nature portfolio

https://doi.org/10.1038/s41586-026-10272-6

Supplementary information

Quantifying climate loss and damage consistent with a social cost of carbon

In the format provided by the authors and unedited

Supplementary Information

Quantifying climate loss and damage consistent with a social cost of carbon

Marshall Burke, Mustafa Zahid, Noah Diffenbaugh, Solomon Hsiang

Supplemental Methods

S1 Framework for estimating loss and damage

Our approach to computing L&D treats CO_2 emissions as generating a flow of future damages for other entities around the world (Figure 1a). Cumulative damage from these emissions is calculated relative to four key time points: the time at which "responsibility" for past emissions begins, which we denote t_0, the time each unit of emissions occurs t_e, the year(s) in which damage occurs t, and the time "settlement" occurs for damages from these emissions t_s. We define t_s as the time when a transfer occurs such that a population i that is receiving damages is given a transfer equal to the cumulative damages (past or future) resulting from the emission at t_e. Settlement year could occur after emissions year, as for loss and damage that has already occurred due to past emissions, or could occur contemporaneously or preceding emissions and damages, as in SC-CO_2 (see Fig 1).

S1.1 Damage from marginal emissions

Let Y_{it} represent income for population i in year t which depends on the climate realization C_{it} in that year, which in turn depends the entire history of anthropogenic CO_2 emissions $\mathbf{E}$. We wish to compute the change of income ΔY_{it} that would result (at an arbitrary "damage year" t) from perturbing this history of emissions to some alternate trajectory $\mathbf{E}'$. Here, the "damage year" t could be any year including and subsequent to the emissions year t_e; a unit of emitted carbon dioxide – and the associated heat added to the climate system – can remain in the atmosphere for centuries[1,2], causing warming and subsequent damage in many future years.

We initially consider a marginal (one year) perturbation at a specific time $t_e \leq t$ in the past, which we denote ΔE_{t_e}. For this marginal emissions case, $\mathbf{E}'$ and $\mathbf{E}$ only differ in year t_e; in other settings of interest, for instance eliminating an entire country's recent history of emissions, trajectories will differ in multiple years. We assume here that an additional ton of emissions is being added to a baseline emissions trajectory, and the goal is to estimate the damages from this additional ton; an approach that instead subtracted a ton of emissions and calculated the avoided damages from this reduction would yield an equivalent estimate to our approach but of opposite sign.

Because CO_2 is well-mixed in the atmosphere, we assume that impacts of this marginal unit of emission do not depend on who emitted it. Damage at a specific moment in time (a flow) from the marginal emissions perturbation ΔE_{t_e} is then the difference between the income expected under the perturbed emissions relative to "baseline" emissions history, which for past years is just the observed emissions and for future years is emissions as projected by the socioeconomic scenarios described below:

$$\Delta Y_{it}(\Delta E_{t_e}) = Y_{it}(C_{it}(\mathbf{E})) - Y_{it}(C_{it}(\mathbf{E}')) \tag{1}$$

This change could be either positive or negative depending on the time scale and population exposed. However, as we will show, if perturbed emissions are higher than baseline emissions, $\Delta Y_{it}(\Delta E_{t_e})$ is typically positive – that is, in the exposed population, income is higher in the baseline setting with lower emissions – and so we use "damage" as shorthand for the change in income.

S1.2 Cumulative damage from marginal emissions

We introduce two more time-points, t_1 and t_2, which represent the beginning and end years over which damages are cumulated. The cumulative damage to population i from a unit of emission in year t_e is the sum of damages experienced in that country between t_1 and t_2, discounted to their value at an arbitrary time of settlement t_s.

$$D_{i,t_e,t_s,t_1,t_2}(\Delta E_{t_e}) = \sum_{t=t_1}^{t_2} (1+r)^{-(t-t_s)} \cdot \Delta Y_{it}(\Delta E_{t_e}) \tag{2}$$

Discounting is done as a function of the difference between the year of damage t and the settlement year. Damage is also potentially a function of the stock of emissions, if a climate variable is non-linear in emissions or damages are non-linear in the climate variable. We do not track this stock notationally, but track it in our application.

Equation 2 can be used to decompose damages from historical and future emissions into three additively separable components of damage, which is helpful both for mapping different popular conceptualizations of "loss and damage" into an aggregate measure of total L&D, and for distinguishing loss and damage that has already occurred from that which is likely to occur in the future, as approaches to liability and compensation could differ across these damage types.

First is the historical damage that has already occurred due to a past marginal emission cumulated through the present day, which we term *HD-CO*$_2$. In this setting, the settlement time t_s equal to the present day (denoted t_p), damages begin cumulating in the emissions year ($t_1 = t_e$) and end in present day ($t_2 = t_p$), emissions are in the past ($t_e < t_p$), and damages are aggregated over populations globally:

$$HD\text{-}CO2_{t_e} = \sum_i \sum_{t=t_e}^{t_p} (1+r)^{-(t-t_p)} \cdot \Delta Y_{it}(\Delta E_{t_e \leq t_p}) \tag{3}$$

This is the quantity depicted in Fig 1e, aggregated over populations. Because the damage year is prior to the year of settlement, the exponent on the discount rate is positive in Eq 2, which means a higher discount rate will amplify the value of past damages and raise estimates of HD-CO_2, as described above.

A past emission will continue to cause damage in future years. This occurs for two reasons. First, carbon dioxide emissions remain in the atmosphere for decades or centuries and will continue to cause warming, and thus damage, unless removed. Second, and somewhat more subtly, past damages drive a wedge between the observed size of the economy and how large the economy would have been without the warming; damage from future warming further expands the size of this existing wedge. Together, these effects create a stream of expected future damages from a past unit of emission, which we term *FD-CO*$_2$, or the future damage from historical emissions. This is the quantity depicted in Fig 1f, aggregated over populations. Calculation is as for HD-CO_2, but impacts of a historical emission are cumulated beginning present day ($t_1 = t_p$) and continued until some distant future year, which in practice is typically 2100 or 2300:

$$FD\text{-}CO2_{t_e} = \sum_i \sum_{t=t_p}^{\infty} (1+r)^{-(t-t_p)} \cdot \Delta Y_{it}(\Delta E_{t_e \leq t_p}) \tag{4}$$

Here discounting works in the typical way, with higher discount rates reducing the value of cumulative damages and lowering FD-CO_2. Empirically, as we show, FD-CO_2 damages are very large relative to HD-CO_2, given the long timescales over which greenhouse gases and associated heat remain in the climate system. The implication is that even if a large historical emitter were to rapidly achieve net-zero, its past emissions would continue to cause substantial uncompensated damage.

The final component is the cumulative future damages expected to occur for a marginal emission in the present year or in future years, discounted to present day. Emissions are now or in the future ($t_e \geq t_p$), damages start cumulating in the year of emission ($t_1 = t_e$) and end in some distant future year as above, and settlement occurs at the time of emission ($t_s = t_e$), which in our example calculations is always assumed to be the present day:

$$SC\text{-}CO2_{t_e} = \sum_i \sum_{t=t_e}^{\infty} (1+r)^{-(t-t_e)} \cdot \Delta Y_{it}(\Delta E_{t_e \geq t_p}) \tag{5}$$

Because this matches how the SC-CO_2 is typically computed, for simplicity we denote this component the SC-CO_2. This is the quantity depicted in Fig 1g, aggregated over populations.

S1.3 Total L&D

Total L&D from CO_2 is then defined as the sum of these three components multiplied by total current past and future CO_2 emissions. Summation of damages begins in starting year t_0, or time at which parties begin to be held accountable for their emissions. Emissions may have occurred before t_0, but it is not necessary to track the resulting damage if parties agree not to hold the emitters responsible for those impacts. Existing agreements have not yet defined t_0, so in our analysis below, we present multiple calculations that vary t_0. We note that in many legal contexts, analogs to t_0 are set to the earliest time when it can be shown that a polluting party became aware that their actions could cause damage to other parties.

The above components correspond to damages from the marginal emission of one ton of carbon dioxide at different points in time. Given total global CO_2 emissions E_{t_e} in each past and future year, then total L&D can be written as:

$$L_{t_0} = \sum_{t_e=t_0}^{\infty} E_{t_e} \cdot (HD\text{-}CO2_{t_e} + FD\text{-}CO2_{t_e} + SC\text{-}CO2_{t_e}) \tag{6}$$

which represents the cumulative present value of all past and future damages from past (since t_0) and future emissions, net of adaptation, discounted to the present day. Here the summation is over damages from each year of emissions, where HD- and FD-CO_2 are non-zero for past emissions and SC-CO_2 is non-zero for future emissions. Again in practice, summation typically stops in some chosen future year (e.g 2100 or 2300).

This definition makes clear the direct link between our definition and computation of L&D and the widely understood concept of the SC-CO_2[3,4]. Most proposals that incorporate the SC-CO_2 charge the emitter the sum total of future damages at the time of emission, equivalent to our definition of SC-CO_2. Full compensation for total L&D, were that a policy or legal goal, would additionally require compensation for past and future damages from emissions that occurred in the past, calculated analogously.

S1.4 Unit-specific L&D, and bilateral attribution to a single emitter

The above approach also enables calculation of total damages experienced by a specific population i, as well as attribution of these damages to a specific emitting party j (e.g. a person, or a firm, or a country). As in total L&D, L&D for population i is the sum of three bilateral components (cumulative past damage in i from a marginal past emission, cumulative future damage in i

from that marginal emission, and future damage in i from a future marginal emission), multiplied by past and future global emissions in each year. The component of these damages attributable to emitter j instead multiplies damages from past and future marginal emissions by yearly emissions from j.

Bilateral attributable L&D from emitter j to receiver i as $L_{j\to i,t_0}$:

$$L_{j\to i,t_0} = \sum_{t_e=t_0}^{t_p}\Bigg(\underbrace{\sum_{t=t_e}^{t_p}(1+r)^{-(t-t_p)}\cdot \Delta Y_{it}(\Delta E_{t_e})\cdot E_{j,t_e}}_{HD\text{-}CO_{2,i,t_e}} + \underbrace{\sum_{t=t_p}^{\infty}(1+r)^{-(t-t_p)}\cdot \Delta Y_{it}(\Delta E_{t_e})\cdot E_{j,t_e}}_{FD\text{-}CO_{2,i,t_e}}\Bigg) + \sum_{t_e=t_p}^{\infty}\underbrace{\sum_{t=t_p}^{\infty}(1+r)^{-(t-t_p)}\cdot \Delta Y_{it}(\Delta E_{t_e})\cdot E_{j,t_e}}_{SC\text{-}CO_{2,i,t_e}} \quad (7)$$

In each term, the outer summation is over emission years, and the inner summation is over the subsequent years in which each emission year creates damages in i.

There are multiple considerations in the calculation of bilateral attributable L&D. First, both j and i are likely emitters, and emissions from i can cause damage in j just as emissions from j can damage i. The net compensation from j to i should then plausibly reflect the net damage flows. For country emitters, we report gross flows (impact of emissions from j on i, impact of emissions from i on j) from which net flows can readily be calculated.

A second consideration is if emissions from j caused estimated benefits in i, either in gross or net terms. While there is no theoretical reason why external damages should be compensated but external benefits not compensated, existing policy discussions regarding loss and damage appear to have only focused on damages – likely because aggregate damages are thought to be substantially larger than benefits, and countries experiencing damage tend to be substantially poorer than either emitting countries or high-latitude countries that might benefit from warming[5]. Nevertheless, we again report gross damage flows between j and i, and we do not constrain damages to be positive (i.e. we compute benefits as well as damages).

S1.5 Ethical and legal considerations in loss and damage estimation

Our proposed approach to computing L&D inherits three key challenges that embody ethical and legal principles, none of which we can fully resolve here.

S1.5.1 Discounting historical damages and selection of a discount rate

We propose that damages which occurred in the past, prior to settlement, are "discounted" similar to how damages in the future are discounted[6,7]. In essentially all financial or decision-making systems, future damages (or benefits) are discounted when they are weighed against damages occurring in the present using a per-period "discount rate" r that behaves similarly to an interest rate. As damages occur further in the future, they are valued relatively less in comparison to current damages when they are transformed into net present value terms. The selection of an appropriate discount rate for climate change policy has generated substantial controversy[8], although it is widely understood and agreed that a nonzero discount rate is crucial to the stability and consistency of inter-temporal decision-making[7,9]. We consider multiple potential discounting approaches.

Discounting damages that have occurred in the past means that their value is *larger* at the time of settlement than at the time when they occurred (Figure 1b). Analogous to the effect of an interest rate on unpaid debt, the inter-temporal discount rate causes the value of past damages to grow exponentially until settlement occurs. This interpretation is conceptually and mathematically identical to how discounting future damage is understood (Figure 1c), the only difference is whether settlement occurs before or after damages are experienced (Figure 1d). From the time value of money perspective, a given dollar amount yesterday is worth more today, because of generally positive market returns. Similarly, from the perspective of time preference, a later-consumed good is still worth less than an equivalent earlier-consumed good, because of our dislike of having to wait to consume things; something that is worth $100 to an individual today is worth less than that to her yesterday if she has to wait to today to consume it. The application of discounting to historical damages is also consistent with standard practice in financial or legal systems, where unpaid costs accrue interest until they are settled.

Discounting of past harm is important for two practical reasons. First, the costs imposed on the harmed party i have an impact that grows with time, since the resources allocated to cope with the harm, or directly lost as a result of the harm, could have otherwise been applied productively (e.g. invested). Second, if discounting is not applied, then an emitting party j is always incentivised to delay settlement, since the resources that would be transferred in settlement can be productively utilized until settlement. Therefore, achieving fair compensation for i and alignment of intertemporal incentives for j require discounting of past damages.

Selecting a discount rate The discount rate r embodies trade-offs in economic valuation between different time periods, and it plays an important role in our proposed calculation of L&D. Different approaches to selecting a discount rate have been widely debated in the context of computing the SC-CO_2[6,7]. To summarize that discussion, at one end of the spectrum, some argue[10] for a low discount rate (e.g. 0.1-2% per year) on ethical grounds because it treats sequential generations more fairly; at the other end, some argue[11] for a higher discount rate (e.g. 4-6% per year) because it more closely reflects how inter-temporal tradeoffs are made in financial markets. We do not present a favored discount rate for calculation of L&D here, but instead present results across the range of values expressed in the literature. These include both fixed discount rates as well as so-called "Ramsey" discounting, which links discount rates to future economic growth. However, importantly, we again note that altering the discount rate changes how past and future climate damages are valued (in the present) in *opposite ways* (Figure 1D): A lower discount rate increases the present value of future damages, but decreases the present value of historical damages; conversely, a higher discount rate decreases the present value of future damages, but increases the present value of historical damages. Estimates in Fig 2 allow analysts to compute past and future damages from past emissions under discount rates of their choice, including a choice that uses different rates for past and future damages.

S1.5.2 When is t_0?

It remains widely debated when to begin counting emissions that parties should be held accountable for. GHG emissions rose rapidly beginning in the mid-20th century, and these emissions were usually the result of activities that benefited the emitters and their descendants. However, many legal systems do not hold parties accountable for generating damages if the party did not know their actions caused harm. Following this logic, previous analyses have worked to establish when some major emitters first understood that GHGs would cause harm. For example, scientists at Exxon warned company executives about potentially damaging global warming beginning in at least 1977, and multiple utilities and car companies were also aware of anthropogenic warming by the 1970s[12]; widely-publicized hearings on the science of anthropogenic climate change were held in US Congress by the mid-1980s. Here, we do not resolve what the correct value of t_0 is, but instead present results for a range of values. For country-level L&D estimates, we set our baseline estimate of the ”year of knowledge” as 1990, or a year after the establishment of the IPCC. Using text-based analysis of United Nations documents, other analyses set the date a decade earlier[13], and we thus compute estimates using 1980 as an alternate start year.

S1.5.3 Consumption or production-based emissions

Should emitting parties be held responsible for the emissions associated with production decisions made by the party, or should they be held accountable for the emissions associated with the goods and services consumed by the entity? This question is relevant for assigning emissions to countries, whose ability to trade means that production- and consumption-based emissions can differ[14]. It is also relevant for assigning emissions to companies, who emit during the production process but who also produce products whose consumption is associated with emissions[15,16]. Views differ on to whom emissions responsibility, and thus damages, should be assigned. Our approach, where possible, is to simply report both production- and consumption-based emissions and their associated damages.

S2 Empirical approach to estimating loss and damage

To estimate the above framework, we seek to quantify how a given quantity of greenhouse gases emitted in an initial year t_0, or a sequence of GHG emissions in multiple years, affects local and global climate in subsequent years, and how these climate changes shape economic output in a specific location relative to a baseline where emissions were unchanged.

S2.1 Calculating changes in global temperature from emissions

To translate any emissions history into a change in local climate, we first use v2 of the Finite Amplitude Impulse Response model (FaIR)[17] to translate emissions into a change in global mean surface temperature (ΔGMST), and then use the CMIP6 ensemble of fully-coupled global climate models[18] to translate global temperature changes into local changes. FaIR is able to simulate equilibrium and impulse-response behavior of more complex global climate models[17].

For any emission pulse in t_0, which we denote E_{t_0}, we use FaIR to estimate ΔGMST in years

$t > 0$, which we denote $\delta T_t = \nu(E_{t_0})$. In such a pulse experiment, we note that this estimate will represent the amount of warming in year t that occurred due to emissions starting in t_0, which will not be equal to the amount of warming since t_0, since the latter is substantially affected by any prior emissions. Similarly, we can use FaIR to compute the change in temperature in any year from a perturbed history of emissions, by differencing estimates run under perturbed versus baseline trajectories: $\delta T_t = \nu(\mathbf{E}^*) - \nu(\mathbf{E})$. Following previous work, we calculate uncertainty in estimates of δT_t by resampling key parameters in FaIR, including the transient climate response, the realized warming fraction used to calculate the equilibrium climate sensitivity, the short thermal adjustment time and the timescale of rapid carbon uptake by the ocean mixed layer, using parameter values from previous work[19].

S2.2 Calculating local changes in temperature

To map changes in global mean surface temperature (GMST) to local temperature changes for population i, we use a pattern scaling approach based on FAQ4.3 of the IPCC WGI AR6[20] and compute ratio of location-specific warming to area-weighted GMST in each of the CMIP6 ensemble of global climate models (GCMs). Following the IPCC, we calculate the forced temperature response as the difference between temperature projected in the late 21st century of the SSP3-7.0 future climate forcing scenario and the early-industrial historical period, for each of the 30 GCMs for which we are able to match a historical and SSP3-7.0 realization in the CMIP6 archive.

For each GCM, we calculate the forced temperature response separately for each model grid cell and for the latitude-weighted average across all grid cells globally. Then, for each grid point, we compute the ratio of the grid-specific warming to the global-average warming. These grid-specific ratios are then applied to the FaIR-estimated change in GMST, and aggregated to the spatial unit of interest (e.g. country); denote this country-model-year specific local temperature change as δT^*_{imt}. We choose the relatively high-emissions SSP3-7.0 scenario to conduct this pattern scaling because using a higher forcing scenario enables a wider range of global warming to be sampled (as shown in FAQ 4.3 in IPCC WGI AR6[20]), and because as the forcing increases, the signal-to-noise ratio of the forced response relative to internal climate variability increases (e.g., ref[21]), enabling more robust quantification of the spatial pattern scaling per unit of global forcing and of the uncertainty in the spatial pattern of the forced response.

Denoting each CMIP6 model as m and grid cell as g, we compute grid- and model-specific warming as the difference in temperature between projected temperature under SSP3-7.0 averaged over 2080-2100 ($\bar{T}^f_{gm}$) and model-estimated temperature in the early-industrial period averaged over 1850-1900 ($\bar{T}^h_{gm}$)

$$\delta \bar{T}_{gm} = \bar{T}^f_{gm} - \bar{T}^h_{gm} \tag{8}$$

$$\delta \bar{T}_m = \sum_g l_g * \delta \bar{T}_{gm} \tag{9}$$

$$\phi_{gm} = \frac{\delta \bar{T}_{gm}}{\delta \bar{T}_m} \tag{10}$$

where l_g are latitude weights that sum to one.

We then apply this ratio of grid-to-global average warming to the estimate of GMST from FaIR to arrive at grid-specific estimates of warming due to the emissions perturbation of interest, and then we calculate warming in spatial unit i as the population-weighted average of warming in each grid cell that falls into i:

$$\delta T^*_{gmt} = \delta T_t * \phi_{gm} \tag{11}$$

$$\delta T^*_{imt} = \sum_{g \in i} pop_g * \delta T^*_{gmt} \tag{12}$$

where pop_g are grid-cell population weights that sum to one[22]. This is computed separately for each year in which the emissions perturbation led to warming, with ϕ_{gm} held fixed, and repeated for each spatial unit i and climate model m.

S2.3 Estimating economic damages

To translate changes in temperature into economic impacts, let $f()$ represent a damage function that maps changes in temperature to changes in an outcome of interest. We calculate impacts in a desired location i and year t as the difference in outcomes between a baseline temperature in that year T_{it} and perturbed temperature in that year $T_{it} + \delta T^*_{imt}$:

$$\delta y_{imt} = f(T_{it} + \delta T^*_{imt}) - f(T_{it}) \tag{13}$$

When damages are being assessed historically – i.e. where $t < 2020$ – baseline temperature T_{it} is simply the observed population-weighted annual average temperature in unit i, using gridded data from ERA5-Land[23] that is aggregated to the country-year level by taking the population-weighted mean over grid cells and months. When t is a future year, baseline temperature is average of observed temperature over the last ten years of the historical dataset (2010-2020) plus the projected change in temperature in each future year relative to 2020, based on historical emissions and projected future emissions following the SSP3-7.0 trajectories and again using FaIR and the GCM ensemble to estimate global warming and pattern-scale local warming. In this setting, future warming is affected both by all past (observed) and all future (projected) global emissions. Future temperature perturbations δT^*_{imt} are calculated similarly, but with the desired emissions perturbation added or subtracted from the "baseline" SSP3-7.0 trajectory before recomputing GMST change and pattern-scaled local warming.

To estimate the damage function $f()$, we focus on an existing empirically-derived damage function that links temperature fluctuations to GDP growth rates, updating previous work[24] with data through 2019 (see below). Because growth is a cumulative process, the effect of a given year's temperature fluctuation on economic output depends on the previous year's temperature fluctuation and its effect on output. In this setting, damage in a given year is then a function of temperature in both current and past years. Relative to a counterfactual with no warming, hot temperatures can reduce output in a given year because growth in that year is slowed due to contemporaneous hot temperatures, but also because contemporaneous growth rates are acting on an economy that was already smaller due to previous hot temperatures.

To capture these dynamics, we amend equation 13 and compute the change in year- and location-specific growth rates resulting from the change in temperature, and then adjust the observed time series of growth in location i and re-calculate damages in each year as:

$$\delta g_{it,m} = f(T_{it} + \delta T^*_{imt}) - f(T_{it}) \tag{14}$$

$$D_{it,m,t_0} = Y_{ik} \prod_{t=k}^{t} (g_{it} + \delta g_{it,m}) - Y_{ik} \prod_{t=k}^{t} g_{it} \tag{15}$$

where Y_{ik} is GDP in initial emissions year t_0 and g_{it} is the growth rate in country i and year t. For historical years, we take growth rates from World Bank data[25]. For future years, we use SSP3 projected growth rates as our main estimates through 2100. For calculations that require growth rates past 2100, our main estimates fix post-2100 rates at SSP-predicted levels in 2100. We check sensitivity by prescribing either 1% or 2% growth rates for all future years; these values are roughly at the 17th percentile or median projected growth rates by 2100 in an independent recent analysis[26].

To calculate the present value of cumulative damages, we compute the discounted sum of damages in each year between the emissions year and a chosen end year. Both the emissions year and end year differ by application, as described above. For the discount rate r, we follow earlier work and either use constant discount rates between 1% and 5%, or use time-varying "Ramsey" discounting. In the latter approach, discount rates are calculated using the Ramsey equation $r_t = \rho + \eta g_t$, where ρ is the pure rate of time preference, η is the elasticity of the marginal utility of consumption, and g_t is the growth rate in consumption in year t. We fix ρ and η at values used in similar recent exercises that were calibrated to near-term 2% discount rate[27] and estimate g_t using the global average annual rate of growth in per capita GDP in the perturbed emissions scenario, as described above. Under Ramsey discounting, higher per capita growth rates have two competing effects: faster growth leads to larger economies for impacts to act upon, yielding higher total damages, but faster growth also yields higher discount rates because a society quickly becoming richer would prefer to consume more today. In our empirical exercise, we find that these competing factors roughly balance (Fig ED8).

S2.3.1 Estimating temperature-output relationships

To estimate the relationship between location-specific warming and output ($f()$ in equation 15) we follow earlier work[24,28] and use panel fixed effects regression to isolate the contribution of annual temperature fluctuations to variation in growth in real per capita GDP, using national accounts data on country-level GDP from 1961-2019. We estimate distributed lag models of the form:

$$g_{it} = \sum_{k=0}^{n} f(T_{i,t-k}, P_{i,t-k}) + \alpha_i + \delta_t + \theta_i * t + \theta_i * t^2 + \varepsilon_{it} \tag{16}$$

where g_{it} is the first difference of the natural log of real per capita GDP in country i and year t, α_i is a vector of country-specific intercepts (country fixed effects) that account for any time-invariant differences between countries, such as differences in average incomes, average temperatures, or in any other time-invariant factor that could be correlated with both differences in average temperature between countries and differences in average growth rates; δ_t is a vector of year-specific intercepts that account for any shocks or trends in either temperature or growth that are common across countries, such as macroeconomic shocks; $\theta_i * t$ and $\theta_i * t^2$ are country-specific quadratic time trends that additionally flexibly control for locally-trending variables correlated with both temperature and growth. For $f()$, we follow earlier work[24] and use a parsimonious quadratic function to allow growth to respond nonlinearly to temperature and precipitation:

$$f(T_{i,t-k}, P_{i,t-k}) = \beta_{1,k}T_{i,t-k} + \beta_{2,k}T^2_{i,t-k} + \lambda_{1,k}P_{i,t-k} + \lambda_{2,k}P^2_{i,t-k} \tag{17}$$

Annual country-average temperature and precipitation data come from ERA5-Land, as described above. We test robustness to alternate approaches for controlling for time-trending unobservables, including region-by-year fixed effects or linear rather than quadratic country time trends; to a restricted sample of countries with at least 20 years of climate and growth data; and to the use of alternate historical climate data from the Climate Research Unit. Estimates are largely similar under alternate specification choices and inputs (Fig ED2a), and the response function has not changed meaningfully over time (ED2b), ruling out the most obvious income or time-driven adaptation stories.

The sum of contemporaneous and lagged effects in Eq 16 offers an empirical test of whether the effects of a temperature shock on output are persistent or transitory[28]. Because we are estimating non-linear relationships, we evaluate this sum of marginal effects at different points in the temperature distribution, i.e.:

$$\frac{\delta g_{it}}{\delta T_{it}} = \sum_{k=0}^{n}(\beta_{1,k} + 2T_i\beta_{2,k}) \tag{18}$$

A sum of marginals not distinguishable from zero suggests "level effects", in which output returns to its previous trajectory in the years following a temperature shock; a sum significantly different than zero suggests "growth effects", in which output is permanently lower following a temperature shock. In contrast to earlier work[24], in which the sum of contemporaneous and lagged effects were negative but not statistically different than zero, our updated data provide clear evidence of negative growth effects for most of our sample, with the sum of contemporaneous and lagged effects (up to 10 lags) negative for nearly all countries in our data and statistically different than zero for countries with current average temperatures above roughly 15C (Fig ED3). This increase in precision is at least in part due to the larger sample of in these updated data, which extends the earlier data through 2019, increasing data size by more than a third.

Because the specification of counterfactual trends plays a key role in determining whether GDP returns to its trend after a temperature shock, we also test an alternative widely used approach, local projections[29], to study the effects of shocks relative to counterfactual trends. Results are consistent with the distributed lag models: in cold countries, a hot year initially boosts output, but that impact turns negative (but not statistically significant) after a few years. In countries with temperatures above about 15°C , a hot year reduces output immediately and that effect persists for at least the next 15 years, although confidence intervals are somewhat wider as time progresses (Figure S8, Figure S9, Figure S10). See "Using local projections" section below for additional detail.

Given strong evidence of persistent (i.e. "growth") effects from both of these approaches, we choose as our baseline temperature-output damage function the model with the number of lags at which damages appear to have stabilized - i.e. the point at which adding additional lags does not change the marginal effect of temperature on growth. Alternate approaches to model selection in this context (such as using customary AIC/BIC criterion) are not very informative in our setting, given the substantial amount of unexplained variation in outcome data (see "Model selection in estimating damage functions" below). Our approach is instead to assume that additional lags should be included until they no longer affect the cumulative effect (i.e. effects stabilize), and that inclusion of additional lags beyond that point will just add noise but not further alter estimated effects. See "Lag selection in distributed lag models" section below for further details.

To study this stabilization, we compute GDP-weighted marginal effects for each lagged model (i.e. for the distributed lag model with one lag, with two lags, etc and for the local projections model after one year, two years etc). GDP-weighted marginal effects are the average of the estimated marginal effects at each country's average temperature weighted by that country's GDP; these weighted marginal effects offer a parsimonious summary statistic for the average global effect of marginal warming. After computing weighted marginal effects for each model, we then fit a breakpoint algorithm[30] to formalize if and where estimates stabilize. Breakpoint estimates are shown in Fig ED4 and suggest that marginal effects stabilize after about three lags in the local projections models and after seven lags in the distributed lag models. The two approaches give nearly identical estimates at 5 lags, so we adopt the 5-lag model as our baseline model. We also compute some of our main estimates using other lag choices (e.g. zero lag) for comparison. Including lags in the response function, as both the data and earlier work[28] suggest is appropriate, has strong implications for our estimates: positive marginal effects in the zero-lag model at the cold end of the temperature distribution become less positive and are no longer statistically significant as more lags are added. This suggests that, with additional warming, we should be much more confident in negative growth effects in hotter regions than we should be in positive growth effects in cooler regions. For nearly all countries on the planet, point estimates on cumulative marginal effects are negative at lag lengths where marginal effects have stabilized. Negative marginal effects for nearly all countries implies large global damages by end of century under moderate-to-high emission scenario (SSP3-7.0), relative to a world that didn't warm (Figure S11).

S2.4 Non-marginal and interdependent emissions reductions

Emissions perturbations from large emitters (e.g. the US) may be non-marginal: the impact of emissions changes in one year could meaningfully depend on emissions and damages from previous years, for instance because the damage from warming in one year is enough to affect the size of the economy that the next year's warming acts upon. Empirically, we show that errors from failing to account for this dependence are likely quite small (Figure S12), and that in practice the calculation of the impact of removing an emitter's entire emissions history can be closely approximated by simply multiplying their schedule of emissions in each year by estimates of HD-CO_2, FD-CO_2, and SC-CO_2 from above.

Nevertheless, to precisely compute the impact of large, multi-year emissions perturbations (e.g. removal of decades of country emissions or carbon major emissions), we calculate damages by feeding the full multi-year emissions perturbation to FaIR rather than estimating marginal damages from pulses in different years and multiplying by tons emitted. We also show that impact of one country reducing its emissions are largely independent of other countries' emissions choices (Figure S12b). The modest impact of accounting for either non-marginal or interdependent emissions is a result of the rough linearity in the link between cumulative emissions and warming, the relatively small impact on global warming of even a large emissions perturbation given cumulative emissions, and the locally linear relationship between modest warming and damages.

S2.5 Sensitivity of damage estimates to analytic choices

To further understand sensitivity of damage estimates to analytic choices, we compute sensitivity to a larger set of analytic choices, including the discount rate, for which we use either fixed discount rates (1,2, or 3%) or "Ramsey" discounting calibrated to a near-term rate of 2%, following ref[27]; different time horizons after which impacts cease, using either 2100 or 2300, or alternatively assuming that there are no impacts of temperature on growth after 2100 but the wedge between the economy with the perturbed temperature and the baseline temperature remains after 2100; different econometric models, including either the 0-lag model or the 5-lag model; or different baseline growth rates, including a baseline rate of 1% or 2% for every country or the use of country-specific growth projections from SSP3.

For ease of comparison, we focus on the sensitivity of estimated SC-CO_2 to these choices (Fig ED8). Under fixed discount rates, assuming higher baseline growth rates yields much higher estimates of the SC-CO_2, as climate change is acting on a much larger future economy; Ramsey discounting undoes this effect, because higher growth in consumption raises discount rates, limiting the present value of future damages. The time horizon of aggregation also has a large impact on estimates, with damages through 2300 many fold higher than damages aggregated through 2100, depending on the discount rate. Finally, as noted above, the choice of econometric model is also consequential, with models that allow a temperature shock to affect growth in subsequent years generating much higher damage estimates than models that only allow temperature to affect growth in the contemporaneous year.

We estimate two additional scenarios. In an "adaptation" scenario, and despite limited evidence that damage functions have changed over time[31] (Fig ED2b, Fig ED3d), we assume that marginal effects of temperature on growth linearly shrink in absolute value each year until growth is independent of temperature in the year 2100. In "rebound" scenarios, we estimate models that assume that negative impacts on growth stabilize after 5 years and are constant through year 10 (consistent with our data), and then recover linearly over the following 5, 10, or 20 years – consistent with persistent but not permanent effects of temperature shocks on output. Both "adaptation" and "rebound" reduce impacts relative to our base scenario, with larger reductions the larger assumed adaptation or rebound (Fig ED8).

For a set of analytic choices regarding discounting, time horizon of impacts, counterfactual growth rates, and the statistical model used for the damage function, we can calculate uncertainty in the estimated SC-CO_2 resulting from quantifiable uncertainties in our empirical pipeline. We find that the two largest sources of uncertainty for the SC-CO_2 are uncertainty in the historical relationship between temperature and growth (as estimated by our regression analysis), and uncertainty in the sensitivity of GMST to a marginal emission (as estimated by FaIR; Fig ED6). Less important is uncertainty across the climate model ensemble in the mapping of GMST to local warming. We do not consider uncertainty in emissions inventories.

S3 Additional data and estimation details

S3.1 Emissions data

To calculate damages caused by specific emitters, we collect emissions data from various sources. For country-level emissions, we use the Global Carbon Budget 2022 datasets. The datasets calculate country-year-level CO_2 emissions (1850-2021) and are divided into data on fossil fuel emissions (with production and consumption emissions) and land use change emissions[32]. For carbon majors, we use data on Scope 1 and Scope 3 emissions (emissions from direct operations and from the use of sold products, respectively) published by the Carbon Disclosure Project (CDP) in 2017 and spanning the years (1988-2015)[33]. The CDP builds on earlier efforts collecting data on emissions by carbon majors[16]. The CDP utilize company-reported scope 1 emissions when disclosed. In cases where companies do not disclose their emissions, the CDP uses production data to estimate emissions. Emissions estimates are reported in carbon dioxide equivalent (CO_2e).

To compare loss and damage associated with company emissions to company revenues, we collected data on company revenues in 2021 from company annual reports or ref[34]. Finally, for private jet emissions, we estimate emissions from the flight duration of trips taken by individuals, data on which was scraped from public flight records and posted on Twitter (now X) by user Jack Sweeney. We used Twitter's API to collect data on the each flight's duration and used it to estimate emissions, using a constant 503 gallons/hr to calculate fuel burned during a flight. We utilized Ninja API to collect celebrities net worth. We emphasize that flights taken by private jets associated with an individual do not necessarily represent flights taken by that individual.

Estimating emissions and damages from individual actions Recent papers have sought to quantify the amount of GHG emissions associated with individual actions[35,36]. To estimate the damages associated with a selected list of individual actions' emissions, we utilize estimates reported by these studies to calculate the cumulative and future damages (through 2100) of a decade of individual behaviors (2010-2020). Estimates are reported in carbon dioxide equivalent (CO_2e). The list of behaviors includes taking a long-haul flight (2tCO_2e/yr), installing heat pump (-0.8tCO_2e/yr), switching to a vegetarian diet from an average American diet (-0.8tCO_2e/yr), eating one serving of beef a month (0.08tCO_2e/yr), and recycling (-0.06tCO_2e/yr). For comparability, for each of these actions we express damages as a result of doing the more emitting action (taking an additional flight, not eating vegetarian, not installing a heat pump, etc). For emissions associated with driving we utilize the US EPA's estimate of average annual emissions for passenger vehicles. An increase of 10% in driving is associated with an additional 0.5tCO_2/yr emitted.

S3.2 Damage function estimation

S3.2.1 Model selection in estimating damage functions

Since the publication of both Dell Jones and Olken[28] and Burke Hsiang and Miguel[24], hereafter DJO and BHM, many other articles have revisited the question of the aggregate economic impacts of climate change, using alternate panel data or alternate statistical approaches. Many find evidence of non-linear effects of temperature on growth effects that are qualitatively consistent with our findings[37–40].

One paper by Newell, Prest, and Sexton[41] (henceforth NPS) was more skeptical of a temperature/growth link. NPS revisited BHM data and propose to use of cross-validation (i.e. training a model on one dataset and evaluating on held out data the model was not trained on) to make key specification choices. Cross-validation is a common statistical approach to evaluating models' predictive performance, and specifically to help identify models that make accurate predictions on unseen data; its properties are less well known when the goal is causal inference. It is straightforward to show using simulation that, in our panel setting, a model that performs better on a prediction task can perform worse on a causal inference task, and thus that the approach to model selection used by NPS can select models that yield substantially biased estimates of the effect of temperature on aggregate output. In particular, NPS ultimately advocate for models that omit accounting for time-trends, a choice that, in this context, we show can be both particularly prone to bias and appear to perform well in cross-validation.

Specifically, consider a data generating process where a location's time series of annual temperatures T_{it} is composed of a location specific mean $\bar{T}_i$, a long term location-specific linear trend $\theta_i * year_t$, and a mean-zero interannual fluctuation $\epsilon_{it} \sim N(0, \sigma)$:

$$T_{it} = \bar{T}_i + \theta_i * year_t + \epsilon_{it} \tag{19}$$

Suppose that per capita growth in each country is a quadratic function of annual temperature, a country-specific mean growth rate $\bar{y}_i$, a country-specific growth trend ($\lambda_i * year$), and other (unmodeled) time-varying sources of growth (ν_{it}).

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \bar{y}_i + \lambda_i * year + \nu_{it} \tag{20}$$

The econometric challenge is to identify β_1 and β_2 in a setting where mean growth rates and mean temperatures could be correlated ($cov(\bar{T}_i, \bar{y}_i) \neq 0$, e.g. if high income countries tend to be cooler) or where temperature trends could be correlated with growth trends ($cov(\theta_i, \lambda_i) \neq 0$, e.g if high latitude countries have warmed more quickly but grown more slowly). Given these potential confounds, a standard approach for identifying β_1 and β_2 is to estimate panel models that flexibly account for time-invariant differences at the country level (typically using unit fixed effects) as well as for time-varying differences that could be country-specific (typically through the inclusion of time fixed effects and/or unit-specific time trends). The goal of these models is not to maximize the accuracy of predicted growth $\hat{y_{it}}$, but instead to isolate the effect of temperature fluctuations from other correlated factors that could affect growth, i.e. identify unbiased estimates of β_1 and β_2. A model that accurately predicts income growth need not generate unbiased estimates of the impact of temperature, and a model designed to estimate the impact of temperature might not explain the most variation in income growth.

We show that these prediction and causal inference goals can indeed be in conflict by simulating a setting where country-specific trends in income growth and temperature are positively and spuriously correlated. We estimate β_1 and β_2 using these simulated data and evaluate models predicted values of $\hat{y_{it}}$, based on panel models that do or do not include country-specific time trends in the regression. Following NPS we evaluate the predictive performance of models using "forecast" cross-validation, where for a panel of length t years, models are trained on the first x years in each country and then evaluated on the final $t - x$ years. We set $t = 50$ and $x = 40$, set values of β_1 and β_2 to match the earlier point estimates in BHM, and choose the variance of ν_{it} to yield signal-to-noise ratios (as proxied by "within" r-squared values in the panel regressions) similar to what was reported in BHM. Again, we impose positively correlated unit-specific time trends in growth and temperature and ask whether the approach taken by NPS correctly selects models that account for these correlated trends in estimating the relationship between temperature and growth. We estimate two models:

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \gamma_i * year + \alpha_i + \varepsilon_{it} \tag{21}$$

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \alpha_i + \varepsilon_{it} \tag{22}$$

Noting that Equation 21 is analogous to Equation 16 used in this analysis (although simplified for clarity), while Equation 22 is analogous to the model preferred by NPS. We note that neither model includes year FE, because (as in NPS), the years in the test sample will not have an estimated year intercept and so predictions cannot be made out of sample; in this simulated setting, the lack of year FE does not affect inference as we have not added common time-varying sources of bias into the true data-generating process. We then evaluate recovered marginal effects of temperature on growth from the models with and without trends, as well as calculate root mean squared error (RMSE) between predicted and observed growth rates in the held out years; we run this simulation 1000 times.

Consistent with NPS, we find that the model without time trends (Equation 22) generally have lower prediction error in held out data (Figure S13a, p=0.04 on a one-sided test), which is presumably because time trends are fit in-sample on limited amounts of data in each country and provide a noisy estimate of how both income growth and temperature will continue to trend out of sample. However, the model without time trends is substantially biased: marginal effects are too positive across the temperature distribution, which is because income growth and temperature (by design in our simulation) are spuriously positively trending together. Estimates in the model with time trends, which again would be rejected by NPS in favor of the no time trends model, are unbiased. These results indicate that using predictive performance for model selection in our context can undermine inference.

Nevertheless, as NPS argue, it might still appear odd to then select the model that does worse in predicting income growth in future years, if our goal is predicting income growth in future years. Critically, however, *neither BHM nor the current analysis use estimated time trends from these models to project forward either trends in temperature or trends in income growth*, precisely because past trends might be a poor guide for future trends. Instead, we rely on decades of research in climate science embedded in global climate models to project future temperature changes, and we rely on output from multiple modeling teams to project future secular growth in income. We do not need our historical models to accurately predict future trends in either of these variables; instead, we need them to credibly isolate temperature from other time-invariant or time varying factors that could affect income growth.

S3.2.2 Lag selection in distributed lag models

We also use simulation to study the related problem of selecting the number of lags to include in a distributed lag regression. In our setting, following earlier work[28], the goal in running a distributed lag model is to understand whether the data are consistent with "level effects", where output loss following an initial shock fully rebounds to previous levels in the subsequent year or years, or "growth effects", whether there is incomplete recovery after a shock – i.e. a lost wedge in output that is not recovered in subsequent years. Calculating the sum of contemporaneous and lagged effects in a distributed lag regression of growth on temperature allows a quantitative test of levels versus growth effects.

What is the right number of lags to include? Options taken in the literature include using standard model selection criterion (e.g. AIC or BIC) to choose between models, and/or examining the statistical significance of individual lag coefficients and including those that are statistically significant. However, as pointed out by Greene[42], if j is the true number of lags, then for any candidate model with fewer than j lags, the estimator of the coefficient vector on the lags is biased and inconsistent due to omitted variables, and statistical inference is unlikely to be successful. This is particularly true if temperature is autocorrelated. Greene[42] points out that inference is not a problem in a candidate model with more than j lags, as there is no longer an omitted variables problem and the vector of coefficients can be consistently estimated. A second problem arises when, as in our setting, the independent variables of interest only explain a small share of the variation in the outcome of interest, meaning that standard selection criteria (e.g. AIC) can be underpowered to detect differences between models.

A simple simulation makes these points clear. We simulate a panel setting where the "true" relationship between temperature and GDP growth is a five-lag model (i.e. where a given temperature increase in one year affects GDP growth up to five years in the future), where the autocorrelation structure in the temperature data is exactly how it is in our observed real-world data, and where noise is added to the outcome variable (GDP growth) such that the r2 of the relationship between GDP growth and temperature is what it is in the observed data (namely, fairly low, near 0.01). In essence, we are re-generating our panel data to have similar autocorrelation and noise properties as the true data, but where we know exactly what the temperature-GDP relationship is. We then re-estimate the GDP growth-temperature relationship using models that have zero to 10 lags, and use AIC to choose the preferred number of lags. We also plot and test the significance of the individual coefficients. To simplify the exercise, we parameterize a linear relationship between temperature and growth, with marginal effects on lags 0-5 respectively equal to -0.01,0.005,-0.003,-0.002,-0.001, and -0.0005 (i.e. a contemporaneous negative effect, so catch-up in the subsequent year, and then additional damages that decline to near zero by the fifth lag and are zero thereafter). We re-run this simulation 100 times, with new draws of the noise each time, and study how often the AIC criteria chooses the "true" 5-lag model as the preferred model, and how often a significance filter on individual coefficients would correctly select the true number of lags.

Results are shown in Figures S14 and S15. The histogram in Figure S14 showing (out of 100 simulations) how many times the AIC criterion choses a model with a given number of lags. Given the noise in the outcome data (many things beyond temperature affect GDP growth) and the autocorrelation in the temperature data, an AIC criteria chooses the ”true” 5-lag model only rarely, mostly choosing models with fewer lags given the penalty term.

Figure S15 shows the true and estimated cumulative marginal effect under candidate models ranging from a model with no lags to a model with 10 lags (left column of plots; histogram is the distribution of estimates across 100 simulations for each model); the individual coefficient estimates for each model as well as their uncertainty over 100 simulations, plotted alongside the true estimates in red (center column of plots); and finally the percent of simulations where each distributed lag coefficient has a p-value less than 0.05 (i.e. where it would have passed a standard significance filter; right column of plots). As expected, models with too few lags are biased in their estimates of cumulative effects. Models with too many lags (i.e. more than five) are unbiased, but cumulative effects estimates are increasingly noisy as more lags are added. Individual coefficient estimates are unbiased but also noisy, with a standard significance filter only likely to choose (biased) models with very few numbers of lags. We note that this is of course a direct result of the amount of noise prescribed in our simulation, which is calibrated to the amount of noise in our actual data; model selection on a less noisy outcome or less autocorrelated temperature data might be better powered to detect differences between models or individual lags.

Given failure of these model selection approaches to consistently select the right model in a simulated setting that matches key aspects of our real-world data, we instead follow the implicit advice of Greene[42] and the results in this simulation and test models with large numbers of lags, and then look for lag lengths where cumulative effects no longer change upon inclusion of additional lags (see Fig ED4 and related discussion).

S3.2.3 Using local projections to estimate the impact of temperature shocks

We use the local projections[29] approach as an additional method to help understand whether temperature has persistent effects on growth, and to help understand how many years after a temperature shock the effect of that shock on growth is no longer changing. However, we do not use estimates from local projections directly in any calculation of HD-, FD-, or SC-CO_2 in our paper. Nevertheless, here we consider robustness of our implementation of the local projections approach to alternatives.

Our “base” local projections model considers the effect of temperature in year t on the growth in income over the ensuing j years, allowing this effect to differ based on average temperature:

$$log(y_{i,t+j}) - log(y_{i,t-1}) = \rho \Delta y_{i,t-1} + \beta_{j1} T_{it} + \beta_{j2} T_{it} * \bar{T}_i + \mathbf{Z}_{it} + \varepsilon_{it} \tag{23}$$

where the change in log GDP/cap in country i between a shock year $t = 0$ and future year $t + j$ is regressed on growth in the previous year, temperature and temperature interacted with country average temperature, and the same set precipitation controls, fixed effects, and time trends included in our main distributed lag regression, here represented by $\mathbf{Z}_{it}$, where:

$$\mathbf{Z}_{it} = \lambda_{1,k} P_{i,t-k} + \lambda_{2,k} P_{i,t-k}^2 + \alpha_i + \delta_t + \theta_i * t + \theta_i * t^2$$

Temperature coefficients β_{j1} and β_{j2} are estimated separately for each year j up to $j = 15$ As described above, the vector of controls $\mathbf{Z}_{it}$ includes and the same set precipitation controls, fixed effects, and time trends included in our main distributed lag regression:

$$\mathbf{Z}_{it} = \lambda_{1,k} P_{i,t-k} + \lambda_{2,k} P_{i,t-k}^2 + \alpha_i + \delta_t + \theta_i * t + \theta_i * t^2$$

Resulting impulse response functions are plotted in Fig ED3 for different levels of average temperature $\bar{T}$. Confidence intervals are obtained by bootstrapping equation 23 for each j, sampling countries with replacement (1000 resamples).

Because our model includes both unit fixed effects and unit time trends, the impulse response function is estimated using unit specific deviations of temperature from long run trend. In related work by Nath, Ramey, and Klenow (hereafter NRK) that we build upon, the authors take a related approach in which temperature in a given year is first residualized on temperature in previous years to isolate unexpected temperature shocks, and then impulse response functions are estimated using these temperature residuals. Importantly, impulse responses are estimated both for growth as well as for temperature, and then the "cumulative response ratio" is computed, which is the growth effect through year j from the temperature shock in t divided by the cumulative temperature impulse response function through j from the same shock. This approach accounts for any correlation between the temperature shock in year t and temperature in any ensuing year through year j, and the possibility that a +1°C temperature shock in t could induce more or less than +1°C in total warming in the years up to $t + j$. NRK estimate their model without unit time trends (or precipitation), and find that temperature shocks are indeed positively correlated and thus that cumulative response ratios are smaller in absolute value than the growth impulse response function.

We test robustness of our approach to that taken by NRK. We find that first residualizing temperature on past temperature shocks and then estimating growth impulse responses makes these impulse response functions more negative across all values of average temperature (Figure S8; as before, we evaluate the growth impulse response functions at average temperatures of 5, 15, and 25°C). Second, we find that if unit time trends are included in the residualizing regression, future temperatures in years $t + j$ are no longer positively correlated with the temperature shock in year t, and if anything are negatively correlated. This makes cumulative response ratios larger in absolute value (more negative) rather than smaller (Figure S9). In all specifications, including estimation of the cumulative response ratio in a way that mimics, to the best of our knowledge, the approach taken in NPK (where no time trends are included, temperature is first residualized on past shocks, and the cumulative response ratio computed), our results are quantitatively and qualitatively similar to our base model: a one year temperature shock has persistent negative effects on income growth for at least a decade, these effects are larger for hotter countries, and even colder countries with average temperatures at 5C have negative point estimates.

Finally, we additionally show robustness of the local projections approach to different controls for lags of growth, lags of temperature, and/or leads of temperature (Figure S10). Specifically, we now run the model in four ways: with no lags of growth or temperature, with two lags of each, with five lags of each, or with five lags of each and five leads of temperature. The subplots in the panel again show local projections estimates for countries with average temperatures of 5C, 15C, or 25°C . We find that these alternate controls for temperature and growth lags all generate quantitatively and qualitatively similar estimates.

S3.3 Using carbon dioxide removal to alleviate future damages

For a marginal ton of CO_2 emitted in year $t = 0$ (which we set to 2020, as in the SC-CO_2 calculation), we estimate the resulting global damage if the emitting entity uses CDR to remove that ton at some year $t \geq 0$. To do this, we use FaIR to estimate the combined effect on warming of an initial marginal emission in 2020 and then a symmetric removal of the same quantity of emissions in some year after 2020, relative to the same baseline emissions pathway described for marginal emissions above; this is depicted in Figure ED10. Our approach could overstate the benefits of CDR for reducing global temperature given the asymmetry in carbon cycle response to removing CO_2 as opposed to adding it, but estimates suggest that this asymmetry is modest for small emissions perturbations[43].

Damages increase with delay length both because (1) additional warming (and thus damage) that occurs during the years in which the not-yet-removed CO_2 is in the atmosphere and (2) because of the widening wedge between the emissions-perturbed economy versus the counterfactual economy during these years – a wedge that is sustained (but does not further grow) even as the perturbed economy returns to its counterfactual growth rate after the ton is removed (see Fig ED10c for a schematic representation of this effect). As a consequence, accumulation of damage does not stop when an emission is re-captured and the magnitude of these continuing damages increase with the delay between emission and capture.

References

[1] H Damon Matthews and Ken Caldeira. Stabilizing climate requires near-zero emissions. *Geophysical research letters*, 35(4), 2008.

[2] David Archer, Michael Eby, Victor Brovkin, Andy Ridgwell, Long Cao, Uwe Mikolajewicz, Ken Caldeira, Katsumi Matsumoto, Guy Munhoven, Alvaro Montenegro, et al. Atmospheric lifetime of fossil fuel carbon dioxide. *Annual review of earth and planetary sciences*, 37: 117–134, 2009.

[3] National Academies of Sciences, Engineering, and Medicine. *Valuing climate damages: updating estimation of the social cost of carbon dioxide*. National Academies Press, 2017.

[4] Interagency Working Group et al. Technical support document: social cost of carbon, methane, and nitrous oxide interim estimates under executive order 13990. Technical report, Tech. rep., White House, 2021.

[5] Noah S Diffenbaugh and Marshall Burke. Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences*, 116(20):9808–9813, 2019.

[6] Christian Gollier and James K Hammitt. The long-run discount rate controversy. *Annu. Rev. Resour. Econ.*, 6(1):273–295, 2014.

[7] Ben Groom, Moritz A Drupp, Mark C Freeman, and Frikk Nesje. The future, now: A review of social discounting. *Annual Review of Resource Economics*, 14:467–491, 2022.

[8] Lawrence H Goulder and Roberton Williams III. The choice of discount rate for climate change policy evaluation. *Climate Change Economics*, 3(04):1250024, 2012.

[9] Shane Frederick, George Loewenstein, and Ted O'donoghue. Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2):351–401, 2002.

[10] Nicholas Stern. Stern review: The economics of climate change. Technical report, HM Treasury, 2006.

[11] William D Nordhaus. A review of the stern review on the economics of climate change. *Journal of economic literature*, 45(3):686–702, 2007.

[12] Geoffrey Supran, Stefan Rahmstorf, and Naomi Oreskes. Assessing exxonmobil's global warming projections. *Science*, 379(6628):eabk0063, 2023.

[13] Atousa Tahmasebi Lee Robinson, Ian Mitchell. Valuing climate liabilities: Calculating the cost of countries' historical damage from carbon emissions to inform future climate finance commitments. *CGD Policy Paper 233*, 2021.

[14] Steven J Davis and Ken Caldeira. Consumption-based accounting of co2 emissions. *Proceedings of the national academy of sciences*, 107(12):5687–5692, 2010.

[15] Steven J Davis, Glen P Peters, and Ken Caldeira. The supply chain of co2 emissions. *Proceedings of the National Academy of Sciences*, 108(45):18554–18559, 2011.

[16] Richard Heede. Tracing anthropogenic carbon dioxide and methane emissions to fossil fuel and cement producers, 1854–2010. *Climatic change*, 122(1):229–241, 2014.

[17] N. J. Leach, S. Jenkins, Z. Nicholls, C. J. Smith, J. Lynch, M. Cain, T. Walsh, B. Wu, J. Tsutsui, and M. R. Allen. Fairv2.0.0: a generalized impulse response model for climate uncertainty and future scenario exploration. *Geoscientific Model Development*, 14(5):3007–3036, 2021. doi: 10.5194/gmd-14-3007-2021.

[18] Veronika Eyring, Sandrine Bony, Gerald A Meehl, Catherine A Senior, Bjorn Stevens, Ronald J Stouffer, and Karl E Taylor. Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization. *Geoscientific Model Development*, 9(5):1937–1958, 2016.

[19] Ashwin Rode, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E Kopp, Kelly E McCusker, et al. Estimating a social cost of carbon for global energy consumption. *Nature*, 598(7880):308–314, 2021.

[20] JY Lee, J Marotzke, G Bala, L Cao, S Corti, JP Dunne, F Engelbrecht, E Fischer, JC Fyfe, and C Jones. Scenario-based projections and near-term information/climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change, 2021.

[21] Ed Hawkins and Rowan Sutton. The potential to narrow uncertainty in regional climate predictions. *Bulletin of the American Meteorological Society*, 90(8):1095–1108, 2009.

[22] CIESIN. Gridded population of the world, version 4 (gpwv4): Population count, revision 11., 2018.

[23] Joaquín Muñoz-Sabater, Emanuel Dutra, Anna Agustí-Panareda, Clément Albergel, Gabriele Arduini, Gianpaolo Balsamo, Souhail Boussetta, Margarita Choulga, Shaun Harrigan, Hans Hersbach, et al. Era5-land: A state-of-the-art global reanalysis dataset for land applications. *Earth system science data*, 13(9):4349–4383, 2021.

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

[25] The World Bank. World development indicators., 2022. URL `http://data.worldbank.org/data-catalog/world-development-indicators`.

[26] Ulrich K Müller, James H Stock, and Mark W Watson. An econometric model of international growth dynamics for long-horizon forecasting. *Review of Economics and Statistics*, 104(5):857–876, 2022.

[27] Kevin Rennert, Frank Errickson, Brian C Prest, Lisa Rennels, Richard G Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. Comprehensive evidence implies a higher social cost of co2. *Nature*, 610(7933):687–692, 2022.

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

[29] Òscar Jordà. Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182, 2005.

[30] Vito MR Muggeo. Estimating regression models with unknown break-points. *Statistics in medicine*, 22(19):3055–3071, 2003.

[31] Marshall Burke, Mustafa Zahid, Mariana CM 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? Technical report, National Bureau of Economic Research, 2024.

[32] Global Carbon Budget. Global carbon budget 2022. *Earth System Science Data*, 14:4811–4900, 2022.

[33] Paul Griffin and CR Heede. The carbon majors database. *CDP carbon majors report 2017*, 14, 2017.

[34] Macrotrends LLC. Macrotrends. `https://www.macrotrends.net/`, 2025. Accessed: 2025-01.

[35] Brent F. Kim, Raychel E. Santo, Allysan P. Scatterday, Jillian P. Fry, Colleen M. Synk, Shannon R. Cebron, Mesfin M. Mekonnen, Arjen Y. Hoekstra, Saskia de Pee, Martin W. Bloem, Roni A. Neff, and Keeve E. Nachman. Country-specific dietary shifts to mitigate climate and water crises. *Global Environmental Change*, 62:101926, 2020. doi: https://doi.org/10.1016/j.gloenvcha.2019.05.010.

[36] Diana Ivanova, John Barrett, Dominik Wiedenhofer, Biljana Macura, Max Callaghan, and Felix Creutzig. Quantifying the potential for climate change mitigation of consumption options. *Environmental Research Letters*, 15(9):093001, 2020.

[37] Tatyana Deryugina and Solomon Hsiang. The marginal product of climate. Technical report, National Bureau of Economic Research, 2017.

[38] Marshall Burke and Vincent Tanutama. Climatic constraints on aggregate economic output. Technical report, National Bureau of Economic Research, 2019.

[39] Matthias Kalkuhl and Leonie Wenz. The impact of climate conditions on economic production. evidence from a global panel of regions. *Journal of Environmental Economics and Management*, 103:102360, 2020.

[40] Ishan B Nath, Valerie A Ramey, and Peter J Klenow. How much will global warming cool global growth? Technical report, National Bureau of Economic Research, 2024.

[41] Richard G Newell, Brian C Prest, and Steven E Sexton. The gdp-temperature relationship: implications for climate change damages. *Journal of Environmental Economics and Management*, 108:102445, 2021.

[42] William H Greene. Econometric analysis, 5th edition. *New Jersey: Prentice Hall*, 2003.

[43] Kirsten Zickfeld, Deven Azevedo, Sabine Mathesius, and H Damon Matthews. Asymmetry in the climate–carbon cycle response to positive and negative co2 emissions. *Nature Climate Change*, 11(7):613–617, 2021.

Figure S1: **Schematic of SC-CO_2 calculation under different time horizons**. A marginal emission in year 2020 generates subsequent warming and slows growth in a hypothetical economy, relative to a baseline where that emission did not occur. Slower growth through 2100 generates a ”wedge” of damages to total GDP shown in black. In our first scenario, where we assume no damages after 2100, the SC-CO_2 is the black wedge, discounted back to 2020. We then consider two additional scenarios where damages occur past 2100. The first of these, ”no growth impacts after 2100”, assumes that growth returns to its baseline rate in 2100 and thereafter (for instance, because a new technology is invented or practice adopted that eliminates the impact of temperature on growth), but this resumed growth is now acting on an economy that is smaller in 2100 than it would have been without the emission; this generates the blue shaded regions of damages. The SC-CO_2 in this scenario is then the discounted sum of the black and blue wedges. In the last scenario, ”growth impacts through 2300”, we assume warming continues to affect the growth rate through 2300 and impacts end in that year. That generates the additional damages shown in green, and the SC-CO_2 in this scenario is the discounted sum of the black, blue, and green regions. The black lines provide the trajectory of GDP under each scenario. The drawing is schematic and the relative size of the wedges is not to scale. Quantitative estimates of the SC-CO_2 under these different scenarios are provided in Fig 2e.

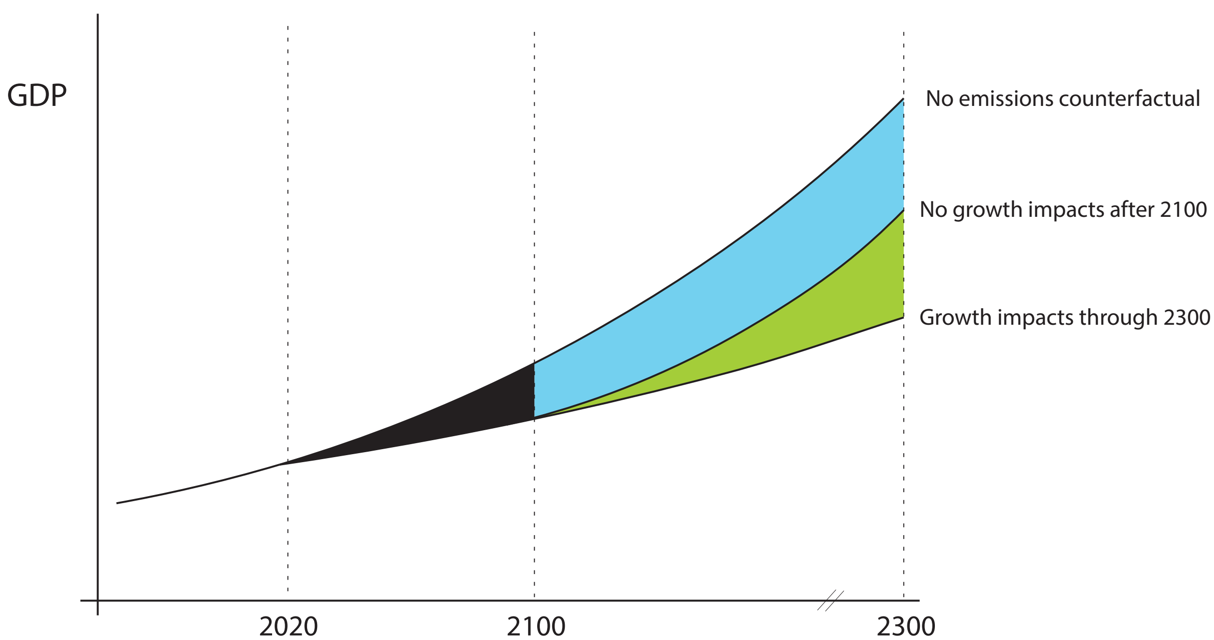

Figure S2: **Estimated damages from emissions related to individual behaviors or firm output over varying time periods**. Estimates match those in Fig 3.

a CUMULATIVE DAMAGES (THROUGH 2100) OF A DECADE (2010-2020) OF INDIVIDUAL BEHAVIORS

	DAMAGES THROUGH 2020	DAMAGES 2021-2100
ADDITIONAL LONG-HAUL (8000KM) FLIGHT PER YEAR	$165 ($55,$385)	$24482 ($6615,$76905)
EATING AND AVERAGE AMERICAN DIET INSTEAD OF VEGETARIAN DIET	$67 ($22,$154)	$9793 ($2646,$30762)
USING A GAS FURNACE INSTEAD OF A HEAT PUMP	$67 ($22,$154)	$9793 ($2646,$30762)
DRIVING 10% MORE THAN AVERAGE AMERICAN	$42 ($14,$96)	$6120 ($1654,$19226)
A SERVING OF BEEF/MONTH	$7 ($2,$15)	$961 ($260,$3018)
NOT RECYCLING	$6 ($2,$12)	$734 ($198,$2307)

b PRESENT VALUE OF FUTURE CUMULATIVE DAMAGES (THROUGH 2100) OF CELEBRITIES PRIVATE JET EMISSIONS IN 2022

	DAMAGES THROUGH 2100
BILL GATES	$1756K ($503K,$5810K)
JEFF BEZOS	$1278K ($366K,$4228K)
FLOYD MAYWEATHER	$1235K ($354K,$4088K)
ELON MUSK	$1169K ($335K,$3868K)
PUMA/JAY-Z	$1139K ($326K,$3770K)
TAYLOR SWIFT	$1015K ($291K,$3359K)
STEVEN SPIELBERG	$844K ($242K,$2794K)
KIM KARDASHIAN	$844K ($242K,$2794K)
A-ROD	$818K ($234K,$2708K)
MARK WAHLBERG	$725K ($208K,$2398K)
TRAVIS SCOTT	$656K ($188K,$2170K)
DAN BILZERIAN	$594K ($170K,$1967K)
KYLIE JENNER	$576K ($165K,$1906K)
JACK NICKLAUS	$567K ($163K,$1877K)

c ACCUMULATED DAMAGES BY 2020 OF EMISSIONS OF CARBON MAJORS 1988-2015 (SCOPE 1 AND 3, $T)

	DAMAGES THROUGH 2020	DAMAGES 2021-2100
SAUDI ARABIAN OIL COMPANY	$4.1T ($0.86T,$8.43T)	$61.14T ($16.03T,$194.1T)
GAZPROM OAO	$4.01T ($0.83T,$8.23T)	$54.55T ($14.27T,$173.38T)
NATIONAL IRANIAN OIL CO	$2.55T ($0.43T,$4.21T)	$30.87T ($8.1T,$98.01T)
EXXONMOBIL CORP	$2.6T ($0.44T,$4.39T)	$27.83T ($7.27T,$88.47T)
PETROLEOS MEXICANOS	$2.39T ($0.39T,$3.79T)	$25.86T ($6.77T,$82.17T)
COAL INDIA	$2.13T ($0.32T,$3.06T)	$24.7T ($6.49T,$78.32T)
ROYAL DUTCH SHELL PLC	$2.28T ($0.35T,$3.49T)	$23.23T ($6.08T,$73.83T)
BP PLC	$2.21T ($0.33T,$3.31T)	$21.46T ($5.61T,$68.19T)
CHINA NATIONAL PETROLEUM CORP	$2.1T ($0.31T,$3T)	$21.3T ($5.58T,$67.65T)
CHEVRON CORP	$2.08T ($0.3T,$2.95T)	$18.56T ($4.85T,$59.02T)
PETROLEOS DE VENEZUELA SA	$1.94T ($0.26T,$2.56T)	$17.13T ($4.48T,$54.44T)
ABU DHABI NATIONAL OIL CO	$1.8T ($0.22T,$2.17T)	$16.13T ($4.23T,$51.19T)
PEABODY ENERGY CORP	$1.7T ($0.2T,$1.88T)	$14.76T ($3.88T,$46.83T)
SONATRACH SPA	$1.66T ($0.18T,$1.79T)	$13.45T ($3.53T,$42.69T)
KUWAIT PETROLEUM CORP	$1.64T ($0.18T,$1.74T)	$13.34T ($3.5T,$42.34T)

Figure S3: **Cumulative damages through 2100 for scope 1 emissions of carbon majors**. Estimates show cumulative historical damages (through 2100) from carbon majors' emissions occurring between 1988 and 2015. Cumulative damages are discounted at 2%. Estimates for Scope 1+3 are given in Fig 3.

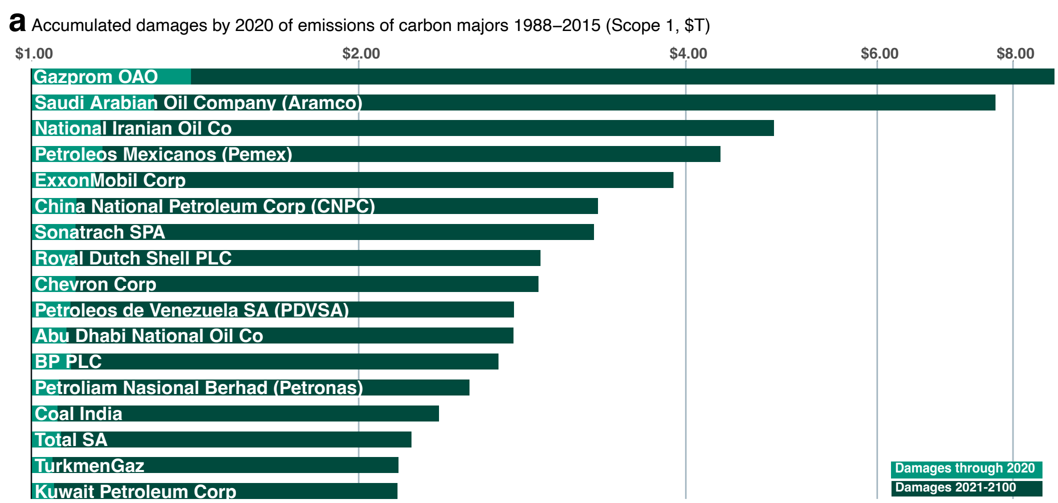

Figure S4: **Attribution of climate damages or benefits since 1980 to specific emitters (Fossil fuel + land use-related emissions).** Same as Figure 4 except valuing damages from emissions starting in 1980 rather than 1990.

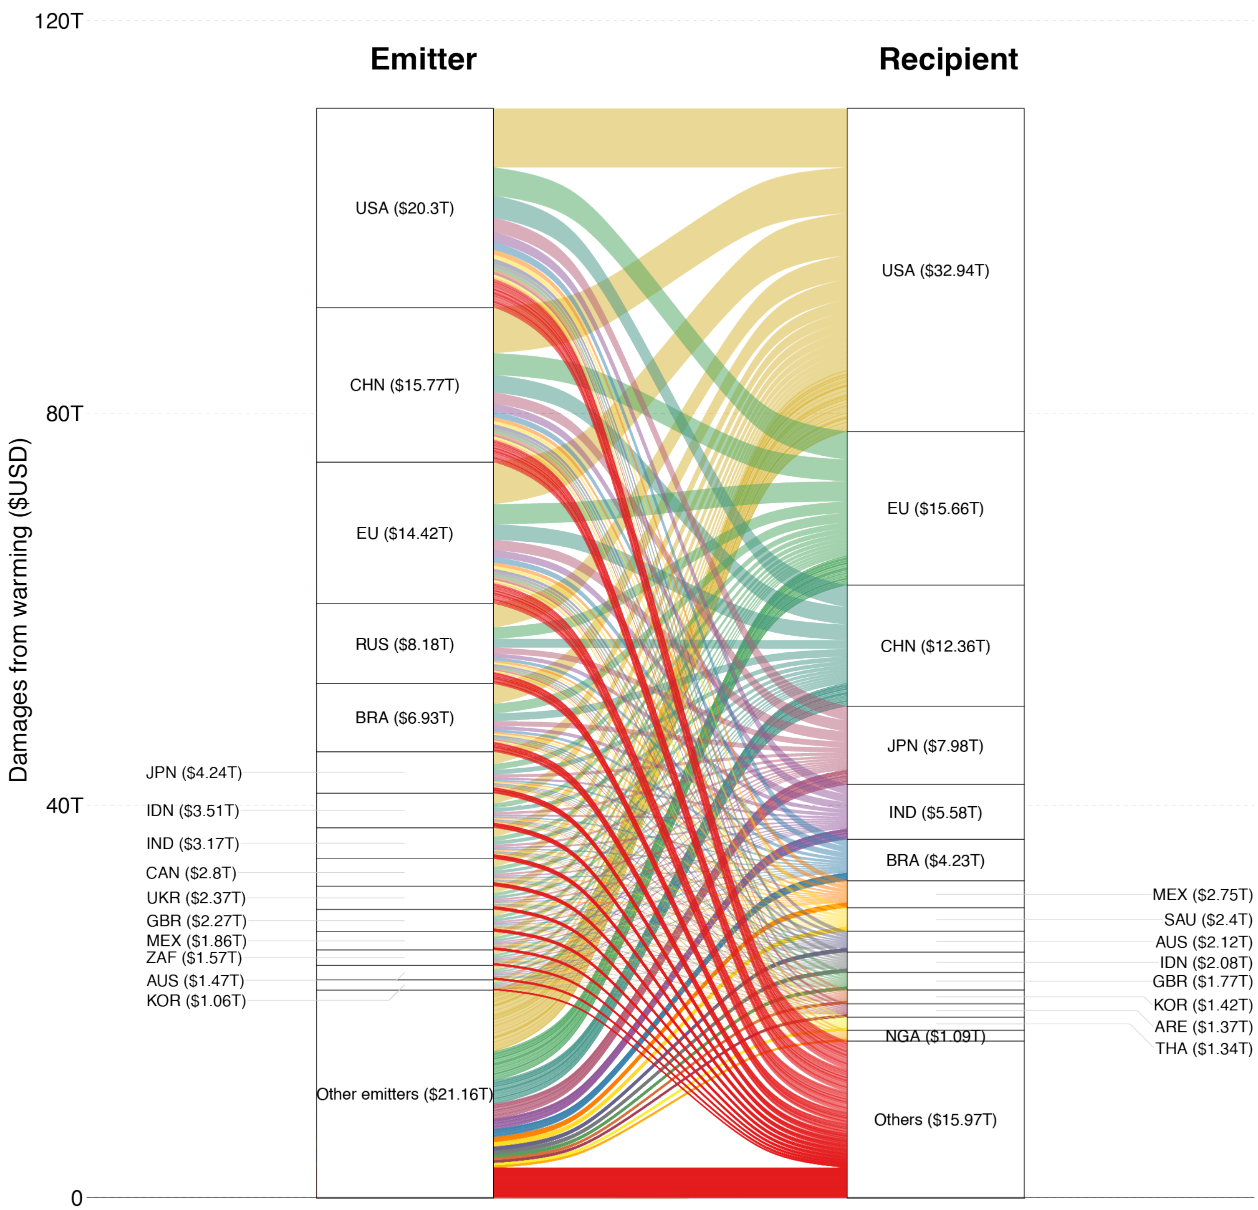

Figure S5: **Attribution of climate damages or benefits since 1960 to specific emitters (Fossil fuel + land use-related emissions).** Same as Figure 4 except valuing damages from emissions starting in 1960 rather than 1990.

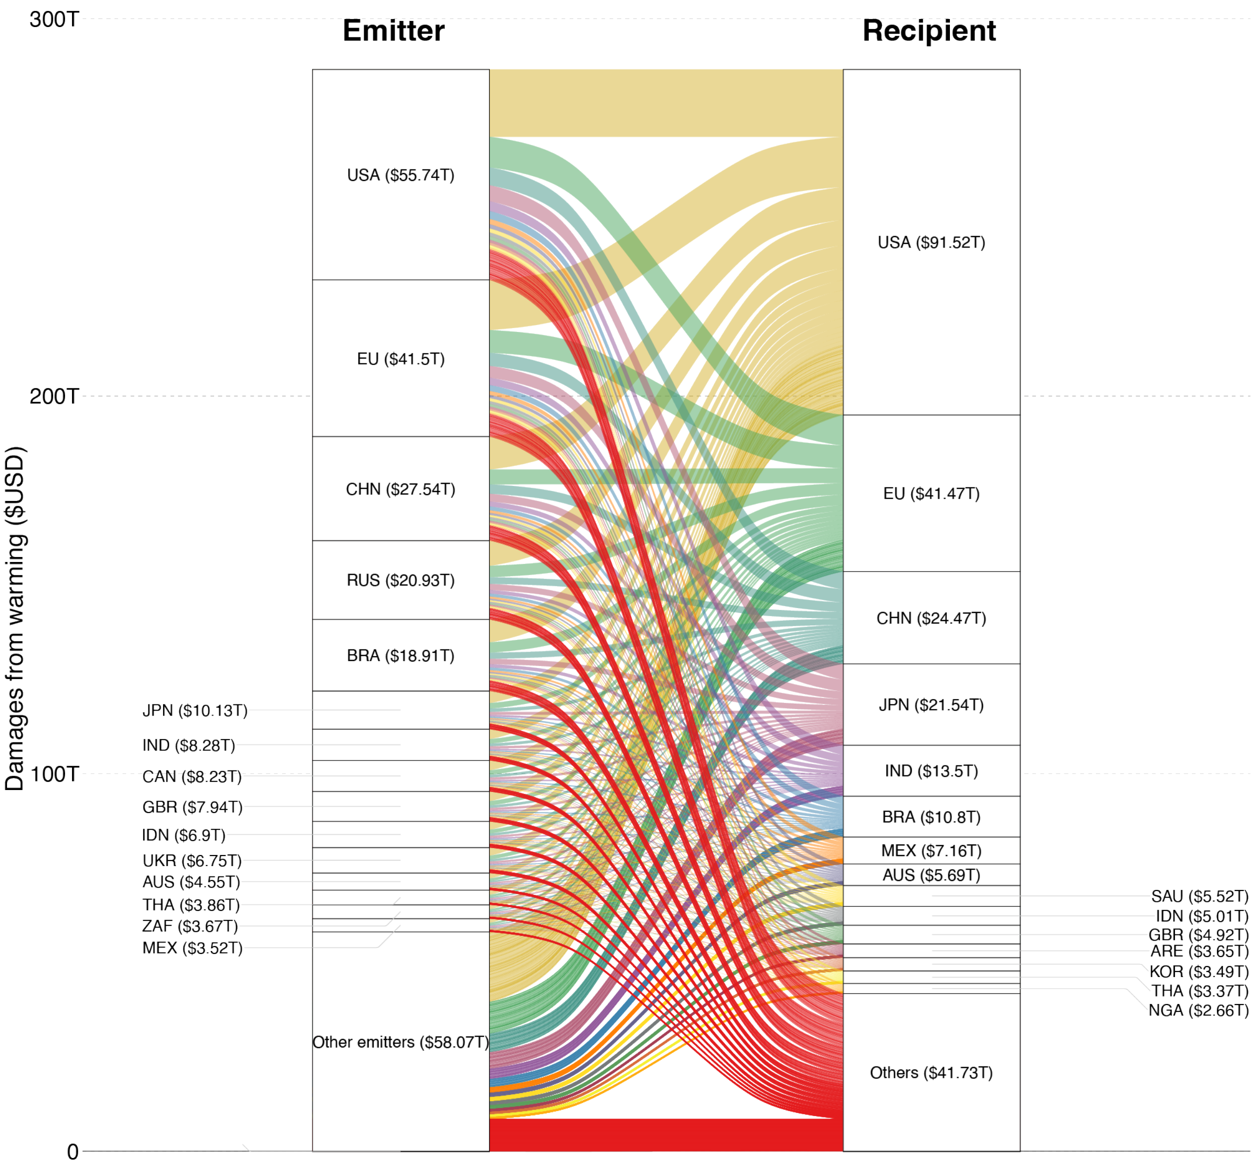

Figure S6: **Attribution of climate damages or benefits since 1990 to specific emitters (consumption emissions) excluding LUCF emissions.** Same as Figure 4 except using consumption-based rather than production-based emissions - i.e. the emissions associated with the products consumed in a country rather than the products produced in that country.

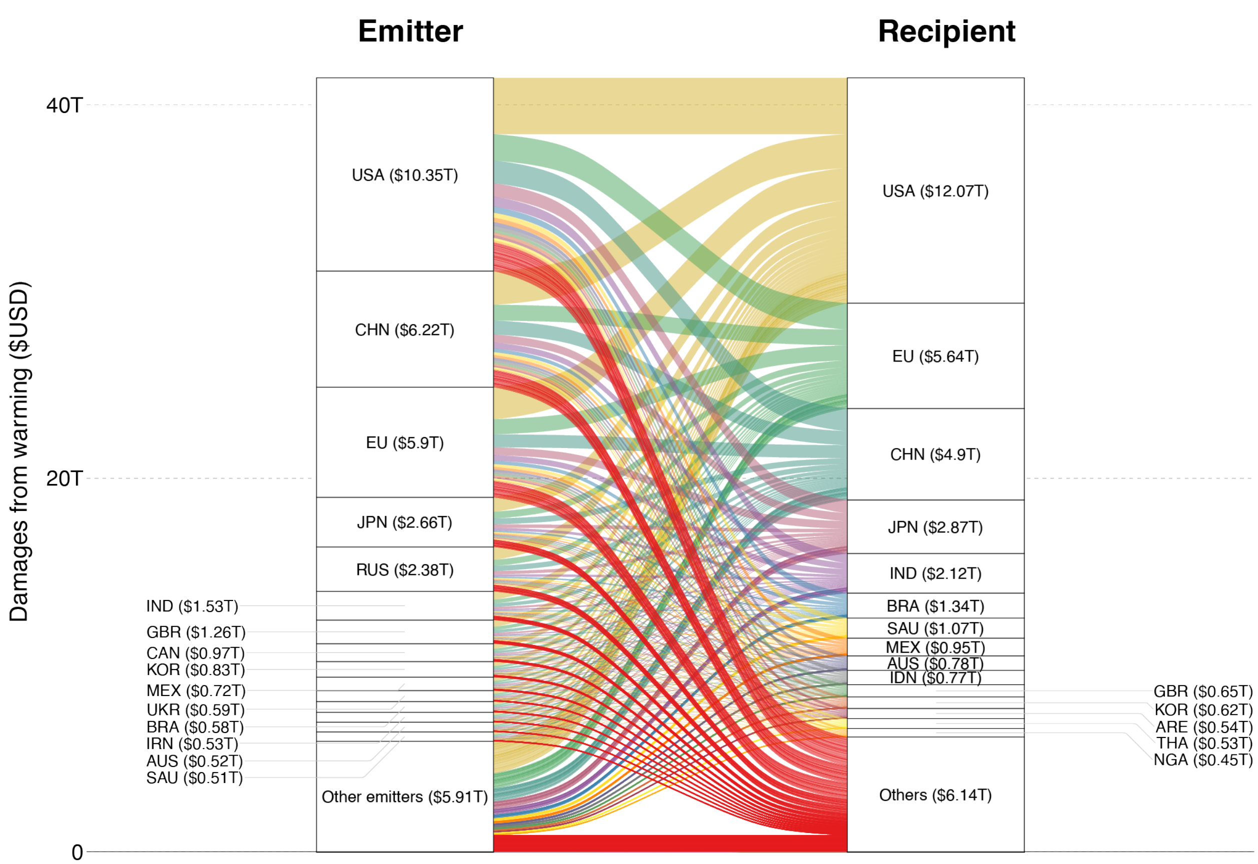

Figure S7: **Attribution of climate damages or benefits since 1990 to specific emitters (production emissions) excluding LUCF emissions.** Same as Figure 4 except excluding land-use-related emissions.

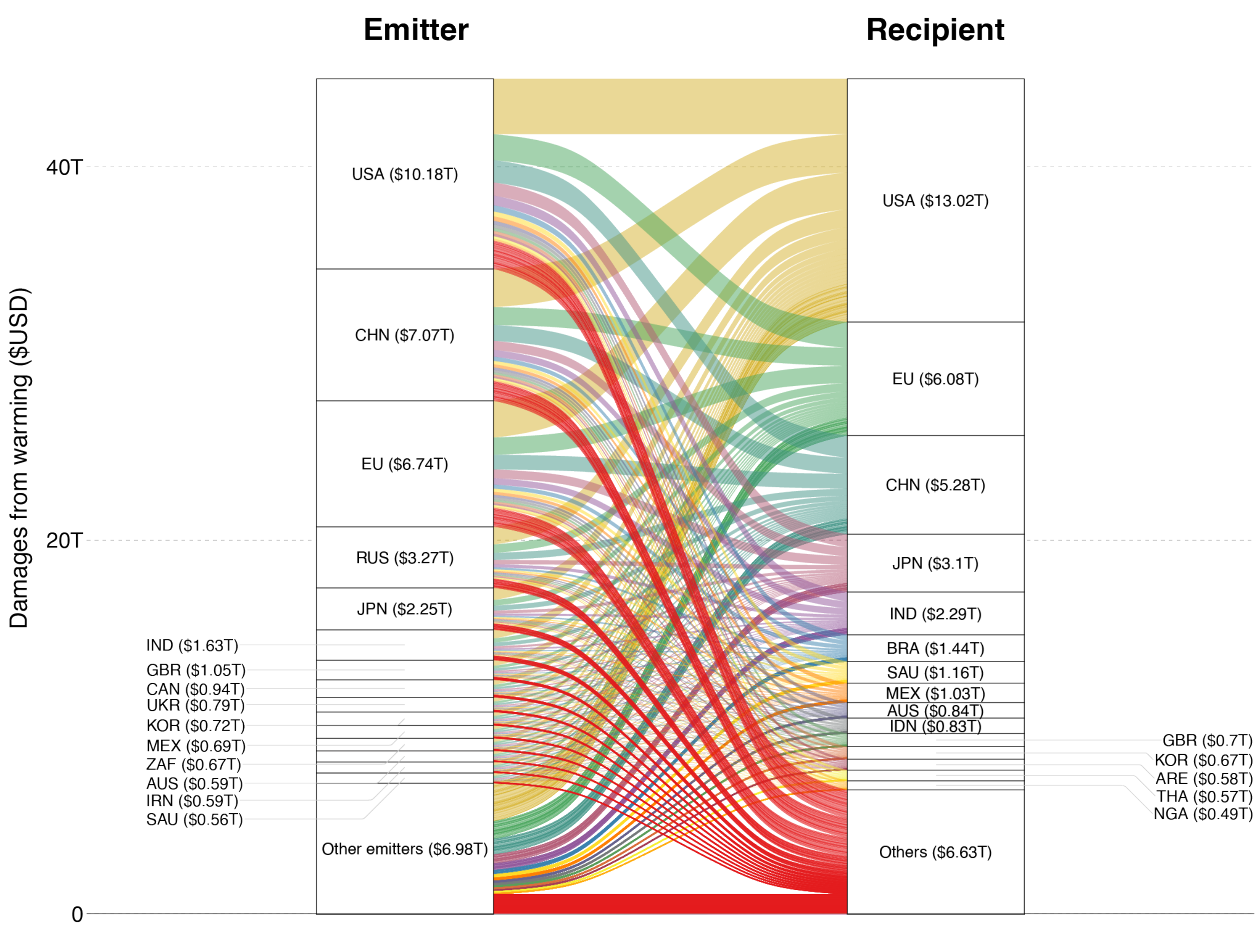

Figure S8: **Sensitivity of local projections estimates to alternate controls**. Estimates plot the impact on growth in the 10 years following a temperature shock in year $t = 0$, for economies with average temperatures at 5 (yellow lines), 15 (orange), and 25°C (red). Top row shows estimates where time trends are included (base model, left) or not included (right) in estimating growth impulse response functions following a 1-year temperature shock in year t. Bottom row shows results with or without time controls where temperature in year t is first residualized on temperature observations from the previous ten years, and growth response estimated as a function of the residualized temperature shock in year t.

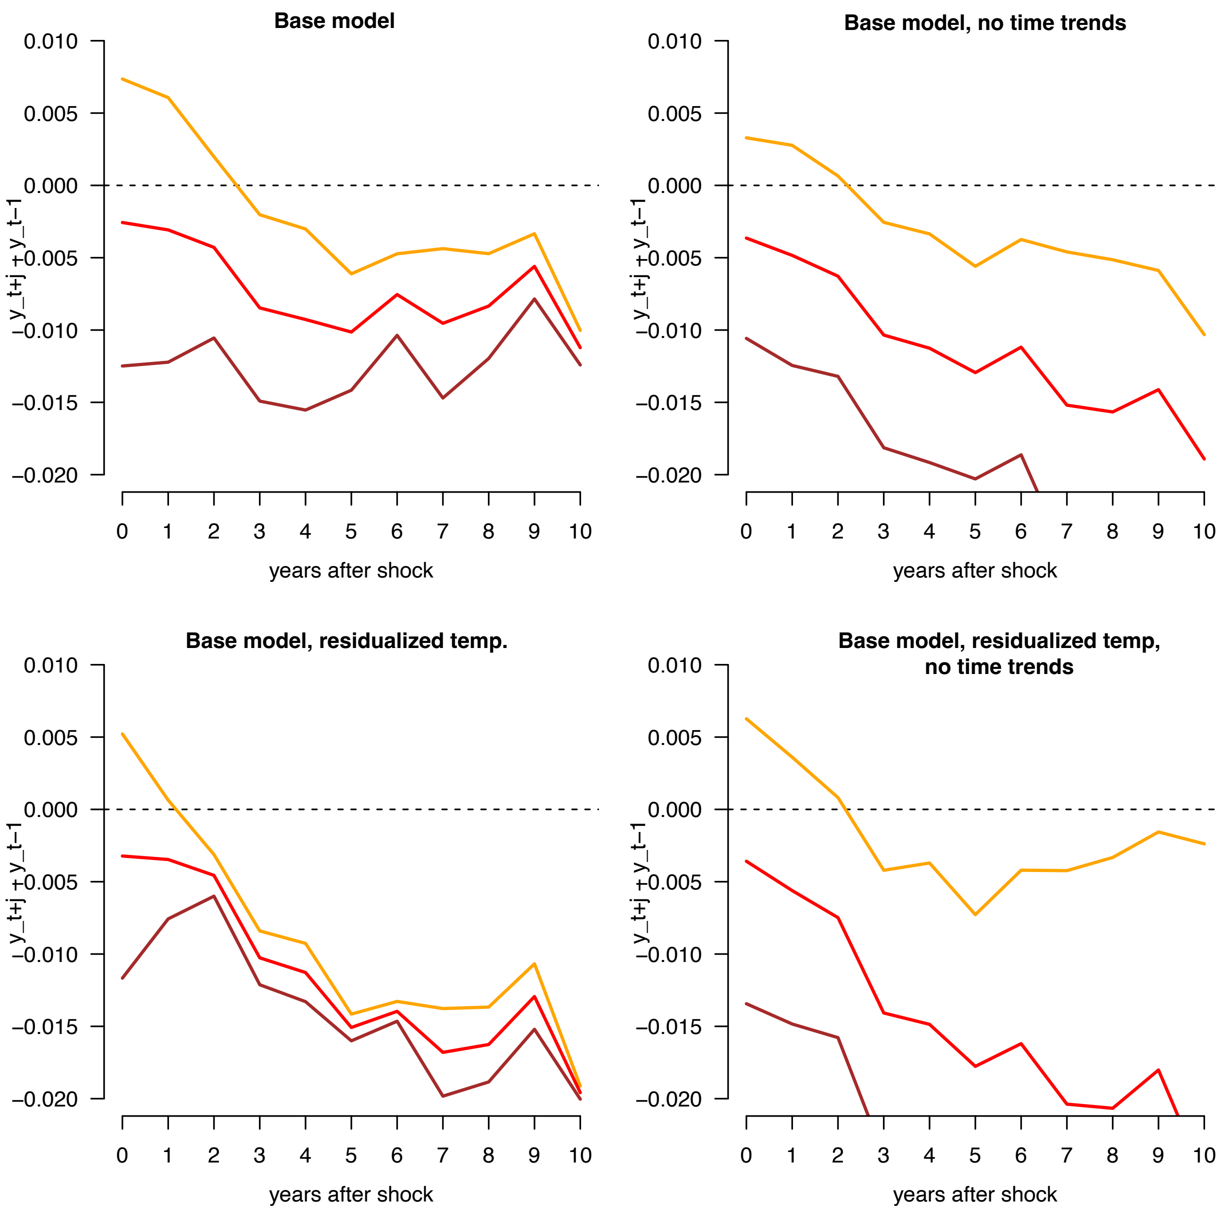

Figure S9: **Growth and temperature impulse response functions, and cumulative response ratios, resulting from a one-year temperature shock under different time controls**. Colors correspond to countries with average temperatures at 5 (yellow lines), 15 (orange), and 25°C (red). Top row of plots mimics Nath et al[40] (NRK) approach but includes unit time trends in both the regression that residualized temperature and in the estimation of impulse response functions; bottom row omits unit time trends. Temperature is positively correlated over time when time trends are omitted (as in NRK) but uncorrelated or negatively correlated with time trends are included. The cumulative response ratio (last column of plots) is the growth impulse response at year $t + j$ from a temperature shock in $t = 0$ divided by the cumulative effect on temperature between $t = 0$ and $t + j$ from that same shock.

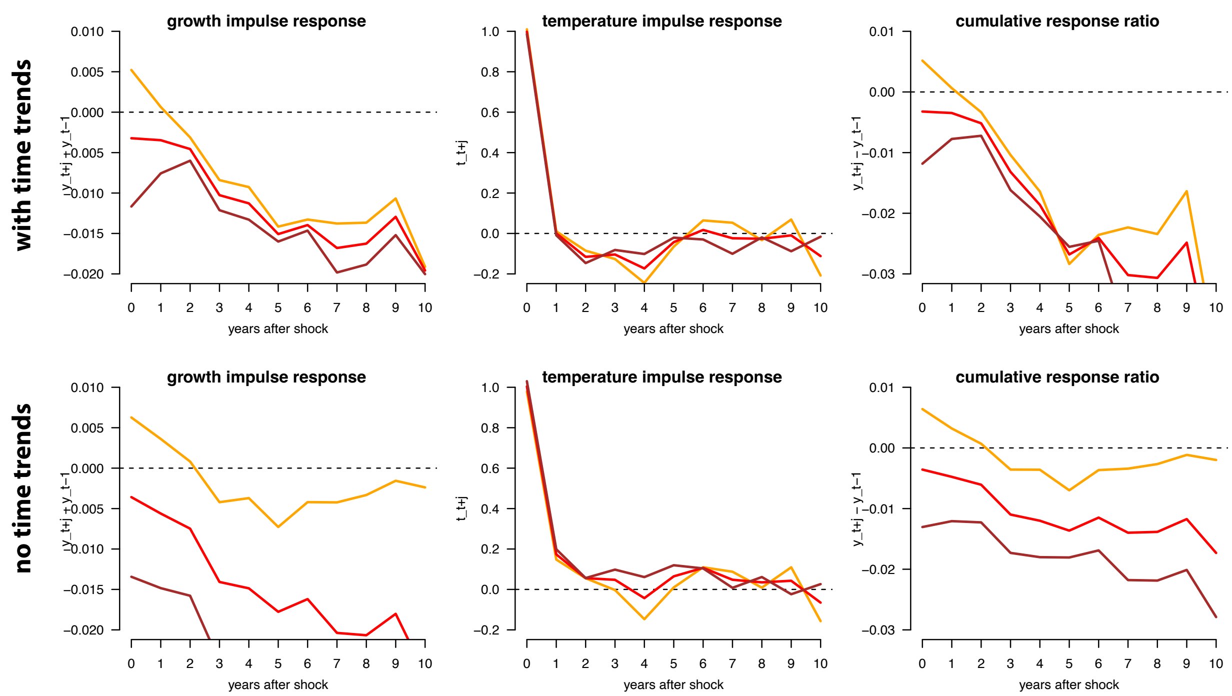

Figure S10: **Local projections estimates are robust to alternate controls.** Results are as in Fig ED3 but under alternate numbers of included lags or leads of temperature and/or growth.

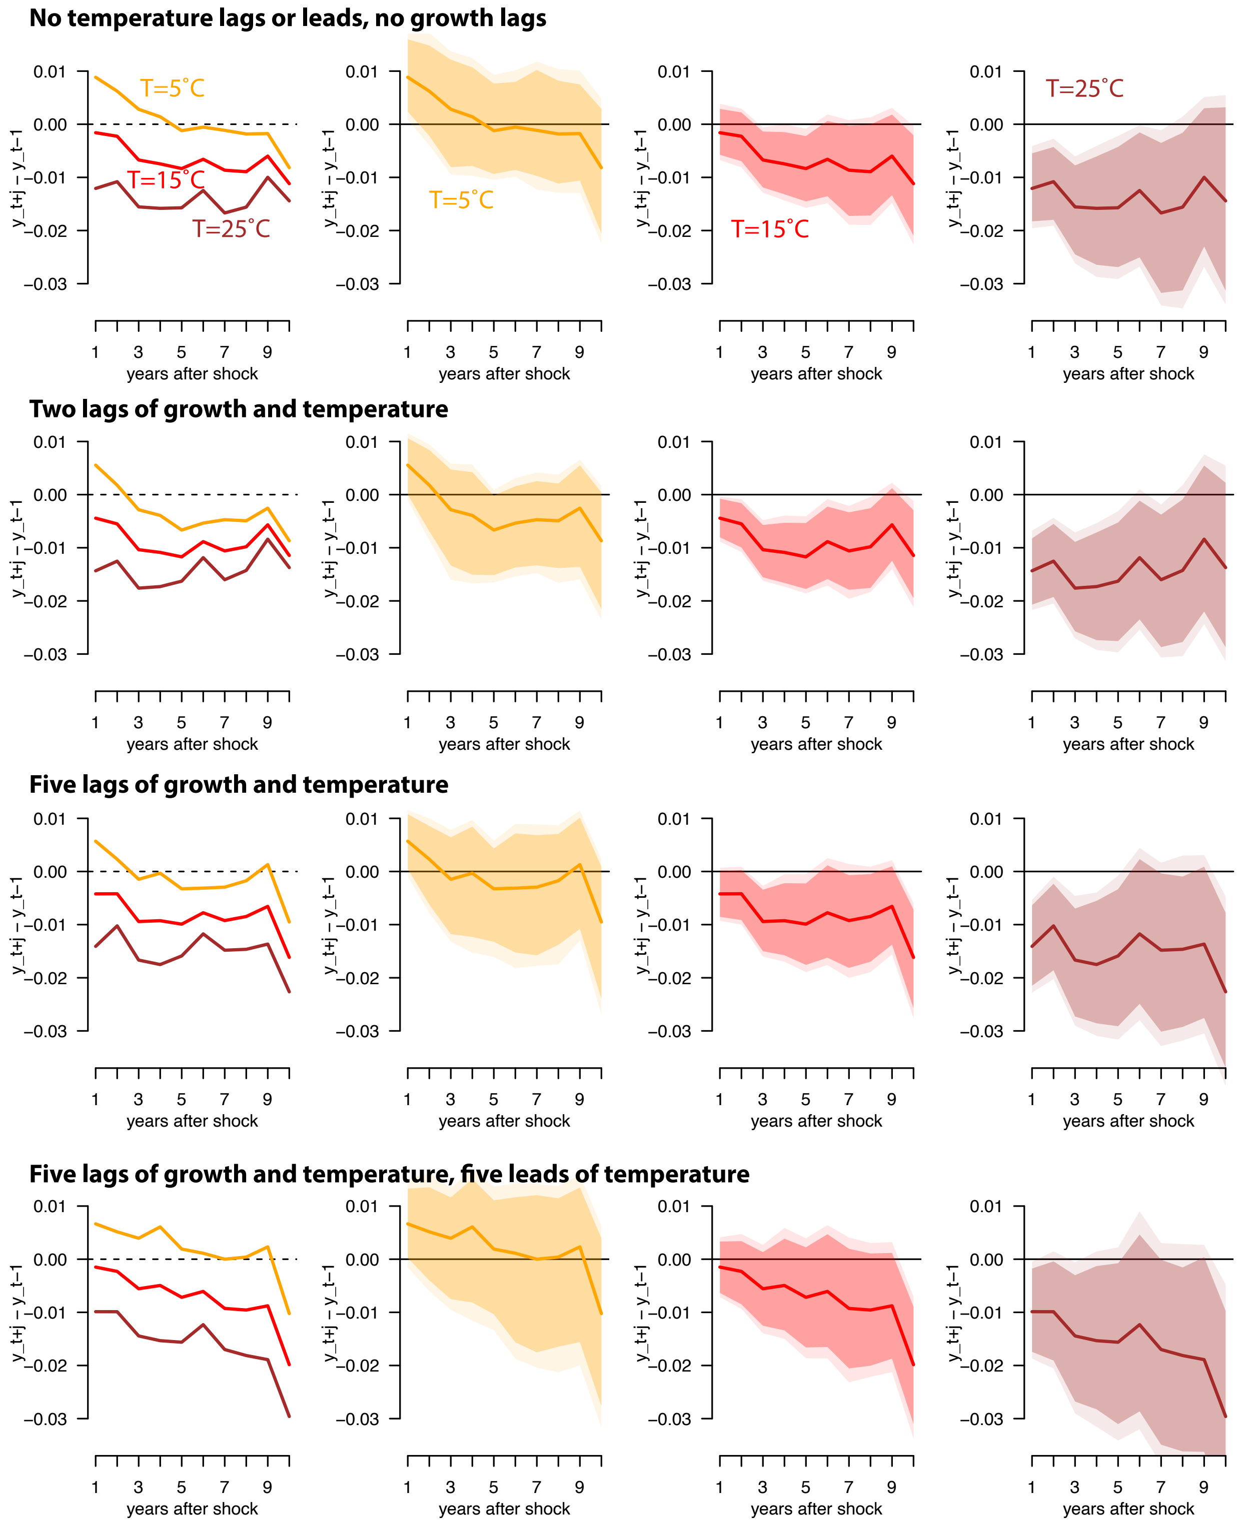

Figure S11: **Global damages to per capita GDP by 2100 under SSP3-7.0, relative to a world with no temperature increase.** Estimates show annual damages under the SSP3-7.0 emissions scenario, with confidence intervals as noted in legend, measured as percent change in global GDP per capita relative to a world where temperatures remained fixed at 2010-2020 values. Negative impacts do not indicate that countries are poorer than they are now, but instead poorer than they would have been absent climate change. Impacts are roughly equivalent to climate change causing the world to grow at 1%/year through 2100 instead of a counterfactual 2%/year over the same period.

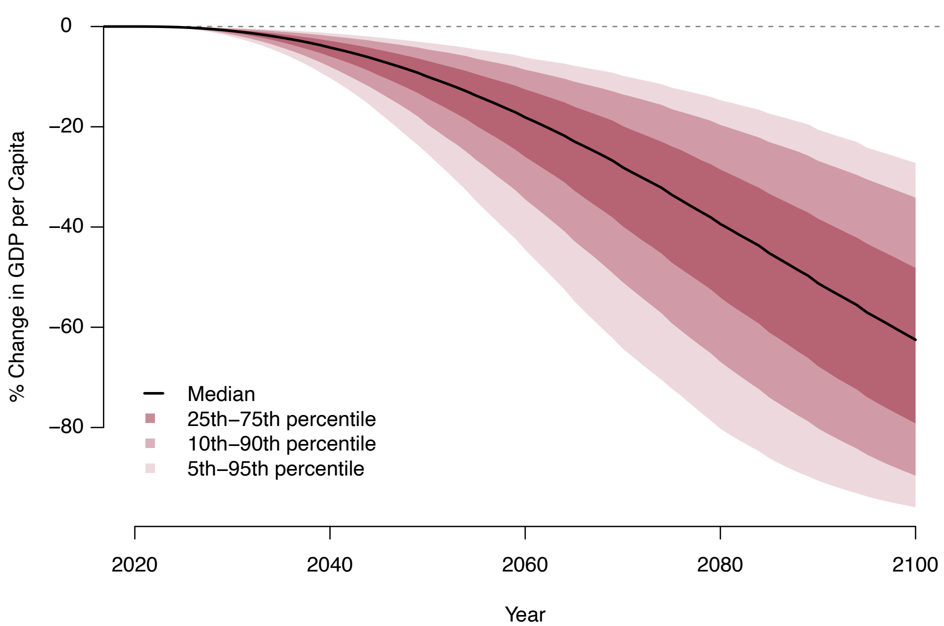

Figure S12: **Estimated damages per ton of CO_2 vary only slightly with the size of the emissions pulse, and damages from country-specific emissions are largely independent of other countries' emissions**. **a** Our baseline approach to calculating marginal damages is with a 1Gt emissions pulse. Using the same 1990 emissions year, estimated per-ton damages using pulses of dramatically smaller or larger size yield damages that differ by only a few percent, both for HD-CO_2 damages (left two columns) and for FD-CO_2 (right two columns). All estimates assume a 2% discount rate, and FD-CO_2 estimates assume impacts end in 2100. **b.** Impacts of historical US emissions (1990-2020) on cumulative global damages through 2020, under alternate assumptions of emissions from all other countries, including no change to other countries' emissions (baseline) or simultaneous reduction in other countries emissions by 10-70%. Values show percentage change in global damages from US emissions, relative to the baseline scenario. Simultaneous reduction in emissions by other countries results in a modest reduction in overall damages from US emissions, and global temperatures are further reduced relative to baseline.

a

	Per tonne HD	% Difference relative to 1GtCO2 pulse	Per tonne FD	% Difference relative to 1GtCO2 pulse
1000tCO2	$185.79	101%	$1822	99%
1MtCO2	$183.88	100%	$1839	100%
1GtCO2	$183.92	100%	$1840	100%
10GtCO2	$184.09	100%	$1844	100%
100GtCO2	$185.83	101%	$1888	103%

b

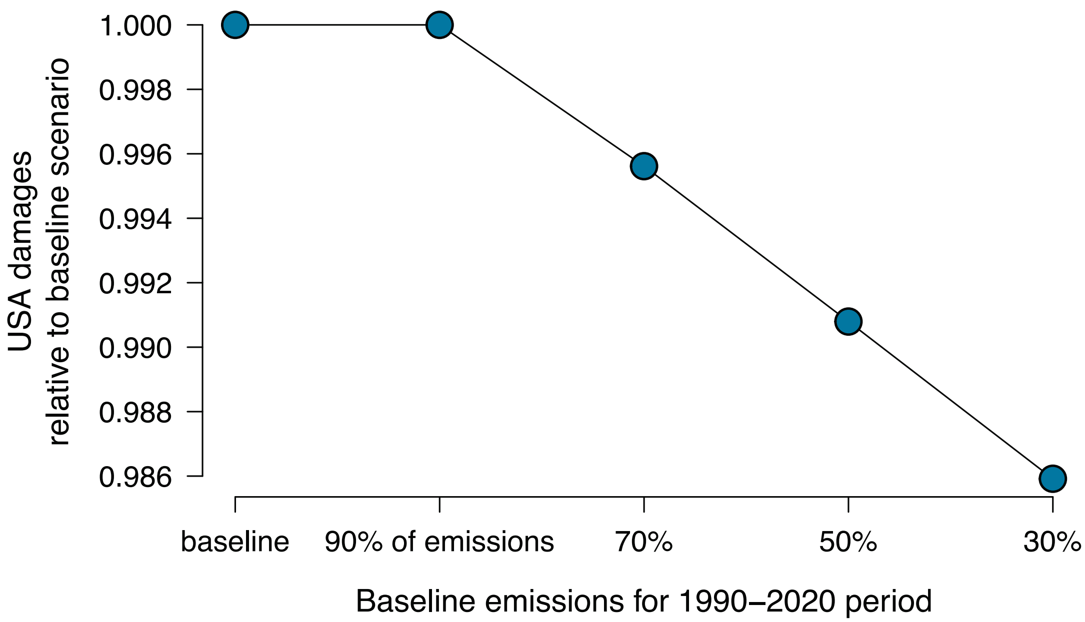

Figure S13: **Selecting regression controls based on model predictive performance can bias estimates of parameters of interest**. Simulation results of estimating regression models with and without time local time trends in a setting where both outcomes (growth) and independent variable (temperature) are spuriously trending by design. Left panel: model with trends has higher RMSE than model without trends on temporally-held-out outcome data. Center panel: model with trends correctly estimates the marginal effect of temperature on growth. Right panel: model without trends has a bias estimate of the marginal effect of temperature on growth, because it did not remove the spurious unit-specific trends. In this example, the bias is positive due to the prescribed positive correlation between local temperature trends and local income growth in the data generating process; the bias in the real world need not be positive. Black link in both panels is the true effect, colored lines are 100 bootstrapped estimates of the two regression models.

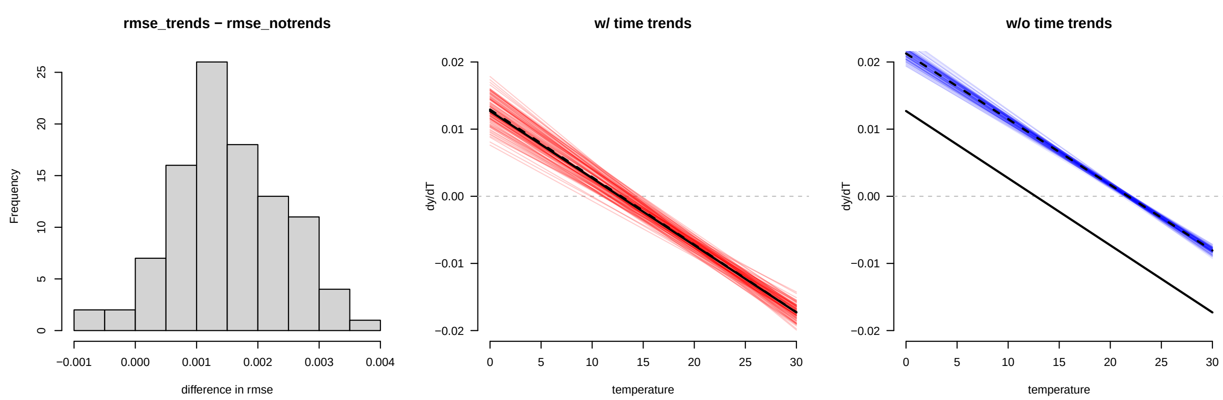

Figure S14: **Model selection via AIC fails to consistently select the model with the true number of lags in simulation**. We simulate a setting in which temperature affects growth in a contemporaneous year and in five subsequent years, and then use panel regression to try to recover this true response by testing candidate models with between 0 and 10 lags. We run the simulation 100 times, and for each simulation selected the model with lowest AIC. See text. The modal choice is a model with three lags, and the "true" 5-lag model is rarely chosen.

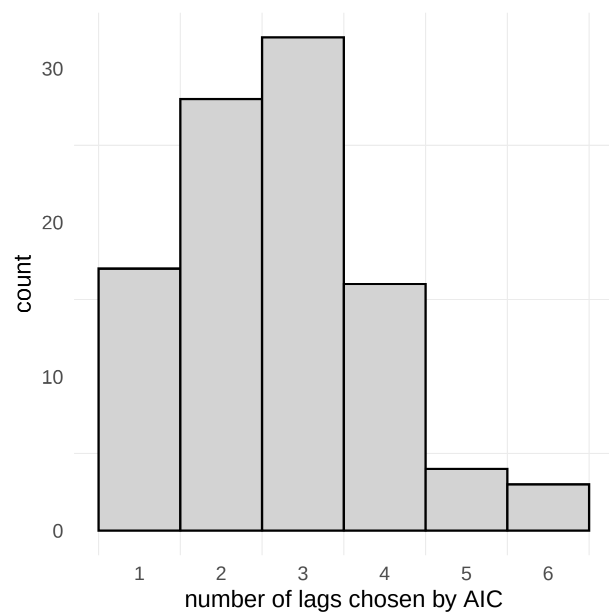

Figure S15: **Distributed lag estimates with too few lags are biased in simulation, and significance filtering does not select the correct number of lags**. Simulation is as in Figure S14; see text. **Left panels**: estimated cumulative effects under different lag lengths; true model has effects out to 5 lags. Red vertical line is the true cumulative effect after 5 lags, black line is the mean effect for a given lag length across simulations, and the density plot is the distribution across simulations. Models with too few lags are biased in their estimates of cumulative effects. Models with too many lags (i.e. more than five) are unbiased, but cumulative effects estimates are increasingly noisy as more lags are added. **Middle panels**: estimates on individual lag coefficients for the same models. Red is true coefficients in the 5-lag model, black is mean estimated coefficients across simulations, gray lines are estimates for each simulation. **Right panels**: for each model and coefficient, the proportion of simulation runs where a 5% significance filter (i.e. the decision to retain a given lag because the p-value on that lag was <0.05) would have selected that lag. The true model has effects out to five lags, but given noise in the outcome, later lags are unlikely to be selected using this significance filter, leading to bias in the estimated cumulative effects (i.e. choosing a model with too few lags).