevailing utilitarian cost-benefit paradigm.

Top-down approaches restrict attention to a widely available outcome of interest. This outcome, most commonly gross domestic production, is taken to be 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 et al. (2015). Alternatively, bottom-up approaches consider impacts to distinct "sectors" of an economy; mortality (Carleton et al., 2022), energy consumption (Rode et al., 2021), and labor (Graff Zivin and Neidell, 2014) are common areas of focus. Results using this approach can be presented as partial costs or combined with other sectoral damages in a integrated model which accounts for their interdependencies. The advantage of the former approach is in its convenience and interpretability: the wide availability and international comparability of economic aggregates such as those provided by national accounts substantially simplifies the task of calculating global climate impacts. The benefit of the latter approach is in its ability to account for disaggregated, distributional, and non-market impacts to which single-dimensional aggregates are largely indifferent.

Conversely, the disadvantages of the two approaches jointly constrain the economic analysis of climate change; the simplicity of the top-down approach and prohibitive data and resource demands of the bottom-up approach pose a critical tradeoff between geographic disaggregation and global coverage. In perhaps the most prominent existing work directly concerned with charac-

terizing global climate inequality, Diffenbaugh and Burke (2019) estimate 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. But this is necessarily calculated by assigning all individuals their country's mean GDP and defining deciles by discrete national population sizes, 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 rare example which makes the opposite compromise between specificity and breadth in order to make progress in this area. 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 inequality implications of global warming, analyses must extend to settings which lack comparably rich or granular data.

The disaggregated distributional focus of this paper complements a much more established body of work reporting unequal impacts specifically between countries. Although estimates of long-term cumulative impacts are a notorious and influential area of disagreement (Newell et al., 2021), there is now general consensus that non-linearities in the temperature response of aggregate variables concentrate climate damage and risk in warmer, generally already poorer countries, exacerbating between-country global inequality (Howard and Sylvan, 2021). These features motivate the methodological refinements described in Section 2 to address the extreme divergence in dynamic effects while maintaining this state-dependency broadly accepted by the literature. Like Diffenbaugh and Burke (2019), a subset of these studies explore the implications of this between-country dimension of inequality directly, for example by highlighting the inverse relationship between country-level contributions to cumulative global greenhouse gas emissions and exposure to their social consequences (Chancel and Piketty, 2015; Ricke et al., 2018; Kotz et al., 2024).

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 (Goodman 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 remains comparatively sparse in large part due to an absence of adequate research infrastructure. Estimation of distributional equivalents to 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 (European Commission et al., 2008). 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 from a wide array of official and survey sources date at least as far back as Deininger and Squire (1996), the validity of these products for cross-country causal analysis has long been known to be limited by inconsistencies across primary sources in construction and availability (Atkinson and Brandolini, 2001). For example, the most recent rigorous evaluation I could find of these secondary datasets concluded that results derived from the conveniently packaged Standardized World Income Inequality Database should be regarded as "not sufficiently credible" and that analyses using the United Nations' long-running World Income Inequality Database should at best be accompanied by justifications for their sample and variable selection criteria (Jenkins, 2015). Nonetheless, these inequality products have been used in at least two recent studies (Cevik and Jalles, 2023; Gilli et al., 2024) relating the inequality measures provided in these products to variations in weather.

In the absence of credible global inequality measures, a handful of works have managed to contribute to the study of within-country inequality impacts by using idiosyncratic data available for individual countries. Hsiang et al. (2017), described earlier, is a particularly impressive example. 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 by 2100 under a moderate warming scenario.

The research programs of Thomas Piketty, Emmanuel Saez, and their coauthors 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 "distributional national accounts" (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 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 (Office of 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. Weather shocks and within-country inequality

2.1 Identifying temperature shocks

Until recently, it has been customary in the empirical climate impacts literature to use raw levels of weather variables such as temperature or rainfall as primary explanatory variables of interest. This convention is motivated by the intuition that annual fluctuations in temperature are plausibly exogenous with respect to most economic activity. Recent work such as Kahn et al. (2021), Bilal and Känzig (2024), and Nath et al. (2024) challenge this convention on the grounds that increased production correlates with weather through its association with greater emissions of greenhouse gases and that temperature levels exhibit substantial autocorrelation which confounds estimation. In place of raw temperature, they each propose different constructions of a temperature *shock* intended to better represent adaptive expectations and satisfy exogeneity conditions for valid dynamic inference.

Ramey (2016) outlines these conditions for general time series settings. For a candidate treatment variable to represent a valid shock, it must correspond to unanticipated movements in exogenous variables, be exogenous with respect to current and lagged endogenous variables, and be uncorrelated with other exogenous shocks included in the regression model. We follow the construction of Kahn et al. (2021) in defining shocks $\widehat{\tau}_{it}$ with respect to an M-period moving average of local temperature:

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

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

Figure 1 depicts country time series of the resulting temperature shocks with colors and flexible local regression curves corresponding to an even split of all country-years into three groups defined by their average temperature. The positive values of the local regression curves over almost the entire observation period capture the idea that systematic warming is characterized by the increased relative frequent and magnitude of "hot" shocks compared to "cold" shocks even if expectations about weather adapt to recent weather. Intervals where the running average shock is closer to 0 correspond to relative stability in the (lagged) climate where imperfect or lagged adaptation is less damaging. The marginal distribution plots on the right demonstrate that colder seasonal countries experience much larger temperature shocks than their warmer counterparts, a reflection of both greater natural weather variability and the fact that anthropogenic absolute warming is greater in colder

regions furthest from the equator[1].

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

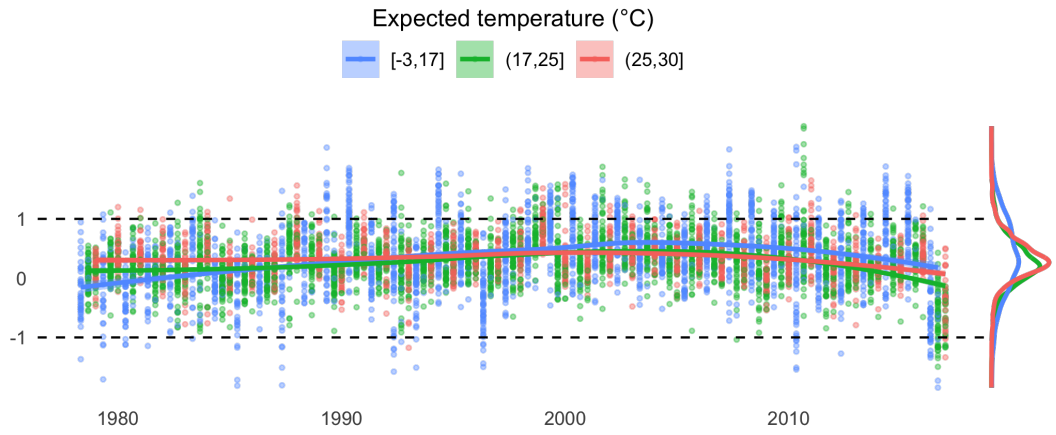

Points correspond to shocks observed in individual country-years. The three colors correspond to an equal-sized binning of country-years by their moving-average temperature. Flexible local regression curves describe time trends split by this binning. Marginal distribution plots on the right show that colder country-years observe greater variance in shocks.

As with the other candidate formulations, shocks defined this way satisfy the exogeneity conditions when lags are included as controls in the regression model being estimated.

I consider the simple definition above to be the most compelling construction of a temperature shock for several reasons. First, for sufficiently large M, this definition of the state variable $\overline{T}_{it}$ has an appealing coherence with the climatological convention of using long-run moving averages of weather realizations to anchor the climate at a fixed point in time. Most commonly, these "climate normals" are defined as 30-year averages to account for the fact that even a stable climate may exhibit cyclical weather patterns spanning several years. 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 across

[1]Figure A1 provides visualizations of the spatial distribution of anthropogenic warming for various global climate models

its three phase cycles which may span 2-7 years. Relatively minor oscillations can span multiple decades.

Secondly, these definitions of the shock and state variables are much more consistent with the empirical literature which generally reports "limited evidence" of adaptive behavior moderating the direct impacts of weather shocks in the context of climate change (Burke et al., 2024). Here, the state $\overline{T}_{it}$ may be taken to represent the time-varying expected value of realized temperature, capturing beliefs about the unobserved probability distribution from which the weather realized in country i at time t is drawn, i.e., the local climate. Since the climate at time t would be best represented by the climate normal centered on time t, which is a function of future realizations, there exists an $M/2$-period lag between the "true" climate and the expected climate $\overline{T}_{it}$. Over the span of our historical data, we find that this difference is non-trivial: the median country observes, on average, a 0.26°C (IQR: 0.18-0.36) increase in the 30-year moving average every 15 years.

This lag between true and expected climate can be interpreted as a bound on the ability for belief or anticipation effects to offset the direct effects of weather shocks, corresponding to climate adaptation in the sense of Kahn et al. (2021) and Hsiang (2016)). Alternatively, these lagged expectations may also be intuitively rationalized in a model of limited or costly adaptive capacity; even a farmer internalizing climate change trends will not transform their land and re-optimize their crop mix for a marginally warmer climate each year. This imperfect adaptability gives rise to the positive trends we observed in Figure 1 during periods of more rapid climate change. To the extent that local expectations about the climate are anchored by recent experience, shocks defined this way may be interpreted as deviations relative to adaptive expectations. It is intuitive that one of the main channels through which systematic climate change imposes economic harms is through the increased frequency and magnitudes of positive "hot" shocks relative to negative "cold" shocks and the implausibility of perfect adaptation.

This contrasts with the constructions of Bilal and Känzig (2024) and Nath et al. (2024) (henceforth, NRK), which define shocks as residuals from autoregressive distributed lag models of temperature. In so doing, they implicitly assume a hyper-responsive model of adaptation wherein beliefs are optimally set each period to costlessly minimize the 'surprisingness' of that period's temperature and thereby reducing direct economic effects of temperature shocks. As a result, their time series of observed temperature shocks are shown to be stationary with zero mean even in the context of systematic climate change. Thus, by the linearity assumption described in Section 2.1 wherein cold shocks have the exact opposite effect of hot shocks, calculating

the accumulated effect of shocks constructed in this manner would net out to zero economic damage. Damages from climate change seem to be recovered in these models only by annualizing long-term systematic warming and assuming the same agents do not internalize the gradual long-term trends despite their perfect and costless adaptability to larger short-term shocks.

Estimating responses to a temperature shock

We adapt the semi-parametric local projections method pioneered by Jordà (2005) to estimate state-dependent impulse responses to a temperature shock. As Jordà (2023) notes, the method has found growing appeal outside its home field of applied macroeconomics as its flexibility (Cloyne et al., 2023), relevance to the potential outcomes framework (Dube et al., 2024), and robustness to misspecification and non-stationarity (Montiel Olea and Plagborg-Møller, 2021; Montiel Olea et al., 2024; Piger and Stockwell, 2023) have been established in recent econometric literature. Berg et al. (2024), Bilal and Känzig (2024), and NRK are other climate papers which apply the method to estimate dynamic economic responses to temperature shocks.

These features are particularly useful for addressing several outstanding identification concerns flagged in reviews of impact estimation methods. The remainder of this section uses an application to country-level GDP impacts to aid exposition of the methods used in this study which will also serves to familiarize the reader with features of the aggregate impacts literatures upon which the inequality results build.

We use panel data to estimate dynamic responses to temperature for several distributional outcomes. For now, we take y_{it} to represent the logarithm of GDP for country i in year t. The local projections method entails estimating $H+1$ iterations of the following single-equation model corresponding to different projection horizons $h \in \{0, 1, ..., H\}$:

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

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 shock $\widehat{\tau}$ accounts for potential autocorrelation in the treatment, validating its use as a shock as described in the previous subsection. Interactions of these lagged shocks with the state variable capture potential state-dependencies in these autocorrelations. Lags of the outcome are included as a guard against trend non-stationarity. Finally, the set μ_i and η_t

of standard linear fixed effects are included to absorb time-invariant country-specific effects and common year-specific effects. Besides the construction of the shock and permitting the state variable to vary over time, this model borrows directly from the static state-dependent model of NRK.

Each projection horizon h thus attributes a main effect β_{1h} representing the impact of a unit pulse of temperature on the "long difference" of y projected h periods after the shock. An interaction term allows the effect to scale linearly with the country-year's long-run average temperature $\overline{T}_{it}$. This state-dependency is motivated by the standard result in the climate impacts literature that the economic effect of a temperature shock is increasingly negative for warmer climates so we should expect this coefficient to be negative. Finally, the inclusion of a main effect λ_h associated with the time-varying state variable allows for the possibility that a systematic change in the climate even within the same country may impact expected growth independent of its impact on the distribution of shocks. For example, if a hot country stabilizes at a climate that is 1°C warmer after an intermediate period of warming, it is plausible that economic productivity may be systematically lower even when given sufficient time to maximally adapt. This will not affect the estimation of an impulse response to a shock of temperature but will be relevant for the historical analysis in Section 3.

It is worth taking stock of the restrictions imposed by the model described in (2). The flexibility of this semi-parametric approach is particularly appealing in the context of a literature where the magnitudes of estimated impacts are known to be largely driven by the specification of the damage function (Newell et al., 2021; Chang et al., 2023). For example, two studies projecting the proportional loss in cumulative global production by 2100 arrive at point estimates of -3.4% (Casey et al., 2023) and -20% (Burke et al., 2015) under the same worst-case warming scenario with almost all this divergence stemming from the decision to specify the outcome in levels (e.g., log GDP) or as growth (e.g., first differences in log GDP). Section A2.1 and Figure A5 describe how this "levels vs. growth" choice amounts to assuming outright that temperature has a one-off or permanent effect.

Compromise persistence structures where level effects are temporary but diminishingly persistent can be accommodated by either model by including p additional lags of the explanatory variables. However, as the persistence being considered increases in horizon, the efficiency gains from parametric specifications are rapidly diminished to accommodate increasingly imprecise dynamic effects. The efficacy of lag augmentation is further constrained by non-linear functional forms and autocorrelation in the explanatory variables, both general features of damage functions where absolute temperature is used

as an explanatory variable. Indeed, Burke et al. (2015) conclude that the identifying variation which remains after including lags "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 parametric choices unfalsifiable.

By comparison, the local projections method imposes no restriction on the evolution of a given IRF over time since the effects associated with each projection horizon are estimated from entirely separate regressions and different subsamples[2] In addition, the local projection specified in long differences has been shown to be particularly robust to biases associated with nonstationarity processes Piger and Stockwell (2023); Jordà (2024) such as those which plague standard GDP regressions specified in levels (Nelson and Plosser, 1982; Campbell and Mankiw, 1987). Altogether, the method is relatively robust to different forms of misspecification by remaining mostly agnostic about the underlying data-generating process and as we shall see, can be used to produce a direct statistical test of this influential persistence.

Collecting estimates of the coefficients of interest $\{\widehat{\beta}_{1h}, \widehat{\beta}_{2h}\}_{h=0}^{H}$ from (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 $\overline{T}$:

$$\mathcal{R}_{\tau\to y}(h;\overline{T}) := \widehat{\beta}_{1h} + \widehat{\beta}_{2h}\cdot\overline{T} \tag{3}$$

For example, the top row of Figure 2 depicts three different IRFs which each result from the same regressions but are distinguished by different assumed values of the state variable. These representative states correspond to the 20th and 80th percentiles of country-year temperatures in our sample (10°C colored blue and 26°C colored red) and their average (18°C, roughly the 40th percentile and colored green). 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 so that point estimates are calculated as $\widehat{\beta}_{1h} \pm 8\widehat{\beta}_{2h}$[3]. This coloring scheme and these reference temperatures will be

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

[3]To be specific, we redefine the state variable as $\overline{T}_{it} := \frac{1}{30}\sum_{m=1}^{30}(T_{i,t-m} - 18)$. Under this transformation, the blue and red IRFs correspond to the states $\overline{T}_i - 8 = 10$°C and $\overline{T}_i + 8 =$

used throughout the rest of this paper.

Figure 2: **GDP impulse response to a 1°C shock**

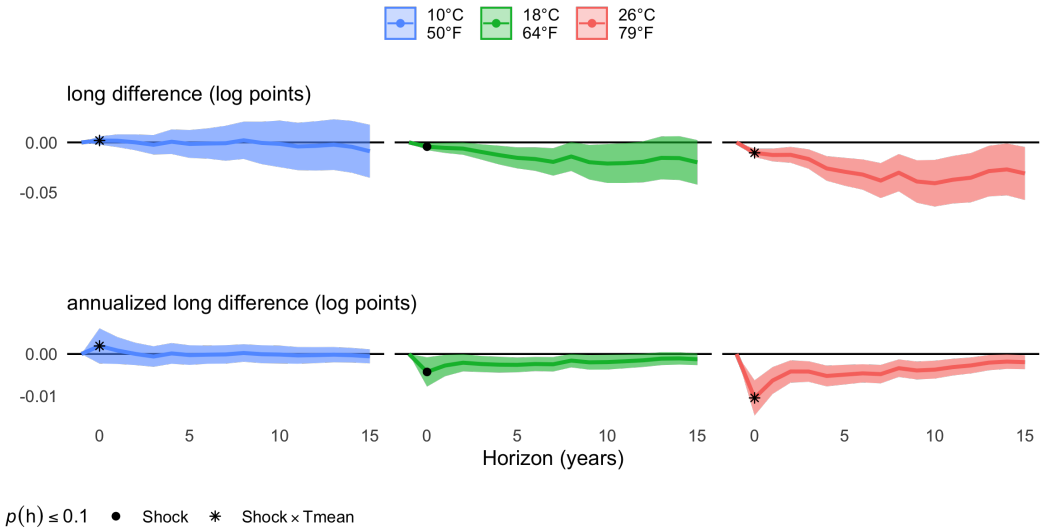

The plots in the top row correspond to estimates of impulse responses of GDP estimated by the $H+1$ local projections regressions specified in (2) for three representative temperatures depicted in different colors. The plots in the bottom row depict the same functions annualized by dividing estimates by $h+1$. All figures in this paper adopt the same color-temperature correspondence where relevant. Points in black correspond to point estimates deemed significant by the joint hypothesis test described in Section 2.2 with filled circles corresponding to significant main effects and asterisks corresponding to significant interaction effects.

Compared to a standard growth regression from which permanent level effects are inferred from observations contemporaneous to the shock, the LP method conservatively estimates persistent effects at each horizon h separately using observations of changes in outcome y over h periods and as depicted by the corresponding IRF $\mathcal{R}_{\tau\to y}(h;\overline{T})$. Leaving aside the question of valid inference for now, interpretation of a resulting IRF proceeds as follows

1. **Level effect**

26°C respectively.

- One-period level effect: $\mathcal{R}_{\tau \to y}(h; \overline{T}) = 0 \ \forall \ h > 0$
- Persistent level effect: $\exists \ h^* > 0$ such that $\mathcal{R}_{\tau \to y}(h; \overline{T}) = 0 \ \forall \ h > h^*$

2. **Growth effect**

- Simple growth effect: $\mathcal{R}_{\tau \to y}(h; \overline{T})$ converges to a non-zero constant
- Divergent growth effect: $\mathcal{R}_{\tau \to y}(h; \overline{T})$ never converges to a constant

The contemporaneous estimates in the IRFs in the first row of Figure 2 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 in NRK wherein per capita GDP appears to stay persistently depressed over the entire projection horizon. The second row depicts the same IRFs 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.

Of course, the standard bias-variance tradeoff applies so that the attractive flexibility features described above come at the cost of efficiency; IRFs estimated by single-equation local projections are fundamentally noisier than their fully parametric and structural counterparts such as vector autoregressions[4].

Importantly, the LP model imposes a linearity assumption on effects estimated within the same horizon. Since the shock enter additively into (2), responses to the shock are assumed to scale proportionally so that a 1°C shock has exactly twice the effect of a 0.5° shock and the exact negative of the effect of a -1°C shock. In reality, non-linearities abound in the interactions between social and climate system (Dietz et al., 2021). Additionally, the linearity assumption imposes the same symmetry in the state interaction term so that relative to the reference temperature, estimated interaction effects for states a given number of degrees warmer than the reference temperature will always be the exact negative of those for states the same number of degrees cooler. In Figure 2, this is reflected in the fact that for a given horizon h, the point estimates in green will always be the exact average of those in blue and red since they are estimated from the same linear regression model.

[4]Different compromises in this tradeoff can be obtained by imposing additional restrictive parametric assumptions such as the functional approximation approach of Barnichon and Matthes (2018) which can substantially narrow the error bands.

It is worth noting that our empirical estimates are expressed in terms of responses to a 1°C shock, which is substantially larger than a shock we would expect to observe in a single year as we showed in Figure 1. The median absolute value of a shock (whether positive or negative) varies by the rate of climate change which will in turn vary across states and time. For the warmest 33% of country-years, 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 and so hot-country effects can be understood as corresponding to a shock 3-5 times larger than one would expect to occur in those climates in a given year. For the coldest 33% of country-years, the median magnitude of a shock was about 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 just 2-3 times larger in magnitude than the median shocks.

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 3 are valuable because they account for four relevant climate inequalities: inequality of temperature effects across an income distribution, inequality of incidence of temperature effects across different climates, inequality in the global distribution of income, and spatial inequality in the incidence of absolute warming.

2.2 Inferring persistence from local projection estimates

The error bands depicted in Figure 2 reflect pointwise estimation imprecision. This estimation imprecision is particularly important to include in this context since the volume of regressions underlying an LP-IRF precludes convenient summarization in a standard regression table.[5] 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 null hypothesis of level effects which persist for less than 15 years.

However, the uncertainty depicted by error bands correspond 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 hypothe-

[5]For example, the results depicted in Figure 2 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 $\alpha = 0.1$.

sis test not captured by pointwise error except at horizon $h = 0$. Focusing on the impulse response corresponding to the main effect, the significance of the main effect of 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 the null hypothesis, Inoue et al. (2024) 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, corresponding to the fact that for each projection horizon, the available estimation sample is reduced by one period (in the case of a perfectly balanced panel).

These significance bands are specific to the main effect and are represented by the gray regions in the middle panels of Figure A6a. 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 A6b depicts a placebo test where I construct bands

to implement significance tests 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[6], 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 2, rounded points correspond to significant main effects while asterisks correspond to significant interaction effects. By this measure, these results cannot reject the null hypothesis of one-period level effects.

2.3 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 7a 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 2.2 implies autocorrelations which

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

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).[7] 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 multiplier $m_{\tau\to y}(h)$ would be calculated as the ratio $\mathcal{R}^c_{\tau\to y}(h)/\mathcal{R}^c_{\tau\to\tau}(h)$. However, this approach has a couple of inconvenient downsides. First, as noted by Jordà (2024), 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 by the algorithm 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 (3) 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 2.2 are equally available. Figure 3b 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 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 and $[-2.7, -1.1]$ percent for 26°C climates.

[7] 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.

Figure 3: **Accounting for persistence of temperature shocks**

(a) Persistence of temperature shocks

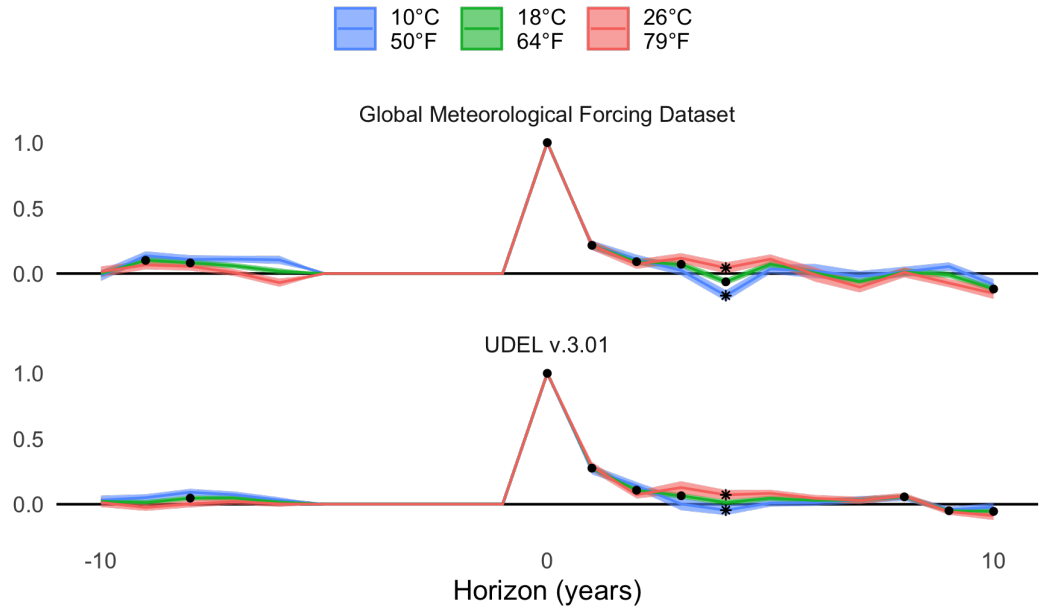

(b) State-dependent temperature multiplier of output

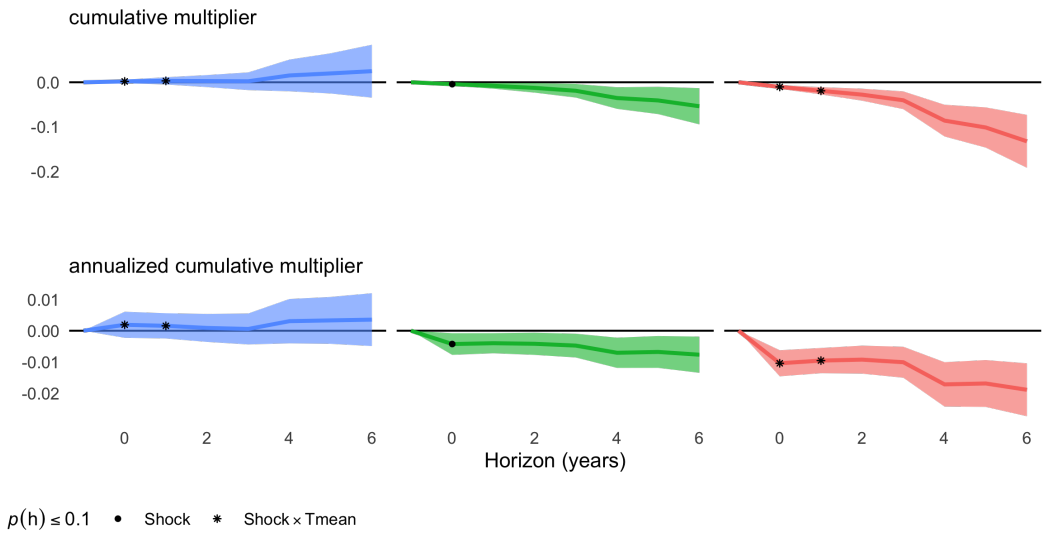

The cumulative multiplier is interpreted as the change in *cumulative* GDP over h periods attributable to a unit shock in $\widehat{\tau}$, net of its impact on future shocks $\widehat{\tau}$. Annualized multipliers represent the same estimates but divided by the projection horizon for visual convenience.

2.4 Data

Having demonstrated how our empirical approach may complement and build on the vast existing literature on the aggregate economic impacts of weather shocks, we now apply the same techniques to a relatively unexplored setting. Providing new evidence on the distributional effects of identified temperature shocks within countries and globally.

Income

This analysis is made possible by grouped income data published by the World Inequality Database (World Inequality Lab, 2024), the successor to the tax data-based World Top Income Database which launched in 2011 with data on over 30 mostly Western countries. A critical difference between the two products was the development of a system of standardizing the integration of tax, national accounts, and survey data (Piketty et al., 2019) modeled after the System of National Accounts (European Commission et al., 2008). This newer generation of inequality data extends coverage within those countries—now claiming to capture "100% of income" in the United States (Piketty et al., 2018)—and rapidly extending coverage and comparability 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 incomes 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. 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 as 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 the most impoverished, it is likely to also include 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 for representing

poverty, we intentionally do not make specific 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).

Climate data

We merge all income data with an annual panel of population-weighted near-surface air temperature data derived from the Global Meteorological Forcing Dataset for Land Surface Modeling (GMFD). This data corrects model-based biases in NCEP/NCAR reanalysis product using observational meteorological station data (Department of Civil and Environmental Engineering/Princeton University, 2006). The GMFD reports historical weather data for the entire global surface at 0.25-degree resolution[8] for every three hours from the start of January 1, 1948 to the end of December 31, 2015.

2.5 Within-country distributional effects

The first set of distributional results we present proceeds similarly to the expositional GDP regressions in Section 2.3 but with separate implementations for each within-country quantile q of interest (e.g., country deciles or percentiles). For every discrete income quantile q, we estimate the $H+1$ state-dependent local-projection regressions for horizons $h \in \{0, 1, ..., H\}$ described by the following single-equation model:

$$\Delta_h^c y_{i,t+h}^q = m_{1,h}^q \widehat{\tau}_{i,t+h}^c + m_{2,h}^q \widehat{\tau}_{i,t+h}^c \overline{T}_{it} + \lambda_h \overline{T}_{it} + \mathbf{Z}_{it}^q \boldsymbol{\gamma}_h^q + \mu_i + \eta_t + \varepsilon_{i,t+h}^q \qquad (4)$$

Here, variables denoted by a superscript c are cumulative variables constructed using the one-step process described in Algorithm A1 to calculate cumulative effects which account for the persistence of temperature shocks. 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 $\Delta y_{i,t-\ell}$. 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.

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

Figure 4: **Cumulative income response to a 1°C shock**

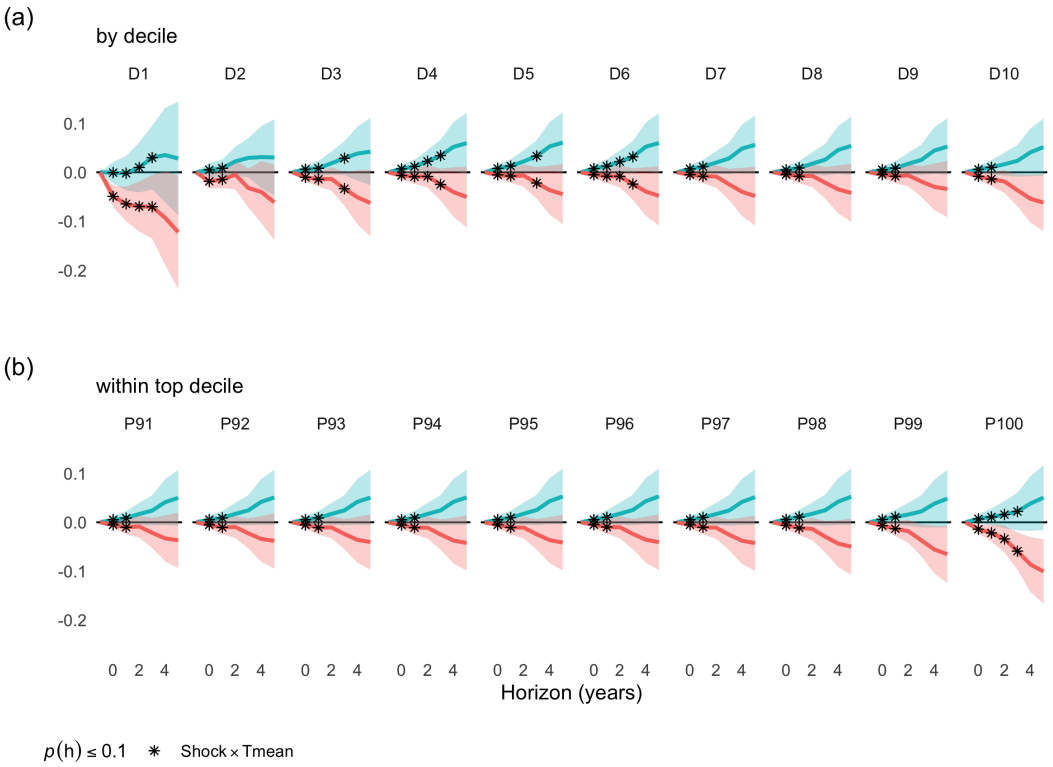

For example, panel (a) of Figure 4 summarizes the results of estimating separate sets of cumulative temperature multiplier functions for each of 10 within-country deciles. Here, color-coding follows the same intuitive scheme as previous plots with blue and red corresponding to representative expected temperatures of 10 and 26°C respectively[9]. 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 Section 2.3. Estimation error is larger than when estimating GDP impacts in Figure 3b, 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 bottom decile compared to $[-3.2, 0.3]$% 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 5a and the bottom panel of Figure 5b 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[10]. To the best of my knowledge, this particular vulnerability of the highest incomes in hot and generally

[9] We omit the green 18°C main effect to avoid over-busy plots, but it is of course implied as the average of the two depicted curves

[10] As mentioned in Section 2.4, 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%.

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 new 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.

While these results strongly imply substantial redistributive effects from a temperature shock, their implied effect on inequality depends both on the initial distribution of income shares across quantiles and the measure of choice, especially given that growth effects are non-monotonic in income quantiles. Figure 6 depicts the annualized effect of a temperature shock on various measures of inequality. Units are not directly comparable across measures but all measures are increasing in (their respective conception of) inequality. Measures are provided in order from the most sensitive to redistribution of top incomes (the Theil index) to the most sensitive to redistribution of bottom incomes (the mean logarithmic deviation index).

The filled circles in the middle column drawn with green error bands correspond to significant main effects of the temperature shock on the given measure of inequality. The asterisks on the plots on either side correspond to significant interactions, measuring linear heterogeneity by expected temperature. In light of the growth incidence results just discussed, the near-zero impacts to cold-country inequality and substantial effects to hot-country inequality are unsurprising. Null effects on the Gini index, known to be relatively insensitive to changes at either end of the income distribution, are consistent with the disproportionate impacts to the bottom 10% and top 1%.

By the Theil index, a transitory temperature shock appears to actually *im-*

Figure 5: **Point estimates of cumulative income growth incidence responses to a 1°C shock**

(a) Incidence patterns for three representative climates

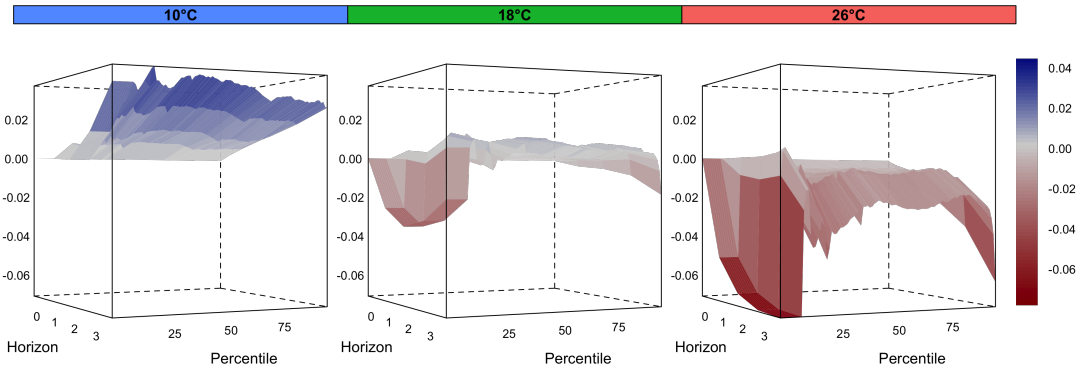

(b) Aggregate vs. distributional responses

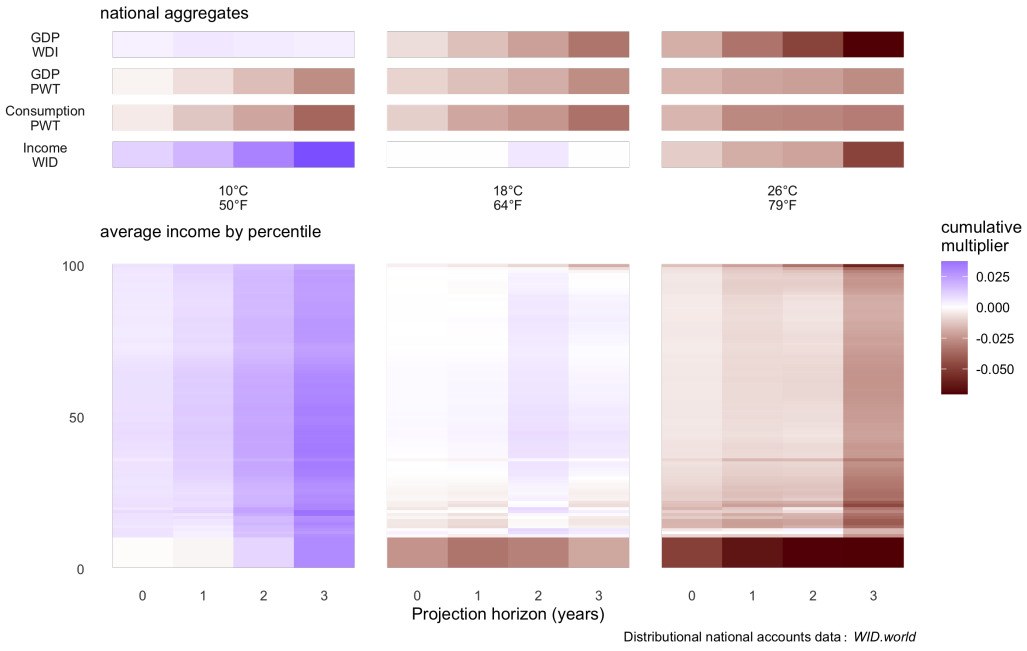

Figure 6: **Inequality index response to a 1°C shock**

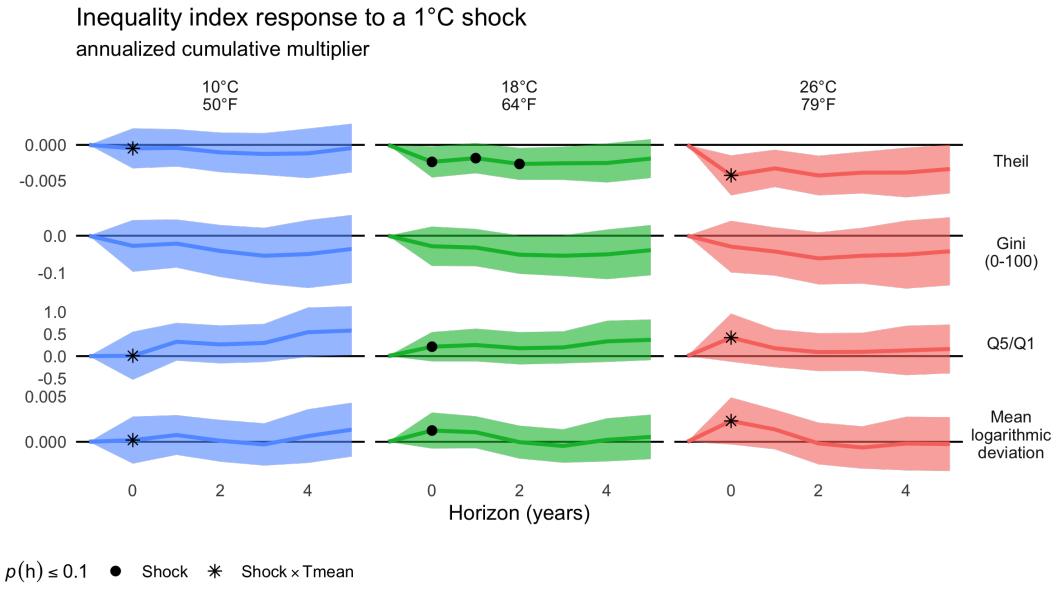

prove inequality while inequality as measured by the ratio of incomes of the top 20% to the bottom 20% and mean logarithmic deviation both suggest the opposite [11]. From the previous analysis, we can surmise that the progressive effect on the Theil index is being driven by the transfer away from top 1% while for the other two measures, the regressive transfer of income shares away from the bottom 10% outweighs this 'improvement' in inequality. While the impacts to the bottom 10% are substantially larger proportionally and represent a population segment 10 times as large, for the Theil measure, the disproportionately large share of income originally held by the top 1% constitutes a large enough share of total national income that these measures disagree on the net effect on inequality[12]. Of course, since the overall effect on income is negative at all percentiles for hot countries, these inconsistencies amount to differently equitable shares of a strictly smaller national income.

3. Climate change and global inequality

Here, we consider the implications of the within-country effects of transitory temperature shocks estimated in Section 2 to estimate the effect of anthropogenic climate change on global inequality. This involves accounting for observed distributions of global income across country-quantiles, the exposure of populations to different climates, and the spatial distribution of global warming. In Section A1, we further decompose these inequality effects by the relative contributions of human activity and natural factors to temperature shocks.

First, for the observed history of country-year weather variables, we calculate the corresponding sequence of shocks and states (1a) and (1b). Then for every within-country quantile from the 11th to the 100th percentile, we collect B bootstrap estimates of the effect of a pulse of temperature on country-quantile incomes estimating (4) through the procedure in Algorithm A1 and

[11] As economic intuition, one may think of the mean logarithmic deviation as a measure of average disutility arising from the given distribution of income compared to a perfectly equal distribution of the same total income, assuming logarithmic utility.

[12] One way to think about the Theil measure is as a fundamentally geometric measure of inequality. For any given income distribution, someone with income \$100,000 giving \$10 to someone with income \$10,000 would improve the Theil index measure of inequality by the exact same amount as someone with a \$1,000 income giving the same \$10 to someone with income \$100 because the ratio of incomes between the donor and recipient are the same (Cowell, 2011). By the same principle, 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.

by the persistence H of this effect using the significance test in Section 2.2. We construct a bootstrap-specific growth differential variable which computes the cumulative income effect of temperature shocks observed contemporaneously and over the preceding H periods. Finally, for every country-quantile-year between 1981 and 2016, we iteratively construct the next year's bootstrap-specific counterfactual income by subtracting the calculated growth differential from the observed level of growth and applying the residual growth rate to the previous year's income.

Upon 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 7a.

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 A1), 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 2.5 is because these groups are dominated by extreme incomes in relatively cold countries. Effects on the global

Figure 7: **Global inequality, observed vs. zero-warming counterfactual**

(a) Global incidence of growth, 1981-2016

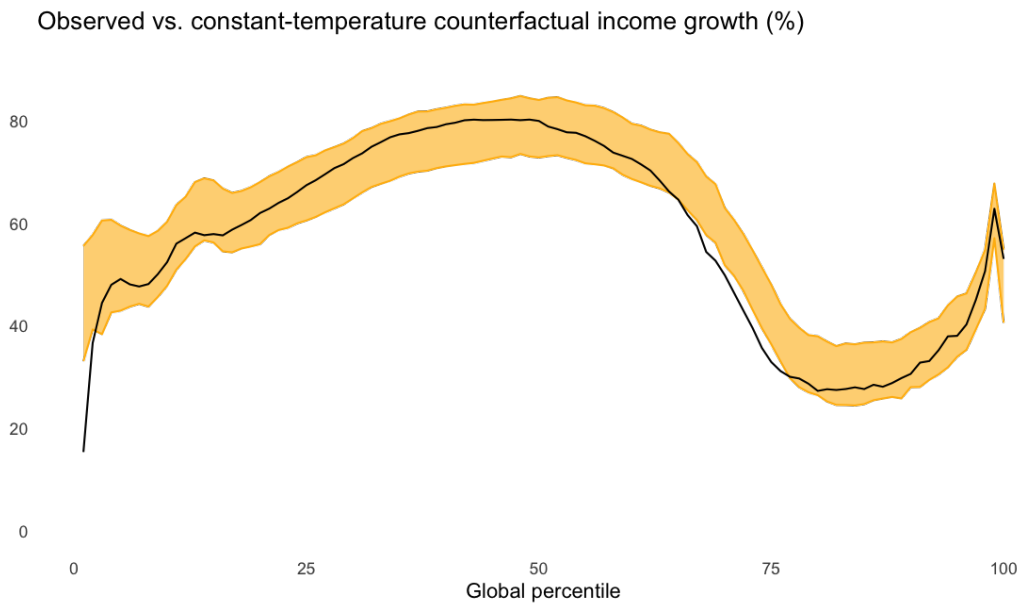

(b) Global inequality indices, 1981-2016

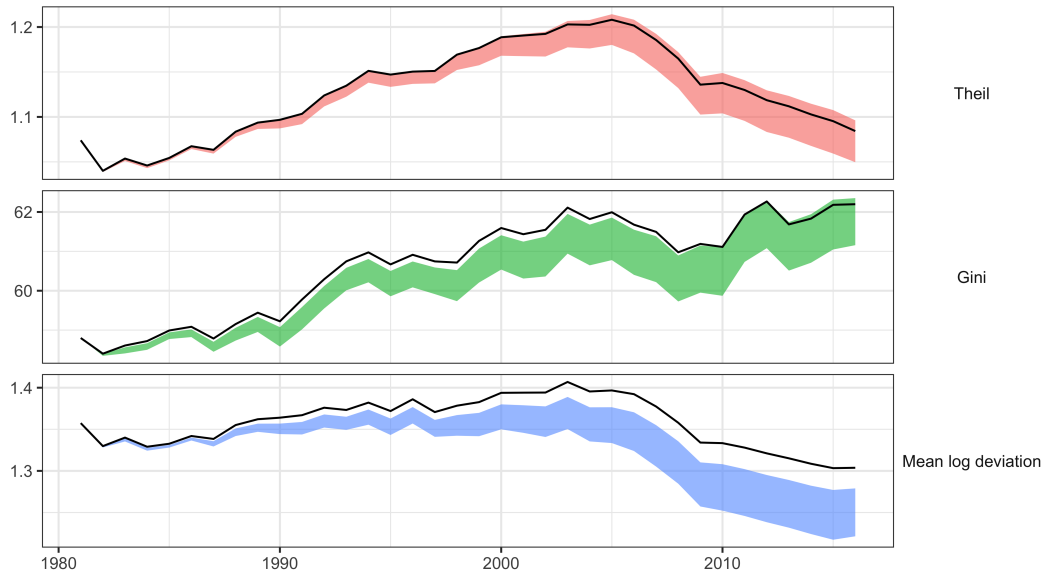

middle class are relatively tempered with slightly positive point estimates estimated for essentially the entire second quartile, also consistent with the within-country incidence curves which found the middle class experienced the mildest economic impacts across all climates.

Despite these different-signed effects, the counterfactual growth incidence band is noticeably flatter than the elephant curve, implying that historical climate change has been globally regressive in addition to being a likely drag on total growth. Figure 7b plots the time series for the global equivalents of the within-country inequality measures introduced earlier. In our within-country analysis, we identified opposing regressive transfers from the bottom decile and progressive transfers from the top decile which opposed one another to produce contradictory effects on within-country inequality depending on the measure of choice. In comparison, the historical effects of these shocks on global inequality are large and regressive across all measures, increasing the global Theil index 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% $[1.6, 6.4]$.

Figure 8 decomposes these effects on global inequality into their within-country and between-country dimensions. Panel (a) compares the observed (black dashed line) and counterfactual (filled colored bands) evolution of these entropy-based decomposable measures in levels. As measured by mean logarithmic deviation, which is relatively sensitive to the welfare of low-income quantiles, observed temperature shocks have unambiguously exacerbated both varieties of inequality decreasing within-country inequality by 2.6% [0.0, 5.6] and between-country inequality by 8.7% [4.9, 13.3]. For the Theil index, the relative sensitivity to negative impacts to top incomes observed in Section 2.5 re-appear, resulting in an estimated 1.3% [-1.3, 3.8] *improvement* in within-country inequality as a result of temperature shocks. These same shocks still worsen between-country inequality by 6.1% [3.6, 9.0].

Panel (b) adopts the same color scheme to visualize their relative contributions to global inequality. The dashed black lines trace the stark observed trend of within-country inequality (teal) increasingly dominating between-country inequality (red) since 1981. The region in white separating the two colors corresponds to the counterfactual shares had average temperatures remained fixed since 1981. That these white regions are entirely under the black dashed lines imply that warming since 1981 has very likely slowed this trend by reducing the within-country share of global inequality by between 1.1 percentage points $[0.1, 2.0]$ by the MLD measure and 1.6 percentage points $[0.7, 2.5]$ by the Theil index. Thus, the incidence of temperature shocks primarily influence the global distribution income by slowing convergence between countries.

Figure 8: **Global inequality between and within countries, 1981-2016**

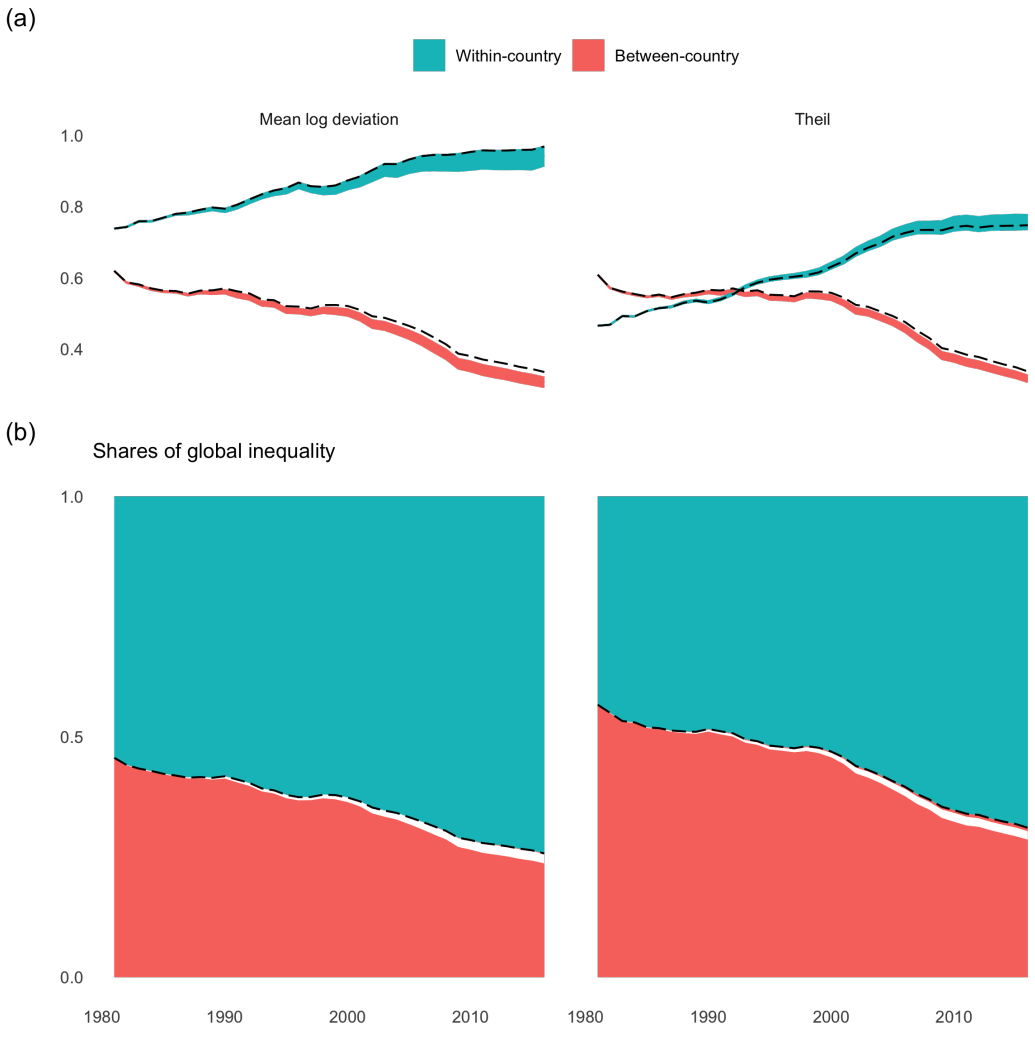

Finally, I note here that these results correspond to impacts of climate change insofar as it is sufficiently well-captured by asymmetries in temperature shocks over time. Section A1 presents additional results decomposing these effects into anthropogenic and natural contributions by applying methods from climate attribution science to output from state-of-the-art climate models, arriving at largely the same conclusions.

4. Conclusion

The results presented in this paper constitute the first comprehensive evidence of regressivity in the global incidence of climate change, complementing an expansive and active economic literature which has remained effectively agnostic on global distributional impacts at subnational levels. By importing methods recently developed in empirical macroeconomics to distributional data thus far largely unused in environmental economics, we were able to overcome several long-standing statistical and practical constraints which have long obstructed progress on this urgent research question.

Given the sparsity of other work on this broad topic, several limitations of the present study suggeset avenues for potentially fruitful future research. Firstly, even the impressive coverage of the World Inequality Database is fundamentally limited in its ability to speak to the most important of distributional considerations, namely the impacts of climate change on global poverty. Our evidence affirms and quantifies the conventional understanding that the world's poorest have been and will continue to be by far the most severely impacted by climate change. But due to the limitations of pretax income as a measure of poverty and welfare, even these impacts are likely to be a conservative lower bound on true impacts.

In revisions in progress, the distributional analysis reported here will be supplemented by an analysis of compositional effects identifying which income groups from which countries have been transitioning up and down the global income distribution as a result of climate change. If income effects are as severe and state-dependent as described here, then it becomes plausible that an important distributional consequence of climate change is that inhospitable environments becomes increasingly deterministic of economic development, severely limiting the upward mobility of populations in hotter climates.

To the best of my knowledge, the vulnerability of the highest incomes in hot and generally poor countries to transitory environmental shocks has not been previously documented or theorized. Weather-induced capital depre-

ciation, capital shares in agricultural income, and the relative unavailability of insurance instruments such as weather derivatives in developing countries are example plausible mechanisms that call for empirical testing.

From a policy perspective, the results presented in this paper suggest that incorporating equity considerations as through a concave utility function would indeed increase the social weighting of the most impacted by climate change. But it is may still be worth taking seriously criticisms that a "profound indifference to inequality" (Sen et al., 2020) remains embedded even in these progressive varieties of cost-benefit analyses predominant in the climate economic literature. For example, Prest et al. (2024) show that one proposed implementation would increase the US government's official estimate of the social cost of carbon by a factor of eight. But somewhat paradoxically, these greater cost estimates are primarily used to prescribe an increased price on carbon which may then disproportionately burden these same segments of the population whose increased welfare weight induced the inequality-exacerbating policy. Equity weighting should of course factor into both the estimation of costs and the design of policy internalizing them.

To this end, it is intuitive that democratic pressures and a richer portfolio of policy instruments should be more effectively mobilized domestically than coordinated across a Westphalian international system (Barrett, 2003; Levin et al., 2009; Keohane and Victor, 2011). If so, levers available domestically but not internationally—such as compensatory fiscal transfers—may yet emerge as vital instruments of climate policy.

References

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, "The Elephant Curve of Global Inequality and Growth," *American Economic Review: Papers and Proceedings*, 2018, *108*, 103–108.

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

Barnichon, Regis and Christian Matthes, "Functional Approximation of Impulse Responses," *Journal of Monetary Economics*, 2018, *99.*

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

Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro, "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," *Journal of Political Economy*, 2016, *124* (1), 105–159.

Barrett, Scott, *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford, UK: Oxford University Press, 2003.

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

Bilal, Adrien and Diego R. Känzig, "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature," Working paper, National Bureau of Economic Research 2024.

Burke, Marshall, Mustafa Zahid, Mariana C. M. Martins, Christopher W. Callahan, Richard Lee, Tumenkhusel Avirmed, Sam Heft-Neal, Mathew Kiang, Solomon M. Hsiang, and David Lobell, "Are We Adapting to Climate Change?," Working paper 32985, National Bureau of Economic Research 2024.

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

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

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

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E. Kopp, Kelly E. McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang, "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," *Quarterly Journal of Economics*, 2022.

Casey, Gregory, Stephie Fried, and Ethan Goode, "Projecting the Impact of Rising Temperatures: The Role of Macroeconomic Dynamics," *IMF Economic Review*, 2023, *71*.

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

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

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

_ , Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, *World Inequality Report 2022*, Paris: World Inequality Lab, 2022.

Chang, Jun-Jie, Zhifu Mi, and Yi-Ming Wei, "Temperature and GDP: A review of climate econometrics analysis," *Structural Change and Economic Dynamics*, 2023.

Cloyne, James, Òscar Jordà, and Alan M. Taylor, "State-Dependent Local Projections: Understanding Impulse Response Heterogeneity," Working paper, National Bureau of Economic Research 2023.

Cowell, Frank, *Measuring Inequality*, Oxford University Press, 2011.

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

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

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

Department of Civil and Environmental Engineering/Princeton University, "Global Meteorological Forcing Dataset for Land Surface Modeling," 2006.

Dietz, Simon, James Rising, Thomas Stoerk, and Gernot Wagner, "Economic impacts of tipping points in the climate system," *PNAS*, 2021, *118*.

Diffenbaugh, Noah S. and Marshall Burke, "Global warming has increased global economic inequality," *PNAS*, 2019.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor, "A Local Projections Approach to Difference-in-Differences Event Studies," Working Paper, National Bureau of Economic Research 2024.

European Commission, International Monetary Fund, Organisation for Economic Co-operation and Development, United Nations, and World Bank, *System of National Accounts 2008*, 1st ed., New York: United Nations, 2008. Accessed: 2024-10-12.

Eyring, V., S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor, "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization," *Geoscientific Model Development*, 2016, *9* (5), 1937–1958.

Gilli, Martino, Matteo Calcaterra, Johannes Emmerling, and Francesco Granella, "Climate change impacts on the within-country income distributions," *Journal of Environmental Economics and Management*, 2024, *127*.

Goodman, Joshua, Michael Hurwitz, Jisung Park, and Jonathan Smith, "Heat and Learning," *American Economic Journal: Economic Policy*, 2020, *12*.

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

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

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

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D.J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser, "Estimating economic damage from climate change in the United States," *Science*, 2017, *356*, 1362–1369.

Inoue, Atsushi, Òscar Jordà, and Guido M. Kuersteiner, "Inference for Local Projections," Working paper, Federal Reserve Bank of San Francisco 2024.

Jenkins, Stephen P., "World income inequality databases: an assessment of WIID and SWIID," *Journal of Economic Inequality*, 2015, *13*.

Jones, Chad, "Pareto and Piketty: The Macroeconomics of Top Income and Wealth Inequality," *Journal of Economic Perspectives*, 2015, *29* (1).

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95.*

__ , "Local projections for applied economics," *Annual Review of Economics*, 2023.

__ , "Local Projections," Working paper, National Bureau of Economic Research 2024.

Jorgenson, Dale W., "Production and Welfare: Progress in Economic Measurement," *Journal of Economic Literature*, 2018, *56.*

Kahn, Matthew E., Kamiar Mohaddes, Ryan N. C. Ng, M. Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang, "Long-Term macroeconomic effects of climate change: A cross-country analysis," *Energy Economics*, 2021, *104.*

Keohane, Robert O. and David G. Victor, "The Regime Complex for Climate Change," *Perspectives on Politics*, 2011, *9* (1), 7–23.

Knutti, R., G. Abramowitz, M. Collins, V. Eyring, P. J. Gleckler, B. Hewitson, and L. O. Mearns, "Good practice guidance paper on assessing and combining multi-model climate projections," in "Meeting Report of the Intergovernmental Panel on Climate Change Expert Meeting on Assessing and Combining Multi-Model Climate Projections," Geneva, Switzerland: IPCC, 2010, pp. 1–13.

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

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

Levin, Simon A., Brian Walker, Scott Barrett, Stephen Polasky, Victor Galaz, Carl Folke, Gustav Engström, Frank Ackerman, Ken Arrow, Stephen Carpenter et al., "Looming Global-Scale Failures and Missing Institutions," *Science*, 2009, *325* (5946), 1345–1346.

Liu, Q., X. Zhang, and P. Zhai, "Evaluation of CMIP6 precipitation simulations across different climatic zones: Uncertainty and model intercomparison," *Advances in Atmospheric Sciences*, 2021, *38* (7), 1180–1193.

Marx, Nicolas Longuet, "Investigating the Impact of Temperature Shocks on Income Disparity," Working paper 32 2024.

Masson-Delmotte, V., P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. R. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. B. R. Matthews, Y. Chen, X. Zhou, M. I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield, eds, *Global warming of 1.5°C: An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty* 2018.

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

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow, "How Much Will Global Warming Cool Global Growth?," Working paper, National Bureau of Economic Research 2024.

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

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

Office of Management and Budget, "Circular No. A-4: Regulatory Analysis," 2023.

Olea, José Luis Montiel and Mikkel Plagborg-Møller, "Local projection inference is simpler and more robust than you think," *Econometrica*, 2021.

_ , _ , Eric Qian, and Christian K. Wolf, "Double Robustness of Local Projections and Some Unpleasant VARithmetic," 2024. Working Paper.

Palagi, Elisa, Matteo Coronese, Francesco Lamperti, and Andrea Roventini, "Climate change and the nonlinear impact of precipitation anomalies on income inequality," *PNAS*, 2022, *119.*

Piger, Jeremy and Thomas Stockwell, "Differences from Differencing: Should Local Projections with Observed Shocks be Estimated in Levels or Differences?," 2023. Working paper.

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

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

_ , _ , and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States," *Quarterly Journal of Economics*, 2018, *133.*

_ , _ , and _ , "Simplified Distributional National Accounts," *AEA Papers and Proceedings*, 2019, *109*.

Pindyck, Robert S., "Climate Change Policy: What Do the Models Tell Us?," *Journal of Economic Literature*, 2013, *51* (3), 860–872.

Prest, Brian C., Lisa Rennels, Frank Errickson, and David Anthoff, "Equity weighting increases the social cost of carbon," *Science*, 2024, *385*.

Ramey, Valerie A., "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2016.

_ and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 2018.

Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni, "Country-level social cost of carbon," *Nature Climate Change*, 2018, *8*, 895–900.

Rode, Ashwin, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E. Kopp, Kelly E. McCusker, Ishan Nath, James Rising, and Jiacan Yuan, "Estimating a Social Cost of Carbon for Global Energy Consumption," *Nature*, 2021.

Roine, Jesper, Jonas Vlachos, and Daniel Waldenström, "The long-run determinants of inequality: What can we learn from top income data?," *Journal of Public Economics*, 2009, *93* (7-8), 974–988.

Schlenker, Wolfram and Michael J. Roberts, "Nonlinear temperature effects indicate severe damages to US crop yields under climate change," *PNAS*, 2009, *106* (37), 15594–15598.

Sen, Amartya, Angus Deaton, and Tim Besley, "Economics with a Moral Compass? Welfare Economics: Past, Present, and Future," *Annual Review of Economics*, 2020, *12*.

Tebaldi, C. and R. Knutti, "The use of the multi-model ensemble in probabilistic climate projections," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2007, *365* (1857), 2053–2075.

United Nations, "Transforming our world: the 2030 Agenda for Sustainable Development," 2015. Resolution adopted by the General Assembly on 25 September 2015, A/RES/70/1.

van Vuuren, D. P., B. Eickhout, P. L. Lucas, and M. G. J. den Elzen, "Long-term multi-model scenarios for global change research: Integrated assessment model comparisons at the crossroads," *Climatic Change*, 2006, *78* (1), 1–21.

World Inequality Lab, "Distributional National Accounts Guidelines: Methods and Concepts used in the World Inequality Database," 2024.

Zelinka, M. D., T. A. Myers, D. A. T. Fisher, K. E. Taylor, T. Andrews, M. V. D. Jackson, and C. Proistosescu, "Causes of Higher Climate Sensitivity in CMIP6 Models," *Geophysical Research Letters*, 2020, *47* (1), e2019GL085782.

Zhou, T., X. Chen, L. Chen, W. Wu, and Z. Lin, "Global surface air temperatures in CMIP6: historical performance and future changes," *Journal of Climate*, 2020, *33* (21), 9731–9743.

A1. Anthropogenic attribution

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

A1.1 Climate simulation data

Climate simulation data come 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[13] 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.

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

Figure A1: **Model-implied spatial distributions of global warming, 1980-2014**

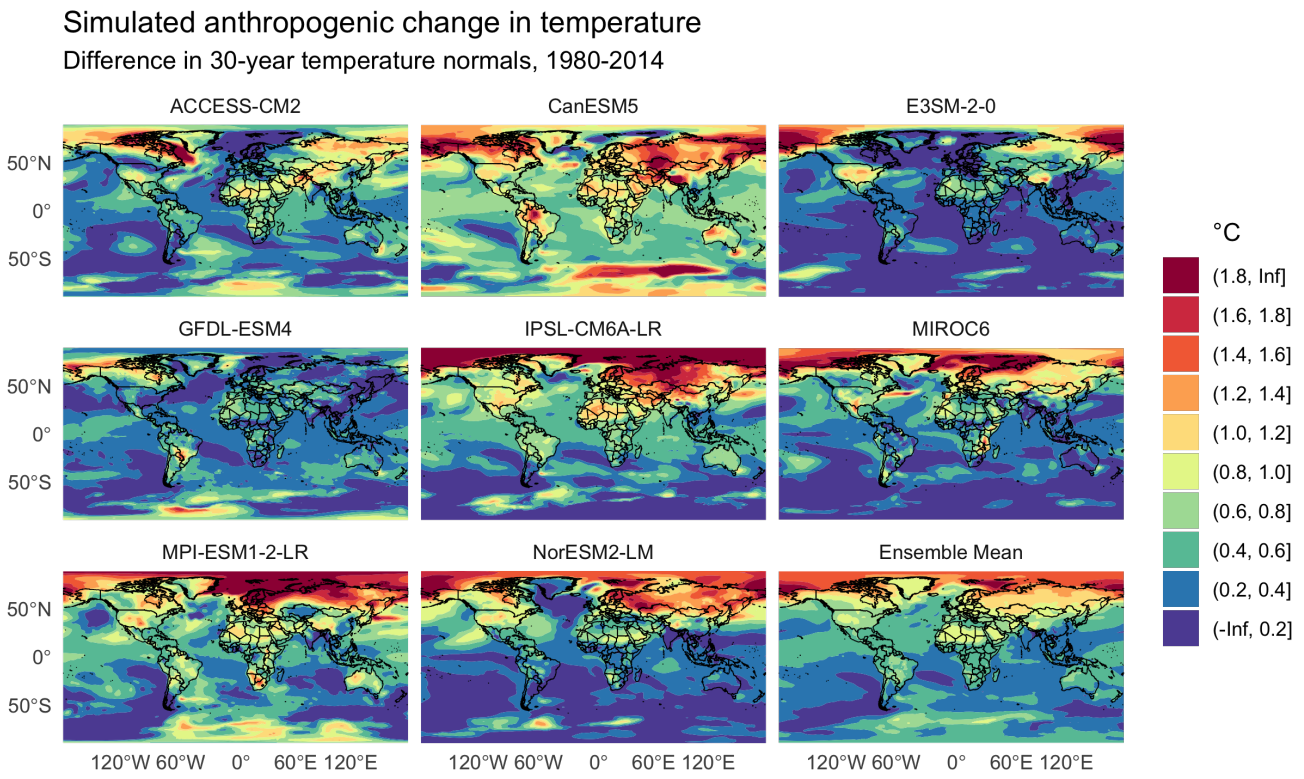

Simulated spatial distribution of anthropogenic warming for eight randomly selected climate models. The bottom-right plots the average distribution for all 13 CMIP6 models which ran both historical and historical-natural simulations.

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.[14] 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.

A1.2 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 (Zhou et al., 2020; Zelinka et al., 2020; Liu 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; van Vuuren et al., 2006).

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 mea-

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

suring 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 A2 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 A2: **Constructing historical counterfactuals holding anthropogenic forcings fixed at 1980 levels**

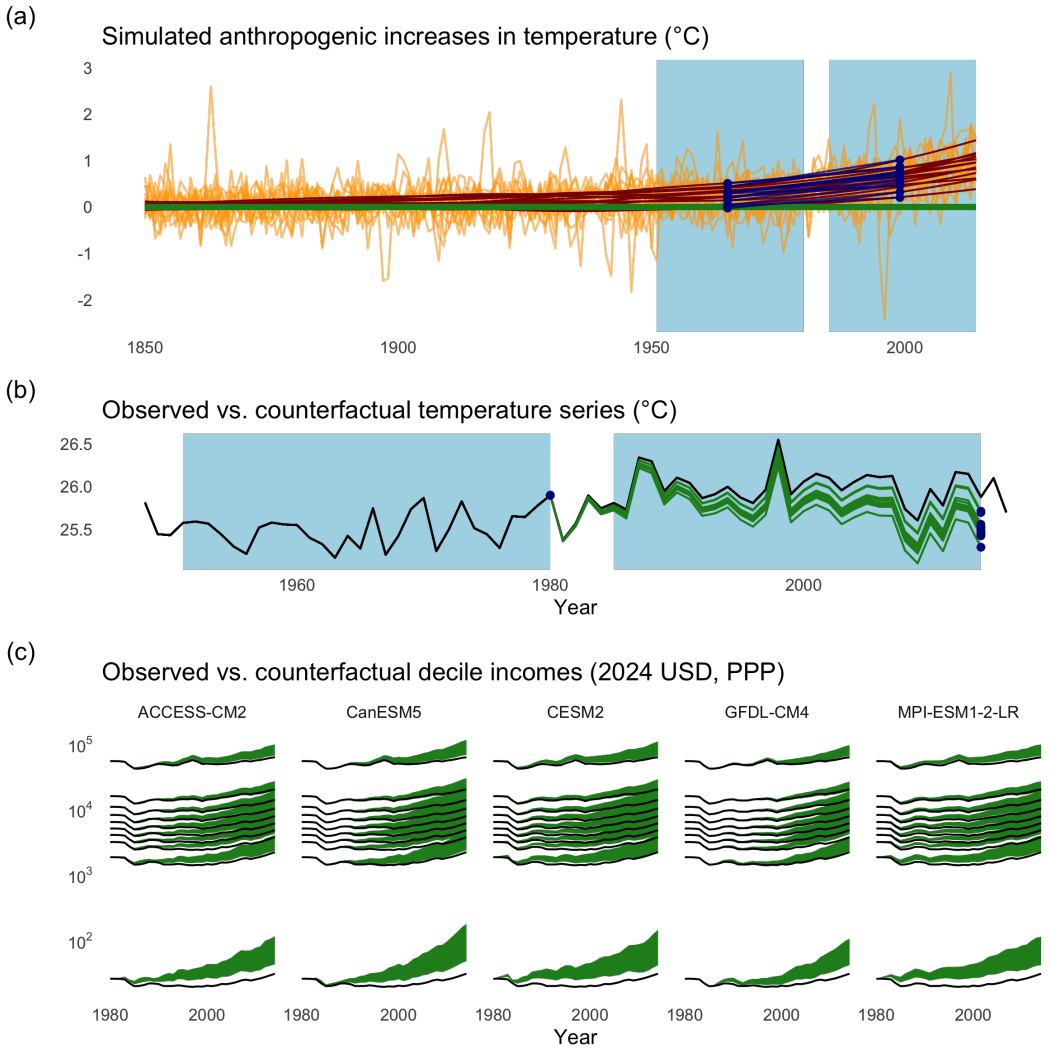

Demonstration of the delta method for constructing bias-corrected counterfactuals using data for the Philippines as an example.

Figure A3: **Anthropogenic growth incidence curves by climate model, 1980-2014**

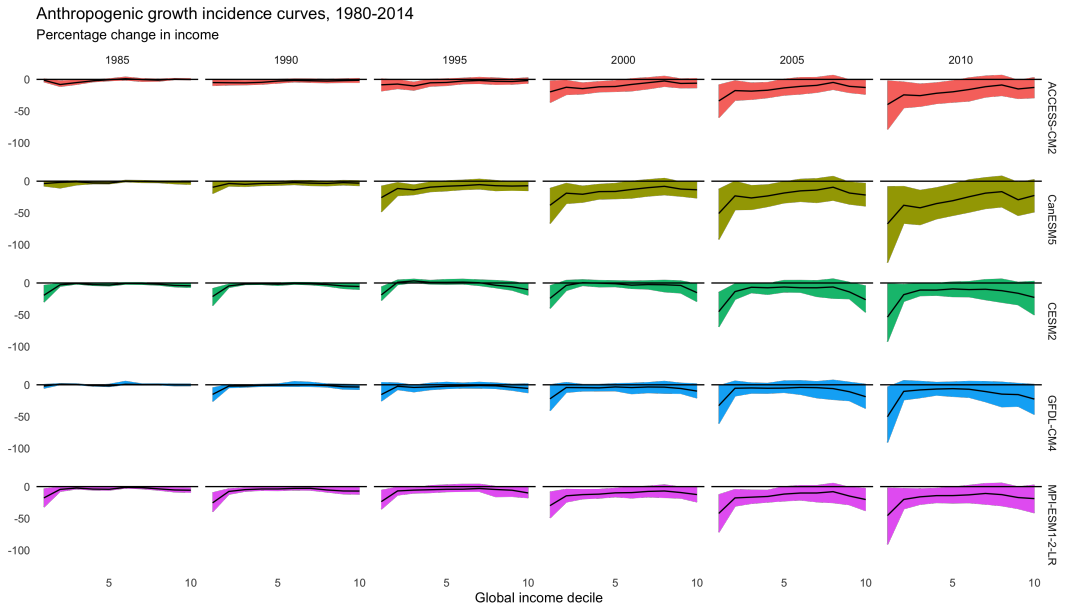

Figure A4: **Anthropogenic influence on within- vs. between-country inequality by climate model**

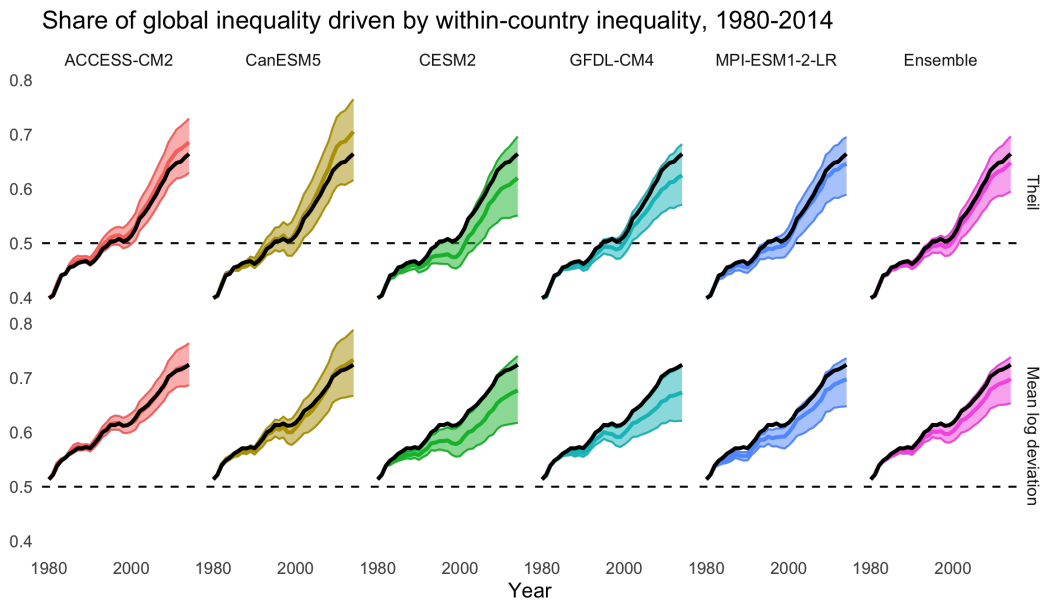

A2. Empirical methods

A2.1 Growth vs. level effects

Figure A5: **Levels and growth under different persistence structures**

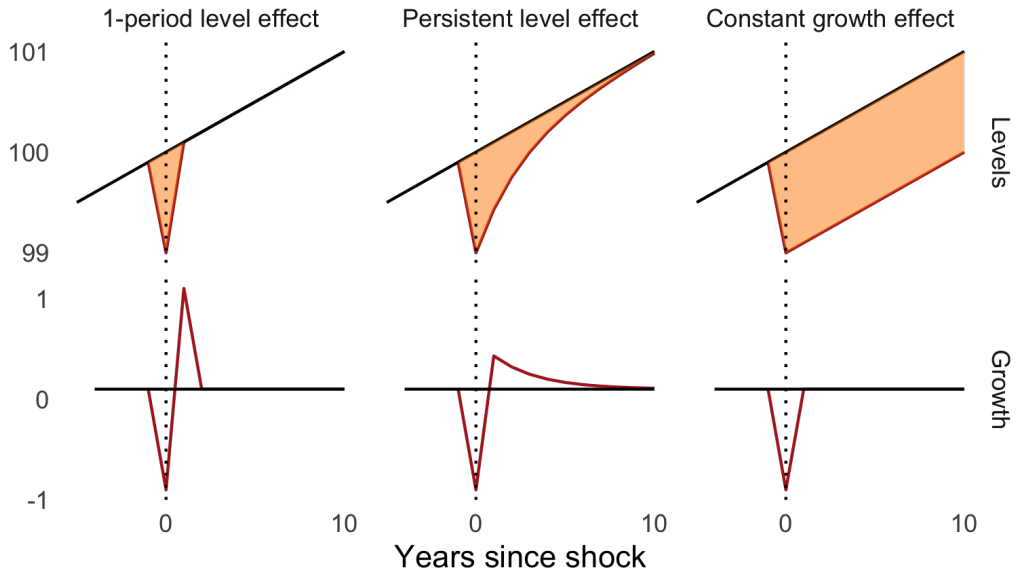

Each column corresponds to a different possible dynamic path of a temperature shock in period 0 on an outcome measured in levels (top panels) and annual growth (bottom). Black and red lines depict the evolution of the outcome variables in the absence and presence of a unit shock in period 0 respectively under the different described dynamics. Filled areas in orange measure the cumulative difference in the outcome attributable to the shock. In the context of climate change, these represent the long-term damages associated with a temperature shock.

Under a one-period level effect, the level of, say, GDP falls by an estimated level $\widehat{\beta}_0$ iin the year of the shock but fully recovers to trend in the next period. All lost production is compensated for in the next period by virtue of an above-trend rate of growth which exactly offsets the negative growth effect in the year of the shock. The shock has no effect on future production and the area in orange represents a one-off economic loss 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 trajectories thus compound over time as observed by the orange area which grows linearly over time. The growth effect depicted in the bottom panel is entirely transitory.

The middle column describes the 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.

$$\log Y_t = \beta^L X_t + v_t \tag{5a}$$

$$\Delta Y_t = \beta^G X_t + u_t \tag{5b}$$

In the levels model, β^L represents the difference between the observed level of GDP and the GDP that would have occurred had a unit shock in period 0 not occurred. Graphically, it is the vertical distance between the black and red lines in the levels panels. Since this is an entirely static model, there is no mechansim for X to affect future values of Y and so the model assumes a perfect return to trend in all ensuing periods. In the growth plots, this is represented by a $\beta^L - 0 = \beta^L$ effect in period 0 and an exactly offsetting $0 - \beta^L = -\beta^L$ effect on growth in period 1 and no effect on growth in all ensuing periods. Level effects which persist for p periods can be accommodated by including p additional lags of X in the model. 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_t; the gap between the black pre-shock trend line and the observed GDP series stays fixed 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

A2.2 Testing for persistence using significance bands

Figure A6: **Impulse responses to a 1°C shock, including significance bands**

(a) GDP impulse response

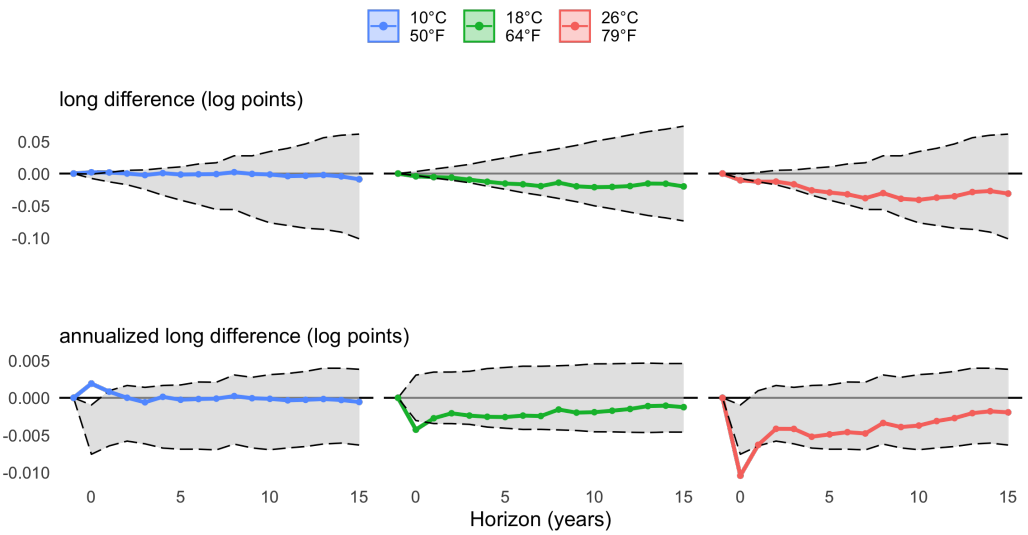

(b) Placebo test: national population response

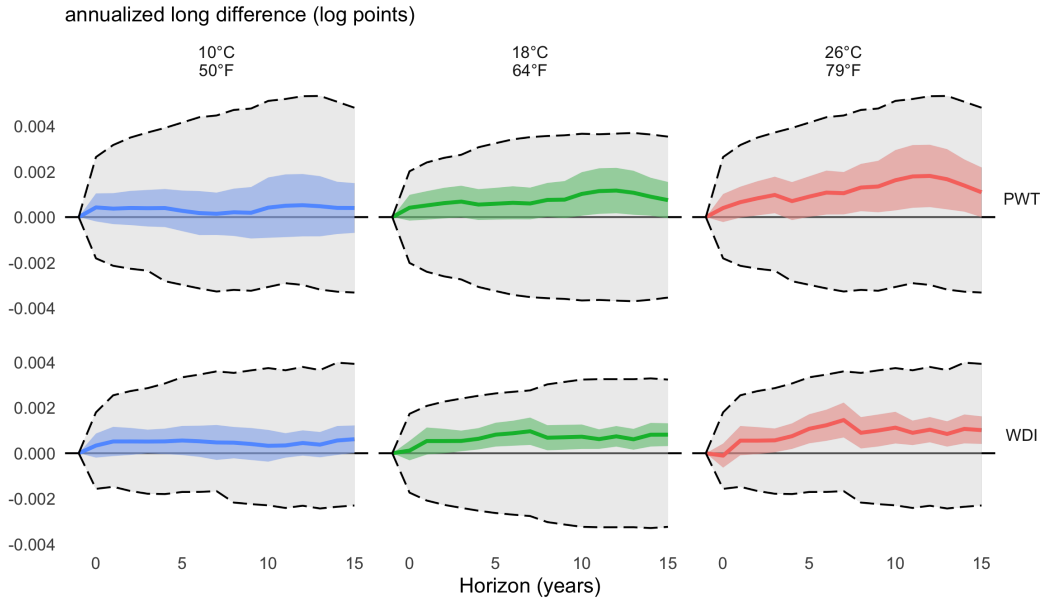

A2.3 Construction of inequality indices

Inequality data from the World Inequality Database is provided 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 impose assumptions about the distribution of household incomes within these groups. This becomes particularly influential for the construction of entropy-based inequality measures such as the mean logarithmic deviation or the Theil index.

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[15].

Figure A7: **Sensitivity of inequality index effects to assumptions about within-percentile inequality**

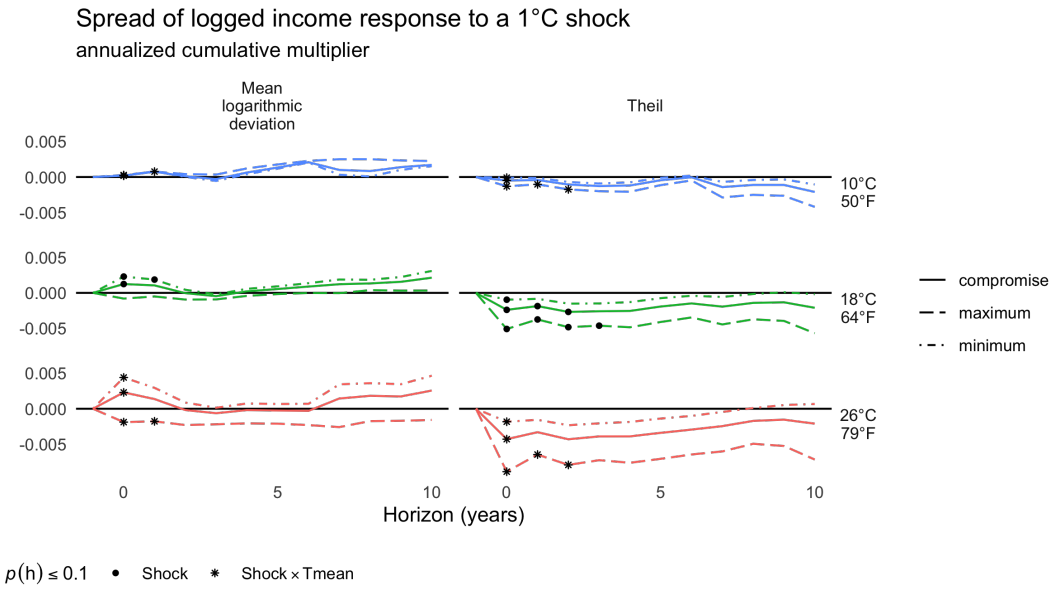

Cowell (2011) describes multiple 'compromise' assumptions between the

[15]We omit the formula but within-group inequality is not generally maximized by an equal split between the two

two extremes, finding that they each approximate an average where the lower bound is given twice the weight of the upper bound. Figure A7 depicts IRFs of the two inequality measures under each of these three assumptions.

A2.4 Data processing

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

Estimate local projection:

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

end

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

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 (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 (1a) and (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 (4);

    Specify maximum projection horizon H using persistence test from Section 2.2;

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

A3. Supplementary results

Figure A8: **Temperature multipliers estimated using survey microdata**

(a) Persistent impacts on income

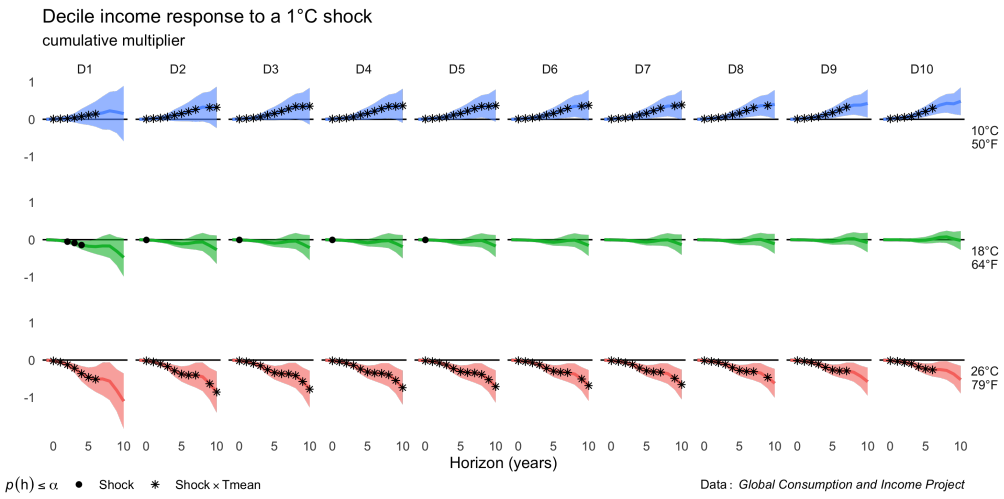

(b) Persistent impacts on consumption

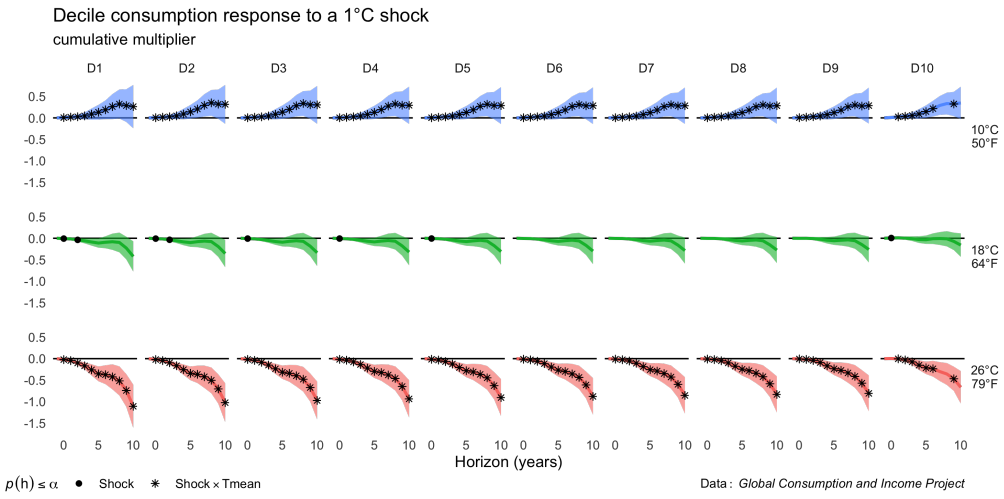

Figure A9: **Population-weighted average temperatures by global percentile**