Global Temperature and Global Mortality*

Adrien Bilal† *Stanford University*

Diego R. Känzig‡ *Northwestern University*

Krzysztof Lisiecki§ *Stanford University*

March, 2026

Abstract

This paper documents large indirect mortality impacts of climate change. Exploiting natural global temperature variability, we find that 1°C of warming increases annual mortality by 2.5 deaths per 1,000 individuals, compared to 0.1 deaths per 1,000 due to the associated local heat exposure. Using external estimates of the impact of global income on mortality, we show that the macroeconomic effects of global temperature changes explain 70% of the gap between these estimates. We use this evidence to estimate damage functions in a neoclassical growth model with finite lifespans. Mortality contributes $420 per ton of CO2 to the Social Cost of Carbon, or one quarter of total climate damages. These results suggest that economic adaptation has a double dividend by also reducing climate-related mortality.

JEL classification: I15, Q51, Q54, Q58, E01, O44

Keywords: Climate change, mortality, global temperature, time-series variation, welfare costs

*We thank Marshall Burke, Tamma Carleton, Solomon Hsiang, Fran Moore, Dev Patel, and Joe Shapiro for helpful comments and suggestions.

†Stanford University, CEPR and NBER. E-mail: adrienbilal@stanford.edu. Web: www.sites.google.com/site/adrienbilal

‡Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com

§Stanford University. E-mail: lisiecki@stanford.edu

1 Introduction

Mortality is one of the most salient consequences of climate change. The associated loss of life and reduced life expectancy can dramatically disrupt the fates of families and communities. A changing climate might affect mortality through a myriad of channels: directly through local heat exposure—the focus of the vast majority of existing studies—but also through other extreme weather realizations, or even indirectly through associated economic contractions. Do the mortality consequences of climate change primarily depend on localized heat exposure, or are wider impacts quantitatively relevant?

In this paper, we construct a comprehensive estimate of the impact of temperature on mortality by leveraging time-series variation in global mean temperature. Our estimates imply that a permanent 1°C increase in global temperature leads to 2.5 additional deaths per 1,000 individuals annually, compared to 0.1 deaths per 1,000 when mediated through local heat exposure only. We show that our global temperature estimate is primarily mediated by the income losses caused by global temperature fluctuations by combining external estimates of global output losses associated with global temperature losses, and external estimates linking income and mortality. We then use our reduced-form results to estimate structural damage functions in a simple model that incorporates both mortality and economic losses. Mortality contributes over $420 per ton of CO2 to the Social Cost of Carbon (SCC) and represents one quarter of overall losses, instead of $10 and less than one tenth of overall losses when based on local heat only.

In the first part of the paper, we develop our empirical approach. We construct a new climate-mortality dataset with comprehensive coverage across the globe. We combine worldwide annual country-level mortality data from the Human Mortality Database and World Health Organization, complemented with direct sources for India and China. This dataset spans 1960–2019 and covers over 100 countries and 70% of the world's population. We link these data with Gross Domestic Product (GDP) information from the Penn World Tables, global temperature data from Berkeley Earth, and gridded extreme weather indicators from the Inter-Sectoral Impact Model Intercomparison Project.

The joint trending behavior of temperature and mortality complicates the identification of causal impacts. To isolate exogenous variation, we construct global temperature *shocks* as innovations to the temperature process that are orthogonal to long-run trends, following Hamilton (2018). The geoscience literature indicates that internal climate phenomena such as the El Niño cycle along with external factors like solar cycles and volcanic

eruptions generate medium-run fluctuations in global mean temperature that are plausibly exogenous to contemporaneous economic and health conditions.

We estimate the dynamic causal effects of global temperature shocks on mortality using panel local projections (Jordà and Taylor, 2025). A global temperature shock scaled to 1°C leads to a gradual increase in the crude death rate that peaks at 1 death per 1,000 after 12 years, with a 95% confidence interval of (0.56, 1.56). The effect is statistically significant over medium- to long-run horizons and gradually mean-reverts after 20 years. Peak effects are four times larger for individuals above 65 years old.

Temperature *shocks* are a useful identification device, but the fundamental determinant of mortality is the *level* of global temperature. Because the temperature shock serves as our instrument, we account for the partially persistent effect of global temperature shocks on the level of global temperature, which remains elevated near 0.5°C for several years.

We convert our estimates to the effect of a permanent increase in global temperature by taking the ratio of cumulative impulse responses of mortality and temperature. We obtain that a permanent 1°C rise in global temperature increases annual mortality by 2.5 deaths per 1,000 people in the long run. This estimate relies on similar identifying assumptions to existing work that focuses on local heat: the relationship between temperature and mortality extends outside the range of shocks in sample and for expected permanent changes.

Next, we address three main identification concerns in a series of robustness exercises. First, omitted variables could bias our estimates if temperature shocks coincide with other adverse global health shocks. Our results remain stable across different sets of controls, including commodity price indices, regional or global trends, the age structure of countries, major wars and epidemics, and indicators for global recessions. The estimates do not appear to be driven by any particularly influential years and remain largely unchanged under alternative definitions of temperature shocks.

Second, reverse causality may bias our estimates if contemporaneous economic activity changes affect emissions and thereby temperature. While this channel is comparatively more relevant for output responses than for mortality outcomes, we nonetheless adjust our mortality estimates for the feedback between contemporaneous output, associated changes in carbon dioxide, methane, and sulfur dioxide emissions, and temperature. Our results remain largely unchanged.

Third, we verify that our estimates are likely to be externally valid. They remain stable in shorter samples (1980–2019 and 1960–2000) and across alternative country samples—excluding China and India or expanding coverage to include African countries—and after removing specific sources of temperature variation such as El Niño events or volcanic eruptions. While no set of sensitivity checks is ever fully exhaustive, the stability of our results across specifications addressing omitted variables, reverse causality and external validity suggests that our specification credibly identifies the causal effect of global temperature on worldwide mortality.

In the second part of the paper, we compare the impacts of global temperature to those of local heat exposure. To maximize comparability, we estimate mortality responses to local heat shocks—innovations in the fraction of days when country-average temperature exceeds historical thresholds—using the same empirical specification and data but replacing global temperature shocks with local heat shocks. We find that 1°C of warming increases mortality by approximately 0.1 deaths per 1,000 per 1°C when local heat is the only mediator. This magnitude is consistent with estimates in Carleton et al. (2022), but is one to two orders of magnitude below our global temperature estimate.

Why, then, does global temperature affect mortality more than local heat? We consider two potential channels through which global temperature affects mortality beyond local heat. First, additional extreme events that correlate with global temperature but that are not captured by local heat may directly influence mortality. However, we show that accounting for the direct effect of extreme cold, droughts, extreme wind and extreme precipitation on mortality falls short of closing the gap.

The second channel we consider is that global temperature changes cause income losses which in turn increase mortality, following the literature on development and health (Preston, 1975; Cutler, Deaton, and Lleras-Muney, 2006). We find strong evidence supporting this income channel by constructing an indirect estimate of the impact of global temperature on mortality that isolates the role of income.

To construct our indirect estimate, we start from the GDP impact of global temperature fluctuations estimated in Bilal and Känzig (2026): a permanent 1°C increase in global temperature reduces world GDP per capita by 22% in the long run. We then combine this response of GDP to global temperature with the response of mortality to global GDP estimated in Bilal, Känzig, and Lisiecki (2026). They exploit innovations in global GDP instrumented with monetary policy surprises to assess the impact of global income on

mortality. Combining the estimate of global temperature shocks on GDP with the estimate of GDP on mortality, we find that global GDP losses mediate 70% of the direct mortality response to global temperature shocks. These findings suggest that the primary reason why global temperature fluctuations increase mortality is because of the induced macroeconomic downturn and the associated worsening of health outcomes rather than direct environmental exposure.

We corroborate this mechanism by exploring heterogeneity in mortality responses along several dimensions. Lower-income regions and countries with low levels of public health care spending exhibit larger mortality responses. Cardiovascular disease and HIV-related deaths drive the aggregate results, while respiratory, tropical, and heat- or cold-related mortality contribute substantially less. This pattern is consistent with mortality operating primarily through economic disruptions and access to healthcare rather than direct heat exposure.

In the third part of the paper, we develop a structural model to evaluate the welfare costs of climate change incorporating both mortality and economic damages. We use a neoclassical growth model with finite lifespans in the spirit of Blanchard and Katz (1992), where households face stochastic mortality and firms produce using capital and labor. Excess temperature relative to baseline affects both mortality rates and total factor productivity through damage functions. We estimate these damage functions by matching the estimated impulse response functions of mortality and output to temperature shocks. We calibrate the value of end-of-lifespan by matching a Value of Statistical Life (VSL) adjusted for age and world income per capita.

We conduct three main counterfactuals. Our first counterfactual is a gradual increase in global mean temperature that starts in 2026 and reaches 3°C above pre-industrial levels by 2100, so 2°C increase above current levels. Mortality damages alone from 2°C warming imply a present-value welfare loss of approximately 15%. Combining mortality with economic losses, the total welfare loss exceeds 40%.

Our second counterfactual evaluates the cost of carbon associated with the warming trajectory caused by the release of one ton of carbon in 2026. We find a Social Cost of Carbon (SCC) of approximately $1,600 per ton of CO2. Mortality contributes $420 per ton to this total under damage functions based on global temperature, with a 95% confidence interval of ($200, $1,050). By contrast, the mortality-only SCC is below $10 per ton for damage functions calibrated to estimates based on local heat exposure, and is not statis-

tically different from zero at the 5% level, consistent with prior work based on local heat. The overall SCC is $160 per ton under local heat. Thus, mortality represents one-quarter of the total SCC when accounting for global temperature damages, compared to less than one tenth of a significantly smaller SCC under local heat estimates.

We also calculate mortality-based damages associated to individual regions following the release of one ton of carbon anywhere in the world: Domestic Costs of Carbon (DCC). We find substantial heterogeneity: populous or wealthy countries face higher domestic costs. For instance, China faces a DCC of just under $50 per ton due to its large population, while the United States faces a DCC of nearly $100 per ton due to their high income levels. As a continent, Africa's DCC is just over $50 due to its low income per capita. Adjusting the VSL for regional income per capita is key for this ranking. Without this adjustment, China's and India's DCCs would exceed $200 per ton, and Africa's would exceed $600 per ton. The worldwide mortality SCC would exceed $1,400.

Our results have two distinct implications for policy and adaptation. First, domestic costs suggest that unilateral decarbonization policies may be cost-effective for populous or wealthy economies even without international coordination and on the basis of mortality alone. Second, the mechanism we identify—macroeconomic disruptions leading to reduced access to care—suggests that policies that adapt the economy to climate change have a double-dividend: they also reduce the mortality consequences of climate change. Thus, society's willingness to pay for economic adaptation is larger than on the basis of economic losses alone.

Related literature. Our findings contribute to several literatures. First, we contribute to the strand of work that evaluates the health and mortality impacts of climate change. Deschênes and Greenstone (2011), Barreca et al. (2016), Burgess et al. (2017), Hsiang et al. (2017), and Carleton et al. (2022) focus on contemporaneous local heat exposure as the main mediator of climate-related mortality. Recent work further examines the role of air pollution caused by wildfires (Burke et al., 2021; Burke et al., 2022) and air pollution more broadly (Aguilar-Gomez et al., 2022). We show that global temperature fluctuations appear to have a larger impact on mortality than implied by local heat and air pollution, largely mediated by associated economic losses.

Second, we relate to work estimating temperature impacts on economic activity (Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015; Newell, Prest, and Sexton, 2021; Nath, Ramey, and Klenow, 2024). Our empirical approach and our analysis of the

income channel builds on Bilal and Känzig (2026): we directly estimate the impact of global temperature fluctuations on mortality, and we separately leverage their estimates of the impact of global temperature fluctuations on output.

Third, our analysis of the income channel of climate-related mortality connects to the broader literature on development and health (Preston, 1975; Pritchett and Summers, 1996; Cutler, Deaton, and Lleras-Muney, 2006; Cutler, Huang, and Lleras-Muney, 2016; Ma et al., 2022). The estimated effects of global GDP on mortality in Bilal, Känzig, and Lisiecki (2026) that we use are qualitatively in line with the well-known negative relationship between economic development and mortality. Quantitatively, their estimates are based on world output and are larger than estimates linking local income and mortality in the cross-section or panel of countries due to global longevity spillovers.[1]

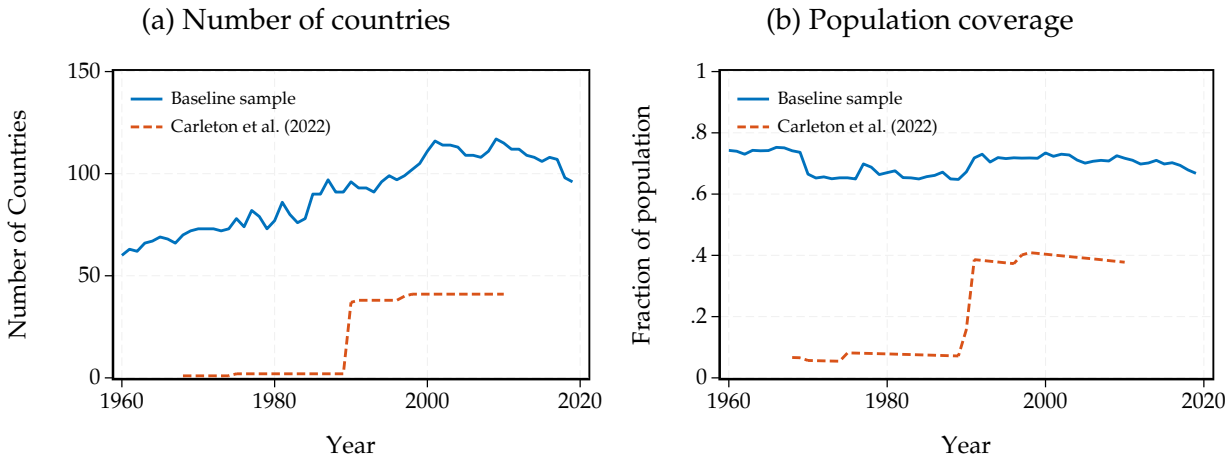
The remainder of the paper proceeds as follows. Section 2 describes the data and construction of temperature shocks. Section 3 presents the baseline estimates of global temperature impacts on mortality. Section 4 compares global temperature to local heat and explores mechanisms. Section 5 develops the structural model and estimates damage functions. Section 6 studies the mortality, economic and welfare consequences of climate change. Section 7 concludes.

2 Data and Empirical Specification

Our goal is to estimate the impact of climate change on mortality. Climate change represents shifts in a broad range of phenomena relative to preindustrial times: heat, cold, precipitation, wind, ocean currents and many others. The scalar measure used by the Intergovernmental Panel for Climate Change (IPCC) to summarize this collection of phenomena is global mean temperature. Therefore, we directly assess the relationship between global mean temperature and mortality.

[1] In advanced economies, mortality has been found to correlate negatively with the local unemployment rate (Ruhm, 2000; Stevens et al., 2015; Finkelstein et al., 2025), although this relationship appears to have weakened over time (Ruhm, 2015). Bilal, Känzig, and Lisiecki (2026) show that the focus on hours or unemployment, as opposed to output, accounts for the change in sign. Our results on the income channel may also help rationalize recent findings by Young and Hsiang (2024) who show that tropical cyclones have delayed mortality effects that appear unrelated to initial accidental deaths.

Figure 1: Mortality Data Coverage



Notes: Panel (a): number of countries with mortality data over time. Panel (b): fraction of world population covered by mortality data. Mortality data for Europe, North America, Asia, Latin America, and Oceania based on Human Mortality Database (HMD), supplemented by WHO when HMD is missing, with direct sources for India and China.

2.1 Mortality Data

To study the effects of temperature on mortality, we assemble a comprehensive dataset of country-level mortality outcomes. Annual mortality data are less widely available than measures of economic activity. We compile information for Europe, North America, and a subset of countries in Asia and South America using the Human Mortality Database (HMD) and the World Health Organization (WHO). We complement these data with information for India and China for 1960–2019 from national statistics. We measure mortality using the crude death rate. We use deaths per 1,000 individuals as our units.

Because our analysis only requires country-level mortality data, we are able to expand coverage relative to Carleton et al. (2022). Our baseline sample coverage rises from 60 countries in 1960, to over 100 countries in the second half of our sample. These countries represent an approximately constant share of world population over the period 1960–2019 that fluctuates between 60 and 80%. Figure 1 illustrates the evolution of country and population coverage over our sample period.

Despite this broad coverage, reliable annual mortality data remain unavailable for many countries, particularly in Africa, where civil registration systems are often incomplete or absent. The United Nations provides comprehensive mortality estimates through the World Population Prospects, constructed by combining censuses, household surveys, and partial registration systems, and by applying demographic models to interpolate

missing observations and ensure internal consistency. While well suited for measuring long-run trends and cross-country differences, these estimates smooth high-frequency variation and are therefore less appropriate for our empirical design.

To nonetheless provide evidence for regions not covered in the baseline sample, we use the United Nations data to examine the mortality implications of global temperature changes in a broader set of countries, including Africa, in Section 4.3.

2.2 Economic and Climate Data.

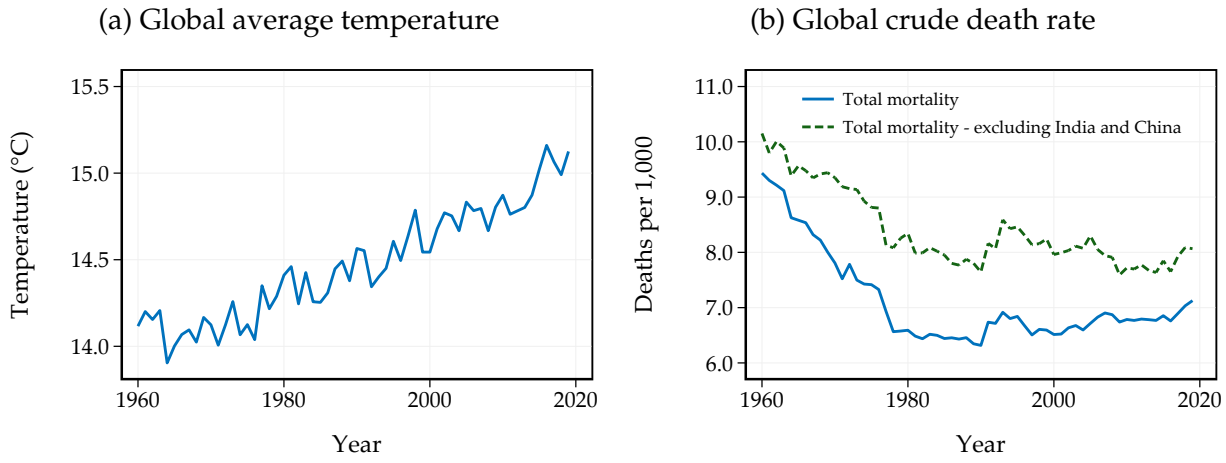
Our economic and climate data builds on the dataset constructed in Bilal and Känzig (2026). GDP and population data are drawb from the Penn World Tables, which provide comprehensive coverage for our sample of countries over the period 1960–2019. We obtain a global commodity price index from Bloomberg. World recession dates are taken from the World Bank (Kose, Sugawara, and Terrones, 2020). Data on armed conflicts are drawn from the Uppsala Conflict Data Program, while information on major epidemics and pandemics is compiled from publicly available historical sources.

We construct global mean temperature data from the Berkeley Earth Surface Temperature Database, supplemented with direct measures from National Oceanic and Atmospheric Administration (NOAA). We construct country-level indicators of exposure extreme weather events using the Inter-Sectoral Impact Model Intercomparison Project (ISIMIP) observed climate dataset. This source provides global, daily reanalysis data on temperature, wind speed, and precipitation from 1901 to 2019 at a $0.5^\circ \times 0.5^\circ$ resolution. For extreme heat, cold, wind, precipitation, and droughts, we construct exposure indices as in Bilal and Känzig (2026) by recording, for each country and year, the fraction of days in which country-wide average weather exceeds or falls below fixed percentiles of its 1950–1980 daily weather distribution. See Appendix A for details.

2.3 Global Temperature Shocks

Figure 2 displays the evolution of global mean temperature and global mortality since 1960. Global mean temperature remains relatively stable around 14°C until the mid-1970s. From the 1970s onward, global mean temperature begins to steadily rise. At the same time, the global crude death rate declines from approximately 9 deaths per 1,000 in 1960 to around 7 deaths per 1,000 by 2019.

Figure 2: Global Mean Temperature and Mortality Since 1960



Notes: Panel (a): evolution of global average temperature (including sea surface) from NOAA since 1960. Panel (b): evolution of global crude death rate (deaths per 1,000 population) computed from HMD, WHO, and national statistics weighted by country population.

Figure 2 indicates that both temperature and mortality trend over our sample. These trends create a well-understood challenge, as simple correlations will deliver a spurious link between the variables (Granger and Newbold, 1974). Therefore, we follow the approach in Bilal and Känzig (2026) and focus on *temperature shocks* for identification. We define such shocks as possibly persistent deviations from the long-run trend in global mean temperature.

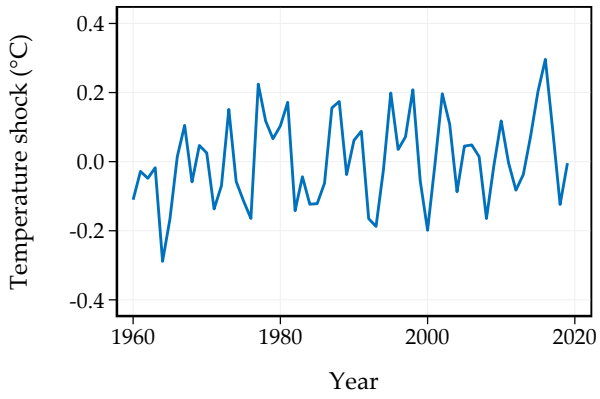
The geoscience literature suggests that two types of phenomena cause global mean temperature to fluctuate around its trend independently from human activity: external causes such as solar cycles and volcanic eruptions (National Oceanic and Atmospheric Administration, 2009; National Oceanic and Atmospheric Administration, 2005), and internal climate variability such as the El Niño-La Niña cycle (Kaufmann, Kauppi, and Stock, 2006; National Oceanic and Atmospheric Administration, 2023).

We isolate the trend and innovations of the temperature process with the method proposed by Hamilton (2018) and applied in Bilal and Känzig (2026). We project global mean temperature on its lags as of period $t-h$ and construct the temperature shock as the corresponding innovation:

$$\widehat{T_t^{\text{shock}}} = T_t - (\widehat{\alpha} + \widehat{\beta}_1 T_{t-h} + \ldots + \widehat{\beta}_{p+1} T_{t-h-p}), \tag{1}$$

where $\widehat{\beta}_i$ denotes the coefficient estimates of the regression of temperature on its lag $j = (h+i-1)$ and $\widehat{\alpha}$ is the estimated intercept. This exercise amounts to isolating shocks that

Figure 3: Global Temperature Shocks



Notes: Global temperature shocks for 1960-2019, computed using Hamilton (2018) filter with $h = 2$ and $p = 2$. Shocks represent deviations from the long-run temperature trend driven by solar cycles, volcanic eruptions, and internal climate variability such as El Niño events.

persist typically for h periods.

Motivated by the observation that phenomena such as El Niño or solar cycles can last for several years, we select a horizon of $h = 2$. We set the number of lags to $p = 2$. Varying these values leaves our results virtually unchanged. Results are also robust to using alternative detrending methods such as the Hodrick-Prescott filter, as we show in Appendix C.3.

Figure 3 shows the resulting global temperature shocks over our sample. They fluctuate around zero, are usually less than 0.2°C, and are weakly autocorrelated because we allow for relatively persistent deviations from the long-run temperature trend. In our empirical specification, we therefore control for lagged temperature shocks to avoid bias from serial correlation.

2.4 Empirical Specification

The mortality effects of temperature shocks may take time to materialize. Therefore, we estimate the dynamic effects of temperature shocks up to 20 years following the initial shock. We do so with panel local projections (Jordà, 2005; Jordà and Taylor, 2025), which allow us to estimate medium-run effects directly without extrapolating from short-term impacts.[2]

[2]We use local projections rather than VARs or distributed lag models because they tend to be more robust at longer horizons (Montiel Olea et al., 2024). Local projections directly estimate the effects of interest at each horizon rather than extrapolating from the first few autocovariances, and they allow for more

We estimate the following series of panel regressions, one for each horizon $h = 0, \ldots, 20$:

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \theta_h T_t^{\text{shock}} + \mathbf{x}_t' \boldsymbol{\beta}_h + \mathbf{x}_{i,t}' \boldsymbol{\gamma}_h + \varepsilon_{i,t+h}, \tag{2}$$

where $y_{i,t}$ is the mortality rate (deaths per 1,000 individuals) for country i in year t, T_t^{shock} is the global temperature shock constructed in Section 2.3, and θ_h is the dynamic causal effect of interest at horizon h. We refer to the sequence $\{\theta_h\}_{h=0}^{20}$ as the impulse response function. Because we use the cumulative change in mortality as the dependent variable, we estimate a possibly persistent level effect on the mortality rate. The country fixed effects $\alpha_{i,h}$ absorb time-invariant differences in mortality levels across countries, such as baseline health system quality or demographic composition. The vector $\mathbf{x}_t$ contains global controls and $\mathbf{x}_{i,t}$ contains country-specific controls. The error term $\varepsilon_{i,t+h}$ captures residual variation in mortality outcomes.

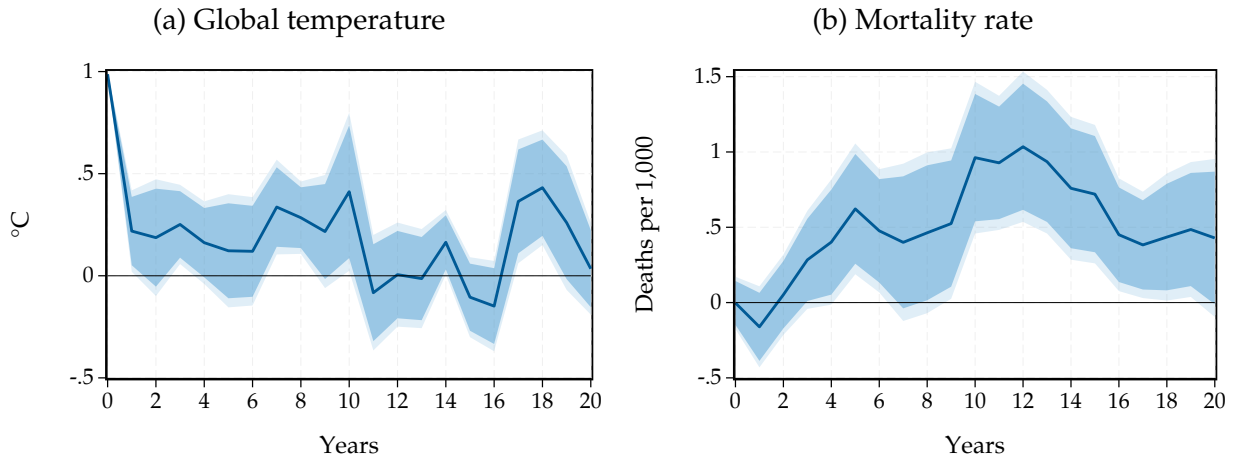
Our specification differs from standard panel studies of local temperature or heat in one critical respect: the identifying variation in global temperature shocks T_t^{shock} is common across all countries. Consequently, we cannot include year fixed effects, as these would absorb all variation in global temperature. Instead, we control for global economic shocks that may correlate with both temperature and mortality. Specifically, $\mathbf{x}_t$ includes lags of the global temperature shock, lags of world GDP growth, commodity price inflation, and global recession indicators.[3] At the country level, $\mathbf{x}_{i,t}$ includes lags of the country's own mortality rate, local GDP growth, and regional linear trends. We include two lags of the shock and four lags of all other control variables. We weight regressions by country population in period $t-1$ to reflect that larger countries contribute more to global mortality outcomes.

Because global temperature variation is common across countries, the error terms $\varepsilon_{i,t+h}$ are potentially serially and cross-sectionally correlated. We therefore compute standard errors using the approach of Driscoll and Kraay (1998), which is robust to both serial and cross-sectional dependence of arbitrary form.

flexible control specifications.

[3] We include a dummy for global recession dates as well as some lags to allow for persistent effects of past global economic contractions on mortality.

Figure 4: The Effect of Global Temperature Shocks on Temperature and Mortality



Notes: Impulse responses of global mean temperature and mortality rate (deaths per 1,000 individuals) to a 1°C global temperature shock, estimated using panel local projections in equation (2). Sample: 1960-2019. Controls: four lags of global temperature shock, world GDP growth, global recession indicators, commodity price inflation, local mortality rate, local GDP growth, regional linear trends, and subregion-specific linear trends. Shaded areas: 90% and 95% confidence bands based on Driscoll-Kraay standard errors.

3 Global Temperature and Mortality

We now present our main empirical results on the impact of global temperature shocks on mortality. We first report our baseline estimates before discussing three identification concerns: omitted variable bias, reverse causality, and external validity.

3.1 Baseline Estimates

Figure 4 displays the impulse responses of global temperature and world mortality to a global temperature shock scaled to 1°C. The left panel shows that by construction, global temperature increases by 1°C on impact. The effect is highly persistent: after 10 years, global temperature remains elevated by nearly 0.5°C, before gradually decaying over time.

The right panel shows the mortality response. Mortality responds with a meaningful lag, and the effect builds up over time. The peak effect occurs 10 years after the shock, when mortality is elevated by approximately 1 deaths per 1,000. The effect remains statistically significant at the 5% level in years 5 and 6, and then again in years 9 to 17. The point estimates eventually revert back to zero over time.

The mortality response is not uniform by age. Appendix Figure B.1 indicates that individuals aged 65 years or more exhibit much stronger responses that peak near 5 excess

deaths per 1,000 five years after the global temperature shocks.

The gradual increase in mortality represents the accumulating effects of persistently elevated temperature over time: the fundamental determinant of mortality is the level of temperature; the temperature shock only serves as an instrument. Hence, the effects are larger than implied by a fully transitory 1°C temperature increase, but are lower than the effect of a permanent 1°C increase.

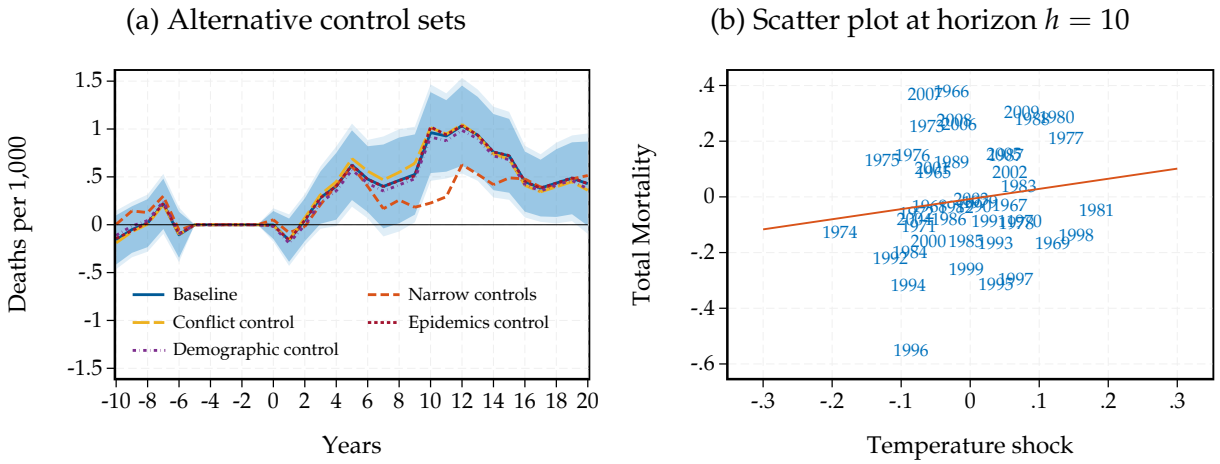
We assess the role of temperature persistence by constructing a counterfactual mortality path that corresponds to a fully transitory temperature change using the method of Sims (1986). Appendix B.2 shows that accumulated temperature persistence accounts for a substantial portion of the peak effect: under a purely transitory shock, the peak mortality effect would be approximately 0.5 deaths per 1,000.

To quantify the long-run impact of a permanent temperature increase, we calculate the ratio of the cumulative impulse response of mortality to the cumulative impulse response of temperature. This ratio corresponds to the steady-state effect of a sustained 1°C rise in global temperature. We find that a permanent 1°C increase in global temperature leads to a long-run increase of approximately 2.5 deaths per 1,000 population. This magnitude coresponds to twice the current global mortality burden from all forms of air pollution combined.

Scaling our estimates to a permanent 1°C increase relies on two conditions that we share with the literature on climate impacts. First, that the relationship between temperature and mortality holds beyond the range of shocks observed in our sample. A 1°C shock is large relative to historical variation—we observe shocks of 0.1-0.2°C in practice. However, among the shocks we do observe, we find limited evidence of non-linearities. Of course, future climate change may involve adaptation, tipping points or other non-linearities that have not yet materialized (Dietz et al., 2021b), which could make actual impacts smaller or larger than our linear exercise suggests. The second condition is that permanent changes are well approximated by the cumulative effects of partially transitory shocks. This condition is more difficult to verify empirically, and is also subject to the same caveats about adaptation and non-linear impacts.

While we estimate effects using country-level mortality data, our identifying variation comes from global temperature shocks common to all countries. This time-series variation warrants additional care in interpreting our results.

Figure 5: Sensitivity to Control Variables and Specification Choices



Notes: Panel (a): impulse responses under alternative control specifications. Baseline: four lags of mortality, global and local GDP growth, commodity price inflation, and global recession dummy, two lags of temperature shocks, and subregion trends. Narrow controls: excludes commodity price inflation and global recession dummy. Conflict control: number of conflicts a country is involved in a given year. Epidemics control: adds indicator for pandemic and epidemic outbreaks. Demographic control: population share 65 years and over. Panel (b): mortality change from year -1 to year $+10$ against temperature shock for each country-year observation, with regression line from baseline specification. Shaded areas: 90% and 95% confidence bands.

3.2 Omitted Variable Bias

The first identification concern is that global temperature shocks may coincide with other global or regional shocks affecting mortality. For instance, if global or regional mortality shocks just so happen to be correlated with global temperature variation over time, we would mistakenly attribute these effects to temperature.

To address this concern, we already include a rich set of global and country-level controls in our baseline specification. Figure 5(a) shows that the results are robust to alternative control sets. First, the results hold true we restrict attention to a narrower set of global controls, excluding commodity price inflation and the global recession indicator. Second, we show that results are robust when we additionally control for demographics, armed conflicts, or epidemics. Across all these specifications, estimates remain similar and largely overlap with the baseline confidence intervals.

We also examine the evolution of mortality prior to the shock to assess the presence of pre-trends. In particular, the concern is that the estimated effects may capture pre-existing trends rather than a causal response to the shock. Reassuringly, we find no evidence of economically meaningful or statistically significant pre-trends across specifications. Estimates are close to zero throughout the pre-treatment period.[4]

[4]Note that the coefficients for horizons $h = -(p+1), \ldots, -1$ are mechanically equal to zero due to the

We further assess whether our results are driven by spuriously correlated mortality shocks by examining the influence of individual years. For each year, Panel 5(b) plots the 10-year-ahead change in mortality against the contemporaneous temperature shock, after residualizing both variables with respect to the baseline controls. The positive relationship does not appear to be driven by a small number of influential years or outliers.

We conduct two additional exercises in Appendix C. First, Figure C.2 reports a jackknife exercise in which we sequentially set each temperature shock to zero and re-estimate the impulse responses. The resulting estimates are highly stable, indicating that our findings are not driven by any single shock. Second, Figure C.3 verifies that our results are robust to alternative methods of constructing temperature shocks, including HP filtering and different choices of the Hamilton filter parameters (h, p). Overall, these results suggest that our estimates are unlikely to be driven by mortality shocks spuriously correlated to temperature shocks nor by the influence of specific countries or specification choices.

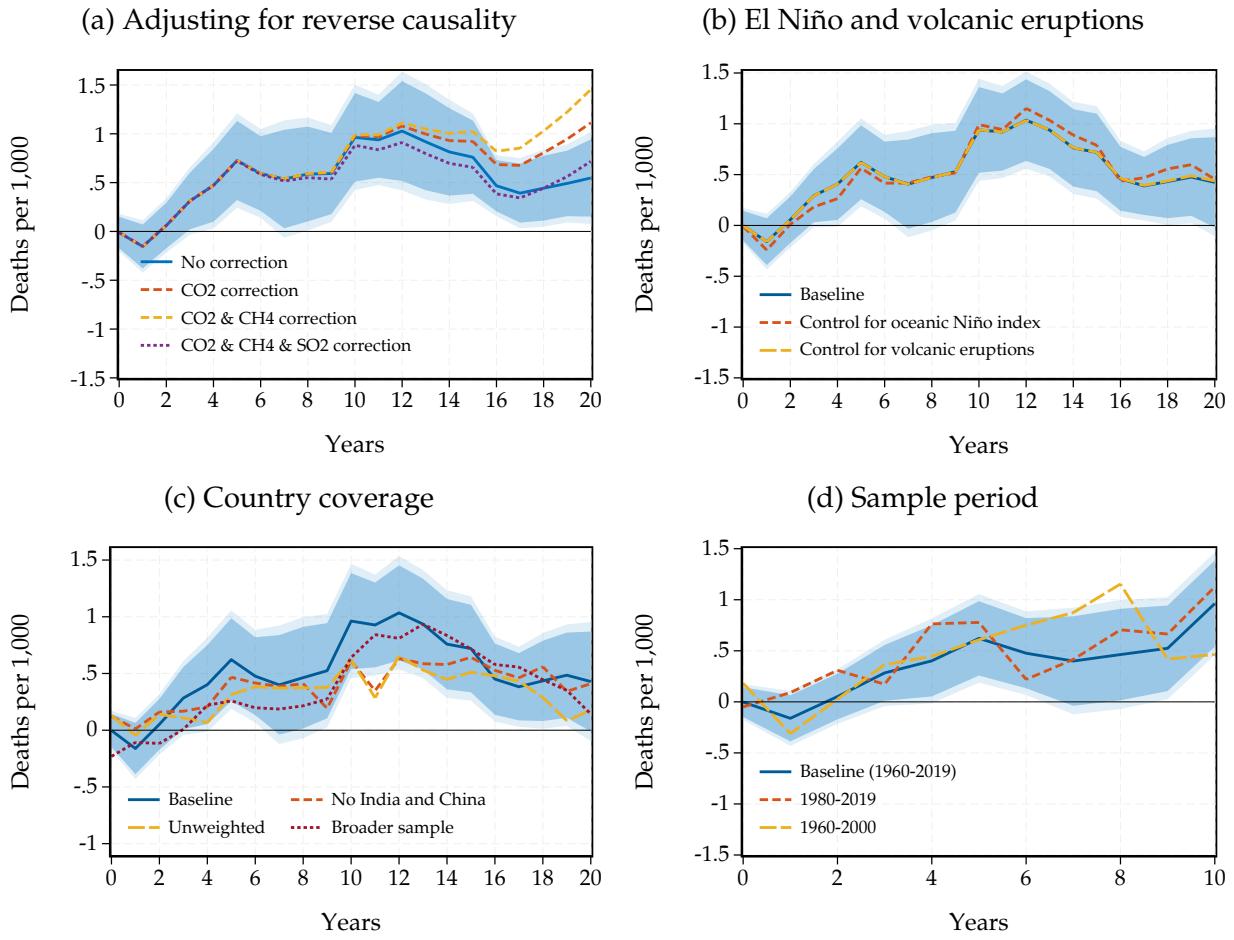
3.3 Reverse Causality

The second identification concern is reverse causality. Global temperature shocks reduce economic activity (Bilal and Känzig, 2026), and thus greenhouse gas emissions. Reductions in emissions alter the future evolution of temperature, thereby potentially biasing our estimates. To address this concern, we follow Bilal and Känzig (2026) and adjust the estimated impulse response function to remove the component attributable to anthropogenic emissions. To do so, we separately estimate the impact of global temperature on GDP per capita. We use these estimates and standard emissions elasticities to obtain the dynamic response of emissions to global temperature shocks. We then use a simple climate model to predict the temperature contribution from observed CO_2, CH_4, and SO_2 emissions over our sample. Finally, we adjust our estimates by removing the greenhouse gas feedback between global temperature shocks and subsequent temperature in a bivariate analysis.

Figure 6(a) shows that adjusting for emissions has minimal impact on our estimates up to horizon 15. Thereafter, the adjustment is more consequential. Adjustments are more pronounced at later horizons because of three sources of delays. First, GDP impacts of global temperature—and hence, emissions—achieve a first peak after at least 5 years and effects remain large up to 20 years. Second, carbon dioxide warming effects take 10 years

inclusion of lagged controls in the local projection specification.

Figure 6: Reverse Causality and External Validity



Notes: Impulse response of mortality to a 1°C global temperature shock. Panel (a): impact of temperature shocks after removing predicted temperature contribution from CO_2, CH_4, and SO_2 emissions using a simple climate model. Panel (b): impulse responses using only temperature shocks coinciding with El Niño events or volcanic eruptions versus all shocks. Panel (c): impulse responses based on different samples of countries and weighting schemes. Broader sample includes mortality data from UNWPP. Panel (d): impulse responses estimated over different sample periods. Full sample: 1960-2019. Early period: 1960-1989. Late period: 1990-2019. Shaded areas: 90% and 95% confidence bands.

to peak.

Carbon dioxide and methane are greenhouse gases and therefore raise global temperatures. To the extent that economic activity increases emissions of these gases, reverse causality would attenuate our estimated temperature effects, biasing them downward rather than upward. By contrast, sulfur dioxide is a major anthropogenic pollutant with a cooling effect, implying an offsetting bias in the opposite direction. Consistent with these mechanisms, our adjustments for each gas have the expected sign. Taken together, however, they largely offset one another, and the impulse response function that accounts for all three gases is almost identical to the baseline.

3.4 External Validity

The third concern is external validity. Do our estimates vary across different types of temperature shocks, geographic coverage, or time periods? For example, if the mortality–temperature relationship has changed over time due to adaptation, estimates based on earlier decades may not generalize to more recent or future climate conditions. Similarly, geographic coverage may matter: if countries that are highly exposed or vulnerable to climate change—such as many African economies—are underrepresented in our baseline sample, our estimates may not capture the full extent of temperature-related mortality effects. Finally, if the relationship depends on the underlying source of temperature variation, our average estimates may not be broadly applicable.

We begin by examining whether the type of temperature variation matters. Figure 6(b) compares our baseline estimates to specifications that control for major sources of global temperature fluctuations, including El Niño events and volcanic eruptions. The impulse response functions are virtually identical across specifications. This suggests that the estimated mortality effects reflect temperature changes broadly, rather than factors specific to particular climate phenomena.

Next, we assess the role of country composition. Panel (c) shows that our results are robust to alternative country samples. We begin by expanding the sample using UNWPP mortality data to include a broader set of countries, particularly in Africa. While this entails a trade-off between external validity and data quality, the resulting estimates are very similar to the baseline. We then examine whether the results are driven by large countries. Excluding India and China—two countries that extend our sample relative to earlier work and receive considerable weight in the baseline—produces comparable estimates, with slightly lower peak effects of around 0.6 deaths per 1,000. Consistent with this, estimating unweighted local projections produces very similar results. Taken together, these findings indicate that our estimates are robust to broader geographic coverage and are not driven by large countries or weighting choices.

Finally, Panel (d) examines stability over time by comparing the full sample (1960–2019) to subsamples covering the first four decades (1960–2000) and the most recent four decades (1980–2019). Given our limited sample, we necessarily lose precision when splitting the sample and restricting the analysis to 40 years. For this reason, we report the first ten years of the impulse response function rather than twenty years as in the baseline estimation. We find that the impulse response functions remain close across samples.

Third, we verify that our results are not specific to the major countries—India and China—that we include in our longer sample relative to earlier work. Figure C.1 shows that excluding either India or China leads to comparable estimates that peak somewhat lower at approximately 0.6 deaths per 1,000.

4 Global Temperature and Local Heat

How do our estimates based on global temperature fluctuations compare with conventional estimates based on variation in local exposure to extreme heat? We answer this question in three steps. First, we estimate mortality responses to local extreme heat in our sample with a comparable empirical specification and find that they are smaller than global temperature impacts. Second, we show that additional extreme events that correlate with global temperature fluctuations cannot account for the difference.

4.1 Mortality Effects of Local Heat

To maximize comparability with our global temperature estimates, we estimate mortality responses to local extreme heat using the following empirical specification:

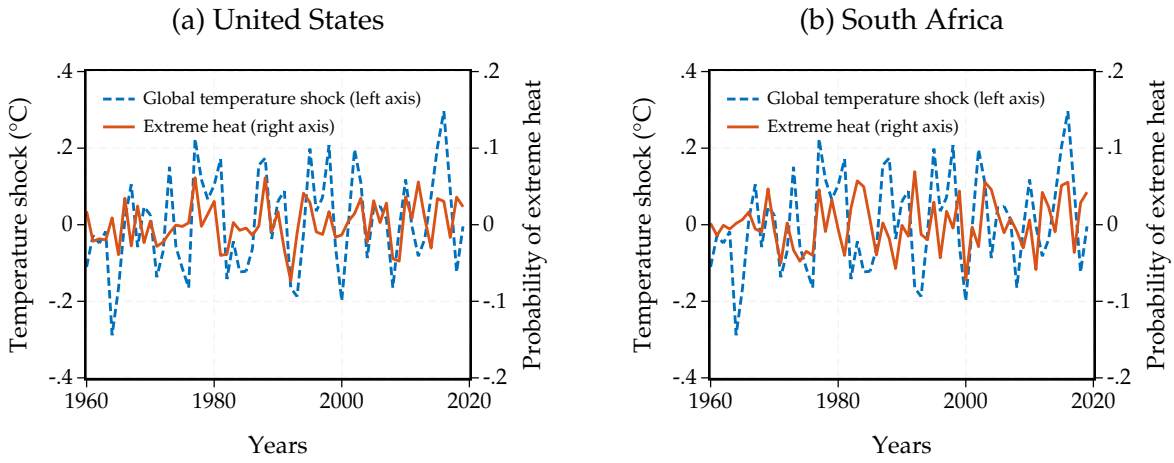
$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \delta_{t,h} + \theta_h^{\text{heat}} H_{i,t}^{\text{shock}} + \mathbf{x}_{i,t}' \boldsymbol{\beta}_h + \varepsilon_{i,t+h}. \tag{3}$$

The panel local projections specification in (3) is similar to (2), but features several important distinctions.

The first difference is that we replace global temperature shock with a national measure of exposure to extreme heat for country i in year t. To construct exposure to extreme heat, we first define average daily national temperature in country i. We then define exposure to extreme heat $H_{i,t}$ as the fraction of days within year t and country i when daily temperature exceeds a national threshold. The threshold is such that extreme heat exposure $H_{i,t}$ has an average value of 0.05 during the historical period 1950-1980. Finally, we construct extreme heat shocks $H_{i,t}^{\text{shock}}$ by applying the same Hamilton (2018) filter to extreme heat exposure $H_{i,t}$.

The second difference is that we include year fixed effects $\delta_{t,h}$ to follow standard practice and control for common global variation in mortality (Deschênes and Greenstone, 2011; Carleton et al., 2022). In fact, our results based on local heat are largely independent

Figure 7: Local Extreme Heat Shocks



Notes: Local extreme heat shocks constructed using Hamilton filter ($h = 2$, $p = 2$) applied to the fraction of days with temperature above the 95th percentile of the 1950-1980 distribution. Left: United States. Right: South Africa.

from the inclusion of year fixed effects.

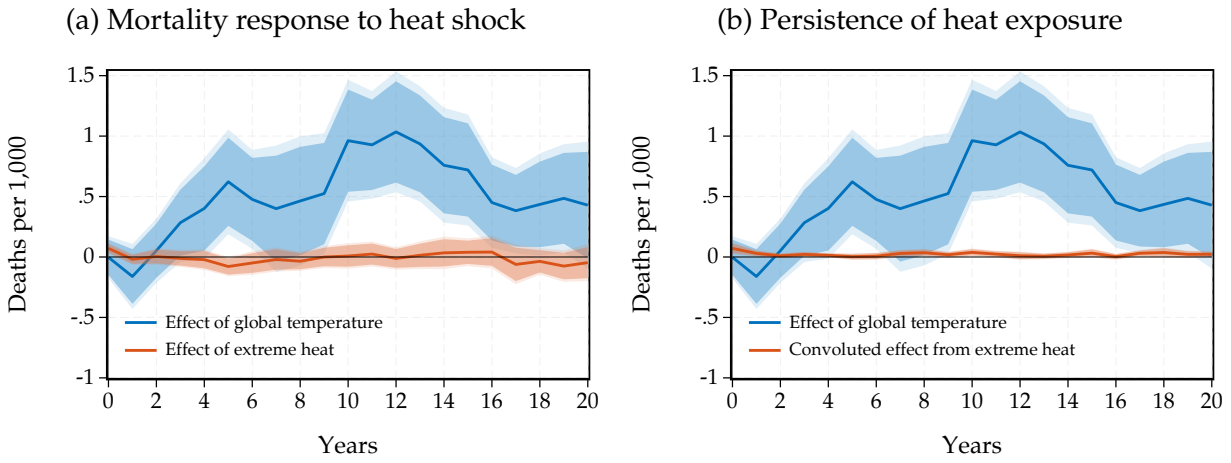
Figure 7 illustrates the variation in local heat shocks for selected countries. These local shocks are only weakly correlated with global temperature fluctuations. For example, a country may experience extreme heat in a given year without a corresponding global temperature shock, and conversely. This pattern reflects the fact that local heat fluctuations are driven by regional weather dynamics that are largely distinct from movements in global mean temperature.

Figure 8 presents the impact of local heat on mortality. Panel (a) shows the impulse response of mortality to local extreme heat shocks. The underlying heat shock is scaled to coincide with the peak response of extreme heat to global temperature shocks, reported in Appendix Figure D.1. This peak response represents a 0.13 rise in the fraction of extreme heat days, from 0.05 on average to 0.18.

We find that rising extreme heat causes mortality to rise. The contemporaneous effect stands at 0.1 excess deaths and is statistically significant at the 5% level. However, the effect decays quickly and becomes statistically insignificant from the second year onward. The effect of local heat on mortality is an order of magnitude smaller than the mortality response to global temperature shocks, which we reproduce in the Figure for comparability.

Is the difference between global temperature and local heat due to specification choices? As dicussed above, we scale the underlying heat shock to match the peak re-

Figure 8: Mortality Impacts of Local Heat and Global Temperature



Notes: Panel (a): impulse response of mortality to a local extreme heat shock that raises the heat index (fraction of days above 95th percentile) from 0.05 to 0.25, representing a five-fold increase in extreme heat frequency. Panel (b): cumulative impulse response showing the persistence of heat exposure following a heat shock. Based on the ratio of cumulative impulse responses, a permanent 1°C increase operating through local heat alone would increase mortality by approximately 0.1 deaths per 1,000, compared to 2.5 for global temperature. Shaded areas: 90% and 95% confidence bands.

sponse of extreme heat to global temperature shocks. Therefore, the difference in mortality responses cannot be directly attributed to different instantaneous magnitudes of the underlying treatments. However, the timing of the underlying treatments may differ: the dynamic response of local heat to a local heat shock may differ from the dynamic response of local heat to a global temperature shock.

Panel 8(b) imposes the same persistence of the underlying treatment using the method from Sims (1986). In this exercise, we convolve the impulse response of mortality to local heat shocks with the impulse response of local heat to global temperature shocks. The resulting mortality response captures what would occur if local heat exposure followed the same magnitude and persistence as under global temperature shocks. After this adjustment, the mortality response to local heat remains an order of magnitude smaller than the response to global temperature shocks.

Is the difference between global temperature and local heat due to how we construct local heat exposure? Existing work often uses local high-degree-day exposure—exposure to days with temperature above a certain threshold—rather than the country-level fraction of extreme heat days that we use (Deschênes and Greenstone, 2011; Carleton et al., 2022). To compare our estimates to existing work, we convert them into the mortality effect of a permanent 1°C temperature increase operating solely through local extreme heat. To do so, we take the ratio of the cumulative impulse response of mortality to local

heat, and local heat to global temperature.

We obtain that a permanent 1°C warming that operates solely through local extreme heat increases mortality by approximately 0.1 deaths per 1,000 annually. This estimate is quantitatively consistent with previous work: Carleton et al. (2022) also obtain 0.1 deaths per 1,000 individuals for 1°C warming. This estimate is an order of magnitude below the global temperature impact of 2.5 deaths per 1,000.

4.2 Extreme Weather and Mortality

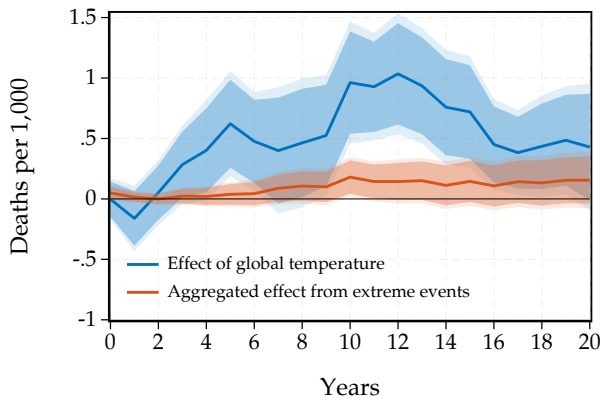
Global temperature shocks may affect mortality not only through local heat but also through other extreme weather events. The geoscience literature indicates that global temperature—in particular ocean temperature—determines the frequency and intensity of multiple climatic phenomena including droughts, extreme precipitation, and extreme wind (Seneviratne et al., 2016; Wartenburger et al., 2017; Seneviratne et al., 2021; Domeisen et al., 2023).

We examine whether this extreme event channel can account for the gap between local heat and global temperature impacts. Following Bilal and Känzig (2026), we construct exposure indices for four extreme weather events: extreme heat (national temperature above 95^{th} percentile, as in Section 4.1), extreme cold (national temperature below 5^{th} percentile), drought (national precipitation below 10^{th} percentile), and extreme wind (national wind average above 99^{th} percentile).

Figure D.1 in Appendix D.1 displays the impulse responses of these extreme event indices to global temperature shocks. Consistently with the geoscience literature, we find that global temperature shocks increase the frequency of extreme heat and cold, droughts, and extreme wind, with the latter three imprecisely estimated.

To assess whether the increase in local extreme events associated with global temperature changes can account for the gap between global and local temperature estimates, we aggregate local effects in three steps. First, we estimate the direct mortality impact of each extreme event by regressing mortality on extreme event indices, controlling for the same set of covariates as in our baseline specification. We estimate these impacts jointly to account for correlation between events. Figure D.2 in Appendix D.1 displays the marginal effect of each extreme event. Extreme heat and extreme cold increase mortality on impact, but has satistically insignificant effects later on. Droughts and extreme wind increase mortality persistently.

Figure 9: Mortality Effect of Extreme Events



Notes: Comparison of direct mortality response to global temperature shocks (blue solid line) with indirect estimate based on aggregating extreme events (red dashed line). The aggregation accounts for extreme heat, extreme cold, droughts, extreme precipitation, and extreme wind. Extreme events explain approximately 20% of the total mortality effect of global temperature, indicating that other channels must be operating. Shaded areas: 90% and 95% confidence bands.

Second, we construct the indirect mortality effect of global temperature operating through extreme events. We interact the impulse response of each extreme event to global temperature shocks from Figure D.1 with the mortality impact of that extreme event from Figure D.2. This interaction gives us the mortality effect that would occur if global temperature affected mortality solely through its impact on each specific extreme event.

Third, we aggregate across all extreme events, accounting for their dynamic correlation with global temperature shocks. Specifically, we compute: $\theta_h^{\text{extremes}} = \sum_X \sum_{s=0}^{h} \phi_s^X \psi_{h-s}^X$, where ϕ_s^X is the response of extreme event X to a global temperature shock at horizon s, and ψ_{h-s}^X is the mortality impact of a transitory increase in extreme event X at horizon $h-s$. The sum over X includes extreme heat, extreme cold, droughts, extreme precipitation, and extreme wind.

Figure 9 presents the results. The blue line shows our baseline direct estimate of mortality responding to global temperature shocks. The red line shows the indirect estimate based on aggregating extreme events. We find that extreme events can account for approximately 0.1-0.2 deaths per 1,000 at long horizons. They represent less than one quarter of the cumulative effect of global temperature. While this effect is not negligible, it falls short of accounting for the full gap between global temperature and local heat estimates.

4.3 The Income Channel

Why does global temperature affect mortality more than local heat and other extreme weather events? In this section, we ask whether macroeconomic disruptions caused by global temperature fluctuations increase mortality due to a link between mortality and income. We first propose two additional sets of results suggesting that income may play an important role. Second, we assess the role of income quantitatively by combining separate estimates of the impact of temperature on income and of the impact of income on mortality.

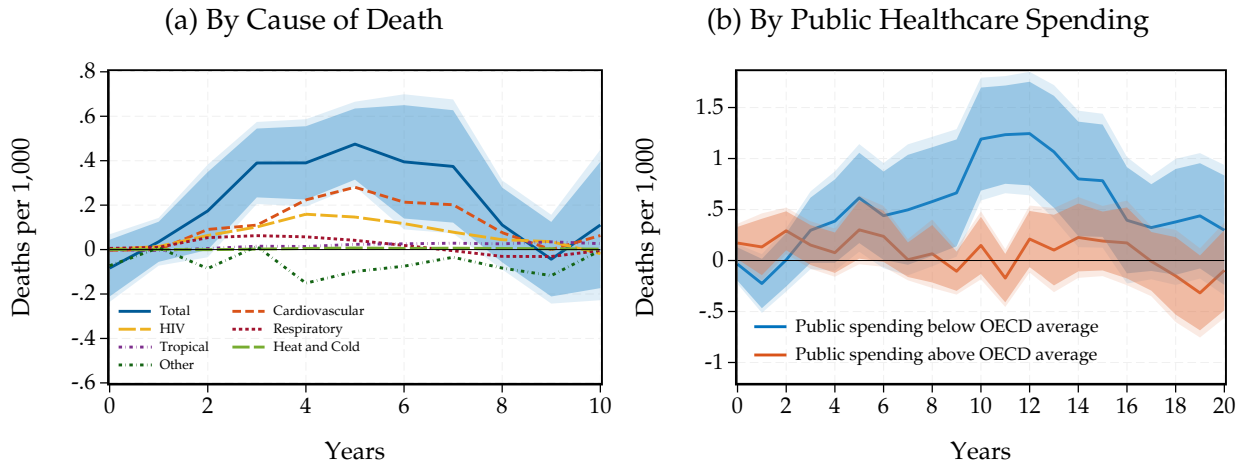
Heterogeneity by cause of death and country. We discuss two additional sets of results that support the interpretation that global temperature shocks lead to an increase in mortality because of associated income losses. First, we show that deaths following global temperature shocks are primarily cardiovascular- and HIV-related, not heat- and cold-related. Second, we show that countries with a lower fraction of public health care spending display larger mortality responses.

We first describe which causes of death are responsible for the estimates response of mortality to global temperature. For this purpose, we use data from the Institute for Health Metrics and Evaluation covering 31 years and 198 countries. Figure 10(a) starts by reproducing our estimates in this dataset. The estimates in this different dataset are similar to our baseline in Figure 4, although of course they do not exactly coincide. We limit our analysis to a horizon of 10 years due to the shorter sample.

Figure 10(a) then disaggregates the rise in the crude death rate by cause of death. Cardiovascular-related reasons are the leading cause of death following global temperature shocks. As for the overall impulse response function, these deaths peak several years after the shock. The magnitude exceeds the estimated impact of heat on mortality from Figure 8 by an order of magnitude, suggesting that most of these cardiovascular excess deaths are not due to heat exposure.

HIV-related deaths are the second most important cause of death after global temperature shocks. As for cardiovascular deaths, they peak several years after the shock occurs. It is not obvious how to directly link HIV deaths to extreme weather exposure. By contrast, HIV treatments are expensive and can be tightly linked to worldwide income. For instance, the United States President's Emergency Plan for AIDS Relief (PEPFAR) has historically funded antiretrovial therapy for two thirds of all antiretrovial patients worldwide, which themselves represent three quarters of all infected patients.

Figure 10: Mortality Impacts By Cause of Death and By Public Healthcare Spending



Notes: Panel (a): Mortality impacts of global temperature shocks by cause of death, deaths per 1,000. Data: Health Metrics and Evaluation, covering 31 years and 198 countries. Solid blue and shaded areas: point estimates, 90 and 95% confidence bands. Dashed and dotted lines: point estimates by cause. Panel (b): Mortality impacts of global temperature shocks by public spending as a fraction of total healthcare spending, deaths per 1,000. Solid lines: point estimates. Shaded areas: 90 and 95% confidence bands.

By contrast, heat- and cold-related deaths represent a negligble fraction of overall deaths. This result is consistent with those in Figure 8, showing that direct effects of extreme heat exposure is small relative to the overal effects of global temperature shocks.

Next, we describe which countries display stronger mortality responses to global temperature shocks. Many countries in Europe and Oceania have near-universal health care systems that offer cheap access for individuals, and may thus be less dependent on the global economy. By contrast, the United States' health care system is more reliant on employer-based private insurance and a sizable fraction of the population is not insured, thereby relying more directly on individual income. Similarly, in lower-income countries, public health care coverage tends to be more limited and out-of-pocket costs are more important.

We first estimate heterogeneous effects by broad region of the world. Appendix Figure D.3 reveals that Europe and Oceania exhibit weak and largely statistically insignificant mortality responses to global temperature shocks. North America, Asia and Latin America display mortality responses that are all close to the world average and peak between 1 and 2 excess deaths per 1,000. Despite data limitations, Africa exhibits the largest mortality response of all regions, that peaks above 2 excess deaths per 1,000.

Second, we directly test whether countries with higher shares of public health care spending exhibit different responses to global temperature shocks. We classify coun-

tries as high public spending in health care if their fraction of public spending in total healthcare expenditure exceeds the OECD average value of 72%. Figure 10(b) shows that countries with a share of public health care spending above the OECD average exhibit no significant response of mortality to global temperature shocks. By contrast, countries with public health care spending below the OECD average exhibit larger and statistically significant mortality responses to global temperature shocks.

That the mortality impacts from global temperature shocks transmit primarily through causes of death and countries where healthcare are income-sensitive is qualitatively consistent with a transmission channel through worldwide income. Next, we assess whether this mechanism is also plausible quantitatively.

GDP losses and mortality. We now investigate quantitatively whether global temperature affects mortality primarily through its impact on economic output and income. Bilal and Känzig (2026) document that global temperature shocks cause large and persistent declines in global GDP per capita. These estimates imply that a permanent increase in global mean temperature of 1°C leads to a 22% reduction in global GDP. If income losses in turn increase mortality—as suggested by the development and health literature (Preston, 1975; Cutler, Deaton, and Lleras-Muney, 2006)—then this income channel may help rationalize the effect of global temperature on mortality.

To assess this mechanism quantitatively, we construct an indirect measure of the effect of global temperature on mortality through income. We interact two components. The first component is the impact of global temperature shocks on global GDP per capita. Appendix Figure D.4 reproduces the analysis from Bilal and Känzig (2026) in our sample of countries, with comparable controls to our mortality specification and up to a 20 year horizon. We obtain similar results.

The second component is the effect of GDP on mortality. We address two challenges in constructing this component. The first challenge is to chose the relevant measure of GDP. Is country-level GDP or worldwide GDP relevant for mortality across countries? A large literature going back to Preston (1975) suggests that mortality reductions over time arise partly because of country-level income growth. However, most of mortality reductions appear to take place independently from country-level income changes, but may instead be associated with income changes at the economic frontier, and thus might be more tightly linked to worldwide income changes. For this reason and for consistency with the estimates in Bilal and Känzig (2026), we focus on global GDP as our main mediating

variable.

The second challenge is to obtain an estimate of the effect of global GDP on mortality. We use the estimates from Bilal, Känzig, and Lisiecki (2026) who leverage global GDP fluctuations in the time series. They show that global GDP shocks—unconditionally and when instrumented with U.S. monetary policy surprises to address confounding forces such as epidemics or geopolitical events—lead to larger long-run mortality elasticities than local GDP shocks. Their estimates imply that a 100 log point increase in global GDP per capita reduces mortality by approximately 6 deaths per 1,000 in the long run. For reference, we reproduce their estimates in Appendix Figure D.5.

With these estimates at hand, we construct the indirect mortality response to a global temperature shock by convolving the GDP response to temperature with the mortality response to GDP:

$$\theta_h^{\text{indirect}} = \sum_{s=0}^{h} \phi_s^{\text{GDP}} \psi_{h-s}^{\text{mort}\,|\,\text{GDP}}. \tag{4}$$

Here, ϕ_s^{GDP} is the response of world GDP per capita to a global temperature shock at horizon s from Figure D.4. The coefficients $\psi_{h-s}^{\text{mort}\,|\,\text{GDP}}$ are not estimated directly: they represent the mortality impact of a transitory GDP shock at horizon $h-s$ as implied by our GDP shock or monetary policy IV estimates.[5]

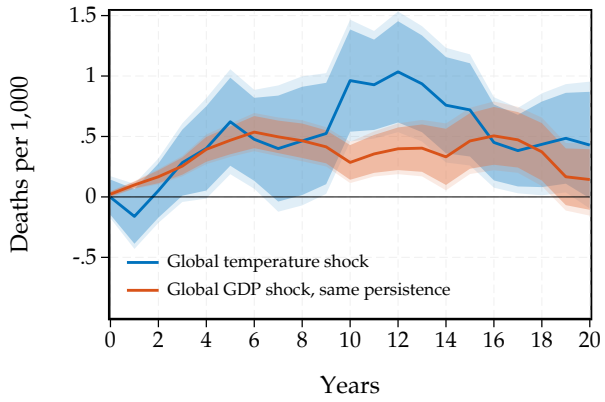
Figure 11 presents the results. The indirect estimate tracks the direct temperature–mortality response closely. Using the full dynamic convolution, global GDP losses account for over 70% of the cumulative mortality impact of global temperature shocks. This result is robust across both the IV-based and reduced-form GDP specifications.

We also assess the additional contribution of extreme weather channels—including heat, cold, droughts, precipitation, and wind—beyond income-mediated effects that are partly captured in the global GDP contraction that follows a global temperature shock. These direct channels account for only about 5 percentage points of the total mortality response.

Overall, these results indicate that global temperature shocks have large mortality effects primarily because they generate persistent global income losses. Estimates based solely on exposure to extreme heat substantially understate the overall mortality conse-

[5] In principle, this linear system admits a solution that recovers $\psi_{h-1}^{\text{mort}\,|\,\text{MP}}$ as in Sims (1986). In practice, the weak response of GDP per capita to monetary policy shocks at early horizons leads to an ill-conditioned linear system. Thus, we use a Ridge regression to regularize the inversion in our monetary policy shock specification. See Appendix Figure **??** for results under alternative Ridge regularization parameters.

Figure 11: Indirect effect of global temperature on mortality through global GDP



Notes: Solid blue line: direct estimate θ_h of the effect of global temperature on mortality. Solid red line: indirect estimate $\theta_h^{\text{indirect}}$ of the effect of global temperature on mortality, mediated by global GDP per capita. Shaded areas: 90 and 95% confidence bands based on Driscoll-Kray standard errors.

quences of warming because they omit mortality impacts mediated by global income.

5 A Model of Climate Change, Mortality and the Economy

To study the long-run economic and health consequences of climate change, we develop a model of climate change, mortality, and the economy. Our framework closely follows the standard neoclassical growth model that underpins integrated assessment models in the spirit of Nordhaus (1992). Relative to Bilal and Känzig (2026), we extend the model by introducing stochastic lifespans and an explicit valuation of survival. We discipline the model using our reduced-form estimates of the impact of global temperature shocks on output and mortality, which we use to structurally identify temperature damage functions.

5.1 Model Description

Setup. Time is continuous and runs forever. The economy is populated by a unit continuum of households who face stochastic lifespans, as in Blanchard (1985). Each household dies with Poisson arrival rate M_t and is replaced by a single offspring valued by the parent. Households have constant relative risk aversion flow preferences $U(C) = \frac{C^{1-\gamma}-1}{1-\gamma}$ and discount the future at rate ρ. In addition to flow utility from consumption, households value being alive. Let $\overline{V}$ denote the value of the end of lifespan.

A household born with initial assets K_0 solves

$$V_0(K_0) = \max_{\{C_t, K_t\}_t} \mathbb{E}_0 \left[\int_0^\tau e^{-\rho t} U(C_t) dt + e^{-\rho \tau} (V(K_\tau) - \overline{V}) \right] \tag{5}$$

subject to

$$C_t + \dot{K}_t = w_t + r_t K_t, \tag{6}$$

where τ denotes the stochastic time of death and M_t is the mortality rate. Labor supply is exogenous and normalized to $L_t = 1$.[6]

Firms produce according to a Cobb–Douglas production function in capital K_t^D and labor L_t^D with time-dependent total factor productivity Z_t: $Y_t = Z_t (K_t^D)^\alpha (L_t^D)^{1-\alpha}$. They hire labor and rent capital in competitive factor markets and solve

$$\max_{K_t^D, L_t^D} Z_t (K_t^D)^\alpha (L_t^D)^{1-\alpha} - (r_t + \delta) K_t^D - w_t L_t^D. \tag{7}$$

Capital depreciates at constant rate δ, covered by firms.

Climate change and damages. We model climate change as changes in productivity Z_t and mortality M_t over time, relative to their baseline values Z_0 and M_0. This representation follows existing work (Nordhaus, 1992; Moore and Diaz, 2015; Carleton et al., 2022) and parsimoniously but comprehensively captures the combined economic and health consequences of excess global temperature documented in the data.

We take the path of global mean temperature T_t relative to a reference level T_0 as given, and denote by $\widehat{T}_t \equiv T_t - T_0$ the path of excess temperature. Global mean temperature affects productivity and mortality through structural damage functions $\{\zeta_s, \xi_s\}_{s \geq 0}$:

$$Z_t = Z_0 \exp\left(\int_0^t \zeta_s \widehat{T}_{t-s} ds \right) \tag{8}$$

$$M_t = M_0 + \int_0^t \xi_s \widehat{T}_{t-s} ds. \tag{9}$$

The structural damage functions $\{\zeta_s, \xi_s\}$ govern the persistence of the effects of transitory global temperature shocks on productivity and mortality. When temperature is constant at $T_t \equiv \overline{T}$, the economy converges to a steady state with corresponding pro-

[6] We adopt this dynastic formulation for analytical convenience. For our purposes, it is equivalent to alternative formulations based on annuities or flow utility shifts.

ductivity and mortality levels determined by the cumulative damage functions, $\overline{Z} = Z_0 \exp\left((\overline{T} - T_0)\int_0^\infty \zeta_s ds\right)$ and $\overline{M} = M_0 + (\overline{T} - T_0)\int_0^\infty \xi_s ds$. The cumulative damage functions $\int_0^\infty \zeta_s ds$ and $\int_0^\infty \xi_s ds$ thus govern the long-run effects of permanent warming.

Competitive equilibrium. A competitive equilibrium is a collection of sequences $\{C_t, K_t, r_t, w_t, M_t, Z_t\}_{t=0}^\infty$ such that households and firms optimize given prices, mortality and productivity evolve according to the temperature damage functions, and factor markets clear: $K_t = K_t^D$ and $1 = L_t^D$.

We abstract from feedback between economic activity, emissions, and temperature dynamics, as our focus is on climate damages. Absent environmental externalities, the competitive equilibrium is efficient as in the standard neoclassical growth model. As of now, we do explicitly not model feedback from income to mortality, although our empirical evidence suggests that income losses amplify mortality responses. We plan to incorporate this channel in the next iteration of the project, as it is central for optimal policy analysis. However, omitting it does not affect our positive evaluation of climate damages under a given temperature path.

Social Cost of Carbon. In our framework, we define the Social Cost of Carbon (SCC) as the one-time dollar amount $\mathcal{C}$ that households would pay at time 0 that would make them indifferent between a world with an additional ton of CO2 emitted at time 0, and a world starting in steady state, without emissions, but having paid $\mathcal{C}$.

Given that we do not model emissions directly, we must map a one-time CO2 pulse into a temperature path in order to calculate the SCC. We follow Folini et al. (2024) and use the temperature response of global mean temperature to a CO2 pulse from Dietz et al. (2021a) and Joos et al. (2013).

We denote by $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$ the path of excess warming implied by a one-time pulse of one ton of CO2 emitted at time 0. The median response in Dietz et al. (2021a) indicates that after a 1 gigaton pulse of carbon, temperature rises steadily and eventually stabilizes at 0.002°C after 15 years.[7] Our welfare numbers do not depend on the temperature response to a CO2 pulse, but instead on a particular warming scenario.

We then construct productivity and mortality paths $\{Z_t^{\text{SCC}}, M_t^{\text{SCC}}\}_{t\geq 0}$ according to equations (8)-(9) in which we use the temperature path $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$ rather than a global warming scenario. The model delivers a path of value functions $\{V_t^{\text{SCC}}(K)\}_{t\geq 0}$, equi-

[7] For brevity, we refer to the 'best fit CMIP5 ensemble' response in Dietz et al. (2021a) as 'median response'.

librium capital stocks $\{K_t^{\text{SCC}}\}_{t\geq 0}$ with initial condition $K_0^{\text{SCC}} = K^{ss}$, leading to a path of realized values $\{V_t^{\text{SCC}}(K_t^{\text{SCC}})\}_{t\geq 0}$, in response to this CO2 pulse-induced warming. Our definition requires that the SCC $\mathcal{C}$ be given implicitly by:

$$V^{ss}(K^{ss} - \mathcal{C}) = V_0^{\text{SCC}}(K^{ss}), \tag{10}$$

where *ss* superscripts denote initial steady-state quantities.

5.2 Estimation Strategy

Our next step is to estimate the structural damage functions $\{\zeta_s, \xi_s\}_{s\geq 0}$ by matching the reduced-form impulse response functions of output and mortality to global temperature shocks documented in Section 3.1. All non-damage parameters are set to standard values from the literature. Specifically, we set the risk-aversion to $\gamma = 1$, the capital share to $\alpha = 0.33$ and the annual capital depreciation rate to $\delta = 0.08$. Our choice of annual rate of time preference is $\rho = 0.02$, following Rennert et al. (2022), consistent with a 2% annual interest rate in steady-state.

Identification of mortality damages is immediate. Given the empirical impulse response of mortality to a temperature shock and the observed path of temperature following that shock, we directly recover the implied sequence of mortality shocks based on (9), and estimate $\{\xi_s\}$ as innovations to mortality consistent with that equation.

Identifying productivity damages is more involved because output responds both directly to productivity and indirectly through capital accumulation. Since the temperature shocks observed in the data are small and generate output fluctuations on the order of 1%, we log-linearize the model around its initial steady state for estimation. As shown in Bilal and Känzig (2026), there exists a sequence-space mapping from productivity shocks to output such that, to a first order,

$$\widehat{y}_t = \widehat{z}_t + \alpha \int_0^\infty \mathcal{K}_{t,s}\widehat{z}_s ds,$$

where $\widehat{z}_t$ and $\widehat{y}_t$ denote log deviations of productivity and output. The first term captures the direct effect of productivity, while the second term captures the equilibrium response through capital accumulation. Given the empirical output response to a temperature shock, we invert this mapping following Bilal and Känzig (2026) to recover the

implied sequence of productivity shocks $\{\widehat{z}_t\}$. We emphasize that log-linearization is used for identification only; all counterfactual exercises are conducted in the fully non-linear model. We then identify $\{\zeta_s\}$ as the innovations to productivity consistent with equation (8).

This approach is consistent with rational expectations. After a temperature shock, agents anticipate that temperature remains persistently elevated, as observed in the data. Thus, the identified damage functions are disciplined by impulse responses to shocks that are themselves persistent, rather than purely transitory. Importantly, however, our estimation takes into account the internal persistence of realized temperature.

In practice, we face two challenges. First, the empirical impulse responses are observed only over a finite horizon, whereas the inversion requires the entire path. Second, unconstrained estimation could imply implausible long-run effects of transitory temperature shocks. We address both issues by imposing smooth parametric forms for the structural damage functions. Specifically, we constrain $x_s = A_x\left(e^{-B_x s} - e^{-C_x s}\right)$ for $x \in \{\zeta, \xi\}$. We estimate the parameters by Non-Linear Least Squares, matching model-implied and empirical impulse responses over the observed horizon.

By restricting the damage kernels to be integrable and mean-reverting, we ensure that transitory temperature shocks have bounded long-run effects. In this sense, our approach is conservative: it disciplines the very long-run impact of short-lived temperature fluctuations while remaining consistent with the empirical evidence.

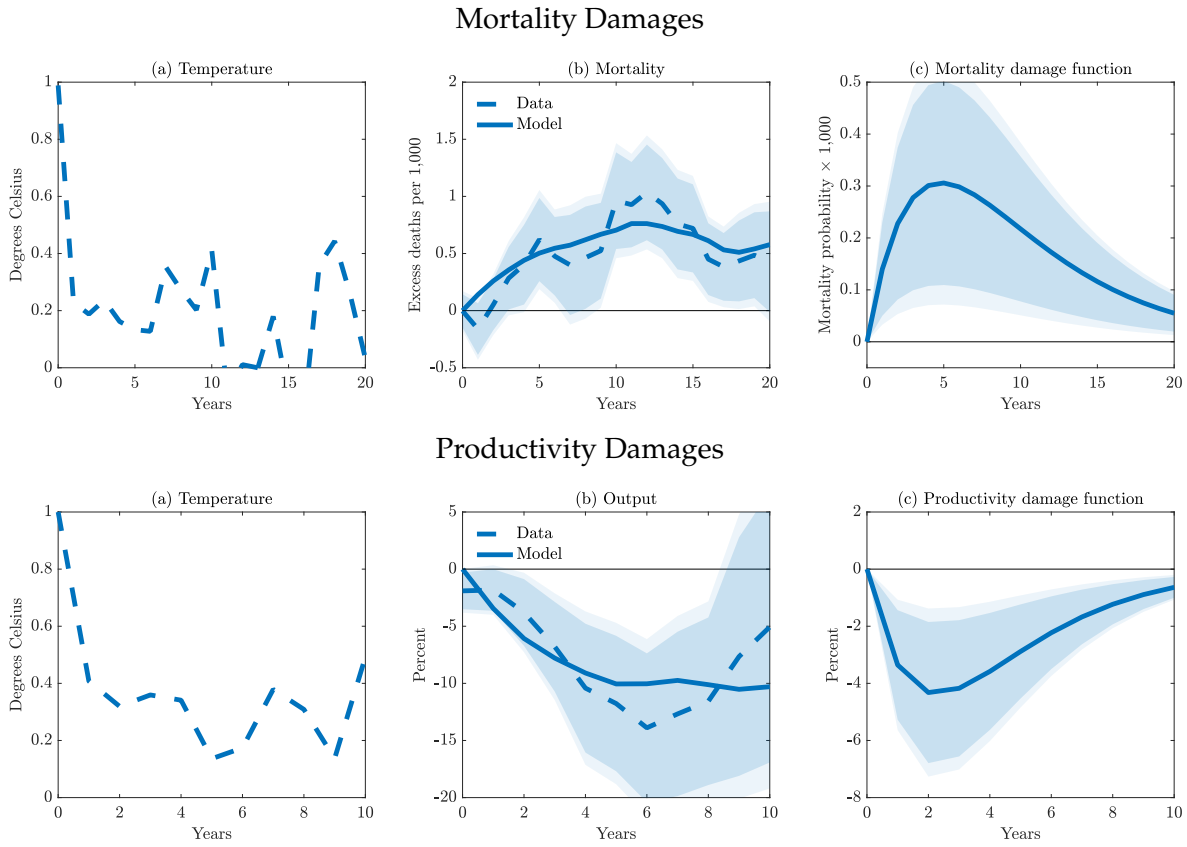
To discipline the value of life in the model, we estimate the end-of-life utility shifter $\overline{V}$ by matching a Value of Statistical Life (VSL). In the model, the empirical VSL satisfies $\overline{V} = u'(C)\text{VSL}$. Under logarithmic utility, this implies $\overline{V} = \frac{\text{VSL}}{C}$, so that the calibration is scale-independent. We proceed in three steps. First, we start from a benchmark U.S. VSL of \$7 million. Second, because climate-related mortality is disproportionately concentrated among individuals aged 65 and above, we adjust for age composition, yielding an effective U.S. VSL of \$2.6 million. Third, we scale this value to world income per capita, which implies an effective global VSL of approximately \$400,000. This calibration pins down $\overline{V}$ and implies a unit elasticity of the VSL with respect to income in the model. In particular, for any country i, the implied VSL satisfies $\text{VSL}_i = \overline{V} \times C_i$.

5.3 Estimation Results

Figure 12 presents the estimated damage functions. Because the empirical mortality response is highly persistent, we target a longer impulse horizon when estimating ξ_s. For productivity, we focus on the first ten years, consistent with Bilal and Känzig (2026), reflecting the horizon over which the empirical output response is precisely estimated.

The first column reproduces the underlying temperature path from Figure 4. The second column shows the mortality and output responses to the temperature shock—in the data and based on our model. The estimated model closely fits the empirical responses of output and mortality, despite its limited degrees of freedom. Of course, the model fit relies on our constrained functional form: if we did not constrain the damage function, the fit would be one-to-one.

Figure 12: Estimated Damage Functions



Notes: Estimated mortality and productivity damage functions obtained by matching model-implied impulse responses to the empirical responses to global temperature shocks. The first row shows the mortality damage function $\{\xi_s\}_s$. The second row shows the productivity damage function $\{\zeta_s\}_s$. Dashed lines denote empirical responses; solid lines denote model fits. Dark and light shaded areas represent 90 and 95% confidence bands, computed using the Delta method.

The last column of the figure depicts the estimated structural damage functions

$\{\zeta_s, \xi_s\}$. These functions represent the mortality and productivity responses to a one-time transitory global temperature shock of 1°C. The mortality consequences of such a shock are substantial. The estimated damage function implies an increase in mortality of about 0.3 deaths per 1,000 at its peak. The effect gradually declines thereafter but remains positive even 20 years after the initial, transitory temperature pulse. We also confirm the productivity damages documented in Bilal and Känzig (2026). The productivity damage function implies a 4% decline in productivity two years after the shock, which then fades gradually but remains negative even ten years out. We compute confidence bands with the Delta-method. The estimated damage functions are statistically significant at the 5% level.

How do the mortality effects of global temperature shocks compare to those associated with local heat shocks? Given that the empirical responses for local heat shocks are closer to zero as shown in Figure 8, such shocks likely also imply smaller damages.

To answer this question quantitatively, we re-estimate our damage functions targeting the impulse response of output to local heat shocks. Doing so yields a mortality damage function that peaks at approximately 0.1 deaths per 1,000, is largely transitory, and is not statistically different from zero at the 5% level (see Appendix Figure E.1). The estimated magnitude is in line with previous evidence on local heat (see e.g. Deschênes and Greenstone 2011; Carleton et al. 2022). We conclude that mortality damages associated with global temperature shocks are larger and more persistent than those implied by local heat shocks.

6 Climate Change, Mortality, and Welfare

With the estimated damage functions at hand, we are now in a position to evaluate the effects of climate change on mortality, the economy, and welfare. We begin by quantifying the impact of the historical warming that the world has experienced to date before studying the effects of future climate change.

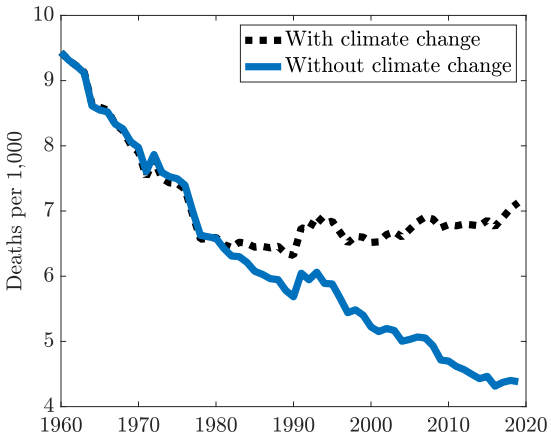
6.1 The Mortality Effect of Historical Warming

Given that the world has warmed by almost 1°C over the period covered by our sample, we begin by quantifying the mortality consequences of the historical warming experienced to date. In particular, we ask whether climate change may have contributed to the

slowdown in the declining mortality trend observed in recent decades.

To do so, we initialize the model in 1960 and feed in the realized path of excess global temperature through 2019. After 2019, we hold temperature constant at its 2019 level. We then compare the resulting mortality path to a counterfactual in which excess temperature remains fixed at its 1960 level throughout. This exercise isolates the contribution of historical warming to observed mortality dynamics.

Figure 13: Mortality Impacts of Historical Warming



Notes: Estimated mortality and productivity damage functions obtained by matching model-implied impulse responses to the empirical responses to global temperature shocks. The first row shows the mortality damage function $\{\xi_s\}_s$. The second row shows the productivity damage function $\{\zeta_s\}_s$. Dashed lines denote empirical responses; solid lines denote model fits. Dark and light shaded areas represent 90 and 95% confidence bands, computed using the Delta method.

Figure 13 presents the results. In the early decades, when warming is limited, the difference between the two paths is small. As excess temperature rises more systematically after 1980, the gap widens. Although annual effects are moderate—reflecting the gradual nature of historical warming—they accumulate over time due to the persistence embedded in the estimated mortality damage function.

By 2019, the crude death rate would be approximately 38% lower in the absence of global warming. Our estimate thus suggest that historical warming materially contributed to offsetting improvements in global health outcomes. Moreover, because the mortality damage function is persistent, part of the mortality burden associated with past warming continues to unfold even if temperature were to remain constant going forward.

In sum, the counterfactual highlights that gradual but sustained increases in global temperature can meaningfully slow the long-run decline in mortality, even when short-run effects appear modest relative to broader demographic and medical trends.

6.2 Climate Change and Mortality in the 21st Century

To evaluate the consequences of future climate change, we specify an exogenous path for global mean temperature. The baseline year $t = 0$ corresponds to 2024. The world subsequently warms by 3°C above preindustrial levels by 2100, after which temperature asymptotes to 3.3°C. This scenario is broadly consistent with IPCC business-as-usual pathways implying 3 to 4°C of warming by 2100 (Hausfather and Peters, 2020; Lee et al., 2023). Given that the world has already warmed by approximately 1°C since preindustrial times, this scenario implies an additional 2°C of warming between 2024 and 2100.

We focus here on mortality effects in isolation, abstracting from productivity damages. Section 6.4 integrates these mortality impacts with the corresponding productivity consequences.

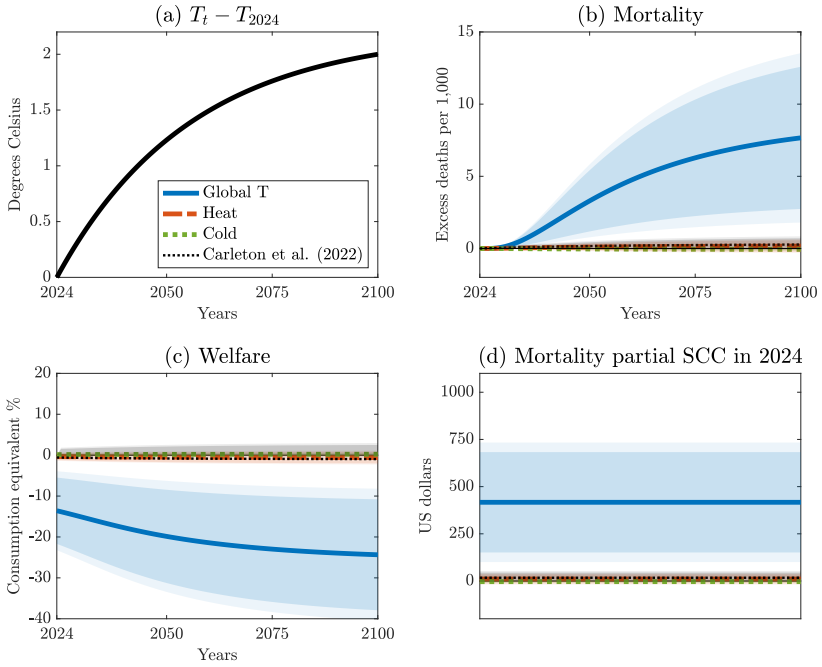
We evaluate the transitional dynamics given the specified warming path under three alternative damage specifications. First, we use the structural damage functions estimated from global temperature shocks. Second, we construct counterfactuals using the damage functions estimated from local heat and cold shocks within our framework. Third, we incorporate mortality damages implied by Carleton et al. (2022), mapped into our model environment. This comparison isolates how different empirical identification strategies translate into long-run mortality, welfare, and SCC implications under the same temperature trajectory.

In each case, we compare allocations and welfare in an economy that warms according to $\widehat{T}_t$ to an economy that remains in steady state under $\widehat{T}_t = 0$. Welfare losses are reported as equivalent percent declines in steady-state consumption. The SCC, expressed in 2024 international dollars, is defined in equation (10). It depends on the temperature response to a marginal CO2 pulse and is therefore independent of the specific warming scenario considered here. All counterfactuals are computed using the fully non-linear model; log-linearization is used only for estimation.

These exercises assess the consequences of sustained multi-degree warming using damage functions estimated from smaller temperature shocks. While households optimally adjust consumption and investment in anticipation of future warming, the framework abstracts from additional behavioral and institutional adaptations beyond those embedded in the neoclassical structure.

Figure 14 presents the results. Panel (a) shows the temperature path under our warming scenario. Panel (b) displays the implied excess mortality per 1,000 individuals. Under

Figure 14: Transitional Dynamics Under Climate Change: Mortality Damages



Notes: Transitional dynamics of the estimated model under the scenario in Panel (a) allowing for mortality damages only. Solid blue lines: model estimated under global temperature. Shaded blue: 90 and 95% confidence intervals based on the Delta-method. Dashed lines: model estimated under alternative damage functions: local heat (red dashed), local cold (green dashed), Carleton et al. (2022) damage function (black dotted).

global temperature damages, mortality rises steadily throughout the century, reaching around 7.5 excess deaths per 1,000 by 2100. The increase reflects both the continued rise in temperature and the persistence embedded in the estimated mortality damage function. Even though annual temperature increases are gradual, their cumulative effects generate substantial additional mortality by the end of the century.

In contrast, when damages are constructed using local heat and cold estimates, the mortality effects are considerably smaller and statistically indistinguishable from zero. (Appendix Figure E.2 zooms into the transitional impacts under these alternative damage functions). This difference mirrors our reduced-form estimates: global temperature shocks generate large and persistent mortality responses, whereas local heat and cold shocks imply much smaller and largely transitory effects. Reassuringly, our mortality impacts under local heat and cold are broadly consistent with existing estimates (see e.g. Carleton et al., 2022).

Panel (c) reports welfare losses as consumption-equivalent declines associated with the mortality impacts. Under global temperature damages, the welfare impact of climate

change amounts to a 15% loss in consumption equivalent in 2024, which continues to increase monotonically over the century as mortality damages accumulate. In contrast, welfare losses under local heat and cold damages are much more modest.

Panel (d) reports the mortality component of the social cost of carbon in 2024. The SCC implied by global temperature damages is $420 per ton of CO2. This value is multiple orders of magnitude larger than the implied social cost under conventional damage function specifications based on local heat and cold ($10-$20 per ton).

Overall, the results indicate that global temperature variation implies large and persistent mortality consequences over the 21st century, with corresponding welfare losses and a substantial mortality component of the SCC. Importantly, damages operate primarily through the decline in global income induced by warming, which generates large and highly persistent mortality implications that estimates based solely on local heat or cold variation fail to capture.

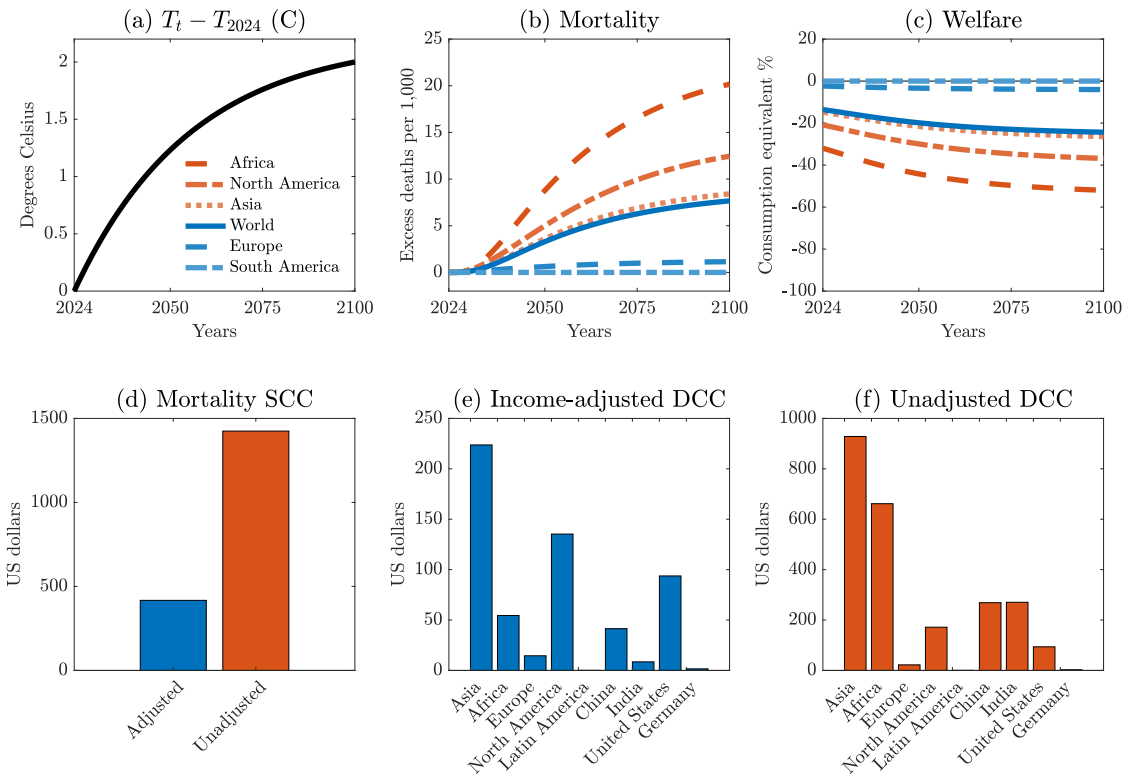
6.3 Regional Impacts

We next examine how the mortality consequences of future warming differ across regions. To do so, we combine the regional reduced-form mortality responses estimated in Appendix D.2.1 with the global temperature path described above. We use these regional impulse responses to construct region-specific mortality damage functions and embed them in the model to evaluate mortality, welfare, and carbon pricing implications across regions.

Figure 15 presents the results. Panel (a) shows the common global temperature path. Panel (b) displays excess mortality per 1,000 individuals across regions. Mortality effects are largest in Africa, North America, and Asia, and substantially smaller in Europe. These differences reflect heterogeneity in the estimated regional temperature sensitivities embedded in the damage functions. Panel (c) reports the corresponding welfare losses from mortality in consumption-equivalent terms. Regions with larger mortality responses experience correspondingly larger welfare declines.

Panel (d) reports the mortality component of the global social cost of carbon. The magnitude of the SCC depends crucially on how regional damages are aggregated. When damages are income-adjusted, poorer regions receive less weight, reducing the aggregate SCC relative to an unadjusted specification. In contrast, without income adjustment, regions with lower income but larger population and larger mortality impacts receive

Figure 15: Mortality Impacts Across Regions



Notes: Transitional dynamics of the model given the warming scenario in Panel (a) based on alternative damage function specifications: local heat (red dashed), local cold (green dashed), Carleton et al. (2022) damage function (black dotted). Shaded areas: 90 and 95% confidence intervals based on the Delta-method.

greater weight, substantially increasing the implied SCC.

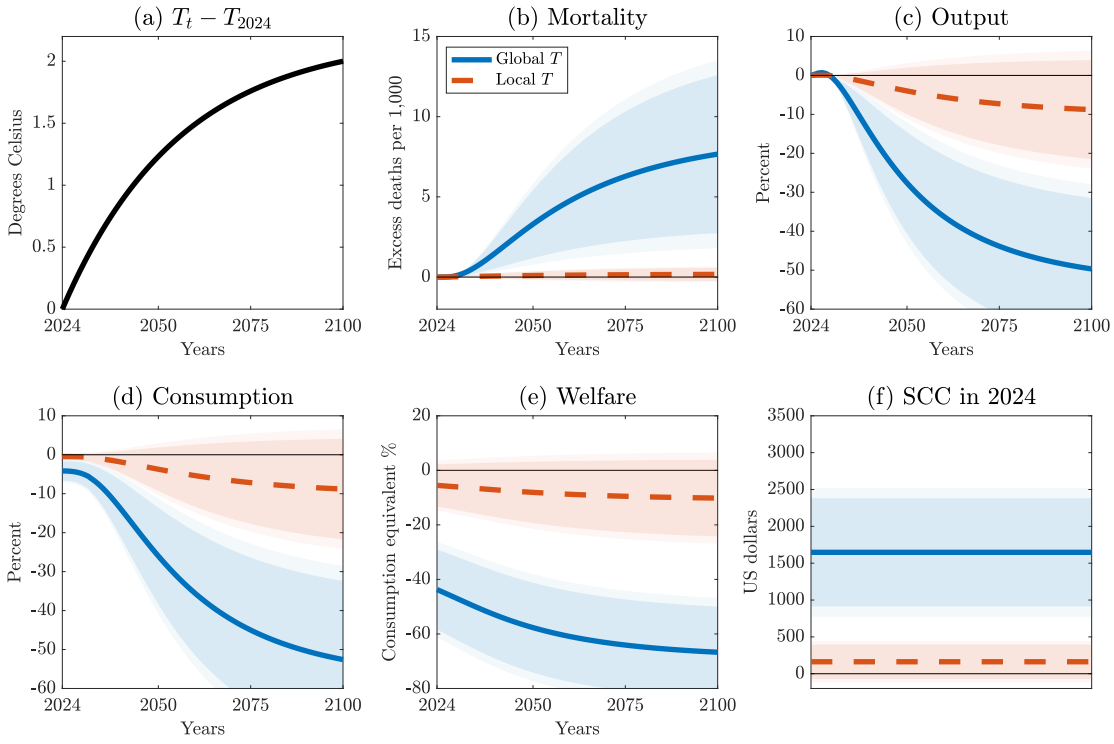
Panels (e) and (f) report regional domestic costs of carbon (DCCs), with and without income adjustment. The domestic cost of carbon is particularly high in populous regions with large projected mortality effects, reflecting the local benefits of emissions reductions. The contrast between income-adjusted and unadjusted DCCs highlights the normative importance of distributional weights in aggregating climate damages.

By way of summary, the regional analysis shows that the mortality consequences of warming are highly uneven across the globe. Regions with higher mortality impacts to global temperature shocks and large populations bear disproportionately large mortality burdens.

6.4 The Total Welfare Impact of Climate Change

We now allow for both mortality and productivity damages and evaluate the total welfare consequences of future climate change. This exercise integrates the health and economic channels within our general equilibrium framework.

Figure 16: Transitional Dynamics Under Climate Change: Total Damages



Notes: Transitional dynamics of the estimated model under the scenario in Panel (a) allowing for both mortality and productivity damages. Solid blue lines: model estimated under global temperature. Shaded blue: 90 and 95% confidence intervals. Dashed red lines: model estimated under local temperature. Shaded red: 90 and 95% confidence intervals. Confidence intervals based on the Delta-method.

Figure 16 presents the results. Panels (c) and (d) show that allowing for productivity damages confirms the substantial economic effects of climate change documented in Bilal and Känzig (2026). Under the global temperature damage function, output and consumption decline sharply over the 21st century. These losses are markedly larger than those implied by conventional damage function specifications based on local temperature, which generates comparatively modest macroeconomic effects.

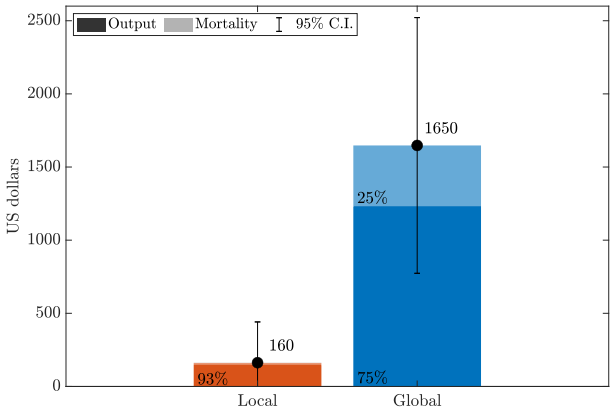
Panel (e) reports total welfare losses in consumption-equivalent terms. Welfare declines are substantially larger under global temperature damages than under local temperature damages, reflecting the combined effect of persistent productivity losses and el-

evated mortality risk. Factoring in mortality, the present-value welfare loss exceeds 40% in 2024 and rises to nearly 70% by the end of the century. Importantly, these figures do not imply that future generations will be poorer. Rather, they indicate that welfare in 2100 would be roughly 70% higher in the absence of climate change. Even under warming, the economy may continue to grow and mortality may continue to decrease; climate change reduces the level of welfare relative to its counterfactual path.

Panel (f) reports the Social Cost of Carbon in 2024. Under the local temperature damage function, the SCC is modest, at roughly $150. Mortality contributes only a small share in this case—on the order of $10—so that the SCC is driven almost entirely by productivity effects.

In contrast, under the global temperature damage function, the SCC is approximately $1,200 when considering productivity damages alone. Allowing for mortality raises the SCC to about $1,620, implying that mortality contributes roughly $420. Thus, mortality makes up a considerable share of the total SCC in this specification.

Figure 17: Social Cost of Carbon: Mortality vs. the Economy



Notes: Transitional dynamics of the estimated model under the scenario in Panel (a) allowing for both mortality and productivity damages. Solid blue lines: model estimated under global temperature. Shaded blue: 90 and 95% confidence intervals. Dashed red lines: model estimated under local temperature. Shaded red: 90 and 95% confidence intervals. Confidence intervals based on the Delta-method.

This split is further illustrated in Figure 17, which decomposes the SCC into output and mortality components. Under local temperature damages, the SCC is driven overwhelmingly by output effects, with mortality accounting for a small share. Under global temperature damages, by contrast, mortality represents roughly 25% of the total SCC. The larger role for mortality reflects the substantial income-mediated health effects identified in our global temperature estimates.

Taken together, these results show that incorporating global temperature variation substantially amplifies both the macroeconomic and mortality consequences of climate change. Not only are output losses larger, but mortality becomes a quantitatively important component of total climate damages. Estimates based solely on local temperature variation therefore understate both the level of the SCC and the relative importance of mortality in driving climate costs.

7 Conclusion

This paper shows that the mortality consequences of climate change are fundamentally global and macroeconomic in nature. While much of the existing literature emphasizes localized heat exposure, we demonstrate that such channels miss large mortality effects of warming that are mediated through income. Global temperature increases trigger sizable and persistent rises in mortality primarily by inducing worldwide economic contractions that erode access to care and other health-preserving resources.

Leveraging plausibly exogenous fluctuations in global mean temperature, we estimate that a permanent 1°C increase raises annual mortality by 2.5 deaths per 1,000 people—an order of magnitude larger than estimates based on local heat alone. We provide supporting evidence that these effects are mediated by income losses. The resulting mechanism reconciles why mortality responds strongly to global temperature even when local heat effects appear modest.

Incorporating these findings into a structural model of climate damages fundamentally reshapes welfare assessments. Mortality contributes roughly one quarter of the total Social Cost of Carbon, amounting to more than $420 per ton of CO2—multiple orders of magnitudes higher than conventional estimates derived from local heat exposure.

Our results carry two broader implications. First, they strengthen the case for unilateral climate mitigation by uncovering a large and previously underappreciated mortality externality operating through global economic linkages. In particular, because the mortality burden scales with population and income exposure, unilateral decarbonization may be cost-effective even for relatively lower-income, but populous, economies. Second, our findings highlight a central role for policies that sustain income growth and stabilize economic activity. By preserving access to health care and other essential resources, such policies reduce not only economic losses but also the human cost of climate change.

References

Aguilar-Gomez, Sandra, Hopi Dwyer, Joshua Graff Heil, and Nathan P. Hendricks (2022). "This is Air: The Non-health Effects of Air Pollution". *Annual Review of Resource Economics* 14, pp. 403–425.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro (2016). "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century". *Journal of Political Economy* 124.1, pp. 105–159.

Bilal, Adrien and Diego R. Känzig (2026). "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature". *Forthcoming, The Quarterly Journal of Economics* 32450.

Bilal, Adrien, Diego R. Känzig, and Krzysztof Lisiecki (2026). "Life Expectancy and Economic Growth: Local or Global GDP?" *Working Paper*.

Blanchard, Olivier Jean and Lawrence F. Katz (1992). "Regional Evolutions". *Brookings Papers on Economic Activity* 23.1, pp. 1–76.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone (2017). "Weather, Climate Change and Death in India". *Working Paper*. University of Chicago.

Burke, Marshall, Anne Driscoll, Sam Heft-Neal, Jiani Xue, Jennifer Burney, and Michael Wara (2021). "The changing risk and burden of wildfire in the United States". *Proceedings of the National Academy of Sciences* 118.2, e2011048118.

Burke, Marshall, Sam Heft-Neal, Jessica Li, Anne Driscoll, Patrick Baylis, Mathew Stiglic, Joakim A. Weill, Jennifer A. Burney, Jeff Wen, Marshall L. Childs, and Carlos F. Gould (2022). "Exposures and behavioural responses to wildfire smoke". *Nature Human Behaviour* 6.10, pp. 1351–1361.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (2015). "Global non-linear effect of temperature on economic production". *Nature* 527.7577, pp. 235–239.

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 (2022). "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". *The Quarterly Journal of Economics* 137.4, pp. 2037–2105.

Cutler, David, Angus Deaton, and Adriana Lleras-Muney (2006). "The Determinants of Mortality". *Journal of Economic Perspectives* 20.3, pp. 97–120.

Cutler, David M., Wei Huang, and Adriana Lleras-Muney (2016). *Economic conditions and mortality: evidence from 200 years of data*. Tech. rep. National Bureau of Economic Research.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2012). "Temperature shocks and economic growth: Evidence from the last half century". *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

Deschênes, Olivier and Michael Greenstone (2011). "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US". *American Economic Journal: Applied Economics* 3.4, pp. 152–85.

Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans (2021a). "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists* 8.5, pp. 895–921.

Dietz, Simon, James Rising, Thomas Stoerk, and Gernot Wagner (2021b). "Economic impacts of tipping points in the climate system". *Proceedings of the National Academy of Sciences* 118.34, e2103081118.

Domeisen, Daniela I. V., Elfatih A. B. Eltahir, Erich M. Fischer, Reto Knutti, Sarah E. Perkins-Kirkpatrick, Christoph Schär, Sonia I. Seneviratne, Weisheimer Antje, and Heini Wernli (2023). "Prediction and Projection of Heatwaves". *Nature Reviews Earth and Environment* 4, 36–50.

Driscoll, John C. and Aart C. Kraay (1998). "Consistent covariance matrix estimation with spatially dependent panel data". *Review of Economics and Statistics* 80.4, pp. 549–560.

Finkelstein, Amy, Matthew J Notowidigdo, Frank Schilbach, and Jonathan Zhang (2025). "Lives Versus Livelihoods: The Impact of the Great Recession on Mortality and Welfare". *The Quarterly Journal of Economics*, qjaf023.

Folini, Doris, Aleksandra Friedl, Felix Kübler, and Simon Scheidegger (2024). "The Climate in Climate Economics". *The Review of Economic Studies*.

Granger, C.W.J. and P. Newbold (1974). "Spurious regressions in econometrics". *Journal of Econometrics* 2.2, pp. 111–120.

Hamilton, James D. (2018). "Why you should never use the Hodrick-Prescott filter". *Review of Economics and Statistics* 100.5, pp. 831–843.

Hausfather, Zeke and Glen P. Peters (2020). "Emissions – the 'business as usual' story is misleading". *Nature* 577.7792, pp. 618–620.

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 (2017). "Estimating Economic Damage from Climate Change in the United States". *Science* 356.6345, pp. 1362–1369.

Joos, F., R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-

K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver (2013). "Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis". *Atmospheric Chemistry and Physics* 13.5, pp. 2793–2825.

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". *American Economic Review* 95.1, pp. 161–182.

Jordà, Òscar and Alan M Taylor (2025). "Local projections". *Journal of Economic Literature* 63.1, pp. 59–110.

Kaufmann, Robert K., Heikki Kauppi, and James H. Stock (2006). "Emissions, concentrations, & temperature: a time series analysis". *Climatic Change* 77, pp. 249–278.

Kose, M. Ayhan, Naotaka Sugawara, and Marco E. Terrones (2020). "Global recessions".

Lee, Hoesung, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. (2023). "IPCC, 2023: Climate Change 2023: Synthesis Report, Summary for Policymakers. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland."

Ma, Lin, Gil Shapira, Damien De Walque, Quy-Toan Do, Jed Friedman, and Andrei A. Levchenko (2022). "The Intergenerational Mortality Trade-Off Of Covid-19 Lockdown Policies". *International Economic Review* 63.3, pp. 1427–1468.

Montiel Olea, José Luis, Mikkel Plagborg-Moller, Eric Qian, and Christian K. Wolf (2024). "Double Robustness of Local Projections and Some Unpleasant VARithmetic". *National Bureau of Economic Research Working Paper Series 32495*.

Moore, Frances C. and Delavane B. Diaz (2015). "Temperature impacts on economic growth warrant stringent mitigation policy". *Nature Climate Change* 5.2, pp. 127–131.

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

National Oceanic and Atmospheric Administration (2005). *Volcanos and Climate*. Online Resource.

— (2009). *Climate Change: Incoming Sunlight*. Online Resource.

— (2023). *What are El Niño and La Niña?* Online Resource.

Newell, Richard G., Brian C. Prest, and Steven E. Sexton (2021). "The GDP-temperature relationship: implications for climate change damages". *Journal of Environmental Economics and Management* 108, p. 102445.

Nordhaus, William D. (1992). "An optimal transition path for controlling greenhouse gases". *Science* 258.5086, pp. 1315–1319.

Preston, Samuel H. (1975). "The Changing Relation between Mortality and Level of Economic Development". *Population Studies* 29.2, pp. 231–248.

Pritchett, Lant and Lawrence H. Summers (1996). "Wealthier is healthier". *The Journal of Human Resources* 31.4, p. 841.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. (2022). "Comprehensive evidence implies a higher social cost of CO2". *Nature* 610.7933, pp. 687–692.

Ruhm, Christopher J. (2000). "Are recessions good for your health?" *The Quarterly Journal of Economics* 115.2, pp. 617–650.

— (2015). "Recessions, healthy no more?" *Journal of Health Economics* 42, pp. 17–28.

Seneviratne, S. I., X. Zhang, M. Adnan, W. Badi, C. Dereczynski, A. Di Luca, S. Ghosh, I. Iskandar, J. Kossin, S. Lewis, F. Otto, I. Pinto, M. Satoh, S.M. Vicente-Serrano, M. Wehner, and B. Zhou (2021). "Chapter 11: Weather and Climate Extreme Events in a Changing Climate". *In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change; Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.). Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA*, 1513–1766.

Seneviratne, Sonia I., Markus G. Donat, Andy J. Pitman, Reto Knutti, and Robert L. Wilby (2016). "Allowable CO2 Emissions Based on Regional and Impact-Related Climate Targets". *Nature* 529, 477–483.

Sims, Christopher A. (1986). "Are forecasting models usable for policy analysis?" *Quarterly Review* 10, pp. 2–16.

Stevens, Ann H, Douglas L Miller, Marianne E Page, and Mateusz Filipski (2015). "The best of times, the worst of times: understanding pro-cyclical mortality". *American Economic Journal: Economic Policy* 7.4, pp. 279–311.

Wartenburger, R., M. Hirschi, M. G. Donat, P. Greve, A. J. Pitman, and S. I. Seneviratne (2017). "Changes in Regional Climate Extremes as a Function of Global Mean Temperature: an Interactive Plotting Framework". *Geoscientific Model Development* 10, 3609–3634.

Young, Rachel and Solomon M. Hsiang (2024). "Mortality caused by tropical cyclones in the United States". *Nature* 635.8038, pp. 121–128.

Online Appendix

Global Temperature and Global Mortality

Adrien Bilal
Diego R. Känzig
Krzysztof Lisiecki

Contents

A. Data **47**
A.1. Mortality Data . . . 47
A.2. Climate Data . . . 48
A.3. Economic Data . . . 49

B. Additional Results for Baseline Estimates **50**
B.1. Age-Specific Mortality . . . 50
B.2. Transitory Temperature Shock . . . 50
B.3. One-step Local Projection . . . 51

C. Additional Results for Omitted Variable Bias **52**
C.1. Excluding India and China . . . 52
C.2. Jackknife Analysis . . . 53
C.3. Alternative Detrending Methods . . . 53

D. Additional Results for Channels **55**
D.1. Extreme Events . . . 55
D.2. The Income Channel . . . 56
D.2.1. Regional Heterogeneity . . . 56
D.2.2. Global GDP and Global Temperature . . . 57
D.2.3. Mortality and GDP . . . 58

E. Additional Model Results **59**
E.1. Mortality Damage Functions From Extreme Heat . . . 59

A Data

A.1 Mortality Data

We construct our mortality variables by combining data from several sources, with the Human Mortality Database (HMD, 2024) serving as the primary input. The HMD is widely regarded as the gold standard for mortality data. It is based on high-quality vital registration data, complemented by census counts, birth records, and population estimates, and applies a harmonized methodology across countries. Crucially for our purposes, it relies on observed annual death counts rather than model-based imputation across years, making it well suited to high-frequency analysis that exploits year-to-year variation. The HMD covers most of Western and Northern Europe, North America, Australia, and a number of Eastern European and Asian countries. We download mortality counts `deaths.zip` and population and live-birth data, `population.zip` and `births.zip`, from `www.mortality.org/Data/ZippedDataFiles`. For each age group—total, ages 0–4, ages 5–64, and ages 65 and above—we divide death counts by population and multiply by 1,000 to obtain mortality rates per 1,000 individuals. This normalization is applied consistently across all data sources to ensure comparability.

Because HMD coverage is largely confined to high-income countries, we supplement the HMD series with data from the World Health Organization (WHO, 2024) Mortality Database to expand coverage and fill gaps. The WHO database compiles death counts reported by national statistical authorities from civil registration systems, with data disaggregated by age, sex, and cause of death and subject to varying degrees of completeness across countries and over time. Relative to the HMD, it provides broader geographic coverage, extending to Latin America and the Caribbean, the Middle East, parts of Asia, Eastern Europe, and Africa. We download mortality counts and population data from `www.who.int/data/data-collection-tools/who-mortality-database`. For countries where the entire HMD series is missing, we use the corresponding WHO mortality counts to construct our measures. For countries with partial gaps, we preserve the level of the HMD series and use the growth rates implied by the WHO data to interpolate and extrapolate missing years.

Despite the expanded coverage, two crucial countries remain missing from the baseline sources: India and China. For these countries, we therefore draw on additional sources. We obtain mortality rates for China directly from the National Bureau of Statis-

tics of China (https://data.stats.gov.cn/english/easyquery.htm?cn=C01), which reports annual crude death rates based on official statistics. For India, we obtain mortality counts from Vital Statistics India (VSI) and combine them with population data from the United Nations World Population Prospects. Starting from 2009, VSI reports are available at https://censusindia.gov.in/census.website/data/VSREPORT. For earlier years we digitize this data from University library sources. Finally, we add mortality rates for Liechtenstein from Carleton et al. (2022). Together, these sources yield a mortality panel with coverage of around 60 countries in 1960, rising to over 100 countries in the second half of our sample period.

Despite this broad coverage, reliable annual mortality data remain unavailable for many countries, particularly in Africa, where civil registration systems are often incomplete or absent. For these countries, we rely on estimates from the United Nations and Division World Population Prospects (UNWPP, 2024), accessed from population.un.org/wpp/. Unlike the HMD and WHO databases, which are based on directly observed deaths from civil registration systems, the UNWPP constructs internally consistent demographic series by combining a wide range of available sources—including population censuses, household surveys, partial registration systems, and international datasets—and applying demographic models to reconcile discrepancies and interpolate missing observations. In particular, the estimates are derived using a cohort-component framework that ensures consistency between population levels and the underlying components of change, namely fertility, mortality, and migration. While this approach provides comprehensive global coverage and is well suited for capturing long-run trends and cross-country differences, it necessarily smooths high-frequency variation and is therefore less appropriate for our empirical design, which exploits annual fluctuations in mortality. Based on this data, we expand our mortality sample. Specifically, for any country where the HMD-WHO-based series contains fewer than 30 observations, we replace it with a mortality rate constructed from UNWPP death counts and population figures. This extended variable covers up to 210 countries and is used in sensitivity checks and for the regional analysis in Africa.

A.2 Climate Data

Our climate dataset builds on Bilal and Känzig (2026) and combines data on global and country-level temperature as well as information on extreme climatic events.

We use annual global temperature anomalies (land and ocean surface) from NOAA National Centers for Environmental Information (2023) and Berkeley Earth (Rohde and Hausfather, 2020), which we convert to temperature levels using their respective climatological baselines.

For spatially disaggregated analysis, we rely on gridded temperature data from Berkeley Earth (Rohde and Hausfather, 2020), available at a $1^\circ \times 1^\circ$ resolution over a long historical period. We aggregate these data to the country level using population weights, although results are robust to area weighting. As a validation exercise, aggregating the gridded data to the global level closely reproduces the official NOAA temperature series.

To capture broader climatic conditions, we complement temperature data with gridded information on temperature, precipitation, and wind from ISIMIP (Lange et al., 2023), available at a $0.5^\circ \times 0.5^\circ$ spatial resolution. To characterize extreme climatic conditions, we construct indicators based on daily gridded data and percentile thresholds of the historical distribution. Specifically, we define extreme events relative to the distribution over a baseline period (1950–1980), distinguishing between country-specific (relative) thresholds and global (absolute) thresholds. Our baseline uses relative thresholds, capturing deviations from local climatic norms. We compute, for each country and year, the fraction of days in which temperature, precipitation, or wind exceed the relevant threshold, and consider both country-level and grid-level aggregation approaches. The results are robust to alternative threshold definitions, baseline periods, and aggregation methods.

A.3 Economic Data

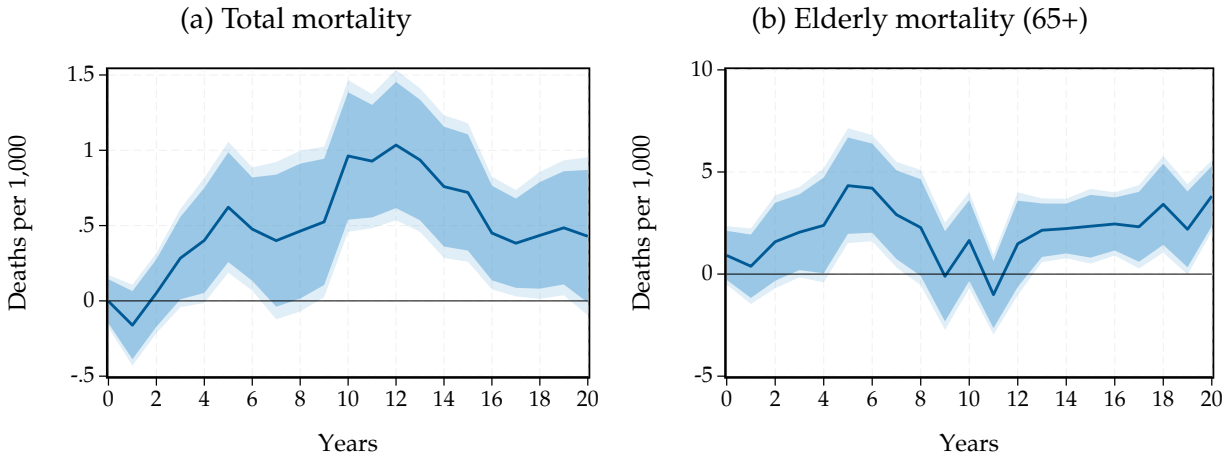
We obtain GDP and population data from the Penn World Tables (PWT; Feenstra, Inklaar, and Timmer, 2015), which provides comprehensive coverage for our sample countries from 1960 to 2019. Our main output measure is real GDP per capita from the national accounts (rgdpna/pop). Based on this data, we construct a measure of world real GDP, by chain-weighting country-level growth rates. We obtain a global commodity price index from Bloomberg. World recession dates are taken from the World Bank Kose, Sugawara, and Terrones (2020). Data on wars are drawn from the Uppsala Conflict Data Program. Information on major epidemics and pandemics is compiled from publicly available historical sources.

B Additional Results for Baseline Estimates

B.1 Age-Specific Mortality

Figure B.1 examines effects on total mortality and mortality among the elderly population (age 65+). The elderly population exhibits particularly large mortality effects, consistent with greater vulnerability to climate-related health risks, and seem to be driving the aggregate impact that we estimate.

Figure B.1: Age-Specific Mortality Responses



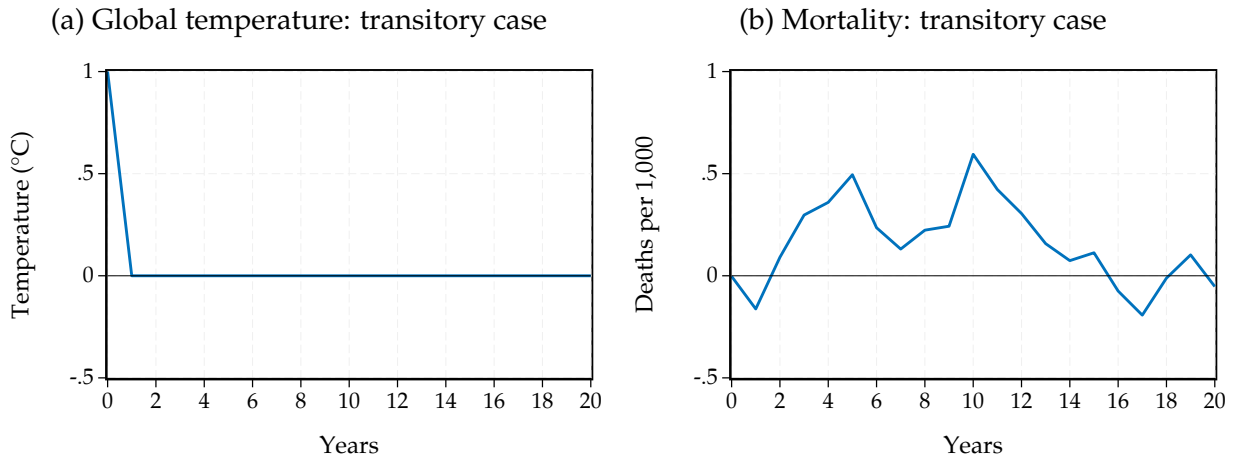
Notes: Impulse responses for age-specific mortality rates. Elderly (65+) show particularly large mortality increases from global temperature shocks, driving the aggregate impact.

B.2 Transitory Temperature Shock

Figure 4 in the main text shows that mortality effects build up over time as global temperature remains persistently elevated following a shock. To isolate the role of temperature persistence, we construct a counterfactual mortality response to a purely transitory temperature shock using the method of Sims (1986).

Figure B.2 shows the results. Panel (a) displays the transitory temperature path: by construction, temperature rises by 1°C on impact and then immediately returns to baseline. Panel (b) shows the corresponding mortality response. Under this transitory shock, the peak mortality effect is approximately 0.6 deaths per 1,000 at year 5, compared to 1.3 deaths per 1,000 in the baseline. This comparison indicates that roughly half of the peak mortality effect in our baseline estimates is due to the accumulated impact of persistently

Figure B.2: Mortality Response to a Transitory Temperature Shock



Notes: Counterfactual impulse responses to a purely transitory 1°C temperature shock, constructed using the method of Sims (1986). The transitory shock raises temperature by 1°C on impact and then immediately returns to baseline, isolating the direct effect of temperature without accumulated persistence.

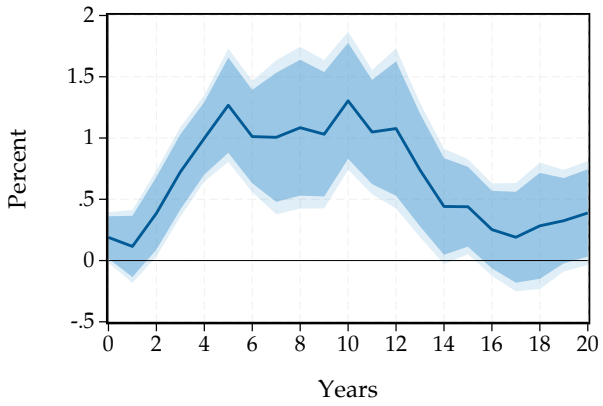
elevated temperature.

B.3 One-step Local Projection

Our baseline specification constructs temperature shocks using the Hamilton filter, which involves estimation uncertainty. To verify that this uncertainty does not affect our inference, we implement a one-step local projection that jointly estimates temperature dynamics and mortality effects with longer lag structures.

Figure B.3 shows that the one-step approach yields virtually identical results to our baseline two-step procedure. This confirms that estimation uncertainty in temperature shocks is not a material concern for our inference.

Figure B.3: One-step Local Projection



Notes: Comparison of baseline two-step procedure (construct temperature shocks, then estimate mortality effects) versus one-step local projection that jointly estimates both with longer lags. The one-step approach accounts for estimation uncertainty in temperature shocks.

C Additional Results for Omitted Variable Bias

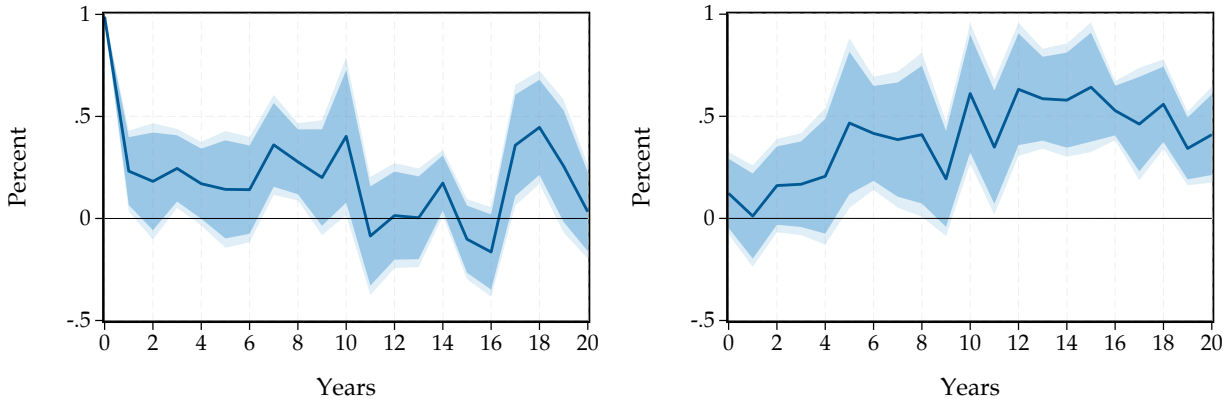
C.1 Excluding India and China

India and China represent a large share of our sample population and have mortality data of potentially lower quality than high-income countries. Figure C.1 shows that excluding these countries yields similar results to our baseline, with a peak effect of approximately 1.2 deaths per 1,000 compared to 1.3 in the full sample, and comparable statistical precision.

Figure C.1: Excluding India and China

(a) Global temperature

(b) Mortality rate

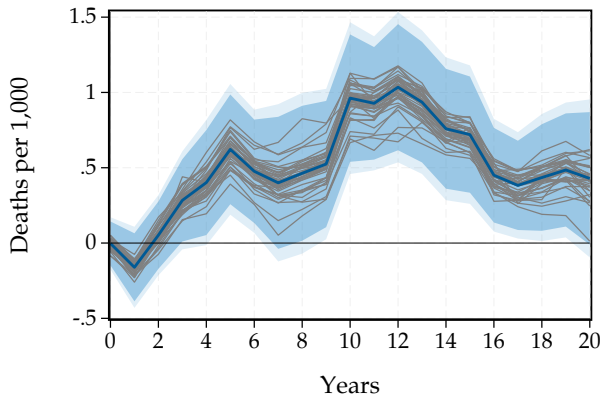


Notes: Impulse responses excluding India and China from the sample. The results are similar to the baseline specification including all 41 countries.

C.2 Jackknife Analysis

To verify that our results are not driven by any single temperature shock, we implement a jackknife procedure. For each temperature shock observation, we set it to zero and re-estimate the impulse responses. Figure C.2 shows that the baseline estimate lies well within the range of jackknife estimates, indicating that no individual shock is driving our findings.

Figure C.2: Jackknife Analysis: Leave-One-Out Temperature Shocks

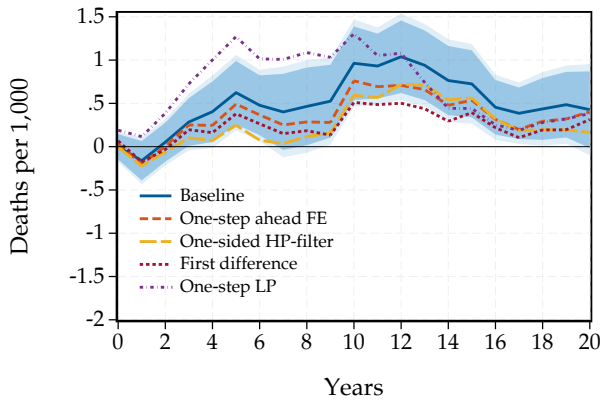


Notes: Jackknife estimates setting each temperature shock to zero one at a time. The baseline estimate (solid line) lies within the range of all jackknife estimates, indicating robustness to excluding any single shock. Shaded area: range of jackknife estimates.

C.3 Alternative Detrending Methods

Our baseline uses the Hamilton filter to construct temperature shocks. Figure C.3 compares results using alternative detrending methods: HP filter, different Hamilton filter parameters, and first differences. The impulse responses are remarkably similar across all methods, confirming that our results are not sensitive to the specific detrending approach.

Figure C.3: Alternative Temperature Shock Construction Methods

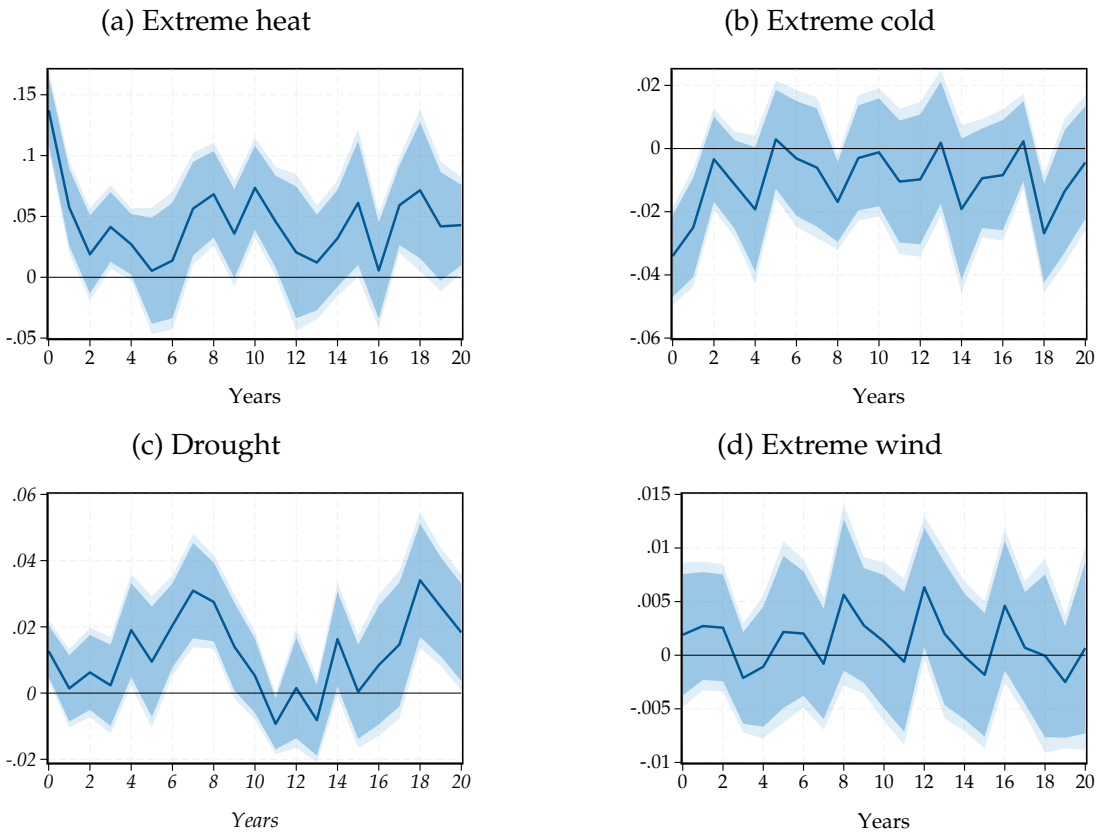


Notes: Impulse responses using alternative methods for constructing temperature shocks: baseline Hamilton filter ($h = 2, p = 2$), HP filter, alternative Hamilton parameters, and first differences. Results are robust across detrending methods.

D Additional Results for Channels

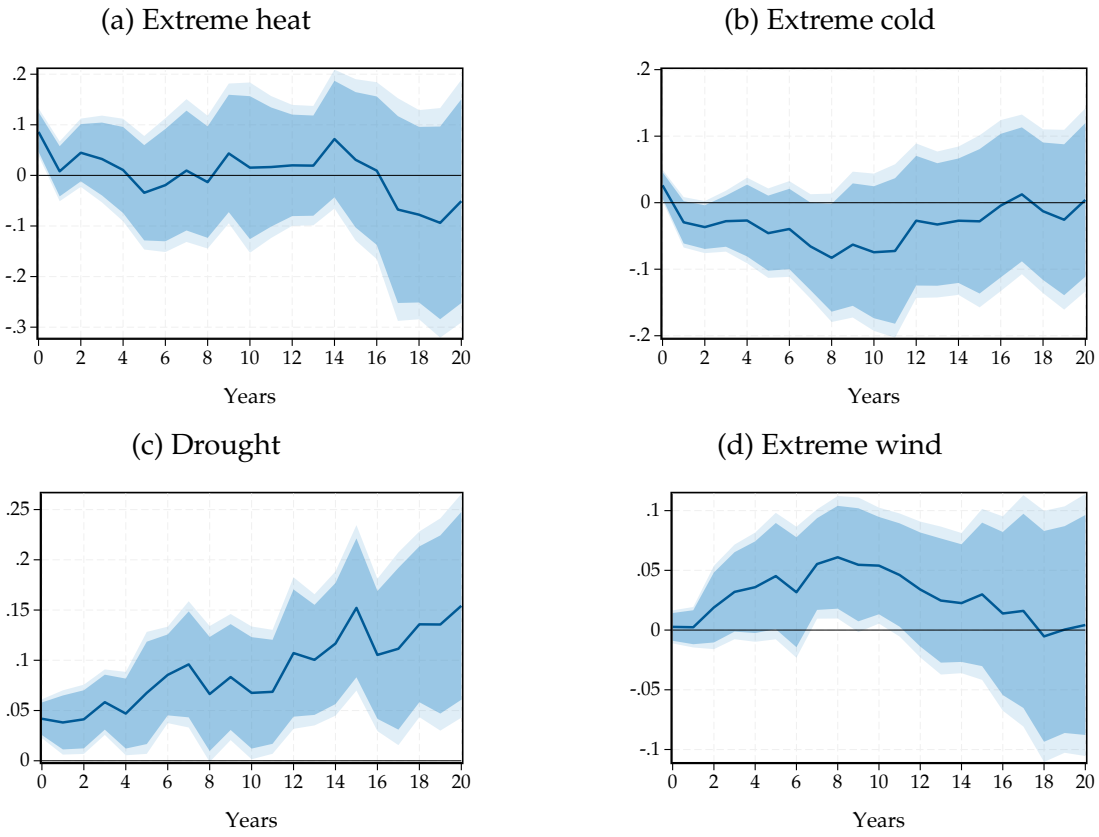
D.1 Extreme Events

Figure D.1: Extreme Weather Events and Global Temperature Shocks



Notes: Impulse responses of extreme weather exposure indices to a 1°C global temperature shock. Each index measures the fraction of days when weather exceeds historical (1950-1980) thresholds. Extreme heat: above 95th percentile (baseline 0.05). Extreme cold: below 5th percentile (baseline 0.05). Drought: below 10th percentile (baseline 0.10). Extreme wind: above 99th percentile (baseline 0.01). Global temperature shocks significantly increase extreme heat, droughts, precipitation, and wind while decreasing extreme cold. Shaded areas: 90% and 95% confidence bands.

Figure D.2: Mortality and Extreme Weather Events



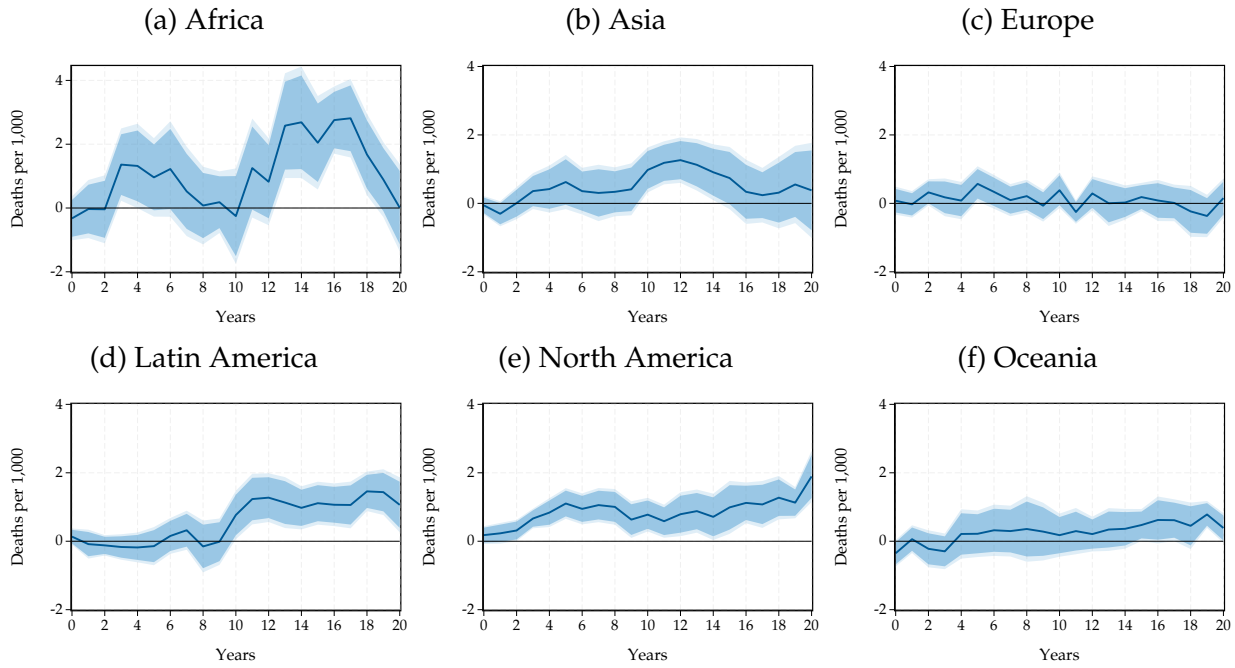
Notes: Impulse responses of mortality to a 1 point increase in extreme weather exposure indices. Each index measures the fraction of days when weather exceeds historical (1950-1980) thresholds. Extreme heat: above 95th percentile (baseline 0.05). Extreme cold: below 5th percentile (baseline 0.05). Drought: below 10th percentile (baseline 0.10). Extreme wind: above 99th percentile (baseline 0.01). Responses are normalized to the peak increase of frequency from Figure D.1. Shaded areas: 90% and 95% confidence bands.

D.2 The Income Channel

D.2.1 Regional Heterogeneity

Figure D.3 shows mortality responses by region. All regions exhibit positive mortality effects from global temperature shocks, though magnitudes vary. Africa and Asia show particularly large responses, while Europe and North America show smaller but still significant effects. This heterogeneity likely reflects differences in baseline climate, healthcare infrastructure, and economic development.

Figure D.3: Regional Heterogeneity in Mortality Responses

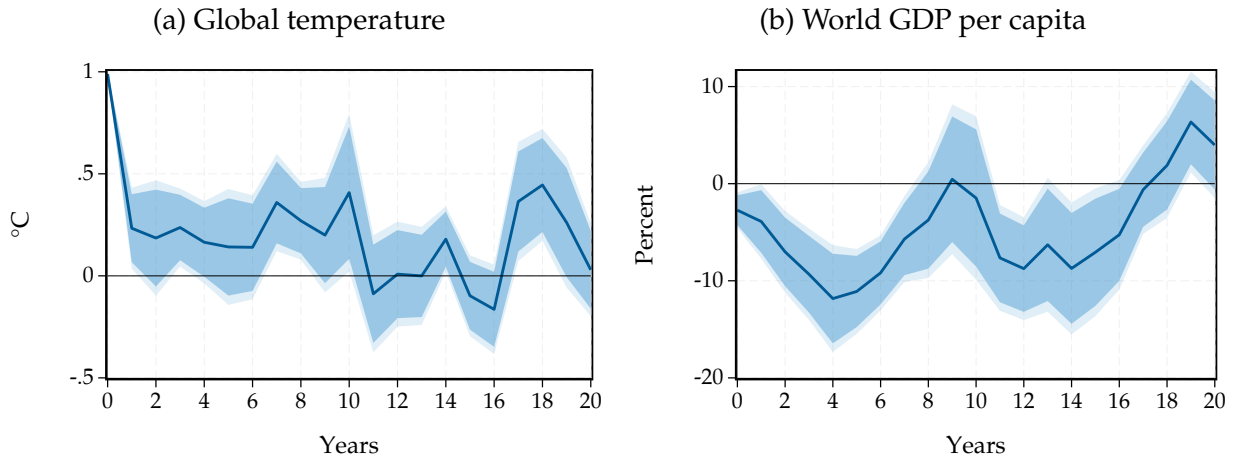


Notes: Impulse responses of mortality to global temperature shocks estimated separately by region. All regions show positive mortality effects, with heterogeneity in magnitudes reflecting differences in baseline climate, healthcare, and economic development.

D.2.2 Global GDP and Global Temperature

Figure D.4 reproduces the results from Bilal and Känzig (2026) extended to a 20-year horizon. Panel (a) shows the impulse response of global temperature to a temperature shock, while Panel (b) shows the corresponding GDP response.

Figure D.4: Global Temperature and World GDP (Bilal and Känzig, 2026)



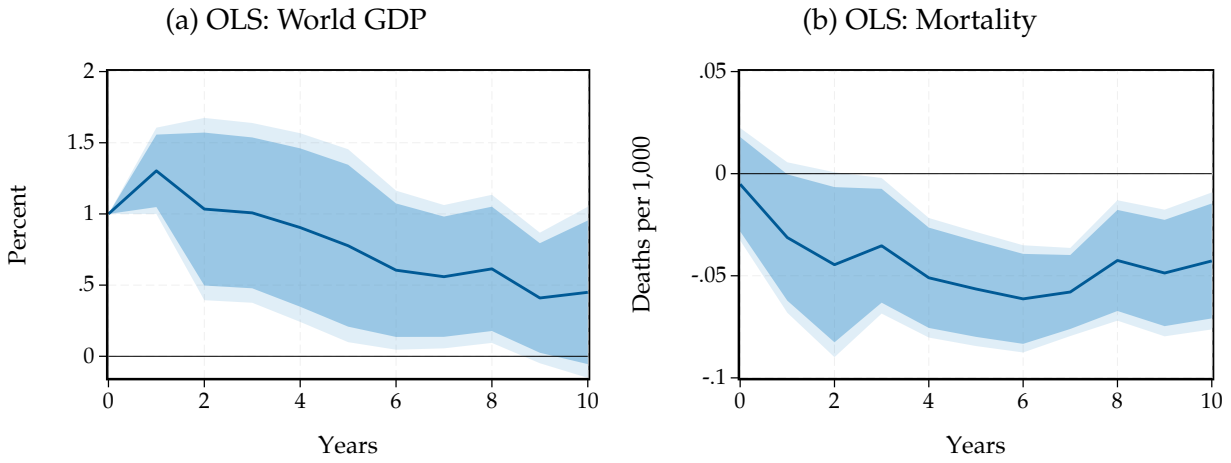
Notes: Impulse responses of global temperature and world GDP per capita to a 1°C global temperature shock, from Bilal and Känzig (2026) extended to 20 years. Solid line: point estimates. Shaded areas: 90% and 95% confidence bands based on Driscoll-Kray standard errors.

D.2.3 Mortality and GDP

Figure D.5 reproduces the results from Bilal, Känzig, and Lisiecki (2026). The top panels display Ordinary Least Squares (OLS) estimates using global shocks extracted with Hamilton (2018) filter. Panel (a) shows the dynamic correlation between a 1% GDP per capita shock on GDP per capita. Panel (b) shows the dynamic correlation between the same shock on mortality.

We compute the effect of a permanent increase in global GDP per capita on mortality by taking the ratio of cumulative impulse responses. The OLS estimates imply that a 100 log points increase in world GDP per capita lowers mortality by 4.5 deaths per 1,000 individuals in the long run.

Figure D.5: GDP and Mortality: OLS Estimates (Bilal, Känzig, and Lisiecki, 2026)

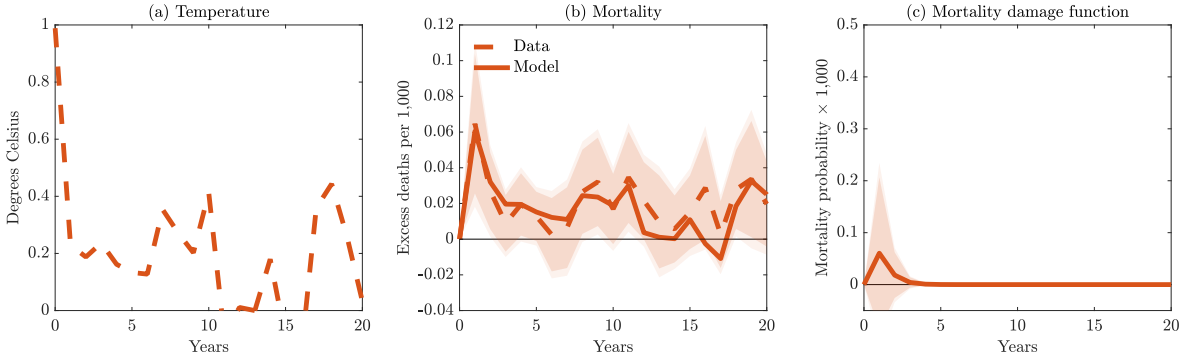


Notes: OLS estimates using global shocks extracted with Hamilton (2018) filter. Panel (a): dynamic correlation between a 1% GDP per capita shock and GDP per capita. Panel (b): dynamic correlation between the same shock and mortality. Solid line: point estimates. Shaded areas: 90% and 95% confidence bands based on Driscoll-Kray standard errors.

E Additional Model Results

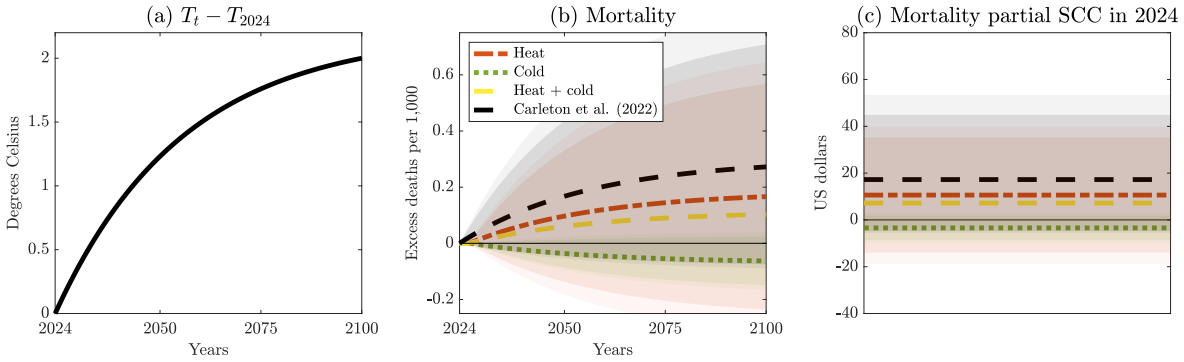
E.1 Mortality Damage Functions From Extreme Heat

Figure E.1: Mortality Damage Functions From Local Heat



Notes: Estimated mortality damage functions obtained by matching model-implied impulse responses to the empirical responses to local heat shocks. Dashed lines denote empirical responses; solid lines denote model fits. Dark and light shaded areas represent 90 and 95% confidence bands, computed using the Delta method.

Figure E.2: Mortality Damage Functions From Local Heat



Notes: Transitional dynamics of the model given the warming scenario in Panel (a) based on alternative damage function specifications: local heat (red dashed), local cold (green dashed), Carleton et al. (2022) damage function (black dotted). Shaded areas: 90 and 95% confidence intervals based on the Delta-method.

References Appendix

Bilal, Adrien and Diego R. Känzig (2026). "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature". *Forthcoming, The Quarterly Journal of Economics* 32450.

Bilal, Adrien, Diego R. Känzig, and Krzysztof Lisiecki (2026). "Life Expectancy and Economic Growth: Local or Global GDP?" *Working Paper*.

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 (2022). "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". *The Quarterly Journal of Economics* 137.4, pp. 2037–2105.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (2015). "The Next Generation of the Penn World Table". *American Economic Review* 105.10, pp. 3150–3182.

Hamilton, James D. (2018). "Why you should never use the Hodrick-Prescott filter". *Review of Economics and Statistics* 100.5, pp. 831–843.

Human Mortality Database (2024). *Human Mortality Database*. Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France). Accessed: 2024-10-11.

Kose, M. Ayhan, Naotaka Sugawara, and Marco E. Terrones (2020). "Global recessions".

Lange, Stefan, Matthias Mengel, Simon Treu, and Matthias Büchner (2023). *ISIMIP3a atmospheric climate input data (v1.2)*. ISIMIP Repository. https://doi.org/10.48364/ISIMIP.982724.2.

NOAA National Centers for Environmental Information (2023). *Climate at a Glance: Global Time Series*. Online Resource.

Rohde, Robert A. and Zeke Hausfather (2020). "The Berkeley Earth Land/Ocean Temperature Record". *Earth System Science Data* 12.4, pp. 3469–3479.

Sims, Christopher A. (1986). "Are forecasting models usable for policy analysis?" *Quarterly Review* 10, pp. 2–16.

United Nations, Department of Economic and Social Affairs Population Division (2024). *World Population Prospects 2024: Online Edition*. https://population.un.org/wpp/. Accessed: 2024-10-11.

World Health Organization (2024). *WHO Mortality Database*. https://www.who.int/data/data-collection-tools/who-mortality-database. Accessed: 2024-08-19.