DISCUSSION PAPER SERIES

DP15139
(v. 2)

Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits

Tamma Carleton, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon M. Hsiang, Andrew Hultgren, Robert Kopp, Kelly McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon (Samuel) Seo, Arvid Viaene, Jiacan Yuan and Alice Tianbo Zhang

DEVELOPMENT ECONOMICS

CEPR

ISSN 0265-8003

Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits

Tamma Carleton, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon M. Hsiang, Andrew Hultgren, Robert Kopp, Kelly McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon (Samuel) Seo, Arvid Viaene, Jiacan Yuan and Alice Tianbo Zhang

Discussion Paper DP15139
First Published 06 August 2020
This Revision 15 August 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits

Abstract

Using 40 countries' subnational data, we estimate age-specific mortality-temperature relationships and extrapolate them to countries without data today and into a future with climate change. We uncover a U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for the elderly. Critically, this relationship is flattened by both higher incomes and adaptation to local climate. Using a revealed preference approach to recover unobserved adaptation costs, we estimate that the mean global increase in mortality risk due to climate change, accounting for adaptation benefits and costs, is valued at roughly 3.2% of global GDP in 2100 under a high emissions scenario. Notably, today's cold locations are projected to benefit, while today's poor and hot locations have large projected damages. Finally, our central estimates indicate that the release of an additional ton of CO2 today will cause mortality-related damages of $36.6 under a high emissions scenario and using a 2% discount rate, with an interquartile range accounting for both econometric and climate uncertainty of [-$7.8, $73.0]. Under a moderate emissions scenario, these damages are valued at $17.1 [-$24.7, $53.6]. These empirically grounded estimates exceed the previous literature's estimates by an order of magnitude.

JEL Classification: Q51, Q54, H23, H41, I14

Keywords: N/A

Tamma Carleton - tcarleton@uchicago.edu
University of Chicago

Amir Jina - amirjina@uchicago.edu
University of Chicago

Michael Delgado - mdelgado@rhg.com
Rhodium Group

Michael Greenstone - mgreenst@uchicago.edu
University of Chicago, Becker Friedman Institute and CEPR

Trevor Houser - tghouser@rhg.com
Rhodium Group

Solomon M. Hsiang - shsiang@berkeley.edu
University of California Berkeley and Stanford University and CEPR

Andrew Hultgren - hultgren@berkeley.edu
UC Berkeley

Robert Kopp - robert.kopp@rutgers.edu
Rutgers University

Kelly McCusker - kmccusker@rhg.com
Rhodium Group

Ishan Nath - inath@uchicago.edu
University of Chicago

James Rising - jarising@gmail.com
London School of Economics

Ashwin Rode - aarode@uchicago.edu
University of Chicago

Hee Kwon (Samuel) Seo - hksseo@gmail.com
Chicago Booth School of Business

Arvid Viaene - arvid.viaene.db@gmail.com
Analysis Group

Jiacan Yuan - jiacan.yuan@gmail.com
Rutgers University

Alice Tianbo Zhang - alice.zhang@columbia.edu
Columbia University

Acknowledgements
This project is an output of the Climate Impact Lab that gratefully acknowledges funding from the Energy Policy Institute of Chicago (EPIC), International Growth Centre, National Science Foundation, Sloan Foundation, Carnegie Corporation, and Tata Center for Development. Tamma Carleton acknowledges funding from the US Environmental Protection Agency Science To Achieve Results Fellowship (#FP91780401). We thank Trinetta Chong, Greg Dobbels, Diana Gergel, Radhika Goyal, Simon Greenhill, Hannah Hess, Dylan Hogan, Azhar Hussain, Stefan Klos, Theodor Kulczycki, Brewster Malevich, Sebastien Annan Phan, Justin Simcock, Emile Tenezakis, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and Megan Landin, Terin Mayer, and Jack Chang for excellent project management. We thank David Anthoff. Max Auffhammer, Olivier Deschenes, Avi Ebenstein, Nolan Miller, Wolfram Schlenker, and numerous workshop participants at University of Chicago, Stanford, Princeton, UC Berkeley, UC San Diego, UC Santa Barbara, University of Pennsylvania, University of San Francisco, University of Virginia, University of Wisconsin-Madison, University of Minnesota Twin Cities, NBER Summer Institute, LSE, PIK, Oslo University, University of British Columbia, Gothenburg University, the European Center for Advanced Research in Economics and Statistics, the National Academies of Sciences, and the Econometric Society for comments, suggestions, and help with data.

Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*

Tamma Carleton[1,2], Amir Jina[3,2], Michael Delgado[4], Michael Greenstone[3,2], Trevor Houser[4], Solomon Hsiang[5,2], Andrew Hultgren[3], Robert Kopp[6], Kelly E. McCusker[4], Ishan Nath[7], James Rising[8], Ashwin Rode[3], Hee Kwon Seo[9], Arvid Viaene[10], Jiacan Yuan[11], and Alice Tianbo Zhang[12]

[1]University of California, Santa Barbara
[2]NBER
[3]University of Chicago
[4]Rhodium Group
[5]University of California, Berkeley
[6]Rutgers University
[7]Princeton University
[8]University of Maryland
[9]The World Bank
[10]E.CA Economics
[11]Fudan University
[12]Washington and Lee University

14th August, 2021

*This project is an output of the Climate Impact Lab that gratefully acknowledges funding from the Energy Policy Institute of Chicago (EPIC), International Growth Centre, National Science Foundation, Sloan Foundation, Carnegie Corporation, and Tata Center for Development. Tamma Carleton acknowledges funding from the US Environmental Protection Agency Science To Achieve Results Fellowship (#FP91780401). We thank Trinetta Chong, Greg Dobbels, Diana Gergel, Radhika Goyal, Simon Greenhill, Hannah Hess, Dylan Hogan, Azhar Hussain, Stefan Klos, Theodor Kulczycki, Brewster Malevich, Sébastien Annan Phan, Justin Simcock, Emile Tenezakis, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and Megan Landín, Terin Mayer, and Jack Chang for excellent project management. We thank David Anthoff, Max Auffhammer, Olivier Deschênes, Avi Ebenstein, Nolan Miller, Wolfram Schlenker, and and numerous workshop participants at University of Chicago, Stanford, Princeton, UC Berkeley, UC San Diego, UC Santa Barbara, University of Pennsylvania, University of San Francisco, University of Virginia, University of Wisconsin-Madison, University of Minnesota Twin Cities, NBER Summer Institute, LSE, PIK, Oslo University, University of British Columbia, Gothenburg University, the European Center for Advanced Research in Economics and Statistics, the National Academies of Sciences, and the Econometric Society for comments, suggestions, and help with data.

Abstract

Using 40 countries' subnational data, we estimate age-specific mortality-temperature relationships and extrapolate them to countries without data today and into a future with climate change. We uncover a U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for the elderly. Critically, this relationship is flattened by both higher incomes and adaptation to local climate. Using a revealed preference approach to recover unobserved adaptation costs, we estimate that the mean global increase in mortality risk due to climate change, accounting for adaptation benefits and costs, is valued at roughly 3.2% of global GDP in 2100 under a high emissions scenario. Notably, today's cold locations are projected to benefit, while today's poor and hot locations have large projected damages. Finally, our central estimates indicate that the release of an additional ton of CO_2 today will cause mortality-related damages of $36.6 under a high emissions scenario and using a 2% discount rate, with an interquartile range accounting for both econometric and climate uncertainty of [-$7.8, $73.0]. Under a moderate emissions scenario, these damages are valued at $17.1 [-$24.7, $53.6]. These empirically grounded estimates exceed the previous literature's estimates by an order of magnitude.

JEL Codes: Q51, Q54, H23, H41, I14.

1 Introduction

Understanding the likely global economic impacts of climate change is of tremendous practical value to both policymakers and researchers. On the policy side, decisions are currently made with incomplete and inconsistent information on the benefits of greenhouse gas emissions reductions. These inconsistencies are reflected in global climate policy, which is at once both lenient and wildly inconsistent. To date, the economics literature has struggled to mitigate this uncertainty, lacking empirically founded and globally comprehensive estimates of the total burden imposed by climate change that account for the benefits and costs of adaptation. This problem is made all the more difficult because emissions today influence the global climate for hundreds of years. Thus, any reliable estimate of the damage from climate change must include projections of economic impacts that are both long-run and at global scale.

Decades of study have accumulated numerous theoretical and empirical insights and important findings regarding the economics of climate change, but a fundamental gulf persists between the two main types of analyses. On the one hand, there are stylized models able to capture the multi-century and global nature of climate change, such as "integrated assessment models" (IAMs) (e.g., Nordhaus, 1992; Tol, 1997; Stern, 2006); their great appeal is that they provide an answer to the question of what the global costs of climate change will be. However, IAMs require many assumptions and this weakens the authority of their answers. On the other hand, there has been an explosion of highly resolved empirical analyses whose credibility lies in their use of real world data and careful econometric measurement (e.g., Schlenker and Roberts, 2009; Deschênes and Greenstone, 2007). Yet these analyses tend to be limited in geographic extent and/or rely on short-run changes in weather that are unlikely to fully account for adaptation to gradual climate change (Hsiang, 2016). At its core, this dichotomy persists because researchers have traded off between being complete in scale and scope or investing heavily in data collection and analysis.

This paper aims to resolve the tension between these approaches by providing empirically-derived estimates of climate change's impacts on global mortality risk. Importantly, these estimates are at the scale of IAMs, yet grounded in detailed econometric analyses using high-resolution globally representative data, and account for adaptation to gradual climate change. The analysis proceeds in three steps that lead to the paper's three main findings.

First, we estimate regressions to infer age-specific mortality-temperature relationships using historical data. These regressions are fit on the most comprehensive dataset ever collected on annual, subnational mortality statistics from 40 countries that cover 38% of the global population. The benefits of adaptation to climate change and the benefits of projected future income growth are estimated by allowing the mortality-temperature response function to vary with long-run climate (e.g., Auffhammer, 2018) and income per capita (e.g., Fetzer, 2014). This modeling of heterogeneity allows us to predict the structure of the mortality-temperature relationship across locations where we lack mortality data, yielding estimates for the entire world.

These regressions uncover a plausibly causal U-shaped relationship where extremely cold and hot temperatures increase mortality rates, especially for those aged 65 and older. Moreover, this relationship is quite heterogeneous across the planet: we find that both income and long-run climate substantially moderate mortality sensitivity to temperature. When we combine these results with current global data on climate, income, and population, we find that the effect of an additional very hot day (35°C / 95°F) on mortality in the >64 age group is ∼50% larger in regions of the world where mortality data are unavailable. This finding suggests that prior estimates may understate climate change impacts, because they disproportionately rely on data from wealthy economies and temperate climates. However, we note that because modern populations have not experienced multiple alternative climates, the estimates of heterogeneity rely on cross-sectional variation and they must be considered associational.

Second, we combine the regression results with standard future predictions of climate, income and population to project future climate change-induced mortality risk both in terms of fatality rates and its monetized value. The paper's mean estimate of the projected increase in the global mortality rate due to climate change is 73 deaths per 100,000 at the end of the century under a high emissions scenario (i.e., Representative Concentration Pathway (RCP) 8.5), with an interquartile range of [6, 101] due both to econometric and climate uncertainty. This effect is similar in magnitude to the current global mortality burden of all cancers or all infectious diseases. It is noteworthy that these impacts are predicted to be unequally distributed across the globe: for example, mortality rates in Accra, Ghana are projected to increase by 17% at the end of the century under a high emissions scenario, while in London, England, mortality rates are projected to decrease by 8% due to milder winters. Importantly, a failure to account for climate adaptation and the benefits of income growth would lead to overstating the mortality costs of climate change by a factor of about 3.

Of course, adaptation is costly; we develop a stylized revealed preference model that leverages observed differences in temperature sensitivity across space to infer these costs. When monetizing projected deaths due to climate change with the value of a statistical life (VSL) and adding the estimated costs of adaptation, the total mortality burden of climate change is equal to roughly 3.2% of global GDP at the end of the century under a high emissions scenario. We find that poor countries are projected to disproportionately experience impacts through deaths, while wealthy countries experience impacts largely through costly adaptation investments.

Third, we use these estimates to compute the global marginal willingness-to-pay (MWTP) to avoid the alteration of mortality risk associated with the temperature change from the release of an additional metric ton of CO_2. We call this the excess mortality "partial" social cost of carbon (SCC); a "full" SCC would encompass impacts across all affected outcomes. Our estimates imply that the excess mortality partial SCC is roughly \$36.6 [-\$7.8, \$73.0] (in 2019 USD) with a high emissions scenario (RCP8.5) under a 2% discount rate and using an age-varying VSL. This value falls to \$17.1 [-\$24.7, \$53.6] with a moderate emissions scenario (RCP4.5). The excess mortality partial SCC is lower in this scenario because the relationship between mortality risk and temperature is convex, meaning that marginal damages are greater under higher baseline emissions.

Overall, this paper's results suggest that the temperature related mortality risk from climate change is substantially greater than previously understood. For example, the estimated mortality partial SCC is more than an order of magnitude larger than the partial SCC for all health impacts embedded in the FUND IAM. Further, under the high emissions scenario, the estimated excess mortality partial SCC is ∼72% of the Biden

Administration's *full* interim SCC.[1]

In generating these results, this paper overcomes multiple challenges that have plagued the previous literature. The first challenge is that CO_2 is a global pollutant, so it is necessary to account for the heterogeneous costs of climate change across the entire planet. The second challenge is that today, there is substantial adaptation to climate, as people successfully live in both Houston, TX and Anchorage, AK, and climate change will undoubtedly lead to new adaptations in the future. The extent to which investments in adaptation can limit the impacts of climate change is a critical component of damage estimates. We address both of these challenges by combining extensive data with an econometric approach that models heterogeneity in the mortality-temperature relationship, allowing us to predict mortality-temperature relationships at high resolution globally and into the future as climate and incomes evolve. Specifically, we develop estimates of climate change impacts at high resolution, effectively allowing for 24,378 representative agents. In contrast, the previous literature has assumed the world is comprised of, at maximum, 170 heterogenous regions (Burke, Hsiang, and Miguel, 2015), but typically far fewer (Nordhaus and Yang, 1996; Tol, 1997).

A final challenge is that adaptation responses are costly, and these costs must be accounted for in a full assessment of climate change impacts. While our revealed preference approach to inferring adaptation costs relies on a strong set of simplifying assumptions, it can be directly estimated with available data and represents an important advance on previous literature, which has either quantified adaptation benefits without estimating costs (e.g., Heutel, Miller, and Molitor, 2017) or tried to measure the costs of individual adaptive investments in selected locations (e.g., Barreca et al., 2016), an approach that is poorly equipped to capture the wide range of potential responses to warming.

The rest of this paper is organized as follows: Section 2 provides definitions and some basic intuition for the economics of adaptation to climate change in the context of mortality. Section 3 details data used throughout the analysis. Section 4 describes our empirical model and estimations results. Section 5 presents projections of climate change impacts with and without the benefits of adaptation. Section 6 outlines a revealed preference approach that allows us to infer adaptation costs and uses this framework to present empirically-derived projections of the mortality risk of climate change accounting for the costs and benefits of adaptation. Section 7 constructs a partial SCC, Section 8 discusses key limitations of the analysis, and Section 9 concludes.

2 Conceptual framework

This section sets out a simple conceptual framework that guides the empirical model the paper uses to estimate society's willingness to pay (WTP) to avoid the mortality risks from climate change. In estimating these mortality risks, it is critical to account for individuals' compensatory responses, or adaptations, to climate change, such as investments in air conditioning. These adaptations have both benefits that reduce the risks of extreme temperatures and costs in the form of foregone consumption. Thus, the full mortality risk of climate change is the sum of changes in mortality rates after accounting for adaptation and the costs

[1] This comparison is made using our preferred valuation scenario, which includes an age-adjusted VSL and a discount rate of 2%. The Biden Administration's interim SCC uses a 3% discount rate and an age-invariant VSL. Under these valuation assumptions, the estimated excess mortality partial SCC is 44% of the Biden Administration's full interim SCC.

of those adaptations. Here, we define some key objects that the paper will estimate, including the full value of mortality risk due to climate change.

We define the *climate* as the joint probability distribution over a vector of possible conditions that can be expected to occur over a specific interval of time. Following the notation of Hsiang (2016), let $\boldsymbol{C}$ be a vector of parameters describing the entire joint probability distribution over all relevant climatic variables (e.g., $\boldsymbol{C}$ might contain the mean and variance of daily average temperature and rainfall, among other parameters). We define weather realizations as a random vector $\boldsymbol{c}$ drawn from a distribution characterized by $\boldsymbol{C}$. Mortality risk is a function of both weather $\boldsymbol{c}$ and a composite good $\boldsymbol{b} = \xi(b_1, ..., b_K)$ comprising all choice variables b_k that could influence mortality risk, such as installation of air conditioning and time allocated to indoor activities. The endogenous choices in $\boldsymbol{b}$ are the outcome of a stylized model in which individuals maximize expected utility by trading off consumption of a numeraire good and $\boldsymbol{b}$, subject to a budget constraint, as outlined in detail in Section 6. Mortality risk is then captured by the probability of death $f = f(\boldsymbol{b}, \boldsymbol{c})$.

Climate change will influence mortality risk through two pathways.[2] First, a change in $\boldsymbol{C}$ will directly alter realized weather draws, changing $\boldsymbol{c}$. Second, a change in $\boldsymbol{C}$ can alter individuals' beliefs about their likely weather realizations, shifting how they act, and ultimately changing their endogenous choice variables $\boldsymbol{b}$. Endogenous adjustments to $\boldsymbol{b}$ therefore capture all long-run adaptation to the climate (e.g., Mendelsohn, Nordhaus, and Shaw, 1994; Kelly, Kolstad, and Mitchell, 2005). Since the climate $\boldsymbol{C}$ determines both $\boldsymbol{c}$ and $\boldsymbol{b}$, the probability of death at an initial climate $\boldsymbol{C}_{t_0}$ is written as:

$$\Pr(\text{death} \mid \boldsymbol{C}_{t_0}) = f(\boldsymbol{b}(\boldsymbol{C}_{t_0}), \boldsymbol{c}(\boldsymbol{C}_{t_0})), \tag{1}$$

where $\boldsymbol{c}(\boldsymbol{C})$ is a random vector $\boldsymbol{c}$ drawn from the distribution characterized by $\boldsymbol{C}$.

Many previous empirical estimates assume that individuals do not make any adaptations or compensatory responses to an altered climate (e.g., Deschênes and Greenstone, 2007; Houser et al., 2015). Under this approach, the change in mortality risk incurred due to a change in climate from $\boldsymbol{C}_{t_0}$ to $\boldsymbol{C}_t$ is calculated as:

$$\textit{mortality effects of climate change without adaptation} = f(\boldsymbol{b}(\boldsymbol{C}_{t_0}), \boldsymbol{c}(\boldsymbol{C}_t)) - f(\boldsymbol{b}(\boldsymbol{C}_{t_0}), \boldsymbol{c}(\boldsymbol{C}_{t_0})), \tag{2}$$

which ignores the fact that individuals will choose new values of $\boldsymbol{b}$ as their beliefs about $\boldsymbol{C}$ evolve.

A more realistic estimate for the change in mortality due to a change in climate is:

$$\textit{mortality effects of climate change with adaptation} = f(\boldsymbol{b}(\boldsymbol{C}_t), \boldsymbol{c}(\boldsymbol{C}_t)) - f(\boldsymbol{b}(\boldsymbol{C}_{t_0}), \boldsymbol{c}(\boldsymbol{C}_{t_0})). \tag{3}$$

If the climate is changing such that the mortality risk from $\boldsymbol{C}_t$ is higher than $\boldsymbol{C}_{t_0}$ when holding $\boldsymbol{b}$ fixed, then the endogenous adjustment of $\boldsymbol{b}$ will generate benefits of adaptation weakly greater than zero, since these damages may be partially mitigated. In practice, the sign of the difference between Equations 2 and 3 will depend on the degree to which climate change reduces extremely cold days versus increases extremely hot days, and the optimal adaptation that agents undertake in response to these competing changes.

Several analyses have estimated reduced-form versions of Equation 3, confirming that accounting for

[2] Hsiang (2016) describes these two channels as a "direct effect" and a "belief effect."

endogenous changes to technology, behavior, and investment mitigates the direct effects of climate in a variety of contexts (e.g., Barreca et al., 2016).[3] Importantly, however, while this approach accounts for the *benefits* of adaptation, it does not account for its *costs*. If adjustments to $\boldsymbol{b}$ were costless and provided protection against the climate, then we would expect universal uptake of highly adapted values for $\boldsymbol{b}$ so that temperature would have no effect on mortality. But we do not observe this to be true: for example, Heutel, Miller, and Molitor (2017) find that the mortality effects of extremely hot days in warmer climates (e.g., Houston) are much smaller than in more temperature climates (e.g., Seattle).[4] We denote the costs of achieving adaptation level $\boldsymbol{b}$ as $A(\boldsymbol{b})$, measured in dollars of forgone consumption.

A full measure of the economic burden of climate change must account not only for the benefits generated by compensatory responses to these changes, but also their cost. Thus, the total cost of changing mortality risks that result from a climate change $\boldsymbol{C}_{t_0} \rightarrow \boldsymbol{C}_t$ is:

$$\textit{full value of mortality risk due to climate change} = \\ VSL \underbrace{\left[f(\boldsymbol{b}(\boldsymbol{C}_t), \boldsymbol{c}(\boldsymbol{C}_t)) - f(\boldsymbol{b}(\boldsymbol{C}_{t_0}), \boldsymbol{c}(\boldsymbol{C}_{t_0}))\right]}_{\text{observable change in mortality rate}} + \underbrace{A(\boldsymbol{b}(\boldsymbol{C}_t)) - A(\boldsymbol{b}(\boldsymbol{C}_{t_0}))}_{\text{adaptation costs}}, \tag{4}$$

where VSL is the value of a statistical life. It is apparent that omitting the costs of adaptation, $A(\boldsymbol{b})$, would lead to an incomplete measure of the full costs of mortality risk due to climate change.

This paper develops an empirical model to quantify climate change's impact on mortality risk at global scale, accounting for the benefits of adaptation, consistent with Equation 3. Throughout the analysis, we consider the effects of climate change induced changes in daily average temperature, such that the mortality risk of climate change implies effects of temperature only (as opposed to other climate variables, such as precipitation). Because income may also influence the choice variables in $\boldsymbol{b}$, we include the benefits of income growth in this empirical model, in addition to the benefits of climate adaptation. This empirical approach and the resulting climate change impact projections are detailed in Sections 4 and 5, respectively.

However, an empirical estimation of the full value of mortality risk due to climate change, shown in Equation 4, is more difficult, as total changes in adaptation costs between time periods cannot be observed directly. In principle, data on each adaptive action could be gathered and modeled (e.g., Deschênes and Greenstone, 2011), but since there exists an enormous number of possible adaptive margins that together make up the vector $\boldsymbol{b}$, computing the full cost of climate change using such an enumerative approach quickly becomes intractable. To make progress on quantifying the full value of mortality risk due to climate change, we develop a stylized revealed preference approach that leverages observed differences in climate sensitivity across locations to infer adaptation costs associated with the mortality risk from climate change. This approach, and resulting estimates of the full (monetized) value of the mortality risk due to climate change, are reported in Section 6.

Section 7 uses these estimates to compute the global marginal willingness-to-pay (MWTP) to avoid the alteration of mortality risk associated with the release of an additional metric ton of CO_2. We call this the

[3]For additional examples, see Schlenker and Roberts (2009); Hsiang and Narita (2012); Hsiang and Jina (2014); Barreca et al. (2015); Heutel, Miller, and Molitor (2017); Auffhammer (2018).

[4]Carleton and Hsiang (2016) document that such wedges in observed sensitivities to climate—which they call "adaptation gaps"—are a pervasive feature of the broader climate damages literature.

excess mortality "partial" social cost of carbon (SCC); a "full" SCC would encompass impacts across all affected sectors (e.g., labor productivity, damages from sea level rise, etc.).

3 Data

To estimate the mortality risks of climate change at global scale, we assemble a novel dataset composed of rich historical mortality records, high-resolution historical climate data, and future projections of climate, population, and income across the globe. Section 3.1 describes the data necessary to estimate $f(\boldsymbol{b}, \boldsymbol{c})$, the relationship between mortality and temperature, accounting for differences in climate and income. Section 3.2 outlines the data we use to predict the mortality-temperature relationship across the entire planet today and project its evolution into the future as populations adapt to climate change. Appendix B provides a more extensive description of each of these datasets.

3.1 Data to estimate the mortality-temperature relationship

3.1.1 Mortality data

Our mortality data are collected independently from 40 countries.[5] Combined, this dataset covers mortality outcomes for 38% of the global population, representing a substantial increase in coverage relative to existing literature; prior studies investigate an individual country (e.g., Burgess et al., 2017) or region (e.g., Deschenes, 2018), or combine small nonrandom samples from across multiple countries (e.g., Gasparrini et al., 2015). Table 1 summarizes each dataset, while spatial coverage, resolution, and temporal coverage are shown in Figure B1. We harmonize all records into a single multi-country unbalanced panel dataset of age-specific annual mortality rates, using three age categories: <5, 5-64, and >64, where the unit of observation is ADM2 (e.g., a county in the U.S.) by year.

3.1.2 Historical climate data

The analysis is performed with two separate groups of historical data on precipitation and temperature. First, we use the Global Meteorological Forcing Dataset (GMFD) (Sheffield, Goteti, and Wood, 2006), which relies on a weather model in combination with observational data. Second, we repeat our analysis with climate datasets that strictly interpolate observational data across space onto grids, combining temperature data from the daily Berkeley Earth Surface Temperature dataset (BEST) (Rohde et al., 2013) with precipitation data from the monthly University of Delaware dataset (UDEL) (Matsuura and Willmott, 2007). Table 1 summarizes these data; full data descriptions are provided in Appendix B.2. We link climate and mortality data by aggregating gridded daily temperature data to the annual measures at the same administrative level as the mortality records (i.e., ADM2) using a procedure detailed in Appendix B.2.4 that allows for the recovery of potential nonlinearities in the mortality-temperature relationship.

[5]We additionally use data from India as cross-validation of our main results, as the India data do not have records of age-specific mortality rates. The inclusion of India increases our data coverage to 55% of the global population.

Table 1: Historical mortality & climate data

Mortality records

Country	N	Spatial scale×	Years	Age categories	Average annual mortality rate*† All-age	>64 yr.	Global pop. share◇	Average covariate values*□ GDP per capita⊗	Avg. daily temp.⊘	Annual avg. days > 28°C
Brazil	228,762	ADM2	1997-2010	<5, 5-64, >64	525	4,096	0.028	11,192	23.8	35.2
Chile	14,238	ADM2	1997-2010	<5, 5-64, >64	554	4,178	0.002	14,578	14.3	0
China	7,488	ADM2	1991-2010	<5, 5-64, >64	635	7,507	0.193	4,875	15.1	25.2
EU	13,013	NUTS2‡	1990▷-2010	<5, 5-64, >64	1,014	5,243	0.063	22,941	11.2	1.6
France⊕	3,744	ADM2	1998-2010	0-19, 20-64, >64	961	3,576	0.009	31,432	11.9	0.3
India^	12,505	ADM2	1957-2001	All-age	724	–	0.178	1,355	25.8	131.4
Japan	5,076	ADM1	1975-2010	<5, 5-64, >64	788	4,135	0.018	23,241	14.3	8.3
Mexico	146,835	ADM2	1990-2010	<5, 5-64,>64	561	4,241	0.017	16,518	19.1	24.6
USA	401,542	ADM2	1968-2010	<5, 5-64, >64	1,011	5,251	0.045	30,718	13	9.5
All Countries	**833,203**	–	–	–	**780**	**4,736**	**0.554**	**20,590**	**15.5**	**32.6**

Historical climate datasets

Dataset	Citation	Method	Resolution	Variable	Source
GMFD, V1	Sheffield, Goteti, and Wood (2006)	Reanalysis & Interpolation	0.25°	temp. & precip.	Princeton University
BEST	Rohde et al. (2013)	Interpolation	1°	temp.	Berkeley Earth
UDEL	Matsuura and Willmott (2007)	Interpolation	0.5°	precip.	University of Delaware

*In units of deaths per 100,000 population.
†To remove outliers, particularly in low-population regions, we winsorize the mortality rate at the 1% level at high end of the distribution across administrative regions, separately for each country.
□ All covariate values shown are averages over the years in each country sample.
× ADM2 refers to the second administrative level (e.g., county), while ADM1 refers to the first administrative level (e.g., state). NUTS2 refers to the Nomenclature of Territorial Units for Statistics 2^{nd} (NUTS2) level, which is specific to the European Union (EU) and falls between first and second administrative levels.
◇ Global population share for each country in our sample is shown for the year 2010.
⊗ GDP per capita values shown are in constant 2005 dollars purchasing power parity (PPP).
⊘ Average daily temperature and annual average of the number of days above 28°C are both population weighted, using population values from 2010.
‡ EU data for 33 countries were obtained from a single source. Detailed description of the countries within this region is presented in Appendix B.1.
▷ Most countries in the EU data have records beginning in the year 1990, but start dates vary for a small subset of countries. See Appendix B.1 and Table B1 for details.
⊕ We separate France from the rest of the EU, as higher resolution mortality data are publicly available for France.
^ It is widely believed that data from India understate mortality rates due to incomplete registration of deaths.

3.1.3 Covariate data

The analysis allows for heterogeneity in the age-specific mortality-temperature relationship as a function of two long-run covariates: a measure of climate (in our main specification, long-run average temperature) and income per capita. We assemble time-invariant measures of both these variables at the ADM1 unit (e.g., state) level using GMFD climate data and a combination of the Penn World Tables (PWT), Gennaioli et al. (2014), and Eurostat (2013). These covariates are measured at ADM1 scale (as opposed to the ADM2 scale of the mortality records) due to limited availability of higher resolution income data. The construction of the income variable requires some estimation to downscale to ADM1 level; details on this procedure are provided in Appendix B.3.

In a set of robustness checks detailed in Section 4.2 and Appendix D.6, we analyze five additional sources of heterogeneity, each of which has been suggested in the literature as an important driver of long-run wellbeing (Alesina and Rodrik, 1994; Glaeser et al., 2004; La Porta and Shleifer, 2014; Bailey and Goodman-Bacon, 2015; World Bank, 2020). These data include country-by-year obvservations of institutional quality from the Center for Systemic Peace (2020), access to healthcare services and labor force informality from the World Bank (2020), educational attainment from the World Bank (2020) and Organization of Economic Cooperaton and Development (2020), and within-country income inequality from the World Inequality Lab (2020).

3.2 Data for projecting the mortality-temperature relationship around the world & into the future

3.2.1 Unit of analysis for projections

We partition the global land surface into a set of 24,378 regions and for each region we generate location-specific projected damages of climate change. The finest level of disaggregation in previous estimates of global climate change damages divides the world into 170 regions (Burke, Hsiang, and Miguel, 2015), but most papers account for much less heterogeneity (Nordhaus and Yang, 1996; Tol, 1997). These regions (hereafter, *impact regions*) are constructed such that they are either identical to, or are a union of, existing administrative regions. They (i) respect national borders, (ii) are roughly equal in population across regions, and (iii) display approximately homogenous within-region climatic conditions. Appendix C details the algorithm used to create impact regions.

3.2.2 Climate projections

We use a set of 21 high-resolution, bias-corrected, global climate projections produced by NASA Earth Exchange (NEX) (Thrasher et al., 2012)[6] that provide daily temperature and precipitation through the year 2100. We obtain climate projections based on two standardized emissions scenarios: Representative Concentration Pathways 4.5 (RCP4.5, an emissions stabilization scenario) and 8.5 (RCP8.5, a scenario with intensive growth in fossil fuel emissions) (Van Vuuren et al., 2011; Thomson et al., 2011)).

These 21 climate models systematically underestimate tail risks of future climate change (Tebaldi and Knutti, 2007; Rasmussen, Meinshausen, and Kopp, 2016).[7] To correct for this, we follow Hsiang et al. (2017) by assigning probabilistic weights to climate projections and use 12 surrogate models that describe local climate outcomes in the tails of the climate sensitivity distribution (Rasmussen, Meinshausen, and Kopp, 2016). Figure B2 shows the resulting weighted climate model distribution. The 21 models and 12 surrogate models are treated identically in our calculations and we describe them collectively as the surrogate/model mixed ensemble (SMME). Gridded output from these 33 projections are aggregated to impact regions; full details on the climate projection data are in Appendix B.2.

Only 6 of the 21 models we use to construct the SMME provide climate projections after 2100 for both high and moderate emissions scenarios, and none simulate the impact of a marginal ton of CO_2.

[6] The dataset we use, called the NEX-GDDP, downscales global climate model (GCM) output from the Coupled Model Intercomparison Project Phase 5 (CMIP5) archive (Taylor, Stouffer, and Meehl, 2012), an ensemble of models typically used in national and international climate assessments.

[7] The underestimation of tail risks in the 21-model ensemble is for several reasons, including that these models form an ensemble of opportunity and are not designed to sample from a full distribution, they exhibit idiosyncratic biases, and have narrow tails. We are correcting for their bias and narrowness with respect to global mean surface temperature (GMST) projections, but our method does not correct for all biases.

Therefore, to include post-2100 years in our estimates of the mortality partial SCC, we rely on the Finite Amplitude Impulse Response (FAIR) simple climate model, which has been developed especially for this type of calculation (Millar et al., 2017).[8] Details on our implementation of FAIR are in Appendix G.

3.2.3 Socioeconomic projections

Projections of population and income are a critical ingredient in the analysis, and for these we rely on the Shared Socioeconomic Pathways (SSPs), which describe a set of plausible scenarios of socioeconomic development over the 21^{st} century. We use SSP2, SSP3, and SSP4, which yield emissions in the absence of mitigation policy that fall between RCP4.5 and RCP8.5 in integrated assessment modeling exercises (Riahi et al., 2017). For population, we use the International Institute for Applied Systems Analysis (IIASA) SSP population projections, which provide estimates of population by age cohort at country-level in five-year increments (IIASA Energy Program, 2016). National population projections are allocated to impact regions based on current satellite-based within-country population distributions from Bright et al. (2012) (see Appendix B.3.3). Projections of national income per capita are similarly derived from the SSP scenarios, using both the IIASA projections and the Organization for Economic Co-operation and Development (OECD) Env-Growth model (Dellink et al., 2015) projections. We allocate national income per capita to impact regions using current nighttime light satellite imagery from the NOAA Defense Meteorological Satellite Program (DSMP). Appendix B.3.2 provides details on this calculation.

Because SSP projections are not available after the year 2100, our calculation of the mortality partial SCC relies on an extrapolation of the relationship between climate change damages and global temperature change to later years; see Section 7 for details.

4 Empirical estimates of the mortality-temperature relationship, accounting for income and climate heterogeneity

Here we describe an empirical approach to quantify the heterogeneous impact of temperature on mortality across the globe using historical data. This method allows us to capture differences in temperature sensitivity across distinct populations in our sample, and thus to quantify the benefits of adaptation as observed historically. The following section details how we combine this empirical information with standard projection data to construct estimates of the mortality risk of climate change, accounting for the benefits of adaptation.

4.1 Empirical model

We estimate the mortality-temperature relationship using a pooled sample of age-specific mortality rates across 40 countries. The effect of temperature on mortality rates is identified using year-to-year variation in the distribution of daily weather following, for example, Deschênes and Greenstone (2011). Additionally,

[8] FAIR is a zero-dimensional structural representation of the global climate designed to capture the temporal dynamics and equilibrium response of global mean surface temperature to greenhouse gas forcing. Appendix G shows that our simulation runs with FAIR create warming distributions that match those from the climate projections in the high-resolution models in the SMME.

we allow the effect of temperature to vary with average temperature (i.e., long-run climate) and average per capita incomes.[9] This approach provides separate estimates for the effect of climate-driven adaptation and income growth on the shape of the mortality-temperature relationship, as they are observed in the historical record.

The two factors defining this interaction model reflect the economics governing adaptation. First, a higher long-run average temperature incentivizes investment in heat-related adaptive behaviors, as the return to any given adaptive mechanism is higher the more frequently the population experiences days with life-threatening temperatures. Second, higher incomes relax agents' budget constraints and hence facilitate adaptive behavior. In other words, people live successfully in both Anchorage, AK and Houston, TX due to compensatory responses to their climate, while the wealthy purchase more safety. To capture these effects, we interact a nonlinear temperature response function with location-specific measures of climate and per capita income.

We fit the following model:

$$M_{ait} = g_a(\boldsymbol{T}_{it}\ ,\ TMEAN_s, \log(GDPpc)_s) + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait}, \tag{5}$$

where a indicates age category with $a \in \{<5, 5\text{-}64, >64\}$, i denotes the second administrative level (ADM2, e.g., county),[10] s refers to the first administrative level (ADM1. e.g., state or province), c denotes country, and t indicates years. Thus, M_{ait} is the age-specific all-cause mortality rate in ADM2 unit i in year t. α_{ai} is a fixed effect for $age \times ADM2$, and δ_{act} a vector of fixed effects that allow for shocks to mortality that vary at the $age \times country \times year$ level.

Our focus in Equation 5 is the effect of temperature on mortality, conditional on average climate and income, which is represented by the age-specific response function $g_a(\cdot)$. Before describing the functional form of this response, we note that our climate data are provided at the grid-cell-by-day level. To align gridded daily temperatures with annual administrative mortality records, we first take nonlinear functions of grid-level daily average temperature and sum these values across the year. We then collapse annual observations across grid cells within each ADM2 using population weights in order to represent temperature exposure for the average person within an administrative unit.[11] This process allows for the recovery of a nonlinear relationship between mortality and temperature at the grid cell level, even though Equation 5 is estimated at a higher level of aggregation (Hsiang, 2016). The nonlinear transformations of daily temperature

[9] These two factors have been the focus of studies modeling heterogeneity across the broader climate-economy literature. For examples, see Mendelsohn, Nordhaus, and Shaw (1994); Kahn (2005); Auffhammer and Aroonruengsawat (2011); Hsiang, Meng, and Cane (2011); Graff Zivin and Neidell (2014); Moore and Lobell (2014); Davis and Gertler (2015); Heutel, Miller, and Molitor (2017); Isen, Rossin-Slater, and Walker (2017).

[10] This is usually the case. However, as shown in Table 1, the EU data is reported at Nomenclature of Territorial Units for Statistics 2^{nd} (NUTS2) level, and Japan reports mortality at the first administrative level.

[11] Specifically, we summarize gridded daily average temperatures T_{zd} across grid cells z and days d to create the annual ADM2-level vector $\boldsymbol{T}_{it}$ as follows:

$$\boldsymbol{T}_{it} = \left[\sum_{z \in i} w_{zi} \sum_{d \in t} T_{zd}, \sum_{z \in i} w_{zi} \sum_{d \in t} T_{zd}^2, \sum_{z \in i} w_{zi} \sum_{d \in t} T_{zd}^3, \sum_{z \in i} w_{zi} \sum_{d \in t} T_{zd}^4 \right]$$

Aggregation across grid cells within an ADM2 is conducted using time-invariant population weights w_{zi}, which represent the share of i's population that falls into grid cell z (see Appendix B.2.4 for details).

are captured by the annual, ADM2-level vector $\boldsymbol{T}_{it}$, and we then choose $g_a(\cdot)$ to be a *linear* function of the *nonlinear* elements of $\boldsymbol{T}_{it}$.

In our main specification, $\boldsymbol{T}_{it}$ contains fourth order polynomials of daily average temperatures, summed across the year. We emphasize results from the polynomial model because it strikes a balance between providing sufficient flexibility to capture important nonlinearities, parsimony, and limiting demands on the data. Analogous to temperature, we summarize daily grid-level precipitation in the annual ADM2-level vector $\boldsymbol{R}_{it}$. We construct $\boldsymbol{R}_{it}$ as a second-order polynomial of daily precipitation, summed across the year, and estimate an age- and country-specific linear function of this vector, represented by $q_{ac}(\cdot)$.

In a set of robustness checks we explore the sensitivity of the results to alternative functional forms for temperature. Specifically, we alternatively define $\boldsymbol{T}_{it}$ as a vector of binned daily average temperatures, as a vector of restricted cubic splines of daily average temperatures, and as a 2-part linear spline of daily average temperatures.[12]

The impact of weather realizations $\boldsymbol{T}_{it}$ on mortality is identified from the plausibly random year-to-year variation in temperature within a geographic unit. Specifically, the $age \times ADM2$ fixed effects α_{ai} ensure that we isolate within-location year-to-year variation in temperature and rainfall exposure, which is as good as randomly assigned. The $age \times country \times year$ fixed effects δ_{act} account for any time-varying trends or shocks to age-specific mortality rates which are unrelated to the climate. We explore robustness to alternative sets of fixed effects in Table D2.

The mortality-temperature response function $g_a(\cdot)$ depends on $TMEAN$, the sample-period average annual temperature, and the logarithm of $GDPpc$, the sample-period average of annual GDP per capita. The model does not include uninteracted terms for $TMEAN$ and $GDPpc$ because they are collinear with α_{ai}, which effectively shuts down the possibility of the climate influencing the mortality rate equally on all days, regardless of daily temperature. This is because we define climate adaptation to be actions or investments that reduce the risk of temperatures that threaten human well-being, as is common in the literature (e.g., Hsiang (2016)). The paper's analysis therefore allows the benefits (and, as discussed later, the costs) of adaptation to influence the shape of the mortality-temperature relationship, but not its level.

We implement a form of $g_a(\cdot)$ that exploits linear interactions between the ADM1-level covariates and all nonlinear elements of the temperature vector $\boldsymbol{T}_{it}$. While long-run climate and GDP per capita enter linearly, they are interacted with all the terms of the fourth order polynomial $\boldsymbol{T}_{it}$. More details on implementation of this regression are given in Appendix D.1.[13] We estimate Equation 5 without any regression weights since

[12] In the binned specification, annual values are calculated as the number of days in region i in year t that have an average temperature that falls within a fixed set of 5°C bins. The bin edges are positioned at the locations $\{-\infty, -15, -10, -5, 0, 5, 10, 15, 20, 25, 30, 35, +\infty\}$ in °C. In the restricted cubic spline specification, daily spline terms are summed across the year and knots are positioned at the locations $\{-12, -7, 0, 10, 18, 23, 28, 33\}$ in °C. In the linear spline specification, heating degree days below 0°C and cooling degree days above 25°C are summed across the year.

[13] To see how we implement Equation 5 in practice, let $\boldsymbol{\beta}_a$ indicate the vector of four coefficients that describes the age-specific fourth-order polynomial mortality-temperature response function. In estimating Equation 5, we allow $\boldsymbol{\beta}_a$ to change with climate and income by modeling each element of $\boldsymbol{\beta}_a$ as a linear function of these two variables. Using this notation, our estimating equation is:

$$M_{ait} = \underbrace{(\boldsymbol{\gamma}_{0,a} + \boldsymbol{\gamma}_{1,a} TMEAN_s + \boldsymbol{\gamma}_{2,a} \log(GDPpc)_s)}_{\boldsymbol{\beta}_a} \boldsymbol{T}_{it} + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait},$$

where $\boldsymbol{\gamma}_{0,a}, \boldsymbol{\gamma}_{1,a}$, and $\boldsymbol{\gamma}_{2,a}$ are each vectors of length four, the latter two describing the effects of $TMEAN$ and $\log(GDPpc)$ on the sensitivity of mortality M_{ait} to temperature $\boldsymbol{T}_{it}$.

we are explicitly modeling heterogeneity in treatment effects rather than integrating over it (Solon, Haider, and Wooldridge, 2015).

A central challenge in understanding the extent of adaptation is that there exists no experimental or quasi-experimental variation in *climate* as opposed to *weather*. Put simply, meaningful variation in climate within a location is not available in recorded history. So, while plausibly random year-to-year fluctuations in temperature within locations are used to identify the effect of weather events in Equation 5, we must use *cross-sectional variation* in climate and income between locations to estimate heterogeneity in the mortality-temperature relationship. We therefore interpret our heterogeneity results as associational.

Nevertheless, we believe this model generates informative estimates of the impact of climate change on mortality for several reasons, including: alternative sources of heterogeneity in mortality sensitivity to temperature have little effect on the estimated response functions; the model performs well out-of-sample on a variety of cross-validation tests; and estimated response functions are robust to a host of alternative specifications. These tests are discussed in detail in Section 4.2.

4.2 Empirical results

Tabular results for the estimation of Equation 5 are reported in Table D1 for each of the three age groups. As these terms are difficult to interpret, we visualize this heterogeneity by dividing the sample into terciles of income and terciles of climate (i.e., the two interaction terms), and then further dividing the sample into the intersection of these two groups of three. This partitions the $\log(GDPpc) \times TMEAN$ space into nine subsamples. We plot predicted response functions at the mean value of climate and income within each of these nine subsamples, using the coefficients in Table D1. The result is a set of predicted response functions that vary across the joint distribution of income and average temperature within the sample data. The resulting response functions are shown in Figure 1 for the >64 age category (other age groups are shown in Appendix D.1), where average incomes are increasing across subsamples vertically and average temperatures are increasing across subsamples horizontally.

The Figure 1 results are broadly consistent with the economic prediction that people adapt to their climate and that income is protective. For example, within each income tercile in Figure 1, the effect of hot days (e.g., days >35°C) declines as one moves from left (cold climates) to right (hot climates). This finding reflects that individuals and societies make compensatory adaptations in response to their climate (e.g., people install air conditioning in hot climates more frequently than in cold ones). With respect to income, Figure 1 reveals that moving from the bottom (low income) to top (high income) within a climate tercile causes a substantial flattening of the response function, especially at high temperatures. Thus, protection from extreme temperatures appears to be a normal good.

Two statistics help to summarize the findings from Figure 1. First, in the >64 age category across all income values, moving from the coldest to the hottest tercile saves on average 7.9 (p-value=0.06) deaths per 100,000 at 35°C. Second, moving from the poorest to the richest tercile across all climate values in the sample saves approximately 5.0 (p-value=0.1) deaths per 100,000 at 35°C for the > 64 age category.

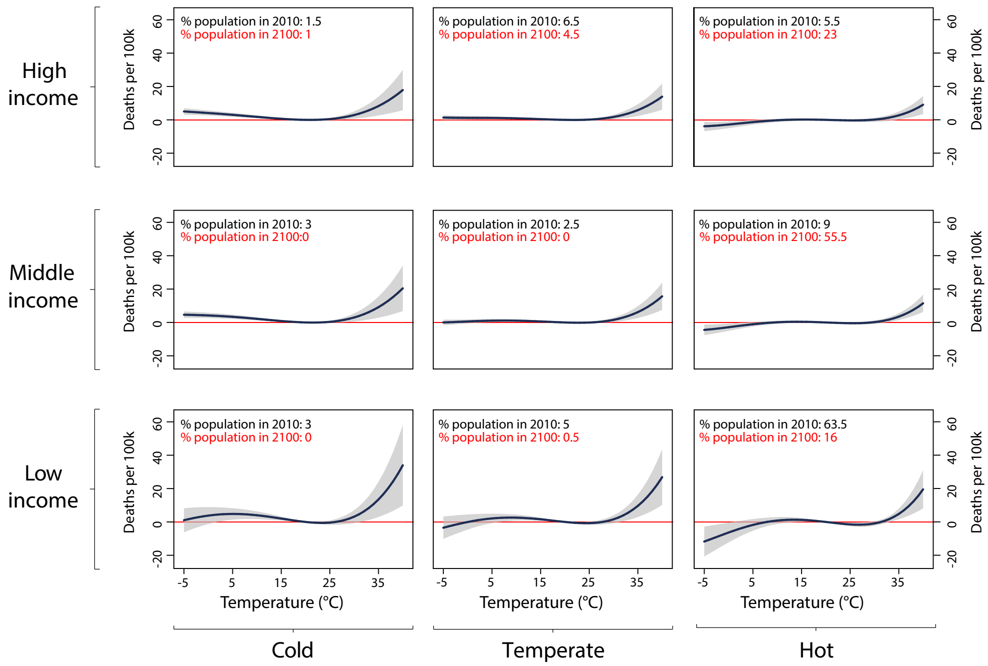

Figure 1: Heterogeneity in the mortality-temperature relationship (age >64 mortality rate). Each panel represents a predicted mortality-temperature response function for the >64 age group for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left apply to the low-income, cold regions of our sample, while those in the upper right apply to the high-income, hot regions of our sample. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level on the top end of the distribution only. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects, and where each temperature variable is interacted with each covariate. Values in the top left-hand corner of each panel show the percentage of the global population that reside within each in-sample tercile of average income and average temperature in 2010 (black text) and as projected in 2100 (red text, SSP3). Other age groups are shown in Figures D1 and D2.

4.3 Sensitivity analyses

4.3.1 Age group heterogeneity

Consistent with prior literature (e.g., Deschênes and Moretti, 2009; Heutel, Miller, and Molitor, 2017), we uncover substantial heterogeneity across age groups within our multi-country sample. Figure 2 displays the average mortality-temperature response for each of our three age categories (<5, 5-64, >64),[14] while Appendix D.1 shows the influence of income and climate on the mortality-temperature relationships for each age group. On average across the globe, we find that people over the age of 64 experience approximately 4.7 extra deaths per 100,000 for a day at 35°C (95°F) compared to a day at 20°C (68°F), a substantially larger effect than that for younger cohorts, which exhibit little response. This age group is also more severely affected by cold days; estimates suggest that people over the age of 64 experience 3.4 deaths per 100,000 for a day at −5°C (23°F) compared to a day at 20°C, while there is a relatively weak mortality response to these cold days for other age categories. Overall, these results demonstrate that the elderly are disproportionately harmed by additional hot days and disproportionately benefit from reductions in cold days.

[14] Age-specific regression estimates in Figure 2 are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects and has no income or climate interaction terms (Equation D.17). See Appendix D.2.1 for details.

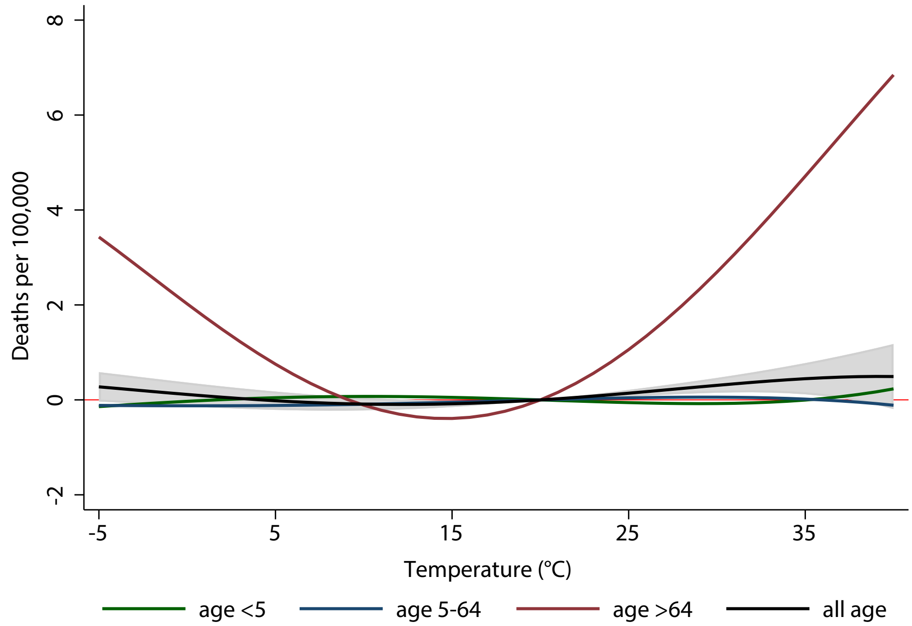

Figure 2: Mortality-temperature response function with demographic heterogeneity. Mortality-temperature response functions are estimated for populations <5 years of age (green), between 5 and 64 years of age (blue), >64 years of age (red), and pooled across all ages (black, with associated 95% confidence intervals shaded in grey). Regression estimates shown are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level. All age-specific response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects (Equation D.17). Confidence intervals are shown only for the all-age response function; statistical significance for age-specific response functions can be seen in Table D2.

4.3.2 Alternative fixed effects

Table D2 reports on the robustness of the estimated mortality-temperature relationship to alternative spatial and temporal controls. Tabular results show the average multi-country marginal effect of temperature evaluated at various temperatures. These estimates can be interpreted as the change in the number of deaths per 100,000 per year resulting from one additional day at each temperature, compared to the reference day of 20°C (68°F). Columns (1)-(3) increase the saturation of temporal controls in the model specification, ranging from country-year fixed effects in column (1) to country-year-age fixed effects in column (2), and adding age-specific state-level linear trends in column (3). Our preferred specification is column (2), as column (1) does not account for differential temporal shocks to mortality rates by age group, while in column (3) we cannot reject the null of equal age-specific, ADM1-level trends. However, estimated age-specific responses are similar across all specifications. This result is robust to alternative functional form assumptions (i.e., different nonlinear functions of $\boldsymbol{T}_{it}$), including a non-parametric binned regression, as well as to the use of alternative, independently-sourced, climate datasets (Figure D3).

4.3.3 Alternative specifications

In Table D2, columns (4) and (5) provide results for the average mortality-temperature relationship under alternative specifications. In column (4), we address the fact that some of the data are drawn from countries which may have less capacity for data collection than others in the sample. Because the mortality data are collected by institutions in different countries, it is possible that some sources are systematically less precise. To account for this, we re-estimate the model using Feasible Generalized Least Squares (FGLS) under the assumption of constant variance within each ADM1 unit.[15] In column (5), we allow for the possibility that temperatures can exhibit lagged effects on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). Lagged effects within and across months in the same calendar year are accounted for in the net annual mortality totals used in all specifications. However, it is possible that temperature exposure in December of each year affects mortality in January of the following year. To account for this, in column (5) we define a 13-month exposure window to additionally account for temperatures previous December.[16] Table D2 shows that the results for both of these alternative specifications are similar in sign and magnitude to those from column (2).

Figure D3 displays the results of estimating the mortality-temperature relationship using a set of alternative functional forms of temperature (i.e., different formulations of the temperature vector $\boldsymbol{T}_{it}$) and using two different climate datasets to obtain those temperatures (see Appendix B.2 for details on these climate datasets). We explore three functional forms in addition to the main fourth-order polynomial specification: bins of daily average temperature, restricted cubic splines, and piecewise linear splines. The first two are especially demanding of the data, particularly in the context of Equation 5, which allows for heterogeneity in temperature sensitivity. Overall, the results for these alternative functional form specifications are similar to the fourth-order polynomial when using both climate datasets (see Appendix D.2 for details).

Finally, we find that the coefficients in Equation 5 are qualitatively unchanged when we use alternative characterizations of the climate (see Appendix D.4) or if we omit precipitation controls (see Appendix D.5).

4.3.4 Additional sources of heterogeneity

In order to predict responses around the world and inform projections of damages in the future, it is necessary for all covariates in Equation 5 to be available globally today, at high spatial resolution, *and* that credible projections of their future evolution are available. One reason we use average incomes and climate in Equation 5 is that both variables meet these criteria.

However, a valid critique of this model is that other factors likely explain heterogeneity in the mortality-temperature relationship, yet are omitted from Equation 5. To address this possibility, we collect data on five other candidate variables that could explain heterogeneity in mortality sensitivity to temperature, such

[15]To do this, we estimate the model in Equation D.17 using population weights and our preferred specification (column (2)). Using the residuals from this regression, we calculate an ADM1-level weight that is equal to the average value of the squared residuals, where averages are taken across all ADM2-age-year level observations that fall within a given ADM1. We then inverse-weight the regression in a second stage, using this weight. All ADM2-age-year observations within a given ADM1 are assigned the same weight in the second stage, where ADM1 locations with lower residual variance are given higher weight. For some ADM2s, there are insufficient observations to identify age-specific variances; to ensure stability, we dropped the ADM2s with less than 5 observations per age group. This leads us to drop 246 (of >800,000) observations in this specification.

[16]The specification in column (5) defines the 13-month exposure window such that for a given year t, exposure is calculated as January to December temperatures in year t *and* December temperature in year $t-1$.

as institutional quality, doctors per capita, and educational attainment.[17] Appendix D.6 shows that adding these variables as additional interaction terms when estimating Equation 5 generates very similar predicted response functions in historical data. This suggests that a model which employs only income and climate explains a large amount of the heterogeneity across space.

Further, we find that including only climate and income as interaction effects out-performs a model that includes additional interaction terms when those variables are not available in future projections. Appendix D.6 shows that including these potential determinants of heterogeneity when estimating Equation 5, but omitting them when generating predictions (as would be necessary when making climate change impact projections), substantially increases prediction error.

4.3.5 Out-of-sample performance

In the next section, we use coefficients estimated from Equation 5, in combination with local-level observations and projections of $TMEAN$ and $\log(GDPpc)$, to generate predicted response functions in all regions of the world, including where mortality data are unavailable, both in the present and into future. To assess the performance of our model in predicting mortality-temperature relationships out-of-sample, we implement multiple custom cross-validation exercises designed to mimic the spatial and temporal extrapolation that is required when using available historical data to generate global climate change projections decades into the future. These tests are described in detail in Appendix D.7, but we summarize their results here.

We perform three cross-validation exercises. In each case, we compare the performance of Equation 5 to the performance of a benchmark model without $TMEAN$ and $\log(GDPpc)$ interactions; that is, a model that ignores adaptation and benefits of income. We do so because most prior literature has estimated impacts of climate change using spatially and/or temporally homogeneous response functions (e.g., Hsiang et al., 2017; Deschênes and Greenstone, 2011). The first exercise uses standard k-fold cross-validation (Friedman et al., 2001), but constrains all observations within an ADM1 (e.g., state) to remain in either the "testing" or the "training" sample within each fold, in order to account for spatial and temporal correlation within these regions. The second exercise subsamples the data based on the in-sample distributions of $TMEAN$ and $\log(GDPpc)$ and tests the model's ability to predict mortality rates in populations with different incomes and climates than the estimation sample. The final exercise subsamples data based on time, testing the model's ability to predict future mortality-temperature relationships.

In all three cases, we find that the model in Equation 5 performs well, both when compared to measures of in-sample model fit and when compared to the out-of-sample performance of a model that omits interaction effects. In particular, Equation 5 performs well in predicting mortality rates in the lowest income and hottest locations, even when those locations are omitted from the estimating sample (see Panel B of Table D5). This is an important result, given the under-representation of low income and hot climates in our mortality records, relative to the global population (see Figure 3). We investigate this finding further in Appendix D.8, where we show strong predictive performance in India, a hot and relatively poor country that is not used in estimation due to its lack of age-specific mortality rates. We do find that Equation 5

[17] In collecting these data, we note that obtaining any of them at subnational scales is a substantial challenge and in most cases not possible. See Appendix D.6 for details.

occasionally over-estimates or under-estimates future mortality sensitivity to hot days in some age groups and for some income levels (see Figure D9). To address this concern, we explore in Appendix F.4 the sensitivity of our main climate change projections to alternative assumptions about the rates of adaptation.

5 Projections of climate change impacts on future mortality rates

This section begins by using the empirical results from Section 4 to extrapolate mortality-temperature relationships to the parts of the world where historical mortality data are unavailable. We then combine these estimates with projected changes in climate exposure and income growth to quantify expected climate change induced mortality risk, accounting for climate model and econometric uncertainty. The paper's ultimate aim is to develop an estimate of the full mortality-related costs of climate change (i.e., the sum of the increase in deaths and adaptation costs shown in Equation 4), but adaptation costs are not observed directly (see Section 2). Therefore, here we display empirically-derived estimates of changes in mortality rates due to climate change, highlighting the difference between projections that do and do not account for the benefits of adaptation. In the following section, we use a stylized revealed preference approach to infer adaptation costs, which allows for a complete measure.

5.1 Defining three measures of climate change impacts

Here we define three measures of climate change impacts that elucidate the roles of adaptation and income growth in determining the full mortality-related costs of climate change. The empirical estimation of each of these measures is first reported in units of deaths per 100,000 using the estimates of $\hat{g}_a(\cdot)$ reported in Section 4, although it is straightforward to monetize these measures using estimates of the value of a statistical life (VSL), and we do so in the next section.

The first measure is the *mortality effects of climate change with neither adaptation nor income growth*, which provides an estimate of the increases in mortality rates when each impact region's response function in each year t is a function of their initial period (indicated as t_0) level of income and average climate (recall Equation 2). In other words, mortality sensitivity to temperature is assumed not to change with future income or temperature. This is a benchmark model often employed in previous work. Specifically, the expected climate induced mortality risk that we estimate for an impact region and age group in a future year t under this measure are (omitting subscripts for impact regions and age groups for clarity):[18]

(i) *Mortality effects of climate change with neither adaptation nor income growth*:

$$\underbrace{\hat{g}(\boldsymbol{T}_t, TMEAN_{t_0}, \log(GDPpc)_{t_0})}_{\substack{\text{mortality risk with climate change} \\ \text{and without adaptation}}} - \underbrace{\hat{g}(\boldsymbol{T}_{t_0}, TMEAN_{t_0}, \log(GDPpc)_{t_0})}_{\text{current mortality risk}}$$

The second measure is the *mortality effects of climate change with benefits of income growth*, which allows response functions to change with future incomes. This measure captures the change in mortality rates that

[18]Note that in all estimates of climate change impacts, population growth is accounted for as an exogenous projection that does not depend on the climate.

would be expected from climate change if populations became richer, but they did not respond optimally to warming by adapting above and beyond how they would otherwise cope with their historical climate. This measure is defined as:

(ii) *Mortality effects of climate change with benefits of income growth*:

$$\underbrace{\hat{g}(\boldsymbol{T}_t, TMEAN_{t_0}, \log(GDPpc)_t)}_{\substack{\text{mortality risk with climate change} \\ \text{and benefits of income growth}}} - \underbrace{\hat{g}(\boldsymbol{T}_{t_0}, TMEAN_{t_0}, \log(GDPpc)_t)}_{\substack{\text{mortality risk without climate change} \\ \text{and with benefits of income growth}}}$$

Note that in expression (ii), the second term represents a counterfactual predicted mortality rate that would be realized under current temperatures, but in a population that benefits from rising incomes over the coming century. This counterfactual includes the prediction, for example, that air conditioning will become much more prevalent in a country like India as the economy grows, regardless of whether climate change unfolds or not.

The third measure is the *mortality effects of climate change with benefits of income growth and adaptation*, and in this case populations adjust to experienced temperatures in the warming scenario (recall Equation 3). This metric is an estimate of the observable deaths that would be expected under a warming climate, accounting for the benefits of optimal adaptation and income growth:

(iii) *Mortality effects of climate change with benefits of income growth and adaptation*:

$$\underbrace{\hat{g}(\boldsymbol{T}_t, TMEAN_t, \log(GDPpc)_t)}_{\substack{\text{mortality risk with climate change, benefits} \\ \text{of income growth, and adaptation}}} - \underbrace{\hat{g}(\boldsymbol{T}_{t_0}, TMEAN_{t_0}, \log(GDPpc)_t)}_{\substack{\text{mortality risk without climate change} \\ \text{and with benefits of income growth}}}$$

As above, expression (iii) includes the subtraction of a counterfactual in which incomes rise but climate is held fixed.

Year t_0 is treated as the baseline period, which we define to be the years 2001-2010, so we are measuring the impact of climate change since this period.[19] These three measures are all reported below in units of deaths per 100,000, using the estimates of $\hat{g}(\cdot)$ shown in Section 4.

5.2 Methods for climate change projection: spatial extrapolation

The fact that carbon emissions are a global pollutant requires that estimates of climate damages used to inform an SCC must be global in scope. A key challenge for generating such globally-comprehensive estimates in the case of mortality is the absence of data throughout many parts of the world. Often, registration of births and deaths does not occur systematically. Although we have, to the best of our knowledge, compiled the most comprehensive mortality data file ever collected, our 40 countries only account for 38% of the global population (55% if India is included, although it only contains all-age mortality rates). This leaves more than 4.2 billion people unrepresented in the sample of available data, which is especially troubling because

[19] While anthropogenic warming has been detected in the climate record far earlier than 2001-2010, we estimate impacts of climate change only since this period.

these populations have incomes and live in climates that may differ from the parts of the world where data are available.

To achieve the global coverage essential to understanding the costs of climate change, we use the results from the estimation of Equation 5 on the observed 38% global sample to estimate the sensitivity of mortality to temperature everywhere, including the unobserved 62% of the world's population. Specifically, the results from this model enable us to use two observable characteristics – average temperature and income – to predict the mortality-temperature response function for each of our 24,378 impact regions. Importantly, it is not necessary to recover the overall mortality rate for these purposes.

To see how this is done, we note that the projected response function for any impact region r requires three ingredients. The first are the estimated coefficients $\hat{g}_a(\cdot)$ from Equation 5. The second are estimates of GDP per capita at the impact region level.[20] And third is the average annual temperature (i.e., a measure of the long-run climate) for each impact region, where we use the same temperature data that were assembled for the regression in Equation 5.

We then predict the shape of the response function for each age group a, impact region r, and year t, up to a constant: $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt}, TMEAN_{rt}, \log(GDPpc)_{rt})$. The various fixed effects in Equation 5 are unknown and omitted, since they were nuisance parameters in the original regression. This results in a unique, spatially heterogeneous, and globally comprehensive set of predicted response functions for each location on Earth.

The accuracy of the predicted response functions will depend, in part, on its ability to capture responses in regions where mortality data are unavailable. An imperfect but helpful exercise when considering whether our model is representative is to evaluate the extent of common overlap between the two samples. Figure 3A shows this overlap in 2015, where the grey squares reflect the joint distribution of GDP and climate in the full global partition of 24,378 impact regions and orange squares represent the analogous distribution only for the impact regions in the sample used to estimate Equation 5. It is evident that temperatures in the global sample are generally well-covered by our data, although we lack coverage for the poorer end of the global income distribution due to the absence of mortality data in poorer countries. As discussed in Section 4, we explore this extrapolation to lower incomes with a set of robustness checks in Appendix D.

5.3 Methods for climate change projection: temporal extrapolation

As discussed in Section 2, a measure of the full mortality risk of climate change must account for the benefits that populations realize from optimally adapting to a gradually warming climate, as well as from income growth relaxing the budget constraint and enabling compensatory investments. Thus, we allow each impact region's mortality-temperature response function to evolve over time, reflecting projected changes in climate and incomes that come from a set of internationally standardized and widely used scenarios. Specifically, we model the evolution of response functions in region r and year t based on these projections and the estimation results from fitting Equation 5.

Some details about these projections are worth noting. First, a 13-year moving average of income per

[20] The procedure is described in Section 3.2 and Appendix B.3.2

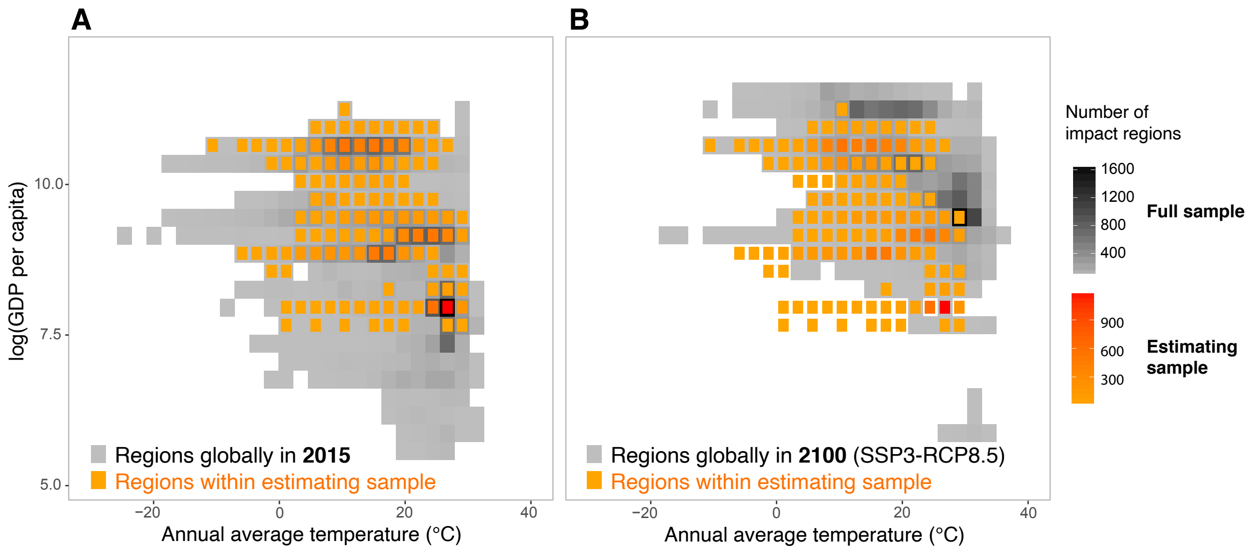

Figure 3: Joint coverage of income and long-run average temperature for estimating and full samples. Joint distribution of income and long-run average annual temperature in the estimating sample (red-orange), as compared to the global sample of impact regions (grey-black). Panel A shows in grey-black the global sample for regions in 2015. Panel B shows in grey-black the global sample for regions in 2100 under a high-emissions scenario (RCP8.5) using climate model CCSM4 and a median growth scenario (SSP3). In both panels, the in-sample frequency in red-orange indicates coverage for impact regions within our data sample in 2015.

capita in region r is calculated using national forecasts from the Shared Socioeconomics Pathways (SSP), combined with a within-country allocation of income based on present-day nighttime lights (see Appendix B.3.2), to generate a new value of $\log(GDPpc)_{rt}$. The length of this time window is chosen based on a goodness-of-fit test across alternative window lengths (see Appendix E.1). Second, a 30-year moving average of temperatures for region r is updated in each year t to generate a new level of $TMEAN_{rt}$. Finally, the response curves $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt}, TMEAN_{rt}, \log(GDPpc)_{rt})$ are calculated for each region for each age group in each year with these updated values of $TMEAN_{rt}$ and $\log(GDPpc)_{rt}$.

Figure 3B shows that over the coming decades, temperatures and incomes are predicted to rise beyond the support of the global cross-section in our data. Thus, we must impose two constraints, guided by economic theory and by the physiological literature, to ensure that future response functions are consistent with the fundamental characteristics of mortality-temperature responses in the historical record and demonstrate plausible out-of-sample projections.[21] First, we impose the constraint that the response function must be weakly monotonic around an empirically estimated, location-specific, optimal mortality temperature, called the *minimum mortality temperature* (MMT). That is, we assume that temperatures farther from the MMT (either colder or hotter) must be at least as harmful as temperatures closer to the MMT. This assumption is important because Equation 5 uses within-sample variation to parameterize how the U-shaped response function flattens; with extrapolation beyond the support of historically observed income and climate, this behavior could go "beyond flat", such that extremely hot and cold temperature days reduce mortality relative to the MMT (Figure E1). In fact, this is guaranteed to occur mechanically if enough time elapses, because income and climate interact with the response function linearly in Equation 5. However, such behavior is

[21] See Appendix E.2 for details on these assumptions and their implementation.

inconsistent with a large body of epidemiological and econometric literature recovering U-shaped mortality-temperature relationships under many functional form assumptions and in diverse locations (Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011), as well as what we observe in our data. As a measure of its role in our results, the weak monotonicity assumption binds for the >64 age category at 35°C in 9% and 18% of impact regions in 2050 and 2100, respectively.[22,23]

Second, we assume that rising income cannot make individuals worse off, in the sense of increasing the temperature sensitivity of mortality. Because increased income per capita strictly expands the choice set of individuals considering whether to make adaptive investments, it should not increase the effect of temperature on mortality rates. Consistent with this intuition, we find that income is protective against extreme heat for all age groups. However, for some age groups, the estimation of Equation 5 recovers statistically insignificant but positive effects of income on mortality sensitivity to extreme cold (Table D1). Therefore, we constrain the marginal effect of income on temperature sensitivity to be weakly negative in future projections, although we place no restrictions on the cross-sectional effect of income when estimating Equation 5.[24]

With these two constraints, we project annual impacts of climate change separately for each impact region and age group from 2001 to 2100. Specifically, we apply projected changes in the climate to each region's response function, which is evolving as climate and income evolve. The nonlinear transformations of daily average temperature that are used in the function $g_a(\boldsymbol{T}_{rt})$ are computed under both the RCP4.5 and RCP8.5 emissions scenarios for all 33 climate projections in the SMME (as described in Section 3.2). This distribution of climate models captures uncertainties in the climate system through 2100.

5.4 Methods for accounting for uncertainty in projected mortality effects of climate change

An important feature of the analysis is to characterize the uncertainty inherent in these projections of the mortality impacts of climate change.[25] As discussed in Section 5.3, we construct estimates of the mortality risk of climate change for each of 33 distinct climate projections in the SMME that together capture the uncertainty in the climate system.[26] Additionally, uncertainty in the estimates of $\hat{g}_a(\cdot)$ is an important second source of uncertainty in our projected impacts that is independent of physical uncertainty.

In order to account for both of these sources of uncertainty, we execute a Monte Carlo simulation following the procedure in Hsiang et al. (2017). First, for each age category, we randomly draw a set of parameters,

[22] The frequency with which the weak monotonicity assumption binds will depend on the climate model and the emissions and socioeconomic trajectories used; reported statistics refer to the CCSM4 model under RCP8.5 with SSP3.

[23] In imposing this constraint, we hold the MMT fixed over time at its baseline level in 2015 (Figure E1D). We do so because the use of spatial and temporal fixed effects in Equation 5 implies that response function levels are not identified; thus, while we allow the *shape* of response functions to evolve over time as incomes and climate change, we must hold fixed their *level* by centering each response function at its time-invariant MMT. Note that these fixed effects are by definition not affected by a changing weather distribution. Thus, their omission does not influence estimates of climate change impacts.

[24] The assumption that rising income cannot increase the temperature sensitivity of mortality binds for the >64 age category under realized temperatures in 30% and 24% of impact region days in 2050 and 2100, respectively.

[25] See Burke et al. (2015) for a discussion of combining physical uncertainty from multiple models in studies of climate change impacts.

[26] Note that while the SMME fully represents the tails of the climate sensitivity distribution as defined by a probabilistic simple climate model (see Appendix B.2.3), there remain important sources of climate uncertainty that are not captured in our projections, due to the limitations of both the simple climate model and the GCMs. These include some climate feedbacks that may amplify the increase of global mean surface temperature, as well as some factors affecting local climate that are poorly simulated by GCMs.

corresponding to the terms composing $\hat{g}_a(\cdot)$, from an empirical multivariate normal distribution characterized by the covariance between all of the parameters from the estimation of Equation 5.[27] Second, using these parameters in combination with location- and time-specific values of income and average climate provided by a given SSP scenario and RCP-specific climate projection from each of the 33 climate projections in the SMME, we construct a predicted response function for each of our 24,378 impact regions. Third, with these response functions in hand, we use daily weather realizations for each impact region from the corresponding simulation to predict an annual mortality impact. Finally, this process is repeated until approximately 1,000 projection estimates are complete for each impact region, age group, and RCP-SSP combination.

With these ∼1,000 response functions, we calculate the mortality effects of climate change (i.e., expressions (i)-(iii) above) for each impact region for each year between 2001 and 2100. The resulting calculation is computationally intensive, requiring ∼94,000 hours of CPU time across all scenarios reported in the main text and Appendix. When reporting projected impacts in any given year, the reports summary statistics (e.g., mean, median) of this entire distribution.

5.5 Results: spatial extrapolation of temperature sensitivity

Figure 4 reports on our extrapolation of mortality-temperature response functions to the entire globe for the >64 age category (see Figure D4 for other age groups). In panel A, these predicted mortality-temperature responses are plotted for each impact region for 2001-2010 average values of income and climate and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Despite a shared overall shape, panel A reveals substantial heterogeneity across regions in this temperature response. Geographic heterogeneity within our sample is shown for hot days in the map in panel C, where colors indicate the marginal effect of a day at 35°C, relative to a day at a location-specific minimum mortality temperature. Grey areas are locations where mortality data are unavailable.

Panels C and D of Figure 4 show analogous plots, but now extrapolated to the entire globe. We can fill in the estimated mortality effect of a 35°C day for regions without mortality data by using 2001-2010 location-specific information on average income and climate. The predicted responses at the global scale imply that a 35°C day increases the average mortality rate across the globe for the oldest age category by 10.1 deaths per 100,000 relative to a location-specific minimum mortality temperature.[28] It is important to note that the effect in locations without mortality data is 11.7 deaths per 100,000, versus 7.8 within the sample of countries for which mortality data are available, largely driven by the fact that our sample represents wealthier locations where temperature responses are more muted.

Overall, there is substantial heterogeneity across the planet. Additionally, it is evident that the effects of temperature on human well-being are quite different in places where we are and are not able to obtain subnational mortality data.

[27] Note that coefficients for all age groups are estimated jointly in Equation 5, such that across-age-group covariances are accounted for in this multivariate distribution.

[28] This average impact of a 35°C is derived by taking the unweighted average level of the mortality-temperature response function evaluated at 35°C across each of 24,378 impact regions globally.

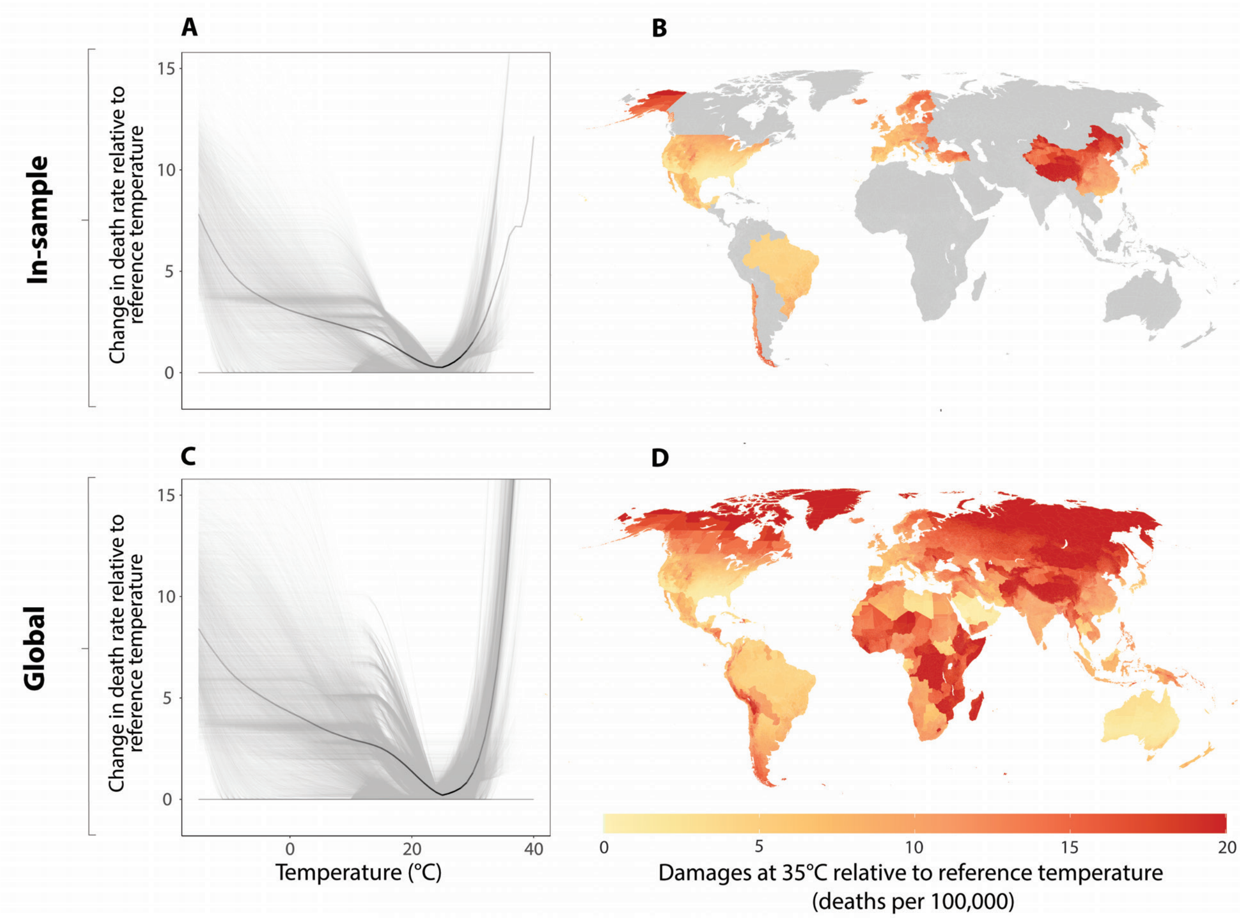

Figure 4: Using income and climate to predict current response functions globally (age >64 mortality rate). In panels A and C, grey lines are predicted response functions for impact regions, each representing a population of 276,000 on average. Solid black lines are the unweighted average of the grey lines, where the opacity indicates the density of realized temperatures (Hsiang, 2013). Panels B and D show each impact region's mortality sensitivity to a day at 35°C, relative to a location-specific minimum mortality temperature. The top row shows all impact regions in the sample of locations with historical mortality data (included in main regression tables), and the bottom row shows extrapolation to all impact regions globally. Predictions shown are averages over the period 2001-2010 using the SSP3 socioeconomic scenario and climate model CCSM4 under the RCP8.5 emissions scenario. Figure D4 shows analogous results for other age groups.

5.6 Results: projection of future climate change impacts

The previous subsection demonstrated that the model of heterogeneity outlined in Equation 5 allows us to extrapolate mortality-temperature relationships to regions of the world without mortality data today. However, to calculate the full global mortality risks of climate change, it is also necessary to allow these response functions to change through time to capture the benefits of adaptation and the effects of income growth. This subsection reports on using our model of heterogeneity and downscaled projections of income and climate to predict impact region-level response functions for each age group and year, as described in Section 5.3. Uncertainty in these estimated response functions is accounted for through Monte Carlo simulation, as described in Section 5.4. Throughout this subsection, we show results relying on income and population projections from the socioeconomic scenario SSP3 because its historic global growth rates in GDP per capita and population match observed global growth rates over the 2000-2018 period much more closely than other SSPs (see Table B3). Appendix F shows results using SSP2 and SSP4, and the methodology we develop can be applied to any available socioeconomic scenario.

5.6.1 Mortality impacts of climate change for 24,378 global regions

Figure 5 shows the spatial distribution of the mortality effects of climate change with benefits of income growth and adaptation (expression (iii)) in 2100 under the emissions scenario RCP8.5, expressed in death-equivalents per 100,000. Other measures of climate change impacts (expressions (i) and (ii)) are mapped in Appendix Figure F1. To construct these estimates, we generate impact-region specific predictions of mortality damages from climate change for all years between 2001 and 2100, separately for each age group. The map displays the spatial distribution of the mean estimate across our ensemble of Monte Carlo simulations, accounting for both climate and statistical uncertainty and pooling across all age groups.[29] The density plots for select cities show the full distribution of impacts across all Monte Carlo simulations, with the white line equal to the mean estimate displayed on the map.

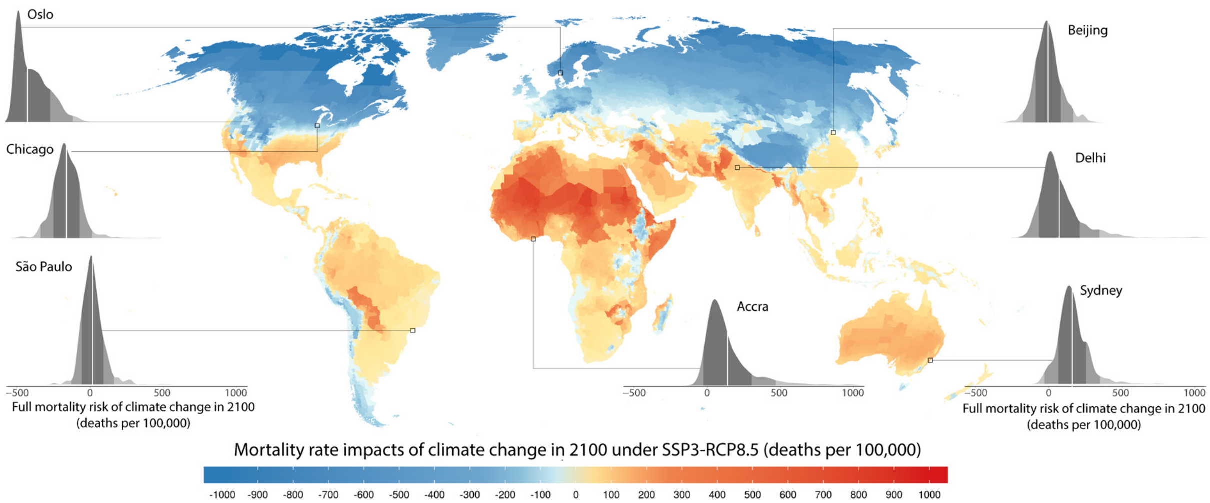

Figure 5: The mortality impacts of future climate change. The map indicates the impact of climate change on mortality rates, measured in units of deaths per 100,000 population, in the year 2100. Estimates come from a model accounting for the benefits of adaptation and income growth, and the map shows the climate model weighted mean estimate across Monte Carlo simulations conducted on 33 climate models; density plots for select regions indicate the full distribution of estimated impacts across all Monte Carlo simulations. In each density plot, solid white lines indicate the mean estimate shown on the map, while shading indicates one, two, and three standard deviations from the mean. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario. See Figure F6 for an analogous map of impacts for RCP4.5 and SSP3.

Figure 5 makes clear that the costs of climate change-induced mortality risks are distributed unevenly around the world, even when accounting for the benefits of income growth and adaptation. Despite the gains from adaptation shown in Figure E2, there are large increases in mortality risk in the global south. For example, in Accra, Ghana, climate change is predicted to lead to approximately 100 more days >32°C (∼90°F) per year and cause 140 additional deaths per 100,000 annually under RCP8.5 in 2100. If adaptation to climate and benefits of income growth were ignored, climate change would be predicted to cause 260

[29]When calculating mean values across estimates generated for each of the 33 climate models that form our ensemble, we use model-specific weights. These weights are constructed as described in Appendix B.2.3 in order to accurately reflect the full probability distribution of temperature responses to changes in greenhouse gas concentrations.

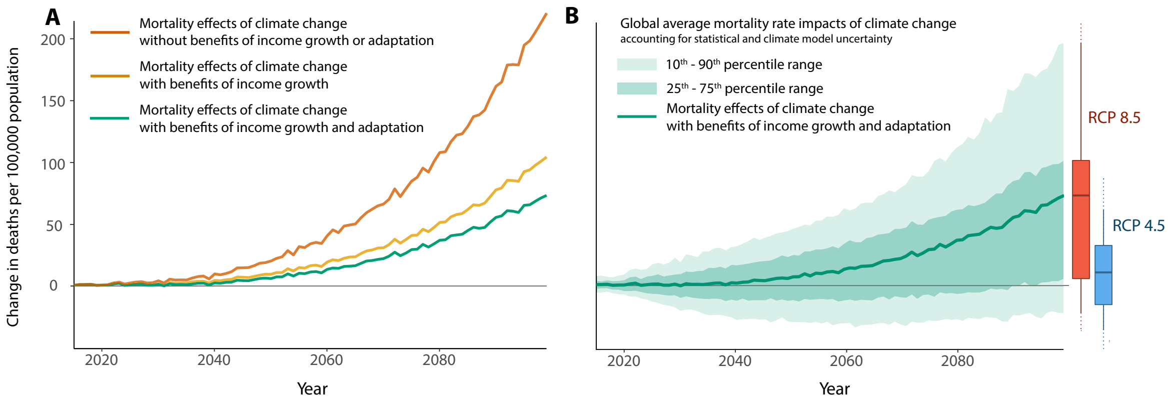

Figure 6: Time series of projected mortality rate impacts of climate change. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents a partial mortality effect. Orange (expression (i)): mortality effects without adaptation. Yellow (expression (ii)): mortality effects with benefits of income growth. Green (expression (iii)): mortality effects with benefits of income growth and adaptation. Panel B shows the 10^{th}-90^{th} percentile range of the Monte Carlo simulations for the mortality effects with benefits of income growth and adaptation (equivalent to the green line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of mortality rate impacts in 2100 under both RCPs. All line estimates shown refer to the RCP8.5 emissions scenario and all line and boxplot estimates refer to the SSP3 socioeconomic scenario. Figure F7 shows the equivalent for SSP3 and RCP4.5.

additional deaths per 100,000 in this scenario. In contrast, there are gains in many impact regions in the global north, including in London, England, where climate change is predicted to save approximately 70 lives per 100,000 annually. When the benefits of adaptation and income growth are included, these changes amount to a 17% increase in Accra's annual mortality rate and an 8% decline in London's.

5.6.2 Aggregate global mortality consequences of climate change

Figure 6 plots predictions of global increases in the mortality rate (deaths per 100,000) for all three measures of climate change impacts, under emissions scenario RCP8.5. The measures are calculated for each of the 24,378 impact regions and then aggregated to the global level. In panel A, each line shows a mean estimate for the corresponding climate change impact measure and year. Averages are taken across the full set of Monte Carlo simulation results from all 33 climate models, and all draws from the empirical distribution of estimated regression parameters, as described in Section 5.4. In panel B, the 25^{th}-75^{th} and 10^{th}-90^{th} percentile ranges of the Monte Carlo simulation distribution are shown for the mortality effects of climate change with benefits of income growth and adaptation (expression (iii)); the black line represents the same average value in both panels. Boxplots to the right summarize the distribution of mortality impacts for both RCP8.5 and the moderate emissions scenario of RCP4.5, and Figure F7 replicates the entire figure for RCP4.5.

Figure 6A illustrates that the mortality cost of climate change would be 221 deaths per 100,000 by 2100, on average across simulation runs (orange line), if the beneficial impacts of adaptation and income are shut down. This is a large estimate; if it were correct, the mortality costs of climate change would be roughly equivalent in magnitude to all global deaths from cardiovascular disease today (WHO, 2018).

However, we estimate that future income growth and adaptation to climate substantially reduce these impacts, a finding that follows directly from the large gains to adaptation and income recovered in the historical record in Section 4. Higher incomes lower the mortality effect of climate change to an average of 104 deaths per 100,000 in 2100 (yellow line), although this estimate exhibits substantial uncertainty (Table D1, Figure F3). Climate adaptation reduces this further to 73 deaths per 100,000 (green line). Although much lower than the *no adaptation* projection, these smaller counts of direct mortality remain economically meaningful—for comparison, the 2019 mortality rate from automobile accidents in the United States was 11 per 100,000.

These large benefits of income growth and climate adaptation are driven by substantial changes in the mortality-temperature relationship over the 21^{st} century. For example, for the >64 age group, the average global increase in the mortality rate on a 35°C day (relative to a day at location-specific minimum mortality temperatures) declines by roughly 75% between 2015 and 2100, going from 10.1 per 100,000 to just 2.4 per 100,000 in 2100 (see Figure E2). Increasing incomes account for 77% of the decline, with adaptation to climate explaining the remainder; income gains account for 89% and 82% of the decline for the <5 and 5-64 categories, respectively.[30]

The values in Figure 6A are mean values aggregated across results from the 33 high-resolution climate models and all Monte Carlo simulation runs, but the full distribution of our estimated damages across climate models (panel B of Figure 6) is right-skewed. Indeed, there is meaningful mass in the "right" tail of potential mortality risk. As evidence of this, the *median* value of the mortality effects of climate change with benefits of income growth and adaptation under RCP8.5 at end of century is 42 deaths per 100,000, as compared to a *mean* value of 73, and the 10^{th} to 90^{th} percentile range is [-22, 197].

Figure 6B and Appendix Figure F5 display the expected implications of emissions mitigation. The average estimate of the mortality effects of climate change with benefits of income growth and adaptation of 73 deaths per 100,000 by the end of the century under RCP8.5 falls to 11 under the emissions stabilization scenario of RCP4.5 (where emissions decline after 2050). For RCP4.5, the median end-of-century estimate is 4, and the 10^{th} to 90^{th} percentile range is [-36, 62].

As a point of comparison to the limited literature estimating the global mortality consequences of climate change, we contrast these results to the FUND model, which is unique among the IAMs for calculating separate mortality impacts as a component of its SCC calculation. Although it is difficult to make a direct comparison due to differences in socioeconomic and emissions scenarios, different treatments of adaptation, and the inclusion of diarrhea and vector-borne diseases in FUND, the closest analog is to compare our estimates of the mortality impacts of climate change including adaptation benefits, a change of 73 deaths per 100,000 by 2100 under RCP8.5, to FUND's reference scenario change of 0.33 deaths per 100,000 in the same year (Anthoff and Tol, 2014).[31] The FUND model was calibrated decades ago based on limited mortality data from just 20 cities largely in wealthy and temperate locations; it is apparent that modern econometric tools and large-scale datasets provide a substantially different picture of the mortality consequences of climate

[30] These values apply to socioeconomic scenario SSP3.

[31] This value was calculated by running the MimiFUND model (v3.12.1) and extracting global additional deaths from all modeled causes. Additional deaths are calculated as the difference between the reference scenario in MimiFUND and a baseline in which both temperature and CO_2 are held constant at their 2005 levels. See Table B4 for details on the differences between our approach, that of FUND, and that of other empirical estimates of the impacts of climate change on mortality.

change.

Before proceeding, we note that a limitation of our empirical approach is that we must sometimes extrapolate response functions to temperatures outside of those historically observed within our data. To address the concern that out-of-sample behavior is disproportionately influencing our results, we repeat the projections of mortality risk changes with two extra sets of restrictions imposed upon our empirically-estimated response functions. These two restrictions, described in detail in Appendix F.3, are: *i)* forcing the response function to be flat for all temperatures outside the observed range, so that, for example, a 42°C day is no more damaging than a 40°C day; or *ii)* setting the marginal effect to be linearly increasing in the out-of-sample regions with a slope equal to the slope at the edge of the observed range. Figure F10 reveals that these two restrictions on out-of-sample behavior have negligible effects on our overall impacts. The value of the mortality impact of climate change including benefits of income growth and adaptation is approximately 1 death per 100,000 smaller by 2100 under RCP 8.5 in the case of the flat out-of-sample restriction (see Appendix F.3 for details).

6 The full value of the mortality risk of climate change

The empirical results above demonstrate that populations with similar incomes but different climates experience strikingly different mortality sensitivity to warming, with warmer populations benefiting from lower sensitivity to increasing heat. These differences reflect a wide variety of compensatory actions and, as highlighted in Equation 4 in Section 2, a full measure of the economic burden of climate change must account for these costs of adaptation. However, it is impossible to enumerate and observe all of the actions individuals take to modify their mortality risk of climate change.

This section develops a revealed preference approach that uses the observed differences in temperature sensitivity to infer measures of location-specific adaptation costs. Specifically, we assume that the differential mortality sensitivities to temperature are due to differential uptake of adaptive technologies, behaviors, or other investments. After all, if these investments were costless, we would expect universal uptake, such that mortality rates would exhibit little to no response to temperature across the globe. The approach therefore assumes that differences in the mortality sensitivity to temperature between locations can be the basis for inferring adaptation costs. This revealed preference approach relies on a strong set of simplifying assumptions, but it can be directly estimated with available data, even when the many dimensions of adaptation remain unobservable.

After outlining our approach for recovering adaptation costs, this section presents projections of the full mortality risk of climate change into the future, accounting for the benefits and costs of adaptation. We additionally demonstrate how the impacts of climate change on mortality and on mortality-related adaptation costs are projected to occur unequally across the globe.

6.1 Revealed preference approach to infer adaptation costs

As in Section 2, we define the the climate as the joint probability distribution over a vector of possible conditions that can be expected to occur over a specific interval of time. $\boldsymbol{C}_t$ describes this probability distribution in time period t and $\boldsymbol{c}(\boldsymbol{C}_t)$ is a random vector of weather realizations drawn from the distribution characterized by $\boldsymbol{C}_t$.

Consider a single representative agent who derives utility in each time period t from consumption of a numeraire good x_t. This agent faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$, which depends both on the weather and on adaptive behaviors and investments captured by the composite good $\boldsymbol{b}_t$. As discussed in Section 2, changes in the climate $\boldsymbol{C}$ influence mortality risk through altering weather realizations $\boldsymbol{c}$ and through changing beliefs about the weather, hence changing adaptive behaviors $\boldsymbol{b}$.

In bringing this framework to our empirical analysis (see Section 6.2 for details), we allow for 24,378 representative agents, one for each of the impact regions that together span the globe. We see this as a substantial improvement upon the existing estimates of global climate change damages that inform the SCC, even though there is heterogeneity in preferences, climate, and income within these regions. For example, the DICE IAM assumes a single homogeneous global region (Nordhaus, 1992), the RICE IAM assumes 10 homogeneous regions (Nordhaus and Yang, 1996), the FUND IAM assumes 16 homogeneous regions (Tol, 1997), and the empirically-derived SCC estimates in Ricke et al. (2018) are country-level.

Each region's representative agent simultaneously chooses consumption of the numeraire x_t and of the composite good $\boldsymbol{b}_t$ in each period to maximize utility given her *expectations* of the weather, subject to an exogenous budget constraint and conditional on the climate. We let $\tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t) = \mathbb{E}_{\boldsymbol{c}_t}[f(\boldsymbol{b}_t, \boldsymbol{c}(\boldsymbol{C}_t)) \mid \boldsymbol{C}_t]$ represent the expected probability of death. This agent therefore solves:

$$\max_{\boldsymbol{b}_t, x_t} \; u(x_t)\left[1 - \tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t)\right] \quad s.t. \;\; Y_t \geq x_t + A(\boldsymbol{b}_t), \tag{6}$$

where $A(\boldsymbol{b}_t)$ represents expenditures for all adaptive investments, and Y is an income we take to be exogenous. Under these assumptions, the first order conditions of Equation 6 define optimal adaptation as a function of income and the climate: $\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)$, which we sometimes denote below as $\boldsymbol{b}_t^*$ for simplicity.[32]

We use this framework to derive an empirically tractable expression for the full value of mortality risk due to climate change, following Equation 4. We begin by rearranging the agent's first order conditions and using the conventional definition of the VSL (i.e., $VSL = \frac{u(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x}$ following, for example, Becker (2007) and Viscusi and Aldy (2003)[33]) to show that in any time period t,

$$\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}} = \frac{-u(x_t^*)}{\partial u/\partial x[1 - \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = -VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \tag{7}$$

That is, marginal adaptation costs (lefthand side) equal the value of marginal adaptation benefits (righthand side), when evaluated at the optimal level of adaptation $\boldsymbol{b}^*$ and consumption x^*. This expression enables

[32] Note that income was omitted in the simplified motivation in Section 2, but enters as an argument of $\boldsymbol{b}_t^*$ here via the budget constraint.

[33] Note that this definition assumes the utility and marginal utility of consumption when dead is zero.

us to use estimates of marginal adaptation benefits, which we obtain from the previous section's estimation results, to infer estimates of marginal adaptation costs.

To make the expression in Equation 7 of greater practical value, we note that the total derivative of expected mortality risk with respect to a change in the climate is the sum of two terms:

$$\frac{d\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} = \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \tag{8}$$

The first term on the righthand side of Equation 8 represents the expected impacts on mortality of all changes in adaptive investments induced by the change in climate; in practice, this term cannot be observed or estimated because of the countless elements of the $\boldsymbol{b}$ vector.[34] The second term is the direct effect that the climate would have if individuals did not adapt (i.e., the partial derivative).[35] If, as is expected, climate change produces an increase in the frequency of heat events that threaten human health, it would be natural to expect the first term to be negative, as people make adjustments that save lives. In this case, we expect the second term to be positive, reflecting the impacts of heat on fatalities absent those adjustment.

Equation 8 makes clear that we can express the unobservable mortality benefits of adaptation (i.e., $\frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}}$) as the difference between the total and partial derivatives of the expected probability of death with respect to climate. This has important practical value because both of these terms can be estimated, as we describe below.

The combination of this algebraic manipulation with Equation 7 allows us to develop an expression for the *total* adaptation costs incurred as the climate changes gradually from t_0 to t that is entirely composed of elements which can be estimated:[36]

$$A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) = \int_{t_0}^{t} \frac{\partial A(\boldsymbol{b}_s^*)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds = -\int_{t_0}^{t} VSL_s \left[\frac{d\tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \tag{9}$$

Equation 9 outlines how we can use estimates of the total and partial derivatives of mortality risk—with respect to the climate—to infer marginal adaptation costs, even though adaptation itself is not directly observable. In the next subsection, we show how we use the empirical model described in Section 4 to separately identify the total derivative $\frac{d\tilde{f}}{d\boldsymbol{C}}$ and the partial derivative $\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$. We empirically quantify these values globally in Section 6.3.

A few details of this approach are worth underscoring. First, the *total* adaptation costs associated with the climate shifting from $\boldsymbol{C}_{t_0}$ to $\boldsymbol{C}_t$ are calculated by integrating marginal benefits of adaptation for a series of infinitesimal changes in climate (Equation 9), where marginal benefits continually evolve with the changing climate $\boldsymbol{C}$. Thus, total adaptation costs in a given period, relative to a base period, are the sum of the adaptation costs induced by a series of small changes in climate in the preceding periods (see Appendix A.1 for a visual description).

[34] This term is often known in the environmental health literature as the effect of "defensive behaviors" (Deschênes, Greenstone, and Shapiro, 2017) and in the climate change literature as "belief effects" (Deryugina and Hsiang, 2017); in our context these effects result from changes in individuals' defensive behaviors undertaken because their beliefs about the climate have changed.

[35] This term is known in the climate change literature as the "direct effect" of the climate (Deryugina and Hsiang, 2017).

[36] Note that x is fully determined by $\boldsymbol{b}$ and income Y through the budget constraint.

Second, the *total* adaptation benefits associated with the climate shifting from $\boldsymbol{C}_{t_0}$ to $\boldsymbol{C}_t$ are defined as the dollar value of the difference between the effects of climate change with optimal adaptation and without any adaptation: $-VSL_t[\tilde{f}(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t), \boldsymbol{C}_t) - \tilde{f}(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0}), \boldsymbol{C}_t)]$. In contrast to total adaptation costs, this expression relies on the relationship between mortality and temperature that holds *only* at the final climate, $\boldsymbol{C}_t$. Therefore, when the marginal benefits of adaptation are greater at the final climate than at previous climates, the total benefits of adaptation will exceed total adaptation costs, generating an adaptation "surplus".[37] For example, at a climate between $\boldsymbol{C}_{t_0}$ and $\boldsymbol{C}_t$, the marginal unit of air conditioning (a key form of adaptation) purchased will have benefits that are exactly equal to its costs. However, at the warmer climate $\boldsymbol{C}_t$, this same unit of air conditioning becomes inframarginal, and is likely to have benefits that exceed its costs. Appendix A.2 derives a formal expression for this adaptation surplus.

Third, while we integrate over changes in climate in Equation 9, we hold income fixed at its endpoint value. This is because the goal is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate only. In contrast, changes in expenditures due to rising income will alter mortality risk under climate change, but are not a consequence of the changing climate; therefore they are not included in our calculation of the total mortality-related costs of climate change.

Finally, this revealed preference approach is purposefully parsimonious so that it can be tightly linked to available data, but such simplification necessarily involves several strong assumptions. We assume that adaptation costs are a function of technology and do not depend on the climate, so that, for example, individuals in Seattle can purchase the same air conditioners as individuals in Houston can. We assume that $\tilde{f}(\cdot)$ is continuous and differentiable, that markets clear for all technologies and investments represented by the composite good $\boldsymbol{b}$, as well as for the numeraire good x, and that all choices $\boldsymbol{b}$ and x can be treated as continuous. We assume that neither adaptation investments nor the climate directly enter the utility function, because the paper's focus is limited to the mortality risks of climate change.[38] Perhaps most importantly, the problem in Equation 6 is static. That is, we assume that there is a competitive and frictionless rental market for all capital goods (e.g., air conditioners), so that fixed costs of capital can be ignored, and that all rental decisions are contained in $\boldsymbol{b}$. While this rules out complementarities between adaptation decisions made by the representative agent in different time periods by assuming that such complementarities can be accommodated by sellers of adaptation services, it has to date been standard in the literature (e.g., Deryugina and Hsiang, 2017; Deschênes and Greenstone, 2011) and accounting for dynamic decision-making would necessitate an ambitious extension of the current paper.[39]

[37] Note that we derive an adaptation surplus assuming continuous adaptation investments $\boldsymbol{b}$; Guo and Costello (2013) find that adaptation surplus is higher when forward-looking agents invest in discrete adaptation behaviors or technologies.

[38] In an alternative specification detailed in Appendix A.4, we allow agents to derive utility both from x and from the choice variables in $\boldsymbol{b}$; for example, air conditioning may increase utility directly, in addition to lowering mortality risk. Under this alternative framework, the costs of adapting to climate change that we can empirically recover, $A(\boldsymbol{b})$, are *net* of any changes in direct utility benefits or costs. Similarly, a model that assumes that climate enters utility directly would also lead to any adaptation costs associated with the direct effects of climate change being "netted out" in our approach to recovering adaptation costs.

[39] For example, the central contribution of Lemoine (2018) is to incorporate complementarity in adaptation actions across periods in a standard model of climate change impact estimation. This paper analyzes only a two-period complementarity, yet estimation in our context would require accurate weather forecast data for all locations and years in our estimating sample, a binding constraint for early years and in developing countries. It is also worth noting that the quantitative impacts of adding dynamic decision-making in Lemoine (2018) were minor, changing the end-of-century estimated losses to U.S. agriculture due to climate change from 47% under a static model to 50% under a dynamic model (see Table 2).

6.2 Computing adaptation costs using empirical estimates

To empirically estimate the adaptation costs incurred as the climate changes gradually from t_0 to t, following Equation 9, we calculate the following approximation (see Section A.3 for details):

$$\widehat{A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A}(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) \approx -\int_{t_0}^{t} VSL_s \left[\frac{d\hat{\tilde{f}}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial\hat{\tilde{f}}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{\partial\boldsymbol{C}}\right] \frac{d\boldsymbol{C}_s}{ds} ds$$
$$\approx -\sum_{\tau=t_0+1}^{t} VSL_\tau \underbrace{\left(\left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_\tau, Y_t}\right)}_{\hat{\boldsymbol{\gamma}}_{\mathbf{1}}\mathrm{E}[\boldsymbol{T}]_\tau} (TMEAN_\tau - TMEAN_{\tau-1}), \qquad (10)$$

where the first line of Equation 10 is identical to Equation 9, except that we use "hat" notation to indicate that $\hat{\tilde{f}}(\cdot)$ is an empirical estimate of expected mortality risk. The second (approximate) equality follows from (i) taking the total and partial derivative of our estimating equation (Equation 5) with respect to climate — where the total derivative accounts for adaptation while the partial does not, (ii) substituting terms and simplifying the expression, and (iii) implementing a discrete-time approximation for the continuous integral (see Appendix A.3 for a full derivation). The under-braced object, $\hat{\boldsymbol{\gamma}}_{\mathbf{1}}\mathrm{E}[\boldsymbol{T}]_\tau$, is the product of the expectation of temperature and the coefficient associated with the interaction between temperature and climate from our estimation of Equation 5: it represents our estimate of marginal adaptation benefits.[40] This derivative is then multiplied by the change in average temperature between each period.[41] Finally, we treat the VSL as a function of income, which evolves as incomes increase over time (see Section 7).

These adaptation cost estimates are calculated annually for each impact region and age group, as in Section 5, and for each of the 33 high-resolution climate model projections. These estimates enable us to develop a complete measure of the mortality costs of climate change that captures both the benefits and costs of adaptation. We continue to call this empirical estimate of Equation 4 the *full mortality risk of climate change*:

(iv) *Full mortality risk due to climate change* (including adaptation costs, recall Equation 4):

$$\underbrace{\hat{g}(\boldsymbol{T}_t, TMEAN_t, \log(GDPpc)_t) - \hat{g}(\boldsymbol{T}_{t_0}, TMEAN_{t_0}, \log(GDPpc)_t)}_{\text{mortality effects of climate change with benefits of income growth and adaptation (iii)}}$$
$$+ \underbrace{\frac{1}{VSL_t}\left[\widehat{A(TMEAN_t, GDPpc_t) - A}(TMEAN_{t_0}, GDPpc_t)\right]}_{\text{estimated adaptation costs}}.$$

The adaptation cost term is multiplied by $\frac{1}{VSL}$ to convert it from dollars to lives. This conversion is important because it enables us to report the full mortality risk of climate change in a single unit, lives, rather than in lives and dollars. We note that using human lives serves as a natural numeraire in this revealed preference framework since we estimate adaptation costs based on lives that could be saved via adaptation, but are not.

[40]Recall that the specific functional form we use to estimate mortality risk as a function of temperature, climate, and income is $g(\cdot) = (\boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_1 TMEAN_t + \boldsymbol{\gamma}_2 \log(GDPpc)_t)\, \boldsymbol{T}_t$. Thus, the partial derivative $\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}$ is equivalent to $\hat{\boldsymbol{\gamma}}_1\mathrm{E}[\boldsymbol{T}]_\tau$.

[41]We assume that individuals use the recent past to form expectations about current temperature realizations, so this expectation is computed over the prior 15 years, with weights of historical observations linearly declining in time.

We refer to these as "death equivalents", or the number of avoided deaths equal in value to the adaptation costs incurred.

6.3 Projections of the full mortality risk of climate change, accounting for adaptation benefits and costs

Table 2 summarizes the results for the full mortality risks of climate change at the end of the century, accounting for adaptation benefits and costs. The columns follow expressions (i)-(iv) detailed in Sections 5 and 6.2. Specifically, column 1 reports the mortality cost of climate change when the beneficial impacts of adaptation and income are shut down. Columns 2 and 3 show the change in mortality risk due to the benefits of income growth and climate adaptation, respectively; both reduce mortality, so the entries are negative. Column 4 presents estimates of adaptation costs in units of "death equivalents", following the calculation in Section 6.2. Finally, columns 5a and 5b show the full mortality risk of climate change, measured in deaths per 100,000 and monetized as a proportion of total global GDP in 2100.

6.3.1 Global estimates of the full mortality risk of climate change

Panel A of Table 2 shows mean estimates for the globe, averaging over a set of Monte Carlo simulations accounting for both climate and statistical uncertainty. The interquartile ranges across simulation runs are in brackets. Column 5a shows that, on average across the globe, the estimated full mortality risk due to climate change (i.e., expression (iv)) is projected to equal ∼85 deaths per 100,000 under RCP8.5 by 2100 (Appendix Figure F2 shows annual results over the century and Table F1 shows results for RCP4.5).[42] Of this full mortality risk, climate adaptation costs are estimated at ∼12 death equivalents per 100,000 (column 4), while increases in mortality rates account for the remaining 73 deaths per 100,000 (sum of columns 1 through 3). It is noteworthy that our estimate for the global average benefits of adaptation (column 3; 31 deaths per 100,000) exceeds the costs of these adjustments, demonstrating that the adaptation surplus of 19 deaths per 100,000 is substantial.[43]

[42] We previously noted considerable heterogeneity across age-groups in our results. We display the underlying age group heterogeneity of these projections in Appendix F.

[43] Appendix A.2 details the derivation of adaptation benefits and adaptation surplus.

Table 2: Estimated 2100 full mortality risks of climate change, globally and regionally (high emissions scenario, RCP8.5)

	Mortality effects of climate change				Full mortality risk of climate change	
	No income growth or adaptation	Benefits of income growth	Benefits of climate adaptation	Costs of climate adaptation		
	deaths/100k	*deaths/100k*	*deaths/100k*	*deaths/100k*	*deaths/100k*	*% of GDP*
	(1)	(2)	(3)	(4)	(5a)	(5b)
Panel A: Global estimates						
Mean Impacts	220.6	-116.5	-31.0	11.7	84.8	3.2
IQR	[76.4, 258.8]	[-149.4, -39.2]	[-60.1, 3.8]	[0.2, 19.4]	[17.4, 116.4]	[-5.4, 9.1]
Panel B: Regional estimates						
China	112.0	-81.8	-28.8	17.7	19.1	1.9
USA	14.8	-13.2	-1.8	10.2	10.1	1.0
India	334.4	-248.2	-25.6	2.1	62.7	6.0
Pakistan	589.1	-161.7	-105.0	53.6	376.0	27.5
Bangladesh	382.5	-89.3	-79.3	34.7	248.5	18.5
Europe	-14.3	-6.2	-74.8	90.8	-4.7	0.1
Sub-Saharan Africa	232.5	-77.4	-34.5	10.5	131.8	8.4

All columns show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In the first row, brackets indicate the interquartile range (IQR). Columns 1-4 each indicate a partial mortality effect of climate change, in units of deaths per 100,000. Column 1 (expression (i)): mortality effects of climate change without benefits of income or adaptation to climate change. Column 2 (expression (ii) - expression (i)): benefits of income growth. Column 3 (expression (iii) - expression (ii)): benefits of adaptation to climate change. Column 4 (Equation 10): mortality-related costs of adaptation inferred using a revealed preference approach, measured in "death equivalents". Columns 5a-5b (expression (iv)): the full mortality risk of climate change, measured in deaths per 100,000 (column 5a) and represented as % of 2100 GDP (column 5b) using an age-adjusted value of the U.S. EPA VSL with an income elasticity of one applied to all impact regions. Column 5a is equivalent to the sum of columns 1 through 4. All estimates shown rely on the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario. Table F1 shows equivalent results for SSP3 and RCP4.5 and details the regional definitions for Europe and sub-Saharan Africa.

Column 5b of Table 2 shows the monetized full mortality risk of climate change at the end of the century. To construct these estimates, we use the value of a statistical life (VSL) to convert changes in mortality rates into dollars. Our primary approach relies on the U.S. EPA's VSL estimate of $10.95 million (2019 USD).[44] We transform the VSL into a value per life-year lost using a method described in Appendix H.1, which allows us to compute the total value of expected life-years lost due to climate change, accounting for the different mortality-temperature relationships among the three age groups documented above. We allow the VSL to vary with income, as the level of consumption affects the relative marginal utilities of a small increment of consumption and a small reduction in the probability of death. Consistent with existing literature (e.g., Viscusi, 2015), our primary estimates use an income elasticity of unity to adjust the U.S. estimates of the VSL to different income levels across the world and over time.[45] When computing the mortality partial SCC in Section 7, we provide multiple alternative valuation assumptions in addition to this benchmark case.

The resulting estimates in column 5b are substantial. For example, under RCP8.5, they amount to 3.2% of global GDP in 2100, with an interquartile range of [-5.4%, 9.1%]. Under RCP4.5 (shown in Table F1), they fall to 0.6% [-3.9%, 4.6%] of global GDP. The uncertainty around these estimates is also meaningful and while we leave explicit pricing of this uncertainty to future work, accounting for it with a certainty equivalence-style calculation would only increase the estimated welfare loss from climate change.

These results suggest that the mortality risks from climate change are much greater than had previously been understood. For instance, these mortality-related damages amount to ∼49-135% of the damages reported for *all sectors of the economy* in FUND, PAGE, and DICE, when the damage functions from each model are evaluated at the mean end-of-century warming observed in our multi-model ensemble under RCP8.5. Under RCP4.5, our mortality-related damages amount to 32-61% of the damages from DICE and PAGE, while damages from FUND are negative at RCP4.5 levels of warming.[46]

The results in this and the previous section have relied on a single benchmark emissions and socioeconomic scenario (RCP8.5, SSP3). Appendix F reports on the sensitivity of the full mortality risk of climate change results to alternative choices about the economic and population scenario, the emissions scenario, and assumptions regarding the rate of adaptation. These exercises underscore that the projected impacts of climate change over the remainder of the 21^{st} century depend on difficult-to-predict factors such as policy, technology, and demographics. However, we note that under both emissions scenarios RCP8.5 and RPC4.5, under all SSP scenarios, and under an alternative projection in which the rate of adaptation is deterministically slowed, the average estimate of the full mortality risk due to climate change is positive (both RCPs) and steadily increasing (RCP8.5) throughout the 21^{st} century.

[44] This VSL is from the 2012 U.S. EPA Regulatory Impact Analysis (RIA) for the Clean Power Plan Final Rule, which provides a 2020 income-adjusted VSL in 2011 USD, which we convert to 2019 USD. This VSL is also consistent with income- and inflation-adjusted versions of the VSL used in the U.S. EPA RIAs for the National Ambient Air Quality Standards (NAAQS) for Particulate Matter (2012) and the Repeal of the Clean Power Plan (2019), among many other RIAs.

[45] The EPA considers a range of income elasticity values for the VSL, from 0.1 to 1.7 (U.S. Environmental Protection Agency, 2016b), although their central recommendations are 0.7 and 1.1 (U.S. Environmental Protection Agency, 2016). A review by Viscusi (2015) estimates an income-elasticity of the VSL of 1.1.

[46] To conduct this comparison, we use the damage functions reported for each IAM in the Interagency Working Group on Social Cost of Carbon (2010), which are indexed against warming relative to the pre-industrial climate. We evaluate each damage function at the mean end-of-century warming (4°C for RCP8.5 and 1.8°C for RCP4.5) across the SMME climate model ensemble used in our analysis, after adjusting warming to align pre-industrial temperature anomalies from the IAMs with the anomalies relative to 2001-2010 from our analysis (Lenssen et al., 2019). We note that these leading IAMs use different socioeconomic scenarios and climate models than those used throughout this paper.

6.3.2 Unequal distribution of mortality risk from climate change.

Panel B of Table 2 displays estimates of the end-of-century mortality risk of climate change for select countries and regions of the world. These results indicate that the full mortality risk caused by climate change varies substantially across the globe. Notably, monetized estimates in column 5b are very high in some regions, such as Pakistan and Bangladesh, where impacts amount to 27.5% and 18.5% of GDP, respectively.[47] The share of the full mortality risk that is due to actual deaths (first term in expression (iv)) versus compensatory investments (second term in expression (iv)) also differs across regions. Some locations suffer large increases in mortality rates, such as India, where 97% of the full mortality risk due to climate change is attributable to rising death rates. Other regions avoid excess mortality through expensive adaptation. For example, the U.S. is projected to benefit from a small decline in the mortality rate of -0.2 deaths per 100,000 at end of century, but is also projected to incur adaptation costs amounting to 10 death equivalents per 100,000.

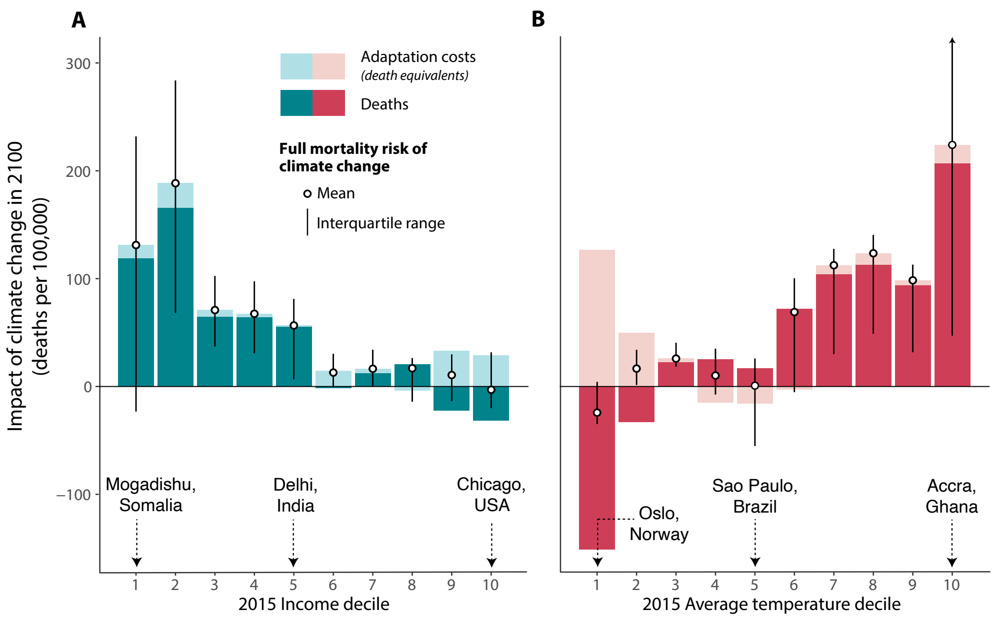

Figure 7: Climate change impacts and adaptation costs are correlated with present-day income and climate. Figure shows the mortality risk of climate change in 2100 (RCP8.5, SSP3) against deciles of 2015 per capita income (A) and average annual temperature (B). Dark colors indicate mean changes in death rates, accounting for the benefits of income growth and climate adaptation, while light colors indicate mean changes in adaptation costs, measured in death equivalents. For all bars shown, means are taken across impact regions falling into the corresponding decile of income or climate and across Monte Carlo simulations that account both for econometric and climate model uncertainty. Black outlined circles indicate the mean estimate of the full mortality risk of climate change, which is the sum of deaths and adaptation costs, and black vertical lines indicate the interquartile range of the distribution across impact regions within each decile. The income and average temperature deciles are calculated across 24,378 global impact regions and are population weighted using 2015 population values.

To visualize these distributional consequences, Figure 7 plots the full mortality risk of climate change in 2100 (dark bars), as well as the mean impact of climate change on adaptation costs (light bars), against deciles of present-day income (panel A) and present-day average temperature (panel B). These results reveal

[47] Note that Table 2 indicates that for Europe, the full mortality risk of climate change as measured in deaths per 100,000 (column 5a) is negative, while it is positive when measured in % of GDP (column 5b). This is because throughout much of Europe, climate change leads to lives being saved due to fewer extremely cold days, particularly for the >64 age group. Under the valuation approach shown in Table 2, an age-adjusted VSL is used, which lowers the relative weight placed on these lives saved in the older age group, as compared to increased mortality risk due to hot days in other age groups.

that the magnitude and composition of future mortality risks under climate change are strongly correlated with current incomes and climate. Panel A shows that the share of the full mortality risk due to adaptation costs is higher at higher incomes, indicating that wealthier locations are predicted to pay for future adaptive investments, while such costs are predicted to be much smaller in poor parts of the globe. In contrast, mortality rates are projected to increase much more dramatically in today's poor countries, indicating that climate impacts in these places will largely take the form of people living shorter lives. Further, the full mortality risk of climate change (shown in black and white circles) is still borne disproportionately by regions that are poor today. Finally, there is substantial variance across impact regions within each income decile, as shown by the interquartile range, underscoring the importance of geographic resolution in projecting climate impacts.

A similar figure in panel B demonstrates that the hottest locations today suffer the largest predicted increases in death rates, while the coldest are estimated to pay the highest adaptation costs. The magnitude of impacts in the top decile of the current long-run climate distribution are noteworthy and raise questions about the habitability of these locations at the end of the century.

7 The mortality partial social cost of carbon

This section uses the estimates of the full mortality risk of climate change to monetize the mortality-related social cost generated by emitting a marginal ton of CO_2. This calculation represents the component of the *total* SCC that is mediated through excess mortality, but it leaves out adverse impacts in other sectors of the economy, such as reduced labor productivity or changing food prices. Hence, it is a mortality *partial* SCC.

7.1 Definition: the mortality partial social cost of carbon

The mortality partial social cost of carbon at time t is defined as the marginal social cost from the change in mortality risk imposed by the emission of a marginal ton of CO_2 in time period t. For a discount rate δ, the mortality partial SCC is:

$$\textit{Mortality partial SCC}_t \ \text{(dollars)} = \int_t^\infty e^{-\delta(s-t)} \frac{dD_s(\boldsymbol{C}_s)}{d\boldsymbol{C}} \frac{\partial \boldsymbol{C}_s}{\partial E_t} ds, \tag{11}$$

where $D_s(\boldsymbol{C}_s)$ represents a "damage function" describing *total* global economic losses (inclusive of both adaptation benefits and costs) due to changes in mortality risk in time period s, as a function of the global climate $\boldsymbol{C}$ (Nordhaus, 1992; Hsiang et al., 2017), and where E_t represents total global greenhouse gas emissions in period t. $D_s(\cdot)$ varies over time, s, because the mortality sensitivity of temperature and total monetized impacts of climate change evolve over time due to changes in per capita income, the climate, and in the underlying population. Thus, the damages from a marginal change in emissions will vary depending on the year in which they are evaluated. In practice, we approximate Equation 11 by combining empirically grounded estimated damage functions $D_s(\cdot)$ with climate model simulations of the impact of a small change in emissions on the global climate, i.e., $\frac{\partial \boldsymbol{C}_s}{\partial E_t}$.

Expressing the mortality partial SCC using a damage function has three key practical advantages. First, the damage function represents a parsimonious, reduced-form description of the otherwise complex dependence of global economic damage on the global climate. Second, as we demonstrate below, it is possible to empirically estimate damage functions from the climate change projections described in Section 6. Finally, because they are fully differentiable, empirical damage functions can be used to compute *marginal* costs of an emissions impulse released in year t by differentiation. The construction of these damage functions, as well as the implementation of the mortality partial SCC, are detailed in the following subsections.

7.2 Constructing damage functions for excess mortality risk

There are two key components of a damage function for excess mortality risk. First, the change in global mean surface temperature, $\Delta GMST_{rmt}$, which indicates the overall magnitude of warming. We compute this value for each of the two emissions scenarios r, each of the 33 climate models m, and each year t.[48] Second, total monetized losses due to changes in mortality risk, inclusive of adaptation benefits and cost, D_{irmt}, captures total damages for a given level of warming. We compute this value by summing projected estimates of the monetized full mortality risk of climate change across all 24,378 global impact regions, separately for each draw i of the uncertain parameters recovered from estimation of the mortality-temperature relationship in Equation 5, emissions scenario r, climate model m, and year t. Therefore, for a given value of $\Delta GMST_{rmt}$, there is variation in damages D_{irmt} due to econometric uncertainty captured by simulation runs i and differential spatial distribution of warming across climate models.

Due to differences in the source of climate projections pre- and post-2100, and lack of available socioeconomic projections after 2100, there are some important methodological differences in how we estimate the relationship between damages D_{irmt} and warming $\Delta GMST_{rmt}$ for years before versus after 2100. This subsection details these differences and also explains the approach to account for damage function uncertainty.

7.2.1 Computing damage functions through 2100.

For each year t from 2020 to 2097, we estimate a set of quadratic damage functions that relate the total global value of mortality-related climate change damages (D_{irmt}) to the magnitude of global warming ($\Delta GMST_{rmt}$):

$$D_{irmt} = \alpha + \psi_{1,t}\Delta GMST_{rmt} + \psi_{2,t}\Delta GMST_{rmt}^2 + \varepsilon_{irmt}. \tag{12}$$

Specifically, to construct the damage function separately for each year t, we combine all 9,750 Monte Carlo simulation runs within a 5-year window centered on t and estimate the regression in Equation 12.[49]This approach allows the recovered damage function $D_t(\Delta GMST)$ to evolve flexibly over the century. We note that pre-2100 damage functions are indistinguishable if we use a third-, fourth- or fifth-order polynomial, and we show robustness of our mortality partial SCC estimates to functional form choice in Appendix H.4.

[48]Our climate change impacts are calculated relative to a baseline of 2001-2010. Therefore, we define changes in global mean surface temperature ($\Delta GMST$) as relative to this same period. Note that the $\Delta GMST_{rmt}$ value in each climate model is a summary parameter, resulting from the complex interaction of many physical elements of the model, including the *equilibrium climate sensitivity*, a number that describes how much warming is associated with a specified change in greenhouse gas emissions.

[49]Because the projections in Section 6 end in 2100, 2097 is the last year for which a centered 5-year window of estimated damages can be constructed, and therefore is the last year for which we estimate Equation 12.

Figure 8A illustrates the procedure for the end-of-century damage function. Each data point plots a value of D_{irmt} from an individual Monte Carlo simulation (vertical axis) against the corresponding value of ΔGMST_{rmt} (horizontal axis), where scatter points for years t=2095 through t=2100 are shown. Red points indicate simulation runs from the high emissions scenario (r=RCP8.5) and blue points indicate runs from the low emissions scenario (r=RCP4.5).[50] The median end-of-century warming relative to 2001-2010 under RCP8.5 across our climate models is +3.7°C, while under RCP4.5 it is +1.6°C. The black line is the end-of-century quadratic damage function, estimated following Equation 12.[51] The estimated damage function recovers total (undiscounted) damages with an age-varying VSL at 3.7°C and 1.6°C of $7.8 and $1.2 trillion USD, respectively. Analogous curves are constructed for all years, starting in 2020.

7.2.2 Computing post-2100 damage functions

Even with standard discount rates, a meaningful fraction of the present discounted value of damages from the release of CO_2 today will occur after 2100 (Kopp and Mignone, 2012), so it is important to develop post-2100 damage functions. The pre-2100 approach cannot be used for these later years because only 6 of the 21 GCMs that we use to build our SMME ensemble simulate the climate after 2100 for both RCP scenarios. Further, the SSPs needed to project the benefits of income growth and changes in demographic compositions also end in 2100.

To estimate post 2100-damages, we develop a method to extrapolate changes in the damage function beyond 2100 using the observed evolution of damages near the end of the 21^{st} century. The motivating principle of the extrapolation approach is that these observed changes in the shape of the damage function near the end of the century provide plausible estimates of future damage function evolution after 2100. This reduced-form approach allows our empirical results to constrain and guide a projection to years beyond 2100. To execute this extrapolation, we pool values D_{irmt} from 2085-2100 and estimate a quadratic model similar to Equation 12, but interacting each term linearly with year t.[52] This allows estimation of a damage surface as a parametric function of year, which can then be used to predict extrapolated damage functions for all years after 2100, smoothly transitioning from our climate model-based damage functions prior to 2100. Appendix G provides a detailed explanation of the approach.

Panel B of Figure 8 illustrates damages functions every 10 years prior to 2100, as well as extrapolated damage functions for the years 2150, 2200, 2250, and 2300. In dollar terms, these extrapolated damages continue to rise post-2100, suggesting larger damages for a given level of warming. This finding comes directly from the estimation of Equation 12 that found that in the latter half of the 21^{st} century the full mortality damages are larger when they occur later, holding constant the degree of warming. This finding that mortality costs rise over time is the net result of countervailing forces. On the one hand, damages

[50] This scatterplot includes simulation runs for RCP4.5 and RCP8.5 for all projections in our 33-member ensemble under our benchmark method of valuation – the age-invariant EPA VSL with an income elasticity of one applied to all impact regions – in the end-of-century years 2095-2100. See Appendix H for results across different valuation assumptions. Due to the dependence of damages on GDP per capita and on demographics, we estimate separate damage functions following Equation 12 separately for every SSP scenario. Results across different scenarios are also shown in Appendix H.

[51] The damage function in Figure 8 is estimated for the year 2097, the latest year for which a full 5-year window of damage estimates can be constructed.

[52] We use 2085-2100 because the time evolution of damages becomes roughly linear conditional on ΔGMST by this period.

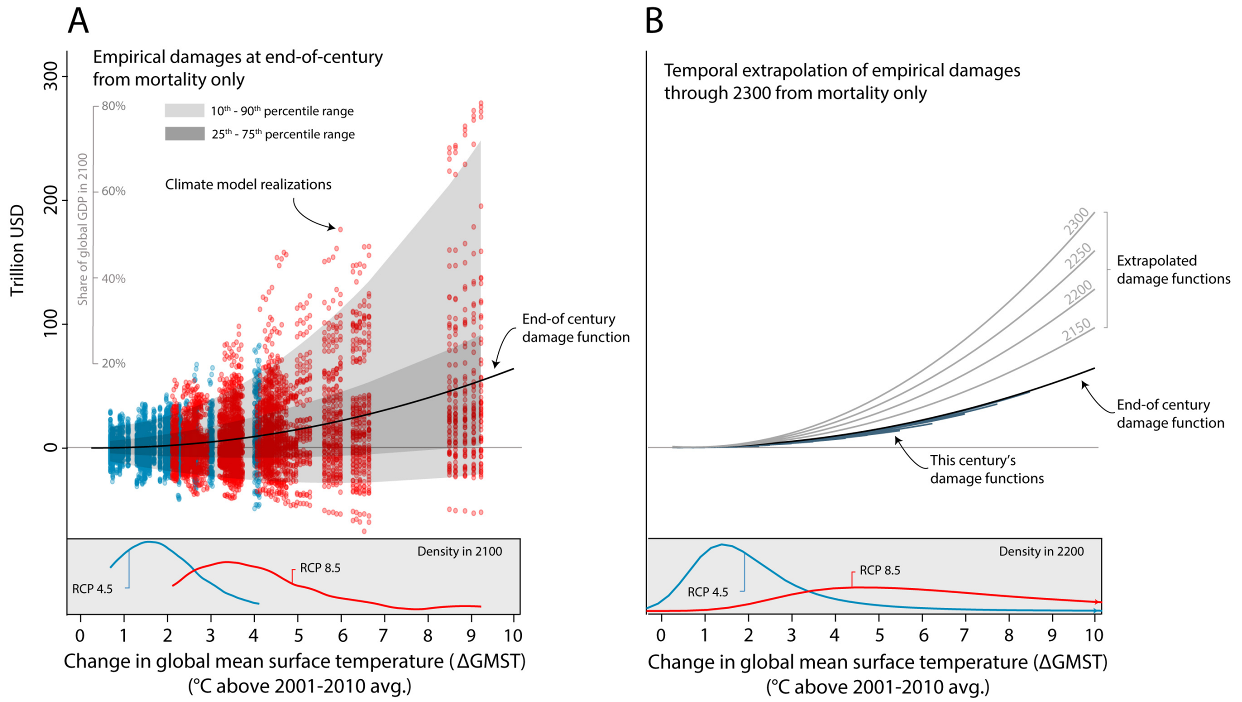

Figure 8: Empirically-derived mortality-only damage functions. Both panels show damage functions relating empirically-derived total global mortality damages to anomalies in global mean surface temperature (ΔGMST) under socioeconomic scenario SSP3. In panel A, each point (red = RCP8.5, blue = RCP4.5) indicates the value of the full mortality risk of climate change in a single year (ranging from 2095 to 2100) for a single simulation of a single climate model, accounting for both costs and benefits of adaptation. The black line is the quadratic damage function estimated through these points. The distribution of temperature anomalies at end of century (2095-2100) under two emissions scenarios across our 33 climate models is in the bottom panel. In panel B, the end-of-century damage function is repeated. Damage functions are shown in dark blue for every 10 years pre-2100, each of which is estimated analogously to the end-of-century damage function and is shown covering the support of ΔGMST values observed in the SMME climate models for the associated year. Our projection results generate mortality damages only through 2100, due to limited availability of climate and socioeconomic projections for years beyond that date. To capture impacts after 2100, we extrapolate observed changes in damages over the 21^{st} century to generate time-varying damage functions through 2300. The resulting damage functions are shown in light grey for every 50 years post-2100, each of which is extrapolated. The distribution of temperature anomalies around 2200 (2181-2200) under two emissions scenarios using the FAIR simple climate model is in the bottom panel. To value lives lost or saved, in both panels we use the age-varying U.S. EPA VSL and an income elasticity of one applied to all impact regions.

are larger in later years because there are larger and older populations[53] with higher VSLs due to rising incomes. On the other hand, damages are smaller in later years because populations are better adapted due to higher incomes and a slower rate of warming, enabling gradual adaptation. The results suggest the former dominates by end of century, causing damages to be trending upward when high-resolution simulations end in 2100. Below and in Appendix H, we explore the sensitivity of the results to alternative extrapolation approaches.

7.2.3 Accounting for uncertainty in damage function estimation.

As discussed, there is substantial uncertainty in projected mortality effects of climate change due to statistical uncertainty in the estimation of mortality-temperature response functions, as well as climate uncertainty arising from differences across the 33 climate models. The approach described above details the estimation

[53] In SSP3, the share of the global population in the >64 age category rises from 8.2% in 2015 to 16.2% in 2100.

of a damage function using the conditional expectation function through the full distribution of simulation results. In addition to reporting the predicted damages resulting from this damage function describing (conditional) expected values, we also estimate a set of quantile regressions to capture the full distribution of simulated mortality impacts.[54] Just as above for the mean damage function, extrapolation past the year 2100 is accomplished using a linear time interaction, estimated separately for each quantile. In the sections below, we use these quantiles to characterize uncertainty in the mortality partial SCC estimates. Thus, central estimates of the mortality partial SCC use the mean regression from Equation 12, while ranges incorporating damage uncertainty use the full set of time-varying quantile regressions.

7.3 Computing marginal damages from a marginal carbon dioxide emissions pulse

We empirically approximate the mortality partial SCC from Equation 11 for emissions that occur in the year 2020 as:

$$Mortality\ partial\ SCC_{2020} \approx \sum_{t=2020}^{2300} e^{-\delta(t-2020)} \frac{\partial \hat{D}_t(\Delta GMST)}{\partial \Delta GMST} \frac{d\Delta GMST_t}{dCO2_{2020}}, \tag{13}$$

where $\Delta GMST$ approximates the multi-dimensional climate vector $\boldsymbol{C}$, and changes in CO_2 represent changes in global emissions E.[55] Additionally, we assume that discounted damages from an emissions pulse in year 2020 become negligible after 2300, and we approximate the integral in Equation 11 with a discrete sum using time steps of one year. The values $\frac{\partial \hat{D}_t(\Delta GMST)}{\partial \Delta GMST}$ are the marginal global damages in each year t that occur as a result of this small change in all future global temperatures; they are computed using the damage functions described in the last subsection. The term $\frac{d\Delta GMST_t}{dCO2_{2020}}$ is the increase in ΔGMST that occurs at each year t along a baseline climate trajectory as a result of a marginal unit of emissions in 2020, which we approximate with an "infinitesimally small" pulse of CO_2 emissions. Because it is computationally infeasible to compute this value and account for uncertainty about the physical magnitude and timing of warming for all 33 climate models in the SMME, we use an alternative, global climate model to estimate $\frac{d\Delta GMST_t}{dCO2_{2020}}$, as detailed below.

7.3.1 Applying a simple climate model to the damage function.

To calculate the change in $\Delta GMST_t$ due to a marginal pulse of CO_2 in 2020, we use the Finite Amplitude Impulse Response (FAIR) simple climate model. We use FAIR to calculate $\Delta GMST_t$ trajectories for emissions scenarios RCP4.5 and RCP8.5, both with and without an exogenous "pulse" of 1 gigaton C (equivalent to 3.66Gt CO_2) in the year 2020, the smallest emission quantity for which a warming signal can be separated from noise within the FAIR climate model. In FAIR, this emissions pulse perturbs the trajectory of atmospheric CO_2 concentrations and $\Delta GMST$ for 2020-2300, with dynamics that are influenced by the baseline RCP scenario. In each emissions scenario, we then predict damages $\hat{D}_t(\Delta GMST_t)$ for $\Delta GMST$ values from

[54] We estimate a damage function for every 5^{th} percentile from the 5^{th} to 95^{th}.

[55] We use CO_2 to represent changes in all global greenhouse gas (GHG) emissions as it is the most abundant GHG and the warming potential of all other GHGs are generally reported in terms of their CO_2 equivalence.

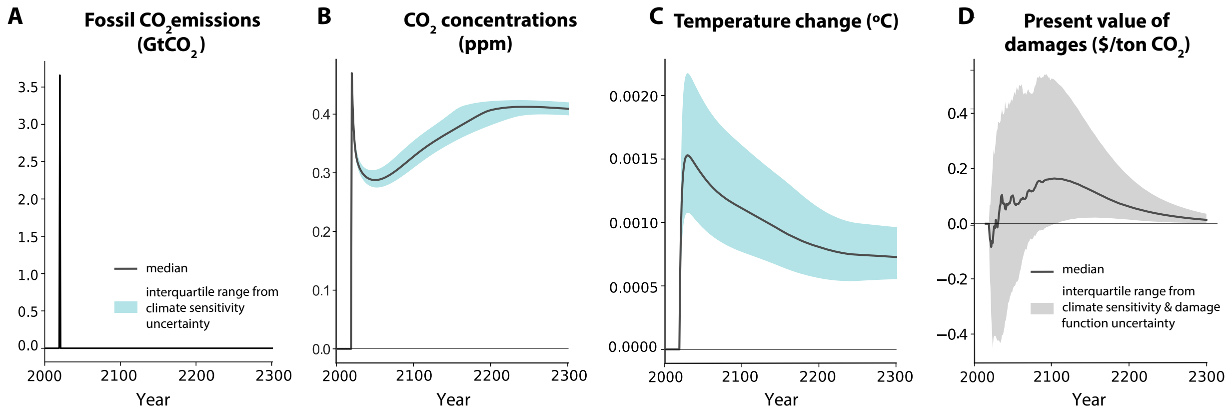

Figure 9: Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020. Panel A shows a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) in 2020 for emissions scenario RCP8.5. Panel B displays the effect of this pulse on atmospheric CO_2 concentrations, relative to the baseline. In panel C, the impact of the pulse of CO_2 on temperature is shown where the levels are anomalies in global mean surface temperature (GMST) in Celsius. In panels A-C, shaded areas indicate the inter-quartile range due to climate sensitivity uncertainty, while solid lines are median estimates. Panel D shows the change in discounted damages over time due to a 1 Gt pulse of CO_2 in 2020 under socioeconomic scenario SSP3, as estimated by our empirically-derived damage functions, using a 2% annual discount rate and the age-varying U.S. EPA VSL with an income elasticity of one applied to all impact regions. The shaded area indicates the inter-quartile range due to climate sensitivity and damage function uncertainty, while the solid line is the median estimate.

the "RCP + pulse" simulation and difference them from predicted damages for $\Delta GMST$ values from the baseline "RCP only" simulation. The resulting damages due to the pulse are converted into USD per one metric ton CO_2. There is naturally uncertainty in these $\Delta GMST$ trajectories, and our approach accounts for uncertainty associated with four key parameters of the FAIR model (i.e., the transient climate response, equilibrium climate sensitivity, the short thermal adjustment time, and the time scale of rapid carbon uptake by the ocean mixed layer). This approach, detailed in Appendix G, ensures that the distribution of transient warming responses we use to generate partial SCC values matches the corresponding distributions from the IPCC Assessment Report 5 (AR5).

7.3.2 Summarizing the impacts of a marginal increase in CO_2 emissions.

Figure 9 graphically depicts the difference between the "RCP + pulse" and baseline RCP trajectories for four key outcomes.[56] The pulse in emissions is shown in panel A. Its influence on CO_2 concentrations is reported in panel B; the immediate decline followed by a century-long increase is largely due to dynamics involving the ocean's storage and release of emissions. Panel C displays the resulting change in temperature, which makes clear that a pulse today will influence temperatures even three centuries later. The solid lines are median estimates, while the shaded blue area in panels A-C depicts the inter-quartile range of each year's outcome, reflecting uncertainty about the climate system (see Appendix G for details).

Panel D plots the discounted (2% discount rate) stream of damages due to this marginal pulse of emissions. The temporal pattern of the present value of mortality damages reflects several factors, including: the nonlinearity of the damage function (e.g., Figure 8), which itself depends on nonlinearities in location-

[56]Using the trajectories in Figure 9 is consistent with the "SCC experiment" that is used in IAMs to calculate an SCC (National Academies of Sciences, Engineering, and Medicine, 2017). We discuss uncertainties in FAIR configuration parameters below and in Appendix G. The median values of parameter-specific distributions used for the central mortality partial SCC estimate include a transient climate response (TCR) of 1.6 and an equilibrium climate sensitivity (ECS) of 2.7.

specific mortality-temperature relationships (e.g., Figure 1); the discount rate; and the dynamic temperature response to emissions (panel C). The peak present value of annual damages from a ton of CO_2 emissions are $0.16 in year 2104; by year 2277, annual damages are always less than $0.02. It is noteworthy that about two-thirds of the present value of damages occur after the year 2100. The shaded grey area represents the inter-quartile range of each year's outcome, and reflects uncertainty in the climate system as well as uncertainty in the damage function. RCP4.5 results are shown in Figure G5 and additional details are in Appendix G.

7.4 Estimates of the partial social cost of carbon due to excess mortality risk

Table 3 reports mortality partial SCC estimates. The columns apply four different annual discount rates – three used in prior estimates of the SCC (2.5%, 3%, and 5%) (Interagency Working Group on Social Cost of Carbon, 2010; National Academies of Sciences, Engineering, and Medicine, 2017), and one lower rate that aligns more closely with recent global capital markets (2%) (Board of Governors of the US Federal Reserve System, 2020). Panel A uses the U.S. EPA's VSL of $10.95 million (2019 USD), transformed into value per life-year lost (see Appendix H.1 for details). This accounts for the different mortality-temperature relationships among the three age groups documented above.[57] Panel B is based on an age-invariant value of $10.95 million (2019 USD) for the VSL. Both approaches then adjust for cross sectional variation in incomes among contemporaries and global income growth. Appendix H presents results under a wide range of additional valuation scenarios, including an alternative and lower Ashenfelter and Greenstone (2004) VSL of $2.39 million (2019 USD),[58] and an approach where the VSL is adjusted only based on global average income such that the lives of contemporaries are valued equally, regardless of their relative incomes.

The estimates in Table 3 utilize the median values of FAIR's four key parameter distributions and the mean global damage function. Interquartile ranges (IQRs) are reported, reflecting uncertainty in climate sensitivity (uncertainty in the simple climate model FAIR) and in the damage function. All values represent the global sum of each impact region's MWTP today (2019 USD) to avoid the release of an additional metric ton of CO_2 in 2020, including both the costs and benefits of adaptation.

Column 1 and panel A of Table 3 reports our preferred estimates of the mortality partial SCC. These are based on a $\delta = 2\%$ discount rate and an age-varying VSL. Under this valuation approach, the mortality partial SCC is $17.1 [-$24.7, $53.6] for the low to moderate emissions scenario and $36.6 [-$7.8, $73.0] for the high emissions scenario. We highlight a 2% discount rate because it conservatively reflects changes in global capital markets over the last several decades: while the Interagency Working Group on Social Cost of Greenhouse Gases (2016) recommends a discount rate of 3% based on the real 10-year Treasury rate calculated in 2003, the average 10-year Treasury Inflation-Indexed Security from 2003 to present is just 1.01% (Board of Governors of the US Federal Reserve System, 2020). We show results for a discount rate

[57] In the main text, a simple life-years calculation that assigns each life-year lost the same economic value is used. In Appendix H, we also show calculations that adjust the value of remaining life by age at death using the estimates of age-specific value of remaining life from Murphy and Topel (2006), which produces results that differ only slightly from those under the primary approach.

[58] See Appendix Table H1 for a comparison of these VSL values with values from the OECD, which are higher than Ashenfelter and Greenstone (2004), but lower than the U.S. EPA's VSL.

Table 3: Estimates of a partial social cost of carbon for excess mortality risk incorporating the costs and benefits of adaptation

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 US Dollars)				
Moderate emissions scenario (RCP4.5)	17.1	11.2	7.9	2.9
Full uncertainty IQR	[-24.7, 53.6]	[-18.9, 36.0]	[-15.2, 26.3]	[-8.5, 11.5]
High emissions scenario (RCP8.5)	36.6	22.0	14.2	3.7
Full uncertainty IQR	[-7.8, 73.0]	[-10.6, 46.8]	[-11.4, 32.9]	[-8.9, 13.0]
Panel B: Globally varying value of a statistical life (2019 US Dollars)				
Moderate emissions scenario (RCP4.5)	14.9	9.8	6.7	1.7
Full uncertainty IQR	[-21.2, 63.5]	[-17.9, 43.5]	[-15.7, 32.1]	[-11.8, 14.7]
High emissions scenario (RCP8.5)	65.1	36.9	22.1	3.5
Full uncertainty IQR	[3.0, 139.0]	[-2.4, 83.1]	[-5.6, 53.4]	[-9.3, 16.0]

In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value (alternative values using the VSL estimate from (Ashenfelter and Greenstone, 2004) are shown in Appendix H). All regions thus have heterogeneous valuation, based on local income. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3 (alternative values using other SSP scenarios are shown in Appendix H). In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details and Appendix H.2 for alternative calculations that allow the value of a life-year to vary with age, based on Murphy and Topel (2006)). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Point estimates rely on the median values of the four key input parameters into the climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] including both climate sensitivity uncertainty and damage function uncertainty (see Appendix G for details).

of 1.5% in Appendix Table H4. We emphasize the age-varying VSL approach because standard economic reasoning implies that valuation of life lost should vary by age (Jones and Klenow, 2016; Murphy and Topel, 2006).

When following the Interagency Working Group on Social Cost of Greenhouse Gases (2016) preference for a discount rate of $\delta = 3\%$ and the use of an age-invariant VSL, the central estimate of the mortality partial SCC is \$6.7 per metric ton of CO_2 for the low to moderate emissions scenario (RCP4.5), with an IQR of [-\$15.7, \$32.1], and \$22.1 [-\$5.6, \$53.4] per metric ton for the high emissions scenario (RCP8.5).

Some other features of these results are worth underscoring. First, mortality partial SCC estimates for RCP4.5 are systematically lower than RCP8.5 primarily because the damage function is convex, so marginal damages increase in the high emissions scenario. Second, the combination of this convexity, which itself is accentuated at higher quantiles of the damage function, and the skewness of the climate sensitivity distribution causes the distribution of partial SCCs to also be right skewed with a long right tail. For example, the 95^{th} and 99^{th} percentiles of the partial SCC for $\delta = 2\%$ and an age-varying VSL for RCP8.5 are \$290.3 and \$704.1, respectively; with $\delta = 3\%$ and an age-invariant VSL these values are \$175.3 and \$391.9. It is apparent that a certainty equivalent estimate of the mortality partial SCC based on standard assumptions of risk aversion, which is beyond the scope of this paper, would be much larger than the mean estimates reported here. Third, in Appendix G we show that uncertainty in the partial SCC is consistently dominated by uncertainty in the damage function, as opposed to uncertainty in climate sensitivity. Fourth, all mortality partial SCC estimates shown in the main text rely on an exogenous socioeconomic trajectory; in Table H6 we show that endogenizing impacts of climate change on income growth based on prior literature (Burke, Hsiang, and Miguel, 2015) has only a small effect on our mortality partial SCC results. Fifth, in Appendix H we show that replacing the extrapolation of damage functions to years beyond 2100 with a damage function frozen at its 2100 shape for all years 2101-2300 lowers our central estimate of the mortality partial SCC by 21%. This indicates that damage function extrapolation has a modest impact on our partial SCC estimates, due in part to the important role of discounting (Table H7). Finally, while Table 3 reflects what we believe to be mainstream valuation and socioeconomic scenarios, Appendix Tables H2-H8 report estimates based on multiple alternative approaches. Naturally, the resulting SCC estimates vary under different valuation assumptions and baseline socioeconomic trajectories, and we point readers to these specifications for a more comprehensive set of results.

8 Limitations

As the paper has detailed, the mortality risk changes from climate change and the mortality partial SCC have many ingredients. We have tried to probe the robustness of the results to each of them, but there are three issues that merit special attention when interpreting the results, because they are outside the scope of the analysis.

8.0.1 Migration Responses.

First, the estimates in the paper do not reflect the possibility of migration responses to climate change. If migration were costless, it seems likely that the full mortality risk and mortality partial SCC would be smaller, as people from regions with high damages (e.g., sub-Saharan Africa) may move to regions with low or even negative damages (e.g., Scandinavia). However, both distant and recent history in the U.S. and around the globe underscores that borders are meaningful and that there are substantial costs to migration which might limit the scale of feasible migrations. Indeed, existing empirical evidence of climate-induced migration, based on observable changes in climate to date, is mixed (Carleton and Hsiang, 2016).

8.0.2 Humidity.

Second, our estimates do not directly incorporate the role of humidity in historical mortality-temperature relationships nor in projections of future impacts. There is growing evidence that humidity influences human health through making it more difficult for the human body to cool itself during hot conditions (e.g., Sherwood and Huber, 2010; Barreca, 2012). While temperature and humidity are highly correlated over time, they are differentially correlated across space, implying that our measures of heterogeneous mortality-temperature relationships may be influenced by the role of humidity. To date, lack of high-resolution historical humidity data and highly uncertain projections of humidity under climate change (Sherwood and Fu, 2014) have limited our ability to include this heterogeneity in our work. However, emerging work on this topic (Yuan, Stein, and Kopp, 2019; Li, Yuan, and Kopp, 2020) will provide opportunities to explore humidity in future research.

8.0.3 Technological Change.

Third, the paper's projections incorporate advancements in technology that enhance adaptive ability, even though we have not explicitly modeled technological change. In particular, we allow mortality-temperature response functions to evolve in accordance with rising incomes and temperatures and do not restrict them to stay within the bounds of the current observed distribution of temperature responses. However, while our estimates reflect technical advancement as it historically relates to incomes and climate, they do not reflect the seemingly high probability of climate-biased technical change that lowers the *relative* costs of goods which reduce the health risks of high temperatures. Therefore, the paper's results will overstate the mortality risk of climate change if directed technological innovations lower the relative costs of adapting to high temperatures.

9 Conclusion

This paper has outlined a new method for empirically estimating the global costs of climate change for a single sector of the economy using micro data. We have implemented this approach in the context of mortality risks associated with temperature change. Specifically, this paper develops a framework for estimating the annual total impact of climate change on mortality risk, accounting for the benefits and costs of adaptation, both globally and for 24,378 regions that comprise the planet. It then uses these estimates to compute a mortality "partial" SCC, defined as the global marginal willingness-to-pay to avoid the changes in mortality risk caused by the release of an additional metric ton of CO_2.

There are three noteworthy methodological innovations and key findings. First, we leverage highly resolved data covering roughly half of the world's population to estimate flexible empirical models relating mortality rates to temperature. These regressions uncover a plausibly causal U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for those aged 65 and older. Moreover, this relationship is quite heterogeneous across the planet as both income and long-run climate substantially moderate mortality sensitivity to temperature. Further when combined with current global data on climate, income, and population, the results imply that the effect of a hot day (35°C / 95°F) on mortality in the >64 age group is ∼50% larger in regions of the world without available mortality data. This suggests that prior estimates from wealthy economies and temperature climates are likely to understate the impacts of climate change on human mortality.

Second, we use these regression results along with future projections of climate, income, and population to estimate future climate change-induced mortality risk both in terms of fatality rates and its monetized value. We find that, under a high emissions scenario, the projected impact of climate change on mortality will be comparable globally to leading causes of death today, such as cancer and infectious disease (Figure 10). We also estimate large benefits from mitigation, as the end of century estimate of the full mortality risk of climate change falls from 85 deaths per 100,000 under the high emissions growth RCP8.5 scenario to 14 per 100,000 under the more moderate RCP4.5 scenario. Importantly, these projected impacts include the benefits of adaptation to gradual climate change; estimates that do not account for adaptation overstate the mortality impacts of climate change in 2100 by more than a factor of about 3. Additionally, we outline and implement a revealed preference method to infer the costs of these adaptation investments, which amount to, on average, 12 death equivalents per 100,000 out of the total of 85 deaths per 100,000 in 2100 in the RCP8.5 scenario.

Further, the estimated mortality-related damages from climate change are distributed unevenly across the world. For example, by 2100, we project that climate change will cause annual damages equivalent to approximately 160 additional deaths per 100,000 in Accra, Ghana, but will also generate annual benefits equivalent to approximately 9 lives per 100,000 in London, England. Notably, the degree to which the full mortality risk of climate change is realized through actual deaths, as opposed to costly adaptation, varies widely across space and time. For example, Figure 10 shows that today's poor locations tend to bear a larger share of the projected burden in the form of direct mortality impacts, while today's rich face large increases in projected adaptation costs.

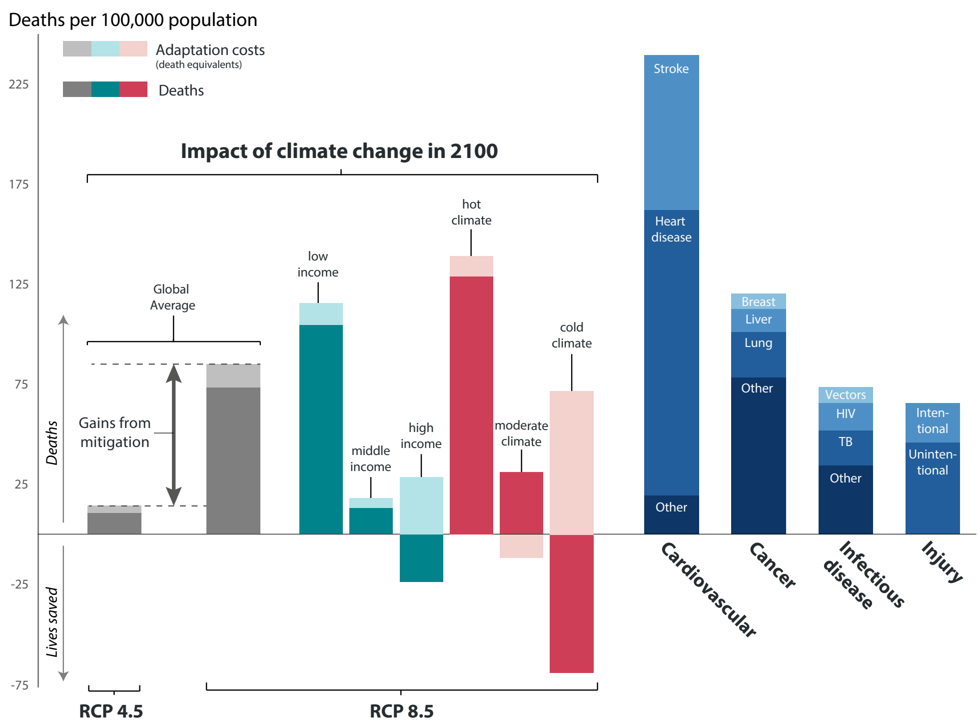

Figure 10: The impact of climate change in 2100 is comparable to contemporary leading causes of death. Impacts of climate change (grey, teal, and coral) are calculated for the year 2100 for socioeconomic scenario SSP3 and include changes in death rates (solid colors) and changes in adaptation costs, measured in death equivalents (light shading). Global averages for RCP 8.5 and RCP 4.5 are shown in grey, demonstrating the gains from mitigation. Income and average climate groups under RCP8.5 are separated by tercile of the 2015 global distribution across all 24,378 impact regions. Blue bars on the right indicate average mortality rates globally in 2018, with values from WHO (2018). Figure F8 provides an RCP4.5 version of this figure.

Finally, we use these projections to develop the first empirically grounded estimates of the mortality partial SCC. Using a 2% discount rate and age-varying VSL, the 2020 mortality partial SCC is roughly

$36.6 (in 2019 USD) with a high emissions scenario and $17.1 with a moderate one. There is substantial uncertainty around these estimates, arising both from climate sensitivity and damage function uncertainty. For example, the interquartile ranges of the mortality partial SCC are [-$7.8, $73.0] and [-$24.7, $53.6], under high and moderate emissions scenarios, respectively. These ranges underscore the nature of the climate challenge and highlight that, while beyond the scope of this paper, valuing this uncertainty under risk aversion will raise the mortality partial SCC.

Overall, the paper's findings suggest that previous research has significantly understated climate change damages due to mortality. For example, the mortality damages we estimate in 2100 account for 49% to 135% of *total* damages across all sectors of the economy according to leading IAMs. Moreover, the mortality partial SCC reported here, under comparable valuation assumptions, is more than 10 times larger than the total health impacts embedded in the FUND IAM (Diaz, 2014).[59]

We believe that this paper has highlighted a key role for systematic empirical analysis in providing a clearer picture of the magnitude of the costs of climate change and how, why, and where they are likely to emerge in the future. Advances in access to data and computing have removed the need to rely so heavily on assumptions when estimating the economic costs of climate change. Looking ahead, the paper's general approach can be applied to other aspects of the global economy besides mortality risk, and doing so is a promising area for future research.

[59] Diaz (2014) computes comparable partial SCC values for FUND ($\delta = 3\%$, "business as usual" emissions) and reports values for three comparable health impacts (diarrhea, vector borne diseases, and cardiopulmonary) that total less than $2 (2019 USD).

References

Alesina, Alberto and Dani Rodrik. 1994. "Distributive politics and economic growth." *The quarterly journal of economics* 109 (2):465–490.

Anthoff, David and Richard SJ Tol. 2014. "The climate framework for uncertainty, negotiation and distribution (FUND): Technical description, version 3.6." *FUND Doc* .

Ashenfelter, Orley and Michael Greenstone. 2004. "Using mandated speed limits to measure the value of a statistical life." *Journal of Political Economy* 112 (S1):S226–S267.

Auffhammer, Maximilian. 2018. "Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data." *NBER Working paper* .

Auffhammer, Maximilian and Anin Aroonruengsawat. 2011. "Simulating the impacts of climate change, prices and population on California's residential electricity consumption." *Climatic Change* 109 (1):191–210. URL http://dx.doi.org/10.1007/s10584-011-0299-y.

Bailey, Martha J and Andrew Goodman-Bacon. 2015. "The War on Poverty's experiment in public medicine: Community health centers and the mortality of older Americans." *American Economic Review* 105 (3):1067–1104.

Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S Shapiro. 2015. "Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004." *The American Economic Review* 105 (5):247–251.

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):105–159. URL http://dx.doi.org/10.1086/684582.

Barreca, Alan I. 2012. "Climate change, humidity, and mortality in the United States." *Journal of Environmental Economics and Management* 63 (1):19–34.

Becker, Gary S. 2007. "Health as human capital: synthesis and extensions." *Oxford Economic Papers* 59 (3):379–410.

Board of Governors of the US Federal Reserve System. 2020. "10-Year Treasury Inflation-Indexed Security, Constant Maturity [DFII10]." Tech. rep., FRED, Federal Reserve Bank of St. Louis. URL https://fred.stlouisfed.org/series/DFII10.

Bright, E. A., P. R. Coleman, A. N. Rose, and M. L. Urban. 2012. "LandScan 2011." Digital dataset: web.ornl.gov/sci/landscan/index.shtml.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2017. "The unequal effects of weather and climate change: Evidence from mortality in India." *NBER Working paper* .

Burke, Marshall, John Dykema, David B Lobell, Edward Miguel, and Shanker Satyanath. 2015. "Incorporating climate uncertainty into estimates of climate change impacts." *Review of Economics and Statistics* 97 (2):461–471.

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

Carleton, Tamma A and Solomon M Hsiang. 2016. "Social and economic impacts of climate." *Science* 353 (6304):aad9837.

Center for Systemic Peace. 2020. "Polity5: Political Regime Characteristics and Transitions 1800-2018." URL http://www.systemicpeace.org/inscrdata.html.

Davis, Lucas W and Paul J Gertler. 2015. "Contribution of air conditioning adoption to future energy use under global warming." *Proceedings of the National Academy of Sciences* 112 (19):5962–5967.

Dellink, Rob, Jean Chateau, Elisa Lanzi, and Bertrand Magné. 2015. "Long-term economic growth projections in the Shared Socioeconomic Pathways." *Global Environmental Change* .

Deryugina, Tatyana and Solomon Hsiang. 2017. "The Marginal Product of Climate." *NBER Working paper* .

Deschenes, Olivier. 2018. "Temperature Variability and Mortality: Evidence from 16 Asian Countries." *Asian Development Review* 35 (2):1–30.

Deschênes, Olivier and Michael Greenstone. 2007. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather." *The American Economic Review* 97 (1):354–385.

———. 2011. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US." *American Economic Journal: Applied Economics* 3 (October):152–185. URL http://www.nber.org/papers/w13178.

Deschênes, Olivier, Michael Greenstone, and Joseph S Shapiro. 2017. "Defensive investments and the demand for air quality: Evidence from the NOx budget program." *American Economic Review* 107 (10):2958–89.

Deschênes, Olivier and Enrico Moretti. 2009. "Extreme weather events, mortality, and migration." *The Review of Economics and Statistics* 91 (4):659–681.

Diaz, Delavane. 2014. "Evaluating the Key Drivers of the US Government's Social Cost of Carbon: A Model Diagnostic and Inter-Comparison Study of Climate Impacts in DICE, FUND, and PAGE." *Working paper* .

Eurostat. 2013. *Europe in Figures: Eurostat Yearbook 2013*. Publications Office of the European Union.

Fetzer, Thiemo. 2014. "Can workfare programs moderate violence? Evidence from India." .

Friedman, Jerome, Trevor Hastie, Robert Tibshirani et al. 2001. *The elements of statistical learning*, vol. 1. Springer series in statistics New York.

Gasparrini, Antonio, Yuming Guo, Masahiro Hashizume, Eric Lavigne, Antonella Zanobetti, Joel Schwartz, Aurelio Tobias, Shilu Tong, Joacim Rocklöv, Bertil Forsberg et al. 2015. "Mortality risk attributable to high and low ambient temperature: A multicountry observational study." *The Lancet* 386 (9991):369–375.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez De Silanes, and Andrei Shleifer. 2014. "Growth in regions." *Journal of Economic Growth* 19 (3):259–309. URL https://ideas.repec.org/a/kap/jecgro/v19y2014i3p259-309.html.

Glaeser, Edward L, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. "Do institutions cause growth?" *Journal of economic Growth* 9 (3):271–303.

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

Guo, Christopher and Christopher Costello. 2013. "The value of adaption: Climate change and timberland management." *Journal of Environmental Economics and Management* 65 (3):452–468.

Guo, Yuming, Antonio Gasparrini, Ben Armstrong, Shanshan Li, Benjawan Tawatsupa, Aurelio Tobias, Eric Lavigne, Micheline de Sousa Zanotti Stagliorio Coelho, Michela Leone, Xiaochuan Pan et al. 2014. "Global variation in the effects of ambient temperature on mortality: a systematic evaluation." *Epidemiology (Cambridge, Mass.)* 25 (6):781.

Heutel, Garth, Nolan H Miller, and David Molitor. 2017. "Adaptation and the Mortality Effects of Temperature Across US Climate Regions." *National Bureau of Economic Research* .

Houser, Trevor, Solomon Hsiang, Robert Kopp, Kate Larsen, Michael Delgado, Amir Jina, Michael Mastrandrea, Shashank Mohan, Roger Muir-Wood, DJ Rasmussen, James Rising, and Paul Wilson. 2015. *Economic risks of climate change: An American prospectus*. Columbia University Press.

Hsiang, Solomon. 2013. "Visually-weighted regression." *SSRN Working Paper* .

———. 2016. "Climate econometrics." *Annual Review of Resource Economics* 8:43–75.

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer et al. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345):1362–1369.

Hsiang, Solomon M and Amir S Jina. 2014. "The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones." Tech. rep., National Bureau of Economic Research.

Hsiang, Solomon M, Kyle C Meng, and Mark A Cane. 2011. "Civil conflicts are associated with the global climate." *Nature* 476 (7361):438.

Hsiang, Solomon M and Daiju Narita. 2012. "Adaptation to cyclone risk: Evidence from the global cross-section." *Climate Change Economics* 3 (02):1250011.

IIASA Energy Program. 2016. "SSP Database, Version 1.1 [Data set]." Tech. rep., National Bureau of Economic Research. URL https://tntcat.iiasa.ac.at/SspDb.Accessed25December,2016.

Interagency Working Group on Social Cost of Carbon. 2010. "Social Cost of Carbon for Regulatory Impact Analysis - Under Executive Order 12866." Tech. rep., United States Government.

Interagency Working Group on Social Cost of Greenhouse Gases. 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis." Tech. rep., United States Government.

Isen, Adam, Maya Rossin-Slater, and Reed Walker. 2017. "Relationship between season of birth, temperature exposure, and later life wellbeing." *Proceedings of the National Academy of Sciences* 114 (51):13447–13452.

Jones, Charles I and Peter J Klenow. 2016. "Beyond GDP? Welfare across countries and time." *American Economic Review* 106 (9):2426–57.

Kahn, Matthew E. 2005. "The death toll from natural disasters: the role of income, geography, and institutions." *Review of Economics and Statistics* 87 (2):271–284.

Kelly, David L, Charles D Kolstad, and Glenn T Mitchell. 2005. "Adjustment costs from environmental change." *Journal of Environmental Economics and Management* 50 (3):468–495.

Kopp, Robert E and Bryan K Mignone. 2012. "The US government's social cost of carbon estimates after their first two years: Pathways for improvement." *Working paper* .

La Porta, Rafael and Andrei Shleifer. 2014. "Informality and development." *Journal of Economic Perspectives* 28 (3):109–26.

Lemoine, Derek. 2018. "Estimating the consequences of climate change from variation in weather." Tech. rep., National Bureau of Economic Research.

Lenssen, N., G. Schmidt, J. Hansen, M. Menne, A. Persin, R. Ruedy, and D. Zyss. 2019. "Improvements in the GISTEMP uncertainty model." *J. Geophys. Res. Atmos.* 124 (12):6307–6326.

Li, Dawei, Jiacan Yuan, and Robert Bob Kopp. 2020. "Escalating global exposure to compound heat-humidity extremes with warming." *Environmental Research Letters* .

Matsuura, Kenji and Cort Willmott. 2007. "Terrestrial Air Temperature and Precipitation: 1900-2006 Gridded Monthly Time Series, Version 1.01." *University of Delaware*. URL http://climate.geog.udel.edu/climate.

Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw. 1994. "The impact of global warming on agriculture: A Ricardian analysis." *The American Economic Review* :753–771.

Millar, Richard J, Zebedee R Nicholls, Pierre Friedlingstein, and Myles R Allen. 2017. "A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions." *Atmospheric Chemistry and Physics* 17 (11):7213–7228.

Moore, Frances C and David B Lobell. 2014. "Adaptation potential of European agriculture in response to climate change." *Nature Climate Change* 4 (7):610.

Murphy, Kevin M and Robert H Topel. 2006. "The value of health and longevity." *Journal of Political Economy* 114 (5):871–904.

National Academies of Sciences, Engineering, and Medicine. 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.* Washington, DC: The National Academies Press. URL https://www.nap.edu/catalog/24651/valuing-climate-damages-updating-estimation-of-the-social-cost-of.

Nordhaus, William D. 1992. "An optimal transition path for controlling greenhouse gases." *Science* 258 (5086):1315–1319.

Nordhaus, William D and Zili Yang. 1996. "A regional dynamic general-equilibrium model of alternative climate-change strategies." *The American Economic Review* :741–765.

Organization of Economic Cooperaton and Development. 2020. "OECD.Stat." URL https://doi.org/10.1787/data-00285-en.

Rasmussen, D. J., Malte Meinshausen, and Robert E. Kopp. 2016. "Probability-weighted ensembles of U.S. county-level climate projections for climate risk analysis." *Journal of Applied Meteorology and Climatology* 55 (10):2301–2322. URL http://journals.ametsoc.org/doi/abs/10.1175/JAMC-D-15-0302.1.

Riahi, Keywan, Detlef P Van Vuuren, Elmar Kriegler, Jae Edmonds, Brian C O'neill, Shinichiro Fujimori, Nico Bauer, Katherine Calvin, Rob Dellink, Oliver Fricko et al. 2017. "The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: an overview." *Global Environmental Change* 42:153–168.

Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni. 2018. "Country-level social cost of carbon." *Nature Climate Change* 8 (10):895.

Rohde, Robert, Richard Muller, Robert Jacobsen, Saul Perlmutter, Arthur Rosenfeld, Jonathan Wurtele, J Curry, Charlotte Wickham, and S Mosher. 2013. "Berkeley Earth temperature averaging process." *Geoinfor Geostat: An Overview* 1 (2):1–13.

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

Sheffield, Justin, Gopi Goteti, and Eric F Wood. 2006. "Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling." *Journal of Climate* 19 (13):3088–3111.

Sherwood, Steven and Qiang Fu. 2014. "A drier future?" *Science* 343 (6172):737–739.

Sherwood, Steven C and Matthew Huber. 2010. "An adaptability limit to climate change due to heat stress." *Proceedings of the National Academy of Sciences* 107 (21):9552–9555.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge. 2015. "What are we weighting for?" *Journal of Human Resources* 50 (2):301–316.

Stern, Nicholas. 2006. *The Economics of Climate Change: The Stern Review.* Cambridge University Press.

Taylor, Karl E, Ronald J Stouffer, and Gerald A Meehl. 2012. “An overview of CMIP5 and the experiment design.” *Bulletin of the American Meteorological Society* 93 (4):485.

Tebaldi, Claudia and Reto Knutti. 2007. “The use of the multi-model ensemble in probabilistic climate projections.” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 365 (1857):2053–2075. URL http://rsta.royalsocietypublishing.org/content/365/1857/2053.

Thomson, Allison M., Katherine V. Calvin, Steven J. Smith, G. Page Kyle, April Volke, Pralit Patel, Sabrina Delgado-Arias, Ben Bond-Lamberty, Marshall A. Wise, Leon E. Clarke, and James A. Edmonds. 2011. “RCP4.5: a pathway for stabilization of radiative forcing by 2100.” *Climatic Change* 109 (1):77. URL https://doi.org/10.1007/s10584-011-0151-4.

Thrasher, Bridget, Edwin P Maurer, C McKellar, and PB Duffy. 2012. “Technical note: Bias correcting climate model simulated daily temperature extremes with quantile mapping.” *Hydrology and Earth System Sciences* 16 (9):3309–3314.

Tol, Richard S.J. 1997. “On the optimal control of carbon dioxide emissions: an application of FUND.” *Environmental Modeling & Assessment* 2 (3):151–163. URL http://dx.doi.org/10.1023/A:1019017529030.

U.S. Environmental Protection Agency. 2016a. “Recommended Income Elasticity and Income Growth Estimates: Technical Memorandum.” Tech. rep., U.S. Environmental Protection Agency Office of Air and Radiation and Office of Policy.

———. 2016b. “Valuing mortality risk reductions for policy: A meta-analytic approach.” Tech. rep., U.S. Environmental Protection Agency Office of Policy, National Center for Environmental Economics.

Van Vuuren, Detlef P, Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque et al. 2011. “The representative concentration pathways: An overview.” *Climatic Change* 109 (1-2):5.

Viscusi, W Kip. 2015. “The role of publication selection bias in estimates of the value of a statistical life.” *American Journal of Health Economics* .

Viscusi, W Kip and Joseph E Aldy. 2003. “The value of a statistical life: a critical review of market estimates throughout the world.” *Journal of risk and uncertainty* 27 (1):5–76.

WHO. 2018. “Global Health Estimates 2016: Deaths by cause, age, sex, by country and by region, 2000–2016.” Tech. rep., World Health Organization.

World Bank. 2020. “World Development Indicators.” .

World Inequality Lab. 2020. “World Inequality Database.” URL https://wid.world/data/.

Yuan, Jiacan, Michael L Stein, and Robert E Kopp. 2019. “The evolving distribution of relative humidity conditional upon daily maximum temperature in a warming climate.” .

Online Appendix for:

Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits

Tamma Carleton[1,2], Amir Jina[3,2], Michael Delgado[4], Michael Greenstone[3,2], Trevor Houser[4], Solomon Hsiang[5,2], Andrew Hultgren[3], Robert Kopp[6], Kelly E. McCusker[4], Ishan Nath[7], James Rising[8], Ashwin Rode[3], Hee Kwon Seo[9], Arvid Viaene[1]0, Jiacan Yuan[1]1, Alice Tianbo Zhang[12]

[1]University of California, Santa Barbara
[2]NBER
[3]University of Chicago
[4]Rhodium Group
[5]University of California, Berkeley
[6]Rutgers University
[7]Princeton University
[8]University of Maryland
[9]The World Bank
[10]E.CA Economics
[11]Fudan University
[12]Washington and Lee University

Appendix Table of Contents

A Using revealed preference to estimate adaptation costs . . . A1
A.1 Graphical solution to inferring unobserved adaptation cost . . . A1
A.2 Surplus generated from compensatory investments . . . A3
A.3 Implementation details for the empirical estimation of adaptation costs . . . A7
A.4 Alternative specification: Including adaptation in the utility function . . . A9

B Data appendix . . . A12
B.1 Mortality data . . . A12
B.2 Climate data . . . A18
B.3 Socioeconomic data and downscaling methodologies . . . A22
B.4 Scale and scope of existing empirically-based estimates of the mortality risk of climate change . . . A27

C Spatial units for projection: "Impact regions" . . . A29
C.1 Algorithm for construction of impact region boundaries . . . A29

D Econometric estimation: Additional results, robustness, out-of-sample validation . . . A33
D.1 Age-specific heterogeneity of the mortality-temperature response function by average income and average climate . . . A33
D.2 Results and robustness with pooled model . . . A34
D.3 Additional results: Spatial extrapolation of temperature sensitivity . . . A40
D.4 Robustness of estimates of subnational heterogeneity in the mortality-temperature response function to an alternative characterization of long-run average climate . . . A41
D.5 Robustness of the mortality-temperature response function to omission of precipitation . . . A44
D.6 Robustness of the mortality-temperature response function to inclusion of additional sources of heterogeneity . . . A44
D.7 Cross-validation to assess out-of-sample performance . . . A49
D.8 Replication of Burgess et al. (2017) and out-of-sample model validation in India . . . A54

E Implementation of projection of future adaptation and benefits of income growth . . . A58
E.1 Determining the temporal dynamics of income effects . . . A58
E.2 Adaptation constraints imposed in the projection of climate change impacts . . . A59
E.3 Projected benefits of adaptation . . . A62

F Climate change projections: Additional results and robustness . . . A63
F.1 Additional climate change projection results . . . A63
F.2 Robustness: Alternative functional form for the mortality-temperature relationship . . . A71
F.3 Sensitivity analysis: Alternative assumptions on out-of-sample extrapolation of response functions . . . A72
F.4 Sensitivity analysis: Alternative assumptions on the rate of adaptation . . . A73

G Calculation of a mortality partial social cost of carbon . . . A75
G.1 Computing post-2100 damage functions . . . A75
G.2 Set up of the climate module using a simple climate model . . . A76

G.3 Converting temperature scenarios to mortality partial SCC A85

H Sensitivity of the mortality partial social cost of carbon A88

H.1 Methodology for constructing value of life-years lost from value of a statistical life (VSL) A88

H.2 Mortality partial social cost of carbon under alternative valuation approaches and socioeconomic scenarios A89

H.3 Alternative approach to estimating post-2100 damages A95

H.4 Robustness of the mortality partial SCC to an alternative functional form of the damage function . . A96

A Using revealed preference to estimate adaptation costs

A.1 Graphical solution to inferring unobserved adaptation cost

In Section 2, we lay out a framework for recovering the costs of adapting to climate change that is micro-founded by a standard utility maximization problem. Figure A1 depicts this optimal adaptation problem faced by individuals and illustrates how we overcome two key empirical challenges to measuring adaptation costs: (1) the universe of adaptation adjustments and their costs are not directly observable and (2) adaptive adjustments are continuous for continuous changes in climate. The problem must be displayed in three dimensions because it involves at least three orthogonal subspaces: climate ($\boldsymbol{C}$), adaptive adjustments to climate ($\boldsymbol{b}$), and an outcome (expressed in dollars of WTP). For illustrative simplicity, here we assume income is held fixed, and we consider a simplified example with univariate climate and univariate adaptation. Further, for this example, higher $\boldsymbol{C} = C$ indicates higher temperatures and higher $\boldsymbol{b} = b$ indicates greater adaptation (i.e., greater protection) from high temperatures, where these terms are unbolded to indicate that they are scalars.

In the lower left panel of Figure A1, the green surface illustrates adaptation costs $A(b)$ which are not directly observable to the econometrician. The height of this surface represents the costs that households would bear to obtain a level of adaptation b. Because we assume markets for adaptive technologies are competitive, $A(b)$ could represent[60] the lower envelope of all firm cost-functions (offer curves) that would supply b, as illustrated by the projection of the surface onto the $A \times b$ plane. Because adaptation costs are a function of technology, they do not depend on the climate and so $\partial A/\partial C = 0$ everywhere, i.e., individuals in Seattle can purchase the same adaptation technology (e.g., air conditioners) as individuals in Houston.

In the lower right panel of Figure A1, the red surface illustrates the expected benefits an individual would accrue for inhabiting some climate C and selecting adaptation b. The height of this surface is a total WTP for adaptation, conditional on the climate: it is equal to the VSL times the expected survival probability $1-\tilde{f}(b, C)$ at each position (b, C). For notational simplicity, we refer to this WTP surface as V. At low levels of adaptation, V declines rapidly with higher temperature C because survival probability declines quickly. At higher levels of adaptation, V declines more gradually with C because adaptation protects individuals against temperature. The solid black lines follow this WTP surface at fixed temperatures, showing how an individual in a given climate would benefit from additional adaptation (bid curves).

Agents at each climate endogenously adapt by selecting the optimal level of b such that the marginal costs equal the marginal benefits. This can be seen on the lower left panel at climates C_{t_0} and C_t, where slices of the benefits surface V are drawn overlaid in red and are tangent to $A(b)$ at the blue circles. Corresponding slices of the adaptation cost surface A are overlaid in green on the benefits surface in the lower right panel. The blue line traces out the equilibrium at different climates. For each climate C there is an optimal level of adaptation $b^*(C)$ endogenously chosen, illustrated by the projection of the equilibrium downward onto the $C \times b$ plane in both panels. The projection of the equilibrium onto the $A \times C$ plane on the left panel illustrates how adaptation expenditures rise with temperature, and the projection onto the $V \times C$ plane

[60] In Appendix A.4 below, A are net costs since they are net any utility benefits or costs of $\boldsymbol{b}$.

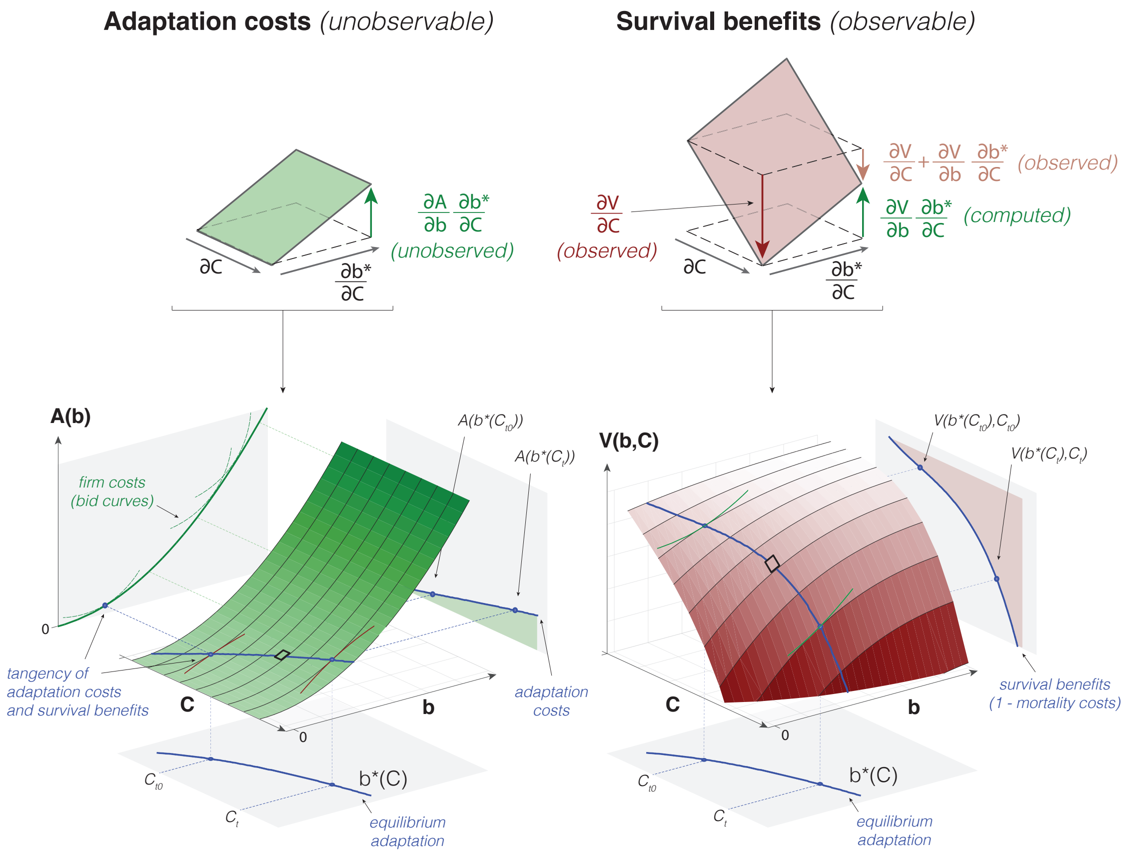

Figure A1: Use of revealed preference to recover WTP for an unobservable adaptation. Horizontal dimensions are climate C, representing temperature, and adaptation level b. Vertical dimensions are adaptation costs $A(b)$ in the left panel and expected survival benefits $V(b,C) = VSL[1-\tilde{f}(b,C)]$ in the right panel, both in units of dollars of WTP. Tangency planes at the top depict infinitesimal surfaces spanning $\partial C \times \frac{\partial b^*}{\partial C}$ at a point along the equilibrium adaptation path $b^*(C)$, which is drawn in blue. Adaptation costs, as a function of the climate, are the height of the green wedge on the $A \times C$ plane in the lower left panel. The value of mortality risk imposed by the climate is the red wedge on the $V \times C$ plane in the lower right panel.

on the right panel illustrates how expected survival benefits decline with temperature, or equivalently, how mortality costs rise with temperature. The sum of changes to these adaptation expenditures and the value of mortality costs is the full cost of changes to the climate.

A key innovation to our analysis is fully accounting for adaptation costs $A(b)$ even though neither $A(.)$ nor b is observed. Indeed, there may be a very large, even infinite, number of ways that populations adapt to climate that cannot be feasibly enumerated by the econometrician. All the econometrician can observe are the effects of adaptation on survival probability $1-\tilde{f}$. If a climate were gradually warmed from C_{t_0} to C_t, individuals would continuously respond by adapting along $b^*(C)$ and traveling up the cost surface in the lower left panel, eventually incurring costs $A(b^*(C_t))$ rather than the initial costs $A(b^*(C_{t_0}))$ that they incurred prior to warming. We point out that the change in this total adaptation cost $A(b^*(C_t)) - A(b^*(C_{t_0}))$ can be inferred based only on the shape of the benefits surface along the equilibrium, information that is recoverable by the econometrician.

To show this, at the top of Figure A1 we draw tangency planes for both the costs and benefits surfaces for a single location along the equilibrium adaptation locus between C_{t_0} and C_t, indicated by black squares on the two surfaces in the lower left and lower right panels. Both tangency planes span an area $\partial C \times \frac{\partial b^*}{\partial C}$, indicating how much additional adaptation populations undertake ($\frac{\partial b^*}{\partial C}$) for an exogenous change in climate (∂C), changes that would cause them to traverse each of these planes from their respective left-most corner to their right-most corner. The corresponding change in survival benefits is $\frac{dV}{dC} = \frac{\partial V}{\partial C} + \frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$ (downward pink arrow on the right), which the econometrician can observe by computing the change in survival probability due to climate between two adjacent locations after allowing them both to fully adapt to their respective climates. If the cooler location is heated by ∂C but not permitted to adapt, its survival benefits change by $\frac{\partial V}{\partial C}$ (downward red arrow), a counterfactual outcome that the econometrician can compute by simulating a warmer environment without allowing for adaptation. The difference between these two changes is equal to the benefits of marginal adaptations $\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$ (upward green arrow, right panel). Along the equilibrium $b^*(C)$, these marginal benefits of adaptation must equal their marginal costs, thus we know the corresponding increase in unobserved adaptation costs $\frac{\partial A}{\partial b}\frac{\partial b^*}{\partial C}$ (upward green arrow, left panel) must be equal in magnitude to $\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$. By continuously computing and differencing the total and partial derivatives of V with respect to an incremental change in climate dC (i.e., $\frac{dV}{dC} - \frac{\partial V}{\partial C}$), we recover the marginal benefits of unobserved incremental adaptations ($\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$), which we know must also equal their marginal costs ($\frac{\partial A}{\partial b}\frac{\partial b^*}{\partial C}$). Then by integrating these marginal costs with respect to the climate (shown in the $A \times C$ plane of the lower left panel) we can compute the total change in adaptation costs $A(b^*(C_t)) - A(b^*(C_{t_0}))$ for the non-marginal change in climate from C_{t_0} to C_t. This intuition holds for an unknown number of margins of adaptation and a climate of arbitrary dimension, which we allow for in the main text and in derivations below.

A.2 Surplus generated from compensatory investments

As discussed in the main text, the equivalence of *marginal* adaptation benefits and *marginal* adaptation costs at each point along the equilibrium pathway $\boldsymbol{b}^*(Y, \boldsymbol{C})$ (Equation 7) does not imply that our estimates of *total* adaptation costs are equivalent to *total* adaptation benefits for any given population at fixed climate $\boldsymbol{C}$. In general, we expect total adaptation benefits to exceed total adaptation costs, generating surplus from compensatory investments. Here, we define this surplus and illustrate why it is not zero. Empirically, we find that this surplus is substantial (see Section 5.5).

We define adaptation surplus as the total benefits of adapting to climate change (i.e., the dollar value of the difference between mortality effects of climate change with and without the benefits of adaptation) minus the total cost of adaptation (i.e., the integral of marginal adaptation costs along the climate change trajectory, as shown in Equation 9). This surplus can be evaluated at any future climate $\boldsymbol{C}_t$. That is, adaptation surplus under a climate changing from time period t_0 to t can be written as:[61]

$$Adaptation\ surplus\ (\boldsymbol{C}_{t_0} \rightarrow \boldsymbol{C}_t) = \underbrace{-VSL[\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t) - \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_t)]}_{\text{total adaptation benefits}} - \underbrace{[A(\boldsymbol{b}^*(\boldsymbol{C}_t)) - A(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}))]}_{\text{total adaptation costs}}$$

$$= -\int_{\boldsymbol{b}^*(\boldsymbol{C}_{t_0})}^{\boldsymbol{b}^*(\boldsymbol{C}_t)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_t)}{d\boldsymbol{b}} d\boldsymbol{b} - \int_{\boldsymbol{b}^*(\boldsymbol{C}_{t_0})}^{\boldsymbol{b}^*(\boldsymbol{C}_t)} \frac{\partial A(\boldsymbol{b})}{\partial \boldsymbol{b}} d\boldsymbol{b}^* \tag{A.1}$$

where both integrals represent line integrals, and where $d\boldsymbol{b}^*$ indicates that the line integral is calculated along the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$.

The first term in the definition of adaptation surplus in Equation A.1 is the total benefits of adaptation, defined as [minus] the mortality effects of climate $\boldsymbol{C}_t$ with optimal adaptation (i.e. $\boldsymbol{b}^*(\boldsymbol{C}_t)$) minus the mortality effects of that same climate, but with adaptation fixed at its initial level (i.e., $\boldsymbol{b}^*(\boldsymbol{C}_{t_0})$). The second term is the total costs of adaptation, defined as the adaptation costs under optimal adaptation in climate $\boldsymbol{C}_t$ minus adaptation costs under optimal adaptation in the initial climate $\boldsymbol{C}_{t_0}$. Adaptation benefits (the first term) can be computed by integrating $\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_t)}{d\boldsymbol{b}}$, the marginal mortality effect of adaptation evaluated at fixed climate $\boldsymbol{C}_t$. Note that this integration is not computed over the optimal pathway, as the climate is fixed at $\boldsymbol{C}_t$ and any $\boldsymbol{b} \neq \boldsymbol{b}^*(\boldsymbol{C}_t)$ is thus off-equilibrium. Adaptation costs (the second term) can be computed by integrating marginal adaptation costs of $\boldsymbol{b}$ along the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$.

The expression for adaptation surplus in Equation A.1 is represented as the difference between two integrals, each computed over the unobserved choice vector $\boldsymbol{b}$. To empirically identify adaptation surplus, we aim to rewrite this expression as a difference between integrals which are computed over the multi-dimensional climate $\boldsymbol{C}$, which changes over time. This is an important step, as changes in the climate $\boldsymbol{C}$ are empirically identifiable, while adjustments to $\boldsymbol{b}$ are unobserved by the econometrician. As shown below (as well as in Section 2 in the main text), total adaptation costs, the second term in Equation A.1, can be rewritten as an integral over time using a simple change of variables. However, rewriting total adaptation benefits, the first term in Equation A.1, as an integral over time (and hence, climate $\boldsymbol{C}$) requires multiple steps, which we outline below.

To see how we construct an empirically tractable expression for total adaptation benefits (first term in Equation A.1), we first consider a visual illustration. Figure A2 shows the construction of total adaptation benefits using the same notation and format as the lower right panel of Figure A1. As in Figure A1, the red surface represents how expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b(C), C)]$ depend on both climate C and adaptation b, in the case where both climate and adaptation are univariate. The basic idea is that we want to quantify the vertical difference between points s and r (i.e., $s - r$), which can be computed empirically as the vertical difference $q - r$ minus the difference $q - s$. To see why, note that the total benefits of adaptation incurred under a climate change from C_{t_0} to C_t are represented by the vertical difference

[61] Note that income only influences the calculation of surplus arising from climate-driven adaptation via changes in the VSL. Therefore, we abstract away from income changes throughout this section, including omitting Y as an argument of $\boldsymbol{b}^*$, for simplicity of exposition.

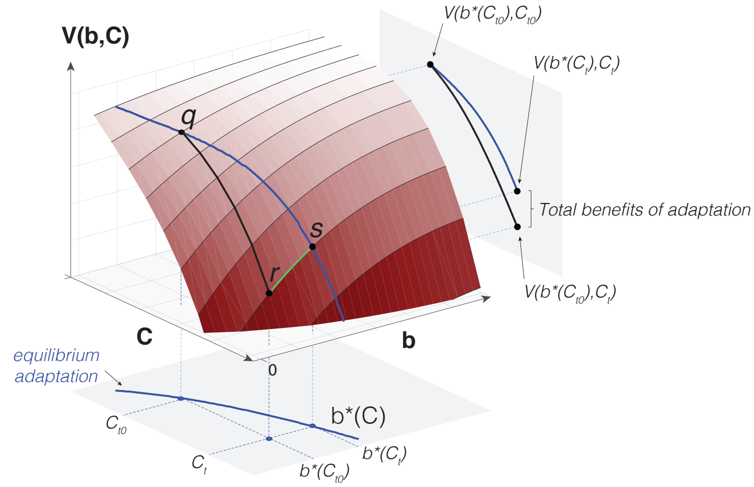

Figure A2: Recovering total benefits of adaptation using revealed preference. Horizontal dimensions are climate C, representing temperature, and adaptation level b. The vertical dimension is expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b, C)]$, in units of dollars of WTP. The equilibrium adaptation path $\{b^*(C), C\}$ is drawn in blue (line $q \rightarrow s$), and the off-equilibrium path $\{b^*(C_{t_0}), C\}$ is drawn in black (line $q \rightarrow r$). To derive the total benefits of adaptation under a change in climate from C_{t_0} to C_t, we integrate the surface along the green line (line $r \rightarrow s$), evaluating changes in survival benefits at a fixed climate C_t, as adaptation evolves from $b^*(C_{t_0})$ to $b^*(C_t)$. The magnitude of total adaptation benefits is shown on the $V \times C$ plane on the right panel.

between points s and r (shown on the $V \times C$ plane on the right panel), because this height measures the total mortality benefits realized from optimally investing in adaptation $b^*(C_t)$ when experiencing climate C_t, instead of holding adaptation fixed at its initial level $b^*(C_{t_0})$. This difference can be computed in two ways. First, total benefits of adaptation can be computed by traversing along the off-equilibrium green line between points r and s; that is, by holding C fixed at C_2 and integrating $V(b, C)$ over b from $b^*(C_{t_0})$ to $b^*(C_t)$. This integration along the green line represents the definition of total adaptation benefits written in Equation A.1. However, this same vertical distance can alternatively be calculated by traversing along the off-equilibrium black line between points q and r (i.e., holding b fixed at $b^*(C_{t_0})$ and integrating $V(b, C)$ over C from C_{t_0} to C_t), and then subtracting off the value of the survival impacts of the optimal pathway from C_{t_0} to C_t (i.e., the height of the surface at point q minus point s). This integration over C (twice) is empirically identifiable, as changes in climate can, in principle, be observed.

Now, consider the construction of total adaptation benefits in an arbitrary multi-dimensional $\boldsymbol{b} \times \boldsymbol{C}$ space. We first note that the Gradient Theorem implies path independence of line integrals on smooth functions; thus, for a continuous and differentiable surface $VSL[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, the integral along any path on this surface depends only on the endpoints of that path. Equation A.1 writes total adaptation benefits using a path along the surface in the $\boldsymbol{b}$ dimension between the end points $\{\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t\}$ and $\{\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_t\}$.[62] However, as discussed above, we cannot compute traversing of this path, as changes in $\boldsymbol{b}$ are unobservable. Thus, we need to define an alternative computable path between the same endpoints. If we can construct a loop on

[62]Note that while Figure A2 illustrates total adaptation benefits using the expected survival *benefits* surface $VSL[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, the definition can be equivalently written using [minus] the expected mortality *costs* surface, $-VSL[\tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, as in Equation A.1. For parsimony, we use the latter notation here and in the subsequent expressions.

the surface that connects the two endpoints, the sum of the desired segment and the remaining segments defining that loop must equal zero, because the line integral over any closed loop L must, by construction, equal zero. We can then rearrange this identity to isolate the computable segments of the loop, allowing us to back out the unobserved segment defining the total benefits of adaptation.

We define such a loop that begins at $\{\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_{t_0}\}$ (analogous to point q in Figure A2) and traverses along the off-equilibrium path from $\boldsymbol{C}_{t_0}$ to $\boldsymbol{C}_t$ with adaptation fixed at $\boldsymbol{b}^*(\boldsymbol{C}_{t_0})$ (analogous to the black line between q and r in Figure A2). In the second segment, it traverses in the $\boldsymbol{b}$ dimension, holding $\boldsymbol{C}$ fixed at $\boldsymbol{C}_t$, to arrive at $\{\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t\}$ (analogous to the green line in Figure A2 and equal to the total benefits of adaptation). Finally, our path arrives back at its starting point by integrating along the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$ (analogous to the blue line between q and s in Figure A2):

$$\begin{aligned}\oint_L \nabla[VSL\tilde{f}(\boldsymbol{b}, \boldsymbol{C})] \cdot \partial\boldsymbol{b}\partial\boldsymbol{C} &= \int_{t_0}^{t} VSL \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds + \int_{\boldsymbol{b}^*(C_{t_0})}^{\boldsymbol{b}^*(C_t)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_t)}{d\boldsymbol{b}} d\boldsymbol{b} \\ &\quad + \int_{t}^{t_0} VSL \frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds \\ &= 0 \end{aligned} \tag{A.2}$$

By rearranging Equation A.2 (including changing the direction of integration for the third segment), we can use this closed loop, which is composed of two computable segments and a third that is unobservable, to calculate the total benefits of adaptation:

$$\begin{aligned}\textit{Total adaptation benefits} &= -\int_{\boldsymbol{b}^*(C_{t_0})}^{\boldsymbol{b}^*(C_t)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_t)}{d\boldsymbol{b}} d\boldsymbol{b} \\ &= -\int_{t_0}^{t} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \end{aligned} \tag{A.3}$$

Using Equation A.3 and a change of variables to rewrite the total costs of adaptation as an integral over $\boldsymbol{C}$, we can rewrite Equation A.1 as:

$$\begin{aligned}\textit{Adaptation surplus } (\boldsymbol{C}_{t_0} \to \boldsymbol{C}_t) &= -\int_{t_0}^{t} VSL \left[\underbrace{\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}}}_{\substack{\text{mortality risk} \\ \text{w/ adaptation}}} - \underbrace{\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}}}_{\substack{\text{mortality risk} \\ \text{w/o adaptation}}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \\ &\quad - \int_{t_0}^{t} \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_s))}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds \end{aligned} \tag{A.4}$$

While the total adaptation benefits term in Equation A.4 (the first term) is composed of values that are, in principle, empirically identifiable, the adaptation cost expression (the second term) remains unobservable because the net cost function $A(\boldsymbol{b}^*(\boldsymbol{C}))$ is unknown. Thus, we take a final step to rewrite the entire adaptation surplus expression in Equation A.4 in terms of objects that are measurable, using Equation 9 from the main text to substitute for the object $\int_{t_0}^{t} \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_s))}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds$:

$$\begin{aligned} \textit{Adaptation surplus } (\boldsymbol{C}_{t_0} \to \boldsymbol{C}_t) = & -\int_{t_0}^{t} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \\ & + \int_{t_0}^{t} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \\ & = \int_{t_0}^{t} VSL \left[\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \end{aligned} \tag{A.5}$$

In Equation A.5, the first term inside the integral represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, but holding adaptation actions fixed at the levels that were optimal under the original climate, $\boldsymbol{C}_{t_0}$. In contrast, the second term represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, allowing adaptation actions $\boldsymbol{b}^*(\boldsymbol{C})$ to evolve optimally with the changing climate. Note that because the second term is a partial derivative, its integral is *not* the total change in the mortality rate. While the two partial derivatives in Equation A.5 will be identical when $\boldsymbol{C} = \boldsymbol{C}_{t_0}$, if they diverge at some point after $\boldsymbol{C}$ warms beyond $\boldsymbol{C}_{t_0}$, then surplus will be nonzero. Thus, a sufficient condition for positive surplus is:

$$\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} > \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \quad \forall s \in (t_0, t] \tag{A.6}$$

This condition says that mortality risk must rise more with changes in the climate at lower levels of adaptation. If this condition holds, the difference between the two partial derivatives in Equation A.6 is weakly positive, and the total adaptation surplus over the climate trajectory $\boldsymbol{C}_{t_0} \to \boldsymbol{C}_t$ is positive.

A.3 Implementation details for the empirical estimation of adaptation costs

In Section 6.2, we describe how we use econometric estimation of Equation 5 in combination with climate model projections to construct empirical estimates of changes in adaptation costs due to climate change, following the theoretical derivation in Section 2. Here, we provide some additional details on this implementation.

Theoretically, adaptation costs can be computed by taking the difference between the total and partial derivative of expected mortality risk with respect to changes in the climate (Equation 9), and integrating this difference. To empirically construct an estimate of these costs, we begin by taking expectations of Equation 5 over weather realizations $\boldsymbol{T}$, to specify our empirically estimated *expected* mortality risk for an age group a in region r for year t:

$$\hat{\tilde{f}}(.)_{art} \equiv \mathrm{E}[\hat{f}(.)_{art}] = \mathrm{E}[\underbrace{\hat{g}_a\left(\boldsymbol{T}_{rt}, TMEAN_{rt}, \log(GDPpc)_{rt}\right)}_{\hat{g}_{art}(\cdot)}] + ... \tag{A.7}$$

where we omit the various estimated terms orthogonal to temperature, which fall out after differentiation.

Recall that the estimates $\hat{g}_{art}(\cdot)$ describe the shape of the annual response function in region r and year t for age group a, taking as inputs the summary climate parameter $TMEAN$ and log income per capita, where the coefficients used to construct $\hat{g}_{art}(\cdot)$ are recovered from the regression in Equation 5. The expectation of $\hat{g}(\cdot)$ is computed over realizations of temperature for region r in year t from the prior 15 years, with weights of historical observations linearly declining in time. Below we omit subscripts for clarity, but the following calculation is conducted yearly for each age and region separately, for each of our 33 high-resolution climate models.

We differentiate expected mortality risk $\hat{\tilde{f}}(.)$ with respect to a small change in climate $\boldsymbol{C}$, by computing how $\hat{\tilde{f}}(.)$ would change if the distribution of daily temperatures shifted due to a change in climate. The climate directly affects mortality by altering the distribution of daily temperatures to which populations are exposed and indirectly affects mortality risk by altering the shape of the mortality-temperature response function. Importantly, our econometric framework allows us to develop estimates of both the partial derivative, which captures the direct effect only where no adaptation is allowed to take place, and the total derivative, which reflects both direct effects and the changing slope of the response function.

In our econometric framework, the partial derivative of expected mortality risk with respect to the climate is captured through a change in events $\boldsymbol{T}$, the argument of $\mathrm{E}[\hat{g}(\cdot)]$, and conditional on climate $\boldsymbol{C}$ ($TMEAN$) and income Y ($\log(GDPpc)$). The partial effect of the climate on expected mortality risk is then:

$$\frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{C}} = \frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} = \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} \tag{A.8}$$

Here, $\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}}$ is the change in the all nonlinear elements of $\boldsymbol{T}$ that describe the daily temperature distribution, resulting from an incremental change in climate.

In contrast, the total derivative of expected mortality risk with respect to a change in climate reflects endogenous adaptations through adjustments to $\boldsymbol{b}$, which in turn change the shape of the response function. Our econometric framework captures these effects through the $TMEAN$ interactions in $g(\cdot)$, which modify the shape of a region's response function based on long run average conditions. When we compute the total derivative of $\hat{\tilde{f}}(.)$ with respect to the climate, we consider both the partial effect of changes to $\boldsymbol{T}$ and the effect of adaptive adjustments captured by the effect of $TMEAN$. The total effect of the climate on expected mortality risk is:

$$\begin{aligned} \frac{d\hat{\tilde{f}}_t}{d\boldsymbol{C}} = \frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} &= \frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\tilde{f}}_t}{\partial TMEAN} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \\ &= \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \end{aligned} \tag{A.9}$$

where $\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}$ captures the ways in which incremental changes in $TMEAN$ affect the shape of the mortality response function, multiplied by the distribution of daily temperatures, $\boldsymbol{T}$. $\frac{\partial TMEAN}{\partial \boldsymbol{C}}$ is the amount that long-run average temperatures are estimated to change during a period of incremental climatic change.

As shown in Equation 10 in the main text, the *difference* between the total and partial derivatives of expected mortality risk with respect to the climate is thus the difference between Equations A.9 and A.8:

$$\frac{d\hat{\hat{f}}_t}{d\boldsymbol{C}} - \frac{\partial\hat{\hat{f}}_t}{\partial\boldsymbol{C}} = \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \tag{A.10}$$

The righthand side of Equation A.10 is fully computable for years in our projection using a combination of empirically estimated parameters, $\hat{g}(\cdot)$, and climate projections, $\{\boldsymbol{T}, TMEAN\}$. Substituting Equation A.10 into Equation 9 from the main text allows us to estimate non-marginal changes in adaptation costs incurred as the climate of each population changes. In each projection, we solve for adaptation costs as a region's climate evolves from time period t_0 to t:

$$\begin{aligned}
\widehat{A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A}(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) &\approx -\int_{t_0}^{t} VSL_s \left[\frac{d\hat{\hat{f}}_s}{d\boldsymbol{C}} - \frac{\partial\hat{\hat{f}}_s}{\partial\boldsymbol{C}}\right] \frac{d\boldsymbol{C}_s}{ds} ds \\
&\approx -\sum_{\tau=t_0+1}^{t} VSL_\tau \left(\left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_\tau, Y_t}\right)(TMEAN_\tau - TMEAN_{\tau-1}) \\
&\approx -\sum_{\tau=t_0+1}^{t} VSL_\tau\, \hat{\boldsymbol{\gamma}}_{\mathbf{1}} \mathrm{E}[\boldsymbol{T}]_\tau \left(TMEAN_\tau - TMEAN_{\tau-1}\right),
\end{aligned} \tag{A.11}$$

where the second equality results from substitution of Equation A.10 into Equation 9 and from employing a discretized approximation of the continuous integral (we use discrete time-steps of one year). As noted in the main text, recall that we hold income fixed at its endpoint value in the calculation of Equation A.11. This is because the goal of the calculation is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate only. Changes in adaptation expenditures due to rising incomes may change mortality risk under climate change, but these changes are voluntary and are not the consequence of the changing climate, and are therefore not included in our calculation of the total mortality-related costs of climate change. These income effects are accounted for econometrically in the estimation of Equation 5 through the interaction with income and they influence predicted temperature-mortality relationships in all of our calculations, but we do not track the cost of these effects and these costs are intentionally excluded from our calculation of climate-change-induced adaptation spending.

As noted in the main text, we treat the VSL as invariant to changes in the climate, although we allow it to be a function of income, which evolves with time. These adaptation cost estimates are calculated for each impact region, age group, and year, using a baseline period t_0 of 2001 to 2010, for each of our 33 high-resolution climate model projections.

A.4 Alternative specification: Including adaptation in the utility function

Throughout the main text, we construct estimates of adaptation costs derived from a representative agent's problem in which utility is a function only of a consumption good x. In this simple model (see Equation 6), there is no direct utility benefit of adaptation behaviors or investments $\boldsymbol{b}$; instead, the actions represented by this composite good influence the agent's problem only through changing mortality risk. In an alternative specification shown here, we allow agents to derive utility both from consumption of x and also possibly from the choice variables in $\boldsymbol{b}$ (for example, air conditioning might increase utility directly, regardless of its effect on mortality risk). We demonstrate that the implications of this alternative model are purely in the *interpretation* of our empirically derived adaptation cost estimates; the calculation described in Section 6.2 of the main text does not change.

As in Section 2 of the main text, we consider a single representative global agent who faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$ in each period t. We further assume there exists some numeraire good x_t for which utility $u(x_t, \boldsymbol{b}_t)$ is quasilinear. As above, this agent maximizes utility conditional on *expected* weather realizations, subject to an exogenous budget constraint and exogenously determined emissions. Letting $\tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t) = \mathbb{E}_{\boldsymbol{c}_t}[f(\boldsymbol{b}_t, \boldsymbol{c}(\boldsymbol{C}_t)) \mid \boldsymbol{C}_t]$ represent the expected probability of death, the agent solves:

$$\max_{\boldsymbol{b}_t, x_t} \; u(x_t, \boldsymbol{b}_t)\left[1 - \tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t)\right] \quad s.t. \;\; Y_t \geq x_t + A(\boldsymbol{b}_t), \tag{A.12}$$

where $A(\boldsymbol{b}_t)$ is the composite price of all adaptive investments and Y is exogenously determined income. As in the main text, we assume that $\tilde{f}(\cdot)$ is continuous and differentiable, that markets clear for all technologies and investments represented by the composite $\boldsymbol{b}$, as well as for the numeraire good x, and that all choices $\boldsymbol{b}$ and x can be treated as continuous.

Rearranging the agent's first order conditions and using the conventional definition of the VSL,[63] we can write:

$$\underbrace{\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}_t} - \frac{\partial u/\partial \boldsymbol{b}}{\partial u/\partial x}}_{\text{net marginal cost of } \boldsymbol{b}} = \frac{-u(x_t^*, \boldsymbol{b}_t^*)}{\partial u/\partial x[1 - \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = \underbrace{-VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}}}_{\substack{\text{marginal survival} \\ \text{benefit of } \boldsymbol{b}}} \tag{A.13}$$

This expression governs expenditures on adaptation. Its righthand side is the product of the negative of the VSL and the marginal change in expected mortality risk due to a change in adaptation, so it represents the expected marginal benefit (in dollar value) of adjusting $\boldsymbol{b}$ through its effect on mortality risk. This object is identical to its counterpart in Equation 7 in the main text. The lefthand side has two parts. The first term represents the marginal cost of all pecuniary expenditures incurred due to a marginal change in adaptation $\boldsymbol{b}$, such as spending on units of air conditioning. The second term represents [minus] the dollar value of all non-mortality marginal utility benefits or costs derived from a marginal change in $\boldsymbol{b}$, such as the utility of enjoying air conditioning or the disutility of exercising at midnight to avoid daytime heat (note that this object is expressed in dollars of WTP by dividing through by the marginal utility of consumption, $\partial u/\partial x$). Together, these two terms can be interpreted as the *net* marginal cost of all adaptive actions composing the composite $\boldsymbol{b}$, because non-mortality marginal benefits and costs are removed from the marginal pecuniary expenditures term $\partial A/\partial \boldsymbol{b}$.

Both terms composing net marginal costs in Equation A.13 are unobservable. In contrast, the marginal survival benefit can be rewritten as the product of the negative of the VSL and the difference between the total and partial derivatives of mortality risk with respect to the climate – i.e., $\frac{d\tilde{f}}{d\boldsymbol{C}} - \frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$ (see Equation 8).

[63] As described in the main text, the value of a statistical life is defined as the willingness to pay for a marginal increase in the probability of survival (Becker, 2007). Mathematically, this object is utility divided by the product of the probability of survival and the marginal utility of consumption: $VSL = \frac{u(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x}$.

As discussed in the main text, we develop an empirical model that allows us to estimate both the total and partial derivates, rendering the marginal survival benefits empirically tractable.

In the main text, we use this insight to develop an expression for the additional adaptation costs incurred as the climate changes gradually, which is composed of observable terms. This expression remains unchanged under the alternative model specification described here, with the exception that the adaptation costs recovered are *net* of utility benefits or costs incurred due to changes in optimal adaptation $\boldsymbol{b}^*$. Here, the additional net adaptation costs incurred as the climate changes gradually from period t_0 to period t are:

$$\begin{aligned} A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) - \frac{1}{\partial u/\partial x}\left[u(x^*(Y_t, \boldsymbol{C}_t), \boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - u(x^*(Y_t, \boldsymbol{C}_{t_0}), \boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0}))\right] \\ = \int_{t_0}^{t} \left[\frac{\partial A(\boldsymbol{b}_s^*)}{\partial \boldsymbol{b}} - \frac{\partial u(x_s^*, \boldsymbol{b}_s^*)/\partial \boldsymbol{b})}{\partial u(x_s^*, \boldsymbol{b}_s^*)/\partial x)}\right] \frac{d\boldsymbol{b}_s^*}{d\boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds \\ = -\int_{t_0}^{t} VSL_s \left[\frac{d\tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{\partial \boldsymbol{C}}\right] \frac{d\boldsymbol{C}_s}{ds} ds, \end{aligned} \tag{A.14}$$

where the last line relies on substitution from Equations A.13 and 8. The righthand side of Equation A.14 can be approximated empirically as shown in Section 6.2 in the main text. Thus, the only implication of this alternative model specification is that adaptation cost estimates should be interpreted as pecuniary expenditures net of direct utility benefits and costs.

Similarly, the mortality "partial" social cost of carbon shown in the main text, which relies on an estimate of adaptation costs, is unchanged under this alternative model specification. However, as in Equation A.14, the mortality partial SCC should be interpreted here as the marginal willingness to pay to avoid the alteration of mortality risk associated with a marginal increase in greenhouse gas emissions inclusive of the benefits and costs of adaptations undertaken to reduce mortality risk. Indeed, the omission of the direct utility benefits and costs of adaptation behaviors and technologies from the mortality partial SCC is intentional, because they are not a response to mortality-related risks. However, these utility effects are caused by climate change and should be included in a full, all-sector SCC.

B Data appendix

B.1 Mortality data

Our mortality data represent 41 countries. In some cases our data represent the universe of reported deaths in those countries, while in others (e.g., China), data are representative samples, as no vital statistics registry system exists. Combined, our dataset covers mortality outcomes for 55% of the global population. Data are drawn from multiple, often restricted, national and international sources, all mortality datasets contain information on deaths per 100,000 population from all causes at a monthly or annual frequency, and all except India contain age-specific mortality rates. Each of the countries' data are drawn from distinct databases, details of which are provided below. Figure B1 displays the spatial coverage and resolution of all mortality records used, as well as their temporal coverage.

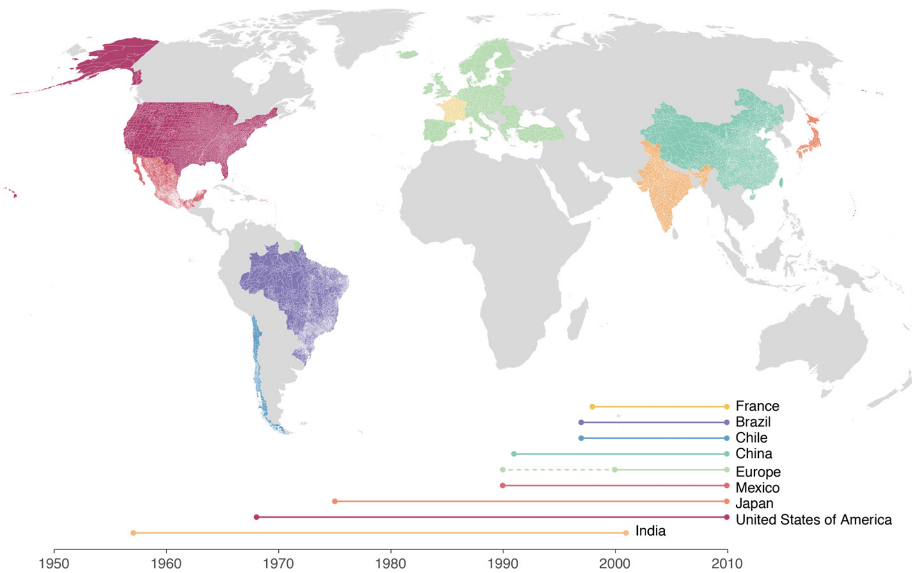

Figure B1: Mortality statistics used to estimate the relationship between mortality, temperature, climate, and income. Figure shows the spatial distribution and resolution of mortality statistics from all countries used to generate regression estimates of the temperature-mortality relationship. Temporal coverage for each country is shown under the map (the dotted line for the European Union (EU) time series indicates that start dates vary for a small subset of countries).

B.1.1 Brazil

Brazilian mortality data at the ADM2-month level were obtained from the Mortality Information System (SIM) of the Ministry of Health in Brazil (Ministry of Health in Brazil, 2019).[64] We use data from 1997-2010 and aggregate the monthly data to annual frequency. Data were provided for both place of death and place of residence. As with all subsequent datasets, we assign weather exposure to deaths in our data at the place of residence, as this is provided for all sources. Data were downloaded in 5-year age groups which were then aggregated to the age groups used in the analysis. ADM2-level populations were obtained from the same source. Administrative boundary files were downloaded from GADM (Global Administrative Areas, 2012).

[64] http://datasus.saude.gov.br/sistemas-e-aplicativos/eventos-v/sim-sistema-de-informacoes-de-mortalidade.

Brazilian death data as downloaded contained a number of ADM2 units with missing values for deaths and no values of zero, implying that these are a mix of true zeros and missing values. To ascertain whether they are more likely to be the former, we examined the relationship between death counts and population in all ADM2 units, and then in only those ADM2 units that ever show a missing value in any year. We found that missing values are more likely to occur in low population ADM2 units, suggesting that these are places that should have recorded zero deaths. We consequently treat these missing values as zeros, but in robustness tests find that treating them as missing does not substantially change any of our results.

B.1.2 Chile

Chilean mortality data at the ADM2 level are obtained from the vital registration system maintained by the Department of Statistics and Information (Departmento de Estadísticas e Información de Salud, DEIS) at the Ministry of Health (Ministry of Health, Chile, 2015).[65] We use data at the ADM2 level for 1997-2012. The vital registration system contains information on individual dates of deaths (often with missing values for days but always containing years) which we aggregate within administrative units to provide the ADM2 total count of deaths in each unit. This also provides data with arbitrarily accurate age grouping, and we aggregate in accordance with the age groups in our analysis. ADM2 population data were downloaded from the National Institute of Statistics (Instituto Nacional de Estadísticas, INE)[66] and merged with the death counts to calculate mortality rates. Administrative boundary files were downloaded from GADM (Global Administrative Areas, 2012).

B.1.3 China

Chinese mortality data are the same as those used in Chen et al. (2013), and were provided by the authors of that paper. The data come from the Chinese Disease Surveillance Points system and are not the universe of mortality as in much of the rest of our sample, but rather a representative sample of the Chinese population benchmarked to the 1990 Chinese census. Locations are given as geographic coordinates relating to the centroid of the surveillance area. Data used in Chen et al. (2013) span from 1991-2000 and cover 145 points to which we assign a climate exposure at the level of the ADM2 unit containing that point. We supplement this with data on a further 161 points from 2004-2012 which were benchmarked to the 2000 census to reflect population changes. This gives us a total of 203 disease surveillance points due to overlap in some points across both periods. Due to the difficulty of establishing consistency between the overlapping points in the two time periods, we include a time-period specific fixed effect in our regressions to allow for unobservable differences in disease and mortality monitoring extent and capacity across time periods. The data record deaths in 5 year age groups, as well as population estimates required to calculate mortality rates. Administrative boundaries for the ADM2 and ADM1 level are obtained from Chen et al. (2013) for the 2000 census boundaries, and points are assigned to an administrative unit based on being contained within those boundaries.

[65] Data are available here: http://www.deis.cl/bases-de-datos-defunciones/.
[66] Data are available here: http://www.ine.cl/estadisticas/demograficas-y-vitales

B.1.4 European Union

The EU maintains a centralized statistical database known as EuroStat (Eurostat, 2013)[67] which contains data on mortality counts and rates for all member countries at EU-specific administrative regions known as "Nomenclature of territorial units for statistics" (NUTS) boundaries.[68] Data on mortality were obtained at NUTS2 level for all member states between the years 1990-2014, though individual countries start and end years vary, as described in Table B1. Population data for each NUTS2 region were obtained through the EuroStat database. We download age-specific data according to the age groups used in the main analysis (<5, 5-64, >64). It is noted in the metadata that populations for NUTS2 regions are estimated to be applicable to the first day of each year, whereas mortality data are counted at the end of that year. Because of this, we offset the assignment of population and mortality by one year, so that, for example, 2005 mortality is matched with 2006 population on January 1^{st}. Administrative shapefiles are downloaded from the same source, and the 2013 version is used in the analysis. We drop the data on France from the EU dataset, as we obtain a higher spatial resolution source directly from the French government.

B.1.5 France

Mortality data for France are obtained at the ADM2-month level from the Institut National D'etudes Demographiques (National Institute for the Study of Demography (INED), 2019)[69] for the years 1998-2010. Data from this source do not have a categorization of mortality for a <5 year old age group, as used in the main analysis. The youngest age group for which there are data is ages 0-19. In the main analysis, we assign the mortality rates in the French data for the 0-19 age group to the <5 age group when pooling across countries. As this introduces some measurement error, we perform a robustness check in which we alternatively assign the deaths in the 0-19 age group to our 5-64 age group; this leads to a minimal change in the multi-country pooled results shown in Table D2. We aggregate the monthly data to the annual level for consistency with other countries' mortality records, and obtain administrative boundary files from GADM (Global Administrative Areas, 2012).

B.1.6 India

Annual data on Indian mortality rates at the district (i.e., ADM2) level were obtained from Burgess et al. (2017). A more thorough description of the data is given by the authors. The Indian data are not used in our main analysis, due to the absence of age-specific mortality rates and the importance of age in defining the mortality-temperature response function (e.g., see Figure 2). However, these data are used to assess the external validity of our extrapolation methods, as discussed in Appendix D.8.

[67] Data are available here: http://ec.europa.eu/eurostat/data/database.

[68] Administrative boundary files were downloaded from: http://ec.europa.eu/eurostat/web/gisco/geodata/reference-data/administrative-units-statistical-units/nuts.

[69] Data are available here: https://www.ined.fr/en/.

B.1.7 Japan

Japanese data on mortality and population at the prefecture-year[70] level were obtained from the National Institute of Population and Social Security Research[71] for the years 1975-2012. Data are available for all 47 prefectures of Japan, with no changes to administrative boundaries in that time. Mortality rates were downloaded as single-year age groups, which were then aggregated into the age groups used in the main analysis (<5, 5-64, >64). Prefecture (i.e., ADM1) boundaries were obtained from GADM (Global Administrative Areas, 2012).

B.1.8 Mexico

Mexican data on municipality-month deaths were obtained for the years 1990-2010 from the National Institute of Statistics and Geographical Information (INEGI), whose open-microdata repository contains the raw mortality files.[72] The data contain detailed information, including the municipality of occurrence and of residence, date, and age at death. We assign locations of deaths based on municipalities of residence. Data were downloaded as monthly mortality counts, then aggregated into municipality-age-year counts, using the age groups from the main analysis (<5, 5-64, >64). These counts were merged with municipality-by-year population values estimated from the Mexican census and as maintained at Minnesota Population Center's Integrated Public Use Microdata Series, International.[73] There were seven municipalities (less than 0.5% of total municipalities) that had inharmonious borders across data sets and years due to administrative splits or mergers; we assigned these municipalities into their respective unions before the splits or after the mergers.

B.1.9 United States

U.S. data on the universe of mortality and population at the county-year level were obtained from the Center for Disease Control (CDC) Compressed Mortality Files (CMF)[74] for the years 1968-2010. CDC removes values for county-year-age totals that are fewer than 10 deaths to preserve anonymity in the data in public files, and we obtain these through a data user agreement with CDC. There is some overlap in years available in the restricted and unrestricted datasets, and where both are available we use the restricted data due to better spatial coverage. In the restricted data, zeros are coded as missing, and so we reassign all missing values to zero. Data were downloaded in 5-year age groups and then aggregated to the age groups used in the main analysis (<5, 5-64, >64). The CMF reports deaths at the county of residence. Administrative boundaries are obtained from the TIGER datasets of the U.S. Census Bureau.[75]

[70] Japanese mortality data are the only data in our sample at first administrative level (i.e., ADM1). Though this is equivalent administratively to states in the U.S., the small size of the prefectures makes them comparable in geographic scale to large U.S. counties or EU NUTS2 regions.

[71] Data are available here: `http://www.ipss.go.jp/index-e.asp`.

[72] The initial link we used was `http://www3.inegi.org.mx/sistemas/microdatos/encuestas.aspx?c=33388&s=est` as of July, 2015. This link has been moved since, and now is being maintained at `http://en.www.inegi.org.mx/proyectos/registros/vitales/mortalidad/` as of June, 2018.

[73] Minnesota Population Center. Integrated Public Use Microdata Series, International: Version 7.0 [dataset]. Minneapolis, MN: IPUMS, 2018. `http://doi.org/10.18128/D020.V7.0`.

[74] Partial data are freely available through the CDC Wonder database.

[75] Data are available here: `https://www.census.gov/geo/maps-data/data/tiger-line.html`.

B.1.10 Aggregate data

Data from each country were standardized as annual rates for the age groups <5, 5-64, and >64, and were merged into a single file. We note that in all cases, place of residence is used for the assignment of temperature exposure to death records. In cases of inharmonious borders between years, we assign exposure based on a temporally consistent set of boundaries that are chosen to be in the most aggregate form, i.e., before administrative units split or after they merge.

[76] France is estimated using data from a different source and the EuroStat version of the France data is not used.

Table B1: Details of the European Union mortality sample

Code	Country	Number of NUTS2 regions	Years
AT	Austria	9	1990-2014 (no data for 1995)
BE	Belgium	11	1990-2014
BG	Bulgaria	6	1990-2014
CH	Switzerland	7	1991-2014
CY	Cyprus	1	1993-2014 (data before 1993 is not disaggregated by age group)
CZ	Czech Republic	8	1992-2014
DE	Germany	50	2002-2014 (2 regions are only available from 2011-2014)
DK	Denmark	5	2007-2014
EE	Estonia	1	1990-2014
EL	Greece	4	1990-2014 (data after 2013 is disaggregated into 13 regions)
ES	Spain	19	1990-2014
FI	Finland	5	1990-2014
FR	France	22	1990-2014 (an additional 4 regions are available in 2014)[76]
HR	Croatia	2	2001-2014
HU	Hungary	7	1990-2014
IE	Ireland	2	1997-2014
IS	Iceland	1	1990-2014
IT	Italy	21	1990-2014 (2 regions only have age-specific information after 2001)
LI	Liechtenstein	1	1994-2014
LT	Lithuania	1	1990-2014
LU	Luxembourg	1	1990-2014
LV	Latvia	1	2002-2014
ME	Montenegro	1	2005-2014
MK	Macedonia	1	1995-2014
MT	Malta	1	1995-2014 (mortality rates for ages <5 are only available from 1995)
NL	Netherlands	12	2001-2014
NO	Norway	7	1990-2014
PL	Poland	16	1991-2014
PT	Portugal	7	1992-2014
RO	Romania	8	1990-2014
SE	Sweden	8	1990-2014
SI	Slovenia	2	2014
SK	Slovakia	4	1997-2014
TR	Turkey	26	2009-2014
UK	United Kingdom	40	1999-2014 (4 regions only have data available after 2000, 2 after 2002, 5 for 2014 only)

B.2 Climate data

This appendix describes the climate data that we use throughout our analysis, as well as the methods that we use to make these data spatially and temporally consistent with the resolution of both historical mortality records and with future projection information. Broadly speaking, we use two classes of climate data: the first is historical data that we use to estimate the mortality-temperature relationship; the second is projected data on future climate, which we use to generate climate change damage estimates under various emissions scenarios. In this appendix, we describe the historical data, describe the projection data, detail our method for constructing a probabilistic ensemble of future climate projections at high resolution using these projection data, and finally we outline our method for spatial and temporal aggregation of both historical and projection climate data.

B.2.1 Historical climate data

Data on historical climate exposure is used to estimate the mortality-temperature response function as well as the heterogeneity in these responses across income and climate spaces. We use two separate groups of historical data on precipitation and temperature from independent sources. First, we use a reanalysis product, the Global Meteorological Forcing Dataset (GMFD) (Sheffield, Goteti, and Wood, 2006), which relies on a climate model in combination with observational data to create globally-comprehensive data on daily mean, maximum, and minimum temperature and precipitation (see Auffhammer et al. (2013) for a discussion of reanalysis data). Second, we repeat our analysis with climate datasets that strictly interpolate observational data across space onto grids. This comparison is important, as the sources of measurement error are likely to differ across reanalysis (which relies in part on a physical climate model) and interpolation (which relies purely on statistical methods such as kriging). For interpolated products, we use the daily Berkeley Earth Surface Temperature dataset (BEST) (Rohde et al., 2013) in combination with the monthly University of Delaware precipitation dataset (UDEL) (Matsuura and Willmott, 2007).

The GMFD dataset serves as our primary historical climate data source for analysis. A primary reason for this choice is that GMFD is used to bias-correct the climate model projections (described below), and using any other estimated relationship with these projection data would consequently be inconsistent. We use BEST and UDEL in order to ensure consistency of our estimated response surfaces across climate datasets.

Global Meteorological Forcing Dataset for Land Surface Modeling The main dataset used in this analysis is the Global Meteorological Forcing Dataset (GMFD) (Sheffield, Goteti, and Wood, 2006). These data provide surface temperature and precipitation information using a combination of both observations and reanalysis. The reanalysis process takes observational weather data and uses a weather forecasting model to interpolate both spatially and temporally in order to establish a gridded dataset of meteorological variables. The particular reanalysis used is the NCEP/NCAR reanalysis, which is downscaled and bias-corrected using a number of station-based observational datasets to remove biases in monthly temperature and precipitation (Sheffield, Goteti, and Wood, 2006). Data are available on a $0.25° \times 0.25°$ resolution grid from 1948-2010. The temporal frequency is up to 3-hourly, but the daily data are used for this analysis. We obtain daily average temperatures and monthly average precipitation for all grid cells globally.

Berkeley Earth Surface Temperature The Berkeley Earth Surface Temperature (BEST) dataset provides temperatures from 1701-2018 over land from a combination of observational records (Rohde et al., 2013), with spatially disaggregated data available from 1753.[77] During the time periods used within this paper (varying between 1957-2014), as many as 37,000 station records, representing 14 separate databases of station data, are incorporated into the BEST data. Station data are incorporated using a kriging methodology that allows for the incorporation of more stations with shorter time series than other well-known global surface temperature interpolation data (like the UDEL temperature dataset). In particular, the spatial averaging method uses close neighbors of a station to identify discontinuities in a particular time series that may be due to instrumental change or re-positioning, and decreases the influence of these changes in the spatially averaged grid (Rohde et al., 2013). This does have the potential drawback of over-smoothing the spatial heterogeneity in temperatures (National Center for Atmospheric Research Staff (Eds), 2015). BEST data are provided at daily frequency on a $1°\times1°$ resolution grid, and we utilize the daily average 2m air temperature variable for each grid cell.

University of Delaware Climate Dataset The University of Delaware climate dataset (UDEL) (Matsuura and Willmott, 2007) is used for precipitation in combination with the BEST data. UDEL provides gridded, interpolated data derived from weather stations on many variables at a monthly frequency and on a $0.5°\times0.5°$ resolution grid. Data are available from 1900-2014. The UDEL data are based on two underlying datasets of stations and have fewer observations underlying the interpolated grid, as compared to BEST. This is likely to lead to some decrease in interpolation accuracy in areas where the spatial coverage of weather stations is low (e.g., sub-Saharan Africa). The interpolation procedure used is based on inverse distance weighting to the central point of each grid cell, and the authors note that other data, like altitude and atmospheric characteristics, are used to improve that interpolation. The monthly average precipitation is obtained for all grid cells globally.

B.2.2 Climate projection data

Data on the future evolution of the climate is obtained from a multi-model ensemble of Global Climate Model (GCM) output. However, two important limitations arise when integrating GCM outputs into the current analysis. First, the relatively coarse resolution ($\sim 1°$ of longitude and latitude) of GCMs limits their ability to capture small-scale climate patterns, which render them unsuitable for climate impact assessment at high spatial resolution. Second, the GCM climate variables exhibit large local bias when compared with observational data.

To address both of these limitations, we use a high-resolution ($0.25°$ X $0.25°$) set of global, bias-corrected climate projections produced by NASA Earth Exchange (NEX): the Global Daily Downscaled Projections (GDDP) (Thrasher et al., 2012).[78] The NEX-GDDP dataset comprises 21 climate projections, which are downscaled from the output of global climate model (GCM) runs in the Coupled Model Intercomparison Project Phase 5 (CMIP5) archive (Taylor, Stouffer, and Meehl, 2012). The statistical downscaling algorithm

[77] Data are available here: http://berkeleyearth.org/data/.

[78] Climate projections used were from the NEX-GDDP dataset, prepared by the Climate Analytics Group and NASA Ames Research Center using the NASA Earth Exchange, and distributed by the NASA Center for Climate Simulation (NCCS).

used to generate the NEX-GDDP dataset is the Bias-Correction Spatial Disaggregation (BCSD) method (Wood et al., 2004; Thrasher et al., 2012), which was developed to address the aforementioned two limitations. This algorithm first compares the GCM outputs with observational data on daily maximum temperature, daily minimum temperature, and daily precipitation during the period 1950-2005. NEX-GDDP uses a climate dataset from GMFD for this purpose (Sheffield, Goteti, and Wood, 2006). A daily, quantile-specific relationship between GCM outputs and observations is derived from this comparison. This relationship is then used to adjust the GCM outputs in historical and in future time periods so that the systemic bias of the GCM is removed. To disaggregate the bias-corrected GCM outputs to higher resolution, this algorithm interpolates the daily changes relative to climatology in GCM outputs into the spatial resolution of GMFD, and merges the fine-resolution changes with the climatology of the GMFD data.

For each GCM, three different datasets are generated. The first uses historical emissions to simulate the response of the climate to historical forcing from 1850 to 2005. The second and third use projected emissions from Representative Concentration Pathways 4.5 and 8.5 (RCP4.5 and RCP8.5) to simulate emissions under those two emissions scenarios up to 2100. RCP 4.5 represents a "stabilization" scenario in which total radiative forcing is stabilized around 2100 (Riahi et al., 2011; Van Vuuren et al., 2011); RCP8.5 simulates climate change under intensive growth in fossil fuel emissions from 2006 to the end of the 21^{st} century. We use daily average temperature and daily precipitation in the RCP4.5 and RCP8.5 scenarios from this dataset, where the daily average temperature is approximated as the mean of daily maximum and daily minimum temperatures.

B.2.3 SMME and model surrogates

The CMIP5 ensemble of GCMs described above is an "ensemble of opportunity", not a systematic sample of possible futures. Thus, it does not produce a probability distribution of future climate change. Moreover, relative to simple climate models designed for probabilistic sampling of the global mean surface temperature (GMST) response to radiative forcing, the CMIP5 ensemble systematically fails to sample tail outcomes (Tebaldi and Knutti, 2007; Rasmussen, Meinshausen, and Kopp, 2016). To provide an ensemble of climate projections with a probability distribution of GMST responses consistent with that estimated by a probabilistic simple climate model, we use the surrogate model mixed ensemble (SMME) method (Rasmussen, Meinshausen, and Kopp, 2016) to assign probabilistic weights to climate projections produced by GCMs and to improve representation of the tails of the distribution missing from the ensemble of GCMs. Generally speaking, the SMME uses (1) a weighting scheme based on a probabilistic projection of global mean surface temperature from a simple climate model (in this case, MAGGIC6) (Meinshausen, Raper, and Wigley, 2011) and (2) a form of linear pattern scaling (Mitchell, 2003) that preserves high-frequency variability to construct model surrogates to fill the tails of probability distribution that are not captured by the GCM ensembles. This method provides us with an additional 12 surrogate models.

The SMME method first divides the unit interval [0,1] into a set of bins. For this analysis, the bins are centered at the 1^{st}, 6^{th}, 11^{th}, 16^{th}, 33^{rd}, 50^{th}, 67^{th}, 82^{nd}, 89^{th}, 94^{th}, and 99^{th} percentiles. Bins are narrower in the tails to ensure samples are created for portions of the GMST probability distribution function that are

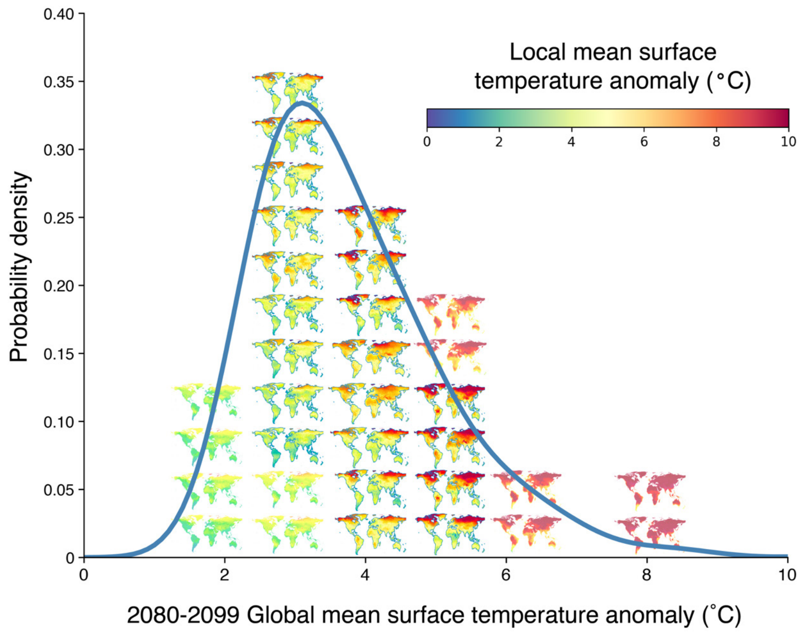

Figure B2: Future climate projections from the surrogate model/mixed ensemble (SMME). Figure shows the 21 climate models (outlined maps) and 12 model surrogates (maps without outlines) that are weighted in climate change projections so that the weighted distribution of the 2080 to 2099 global mean surface temperature anomaly (ΔGMST) exhibited by the 33 total models matches the probability distribution of estimated ΔGMST responses (blue-gray line) under RCP8.5. For this construction, the anomaly is relative to values in 1986-2005.

not captured by CMIP5 models. The bounds and center of each bin are assigned corresponding quantiles of GMST anomalies for 2080-2099 from simple climate model (SCM) output; in the application here and that of Rasmussen, Meinshausen, and Kopp (2016), this output came from the MAGICC6 (Meinshausen, Raper, and Wigley, 2011) model, constrained to match historical temperature observations and the conclusions of the IPCC Fifth Assessment Report regarding equilibrium climate sensitivity. The GMST of CMIP5 models are categorized into bins according to their 2080-2099 GMST anomalies.

If the number of CMIP5 models in a bin is less than 2, surrogate models are generated to raise the total number of models to 2 in that bin. The surrogate models are produced by using the projected annual GMST of the SCM that is consistent with the bin's central quantile to scale the spatial pattern of a selected CMIP5 model, then adding the intercept and residual from the same model. There are two cases of selecting CMIP5 models for pattern and residual. When there is only one CMIP5 model in a bin, an additional model is selected that has a GMST projection close to GMST in the bin and a precipitation projection over the region of interest complementary to the model already in the bin (i.e., if the model in the bin is relatively dry, then a relatively wet pattern is selected, and vice versa.) When there is no CMIP5 model, two models are picked with GMST projections close to that of the bin, with one model being relatively wet and one being relatively dry. In the final probabilistic distribution, the total weight of the bin is equally divided among the CMIP5 models and surrogate models in the bin. For instance, if four models are in the bin centered at the 30th percentile, bounded by the 20th – 40th percentiles, each will be assigned a probability of $20\% \div 4 = 5\%$. The resulting distribution of GMST for all members of the SMME is shown in Figure B2.

B.2.4 Aggregation of gridded climate data to administrative boundaries

We link gridded historical climate data to administrative mortality records by aggregating grid cell information to the same spatial and temporal level as the mortality records (see Table 1). Similarly, to generate future climate change impact projections at each of our 24,378 custom impact regions (impact regions are administrative regions or agglomerations of administrative regions; see Appendix C for details), we aggregate grid cell information to impact region scale. In both cases, nonlinear transformations of temperature and rainfall are computed at the grid cell level before averaging values across space using population weights and finally summing over days within a year. This procedure recovers grid-by-day-level nonlinearities in the mortality-temperature (and mortality-precipitation) relationship, because mortality events are additive (Hsiang, 2016).

To see how this calculation is operationalized, consider the fourth-order polynomial specification for temperature used in our main set of results for estimation of Equations D.17 and 5. In this case, we begin with data on average temperatures for each day d at each grid cell z, generating observations T_{zd}. These grid-level values must then be aggregated to the level of an administrative unit i in year t. To do this, we first raise grid-level temperature to the power p, computing $(T_{zd})^p$ for $p \in \{1, 2, 3, 4\}$. We then take a spatial average of these values over administrative unit i, weighting the average by grid-level population (and accounting for fractional grid cells that fall partially within administrative units). Population weights are time-invariant and calculated from the 2011 Landscan dataset (Bright et al., 2012). We then sum these daily polynomial terms T_{zd}^p over days in the year t. The vector of annual, administrative-level-by-year temperature variables we use for estimation is thus:

$$\boldsymbol{T}_{it} = \left[\sum_{d \in t} \sum_{z \in i} w_{zi} (T_{zd})^1, \sum_{d \in t} \sum_{z \in i} w_{zi} (T_{zd})^2, ..., \sum_{d \in t} \sum_{z \in i} w_{zi} (T_{zd})^P \right]$$

where w_{zi} is the share of i's population that falls into grid cell z, and where superscripts indicate polynomial powers. This nonlinear transformation performed prior to aggregation allows the aggregated measure of temperature to capture grid-by-day level exposure to very hot and very cold temperatures. In the econometric estimation of Equations D.17 and 5, quadratic polynomials in precipitation are similarly calculated and weighted averages are taken over administrative units. In Appendix Figure D3, we show robustness of the mortality-temperature relationship to four different nonlinear functional forms of temperature, all of which undergo an analogous grid-level transformation before averaging across space and summing over time. In future projections, all daily gridded climate projection data from each of the 33 members of the SMME are analogously aggregated across space and time.

B.3 Socioeconomic data and downscaling methodologies

This appendix provides details of the socioeconomic data used throughout our analysis, which includes historical subnational incomes, future projections of incomes, and future projections of population counts and age distributions. Additionally, because we require these variables at high spatial resolution both for econometric estimation and for future projections, we detail the downscaling procedures we use to disaggregate available socioeconomic data, which is generally provided at relatively low resolution.

B.3.1 Historical income data

Our main specification (Equation 5) estimates heterogeneity in mortality-temperature responses as a function of income and long-run average temperature in each location. In order to obtain income data for each subnational region in our mortality records, we draw subnational incomes from three main sources, using a combination of subnational GDP datasets as well as globally-comprehensive national GDP data:

- **Penn World Tables (PWT) national GDP**.[79] This dataset provides national level incomes from 1950 to 2014 for most of the countries in the world. We use Penn World Tables version 9.0 to obtain national level income for all countries in our sample (Brazil, Chile, China, France, India, Japan, Mexico, USA, and the 33 EU countries listed in Table B1).
- **Eurostat (2013) subnational GDP**.[80] This dataset provides national and sub-national level income data for the European countries in our dataset. We use this dataset to obtain subnational income at the NUTS2 level of aggregation, which is the level at which we observe mortality records.
- **Gennaioli et al. (2014) subnational GDP**. This dataset provides national and sub-national income data for 1,503 administrative regions from 83 countries. We use this dataset to obtain subnational level income data for all countries outside the EU: Brazil, Chile, China, France,[81] India, Japan, Mexico, and USA. Data are provided by the authors at the first administrative subdivision for each country (i.e., ADM1).

Using these data, we construct a consistent multi-country panel of subnational incomes at the NUTS2 level for EU countries and ADM1 level for the non-EU countries, which can be used for estimation of Equation 5. To do so, we use Eurostat (2013) and Gennaioli et al. (2014) to downscale the PWT national-level incomes. We prefer this approach to using the subnational data directly, as there are known inconsistencies in measurement of subnational GDP across countries. Thus, we make the assumption that the within-country distributions of GDP recorded in Eurostat (2013) and Gennaioli et al. (2014) are accurate, but the exact levels may not be. We rely on the PWT data as a consistent measure of GDP levels for all countries; thus, our subnational GDP estimates sum to national GDP from PWT for all countries in the sample. For administrative region s in country c in year t we calculate a weight, ν_{sct} that will apportion national income to subnational regions as follows:

[79] Penn World Tables (PWT) database: https://www.rug.nl/ggdc/productivity/pwt/.

[80] Eurostat database: http://ec.europa.eu/eurostat/data/database.

[81] As noted in Appendix B.1, we use higher resolution mortality data from France than that which is available through EuroStat. Therefore, we also rely on administrative income data from Gennaioli et al. (2014) instead of lower resolution income data from EuroStat.

$$\nu_{sct} = \begin{cases} \frac{GDPpc_{sct}^{Eurostat}}{\sum_{s \in c} GDPpc_{sct}^{Eurostat}} & \text{if c} \in \text{ EU} \\ \frac{GDPpc_{sct}^{Gennaioli}}{\sum_{s \in c} GDPpc_{sct}^{Gennaioli}} & \text{otherwise} \end{cases}$$
$$GDPpc_{sct} = \nu_{sct} \times GDPpc_{ct}^{PWT}$$

where $GDPpc^{PWT}$ corresponds to per capita GDP drawn from the PWT dataset. Using these estimates of administrative-level GDP per capita, we construct the time-invariant income covariate $\log(GDPpc)_s$ used for estimation of Equation 5 as follows. First, we take the log of our GDP per capita estimate for year t and region s. Second, we use a Bartlett kernel to compute a weighted average of lagged values of $\log(GDPpc)_{st}$, where the length of the kernel is empirically derived as described in Appendix E.1. We take this approach because changes in income are unlikely to immediately translate into changes in mortality-temperature sensitivity. Finally, we average this Bartlett kernel value across all years in the sample for each region s (note that the length of the panel varies by country, as shown in Figure B1).

Note that data in Eurostat (2013) are an annual panel. However, the data collected by Gennaioli et al. (2014) are drawn from disparate sources, often using census data, which are typically not annual, leading to an unbalanced panel. To construct annual values of income per capita using the Gennaioli et al. (2014) data, we linearly interpolate between years, before constructing the Bartlett kernel and taking averages across all years. For instances where we need to extrapolate backwards in time (i.e., when mortality data are available earlier than income data), we extrapolate backwards logarthmically. All subnational income data are in constant 2005 dollars PPP. A summary of the available years of data before interpolation is given in Table B2.

Country	ISO code	Years in mortality sample	Years in income sample[82]
Brazil	BRA	1997-2009	1995, 2000, 2005, 2010
China	CHN	1991-2012	1990, 1995, 2000, 2005, 2010
Chile	CHL	1997-2012	1995, 2000, 2010
EU		1990-2012	2003-2012
France	FRA	1998-2012	1995, 2000, 2005, 2010
India	IND	1957-2001	1980, 1985, 1990, 1995, 2000, 2005, 2010
Japan	JPN	1975-2012	1975, 1980, 1985, 1990, 1995, 2000, 2005, 2009
Mexico	MEX	1990-2012	1995, 2000, 2005, 2010
USA	USA	1968-2013	1965, 1970, 1975, 1980, 1985, 1990, 1995, 2000, 2005, 2009

Table B2: Temporal coverage of mortality records and years of available subnational income data.

[82] EU subnational income data come from Eurostat (2013). For all other countries, subnational income data are obtained

B.3.2 Income projections and downscaling methodology

Future projections of national incomes are derived from the Organization for Economic Co-operation and Development (OECD) Env-Growth model (Dellink et al., 2015) and the International Institute for Applied Systems Analysis (IIASA) GDP model (Samir and Lutz, 2014), as part of the "socioeconomic conditions" (population, demographics, education, income, and urbanization projections) of the Shared Socioeconomic Pathways (SSPs). The SSPs propose a set of plausible scenarios of socioeconomic development over the 21^{st} century in the absence of climate impacts and policy for use by the Integrated Assessment Modeling (IAM) and Impacts, Adaptation, and Vulnerability (IAV) scientific communities.

While there are many models within the SSP database, only the IIASA GDP model and OECD Env-Growth model provide GDP per capita projections for a wide range of countries. The IIASA GDP model describes incomes that are lower than the OECD Env-Growth model, so we produce results for both of these models to capture uncertainty within each socioeconomic scenario (we compute results for three socioeconomic scenarios: SSP2, SSP3, and SSP4). To construct annual estimates, we smoothly interpolate between the time series data in the SSP database, which are provided in 5-year increments. For each 5-year period, we calculate the average annual growth rate, and apply this growth rate to produce each year's estimate of GDP per capita.[83]

Throughout the main text, we show results relying on SSP3, although sensitivity of all main results to socioeconomic scenario are shown in the Appendix. While the methodology we develop to estimate future impacts of climate change on mortality, as well as a partial mortality-only SCC, can be applied to any available socioeconomic scenario, we emphasize SSP3 because its historic global growth rates in GDP per capita and population match observed global growth rates over the 2000-2018 period much more closely than either SSP2 or SSP4, as shown below in Table B3.

Although the SSP scenarios provide national-level income projections, our high-resolution analysis requires estimates of location-specific GDP within country borders. To generate values of income for each of our 24,378 impact regions over time, we allocate national GDP per capita values from the SSPs across impact regions within a country through a downscaling procedure that relies on nightlights imagery from the NOAA Defense Meteorological Satellite Program (DMSP). This approach proceeds in two steps. First, we use available subnational income data from Gennaioli et al. (2014) in combination with higher-resolution income data from the U.S., China, Brazil, and India, to empirically estimate the relationship between GDP per capita and nightlight intensity.[84] Second, we use this estimated relationship to allocate national-level GDP data across impact regions within each country, based on relative intensity of night lights in the present. While this approach models heterogeneity in income levels across impact regions, each region grows in the future at the same rate as the national country projection from the SSPs. We detail these two steps below.

Estimation of the GDP-nightlights relationship While there exists a growing literature linking

from Gennaioli et al. (2014).

[83] OECD estimates of income are provided for 184 countries and IIASA's GDP projections cover 171 countries. For the remaining countries, we apply the average GDP per capita from the available countries for the baseline period, and allow this income to grow at the globally averaged growth rate.

[84] Due to cross-country inconsistencies in subnational income data, the income data for the US are primarily used to estimate the relationship between GDP per capita and nightlights intensity; other countries' data provide validation only.

Table B3: Comparison of SSP growth rates to observed data in the historical record This table shows global average growth rates in GDP per capita and in population from observational data (World Bank), as well as from each SSP scenario used in our analysis. Note that International Institute for Applied Systems Analysis (IIASA) GDP model (Samir and Lutz, 2014) only provides GDP per capita estimates after 2010. For both GDP per capita and population, and for each historical time period, SSP3 matches historical data more closely; we therefore show climate change projection results using this scenario throughout the main text.

	Reference	*Scenario*		
	World Bank	SSP2	SSP3	SSP4
GDP per capita				
OECD (2000-2018)	2.39%	2.65%	2.57%	2.63%
OECD (2010-2018)	2.37%	3.01%	2.85%	2.98%
IIASA (2010-2018)	2.37%	3.69%	3.17%	3.55%
Population				
IIASA (2000-2018)	1.21%	1.13%	1.18%	1.12%
IIASA (2010-2018)	1.17%	1.04%	1.13%	1.02%

economic output to nightlights intensity, we take an unconventional regression approach to recovering this relationship because our goal is to apportion national income within a country, as opposed to predict the level of income at any given location. In particular, we are interested in the ratio $\frac{GDPpc_{rc}}{\sum_{r\in c} w_{rc}GCPpc_{rc}}$ for impact region r in country c (where w_{rc} is a region-specific population weight), which will allow us to predict income at the impact region level, given projections of national GDP per capita from the SSPs, $\sum_{r\in c} w_{rc}GDPpc_{rc} = GDPpc_c^{SSP}$. Thus, we estimate a regression relating *relative* GDP per capita to *relative* nightlights intensity, where each administrative region's values are calculated as relative to the country mean. The dependent variable for administrative region i in country c and year t is thus $\frac{GDPpc_{ict}}{\sum_{i\in c} w_{ict}GDPpc_{ict}}$.[85] To construct a measure of location-specific relative nightlight intensity, we calculate a z-score of nightlights (ZNL) for each administrative region i within a country c using:

$$ZNL_{ict} = \frac{NL_{ict} - \overline{NL}_{ct}}{\sigma(NL_{ct})}$$

where $\overline{NL}_{ct}$ is the country average nightlights intensity, $\sigma(NL_{ct})$ is the standard deviation of nightlights intensity across all administrative regions within country c, and where the stable nightlights data product from 1992-2012 is used to construct time-varying measures of average nightlights intensity across an administrative region, NL_{ict}.

The regression we estimate is as follows:

$$\frac{GDPpc_{ict}}{\sum_{i\in c} w_{ict}GDPpc_{ict}} = \alpha + \beta ZNL_{ict} + \epsilon_{ict} \tag{B.15}$$

where β represents the impact of a one standard deviation increase in a region's nightlights intensity, relative to its country average, on that region's relative GDP per capita.

Allocation of national GDP to impact regions using relative nightlight intensity We use the

[85] As discussed, the income data available from Gennaioli et al. (2014) are at the first administrative level (i.e. ADM1).

estimated coefficients from Equation B.15 to compute income at impact region level. To do so, we construct values $ZNL_{rct} = \frac{NL_{rct} - \overline{NL}_{ct}}{\sigma(NL_{ct})}$ for each impact region r using the average of stable nightlights from DMSP across the years 2008-2012. We then estimate $GDPpc_{rct}$ as follows:

$$\widehat{GDPpc}_{rct} = \left[\hat{\alpha} + \hat{\beta} ZNL_{rct}\right] \times GDPpc_{ct}^{SSP}$$

where $\sum_{r \in c} w_{rc} GDPpc_{rc}$ comes from one of the SSP projected income scenarios. The result of this approach is that the subnational downscaled incomes will sum to the national income from the SSPs, as these ratios sum to one, by construction.

B.3.3 Population projections and downscaling methodology

Future projections of national populations are derived from the International Institute for Applied Systems Analysis (IIASA) (Samir and Lutz, 2014) population projections as part of the Shared Socioeconomic Pathways (SSPs).[86] The IIASA SSP population projections provide estimates of population by age cohort, gender, and level of education for 193 countries from 2010 to 2100 in five-year increments. Each projection corresponds to one of the five SSPs, as defined in O'Neill et al. (2014). These populations are mapped to impact regions by country code using 3-digit country ISO-codes.

To assemble population projections for each of our 24,378 impact regions, we downscale the country-level projections from the SSPs using 2011 high-resolution LandScan estimates of populations (Bright et al., 2012). Populations for impact regions in countries or areas not given in the SSP database are held constant at the values estimated by LandScan in 2011. Thus, for any given impact region r in year t, population for scenario v (pop_{rtv}) is estimated as:

$$\widehat{pop}_{rtv} = \begin{cases} pop_{ctv}^{SSP} \left(\frac{pop_{r,2011}^{LandScan}}{\sum_{r \in c} pop_{r,2011}^{LandScan}} \right), & \text{if } r \in C \\ pop_{r,2011}^{LandScan}, & \text{if } r \notin C \end{cases} \tag{B.16}$$

where pop_{ctv}^{SSP} is the SSP population given for country c and year t for scenario v, $pop_{r,2011}^{LandScan}$ is the LandScan estimate for impact region r, and C is the set of 193 countries available in the SSP Database. Note that while this approach distributes country-level projections of population heterogeneously to impact regions within a country, it fixes the relative population distribution within each country at the observed distribution today. The division of population totals into the three age categories used throughout the analysis (0-4, 5-64, >64) is assumed to be constant across all impact regions within a country, and is thus taken directly from the SSPs.

B.4 Scale and scope of existing empirically-based estimates of the mortality risk of climate change

[86] The population data are accessed from the SSP database (IIASA Energy Program, 2016).

Table B4: Scale and scope of existing empirically-based estimates of the mortality risk of climate change. Table highlights papers quantifying the mortality risks of climate change that either are broad in spatial scope (e.g., covering multiple countries) or aim to account for at least one driver of adaptation (e.g., income). The first row highlights the present study, while the last row indicates the studies used to calibrate the mortality component of the FUND IAM, which is currently used to inform the U.S. government's social cost of carbon (SCC).

Study authors	**Spatial extent**	**Temporal extent**	**Types of adaptation accounted for**			**Future climate change projections**
			Climate adaptation	*Income-related adaptation*	*Adaptation costs*	
Carleton et al., 2021	Estimated on 40 countries, empirically-based extrapolation to global	1968 - 2010	Yes	Yes	Yes	Yes: ensemble of 33 IPCC-recommended climate models
Burgess et al., 2017	India	1957 - 2000		Yes: rural vs. urban, income, credit access		Yes: 1 climate model; adaptation not projected
Barreca et al., 2016	United States	1900 - 2004	Yes: region-specific models are estimated	Yes: healthcare access, electricity access, A/C adoption		
Deschenes, 2018	16 East, South, and SE Asian countries	1960 - 2015				Yes: 1 climate model
Heutel et al., 2017	United States	1992 - 2011	Yes			Yes: ensemble of 21 IPCC-recommended climate models
Portnykh, 2017	Russia	2006 - 2014	Yes			Yes: 1 climate model
Geruso & Spears, 2018 *(infant only)*	53 developing countries across Africa, Latin America, and Asia	1980 - 2010		Yes: literacy, income		
Guo et al., 2018	412 municipalities in N. America, S. America, Europe, Asia, and Oceania	1984 - 2015	Yes: local heat wave thresholds dependent on future temperatures			Yes: 5 climate models
Martens, 1998 (calibrates the FUND model)	20 cities in N. America, S. America, Europe, Asia, Africa, and Oceania	Meta-analysis, time periods vary				Yes: 3 climate models

C Spatial units for projection: "Impact regions"

We create a set of custom boundaries that define the spatial units for which location-specific projected damages of climate change are computed. To do so, we utilize politically defined regions, as opposed to a regular grid, as socioeconomic data are generally collected at this scale and because administrative regions are relevant to policy-makers. These regions, hereafter referred to as "impact regions", are constructed such that they are identical to existing administrative regions or are a union of a small number of administrative regions. We use version 2 of the Global Administrative Region dataset (GADM) (Global Administrative Areas, 2012), which contains 218,328 spatial units, to delineate boundaries. However, for computational feasibility and greater comparability across regions, we agglomerate these regions to create a set of 24,378 custom impact regions. To conduct this agglomeration, we establish a set of criteria that ensures these impact regions have approximately comparable populations and are internally consistent with respect to mean temperature, diurnal temperature range, and mean precipitation. A map of these regions is shown in Figure C1, and we detail this agglomeration algorithm below.

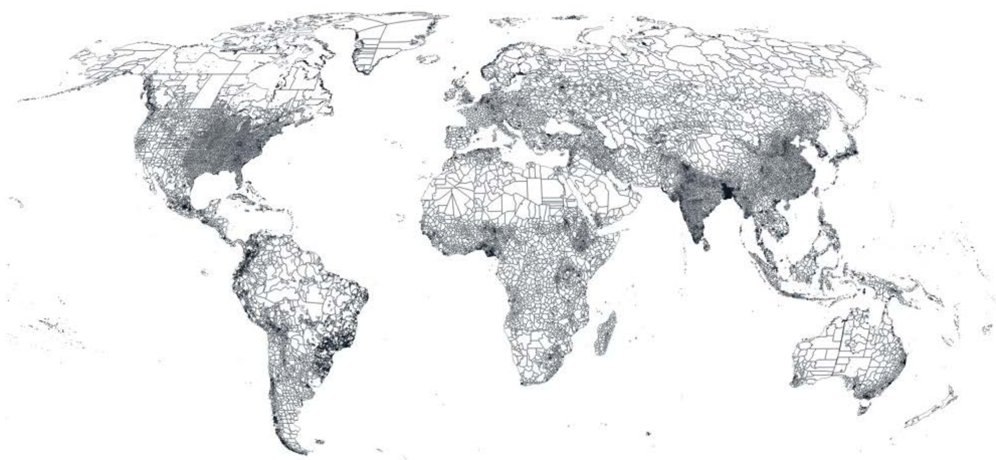

Figure C1: Map of the 24,378 "impact regions" for which location-specific projections are calculated. We use a clustering algorithm to form these regions from the full set of GADM administrative regions, such that they are roughly similar in total population, and so that they are approximately internally homogenous with respect to mean temperature, diurnal temperature range, and mean precipitation.

C.1 Algorithm for construction of impact region boundaries

We develop an algorithm which agglomerates administrative units from GADM into a smaller set of impact regions. Our goal is to create a set of approximately 20,000 impact regions that are spatially compact, of approximately equal population, and exhibit internally homogeneous climates. This procedure is conducted in three steps.

Step 1: Constructing a target region count for each country First, for each country, we generate a target number of regions; this is the number of regions that a country should roughly be divided into, based on its spatial extent, population, and climatic variability, and conforming to the goal of constructing approximately 20,000 global regions. We create this target for country c as the arithmetic mean of a population-based target and a climate-based target:

$$\begin{aligned} target_c &= \frac{1}{2}\left[population\_target + climate\_target\right] \\ &= \frac{1}{2}\left[20000\frac{pop_c}{\sum_c pop_c} + 20000\frac{A_c V_c}{\sum_c A_c V_c}\right] \end{aligned}$$

where pop_c is population of country c in 2011 from Landscan (see Appendix B.3.3) and A_c is the total area of country c. The variable V_c is a measure of a country's internal climate variability, relative to the global average, and is defined as follows:

$$V_c = \frac{Var_z[T]}{\mathrm{E}_c[Var_z[T]]} + \frac{Var_z[D]}{\mathrm{E}_c[Var_z[D]]} + \frac{Var_z[R]}{\mathrm{E}_c[Var_z[R]]} + \frac{Var_z[Q]}{\mathrm{E}_c[Var_z[Q]]}$$

where T is mean daily temperature, D is the diurnal temperature range, R is precipitation in the wettest month of the year, Q is precipitation in the driest month of the year, and where variances are taken over grid cells z within country c and expectations are taken over all countries c.

Step 2: Categorization of countries based on their target region count Second, we categorize countries based on whether there exists an administrative level in the GADM dataset (e.g. ADM1, which are equivalent to U.S. states; ADM2, which are equivalent to U.S. counties) for which the number of administrative units is roughly equivalent to the target number of regions. This categorization process leads to each country being divided into one of three cases, as shown in Figure C2. First, if there exists a GADM administrative level l, in country c, for which N_l, the number of administrative regions at level l, lies within the range $\frac{1}{2}target_c \leq N_l \leq 2target_c$, we simply use the administrative level l as our set of impact regions for country c. Countries which fall into this category are shown in shades of blue in Figure C2. This categorization includes the case where $target_c \leq 1$, in which case the entire country (i.e. ADM0 in GADM) is one impact region (shown in the lightest blue). Second, if the target number of regions for country c exceeds the maximum available region disaggregation in GADM, we simply use the highest resolution administrative level available from GADM. Countries which fall into this category are shown in dark blue in Figure C2. Finally, for all other countries, administrative units from GADM must be agglomerated to construct impact regions at a lower level of spatial resolution; these countries are shown in red in Figure C2. The agglomeration algorithm is described below.

Step 3: Agglomeration algorithm for impact region construction The third step in the process of constructing impact regions is to develop an agglomeration algorithm that will cluster administrative units from GADM into lower spatial resolution regions. Note that this third step only has to be conducted for the countries shown in red in Figure C2, as all other countries have a target number of impact regions that is well approximated by existing GADM administrative regions at some level l. For these remaining counties, the algorithm proceeds as follows.

First, we calculate a set of attributes at the highest administrative level available from GADM within each country. As the agglomerations are performed, the attributes of each new agglomerated region are generated from its component regions. These attributes are as follows:

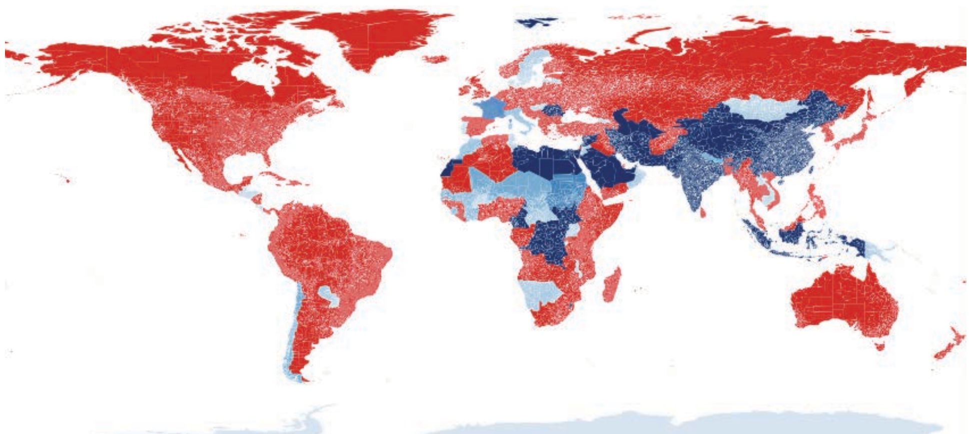

Figure C2: Categorization of countries based on the method used to construct impact regions out of GADM administrative regions. A country's target number of impact regions is $target_c$, as computed in the text. Countries in shades of blue have target values that can be approximated by one of the available GADM administrative levels l, such as ADM1 or ADM2, as there exists a level l such that the total number of administrative regions, N_l, falls within the range $\frac{1}{2}target_c \leq N_l \leq 2target_c$. Darker shades denote higher administrative levels, which have more regions. The ADM0 (country) level is also used if $target_c \leq 1$, and the highest available administrative level is used if $target_c$ is greater than the maximum N_l for country c. Finally, countries in red require agglomeration from the native GADM regions, as there is no administrative level l which satisfies the range criterion above, given the target region count $target_c$. This agglomeration algorithm is described in the text. We make an exception for the United States, shown in red, and represent it at ADM2 (county) level.

- The set of GADM regions within the agglomeration
- The set of neighboring agglomerated regions
- Population (pop),[87] and area (A)
- Socioeconomic and climatic traits ($\{T\}$): population density, average temperature, diurnal temp range, wet season precipitation, and dry season precipitation
- Centroids of all GADM regions contained within the agglomeration ($\{(Lat, Lon)\}$)

The agglomeration process is a greedy algorithm, which performs the following steps:

1. A set of proposed agglomerations is generated. For a given region r within a containing administrative region S_l of administrative level l, these consist of:
    - The combination of r with each of its neighbors within S_l.
    - The next higher administrative region, S_{l+1} (e.g., all counties within the same state).
    - If neither of the above is available (e.g., an island state, with S_l equalling the country), the combination of r and the closest neighbor also at the first administrative level.

[87] Population data are from Landscan (Bright et al., 2012), as in Appendix B.3.3.

2. Each proposed agglomeration from step 1, across all regions, is scored. For a region r containing subregions indexed by j, the scores consist of a weighted sum of the following:

Attribute	Expression	Weight
Area	$(\sum_j A_j/A_0)^2$, where A_0 is the average US county area	0.01
Population	$(\sum_j pop_j/pop_0)^2$, where pop_0 is the average US county population	1
Dispersion	$Var[Lat] + Var[Lon \cos \mathrm{E}[Lat]]$	10
Other traits	$\sum_T Var[T_r]/T_0$, where T_0 is 1 for population density, 100 for elevation, 8.0 for mean temperature, 2.1 for diurnal temperature range, 25.0 for wet season precipitation and 2.6 for dry season precipitation	100
Circumference	$M \frac{n}{6\sqrt{M}}$, where M is the number of contained regions and n is the number of neighboring regions	1

3. The agglomeration with the smallest score from step 2 is identified.

4. The regions within the new agglomeration are merged, and new properties are applied to the new region.

5. This process repeats until the target number of regions $target_c$ for country c is reached.

D Econometric estimation: Additional results, robustness, out-of-sample validation

This appendix shows additional illustrations of and tabular results for the main econometric regressions used and discussed throughout the main text (Figures D1 and D2, and Table D1), results obtained using a pooled version of the main model in which no heterogeneity in the mortality-temperature relationship is modeled (Figure 2 and Table D2), a set of robustness checks for the main empirical results (Figures D3, D5, D6, and D7 and Tables D3 and D4), and a set of out-of-sample validation tests designed to evaluate the accuracy with which our estimates predict mortality-temperature responses in locations and time periods that are not used for estimation (Figures D8, D9, and D10 and Tables D5 and D6).

D.1 Age-specific heterogeneity of the mortality-temperature response function by average income and average climate

The estimation of Equation 5 tests for systematic heterogeneity in the mortality-temperature response function by modeling interactions between the temperature variables ($\boldsymbol{T}$) and the ADM1-level covariates of average climate ($TMEAN$) and average income ($\log(GDPpc)$). To see how we implement Equation 5 in practice, note that in Equation D.17, we estimate $g_a(\cdot)$ as the inner product between the nonlinear functions of temperature $\boldsymbol{T}_{it}$ and a vector of coefficients $\boldsymbol{\beta}_a$; that is, $g_a(\boldsymbol{T}_{it}) = \boldsymbol{\beta}_a \boldsymbol{T}_{it}$. For example, in the polynomial case, $\boldsymbol{T}_{it}$ is a vector of length P and contains the annual sum of daily average temperatures raised to the powers $p = 1, ..., P$ and aggregated across grid cells. The coefficients $\boldsymbol{\beta}_a$ therefore fully describe the age-specific nonlinear response function. In Equation 5, we allow $g_a(\boldsymbol{T}_{it})$ to change with climate and income by allowing each element of $\boldsymbol{\beta}_a$ to be a linear function of these two variables. We do not include a triple interaction between temperature, climate and income. Using this notation, our estimating equation is:

$$M_{ait} = \underbrace{(\boldsymbol{\gamma}_{0,a} + \boldsymbol{\gamma}_{1,a} TMEAN_s + \boldsymbol{\gamma}_{2,a} \log(GDPpc)_s)}_{\boldsymbol{\beta}_a} \boldsymbol{T}_{it} + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait}$$

where $\boldsymbol{\gamma}_{0,a}, \boldsymbol{\gamma}_{1,a}$, and $\boldsymbol{\gamma}_{2,a}$ are each vectors of length P, the latter two describing the effects of $TMEAN$ and $\log(GDPpc)$ on the sensitivity of mortality M_{ait} to temperature $\boldsymbol{T}_{it}$.

Tabular results from this estimation are reported in Table D1 for each of the three age groups of interest. Each coefficient represents the change in the temperature-sensitivity of mortality rates associated with a marginal increase in the relevant covariate (e.g., $TMEAN$), evaluated at the daily temperature shown. All temperature sensitivities are shown relative to a moderate day at 20°C. For example, higher incomes correspond with lower sensitivity of infant mortality to both cold temperatures (coefficient of -0.87 on a -5°C day), and to hot temperatures (coefficient of -0.93 on a 35°C day).[88] Although not all of the coefficients would be judged statistically significant by conventional criteria, it is noteworthy that higher incomes and warmer climates are associated with lower mortality consequences of hot days for all age categories. Income

[88] Because our covariates are linearly interacted with the full vector of temperature variables describing the nonlinear mortality-temperature response, the effect of each covariate depends on the realized daily temperature.

and climate are associated with cold day mortality differentially across age groups, with some evidence that higher income locations exhibit more extreme cold day sensitivity for the oldest age group. This relationship may arise due to age being positively correlated with income within the over 64 category, as older individuals are more susceptible to cold-related death risks (Deschênes and Moretti, 2009).

Table D1: Marginal effect of covariates on temperature sensitivity of mortality rates. Coefficients (standard errors) represent the marginal effect of increasing each covariate by one unit on the temperature sensitivity of mortality, evaluated at each of the shown daily average temperatures. Temperature sensitivity is defined as the impact of a particular temperature on mortality rates, relative to a moderate day at 20°C. Regression is a fourth-order polynomial in daily average temperature, estimated using GMFD weather data with a sample that was winsorized at the top 1% level. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects. Each temperature variable is interacted with each covariate.

	Age < 5		Age 5-64		Age >64	
	log($GDPpc$)	$TMEAN$	log($GDPpc$)	$TMEAN$	log($GDPpc$)	$TMEAN$
35° C	-0.887*	-0.099*	-0.236	-0.031*	-3.881	-0.624*
	(0.536)	(0.053)	(0.160)	(0.018)	(2.380)	(0.331)
30°C	-0.280	-0.044	-0.019	-0.014	-0.189	-0.292**
	(0.277)	(0.028)	(0.068)	(0.009)	(0.910)	(0.141)
20°C	–	–	–	–	–	–
	–	–	–	–	–	–
0°C	-0.973*	0.029	0.050	-0.030*	0.269	-0.731***
	(0.536)	(0.031)	(0.150)	(0.018)	(2.019)	(0.153)
-5°C	-1.165*	0.028	0.216	-0.040**	3.097	-0.920***
	(0.629)	(0.032)	(0.210)	(0.020)	(2.956)	(0.202)

Regression includes $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects. Adjusted R^2 = 0.933; N=820,237. Standard errors clustered at the ADM1 level. *** p<0.01, ** p<0.05, * p<0.1

As these terms are difficult to interpret, we visualize this heterogeneity in the main text in Figure 1 by dividing the sample into terciles of income and climate (i.e., the two interaction terms), creating nine discrete bins describing the log($GDPpc$) $\times$ $TMEAN$ space. We plot the predicted response functions at the mean value of covariates within each of these nine bins, using the coefficients shown in Table D1. This results in a set of predicted response functions that vary across the joint distribution of income and average temperature within our sample data, shown in Figure 1 for the >64 age category. Figures D1 and D2 replicate this figure for the other two age groups in our analysis.

D.2 Results and robustness with pooled model

In the main text, we estimate the mortality-temperature relationship while explicitly modeling heterogeneity due to income and climate. In this sub-section, we instead show a series of results using a model without interactions, yielding average treatment effects within age groups. One key advantage of this simpler model is that it is straightforward to examine robustness of estimates to various datasets and fixed effect specifications. Below we describe the model, show results for average mortality-temperature relationships by age group, demonstrate robustness to other functional forms of temperature and other climate datasets, and show results for alternate model specifications.

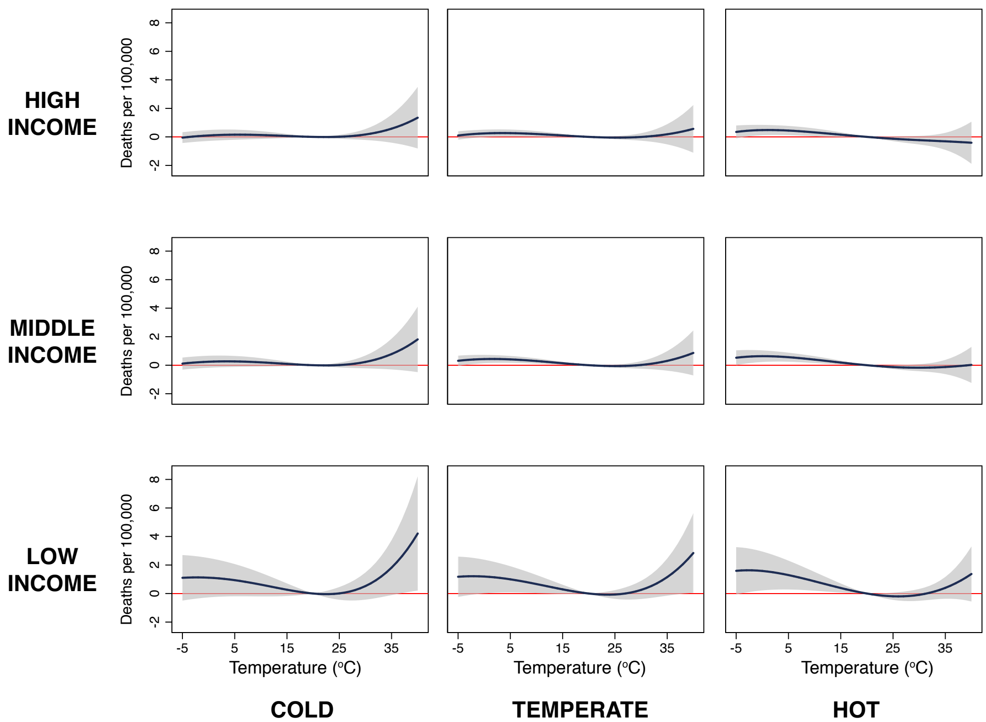

Figure D1: Heterogeneity in the mortality-temperature relationship (ages <5 mortality rate). Each panel represents a predicted response function for the ages <5 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left are the predicted mortality-temperature sensitivities for low income, cold regions of our sample, while those in the upper right apply to the high income, hot regions of our sample. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level on the top end of the distribution only. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects, and where each temperature variable is interacted with each covariate and a dummy for each age category.

D.2.1 Estimating a pooled multi-country mortality-temperature response function

Here we estimate a pooled, multi-country, age-specific, mortality-temperature response function. The model exploits year-to-year variation in the distribution of daily weather to identify the response of all-cause mortality to temperature, following, for example, Deschênes and Greenstone (2011). Specifically, we estimate the following equation on the pooled mortality sample from 40 countries:

$$M_{ait} = g_a(\boldsymbol{T}_{it}) + q_{ac}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{D.17}$$

where a indicates age category with $a \in \{< 5, 5\text{-}64, > 64\}$, i denotes the second administrative level (ADM2, e.g., county),[89] c denotes country, and t indicates years. Thus, M_{ait} is the age-specific all-cause mortality rate in ADM2 unit i in year t. α_{ai} is a fixed effect for $age \times ADM2$, and δ_{act} a vector of fixed effects that allow for shocks to mortality that vary at the $age \times country \times year$ level.

[89]This is usually the case. However, as shown in Table 1, the EU data is reported at Nomenclature of Territorial Units for Statistics 2^{nd} (NUTS2) level, and Japan reports mortality at the first administrative level.

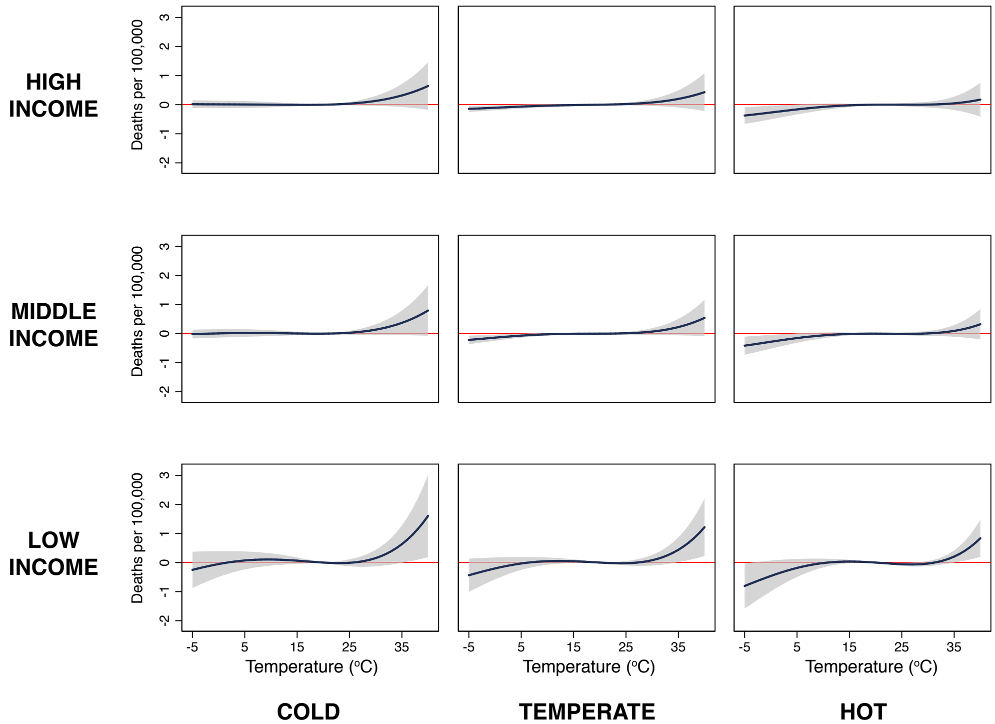

Figure D2: Heterogeneity in the mortality-temperature relationship (ages 5-64 mortality rate). Each panel represents a predicted response function for the ages 5-64 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left are the predicted mortality-temperature sensitivities for low income, cold regions of our sample, while those in the upper right apply to the high income, hot regions of our sample. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level on the top end of the distribution only. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects, and where each temperature variable is interacted with each covariate and a dummy for each age category.

Our focus in Equation D.17 is the effect of temperature on mortality, represented by the response function $g_a(\cdot)$, which varies by age. As in our our main specification, $\boldsymbol{T}_{it}$ contains polynomials of daily average temperatures (up to fourth order), summed across the year. These calculations are performed at the grid cell level before being aggregated up to the level of the administrative units in the data. Results for alternative functional form specifications are shown in Figure D3 and the consequences of alternate functional forms for climate change projection results are shown in Appendix F. Analogous to temperature, we summarize daily grid-level precipitation in the annual ADM2-level vector $\boldsymbol{R}_{it}$. We construct $\boldsymbol{R}_{it}$ as a second-order polynomial of daily precipitation, summed across the year, and estimate an age- and country-specific linear function of this vector, represented by $q_{ac}(\cdot)$.

We fit the multi-country pooled model in Equation D.17 using weighted least squares, weighting by age-specific population so that the coefficients correspond to the average person in the relevant age category and to account for the greater precision associated with mortality estimates from larger populations.[90] Standard

[90] We constrain population weights to sum to one for each year in the sample, across all observations. That is, our weight for an observation in region i in year t for age group a is $\omega_{it}^a = pop_{it}^a / \sum_i \sum_a pop_{it}^a$. This adjustment of weights is important in our context, as we have a very unbalanced panel, due to the merging of heterogeneous country-specific mortality datasets.

errors are clustered at the first administrative level (ADM1, e.g., state), instead of at the unit of treatment (ADM2, e.g., county), to account for spatial as well as temporal correlation in error structure. Robustness of this model to alternative fixed effects and error structures is shown in Table D2, and to alternative climate datasets in Figure D3.

Age-specific pooled multi-country mortality-temperature response functions. As prior work has shown that age cohorts respond differently to temperature, and because we expect considerable demographic transitions in the future, we allow for heterogeneity across age groups in Equation D.17. Specifically, we allow for separate mortality-temperature response functions $g_a(\boldsymbol{T}_{it})$ for each of three age categories (<5, 5-64, >64). Figure 2 in the main text displays the mortality-temperature responses for each age group, estimated from Equation D.17 and using the pooled 40-country sample and our preferred specification (column (2) in Table D2). This reveals substantial heterogeneity across age groups within our multi-country sample: people over the age of 64 experience approximately 4.7 extra deaths per 100,000 for a day at 35°C (95°F) compared to a day at 20°C (68°F), a substantially larger effect than that for younger cohorts, which exhibit little response. This age group is also more severely affected by cold days; estimates indicate that people over the age of 64 experience 3.4 deaths per 100,000 for a day at -5°C (23°F) compared to a day at 20°C, while there is a relatively weak mortality response to these cold days for other age categories. Overall, these results demonstrate that the elderly are disproportionately harmed by additional hot days and disproportionately benefit from reductions in cold days, consistent with prior evidence from the U.S. (Deschênes and Moretti, 2009; Heutel, Miller, and Molitor, 2017). It is important to note, however, that the oldest age group (over 64 years) accounts for just 12% of the population in our historical sample.

Robustness to temperature functional form and climate data. Figure D3 displays the results of estimating a version of Equation D.17 using a set of different functional forms of temperature (i.e., different formulations of the temperature vector $\boldsymbol{T}_{it}$) and using two different climate datasets to obtain those temperatures (see Appendix B.2 for details on these climate datasets). Here we show the mortality response $g_a(\boldsymbol{T}_{it})$ for the >64 age group. The four functional forms estimated are fourth-order polynomials, bins of daily average temperature, restricted cubic splines, and piecewise linear splines. The binned functional form is an important benchmark, as it is closest to being fully non-parametric; the similarity of the binned regression response functions with those from three other functional forms is reassuring. The GMFD climate data (top) and BEST climate data (bottom) are drawn from independent sources, as described in Appendix B.2, and lead to broadly similar response functions across all functional forms. All regressions include $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects, and are population weighted.

Alternative specifications. In Table D2, marginal effects of temperature on age-specific mortality rates are shown for a range of alternative specifications. These estimates can be interpreted as the change in the number of deaths per 100,000 per year resulting from one additional day at each temperature, compared to the reference day of 20°C (68°F). Columns (1)-(3) increase the saturation of temporal controls in the model specification, ranging from country-year fixed effects in column (1) to country-year-age fixed effects in column (2), and adding age-specific state-level linear trends in column (3). Our preferred specification is column (2), as column (1) does not account for differential temporal shocks to mortality rates by age group, while in column (3) we cannot reject the null of equal age-specific, ADM1-level trends. In column

Table D2: Temperature-mortality response function with demographic heterogeneity estimated using pooled subnational data. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the top 1% level. Point estimates indicate the effect of a single day at each daily average temperature value shown, relative to a day with an average temperature of 20°C (68°F).

	Age-specific mortality rate (per 100,000)				
	(1)	(2)	(3)	(4)	(5)
Panel A: <5 years of age					
35° C	2.218***	-0.003	0.041	0.074	-0.060
	(0.487)	(0.252)	(0.157)	(0.212)	(0.252)
30° C	1.303***	-0.077	0.009	0.027	-0.076
	(0.217)	(0.102)	(0.065)	(0.092)	(0.102)
20°C	–	–	–	–	–
	–	–	–	–	–
0° C	-2.098***	-0.030	-0.083	-0.051	-0.094
	(0.312)	(0.122)	(0.108)	(0.044)	(0.118)
-5° C	-2.224***	-0.141	-0.117	-0.011	-0.195
	(0.380)	(0.121)	(0.104)	(0.075)	(0.121)
Panel B: 5 - 64 years of age					
35° C	4.551***	0.017	0.019	0.089	0.035
	(0.656)	(0.110)	(0.067)	(0.182)	(0.110)
30° C	2.583***	0.057	0.034	0.039	0.069
	(0.253)	(0.065)	(0.036)	(0.081)	(0.064)
20°C	–	–	–	–	–
	–	–	–	–	–
0° C	-4.116***	-0.124*	-0.094*	-0.008	-0.126**
	(0.292)	(0.064)	(0.050)	(0.040)	(0.059)
-5° C	-4.689***	-0.116	-0.093*	-0.002	-0.115
	(0.364)	(0.079)	(0.051)	(0.056)	(0.073)
Panel C: >64 years of age					
35° C	-3.686**	4.712**	2.059	4.868***	4.855**
	(1.773)	(1.939)	(1.318)	(1.884)	(1.885)
30° C	-1.870**	2.691***	1.003*	2.446***	2.772***
	(0.770)	(0.828)	(0.587)	(0.706)	(0.800)
20°C	–	–	–	–	–
	–	–	–	–	–
0° C	8.282***	2.023***	1.751***	1.242***	1.691**
	(0.762)	(0.731)	(0.510)	(0.373)	(0.713)
-5° C	10.458***	3.431***	2.493***	2.014***	2.909***
	(0.905)	(0.959)	(0.579)	(0.523)	(0.909)
Adj R-squared	0.982	0.987	0.989	0.999	0.987
N	820697	820237	820237	819991	820237
Age×ADM2 FE	Yes	Yes	Yes	Yes	Yes
Country×Year FE	Yes	–	–	–	–
Age×Country×Year FE	–	Yes	Yes	Yes	Yes
Age×ADM1 linear trend	–	–	Yes	–	–
Precision weighting (FGLS)	–	–	–	Yes	–
13-month exposure	–	–	–	–	Yes

Standard errors clustered at the ADM1 (e.g., state) level.
Regressions in columns (1)-(3), and (5) are population-weighted.
Column (4) weights use a precision-weighting approach (see text).
*** p<0.01, ** p<0.05, * p<0.1

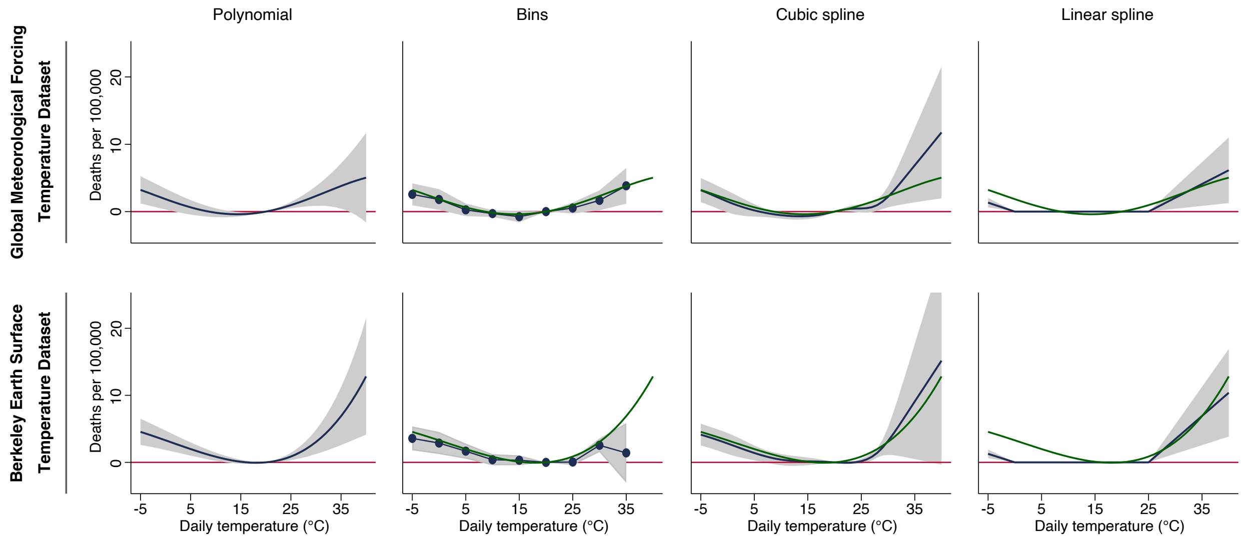

Figure D3: Robustness of the mortality-temperature relationship to alternative functional forms and to different historical climate datasets (age >64). Row 1 shows the mortality-temperature response function as estimated using daily temperature and precipitation data from the Global Meteorological Forcing Dataset (GMFD). Row 2 shows the same response, using daily temperatures from Berkeley Earth Surface Temperature (BEST), and monthly precipitation from the University of Delaware. Each column displays a distinct functional form, with the fourth-order polynomial shown in column 1 overlaid in teal on each subsequent column. See Section 4 for details on each functional form.

(4), we address the fact that some of our data are drawn from countries which may have less capacity for data collection than others in the sample. Because our mortality data are collected by institutions in different countries, it is possible that some sources are systematically less precise. To account for this, we re-estimate our model using Feasible Generalized Least Squares (FGLS) under the assumption of constant variance within each ADM1 unit.[91] In column (5), we address the possibility that temperatures can exhibit lagged effects on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). Lagged effects within and across months in the same calendar year are accounted for in the net annual mortality totals used in all specifications. However, it is possible that temperature exposure in December of each year affects mortality in January of the following year. To account for this, in column (5) we define a 13-month exposure window to additionally account for temperatures previous December.[92] Table D2 shows that the results for both of these alternative specifications are similar in sign and magnitude to those from column (2).

[91]To do this, we estimate the model in Equation D.17 using population weights and our preferred specification (column (2)). Using the residuals from this regression, we calculate an ADM1-level weight that is equal to the average value of the squared residuals, where averages are taken across all ADM2-age-year level observations that fall within a given ADM1. We then inverse-weight the regression in a second stage, using this weight. All ADM2-age-year observations within a given ADM1 are assigned the same weight in the second stage, where ADM1 locations with lower residual variance are given higher weight. For some ADM2s, there are insufficient observations to identify age-specific variances; to ensure stability, we dropped the ADM2s with less than 5 observations per age group. This leads us to drop 246 (of >800,000) observations in this specification.

[92]The specification in column (5) defines the 13-month exposure window such that for a given year t, exposure is calculated as January to December temperatures in year t *and* December temperature in year $t-1$.

D.3 Additional results: Spatial extrapolation of temperature sensitivity

Figure D4 reports on our extrapolation of mortality-temperature response functions to the entire globe for the <5 age group as well as the >64 age group shown in the main text in Figure 4. As in Figure 4, panels A and B show predicted mortality-temperature responses for each impact region for 2001-2010 average values of income and climate and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Geographic heterogeneity within our sample is shown for hot days in the maps in panels C and D, where colors indicate the marginal effect of a day at 35°C, relative to a day at a location-specific minimum mortality temperature. Grey areas are locations where mortality data are unavailable. Figure 4E–H show analogous plots, but now extrapolated to the entire globe.

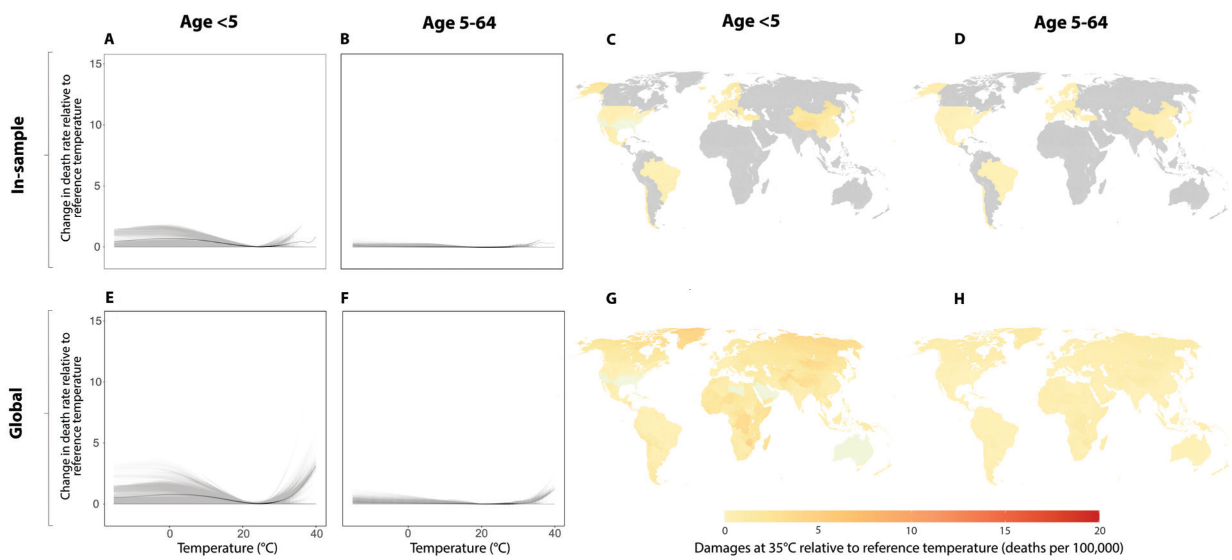

Figure D4: Using income and climate to predict current response functions globally (ages <5 and 5-64). In panels A, B, E and F, grey lines are predicted response functions for impact regions, each representing a population of 276,000 on average. Solid black lines are the unweighted average of the grey lines, where the opacity indicates the density of realized temperatures (Hsiang, 2013). Panels C, D, G and H show each impact region's mortality sensitivity to a day at 35°C, relative to a location-specific minimum mortality temperature. The top row shows all impact regions in the sample of locations with historical mortality data (included in main regression tables), and the bottom row shows extrapolation to all impact regions globally. Column titles indicate corresponding age categories. Predictions shown are averages over the period 2001-2010 using the SSP3 socioeconomic scenario and climate model CCSM4 under the RCP8.5 emissions scenario. Figure 4 shows the analogous figure for age >64.

D.4 Robustness of estimates of subnational heterogeneity in the mortality-temperature response function to an alternative characterization of long-run average climate

Our primary results rely on a parsimonious representation of the climate: to capture adaptation to long-run climate, we interact our nonlinear temperature variables ($\boldsymbol{T}$) with the long run average annual temperature ($TMEAN$), conditioning on income ($\log(GDPpc)$). In this specification, $TMEAN$ acts as a summary statistic of the long-run average climate, and we find that the mortality sensitivity to high temperatures declines as $TMEAN$ rises. To test the robustness of this finding, here we use a richer characterization of the climate, replacing our climate interaction term $TMEAN$ in Equation 5 with two interaction terms: long-run average heating degree days (HDDs), calculated relative to a 20°C threshold, and long-run average cooling degree days (CDDs), also calculated relative to 20°C. We re-estimate Equation 5 with these characterizations of average exposure to cold (HDD) and hot (CDD) days, linearly interacting each climate covariate with each element of $\boldsymbol{T}$, as is done in the main specification using $TMEAN$.

The marginal effect of each climate variable on the temperature sensitivity of mortality is shown in Table D3. Consistent with our main results in Table D1, warmer climates (as captured by higher CDDs) are associated with lower sensitivity of mortality rates to high daily temperatures. This finding is particularly true for the older age group.

Table D3: Marginal effect of covariates on temperature sensitivity of mortality rates using an HDD-CDD interaction model Coefficients (standard errors) represent the marginal effect of increasing each covariate by one unit on the temperature sensitivity of mortality, evaluated at each of the daily average temperatures shown. Temperature sensitivity is defined as the impact of a particular temperature on mortality rates, relative to a moderate day at 20°C. Regression is a fourth-order polynomial in daily average temperature, estimated using GMFD weather data with a sample that was winsorized at the top 1% level. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects. Each temperature variable is interacted with each covariate, and HDDs and CDDs are defined relative to 20°C. Standard errors are clustered at the ADM1 level.

	Age <5			Age 5-64			Age >64		
	$\log(GDPpc)$	HDD	CDD	$\log(GDPpc)$	HDD	CDD	$\log(GDPpc)$	HDD	CDD
35°	-1.07817**	0.00031	-0.00068	-0.28400*	-0.00004	-0.00030*	-4.72093**	-0.00135	-0.00431*
	(0.50360)	(0.00040)	(0.00067)	(0.15853)	(0.00009)	(0.00015)	(2.38923)	(0.00136)	(0.00254)
30°	-0.33327	0.00037**	0.00051	-0.02308	-0.00001	-0.00011	-0.29392	-0.00093*	-0.00403***
	(0.26543)	(0.00018)	(0.00035)	(0.06668)	(0.00005)	(0.00009)	(0.91042)	(0.00054)	(0.00121)
20°	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–
0°	-0.34991	-0.00061***	-0.00164**	0.06318	0.00001	-0.00031	1.58548	-0.00128	-0.01390***
	(0.49705)	(0.00021)	(0.00082)	(0.16742)	(0.00006)	(0.00029)	(2.22197)	(0.00093)	(0.00404)
-5°	-0.34121	-0.00058***	-0.00158*	0.21636	-0.00001	-0.00061	5.14874	-0.00227*	-0.02171***
	(0.57811)	(0.00022)	(0.00085)	(0.24729)	(0.00008)	(0.00038)	(3.32592)	(0.00133)	(0.00568)
Adj R-squared					0.93353				
					820237				
Age×ADM2 FE					Yes				
Age×country×year FE					Yes				

The coefficients in Table D3 determine the spatial and temporal heterogeneity in response functions that we predict at the impact region, age, and year level across the globe. To see a visual example of how this alternative model compares to our primary specification, in Figure D5 we show the slope of the response function evaluated at 35°C under the primary specification (y-axis) and the alternative HDD/CDD specification (x-axis), for each age group. Each scatter point represents one ADM1 region within our estimating sample. Consistent with Tables D1 and D3, we see that across age groups, the more nuanced characterization of the climate using cooling and heating degree days has a minimal effect on our predicted response functions.

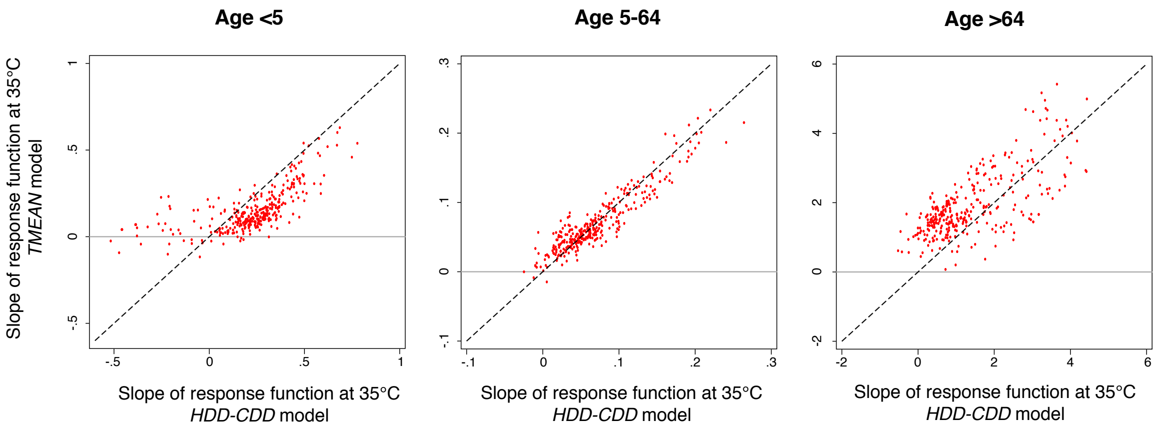

Figure D5: Predicted mortality-temperature response functions in-sample are similar under alternative characterizations of long-run average annual temperature. Each panel contains a scatter plot of the slope (i.e., derivative) of the predicted mortality-temperature response function, evaluated at 35°C, under two distinct characterizations of the long-run average climate. On the y-axis, the response function is predicted using coefficients from a version of Equation 5 in which all nonlinear temperature variables are interacted with long-run annual average temperature (this is the main specification used throughout the analysis). On the x-axis, the response function is predicted using coefficients from a version of Equation 5 in which all nonlinear temperature variables are interacted with long-run annual average heating degree days (HDDs) below 20°C and cooling degree days (CDDs) above 20°C. Predictions shown are for all ADM1 regions within our estimating sample. Each column shows predictions for a different age category.

D.5 Robustness of the mortality-temperature response function to omission of precipitation

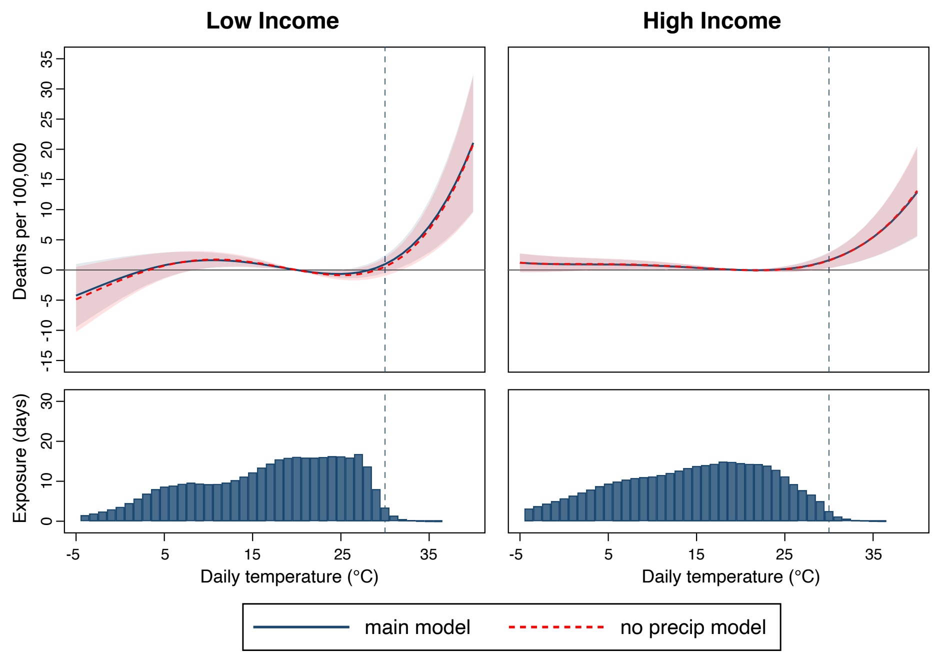

Figure D6: Predicted mortality-temperature response functions are similar with versus without precipitation controls (age >64). Each panel shows the predicted mortality-temperature relationship resulting from the estimation of a version of Equation 5, and evaluated at the population-weighted mean value of the logarithm of GDP per capita and long run average temperature within the lowest income tercile of the estimation sample (left panel), and the highest income tercile of the estimation sample (right panel). The main regression model (Equation 5) is shown in the solid blue line, while a version of Equation 5 omitting the country-specific quadratic precipitation controls is shown in the dashed red line. Vertical dashed lines indicate the 99.5^{th} percentile of the daily temperature distribution in each income group. 95% confidence intervals for both regression models are shown in the shaded areas.

D.6 Robustness of the mortality-temperature response function to inclusion of additional sources of heterogeneity

The analysis implemented in the main text relies on a two-factor model to explain heterogeneity in the mortality-temperature relationship. There are three primary reasons we use just income and long-run average climate to model heterogeneity. First, these covariates are conceptually intuitive determinants of adaptation; looser budget constraints enable more investment in adaptive technologies and behaviors, while increased exposure to a particular weather event leads to updated beliefs and corresponding adaptive investments. Second, both covariates have been shown to be important in explaining heterogeneity in climate impacts across a number of other contexts.[93] Finally, substantial research has been conducted to generate projected

[93] See, for example, Mendelsohn, Nordhaus, and Shaw (1994); Kahn (2005); Auffhammer and Aroonruengsawat (2011); Hsiang, Meng, and Cane (2011); Graff Zivin and Neidell (2014); Moore and Lobell (2014); Davis and Gertler (2015); Heutel, Miller, and Molitor (2017); Isen, Rossin-Slater, and Walker (2017)

scenarios of both of these covariates into the future, as described in Appendix B.

However, a valid critique of this model is that other factors that likely explain heterogeneity in the mortality-temperature relationship are omitted from our main estimating equation (Equation 5). In this section, we examine five additional covariates that plausibly influence the mortality-temperature relationship, but for which projections into the future are not available (potentially due to the impossibility of that exercise). First, we show that the inclusion of these additional variables into the interaction model in Equation 5 has a negligible impact on predicted mortality-temperature relationships. Second, we show that including these covariates in estimation, but omitting them when generating predictions, as would be necessary when generating a climate change impact projection, leads to substantial bias.

The five variables, all of which are only available across our sample at the national level, are:

1. Institutions. We use the polity scores from the Center for Systemic Peace (2020). This measure of the strength of democratic institutions has been widely discussed in the literature on economic growth (e.g., Glaeser et al., 2004). There are numerous ways in which the strength of institutions could moderate the mortality-temperature relationship, for example, by leading to greater responsiveness of politicians to the population of a country and so providing more public goods or healthcare access. These data are available for all countries in our sample.

2. Healthcare. In order to capture variation in the quality of healthcare across countries directly, we use the number of doctors per capita obtained from the World Development Indicators (World Bank, 2020). These data are available for all countries in our sample, although some years are missing.

3. Education. As a proxy for the education level of our sample, we use the *percent of population that have at least completed a secondary education* obtained by combining data from the World Bank (2020) and Organization of Economic Cooperaton and Development (2020). Each source has substantial gaps, with the former providing better coverage for developing countries and the latter providing better coverage for developed countries. The combination provides maximum coverage of our mortality sample. Where the sample overlaps, the data are close in both levels and trends, but a slight discrepancy arises due to differences in how the variable is defined: World Bank (2020) is defined as the measure of adults 25+ who have completed secondary education, while in Organization of Economic Cooperaton and Development (2020) it is constrained to adults 25-64. To combine these two datasets, we use the observations in which both datasets are available and regress World Bank (2020) observations on Organization of Economic Cooperaton and Development (2020) observations. We then add the recovered intercept term to the World Bank (2020) data so that average levels across the two datasets match. We then use the union of these two datasets in our regressions, averaging the two sources for observations with data available from both.

4. Inequality. We capture inequality using national-level Gini coefficients from the World Inequality Database (World Inequality Lab, 2020). Data are available for all countries in our sample with the exception of Bulgaria, Montenegro, Malta, and pre-1990 Japan.

5. Informality. It is plausible that the ability to smooth income or health shocks due to temperature exposures may be affected by the access to stable employment. Informality in the labor force has been pointed out to be an important determinant of growth across countries (e.g., La Porta and Shleifer, 2014). We use the *percent of population self-employed* from the World Development Indicators as it is widely available and is mentioned by La Porta and Shleifer (2014) as being a good proxy for informality. However, no data exist for this variable before 1991, meaning that some of our observations, primarily from the US and Japan, are omitted when this variable is used.

We combine these variables with our data in the same manner as the covariates in our two-factor model, which involves taking the time-invariant average over the period for which we have data for each country. We then run the following regression:

$$M_{aict} = g_a(\boldsymbol{T}_{it}, TMEAN_s, \log(GDPpc)_s, COVAR_c) + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait}, \tag{D.18}$$

where $COVAR_c$ is one of the five alternative covariates mentioned above. All other variables are identical to our main model in Equation 5. Due to missing data for certain countries, merging these data results in sample sizes that are in most cases smaller than our original sample. Only the institutions regression has 100% of the observations as our original sample, while health (91.5%), education (81.7%), inequality (99.7%), and informality (72.8%) all have smaller sample sizes.

We conduct two tests using the results from estimating Equation D.18 for each alternative covariate. First, we assess whether each added interaction variable actually explains substantial heterogeneity that is not already captured by our two-factor model. These results are shown in Figure D7, where we use the coefficients estimated in Equations 5 and D.18 to predict mortality-temperature response functions for our main model and for the model including the alternate covariate. We predict responses using average values of income and climate in the lowest income tercile of our data (left column) and in the highest income tercile of our data (right column). Since the sample sizes differ across variables due to data availability constraints, we re-estimate the main model using a comparable sample for each additional covariate, leading to some variation in the main model response function across rows. The daily distribution of temperatures in each income group is shown in the top row.

Figure D7 shows that differences between the main model (blue solid lines) and the alternative models (red dashed lines) are small across most of the temperature support. Some larger differences do emerge, for example in the mortality-temperature response in poorer countries when including an interaction with education, but across all covariates, none of the differences from our main model are statistically significant. Moreover, the majority of these slightly larger differences occur in the coldest 0.5% of our sample (with the middle 99% of the temperature distribution indicated by vertical dashed lines). The implication of this test is that, while the alternative covariates do in some cases explain some heterogeneity in the temperature-mortality relationship, the difference in predicted mortality-temperature relationships between our main model and these alternatives is never statistically significant and is substantively small across most of the temperature support, reaching a maximum only at the coldest temperatures, which are rare in our sample and are likely to become increasingly rare globally as the Earth warms.

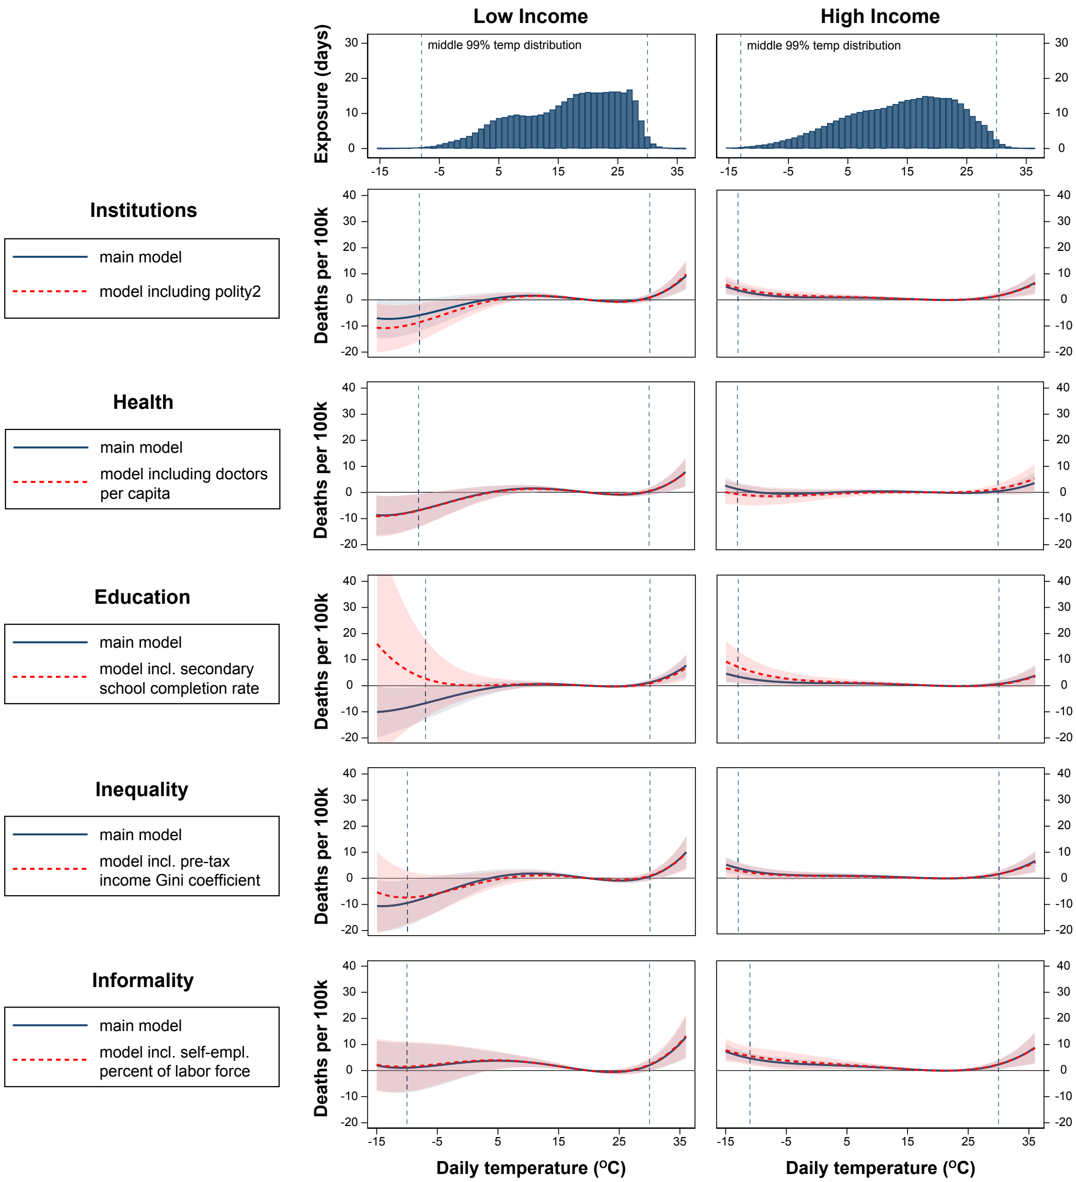

Figure D7: Predicted mortality-temperature response functions are robust to inclusion of additional interaction terms (age >64). Comparison of response functions estimated using Equation 5 ("main model", blue line) and including additional covariates using Equation D.18 ("alternative model", red dashed line) for the >64 age group. Samples are adjusted according to data availability for each covariate. The left panel shows the predicted response evaluated at the population-weighted mean level of income, climate, and the corresponding covariate across all observations that fall into the lowest tercile of GDP per capita in our estimation sample, while the right shows the corresponding predicted response for observations in the highest tercile of GDP per capita in our estimation sample. The area between the vertical dashed lines is the middle 99% of the distribution of temperatures indicated by the histogram. As samples vary slightly (see Table D4) across each comparison, the middle 99% varies slightly across figures, as does the main model line. Results for other age groups are similar.

Second, we perform a test to understand the consequences of explicitly modeling heterogeneity for a variable that does not have values projected into the future. To do so, we predict mortality rates in-sample and compute the Root Mean Squared Error (RMSE) using a model in which an additional covariate is included in estimation, but not in prediction, as would be necessary when generating future climate change impact projections. We compare these RMSE values to those from our two-factor model, where both interaction terms can be included in climate change impact projections.

In Table D4 the RMSE values are shown for both cases. The "Main Model" column shows the RMSE from estimation of Equation 5 using the sample that varies slightly across rows based on data availability of each of the alternate covariates. The "Alternative Model" column shows the RMSE from estimation of Equation D.18 for each additional covariate, where predicted values are generated by setting the coefficient on the alternate covariate to zero. This second case mimics the situation in which we estimate a model with an interaction term for a variable that cannot be projected into the future. The key comparison is between the RMSE for the "Main Model" and the RMSE for the "Alternative Model", shown in the final "Difference" column. For each of these additional covariates, the alternative model exhibits a substantially worse model fit (i.e., the "Difference" column is negative), implying that the inclusion of a covariate in estimation for which there are no data in the future would lead to a model that performs strictly worse than the two-factor model used throughout the paper.

Table D4: Evaluating omission of additional sources of heterogeneity in climate projections. Each row corresponds to model performance metrics from the estimation of Equation D.18 with the inclusion of the named covariate. All covariates other than those in the main model (Equation 5) are observed at country level. Sample sizes differ across rows due to the availability of data for each of the covariates, with sample sizes ranging from 73% to 100% of our main estimating sample. The "RMSE Main Model" column shows the in-sample root mean squared error (RMSE) of our main model estimated using Equation 5. The "RMSE Alternative Model" column shows the in-sample RMSE of a model that is estimated using Equation D.18, but where predictions are generated omitting the impact of the additional covariate. This is done to mimic a situation in which estimation includes historical data on additional determinants of heterogeneity, but climate change projections must be made without projection data for that additional covariate. The differences, all negative, are shown in the difference column and indicate that the RMSEs of our main estimation equation are consistently lower than those for the alternative models.

Model	**Covariate**	**Observations**	**Proportion of Full Sample**	**RMSE Main Model**	**RMSE Alternative Model**	**Difference**
Institutions	Polity 2 Score	820,237	1.000	565.19	567.26	-2.07
Health	Doctors per Capita	750,486	0.915	572.87	574.68	-1.80
Education	Secondary School Completion Rate	670,454	0.817	518.62	559.34	-40.72
Inequality	GINI Coefficient	817,744	0.997	566.03	575.13	-9.10
Informality	Self Employed percent of LF	597,059	0.728	602.65	607.42	-4.76

Sources: Center for Systemic Peace (Polity2), World Development Indicators (Doctors Per Capita, Secondary Completion Rate, Self Employment), World Inequality Database (GINI), OECD (Secondary Completion Rate)

D.7 Cross-validation to assess out-of-sample performance

Throughout our analysis, we use coefficients estimated from Equation 5 in the main text, in combination with local-level observations and projections of $TMEAN$ and $\log(GDPpc)$, to generate predicted response functions in all regions of the world, including where mortality data are unavailable, both in the present and into future (see Section 5.2 for details). In contrast, much prior work generates projected impacts of climate change using spatially and/or temporally homogeneous response functions (e.g., Hsiang et al., 2017; Deschênes and Greenstone, 2011). To assess the performance of our model in predicting mortality-temperature relationships out-of-sample, in this section we implement multiple custom cross-validation exercises designed to mimic the spatial and temporal extrapolation that is required when using available historical data to generate global climate change projections decades into the future.

We perform three cross-validation exercises, each of which provides multiple measures of the out-of-sample performance of Equation 5. In each case, we compare these performance metrics to the performance of a benchmark model that ignores adaptation, and to a measure of in-sample model fit. Because we are assessing the performance of our interaction model in predicting mortality sensitivity to temperature, as opposed to mortality rates overall, all measures of model fit are reported using residualized data, in which all fixed effects and controls are removed from identifying variation prior to estimation.[94] Results from all three tests are shown in Table D5, and are discussed in the following subsections.

D.7.1 K-fold cross-validation with spatial blocking

First, we conduct a standard k-fold cross-validation analysis with 10 folds. Because of the panel structure of our data and because we use $ADM1$ level climate and income variables to determine mortality sensitivity to temperature (see Equation 5), we spatially block when defining these ten folds, ensuring that all $ADM2 \times year \times age$ observations that fall within the same $ADM1$ are removed from the sample in the same fold. This ensures that the strong serial and spatial correlation between observations within an $ADM1$ does not artificially inflate our measure of out-of-sample performance. Panel B of Table D5 shows the root mean squared error (RMSE) from our main interaction model, Equation 5, as well as from a benchmark model that does not account for any form of adaptation. These results show that our interaction model strongly out-performs this benchmark model (RMSE decline of 12.62). Moreover, our interaction model obtains high out-of-sample performance when compared to the in-sample RMSE shown in Panel A (RMSE increase of 0.15).

D.7.2 Cross-validation using blocking by covariate values

Second, we design a custom cross-validation analysis that systematically removes blocks of data based on long-run climate and income, the two covariates determining mortality sensitivity to temperature in Equation 5. This exercise is designed to mimic the spatial extrapolation we conduct to generate mortality-temperature relationships in locations without mortality data, based solely on their long-run climate and income (see

[94] Following the main specification in the paper, we remove $age \times ADM2$ fixed effects, $age \times country \times year$ fixed effects, as well as country-specific quadratic precipitation controls.

Sample	Observations	% of global population (2010)	% of global population (2100)	RMSE (adaption model)	RMSE (no adaption model)	RMSE difference
A: In-sample model fit						
Full Sample In-Sample	820,698	–	–	565.19	577.80	-12.62
B: 10-fold cross-validation at ADM1 level						
Full Sample Out-of-Sample	820,698	–	–	565.34	577.97	-12.62
C: 9-fold cross-validation across income × long-run temperature covariate space						
Full Sample Out-of-Sample	820,698	–	–	565.51	577.88	-12.37
By held-out block						
Low Income - Cold	4156	6.5	0.0	367.46	372.69	-5.23
Low Income - Moderate	15,279	9.0	0.5	644.94	650.49	-5.55
Low Income - Hot	334,968	65.5	46.0	534.72	546.21	-11.49
Middle Income - Cold	3507	1.5	0.0	571.85	566.54	5.31
Middle Income - Moderate	15,108	1.0	0.0	554.08	552.70	1.38
Middle Income - High	78,160	2.0	24.5	530.68	530.93	-0.25
High Income - Cold	125,934	5.0	2.0	617.93	643.54	-25.61
High Income - Moderate	137,706	5.0	5.0	593.86	605.36	-11.50
High Income - Hot	105,880	4.0	22.0	577.63	588.77	-11.14
D: 2-fold cross-validation across time (post-2004 hold-out)						
Pre-2005 In-Sample	607,979	–	–	565.72	577.85	-12.13
Post-2004 In-Sample	212,719	–	–	563.72	578.38	-14.66
Post 2004 Out-of-Sample	212,719	–	–	564.00	578.36	-14.36

Table D5: Evaluation of out-of-sample model performance This table presents results from three separate cross-validation exercises. Panel A shows the in-sample model performance of the interaction model in Equation 5 in the main text. For panels B, C, and D, a section of the data is omitted from the sample and Equation 5 is re-estimated using the remaining observations. The mortality rates of the omitted observations are then predicted out-of-sample, and the root mean squared error (RMSE) is calculated and shown in the column titled "RMSE (adaptation model)". This process is repeated, but with a regression equation that omits any model of heterogeneity; RMSE values for this model are shown in the column titled "RMSE (no adaptation model)". Differences between the model with adaptation (i.e., Equation 5) and the model without adaptation are shown in the "RMSE difference" column, with negative values indicating that our main estimating equation performs better than the benchmark model without adaptation. Panel B reports results from k-fold cross-validation, panel C from a custom cross-validation that blocks data based on long-run income per capita and average temperatures, and panel D from a custom cross-validation that divides the data into pre-2005 and 2005-2010 samples. All results are shown using residualized data in which $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects were removed from all variables before cross-validation was conducted. See text for details.

Section 5.2 for details). To conduct this analysis, we split our sample into 9 "blocks", based on the tercile of the long-run climate and income distributions that each observation falls into.[95] For example, one block corresponds to observations in the lowest income tercile and the coldest climate tercile, while another block

[95] Because long-run climate and income are observed at the $ADM1$ level, terciles are defined using the distribution over unique $ADM1$s. We compute these distributions independently for each variable.

corresponds to observations in the middle income tercile and the hottest climate tercile. We then repeatedly remove each block of the data, re-estimate Equation 5 using the remaining 8 blocks of data, and generate out-of-sample predictions for the removed block. In Panel C we show out-of-sample RMSE results for each block, as well as for the full sample.

We draw three main conclusions from the results in Panel C of Table D5. First, the interaction model overall performs well when predicting mortality rates in locations where income and climate fall outside the estimation sample ranges. The full-sample out-of-sample RMSE from this exercise is only slightly higher than the in-sample RMSE (RMSE increase of 0.32), which relies on the full dataset for estimation. Additionally, this exercise produces only a slightly higher RMSE than the standard k-fold cross-validation shown in Panel A (RMSE increase of 0.17). Note that this climate and income blocking exercise is much more challenging than k-fold cross-validation, as entire portions of the income and climate distributions are omitted from estimation. Second, the interaction model out-performs the no-adaptation benchmark model overall, and in 7 out of 9 blocks of the data, encompassing 100% of global population at 2100. This indicates that the model in Equation 5 provides substantially more accurate mortality rate predictions than a model ignoring heterogeneity based on income and long-run climate. Finally, the interaction model performs particularly well in the low income and hot climate block, conditions that will be experienced by 46% of the global population in 2100 under SSP3. We further demonstrate predictive power in hot and low income conditions in another out-of-sample exercise described in Section D.8 below.

D.7.3 Cross-validation using temporal blocking

Our third out-of-sample analysis assesses our model's ability to predict mortality and mortality-temperature sensitivity into the future. To conduct this analysis, we split our sample into observations that fall before the year 2005, and those after (and including) 2005. We choose this cutoff because it ensures a roughly balanced set of countries in both samples (see Figure B1). We then use the pre-2005 data to estimate the interaction model in Equation 5, and use these regression results to predict mortality rates in the 2005-2010 data. Results are shown in Panel D of Table D5. These RMSE comparisons show that our interaction model performs well when predicting residualized mortality rates in years that were not used in estimation, showing little RMSE increase relative to an in-sample estimation using 2005-2010 data only (RMSE increase of 0.28), and relative to the full sample in-sample RMSE (RMSE increase of 1.19). The interaction model substantially out-performs a model without adaptation in this out-of-sample test (RMSE decline of 14.36).

D.7.4 Visualizing out-of-sample performance

In addition to the tabular results shown in Table D5, here we show visualizations of the predictive performance of the interaction model for the two cross-validation approaches based on spatial or temporal blocking. Specifically, we display the difference between out-of-sample predicted mortality-temperature response functions and estimates of those same response functions using a subset of the sample data. While the tabular results of out-of-sample performance in Table D5 indicate overall measures of fit, these figures help visualize when and to what extent the model accurately extrapolates mortality-temperature sensitivity in different portions of the sample. This exercise differs from the RMSE exercise shown above in that the comparison is between two estimated response functions, neither of which is known with certainty. In contrast, the RMSE table reports error between the true mortality rate observations and predicted mortality rates.

First, Figure D8 plots the differences in the >64 mortality-temperature relationship between an out-of-sample prediction and an "in-sample" estimation using a subset of the data, for four of the 9 blocks in the spatial blocking exercise described above. The solid red line indicates the difference between: *(i)* the mortality-temperature response function predicted for each block (e.g., high income and cold climate in the upper left) by the estimation of Equation 5 using the remaining 8 blocks of data, and *(ii)* the mortality-temperature response function estimated using a model without interactions with data from within that block alone. 95% confidence intervals on the difference are shown in the shaded red area,[96] and vertical dashed lines indicate the middle 99% of the daily temperature distribution within each block.[97]

Figure D8 demonstrates that the differences between in-sample and out-of-sample response functions are rarely statistically significant for the >64 age group (these results are similar for other age groups). Statistically significant differences only arise at the extreme cold end of the temperature distribution, often in blocks where those temperatures are rarely realized (e.g., cold temperatures in the hot climate blocks). Moreover, for the majority of observed temperatures, the magnitude of these differences are small relative to the overall mortality-temperature response function (see Figure 1).

Analogous to Figure D8, the temporal extrapolation performance of our interaction model is shown in Figure D9, which plots the differences in the >64 mortality-temperature relationship between an out-of-sample prediction and an in-sample estimation for the lowest and highest income terciles of the full sample income distribution. The solid red line indicates the difference between: *(i)* the mortality-temperature response function predicted for each income tercile by the estimation of Equation 5 using pre-2005 data only, and *(ii)* the mortality-temperature response function estimated using a model without interactions with data from 2005-2010 only. The 95% confidence intervals on the difference are shown in the shaded red area,[98] and vertical dashed lines indicate the 99.5^{th} percentile of the daily temperature distribution within each income group. As above, this exercise differs from results shown in Table D5 in that the comparison is between two estimated response functions, neither of which is known with certainty.

Figure D9 demonstrates that, for the >64 age group, adaptation to heat is over-estimated in low income regions (left panel) and under-estimated in high income regions (right panel). That is, using data from before 2005 leads to predicted mortality sensitivity to temperature in 2005-2010 that is too *low* in low income regions and too *high* in high income regions, relative to an in-sample estimate. These differences are statistically significant for the low income group and of modest size (∼5 deaths per 100,000 per 30°C day). In contrast, differences between out-of-sample predicted response functions and in-sample estimated response functions are small and statistically insignificant for other age groups. To show the implications

[96]Standard errors on the difference between response functions are calculated by running the interaction model for the 8 blocks of data and the uninteracted model for the remaining 1 block of data in a stacked regression saturated with block-level indicators.

[97]Note that vertical dashed lines are omitted from Figure D8 if the 0.5^{th} or 99.5^{th} percentile of the data do not fall between -5°C and 35°C.

[98]Standard errors on the difference between response functions are calculated by running the interaction model for the pre-2005 data and the uninteracted model for the 2005-2010 data in a stacked regression saturated with pre-2005 indicators.

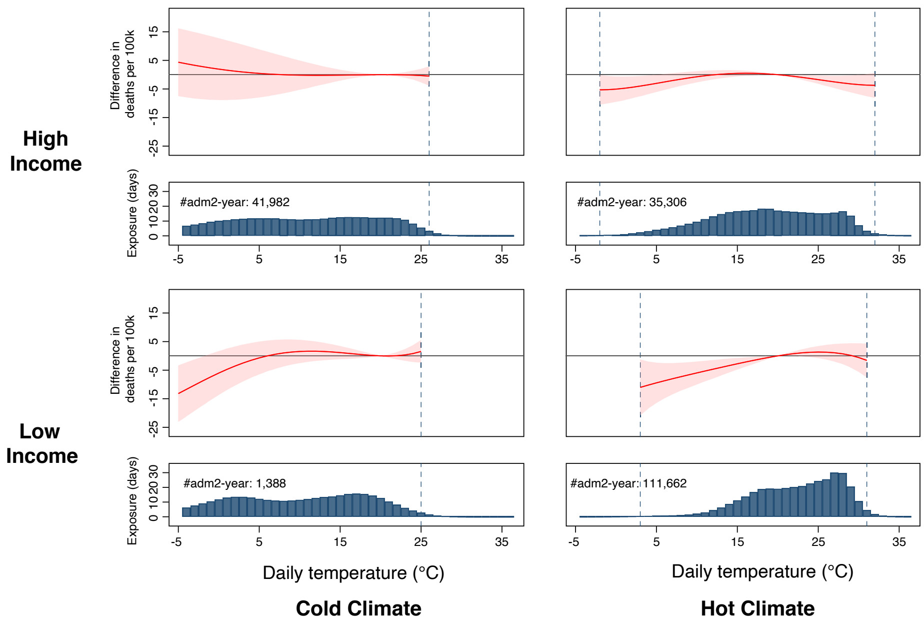

Figure D8: Differences between spatial out-of-sample predicted mortality-temperature response functions and in-sample estimated response functions (age >64mortality rate). Each panel shows the difference between *(i)* an out-of-sample predicted mortality-temperature relationship resulting from the estimation of a version Equation 5 using all data *except* observations falling within the income and climate "block" indicated (e.g., high income and cold, low income and hot), and *(ii)* an in-sample estimated mortality-temperature relationship resulting from the estimation of a similar model without interactions using data from within the income and climate block alone. Positive values indicate the out-of-sample prediction overestimates mortality sensitivity to temperature relative to an in-sample estimation (and vice versa for negative values). Histograms show the distribution of daily temperature for all locations falling within the indicated block and vertical dashed lines indicate the middle 99% of the daily temperature distribution within each block. 95% confidence intervals are computed by running the out-of-sample and in-sample regressions in a stacked regression model saturated with block-level indicators. Results for other age groups are similar.

of both under- and over-estimating rates of adaptation, we show in Section F.4 sensitivity of our projected mortality impacts of climate change to assumptions about the rate of adaptation over the 21^{st} century.

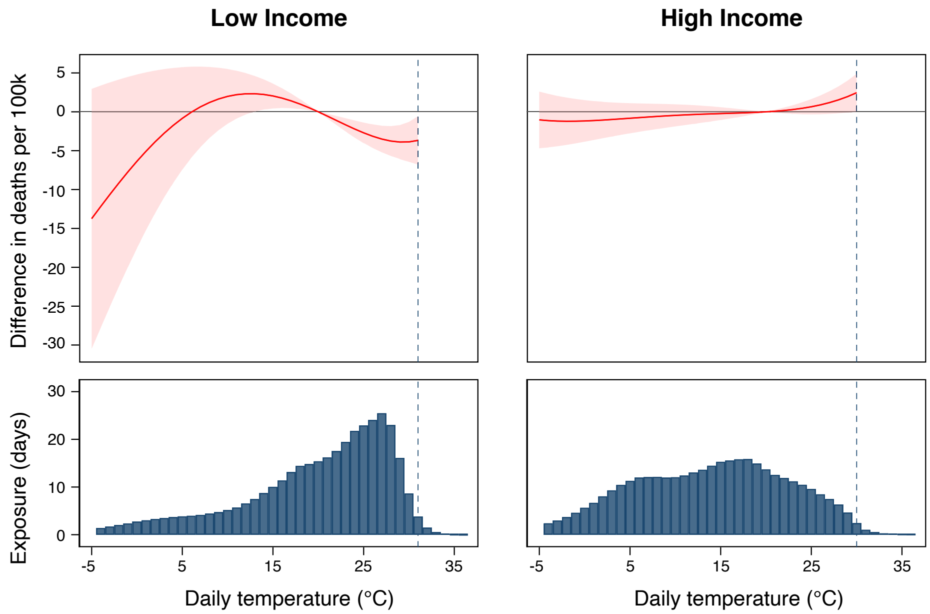

Figure D9: Differences between temporal out-of-sample predicted mortality-temperature response functions and in-sample estimated response functions (age >64 mortality rate). Each panel shows the difference between *(i)* an out-of-sample predicted mortality-temperature relationship resulting from the estimation of Equation 5 using on all data before 2005, and *(ii)* an in-sample estimated mortality-temperature relationship resulting from the estimation of a similar model without interactions using data from 2005-2010 alone. Histograms show the distribution of daily temperature for all locations falling within the indicated block and vertical dashed lines indicate the 99^{th} percentile of the daily temperature distribution within each block. 95% confidence intervals are computed by running the out-of-sample and in-sample regressions in a stacked regression model saturated with block-level indicators. Results for other age groups show smaller and statistically insignificant differences between out-of-sample and in-sample response functions.

D.8 Replication of Burgess et al. (2017) and out-of-sample model validation in India

As discussed above in Appendix D.7, the accuracy of the spatial and temporal extrapolation of response functions conducted in the main text depends in part on the representativeness of the observed sample. In the ideal case, we would have data for countries that cover the full distribution of income and climate, but as shown in Figure 3 in the main text, our observed sample lacks coverage for the poorest and hottest regions of the global income-climate distribution. The results in Appendix D.7 provide some confidence that we are able to extrapolate our estimates to the poorer and hotter regions of the world within our sample. However, a reasonable concern is that countries with subnational, age-specific mortality rate data may be different than the countries without such data. These are often developing countries where mortality data are either inconsistently collected or not collected at all. To address this concern, here we expand upon the cross-validation experiments shown above to test how our model performs in a region that is both lower income and hotter than our estimation sample. India represents the poorest and hottest country for which we have been able to obtain mortality records, and therefore provides an important check on the extrapolation performance of our interaction model.

To execute this validation test, we use data from Burgess et al. (2017) and compare mortality-temperature relationships estimated using these data to those predicted for India from our main estimating equation, Equation 5. The primary reason that we do not use data from India in our main estimation is that they do not contain age-specific mortality rates, and we show that age at death is a key source of heterogeneity (see Figure 2).

We begin by estimating a version of the main specification in Burgess et al. (2017):

$$M_{it} = f(T_{it}) + q(R_{it}) + \alpha_i + \lambda_r^1 t + \lambda_r^2 t^2 + \varepsilon_{it}. \tag{D.19}$$

The outcome M_{it} is the all-age mortality rate for district i in year t, which we estimate in levels rather than in logs, as in Burgess et al. (2017), to ensure direct comparability with our main specification. To generate comparable functional forms, we estimate Equation D.19 using a fourth-order polynomial, denoted by $f(T_{it})$, as we have used this as our main specification (Burgess et al. (2017) use a binned temperature specification). We control for precipitation, denoted by $q(R_{it})$, identically to Burgess et al. (2017) via a set of three dummy variables, each of which takes the value of 1 when total annual rainfall in district i and year t falls within each of three location-specific rainfall terciles. Due to the negligible effect of precipitation on our estimates of the mortality-temperature relationship (see Figure D6), this choice makes little difference. Following Burgess et al. (2017), we also include a set of district fixed effects, α_i, and linear and quadratic trends for each "climate region" (of which the authors note there are four separate regions) r of India. Observations are weighted by district population and standard errors are clustered at the ADM2 level.

We then use the estimates from Equation 5 to predict the mortality-temperature relationship across India. To do so, we first predict mortality-temperature response functions for India for each of the three age groups in our main estimation, using the population-weighted average values of ADM1-level incomes and average temperatures across the country. Once we have the predicted age-specific national responses, we take the age-weighted average of these response functions to generate an all-age average mortality-temperature response function across India. We compute this average using age-specific population values from the year 2015, which are available in the Burgess et al. (2017). We cluster standard errors at the ADM1 level (in India, this is equivalent to the state level), as in all our specifications throughout the main text. For the purposes of assessing out-of-sample performance of our main model as it compares to alternative models estimated in the literature, we also predict an all-India response function using estimates for a version of Equation 5 that models only heterogeneity in long-run income and one that models only heterogeneity in long-run average temperature.

Figure D10 shows the result of this replication and out-of-sample validation exercise for India. The figure compares our predicted responses in India (in blue) to the mortality-temperature response estimated using India's data alone (in red), following Burgess et al. (2017). Our model performs well, despite containing no information on Indian mortality rates: for the hotter end of the response function, where much of the low income world resides, our prediction is, if anything, conservative in extrapolating out-of-sample. Included in the figure are two dashed blue lines which show the predicted mortality-temperature relationship using estimates from models with only one or the other of our two interaction terms. The model using estimates from Equation 5 replicates the country model more closely than both alterative models.[99]

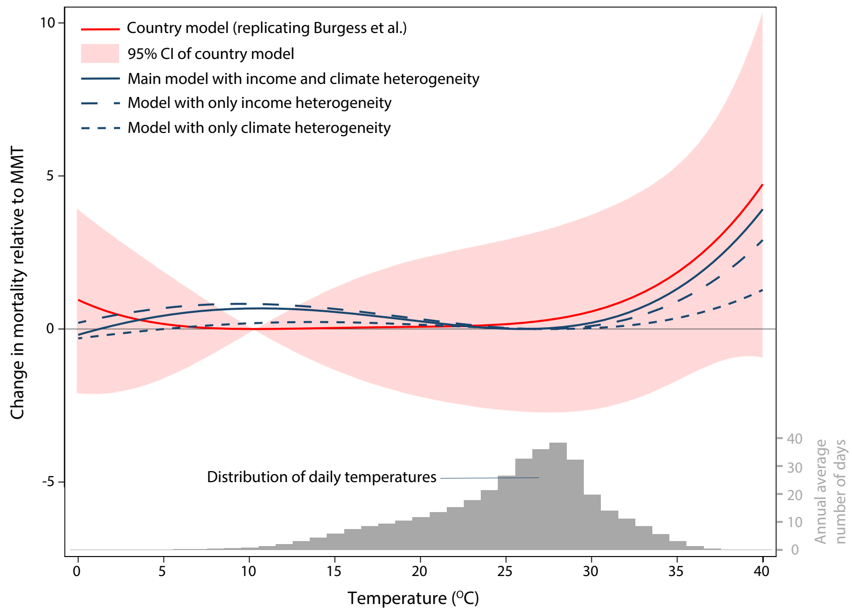

Figure D10: Out-of-sample validation of the mortality-temperature response function in India. Dark blue lines indicate out-of-sample predicted response function using coefficients from the interaction model in Equation 5, as well as versions interacting with only income (long-dashed line) and only climate (short-dashed line). The red line is estimated following Burgess et al. (2017) using all-age mortality data for India, as described in the text. The relative congruence between red and solid dark blue lines shows that our interaction model generates reasonable predicted response functions in the poorest and hottest regions of the world, the subset of the covariate space for which the main estimating sample has the least representation. Note that all curves are centered on their respective minimum mortality temperatures, as we use these curves to compute predicted deaths below in Table D6, and all predicted deaths reported in the paper use location- and model-specific minimum mortality temperatures.

In addition to the comparison of response functions in Figure D10, Table D6 reports quantitative differences in predicted mortality rates across the four models. This table summarizes the quantitative difference between the response functions shown in Figure D10. Mortality rates are predicted in-sample for each model by taking the product of the historical distribution of temperature exposure with the response function.[100] Table D6 shows that predicted mortality rates follow the same pattern as the response functions in Figure D10, with all out-of-sample model predictions falling below Burgess et al. (2017) predicted mortality rates. These out-of-sample predictions under-estimate in-sample predicted mortality rates by 37% (full interaction model in Equation 5), 42% (interaction model with income only), and 78% (interaction model with long-

[99]Note that in Figure D10, each response function is centered such that the predicted change in mortality is zero at the value of its minimum mortality temperature (MMT). This recentering is arbitrary, as the inclusion of fixed effects in Equation D.19 implies that the level of the response function is not recoverable; only slopes are causally identified. However, we choose this approach to match the quantitative exercise shown in Table D6, which uses the MMT as the reference temperature in order to mimic the climate change projection approach we take throughout the main text.

[100]These values should be interpreted as mortality rates relative to a temperature distribution in which the minimum mortality temperature (MMT) is experienced every day. This is analogous to the climate change impact projections we conduct in Section 5, where all impacts are reported relative to location-specific MMTs.

run temperature only). However, all predicted values fall within the large confidence interval given by the in-sample model's prediction. We conclude that our estimates compare well to an in-sample model in India both quantitatively and qualitatively, while noting that the out-of-sample values tend to be conservative.

Model	Predicted mortality	Lower 95% CI	Upper 95% CI	Difference	% difference
Burgess et al. (2017)	119.67	-867.85	1107.20	0.00	0.00
Full interaction (Eq. 5)	74.87	-8.26	158.00	-44.80	-37.44
Income-only interaction	69.22	-6.96	145.40	-50.45	-42.16
Climate-only interaction	26.15	-4.37	56.66	-93.53	-78.15

Table D6: Evaluation of differences in predicted mortality rates when using an in-sample estimation for India versus an out-of-sample predicted response function. Each row refers to a different empirical model of the mortality-temperature relationship in India. The Burgess et al. (2017) model in the first row is estimated following Burgess et al. (2017) using all-age mortality data from India. In all three remaining rows, the mortality-temperature relationship is predicted from a form of Equation 5, relying on data from 40 other countries; no Indian data are used. Predicted mortality represents total deaths per 100,000 per year that are attributable to historical temperature variation. Differences indicate the difference between out-of-sample predicted mortality rates from models in rows 2-4, relative to the model in row 1. Out-of-sample predicted mortality in rows 2-4 are smaller than predicted mortality with the India data, but within the confidence interval.

E Implementation of projection of future adaptation and benefits of income growth

In the main analysis, our estimates of the full mortality risk of climate change account for both the benefits and the costs of adaptation, as well as the benefits of income growth. In this appendix, we provide details on our implementation of adaptation and income benefits in future climate change projections. In Appendix E.1 we detail the procedure we use to determine the temporal dynamics of income effects on the mortality-temperature relationship in future years, in Appendix E.2 we describe the assumptions we impose on the process of adaptation and income benefits over the course of the 21^{st} century, and in Appendix E.3 we show a visual example of how the the mortality-temperature relationship is projected to change over time.

E.1 Determining the temporal dynamics of income effects

We estimate the relationship between long-run average climate, average income, and mortality-temperature sensitivity via the estimation of Equation 5 using cross-sectional variation in climate and income in combination with year-to-year variation in daily average temperatures. In generating future projections of climate change impacts (i.e. results in Section 5.5), we apply the estimated coefficients from Equation 5 over time, allowing impact region response functions to evolve as the climate warms and incomes grow. To do so, we must make an assumption regarding the rate at which the income and average climate covariates update. Here, we detail how we define this speed of adjustment in the case of income growth. While we can derive a duration over which updating occurs in the case of income due to substantial time series variation in incomes in our observed data, the historical trends for temperature have been small to date, making a similar derivation infeasible. Thus, for the case of updating based on long-run average climate, we use the standard definition of "climate" and assume a duration of 30 years.

In future projections, we estimate impact region response functions using time-varying measures of $\log(GDPpc)_{rt}$ (see Section 5.3 for details):

$$\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt}).$$

The temporal structure of the covariate $\log(GDPpc)_{rt}$ mediates the rate of income-based adaptation. If the income covariate were held fixed at historical levels, no income-based adaptation would be implemented. At the other extreme, if the contemporaneous income for year t were applied in each year, then changes in income would be assumed to translate into immediate changes in mortality-temperature sensitivity. This case is also implausible, as benefits of income are likely to take multiple years to manifest, as richer governments and citizens invest in adaptive capital and enjoy greater health. To allow for this intermediate case, we construct the income covariate used for future projections with a weighted average of recent year incomes, according to a Bartlett kernel. Specifically, to calculate the covariate $\log(GDPpc)_{rt}$, we compute:

$$\log(GDPpc)_{rt} = \frac{\sum_{s=1}^{L}(L-s+1)\log(G\mathring{D}Ppc)_{r,t-s}}{\sum_{s=1}^{L}(L-s+1)}$$

where L is the total number of lags (in years) and $\log(\mathring{GDPpc})_{rt}$ is the instantaneous log income for region r in year t.

To find a plausible length L for the Bartlett kernel, we study changes in the response of mortality for people over 64 to temperature in the United States, where we have access to a long panel of mortality rates and income data (1968 to 2010). First, we estimate the following model:

$$M_{ait} - M_{ai,t-1} = \boldsymbol{\beta}_t \left[\boldsymbol{T}_{it} - \boldsymbol{T}_{i,t-1}\right] + q_a(\boldsymbol{R}_{it}) + \varepsilon_{it} \tag{E.20}$$

where M_{ait} is the mortality rate for region i in period t and age group $a > 64$, $\boldsymbol{T}_{it}$ is the vector of polynomials of daily average temperatures (up to the fourth order), $\boldsymbol{R}_{it}$ is the vector of cumulative monthly precipitation (up to the second order), as in the main text (see Equation D.17). Coefficients are estimated for the difference between each pair of years in order to remove the year fixed effect. This produces a series of coefficients, $\boldsymbol{\beta}_t$, and their standard errors, $\boldsymbol{\sigma}_t$. We then use a Bayesian model to estimate the length of the Bartlett kernel that best explains the change in these coefficients over time. Under the model, each coefficient β_{pt} of vector $\boldsymbol{\beta}_t$ is a draw from a Gaussian distribution with a mean that varies with national average income. That is,

$$\beta_{pt} \sim \mathcal{N}(\theta_p + \phi_p \log(GDPpc)_t, \tau_p + \sigma_{pt})$$

In this model, θ_p and ϕ_p correspond to the uninteracted and income-interacted coefficients from our standard model in Equation 5, respectively. τ_p is a hyper-parameter which controls the rate of pooling of the data; if it is 0, inverse-variance weighting is used across individual year estimates.

The covariate $\log(GDPpc)_t$ is calculated as a Bartlett kernel over a maximum of 25 years of delayed income levels. National real income data from the U.S. Bureau of Economic Analysis is used to construct $\log(GDPpc)_t$. The kernel is characterized by the unknown lag parameter L, which is also estimated by the model. The maximum likelihood estimate for the Bartlett kernel length is 13 years, with a 95% confidence interval of 9.7 years. We therefore use a Bartlett kernel of length 13 when constructing the income covariate used to predict future response functions for all impact regions in all years and for all age groups.

E.2 Adaptation constraints imposed in the projection of climate change impacts

As discussed in Section 5.2, we impose two assumptions when applying our econometrically-derived model of adaptation to generate projections of future climate change. These assumptions are guided by economic theory as well as the physiological literature and are used to ensure plausible out-of-sample projections over the 21^{st} century. Graphical intuition for these constraints is shown in Figure E1.

Assumption #1: Weak monotonicity. A large body of epidemiological and econometric literature has recovered U-shaped relationships between mortality rates and daily temperatures, where both extreme cold and extreme heat increase the risk of death. These parabolic response functions have been recovered in studies using a wide range of functional form assumptions (e.g., binned daily temperatures, restricted cubic splines, or polynomials) and across diverse locations globally (e.g., Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011). As shown in Section 6, we also recover U-shaped relationships

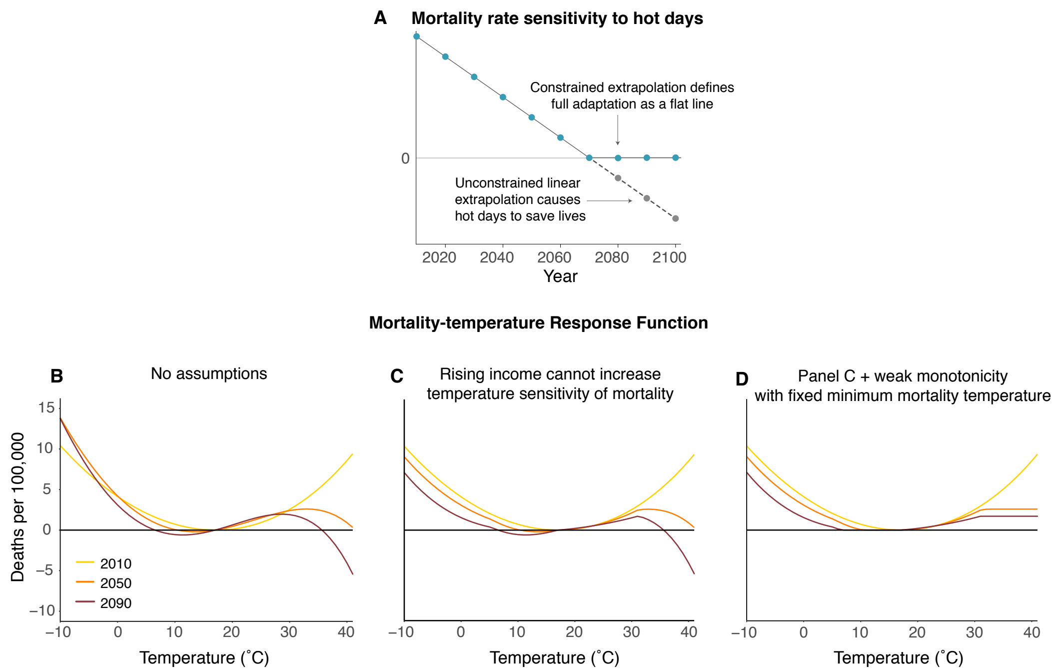

Figure E1: Two assumptions imposed in climate projections ensure that full adaptation is defined as a flat-line response function and that responses conform to basic physical and economic constraints. Panel A demonstrates heuristically the importance of imposing assumptions on the shape of response functions under adaptation over the 21^{st} century. As shown, linearly declining mortality rate sensitivity to hot days occurs over the course of the century as populations adapt. However, linear extrapolation can lead to mortality benefits on hot days, as shown with the dashed line and grey dots. Our assumptions (shown in teal) ensure that full adaptation is realized when hot days impose zero additional mortality risk. Panels B through D represent an empirical example of how the imposition of these constraints can change the shape of the adapted response function, for the Chicago, Illinois impact region. Panel B has no assumptions, panel C imposes the assumption that income is weakly protective, and panel D imposes the assumption of weak monotonicity around a time-invariant minimum mortality temperature (MMT).

between mortality rates and daily temperatures across our multi-country sample. In our projections of future mortality responses to daily temperature, we ensure consistency with this literature and with our own estimates from historical data by imposing the constraint that the response function must remain weakly monotonic around an empirically estimated minimum mortality temperature. That is, we assume that temperatures farther from the minimum mortality temperature (either colder or hotter) must be at least as harmful as temperatures closer to the minimum mortality temperature.

To implement this assumption, we first identify a range of physiologically optimal temperatures. Drawing on extensive research across epidemiology and medicine (e.g., Curriero et al., 2002; Guo et al., 2014), as well as ergonomics (e.g., Seppanen, Fisk, and Lei, 2006; Hancock, Ross, and Szalma, 2007), we let this range of possible minimum mortality risk cover the temperatures 10°C to 30°C. We then search, within this range, for the temperature at which the location-specific response function in each impact region r in the baseline years of 2001-2015 is minimized. Because distinct populations may differ substantially in the temperature at which mortality is minimized,[101] it is important to note that we allow these minimum mortality temperatures

[101] E.g., Guo et al. (2014) demonstrate that mortality risk is smallest around the 75th percentile of local temperatures in 12

(MMTs) to be spatially heterogeneous. With these optimal temperatures in hand, we impose the assumption that mortality rates must remain weakly increasing in daily temperatures to both the left and the right of this minimum. To operationalize this, we calculate impacts along an adjusted response function that is defined as the cumulative maximum to the right and left of the minimum mortality temperature along each region- and year-specific response function derived from our response surface estimated in Equation 5. Consistent with prior literature (Heutel, Miller, and Molitor, 2017; Curriero et al., 2002; Gasparrini et al., 2015), we find that these minimum mortality temperatures are highly correlated both with both long-run average temperature (positively) and with income (negatively).

This assumption is important because Equation 5 parameterizes the flattening of the U-shaped response function such that, with enough warming or sufficiently high income, the mortality-response function could become an inverted-U-shape. This is guaranteed to occur mechanically at some future date, as a result of extrapolating response functions out of the support of historically observed data. To avoid this unrealistic behavior, we impose weak monotonicity. An example of this assumption in practice is given in panel E of Figure E1.[102]

In imposing the weak monotonicity constraint, we fix the MMT at its baseline level in 2015 for each impact region. We do so because the use of spatial and temporal fixed effects in Equation 5 implies that response function levels are not identified; thus, while we allow the *shape* of response functions to evolve over time as incomes and climate change, we must hold fixed their *level* by centering each response function at its time-invariant MMT.[103]

Assumption #2: Rising income cannot increase the temperature sensitivity of mortality. We assume that because increased income per capita strictly expands the choice set of individuals considering whether to make adaptive investments, future increases in income cannot raise the impacts of temperature on mortality rates. While we place no restrictions on the cross-sectional effect of income on the temperature sensitivity as estimated in Equation 5, we do not allow any income gains through time to raise the marginal effect of temperature on mortality. Note that this condition will only be binding if the marginal effect of income estimated in Equation 5 is positive for some nonempty set of temperatures. Further note that we impose this assumption first, before imposing weak monotonicity, as described under assumption #1. An example of this assumption in practice is given in panel C of Figure E1.

A visual example of the influence of these constraints can be seen for one example impact region (Chicago, Illinois) in Figure E1. Under these assumptions, we estimate projected daily impacts separately for each impact region, and then aggregate these high resolution effects to state, country, and global levels, using population weighting.

different countries.

[102]See Appendix F.4 for results in which we explore a scenario with slower rates of adaptation. Under this alternative scenario, Assumption #1 binds much less frequently.

[103]Note that these fixed effects are by definition not affected by a changing weather distribution. Thus, their omission does not influence estimates of climate change impacts.

E.3 Projected benefits of adaptation

Figure E2 provides an examination of projected changes in the mortality-temperature relationship over time, which is a key ingredient for projections of future damages and adaptation. We plot the spatial distribution of the *change* in the mortality-temperature relationship evaluated at 35°C between 2015 and 2050 (panel A) and 2015 and 2100 (panel B) for the >64 age category. Specifically, these values are calculated as:

$$\hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{rt}, \log(GDPpc)_{rt}) - \hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2015}, \log(GDPpc)_{r,2015}),$$

where $\boldsymbol{T}_{35}$ is a fourth order polynomial for a daily temperature of 35°C, a indicates the >64 age group, and t is either 2050 or 2100. The maps reveal that in most regions of the world, there is a clear downward trend in the sensitivity of mortality rates to high temperatures, as locations get both richer and hotter as the century unfolds. For the >64 age group, the average global increase in the mortality rate on a 35°C day (relative to a day at location-specific minimum mortality temperatures) declines by roughly 75% between 2015 and 2100, going from 10.1 per 100,000 to just 2.4 per 100,000 in 2100.

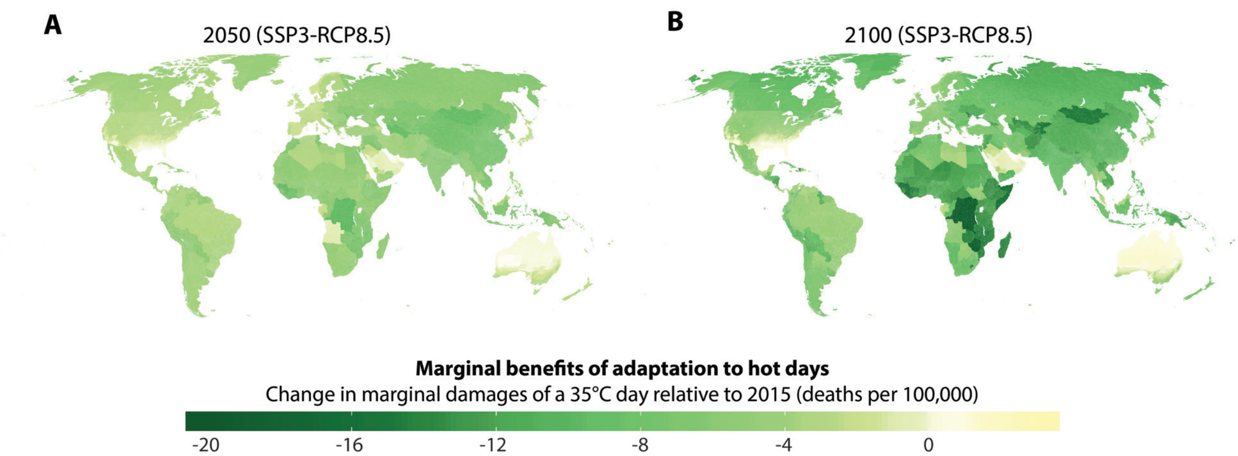

Figure E2: Spatial and temporal heterogeneity in temperature sensitivity. Panels A and B indicate the change in mortality sensitivity to hot days (35°C) for the oldest age category (>64) between 2015 and 2050 (A), and between 2015 and 2100 (B). Specifically, these values are $\hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2050}, \log(GDPpc)_{r,2050}) - \hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$ in panel A and $\hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2100}, \log(GDPpc)_{r,2100}) - \hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$ in panel B, where $\boldsymbol{T}_{35}$ is a fourth order polynomial for a daily temperature of 35°C and where the age group is $a > 64$. Darker colors signify larger predicted adaptation to heat. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

F Climate change projections: Additional results and robustness

This appendix provides additional illustrations of the main climate change projection results used and discussed throughout the main text (i.e., Section 5.5), as well as a robustness check and sensitivity analysis regarding the functional form of the mortality-temperature relationship, different assumptions about the behavior of the relationship outside of the historical sample values, and assumptions regarding the rate of adaptation.

F.1 Additional climate change projection results

Alternative measures of climate change impacts In Figure 5 of the main text, we show a map of impact region-level mean estimates of the mortality effects of climate change, accounting for adaptation and income benefits, but not adaptation costs. However, in Sections 5 and 6.2 we also define three other measures of expected climate change impacts: (i) mortality effects of climate change with neither adaptation nor income growth; (ii) mortality effects of climate change with benefits of income growth; (iv) full mortality risk of climate change, accounting both for adaptation and income benefits as well as adaptation costs. Panels A, B and D in Figure F1 below show projected impacts for each of these alternative measures; for comparison, panel C repeats the mortality effects of climate change with benefits of adaptation and income growth map from the main text.

Section 5 also presents a time series of aggregate global mortality consequences of climate change for measures (i), (ii), and (iii). Figure F2 adds estimates of measure (iv) to the same figure, showing the mean estimate of the full mortality risk of climate change over time, as well as the uncertainty surrounding this mean, as captured by Monte Carlo simulations.

Finally, Figure F3 presents the same time series of aggregate global mortality consequences of climate change for measures (i) and (ii), as in Figure F2, but adds shading to indicate the uncertainty surrounding the mortality effects of climate change with benefits of income growth. Just as in Table D1, the uncertainty in the estimated relationship between income per capita and the temperature sensitivity of mortality is apparent.

Climate change projections by age group In the main text, Figure 6 displays a time series of climate change impacts on the global average mortality rate. This aggregate value represents, in each year, the sum across age-specific projections, where death rates are population weighted by age-specific population values. Below in Figure F4, we show each of these age-specific projections for SSP3 and RCP8.5 (for reference, Table 1 shows that the average mortality rate for the oldest age group is 4,736 deaths per 100,000 in our estimation sample). While all age groups have a mean estimate that is above zero by end-of-century, the oldest age group dominates our projections in terms of death rates. These large demographic differences are taken into account in our valuation steps (see Section 7 and Appendix G).

Climate change projections by socioeconomic scenario Throughout Section 5.5 of the main text, we display climate change projection results under the socioeconomic scenario SSP3. Each SSP scenario models a different possible pathway of economic development, population growth, and demographics; here, we show the global mortality effects of climate change under two alternative scenarios (SSP2 and SSP4,

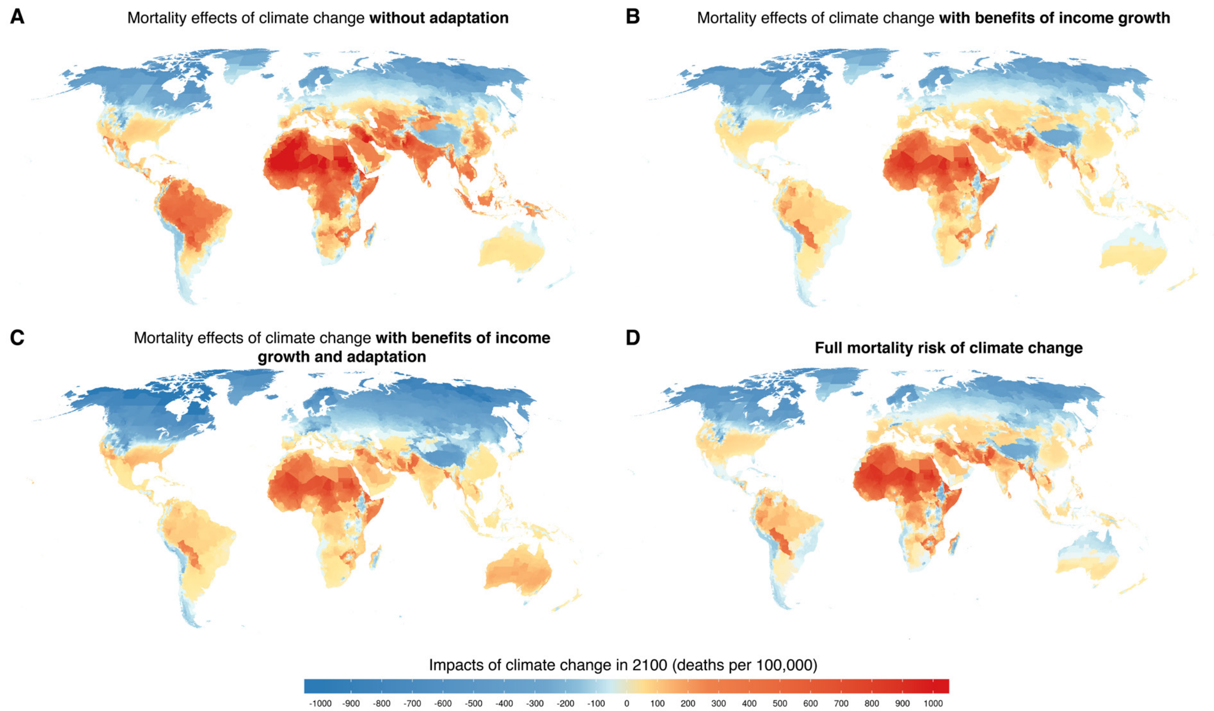

Figure F1: Mortality costs of climate change under alternative adaptation scenarios. All maps show predicted mortality effects of climate change and colors in each impact region represent the mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Panel A shows estimates of the change in mortality rates when each impact region does not adapt. Panel B shows estimates of the change in mortality rates when impact region mortality sensitivity to temperature changes with future income, but not to future temperatures. Panel C allows populations to additionally adjust to experienced temperatures in the warming scenario, showing mortality rate changes when mortality sensitivity to temperature evolves with both future income and temperature. Finally, panel D shows the full mortality risk of climate change. This measure allows the mortality sensitivity to temperature to change with future income and future temperature, while also accounting for the costs of adapting to a warming climate. Adaptation costs are calculated are measured in units of death equivalents. All projections shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario and are calculated as the climate model weighted mean estimate across Monte Carlo simulations conducted on 33 climate models.

alongside SSP3). In each column, we show results for two separate modeling groups that produce projections for each SSP (IIASA and OECD, as discussed in Appendices B.3.2 and B.3.3).

Gains from mitigation spatially and in aggregate. Figure 5 displays climate change impacts spatially under the socioeconomic scenario SSP3 for the entire globe. Figure F6 shows a comparison between impacts under RCP8.5 and RCP4.5, showing the gains from mitigation. As expected, reducing emissions to the level of RCP4.5 is predicted to have substantial benefits in terms of reduced mortality risks from climate change. However, the spatial pattern of impacts remains, with clearly unequal distribution of impacts between places that are relatively poor today versus places that are relatively wealthy. Figure F7 replicates the aggregate time series (Figure 6 in the main text) for RCP4.5 under the SSP3 scenario, and Table F1 replicates the aggregate mortality damage estimates for 2100 (Table 2 in the main text) for RCP4.5 under SSP3.[104] The gains from reducing emissions are evident in both sets of aggregate results.

[104]In Table 2 in the main text and F1 here, Europe includes the Aland Islands, Albania, Andorra, Austria, Belarus, Belgium,

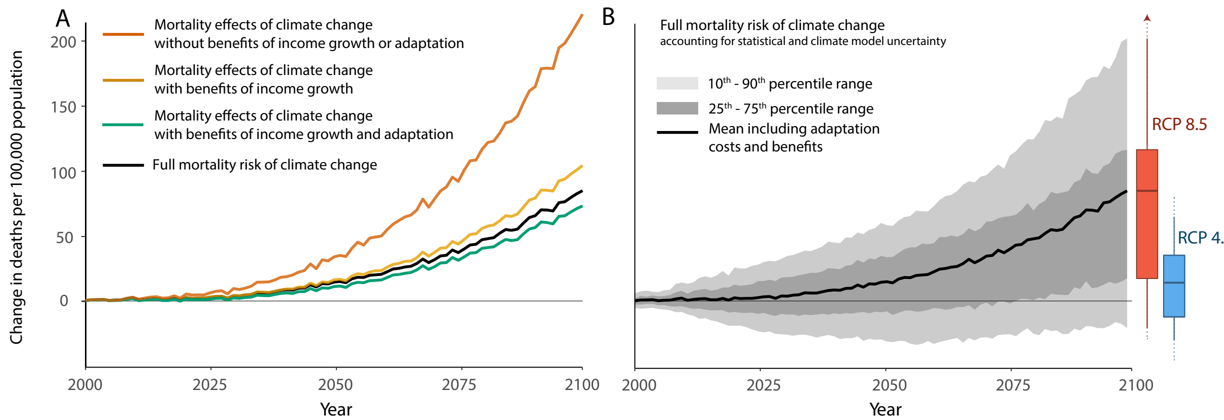

Figure F2: Time series of projected full mortality risk of climate change. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents a partial mortality effect, while the black line shows the full mortality risk due to climate change, accounting for both adaptation costs and benefits. Orange (expression (i)): mortality effects without adaptation. Yellow (expression (ii)): mortality effects with benefits of income growth. Green (expression (iii)): mortality effects with benefits of income growth and adaptation. Black (expression (iv)): full mortality risk calculated as the sum of mortality effects with adaptation and income growth benefits plus estimates of costs incurred to achieve adaptation, measured in units of death equivalents. Panel B shows the 10^{th}-90^{th} percentile range of the Monte Carlo simulations for the full mortality risk of climate change (black line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of full mortality risk impacts in 2100 under both RCPs. All line estimates shown refer to the RCP8.5 emissions scenario and all line and boxplot estimates refer to the SSP3 socioeconomic scenario. Figure F7 shows the equivalent for SSP3 and RCP4.5.

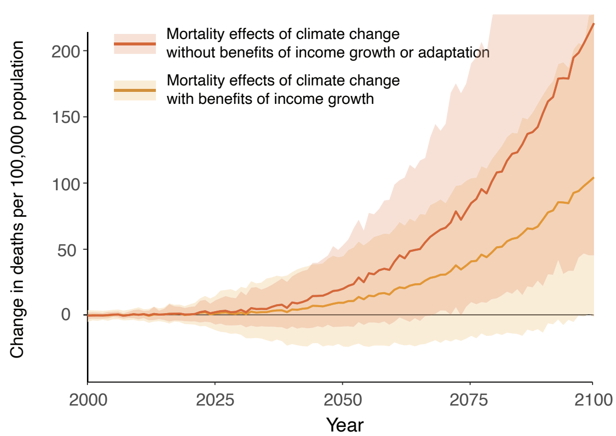

Figure F3: Uncertainty in the mortality effects of climate change including benefits of income growth. Both solid lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Shaded areas indicate the 10^{th}-90^{th} percentile range of the Monte Carlo simulations. The orange line and confidence interval shows the mortality effects without adaptation (expression (i)), while the yellow line and confidence interval shows the mortality effects with benefits of income growth (expression (ii)). Both projection estimates shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

Bosnia and Herzegovina, Bulgaria, Croatia, Czech Republic, Denmark, Estonia, Faroe Islands, Finland, France, Germany, Gibraltar, Greece, Guernsey, Hungary, Iceland, Ireland, Isle of Man, Italy, Jersey, Kosovo, Latvia, Liechtenstein, Lithuania, Luxembourg, Macedonia, Malta, Moldova, Monaco, Montenegro, Netherlands, Norway, Poland, Portugal, Romania, Russia, San Marino, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Ukraine, United Kingdom, and Vatican City. Similarly, sub-Saharan Africa includes Angola, Benin, Botswana, Burkina Faso, Burundi, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Cote d'Ivoire, Democratic Republic of the Congo, Djibouti, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Gambia, Ghana, Guinea, Guinea-Bissau, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritania, Mauritius, Mayotte, Mozambique, Namibia, Niger, Nigeria, Reunion, Republic of Congo, Rwanda, Saint Helena, Sao Tome and Principe, Senegal,

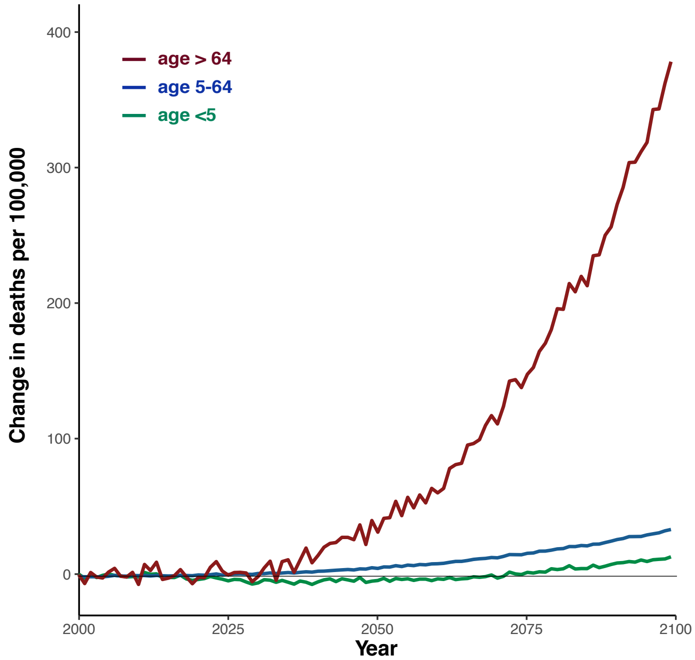

Figure F4: Heterogeneity in climate change impacts on mortality by age group. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Each line represents one of the three age groups used in the analysis: <5, 5-64, and >64. Results are shown for the combination of SSP3 and RCP8.5 with a fourth-order polynomial functional form of temperature. Figure 6 in the main text represents the sum across these age-specific projections, where death rates are population weighted by age-specific population values.

Seychelles, Sierra Leone, Somalia, South Africa, South Sudan, Sudan, Swaziland, Tanzania, Togo, Uganda, Western Sahara, Zambia, and Zimbabwe.

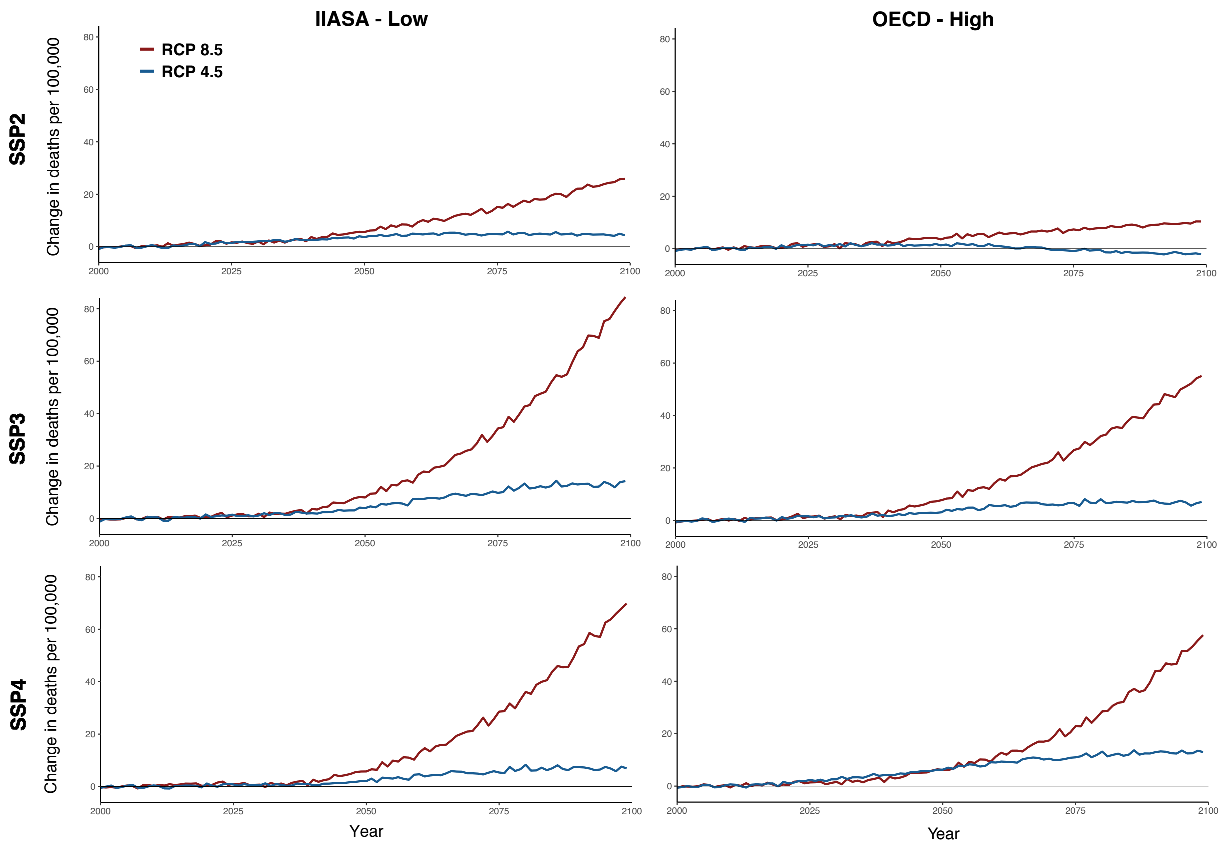

Figure F5: The full mortality risk of climate change under different scenarios of population growth, economic growth, and emissions. Rows denote different Shared Socioeconomic Pathway (SSP) scenarios, columns denote two separate modeling groups that produce data for each SSP, and each panel shows a time series of the total mortality costs of climate change for RCP 4.5 and RCP 8.5. Both lines indicate total predicted mortality costs due to climate change, accounting for both adaptation benefits and costs, and indicate the mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. RCP8.5 is a high-emissions scenario, while RCP4.5 is a scenario with aggressive emissions reductions. The OECD economic projections tends to exhibit slightly higher income growth than the IIASA economic projections. Throughout the main analysis, projection results relying on IIASA and OECD socioeconomic projections are both used and weighted equally.

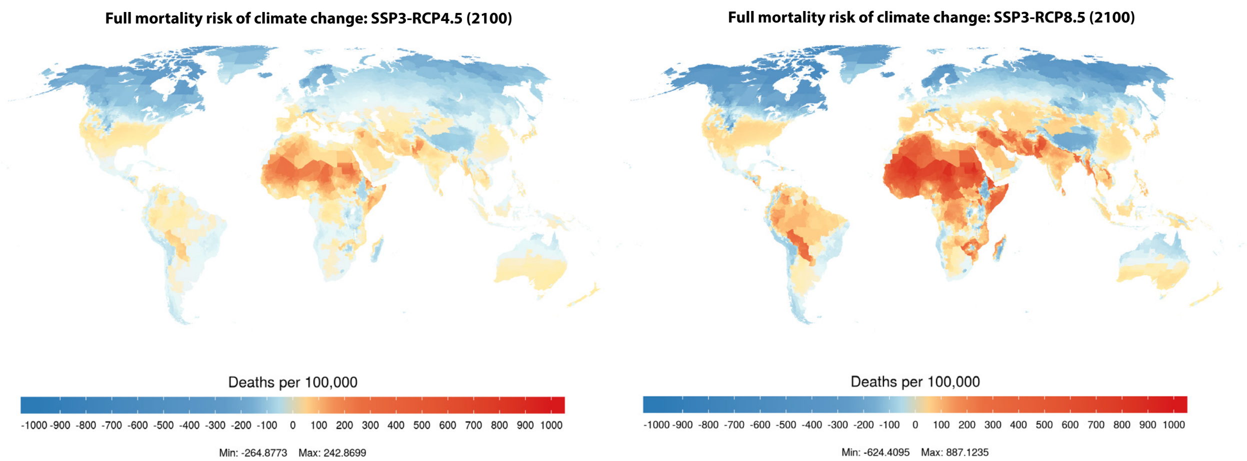

Figure F6: The mortality risk of future climate change under RCP4.5 and RCP8.5 for SSP3. These maps indicate the full mortality risk of climate change, measured in units of deaths per 100,000 population, in the year 2100. Estimates account for both the costs and the benefits of adaptation, and the map shows the weighted mean estimate across Monte Carlo simulations conducted on 33 climate models, accounting for both econometric and climate uncertainty. All values shown refer to the SSP3 socioeconomic scenario.

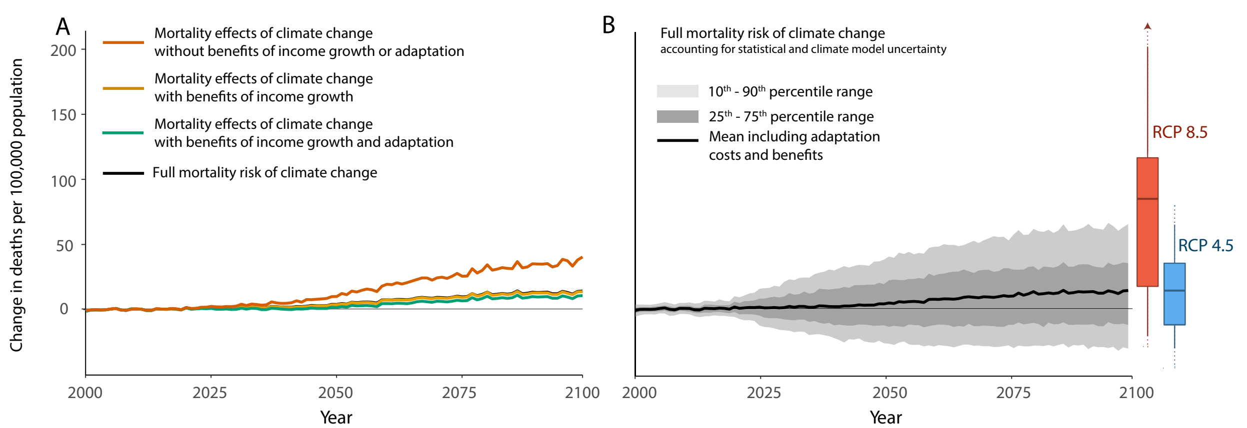

Figure F7: Time series of projected mortality risk of climate change under RCP4.5 for SSP3. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents a partial mortality effect, while the black line shows the full mortality risk due to climate change, accounting for both adaptation costs and benefits. Orange (expression (i)): mortality effects without adaptation. Yellow (expression (ii)): mortality effects with benefits of income growth. Green (expression (iii)): mortality effects with benefits of income growth and adaptation. Black (expression (iv)): full mortality risk calculated as the sum of mortality effects with adaptation and income growth benefits plus estimates of costs incurred to achieve adaptation, measured in units of death equivalents. Panel B shows the 10^{th}-90^{th} percentile range of the Monte Carlo simulations for the full mortality risk of climate change (black line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of full mortality risk impacts in 2100 under both RCPs. All line estimates shown refer to the RCP4.5 emissions scenario and all line and boxplot estimates refer to the SSP3 socioeconomic scenario.

Table F1: Estimates of the global mortality risk of climate change in 2100 (moderate emissions scenario, RCP4.5)

	Mortality effects of climate change				Full mortality risk of climate change	
	No income growth or adaptation	**Benefits of income growth**	**Benefits of climate adaptation**	**Costs of climate adaptation**		
	deaths/100k	*deaths/100k*	*deaths/100k*	*deaths/100k*	*deaths/100k*	*% of GDP*
	(1)	(2)	(3)	(4)	(5a)	(5b)
Panel A: Global estimates						
Globe	40.3	-26.5	-3.0	3.5	14.2	0.6
IQR	[7.8, 57.9]	[-47.8, -2.6]	[-16.3, 8.1]	[-0.9, 7.5]	[-12.3, 35.2]	[-3.9, 4.6]
Panel B: Regional estimates						
China	14.8	-8.6	-12.5	9.9	3.4	0.5
USA	-9.0	-1.0	-2.7	12.7	-0.1	0.2
India	78.0	-69.6	3.6	-1.0	11.1	1.5
Pakistan	144.3	-44.3	-18.6	16.1	97.6	8.0
Bangladesh	77.2	-21.9	-13.3	7.1	49.1	4.5
Europe	-28.0	6.5	-42.0	50.6	-12.8	-0.7
Sub-Saharan Africa	37.6	-16.1	-1.5	0.5	20.5	1.7

All columns show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In the first row, brackets indicate the interquartile range (IQR). Columns 1-4 each indicate a partial mortality effect of climate change, in units of deaths per 100,000. Column 1 (expression (i)): mortality effects of climate change without benefits of income or adaptation to climate change. Column 2 (expression (ii) - expression (i)): benefits of income growth. Column 3 (expression (iii) - expression (ii)): benefits of adaptation to climate change. Column 4 (Equation 10): mortality-related costs of adaptation inferred using a revealed preference approach, measured in "death equivalents". Columns 5a-5b (expression (iv)): the full mortality risk of climate change, measured in deaths per 100,000 (column 5a) and represented as % of 2100 GDP (column 5b) using an age-adjusted value of the U.S. EPA VSL with an income elasticity of one applied to all impact regions. Column 5a is equivalent to the sum of columns 1 through 4. All estimates shown rely on the RCP4.5 emissions scenario and the SSP3 socioeconomic scenario. Note that benefits of income growth are positive for Europe (column 2) because higher incomes lower mortality sensitivity to extreme cold, as well as to extreme heat (see Table D1 and Figure 1). In Europe, many lives are saved under climate change due to reductions in extreme cold, and accounting for income growth lowers the estimated number of lives saved.

The impact of climate change in 2100 under RCP4.5 compared to contemporary leading causes of death. Figure F8 presents the same results as Figure 10 in the main text, but for RCP4.5. As can be seen, despite the overall decrease in the average impact under SSP3 and RCP4.5 when compared to RCP8.5, much of the inequality in both the impacts and the adaptation costs that was evident in Figure 10 remains.

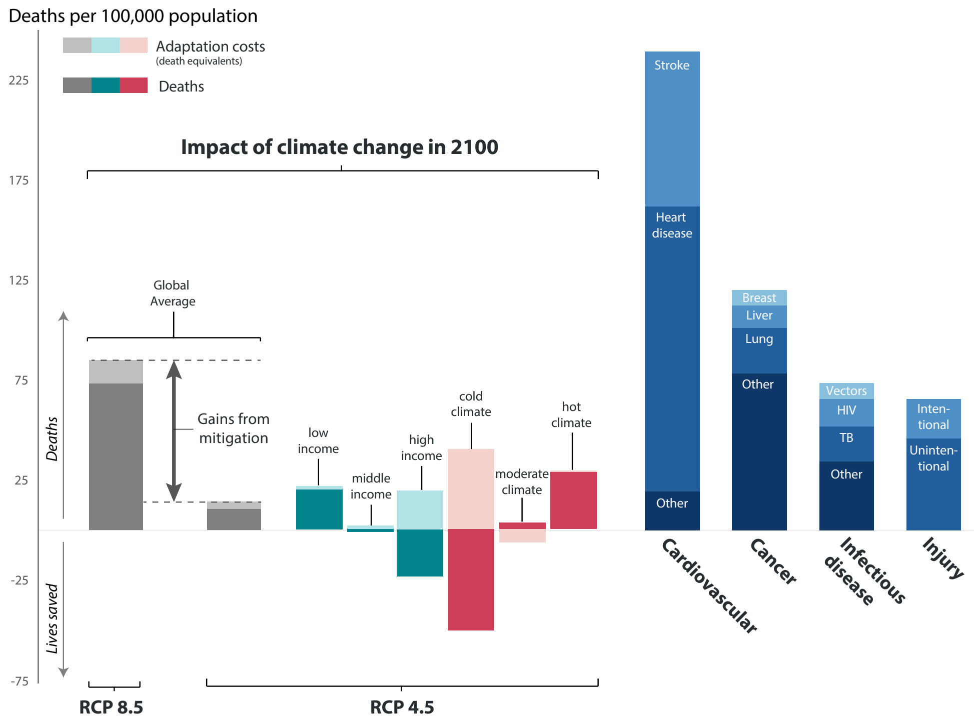

Figure F8: The impact of climate change in 2100 under RCP4.5 compared to contemporary leading causes of death. Impacts of climate change (grey, teal, and coral) are calculated for the year 2100 for SSP3 and include changes in death rates (solid colors) and changes in adaptation costs, measured in death equivalents (light shading). Global averages for RCP 8.5 and RCP 4.5 are shown in grey, demonstrating the gains from mitigation. Income and average climate groups under RCP4.5 are separated by tercile of the 2015 global distribution across all 24,378 impact regions. Blue bars on the right indicate average mortality rates globally in 2018, with values from WHO (2018).

F.2 Robustness: Alternative functional form for the mortality-temperature relationship

As discussed in Section D.2, we experiment with four distinct nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ in Equations D.17 and 5 in the main text. The fourth order polynomial is our main specification because it strikes a balance between providing sufficient flexibility to capture important nonlinearities, parsimony, and limiting demands on the data when covariate interactions are introduced in Equation 5. However, the binned specification, in which $\boldsymbol{T}_{it}$ contains binned daily temperatures with a fixed set of 5°C bins, is the most flexible functional form. In Figure D3, we show that the binned and fourth order polynomial functional forms recover similar mortality-temperature response functions across our pooled multi-country sample. Below in Figure F9, we show that this similarity carries through to generate similar climate change impact projections across the binned and polynomial functional forms. Both projections are constructed using estimation of the interaction model in Equation 5 in combination with high-resolution covariates $TMEAN$ and $\log(GDPpc)$ to generate impact region-specific response functions (see Section 5.2 for details).

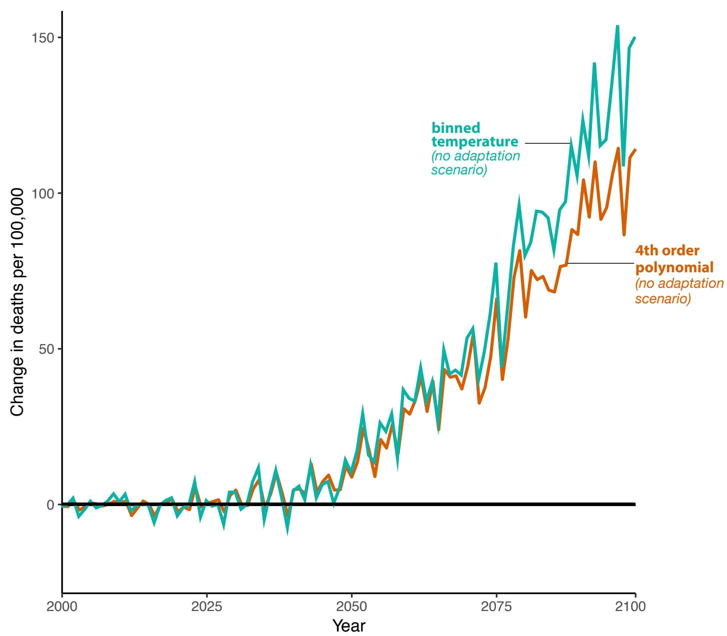

Figure F9: Robustness of impact projections to alternate functional forms of temperature. Each line represents the time series of changes to the mortality rate due to climate change under the socioeconomic scenario SSP3 and the emissions scenario RCP 8.5. Results shown are for a single climate model (CCSM4). Lines shown refer to estimates of mortality effects of climate change without adaptation or benefits of income growth, in which response functions do not evolve over time. In orange is the projected impact of climate change estimated using a fourth-order polynomial functional form of temperature in estimation of the regression model in Equation 5. In green is the same object, but with binned daily temperatures used as a functional form in estimation. While the binned regression imposes far fewer restrictions on the regression than does the polynomial, the projected impacts under these two sets of parameterizations are strikingly similar.

F.3 Sensitivity analysis: Alternative assumptions on out-of-sample extrapolation of response functions

The paper uses historical data to estimate the mortality-temperature response and uses the results to project the impacts of temperatures in the future. A key challenge, however, is that climate change will cause locations to experience temperatures that have not been observed in the historical record (e.g., see Figure 3), thus necessitating out of sample predictions.

Figure F10 probes the sensitivity of the projections of mortality risk changes up to 2100 to alternative assumptions about the relationship between mortality and temperature at temperatures that are not observed in available data sets. Specifically, for all temperatures above the maximum and below the minimum daily temperatures within our dataset, we alter the slope of the impact region-specific response functions in two ways. First for "constant out-of-sample extrapolation", we set the marginal effect of temperature fluctuations to equal the value at the maximum if above the maximum temperature, and vice versa for temperatures below the minimum (Figure F10B). This implies that the response function is flat for all temperatures outside the observed range. For "linear out-of-sample" extrapolation, we set the marginal effect to be linearly increasing in the out-of-sample regions with a slope equal to the slope between the response function evaluated at the maximum (minimum) and the maximum minus 0.1C (plus 0.1C) (Figure F10C). It is apparent that neither of these alternatives have a meaningful effect on the overall projected impacts; looking at projections from a single GCM, the projected impact of climate change on mortality rates, including the benefits of income growth and adaptation, is 13.6 per 100,000 in 2100 under RCP8.5 in the paper's main specification (Panel A) and 12.6 per 100,000 and 13.4 per 100,000 in Panels B and C.

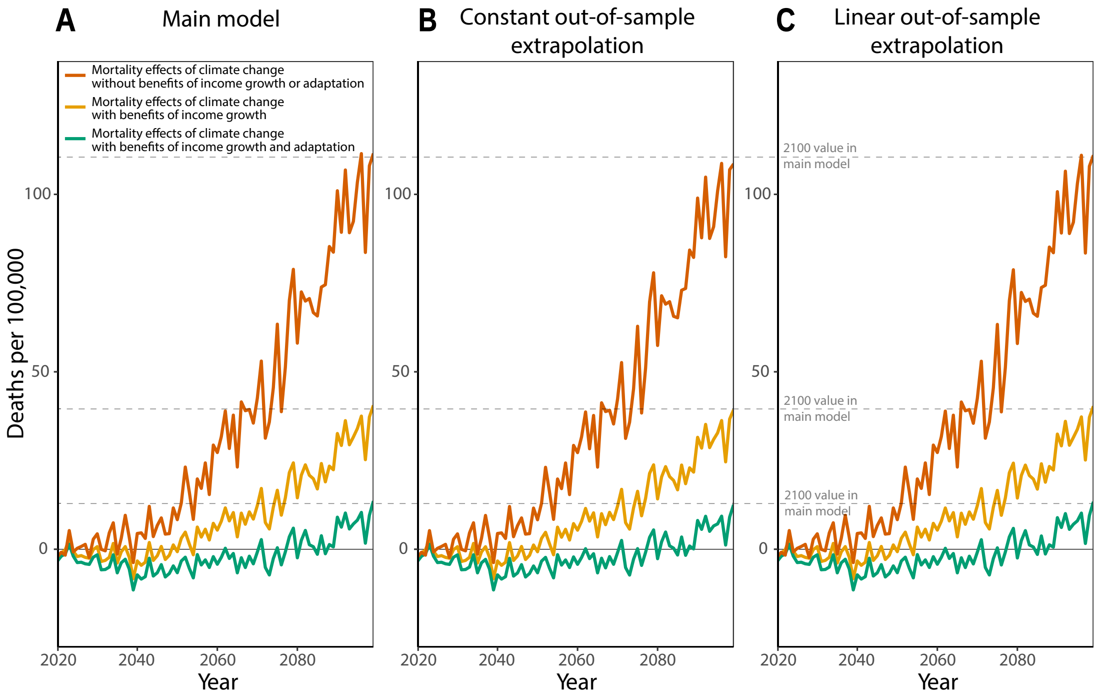

Figure F10: Two alternative assumptions on out-of-sample extrapolation of response functions. A) Time series projection of main model with no out-of-sample restrictions. B) Time series projection of response functions with constant out-of-sample restrictions above the maximum and below the minimum temperatures in our estimating sample. C) Time series projection of response functions with linearly increasing out-of-sample restrictions above the maximum and below the minimum temperatures in our estimating sample, with a slope equal to the slope between the response function evaluated at the maximum (minimum) and the maximum minus 0.1C (plus 0.1C). All projections rely on the CCSM4 climate model under RCP 8.5 and SSP3.

F.4 Sensitivity analysis: Alternative assumptions on the rate of adaptation

In our main results, we use the estimated coefficients from Equation 5 in combination with high-resolution data on the covariates $TMEAN$ and $\log(GDPpc)$ to extrapolate response functions both *across space* (to capture spatial heterogeneity in the mortality-temperature relationship) and *over time* (to capture future changes in the mortality-temperature relationship due to adaptation and benefits of income growth). As discussed in Section 4, the estimation of Equation 5 relies on cross-sectional variation in $TMEAN$ and $\log(GDPpc)$, in combination with plausibly random year-to-year variation in daily temperatures. However, as discussed in Appendix E.1, we apply the estimated coefficients from Equation 5 over time when computing future climate change impacts; in doing so, we must make an assumption regarding the rate at which mortality sensitivity to temperature declines with changing covariates. As discussed previously, our main specification relies on a 13-year Bartlett kernel for $\log(GDPpc)$ and a 30-year Bartlett kernel for $TMEAN$.

Here, we conduct two sensitivity analyses, each of which adjusts the assumed rate of adaptation. In the first, the speed at which the mortality-temperature response function changes with time-varying covariates is deterministically reduced by half. In the second, this rate is increased by 150%. These exercises are used to understand how climate change impact projections change if the evolution of the response function towards zero (see Figure E1) is assumed to occur more slowly or more quickly.

In the main model, income grows for each impact region r according to $GDP_{rt} = \rho_{ct} GDP_{r,t-1}$, where c indicates the country that region r falls into, and ρ_{ct} is a country- and year-specific growth rate given exogenously by the SSP scenarios. The kernel-averaged climatic temperature for region r used in the main model is $TMEAN_{rt} = TMEAN_{r,t-1} + \Delta TMEAN_{rt}$. In this "slow adaptation" alternative approach (see Figure F11B), we replace income growth with $GDP_{rt} = \left(\frac{\rho_{ct}-1}{2}+1\right) GDP_{r,t-1}$ after the year 2015, and we reduce linear growth in temperature by replacing it with $TMEAN_{rt} = TMEAN_{r,t-1} + \frac{\Delta TMEAN_{rt}}{2}$. In the "fast adaptation" alternative approach (see Figure F11C), we similarly replace income growth with $GDP_{rt} = 1.5\rho_{ct} GDP_{r,t-1}$ after the year 2015, and we reduce linear growth in temperature by replacing it with $TMEAN_{rt} = TMEAN_{r,t-1} + 1.5\Delta TMEAN_{rt}$. Note that both the primary specification and reduced growth analyses generate identical covariates (and hence, response functions) in 2015.

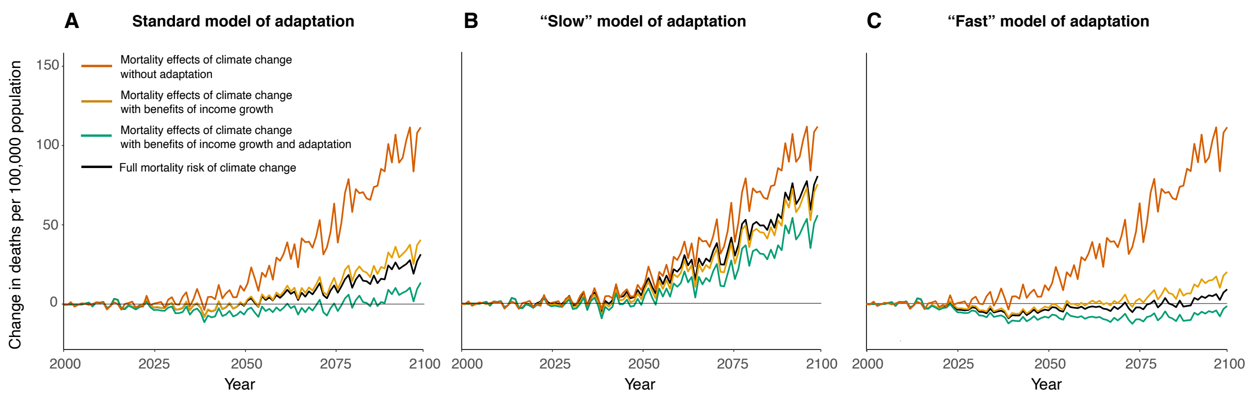

Figure F11: Impacts of climate change on mortality under alternative assumptions about rates of adaptation. Time series of projected full mortality risk of climate change (black line), as compared to partial estimates from incomplete accounting of the costs and benefits of adaptation (other colors). All lines show predicted mortality impacts of climate change across all age categories under the RCP8.5 emissions scenario, for the socioeconomic scenario SSP3, and using a single climate model (CCSM4). Panel A shows results for our standard model of adaptation, as described in Section 5.2. Panel B shows results for an alternative model of adaptation in which the rate of adaptation to both income growth and to a warming climate is cut in half. Panel C shows results for an alternative model in which the rate of adaptation to both income growth and to a warming climate is increased by 150%.

G Calculation of a mortality partial social cost of carbon

In principle, one could compute a mortality partial social cost of carbon (SCC) estimate by perturbing each global climate model (GCM) in the Surrogate Mixed-Model Ensemble (SMME) with a pulse of CO_2 and projecting mortality for each location in both the original and perturbed simulations. However, in practice, such a procedure is both prohibitively costly from a computational standpoint and would also prevent the calculation of an SCC for any climate trajectory that did not exactly coincide with one of the 33 models. Instead, we rely on a "simple climate model",[105] in combination with our empirically-derived damage functions, to construct mortality partial SCC estimates. We detail this implementation below.

G.1 Computing post-2100 damage functions

For data availability reasons, it is necessary to develop an alternative approach to estimate post-2100 damage functions. Only 6 of the 21 GCMs that we use to build our SMME ensemble (see Section 3.2) are run by their respective modeling teams to simulate the climate after the year 2100 for both RCP scenarios and post-2100 data are not available in the NEX-GDDP downscaled and bias-corrected projections that we use for generating high-resolution impact projections. Similarly, the SSPs needed to project the benefits of income growth and changes in demographic compositions also end in 2100. While one approach is to simply end economic cost calculations in 2100, as was done in Hsiang et al. (2017), neglecting post-2100 damages is a substantial omission because a large fraction of costs, in NPV, are thought to occur after 2100 at 3% discount rates (Kopp and Mignone, 2012).

To estimate post 2100-damages, we develop a method to extrapolate changes in the damage function beyond 2100 using the observed evolution of damages near the end of the 21^{st} century. The year-specific damage functions estimated using Equation 12 reveal that in the latter half of the 21^{st} century, full mortality damages are larger for a given level of warming if warming occurs later in time and damage functions become more convex with time at the end of the 21^{st} century. The finding that mortality costs rise over time is the net result of countervailing forces. On the one hand, later years are projected to have larger and older[106] populations with higher VSLs due to rising income, facts that raise damages. On the other hand, populations are better adapted due to higher incomes and a slower rate of warming projected in later years, an effect that would lower damages. Our results suggest the former dominates by end of century, causing damages to be trending upward at the moment that our high-resolution simulations end in 2100.

The motivating principle of our extrapolation approach is that these observed changes in the shape of the damage function near the end of the century provide plausible estimates of future damage function evolution after 2100. To execute this extrapolation, we pool values D_{irmt} from 2085-2100 and estimate a quadratic model similar to Equation 12, but interacting each term linearly with year t (we use 2085-2100 because the evolution of damages over time becomes roughly linear conditional on ΔGMST by this period). The temporal trend over the entire 21^{st} century is convex, implying that our linearization is, if anything,

[105]See Hsiang and Kopp (2018) for a description of climate model classes.

[106]In SSP3, the share of the global population in the most vulnerable >64 age category rises from 8.2% in 2015 to 16.2% in 2100.

conservative. The specific interaction model we estimate is:

$$D_{irmt} = \alpha + \nu_1 \Delta GMST_{rmt} \times t + \nu_2 \Delta GMST^2_{rmt} \times t + \varepsilon_{irmt}$$

This allows us to estimate a damage surface as a parametric function of year. We then predict extrapolated damage functions for all years after 2100, smoothly transitioning from our flexible climate model-based damage functions prior to 2100.

G.2 Set up of the climate module using a simple climate model

A core component of any analysis of the SCC is the climate module used to estimate both the baseline climate and the response of the climate system to a marginal change in greenhouse gas emissions. The Finite Amplitude Impulse Response (FAIR) model (Millar et al., 2017) satisfies key criteria for such a module, including those outlined by the National Academies of Sciences, Engineering, and Medicine (2017). In particular, the National Academies of Sciences, Engineering, and Medicine (2017) recommends that the climate module be transparent, simple, and "consistent with the current, peer-reviewed scientific understanding of the relationships over time between CO_2 emissions, atmospheric CO_2 concentrations, and CO_2-induced global mean surface temperature change, including their uncertainty" (National Academies of Sciences, Engineering, and Medicine, 2017, p.88). For this last criterion, the authors recommend that the module be "assessed on the basis of its response to long-term forcing trajectories (specifically, trajectories designed to assess equilibrium climate sensitivity, transient climate response and transient climate response to emissions, as well as historical and high- and low-emissions scenarios) and its response to a pulse of CO_2 emissions." The authors specifically point to the FAIR model as an example of a model that is structurally capable of meeting all these criteria.

The FAIR model is defined by five equations that represent the evolution of global mean variables over time t. Global mean surface temperature $GMST$ is the sum of two temperature variables, T_0 and T_1, representing the slow and fast climate system response to forcing F:

$$\frac{dT_i}{dt} = \frac{q_i F - T_i}{d_i}, i \in \{0, 1\}, \tag{G.21}$$

where the q_i values collectively define the equilibrium climate sensitivity (ECS), and where the d_i values (the thermal adjustment times) along with q_i define the transient climate response (TCR). The ECS is the sensitivity of the climate (as measured by GMST increases) to a doubling of atmospheric CO_2, relative to some initial state. The TCR is the average temperature response to a doubling of CO_2 in which the CO_2 increases by 1% each year. The ECS is larger than the TCR, as it captures the time taken for the climate system to fully adjust to increased CO_2.

The CO_2 concentration above the pre-industrial baseline, R, is the sum of four fractions, R_j, representing different uptake timescales:

$$\frac{dR_j}{dt} = a_j E - \frac{R_j}{\alpha_j \tau_j}, j \in \{0, 1, 2, 3\} \tag{G.22}$$

where E is the CO_2 emissions rate, a_j values represent the fraction of emissions that enter each atmospheric fraction, τ_j values represent the base uptake time scale for each fraction, and where α_j is a state-dependent coefficient that reflects feedbacks from temperature onto uptake timescales. The remaining three equations describe forcing F as a function of R and of exogenous non-CO_2 forcing, and α as a function of global mean surface temperature and atmospheric CO_2 concentrations (see Millar et al. (2017) for details).

We obtain the latest release of the FAIR model, which was version 1.3.2 at the time of computation, from its online repository.[107] As described below in Section G.2.1, we develop a methodology to generate mortality partial SCC estimates that capture uncertainty in climate sensitivity by varying four core parameters in FAIR: the equilibrium climate sensitivity (ECS), the transient climate response (TCR), the short thermal adjustment time (d_2), and the time scale of rapid carbon uptake by the ocean mixed layer (τ_3). By varying these four parameters across thousands of Monte Carlo simulations, we are able to capture uncertainty in the short and long term response of temperature and the carbon cycle to changes in emissions. The median values across our uncertainty distributions (described in detail below) for each core model parameter are as follows: ECS is 2.72°C per CO_2 doubling, TCR is 1.58°C per CO_2 doubling, d_2 is 3.66 years, and τ_3 is 4.03 years. Throughout our implementation, all other parameters in FAIR are held fixed at their default values.

The two scenarios considered in this analysis, RCP4.5 and RCP8.5, represent two widely divergent emissions and climatic pathways, especially in years beyond 2050. Following the method used in previous estimates of the SCC, including in the National Academies of Sciences, Engineering, and Medicine (2017), we include projections starting in the current period (here defined as 2020) through the year 2300. Due to the long residence times of CO_2 in the atmosphere and the changes in global mean surface temperature associated with CO_2 emissions, SCC estimates can vary significantly depending on the definition of this window, especially when low discount rates are applied. To illustrate the large differences across RCP scenarios, Figure G1 shows fossil CO_2 emissions, CO_2 concentrations, total radiative forcing (the difference between incoming solar radiation and outgoing terrestrial radiation), and temperature as anomalies from FAIR's reference state, which is year 1765, for the median climate parameters listed above and under each emissions scenario.

In order to estimate the marginal effect of CO_2 emissions, we add two additional scenarios to the "control scenarios" of RCP4.5 and RCP8.5. Each additional scenario adds a 1 GtC (3.66 Gt CO_2) pulse of fossil CO_2 emissions in 2020 to each of the control scenarios described above. The FAIR model is then run again for these pulse scenarios, resulting in a new time series of concentrations, forcing, and temperature anomalies. The difference between the control and pulse scenarios, including climate uncertainty (discussed below), is shown in the main text Figure 9; as described below and in Section 7, this difference is used to construct mortality partial SCC estimates.

[107] https://github.com/OMS-NetZero/FAIR/tree/v1.3.2.

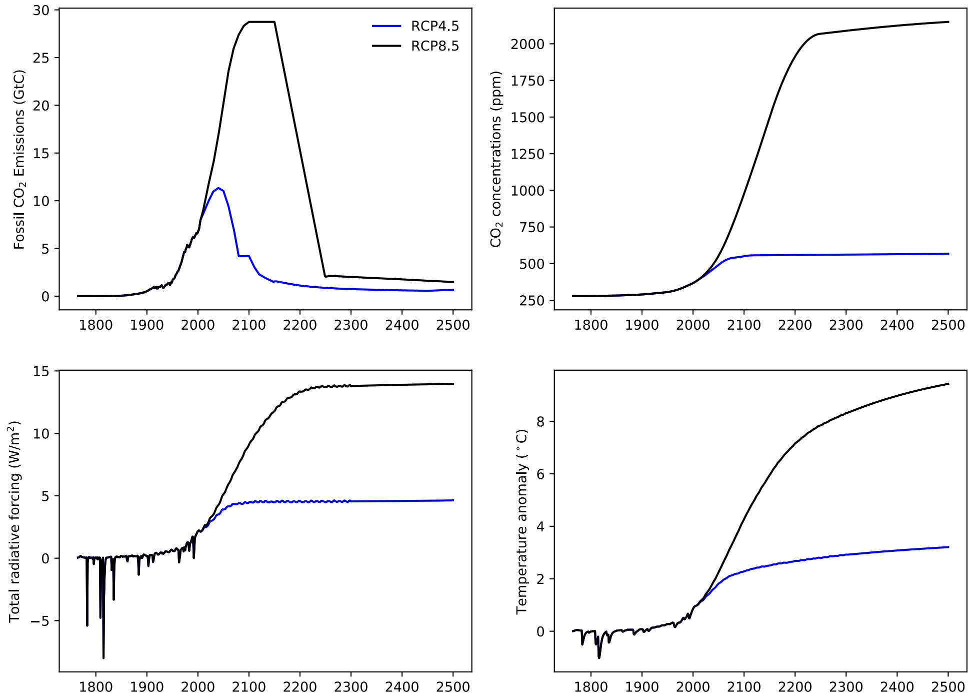

Figure G1: Behavior of key variables in the FAIR simple climate model under median climate parameters. Each panel shows the temporal trajectory of key variables in FAIR that are used in our calculation of the social cost of carbon. The trajectories shown arise under FAIR run with median climate parameter values calculated from our uncertainty distributions for the equilibrium climate sensitivity, transient climate response, short thermal adjustment time, and time scale of rapid carbon uptake by the ocean mixed layer. The values are shown as anomalies from the year 1765, FAIR's reference state.

G.2.1 Methodology for capturing uncertainty in climate sensitivity within the simple climate model FAIR

A complete study of the mortality partial SCC should represent the uncertainty in key model parameters, including the joint probability distribution of the ECS and TCR. We discuss here our approach to modeling this climate sensitivity uncertainty.

The analysis described above relies solely on the simple climate model FAIR with key climate parameters set to median values that are computed from their uncertainty distributions. We now discuss the development of those uncertainty distributions and the representation of climate uncertainties in FAIR. To represent climate uncertainties, we vary TCR, ECS, d_2, and τ_3 such that our climate uncertainties conform to those of the literature. These four parameters represent the behavior of the short and long timescales of response of temperature and the carbon cycle. For TCR and ECS, we draw upon constraints from the IPCC Fifth Assessment Report (AR5) (Collins, Knutti et al., 2013); for d_2 and τ_3 we follow Millar et al. (2017), based on analysis of Joos et al. (2013) and Geoffroy et al. (2013).

In general, we produce initial distributions of these parameters based on the literature constraints. However, a key difference between our approach and those in the existing literature is that we explicitly model the tails of the climate sensitivity uncertainty distributions. The AR5 synthesis generally regards the 5–95%

ranges of variables in the CMIP5 models as representing the "likely" range (central at least 66% probable range) due to structural uncertainty. Previous studies based on CMIP5 results (e.g., Joos et al. (2013); Ricke and Caldeira (2014)) and those using the CMIP5 5–95% range of TCR and ECS as 5-95% input ranges to their models (e.g., Millar et al. (2017)) thus show results that characterize only the central 66% of possibilities. Here we explicitly model the tails of the input and output distributions by generating TCR and ECS distributions with likely ranges as specified by the AR5 report. To preserve the expected correlation between TCR and ECS, rather than sampling ECS directly, we follow Millar et al. (2015) and instead sample the realized warming fraction (RWF, the ratio of TCR/ECS), which is nearly independent of TCR. We subsequently filter the parameter sets to ensure consistency with expectations regarding the initial pulse adjustment timescale (the time it takes the climate system to reach a warming peak following a pulse emission of CO_2).

Below we outline the sources used to construct the distributions of each parameter.

TCR: Collins, Knutti et al. (2013) conclude that "TCR is *likely* in the range 1°C to 2.5°C... is positive and extremely unlikely greater than 3°C" (p. 1112). In IPCC terminology (Mastrandrea et al., 2010), *likely* refers to a probability of at least 66%, *very likely* to a probability of at least 90%, and *extremely likely* to a probability of at least 95%. Thus we construct a log-normal distribution for TCR with the 17th to 83rd range of 1.0-2.5 °C.

RWF: As noted by the National Academies of Sciences, Engineering, and Medicine (2017), a RWF likely range of 0.45 to 0.75 is approximately consistent with the ECS likely range of 1.5 − 4.5°C (Collins, Knutti et al., 2013). We construct a normal distribution for RWF following this central 66% likelihood range, and sample this distribution, along with TCR, to construct the ECS distribution as TCR/RWF.

ECS: Collins, Knutti et al. (2013) conclude that "ECS is positive, *extremely unlikely* less than 1°C (high confidence), and *very unlikely* greater than 6°C (medium confidence)" (p. 1111) and *likely* between 1.5 and 4.5°C. To construct our sampling distribution, we randomly draw samples from the TCR and RWF distributions, and obtain ECS samples by calculating TCR/RWF. The constructed ECS samples follow a log-normal distribution with the 17^{th}-83^{rd} range of 1.60-4.65 °C.

$\boldsymbol{d_2}$**:** The AR5 does not assess the range of d_2. Following Millar et al. (2017), we construct our distribution of d_2 as a log-normal distribution with a 5-95^{th} percentile range of 1.6-8.4 years.

$\boldsymbol{\tau_3}$**:** Joos et al. (2013) summarized τ_3 in three comprehensive Earth System Models (HADGEM2-ES, MPI-ESM, NCARCSM1.4), seven Earth System Models of Intermediate Complexity (EMICs), and four box-type models (ACC2, Bern-SAR, MAGICC, TOTEM). Using the mean (4.03) and standard deviation (1.79) of these values, we construct a normal distribution for τ_3.

After defining these distributions, we generate a 100,000-member ensemble of parameter sets via Monte Carlo sampling. As τ_3 should be larger than 0, we sample from a truncated normal distribution, and discard parameter sets in which $\tau_3 < 0$ or $> 2 \times 4.03$ to keep the mean of τ_3 in parameter sets consistent with the multi-model mean in Joos et al. (2013). About 2.4% of parameter sets are filtered by this constraint. Similarly, RWF must be less than 1. We therefore truncate its distribution at 1, which is the 99.4^{th} percentile, and truncate at the 0.06^{th} percentile to keep symmetry (which also removes unrealistic RWF values near and less than 0 that cause unrealistic, large and/or negative ECS values). About 1.2% of parameter sets are filtered by this constraint. After applying the τ_3 and RWF filters, which have a small overlap, we are left with 96,408 parameter samples. Using these remaining parameter samples, we evaluate model performance according to several criteria.

Our criteria for evaluating model performance are described in detail below, and summarized in Table G1 and Figure G2.

Initial pulse-adjustment timescale (IPT): The National Academies of Sciences, Engineering, and Medicine (2017) report highlights the IPT as a measure that is important for SCC computations, yet does not provide a clear, consistent definition. It "measures the initial adjustment timescale of the temperature response to a pulse emission of CO_2" and is "the time over which temperatures converge to their peak value in response to the pulse." (National Academies of Sciences, Engineering, and Medicine, 2017, p.88). This could either be the time to an initial peak, or the ultimate maximum temperature change over the duration of a simulation, which also depends on simulation length. Here we catalogue multiple versions of a potential IPT metric, comparing with previous literature where appropriate.

To assess the IPT, we set CO_2 concentrations to 2010 levels (389 ppm) and hold them constant throughout the simulation. To provide an emissions baseline to which a pulse will be added, we numerically solve the CO_2 emissions pathway in FAIR to meet the CO_2 concentration pathway for each parameter sample. We then construct a pulse experiment, in which 100 GtC of CO_2 is injected instantaneously in the year 2015. The difference in temperature between the pulse and control run measures the temperature response to a CO_2 pulse. To quantify the time to initial peak, we define the IPT as the time at which the time derivative of the temperature response first becomes negative (noting that, in many simulations, feedbacks between temperature and the carbon cycle mean that the temperature rises again after the initial peak and decline, and reaches the maximum temperature later. Therefore, the time to initial peak is not necessarily the same as the time to maximum temperature). The resulting IPT has a median of 9.0 years, with a central 90% probability range of 0–24.0 years. We drop parameter sets that lead to simulations in which the first negative time derivative of temperature occurs after 100 years post-pulse, indicative of temperatures that only increase throughout the experiment (in contrast to the simulations with an initial post-pulse decrease in temperature that begins increasing again after a time). This results in a filtering out of 112 additional parameter samples on top of the τ_3 and RWF filters, yielding a total number of post-filtering simulations of 96,306 for examination in the remaining discussion.

We also evaluate other potential metrics: the time to maximum temperature considering the full 500 year simulation, the time to maximum temperature considering just the 100 years post-pulse, and the time to maximum temperature considering 100 years post-pulse but excluding simulations reaching max at year 100. We find central 90% probable ranges of 4.0–485 (median 19.0), 4.0–100 (median 12.0), and 3.0–23.0 (median 9.0), respectively. The results of Joos et al. (2013) and subsequent analysis by Ricke and Caldeira (2014) indicate that a peak in warming in response to a pulse emission occurs within about a decade after emission. In particular, Ricke and Caldeira (2014) estimate a central 90% range for time to peak warming of 6.6–30.7 years, with a median of 10.1 years, and 2% of simulations reaching maximum at the end of their 100-year simulations. Ricke and Caldeira (2014), however, do not sample from continuous distributions of ECS and TCR, but rather use narrower discrete distributions of parameters based on individual CMIP5

GCMs; thus, we expect their range to be narrower than that in our analysis. Considering the first 100-years of simulation, our median time to peak warming is comparable to Ricke and Caldeira (2014), but spans a wider range of outcomes, as expected, with 24% of simulations reaching their peak at 100 years post-pulse (44% reach peak warming at simulation's end in year 2500).

Transient climate response to emissions (TCRE): The TCRE measures the ratio of transient warming to cumulative carbon emissions at the time of CO_2 doubling in a simulation with a 1% /year increase (year 70). Collins, Knutti et al. (2013) concluded that TCRE is between 0.8 and 2.5°C per 1000 GtC with at least 66% probability. To assess TCRE, we set up an experiment that increases CO_2 concentrations at 1%/year until CO_2 concentrations double in year 70. Again, for each parameter sample, we numerically solve the CO_2 emissions pathway in FAIR to meet the CO_2 concentration pathway. The resulting TCRE exhibits a likely range of 0.88–2.34°C per 1000 GtC, which is consistent with the central 66% probable range assessed by AR5.

Longevity of pulse warming: The coupled climate-carbon cycle experiments of Joos et al. (2013) indicate that a majority (about 70% in the multimodel mean) of peak warming persists 500 years after emissions. In our IPT experiments, the central 66% probable range is 72.9 – 137.6 percent of initial peak warming persists after 500 years.

Representative Concentration Pathway (RCP) experiments: We assess the warming in the RCP experiments relative to those in the CMIP5 multi-model ensemble, noting that we compare the central 66% probability ranges from our ensemble to those of the CMIP5 5th–95th percentile range (Table G1).

The final reduced sample set constitutes 96,306 samples as noted above, and the diagnostic metrics are essentially unchanged from the pre-filtering distributions (see Table G1). Based on this post-filtering evaluation, we conclude that the resulting distribution is adequately consistent with our target constraints and the recommendations of the National Academies of Sciences, Engineering, and Medicine (2017). We apply the retained parameter sets to FAIR to produce climate projections that represent climactic uncertainties and are further used in calculating the SCC uncertainty, as described in the next section. The interquartile range of the final SCC values across the entire distribution of parameter sets are shown in Table 3 in the main text.

Finally, we assess the reasonableness of the "handoff" between the SMME models, on which the damage function is estimated, and FAIR, with which future damages due to a pulse of CO_2 are calculated using the difference in temperature between the pulse and control runs. A comparison of climate sensitivity uncertainty across these two climate projections is important, as the climate sensitivity uncertainty captured in the empirically-based projections of mortality damages derives from the SMME, while the uncertainty we proliferate through to the SCC relies on the simple climate model FAIR. Figure G3 shows the distribution of GMST changes relative to 2001-2010 (ΔGMST) over time, according to the SMME (top row) and the simple climate model FAIR (bottom two rows). To ensure comparability, here and in damage function estimation we use smoothed values of the ΔGMST realizations from each SMME model, where smoothing is done using a 20-year centered moving average. SMME data are available until the year 2100; thus, the top two rows show a direct comparison between FAIR and the SMME models for these years, showing a strong amount of overlap in both RCP4.5 and RPC8.5 distributions of warming and indicating the handoff is reasonable (as

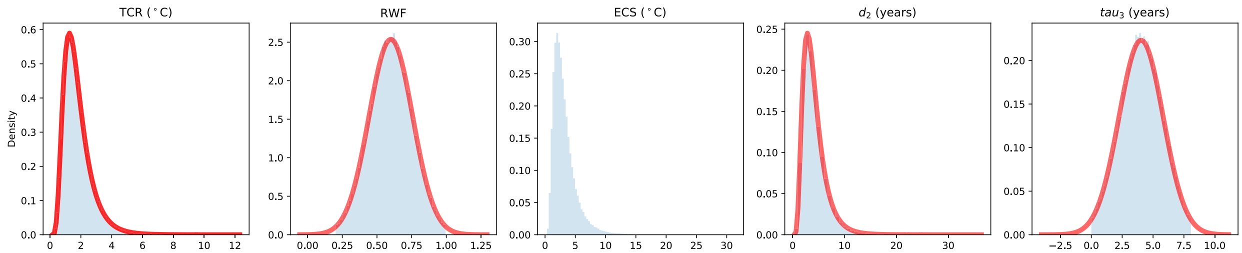

Figure G2: Distributions of key FAIR parameters for climate sensitivity uncertainty both before (red curve) and after (blue shading) applying constraints. Each panel indicates the distribution of a key parameter in the FAIR simple climate model, both before (in red) and after (in blue) the imposition of constraints described in the text. Distributions shown are: **A** transient climate response (TCR); **B** realized warming fraction (RWF) used to define ECS (=TCR / RWF); **C** equilibrium climate sensitivity (ECS) shown only after applying constraints due to unrealistic values in the initial distribution occurring as RWF$\rightarrow 0$; **D** short thermal adjustment time (d_2); **E** time scale of rapid carbon uptake by the ocean mixed layer (τ_3).

would be expected based on the construction of the SMME).

Parameter	*Distribution from literature*	*Pre-IPT distribution*	*Post-IPT distribution*	Distribution	*Source*
TCR (C)	[1.00, 2.50]	[1.00, 2.49]	[1.00, 2.50]	Lognormal	AR5
RWF	[0.45, 0.75]	[0.45, 0.75]	N/A	Normal	NAS (2017)
ECS (C)	[1.5, 4.5]	[1.60, 4.65]	[1.61, 4.61]	Lognormal	AR5
d_2 (years)	(1.6, 8.4)	(1.6, 8.4)	(1.6, 8.3)	Lognormal	Millar et al. (2017)
τ_3 (years)	Joos et al. (2013) point estimates	4.04 (1.07, 6.96)	4.04 (1.25, 6.79)	Normal	Joos et al. (2013)
Key metrics					
TCRE (C/TtC)	[0.8, 2.5]	N/A	[0.88, 2.34]	Normal	AR5
Time to T_{max} (years)	(6.6, 30.7)	(4.0, 100.0)*	(4.0, 100.0)*	N/A	Ricke and Caldeira (2014)
RCP 4.5 GMST					
2046 – 2065	1.4 [0.9, 2.0]	N/A	1.38 [0.73, 1.98] (0.51, 2.88)	Normal	AR5
2081 – 2100	1.8 [1.1, 2.6]	N/A	1.81 [0.93, 2.60] (0.65, 3.88)	Normal	AR5
2181 – 2200	2.3 [1.4, 3.1]	N/A	2.37 [1.13, 3.46] (0.78, 5.41)	Normal	AR5
2281 – 2300	2.5 [1.5, 3.5]	N/A	2.73 [1.24, 4.01] (0.85, 6.45)	Normal	AR5
RCP 8.5 GMST					
2046 – 2065	2.0 [1.4, 2.6]	N/A	2.05 [1.09, 2.90] (0.77, 4.20)	Normal	AR5
2081 – 2100	3.7 [2.6, 4.8]	N/A	3.71 [1.96, 5.31] (1.39, 7.73)	Normal	AR5
2181 – 2200	6.5 [3.3, 9.8]	N/A	7.34 [3.82, 10.60] (2.69, 15.35)	Normal	AR5
2281 – 2300	7.8 [3.0, 12.6]	N/A	8.86 [4.48, 12.84] (3.11, 18.84)	Normal	AR5

Table G1: Comparisons of the distributions of key FAIR parameter values. This table compares the distributions of key FAIR parameter values that pass the initial pulse-adjustment timescale (IPT) constraint against the relevant distributions from the literature (included in the IPT constraint is filtering of τ_3 and RWF as specified in the text). Distributions shown are: transient climate response (TCR); realized warming fraction (RWF); equilibrium climate sensitivity (ECS); short thermal adjustment time (d_2); time scale of rapid carbon uptake by the ocean mixed layer (τ_3); transient climate response to emissions (TCRE); and the change in global mean surface temperature (GMST) from the reference period 1986-2005 at various points in the projections. Note that RWF is only used to create our ECS distribution, and so the post-IPT distribution of RWF is not reported. Distributions reported are determined by the reference values from the literature, so that different parameters have different descriptions of their associated distributions: 5 to 95% ranges are given in ( ), 17 to 83% ranges (*likely* ranges for AR5) are given in [ ], and means are given without ( ) or [ ].

* We only consider the first 100 years post-pulse to be consistent with the length of the simulations in Ricke and Caldeira (2014).

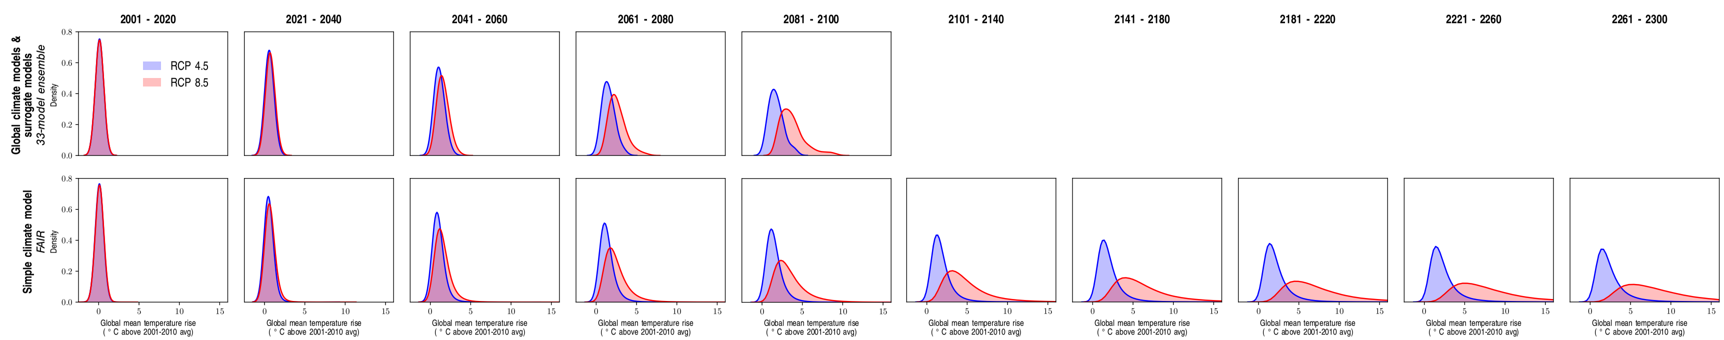

Figure G3: Distribution of changes in global mean surface temperature (ΔGMST) from an ensemble of global climate models and surrogate models (SMME) and from the simple climate model FAIR. Top row: Distribution of ΔGMST from 2001 to 2100, according to an ensemble of 33 GCMs and surrogate models that form the SMME. Second row: Distribution of ΔGMST from 2001 to 2300, according to 96,306 of simulation runs of the simple climate model FAIR.

G.3 Converting temperature scenarios to mortality partial SCC

We convert the temperature scenarios developed in the climate module into estimates of mortality-related damages using the global damage functions described in Section 7. These damage functions characterize valued mortality damages as a function of ΔGMST (changes in GMST relative to 2001-2010). Figure G4 shows these functions in 5-year time steps for each combination of valuation assumptions using the US EPA VSL (see Sections 5.5 and 7 for discussion of valuation of mortality-related costs of climate change). This figure contains the same information as Figure 8 in the main text, while additionally demonstrating substantial heterogeneity across distinct valuation scenarios (our primary valuation method uses an age-varying VSL in which impact region-specific VSLs are constructed using an income elasticity of one; this valuation is shown in the bottom row and second column of Figure G4) .

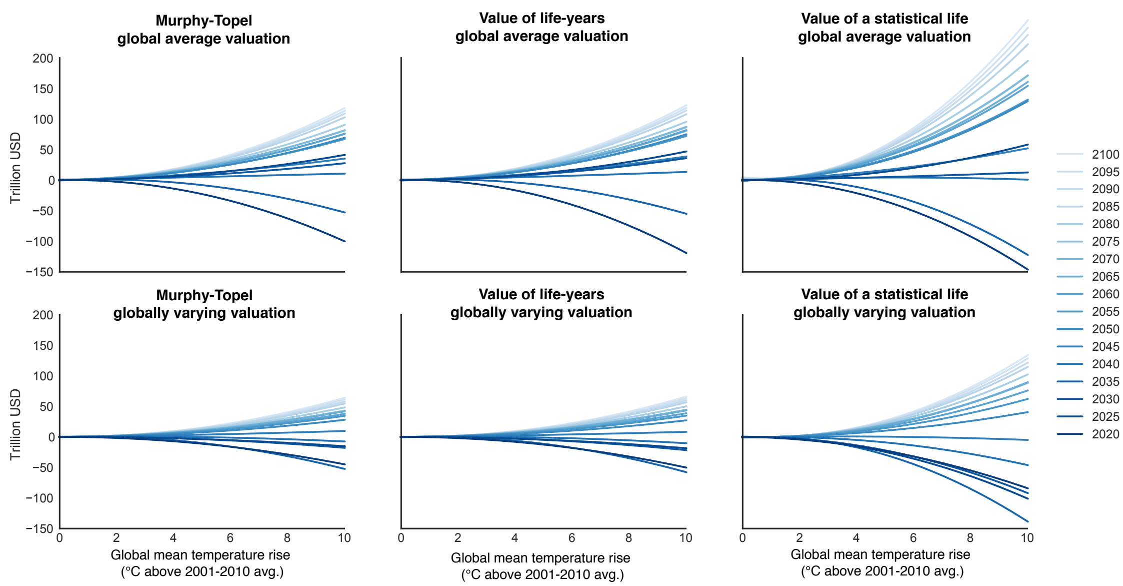

Figure G4: Temporal evolution of empirically derived damage functions (trillion USD) as a function of global mean surface temperature anomaly. Each panel shows estimates of quadratic damage functions estimated independently for each 5-year period from 2015 to 2100 under various valuation assumptions regarding the valuation of lives lost or saved.

The coefficients on these quadratic damage functions are constructed for each year from 2020 to 2300, as described in the main text. We then generate annual estimates of temperature-related mortality damages by applying the ΔGMST values from both the control FAIR scenarios (RCP4.5 and RCP8.5), as well as pulse scenarios, to the empirically derived damage functions. After computing mortality damages associated with each scenario, we subtract each pulse scenario from the corresponding control scenario and divide by the pulse amount to estimate the marginal effect of the pulse. This time series is then discounted using 2.0%, 2.5%, 3% and 5% discount rates, and summed through time to create a net present value, following Equation 13 in Section 7. This final value is the net present value of the full mortality risks caused of a marginal emission of CO_2. An alternative estimate would make use of Ramsey-like discounting, accounting for the relationship between consumption growth and the discount rate, but we leave this for future study.

Figure G5 replicates the SCC calculation graphically shown in Figure 9 for RCP 4.5.

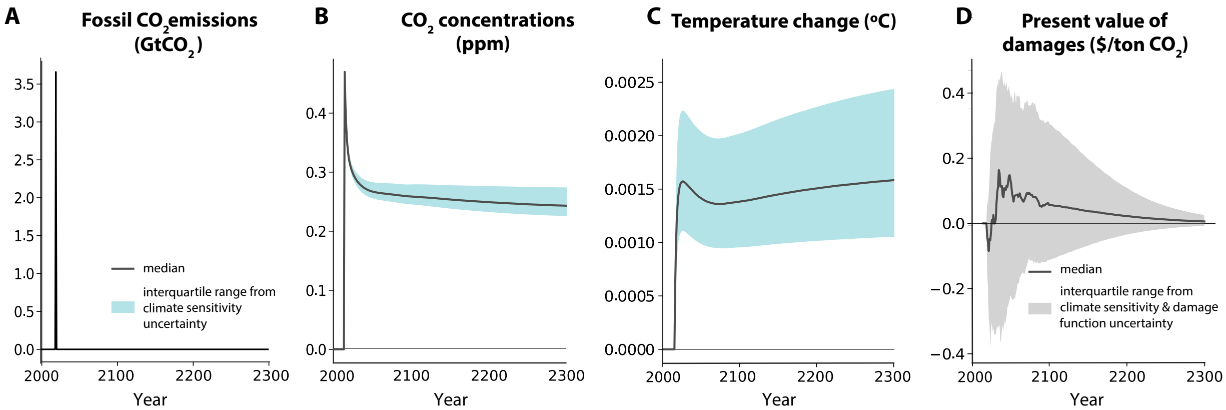

Figure G5: Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020 under RCP4.5. Panel A shows a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) in 2020 for emissions scenario RCP 4.5. Panel B displays the effect of this pulse on atmospheric CO_2 concentrations, relative to the baseline. In panel C, the impact of the pulse of CO_2 on temperature is shown where the levels are anomalies in global mean surface temperature (GMST) in Celsius. In panels A-C, shaded areas indicate the inter-quartile range due to climate sensitivity uncertainty, while solid lines are median estimates. Panel D shows the change in discounted damages over time due to a 1 Gt pulse of CO_2 in 2020, as estimated by our empirically-derived damage functions, using a 2% annual discount rate and the age-varying U.S. EPA VSL with an income elasticity of one applied to all impact regions. The shaded area indicates the inter-quartile range due to climate sensitivity and damage function uncertainty, while the solid line is the median estimate.

In the main text, we report uncertainty in the mortality partial SCC in three ways: accounting for climate sensitivity uncertainty only, damage function uncertainty only, and full uncertainty (both climate and economic). Here we briefly describe how these values are generated.

Mortality partial SCC estimates accounting for both climate sensitivity and damage function uncertainty: Using our Monte Carlo projections of damages, for each year from 2015 to 2100 we pool all Monte Carlo results for the associated 5-year window. We then run quantile regressions to fit quantile-specific damage functions for 19 quantiles (i.e., every 5^{th} percentile from the 5^{th} to 95^{th}). As in the mean damage function estimation, extrapolation past the year 2100 is accomplished using a time interaction model (see Section 7). In this extrapolation, we allow each quantile of the Monte Carlo distribution to evolve over time heterogeneously, based on the observed changes over time that we estimate at the end of the 21^{st} century.

We run each quantile-specific damage function through each of the 96,306 sets of FAIR parameters and up-weight runs in order to reflect probability mass in the damage function uncertainty space. This process reflects a joint sampling from the full space of damage function uncertainty and climate sensitivity uncertainty. The relevant SCC interquartile range (IQR) is resolved from the resulting distribution of mortality partial SCCs.

Mortality partial SCC estimates accounting for climate sensitivity uncertainty only: To isolate uncertainty in the mortality partial SCC that derives from climate sensitivity uncertainty, we run the central estimate of our damage function through each of the 96,306 sets of FAIR parameters. The corresponding SCC IQR is resolved from the resulting distribution of mortality partial SCCs.

Mortality partial SCC estimates accounting for damage function uncertainty only: To isolate uncertainty in the mortality partial SCC that derives from uncertainty in the damage function, we run the set of quantile-year damage functions through FAIR with each climate parameter fixed at its median value (as is done in the central mortality partial SCC estimates). The corresponding SCC IQR is resolved from the resulting distribution of mortality partial SCCs.

H Sensitivity of the mortality partial social cost of carbon

The mortality partial social cost of carbon (SCC) estimates shown in the main text depend upon a set of valuation and functional form assumptions and are reported for a particular socioeconomic scenario (SSP3). In this appendix, we detail our valuation approach and provide a wide range of additional mortality partial SCC estimates under alternative valuation approaches, alternative functional forms and extrapolation approaches for the damage function, and under multiple different socioeconomic scenarios. In all cases, we show multiple discount rates and emissions trajectories.

H.1 Methodology for constructing value of life-years lost from value of a statistical life (VSL)

As described in Section 7, panel A of Table 3 utilizes a valuation approach that adjusts the VSL by the total value of expected life-years lost. We provide this metric in order to accommodate the large heterogeneity in mortality-temperature relationships that we uncover across age groups. To adjust VSL values accordingly (see Table H1 for a set of commonly used VSLs), we first calculate the value of lost life-years by dividing the U.S. EPA VSL by the remaining life expectancy of the median-aged American. This recovers an implied value per life-year. We then apply an income elasticity of one[108] to convert this life-year valuation into a per life-year VSL for each impact region in each year. To calculate life-years lost for a given temperature-induced change in the mortality rate, we use the SSP projected population values, which are provided in 5-year age bins, to compute the implied conditional life expectancy for people in each age bin. We take the population-weighted average of remaining life expectancy across all the 5-year age bins in our broader age categories of <5, 5-64, and >64. This allows us to calculate total expected life-years lost, which we multiply by the impact-region specific VSL per life-year to calculate total damages.

Table H1: Value of statistical life estimates. VSL values are converted to 2019 USD using the Federal Reserve's US GDP Deflator.

	VSL (Millions USD)	
	Unadjusted	2019 Dollars
EPA ($2011)	$9.90	$10.95
Ashenfelter and Greenstone ($1997)	$1.54	$2.39
OECD (OECD Countries; $2005)		
Base	$3.00	$3.82
Range	$1.50 - $4.50	$1.91 - $5.73
OECD (EU27 Countries; $2005)		
Base	$3.60	$4.58
Range	$1.80 - $5.40	$2.29 - $6.88

[108] As noted in the main text, the EPA recommends VSL income elasticities of 0.7 and 1.1 (U.S. Environmental Protection Agency, 2016), while a review by Viscusi (2015) estimates an income-elasticity of the VSL of 1.1.

This procedure assumes that our estimated climate change driven deaths occur with uniform probability for all people within an age category. Without historical data containing information on age-specific mortality rates at higher resolution than our three age categories, or information on other chronic health conditions that may lower the life expectancy of individuals in each age group, we cannot empirically parameterize a more detailed life expectancy calculation. However, it is plausible that older individuals within an age category and those with chronic conditions are more likely to die due to extreme temperatures, which would imply that our mortality risk values, when computed using a value of life-years lost approach, are overstated. While we do not have the data sufficient to test this hypothesis, prior evidence from pollution-related mortality in the United States suggests this bias may be substantial (Deryugina et al., 2019).

The above methodology also values each life-year lost identically. In an alternative set of calculations (see results in Appendix H.2), we adjust the life-year values based on the age-specific value of remaining life derived by Murphy and Topel (2006). Murphy and Topel (2006) provide estimates of the value of remaining life for each age group. The authors do not estimate the level of the VSL, but instead provide age-specific values *relative* to a given population-wide VSL. We use these relative values of remaining life by age to adjust the U.S. EPA VSL, such that life-years lost are heterogeneously valued for each impact region in each year, by age. The resulting SCC calculations are shown in Tables H2 and H3.

H.2 Mortality partial social cost of carbon under alternative valuation approaches and socioeconomic scenarios

In the main text, mortality partial SCC values are shown using a combination of the US EPA VSL, an income elasticity of one, and valuation methods that value deaths using both an age-varying and an age-invariant value of a statistical life calculation (see Appendix H.1). This appendix shows a range of mortality partial SCC estimates under alternative VSL values, alternative assumptions about the role of income in valuation, with a life-years adjustment to the VSL that allows for age-specific values of remaining life, as derived by Murphy and Topel (2006), and under alternative socioeconomic scenarios.

Table H2 provides mortality partial SCC estimates across these distinct valuation approaches under the method shown in the main text Table 3, in which an income elasticity of one is used to adjust VSLs across the globe and over time. Table H3 provides mortality partial SCC estimates across distinct valuation approaches under a globally uniform valuation method in which a globally homogeneous VSL is used in each year, which evolves over time based on global income growth. Under this alternative, the lives of contemporaries are valued equally, regardless of their relative incomes. The method shown in the main text is most consistent with the revealed preference approach we use to estimate costs of adaptation, given that we observe how individuals make private tradeoffs between their own mortality risk and their own consumption (recall Equation 7). However, the latter approach might be preferred by policymakers interested in valuing reductions in mortality risk equally for all people globally, regardless of how individuals value their own mortality risk.

Table H2: Globally varying valuation: Estimates of a mortality partial Social Cost of Carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions have heterogeneous valuation, based on local income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix H.1. The first row of every valuation shows our estimated mortality partial SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] showing the influence of climate sensitivity and damage function uncertainty (see Appendix G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally varying valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	17.1	11.2	7.9	2.9	7.9	5.2	3.7	1.3
Climate sensitivity uncertainty	[8.3, 39.3]	[5.9, 24.1]	[4.4, 15.8]	[2.0, 4.3]	[3.9, 18.3]	[2.8, 11.2]	[2.1, 7.4]	[0.9, 2.0]
Damage function uncertainty	[-21.9, 50.8]	[-19.2, 32.1]	[-12.1, 26.6]	[-6.3, 12.0]	[-10.2, 23.7]	[-9.0, 15.0]	[-5.6, 12.4]	[-2.9, 5.6]
Full uncertainty	[-24.7, 53.6]	[-18.9, 36.0]	[-15.2, 26.3]	[-8.5, 11.5]	[-11.5, 25.0]	[-8.8, 16.8]	[-7.1, 12.2]	[-3.9, 5.3]
RCP 8.5	36.6	22.0	14.2	3.7	17.0	10.2	6.6	1.7
Climate sensitivity uncertainty	[18.8, 76.6]	[11.6, 45.2]	[7.7, 28.3]	[2.4, 6.2]	[8.7, 35.7]	[5.4, 21.0]	[3.6, 13.2]	[1.1, 2.9]
Damage function uncertainty	[-8.4, 74.2]	[-8.7, 48.2]	[-6.4, 35.6]	[-7.3, 14.1]	[-3.9, 34.6]	[-4.0, 22.4]	[-3.0, 16.6]	[-3.4, 6.6]
Full uncertainty	[-7.8, 73.0]	[-10.6, 46.8]	[-11.4, 32.9]	[-8.9, 13.0]	[-3.6, 34.0]	[-5.0, 21.8]	[-5.3, 15.3]	[-4.1, 6.1]
Panel B: Value of statistical life								
RCP 4.5	14.9	9.8	6.7	1.7	7.0	4.6	3.1	0.8
Climate sensitivity uncertainty	[2.4, 52.9]	[2.7, 30.4]	[2.5, 18.3]	[1.0, 2.1]	[1.1, 24.6]	[1.2, 14.2]	[1.2, 8.5]	[0.5, 1.0]
Damage function uncertainty	[-12.8, 44.1]	[-11.8, 33.1]	[-11.1, 25.6]	[-6.8, 12.6]	[-6.0, 20.5]	[-5.5, 15.4]	[-5.2, 11.9]	[-3.2, 5.9]
Full uncertainty	[-21.2, 63.5]	[-17.9, 43.5]	[-15.7, 32.1]	[-11.8, 14.7]	[-9.9, 29.6]	[-8.3, 20.3]	[-7.3, 15.0]	[-5.5, 6.9]
RCP 8.5	65.1	36.9	22.1	3.5	30.3	17.2	10.3	1.6
Climate sensitivity uncertainty	[30.0, 147.0]	[17.5, 82.3]	[10.8, 48.3]	[2.2, 5.6]	[14.0, 68.5]	[8.1, 38.3]	[5.0, 22.5]	[1.0, 2.6]
Damage function uncertainty	[18.4, 98.2]	[8.3, 63.1]	[2.3, 43.7]	[-7.0, 14.5]	[8.6, 45.7]	[3.9, 29.4]	[1.1, 20.3]	[-3.3, 6.7]
Full uncertainty	[3.0, 139.0]	[-2.4, 83.1]	[-5.6, 53.4]	[-9.3, 16.0]	[1.4, 64.7]	[-1.1, 38.7]	[-2.6, 24.9]	[-4.3, 7.5]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	17.5	11.6	8.3	3.1	8.1	5.4	3.9	1.5
Climate sensitivity uncertainty	[8.8, 39.6]	[6.3, 24.5]	[4.7, 16.3]	[2.1, 4.8]	[4.1, 18.4]	[2.9, 11.4]	[2.2, 7.6]	[1.0, 2.2]
Damage function uncertainty	[-16.4, 56.2]	[-16.6, 35.8]	[-12.2, 27.4]	[-6.0, 12.4]	[-7.7, 26.2]	[-7.7, 16.7]	[-5.7, 12.8]	[-2.8, 5.8]
Full uncertainty	[-25.3, 56.6]	[-19.3, 37.9]	[-15.6, 27.7]	[-8.6, 12.2]	[-11.8, 26.4]	[-9.0, 17.7]	[-7.3, 12.9]	[-4.0, 5.7]
RCP 8.5	36.3	22.0	14.3	4.0	16.9	10.3	6.7	1.9
Climate sensitivity uncertainty	[18.8, 75.5]	[11.7, 44.8]	[7.9, 28.3]	[2.6, 6.6]	[8.7, 35.2]	[5.5, 20.9]	[3.7, 13.2]	[1.2, 3.1]
Damage function uncertainty	[-8.2, 67.4]	[-7.5, 46.8]	[-8.3, 33.3]	[-5.5, 14.0]	[-3.8, 31.4]	[-3.5, 21.8]	[-3.9, 15.5]	[-2.6, 6.5]
Full uncertainty	[-8.0, 70.9]	[-11.0, 46.4]	[-11.6, 33.0]	[-8.8, 13.6]	[-3.7, 33.0]	[-5.1, 21.6]	[-5.4, 15.4]	[-4.1, 6.3]

Table H3: Globally uniform valuation: Estimates of a mortality partial Social Cost of Carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions are given the global median VSL, after scaling using income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix H.1. The first row of every valuation shows our estimated mortality partial SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] showing the influence of climate sensitivity and damage function uncertainty (see Appendix G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally uniform valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	37.5	26.4	19.9	9.0	17.5	12.3	9.3	4.2
Climate sensitivity uncertainty	[19.4, 82.2]	[14.5, 54.4]	[11.4, 38.7]	[5.8, 15.1]	[9.0, 38.3]	[6.8, 25.3]	[5.3, 18.0]	[2.7, 7.1]
Damage function uncertainty	[-15.7, 87.9]	[-10.9, 63.1]	[-11.4, 44.2]	[-2.4, 23.1]	[-7.3, 40.9]	[-5.1, 29.4]	[-5.3, 20.6]	[-1.1, 10.7]
Full uncertainty	[-13.3, 101.7]	[-10.2, 68.7]	[-8.4, 50.0]	[-5.0, 21.7]	[-6.2, 47.4]	[-4.8, 32.0]	[-3.9, 23.3]	[-2.3, 10.1]
RCP 8.5	72.3	46.3	32.0	11.5	33.6	21.6	14.9	5.3
Climate sensitivity uncertainty	[37.8, 149.0]	[24.9, 93.4]	[17.7, 62.8]	[7.0, 20.1]	[17.6, 69.4]	[11.6, 43.5]	[8.2, 29.2]	[3.2, 9.4]
Damage function uncertainty	[7.2, 127.1]	[4.3, 86.3]	[-1.9, 59.0]	[-4.8, 24.8]	[3.4, 59.2]	[2.0, 40.2]	[-0.9, 27.5]	[-2.2, 11.5]
Full uncertainty	[4.6, 141.1]	[-0.5, 92.1]	[-2.9, 64.6]	[-4.6, 24.9]	[2.2, 65.7]	[-0.2, 42.9]	[-1.4, 30.1]	[-2.1, 11.6]
Panel B: Value of statistical life								
RCP 4.5	46.2	33.7	25.9	11.9	21.5	15.7	12.1	5.5
Climate sensitivity uncertainty	[15.3, 134.1]	[14.0, 86.6]	[12.3, 60.2]	[7.2, 21.4]	[7.1, 62.4]	[6.5, 40.3]	[5.7, 28.0]	[3.4, 10.0]
Damage function uncertainty	[14.3, 102.0]	[12.9, 75.4]	[3.5, 56.9]	[-0.7, 26.9]	[6.6, 47.5]	[6.0, 35.1]	[1.6, 26.5]	[-0.3, 12.5]
Full uncertainty	[2.8, 148.2]	[-1.8, 98.6]	[-4.1, 71.0]	[-4.2, 30.2]	[1.3, 69.0]	[-0.8, 45.9]	[-1.9, 33.1]	[-2.0, 14.1]
RCP 8.5	143.9	87.5	57.5	17.6	67.0	40.8	26.8	8.2
Climate sensitivity uncertainty	[68.8, 317.6]	[43.1, 189.7]	[29.2, 121.7]	[10.1, 32.9]	[32.0, 147.9]	[20.1, 88.3]	[13.6, 56.7]	[4.7, 15.3]
Damage function uncertainty	[59.5, 197.8]	[38.1, 130.2]	[23.8, 94.9]	[3.1, 33.2]	[27.7, 92.1]	[17.7, 60.6]	[11.1, 44.2]	[1.5, 15.4]
Full uncertainty	[39.0, 287.0]	[21.8, 176.9]	[11.9, 117.4]	[-2.0, 37.9]	[18.2, 133.7]	[10.1, 82.4]	[5.5, 54.7]	[-1.0, 17.6]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	35.8	25.3	19.0	8.6	16.7	11.8	8.8	4.0
Climate sensitivity uncertainty	[18.4, 79.1]	[13.8, 52.1]	[10.9, 37.0]	[5.5, 14.4]	[8.6, 36.8]	[6.4, 24.3]	[5.1, 17.2]	[2.6, 6.7]
Damage function uncertainty	[-15.1, 90.6]	[-8.0, 61.1]	[-4.9, 46.8]	[-4.1, 21.5]	[-7.0, 42.2]	[-3.7, 28.4]	[-2.3, 21.8]	[-1.9, 10.0]
Full uncertainty	[-14.2, 99.9]	[-10.8, 67.2]	[-9.0, 48.8]	[-5.7, 21.3]	[-6.6, 46.5]	[-5.0, 31.3]	[-4.2, 22.7]	[-2.7, 9.9]
RCP 8.5	70.1	44.6	30.7	10.9	32.6	20.8	14.3	5.1
Climate sensitivity uncertainty	[36.6, 144.7]	[24.0, 90.0]	[17.0, 60.2]	[6.6, 19.1]	[17.1, 67.4]	[11.2, 41.9]	[7.9, 28.0]	[3.1, 8.9]
Damage function uncertainty	[7.0, 123.2]	[0.0, 79.4]	[-0.8, 59.5]	[-4.4, 22.5]	[3.3, 57.4]	[0.0, 37.0]	[-0.4, 27.7]	[-2.0, 10.5]
Full uncertainty	[3.7, 134.5]	[-0.8, 87.9]	[-2.7, 61.7]	[-5.2, 24.0]	[1.7, 62.7]	[-0.4, 40.9]	[-1.3, 28.7]	[-2.4, 11.2]

Table H4 shows mortality partial SCC estimates using a 1.5% discount rate, which more accurately reflects recent global capital markets than the discount rates shown in the main text (the average 10-year Treasury Inflation-Indexed Security value from 2003 to present is just 1.01% (Board of Governors of the US Federal Reserve System, 2020)).

Table H4: Estimates of a partial Social Cost of Carbon (SCC) for excess mortality risk incorporating the costs and benefits of adaptation, 1.5% discount rate. In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value. All regions thus have heterogeneous valuation, based on local income. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Point estimates rely on the median values of the four key input parameters into the climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] including both climate sensitivity uncertainty and damage function uncertainty (see Appendix G for details).

	Annual discount rate $\delta = 1.5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 US Dollars)	
Moderate emissions scenario (RCP 4.5)	28.5
Full uncertainty IQR	[-35.6, 88.5]
High emissions scenario (RCP 8.5)	66.4
Full uncertainty IQR	[-2.8, 126.5]
Panel B: Globally varying value of a statistical life (2019 US Dollars)	
Moderate emissions scenario (RCP 4.5)	24.6
Full uncertainty IQR	[-25.5, 102.9]
High emissions scenario (RCP 8.5)	123.9
Full uncertainty IQR	[13.7, 253.6]

Table H5 shows mortality partial SCC estimates under various socioeconomic projections (SSP3 is used throughout the main text; see Appendix B.3.2 for a discussion of this choice). We note that under SSP4 and a moderate emissions scenario (RCP4.5), the central estimate of the partial SCC is negative under all discount rates shown. While SSP4 shows global average increases in the full mortality risk of climate change by 2100 under both emissions scenarios (see Figure F5), the negative SCC is driven by different income and demographic changes projected under SSP4 relative to the other SSPs, both of which influence the valuation of lives lost. In particular, SSP4 projects that today's wealthy and relatively cold locations will experience dramatically higher future incomes, with much older populations, when compared to SSP2 or SSP3. This increase in income and rapid aging of the population leads to many lives saves in cold regions of the world as the climate warms, and each life is valued highly due to income growth raising the VSL (recall that we use an income elasticity of one for the VSL throughout the text). In contrast, SSP4 projects very low income growth in today's hot and poor locations, such that lives lost due to warming in these regions receive little value in this scenario. Note that with sufficiently high emissions (RCP8.5), heat-related deaths outweigh cold-related lives saved even in today's wealthy and relatively cold regions of the world, such that the partial SCC for SSP4 is no longer negative.

Table H5: Estimates of a mortality partial Social Cost of Carbon (SCC) under various socioeconomic projections. In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD. In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Each row shows, for a different SSP scenario, our estimated SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function.

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 USD)				
RCP 4.5				
SSP2	25.7	15.8	10.4	2.9
SSP3	17.1	11.2	7.9	2.9
SSP4	-14.5	-10.0	-7.5	-3.7
RCP 8.5				
SSP2	33.3	18.7	11.0	1.2
SSP3	36.6	22.0	14.2	3.7
SSP4	22.5	13.0	7.9	1.2
Panel B: Globally varying value of a statistical life (2019 USD)				
RCP 4.5				
SSP2	2.0	0.3	-0.9	-3.3
SSP3	14.9	9.8	6.7	1.7
SSP4	-64.3	-46.6	-36.1	-18.5
RCP 8.5				
SSP2	43.9	22.0	10.7	-2.5
SSP3	65.1	36.9	22.1	3.5
SSP4	23.1	8.6	1.2	-6.4

Finally, Table H6 shows mortality partial SCC estimates under both SSP2 (repeating values in Table H5) and a "hybrid" SSP designed to approximate a scenario in which climate change impacts on economic growth are endogenized. Throughout our main analysis, we treat income as exogenously given by the Shared Socioeconomic Pathways (SSPs). However, a growing literature indicates that the level and/or growth rate of income is influenced by temperature (e.g., Burke, Hsiang, and Miguel, 2015; Kalkuhl and Wenz, 2020). Following this literature and allowing income to respond to emissions could influence our mortality partial SCC estimates both by changing location-specific VSLs, and by changing income-driven adaptation in location-specific mortality-temperature relationships. While a full treatment of this topic is beyond the scope of this analysis, here we create a hybrid SSP that is constructed to approximate the impact of endogenous economic growth on the mortality partial SCC. Our analysis involves two steps.

First, we choose two scenarios from the three SSPs included in the main analysis (SSP2, SSP3, and SSP4) for which differences in income across SSPs for each quintile of the global income distribution approximately match the impacts of climate change from Burke, Hsiang, and Miguel (2015). To see this visually, Panel A of Figure H1 shows the estimated impacts of climate change on GDP per capita from Burke, Hsiang, and Miguel (2015), where impacts are shown for each quintile of the 2010 country-level income distribution. The level of these curves indicate the difference between incomes under climate change following RCP8.5 versus without climate change. Panel B of Figure H1 shows the difference between incomes under SSP2 versus under a hybrid SSP in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010. As can be seen by comparing across panels, the difference between SSP2 and our hybrid scenario closely approximates the estimated GDP per capita climate change impacts in Burke, Hsiang, and Miguel (2015).

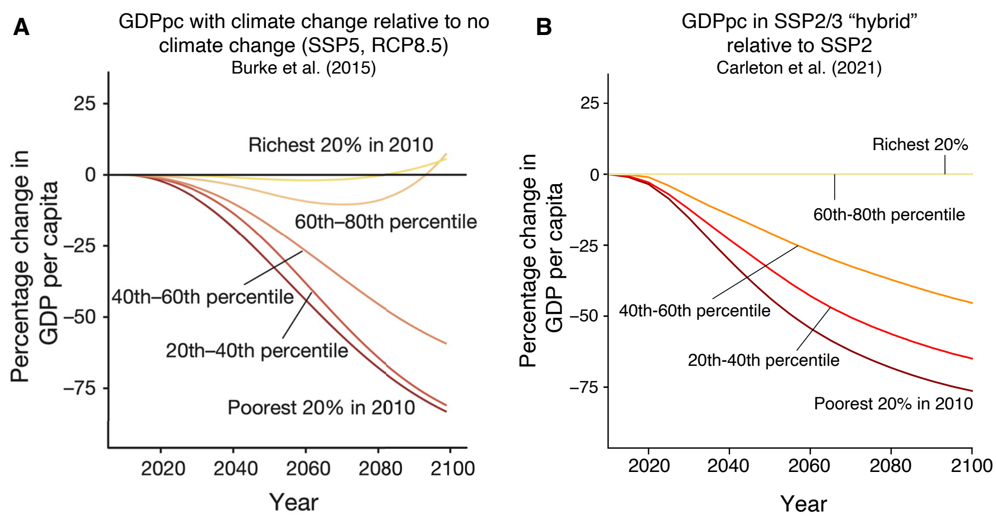

Figure H1: Constructing a hybrid Shared Socioeconomic Scenario (SSP) to approximate an income trajectory that is endogenous to climate change, following Burke, Hsiang, and Miguel (2015). Panel A is reproduced from Burke, Hsiang, and Miguel (2015) and shows mean impacts of climate change by 2010 income quantile for the authors' benchmark empirical model. The right panel shows the mean difference in income between SSP2 and a hybrid socioeconomic scenario in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010.

Second, we compute the mortality partial SCC under SSP2 as well as under our hybrid scenario, and compare SCC estimates. The difference in SCCs across these two scenarios approximates the effect of endogenizing income growth to climate change (as estimated by Burke, Hsiang, and Miguel (2015)) on the mortality partial SCC. Table H6 shows this comparison. Despite the extraordinary income differences shown in Figure H1, under our central valuation approach (δ=2%) and RCP8.5 emissions, the SCC rises by just 6% when using the hybrid socioeconomic scenario, relative to SSP2. Note that we do not report SCCs in Table H6 under RCP4.5, as the hybrid scenario was calibrated to match estimates from Burke, Hsiang, and Miguel (2015) for RCP8.5.

Table H6: Estimates of a partial Social Cost of Carbon (SCC) for excess mortality risk under a hybrid socioeconomic scenario designed to approximate an endogenous growth trajectory. Partial mortality SCC estimates are shown for a reference socioeconomic scenario (SSP2, "Reference"), as well as a hybrid scenario ("Hybrid") in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010. An income elasticity of one is used to scale the U.S. EPA VSL value and all estimates correspond to RCP8.5 emissions. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year. See text for details on the hybrid socioeconomic scenario.

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
SSP2 (Reference)	33.3	18.7	11.0	2.5
Hybrid SSP	35.3	20.6	12.7	1.2

H.3 Alternative approach to estimating post-2100 damages

As discussed in Section 7, we rely on an extrapolation of estimated damage functions to capture mortality impacts of climate change after the year 2100, due to data limitations. In this appendix, we explore the importance of this extrapolation by using an alternative approach to estimating post-2100 damage functions. Here, we calculate mortality partial SCC estimates using a set of damage functions in which the estimated 2100 damage function is applied to all years from 2100-2300. Effectively, this freezes the damage function at its 2100 level for all later years. Values shown are for SSP3, RCP8.5, with a discount rate of 2% and an age-varying VSL. Table H7 shows that this alternative approach to post-2100 damage estimation causes our central estimate of the SCC to fall by 21%.

Table H7: The influence of damage function extrapolation in years after 2100 on estimates of a mortality partial Social Cost of Carbon (SCC). In this table, an income elasticity of one is used to scale the U.S. EPA VSL value, and all SCC values are for the year 2020 under RCP8.5 emissions, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. The VSL is age-varying, so that these values are directly comparable to panel A in Table 3 in the main text. For the first column, damage functions continue to evolve over time in the years after 2100, according to the method described in Section 7. In the second column, the damage function estimated for the year 2100 is used for all years after 2100. All mortality partial SCC estimates use the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function.

	Extrapolating post-2100 damage function	*Holding post-2100 damage function fixed*
Pre-2100 damages	$12.8	$12.8
Post-2100 damages	$23.8	$16.0
Total damages	$36.6	$28.8

H.4 Robustness of the mortality partial SCC to an alternative functional form of the damage function

Throughout the main text, we report mortality partial SCC estimates that rely on a quadratic damage function estimated through all damage projections from all Monte Carlo simulation runs (see Section 7 for details). In Table H8, we show mortality partial SCC estimates for our central valuation approach using a cubic polynomial damage function in place of a quadratic. Across emissions scenarios and discount rates, we find that this alternative functional form has a minimal impact on mortality partial SCC estimates.

Table H8: Estimates of a mortality partial Social Cost of Carbon (SCC) using a cubic polynomial damage function In this table, an income elasticity of one is used to scale the U.S. EPA VSL value. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. Damage functions are estimated as a cubic polynomial, instead of a quadratic (as in the main text). In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Estimates rely on the median values of the four key input parameters into the simple climate model FAIR and a conditional mean estimate of the damage function.

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 USD)				
RCP 4.5	9.4	6.5	4.9	2.4
RCP 8.5	44.5	25.7	16.1	4.0
Panel B: Globally varying value of a statistical life (2019 USD)				
RCP 4.5	18.7	12.5	9.1	3.8
RCP 8.5	68.4	37.6	21.9	2.8

Appendix references

Ashenfelter, Orley and Michael Greenstone. 2004. "Using mandated speed limits to measure the value of a statistical life." *Journal of Political Economy* 112 (S1):S226–S267.

Auffhammer, Maximilian and Anin Aroonruengsawat. 2011. "Simulating the impacts of climate change, prices and population on California's residential electricity consumption." *Climatic Change* 109 (1):191–210. URL http://dx.doi.org/10.1007/s10584-011-0299-y.

Auffhammer, Maximilian, Solomon M Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using weather data and climate model output in economic analyses of climate change." *Review of Environmental Economics and Policy* 7 (2):181–198.

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):105–159. URL http://dx.doi.org/10.1086/684582.

Becker, Gary S. 2007. "Health as human capital: synthesis and extensions." *Oxford Economic Papers* 59 (3):379–410.

Board of Governors of the US Federal Reserve System. 2020. "10-Year Treasury Inflation-Indexed Security, Constant Maturity [DFII10]." Tech. rep., FRED, Federal Reserve Bank of St. Louis. URL https://fred.stlouisfed.org/series/DFII10.

Bright, E. A., P. R. Coleman, A. N. Rose, and M. L. Urban. 2012. "LandScan 2011." Digital dataset: web.ornl.gov/sci/landscan/index.shtml.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2017. "The unequal effects of weather and climate change: Evidence from mortality in India." *NBER Working paper* .

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

Center for Systemic Peace. 2020. "Polity5: Political Regime Characteristics and Transitions 1800-2018." URL http://www.systemicpeace.org/inscrdata.html.

Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li. 2013. "Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy." *Proceedings of the National Academy of Sciences* 110 (32):12936–12941.

Collins, Matthew, Reto Knutti et al. 2013. *Long-term Climate Change: Projections, Commitments and Irreversibility*, chap. 12. Intergovernmental Panel on Climate Change, 1029–1136. URL http://www.ipcc.ch/report/ar5/wg1/.

Curriero, Frank C, Karlyn S Heiner, Jonathan M Samet, Scott L Zeger, Lisa Strug, and Jonathan A Patz. 2002. "Temperature and mortality in 11 cities of the eastern United States." *American Journal of Epidemiology* 155 (1):80–87.

Davis, Lucas W and Paul J Gertler. 2015. "Contribution of air conditioning adoption to future energy use under global warming." *Proceedings of the National Academy of Sciences* 112 (19):5962–5967.

Dellink, Rob, Jean Chateau, Elisa Lanzi, and Bertrand Magné. 2015. "Long-term economic growth projections in the Shared Socioeconomic Pathways." *Global Environmental Change* .

Deryugina, Tatyana, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif. 2019. "The mortality and medical costs of air pollution: Evidence from changes in wind direction." *American Economic Review* 109 (12):4178–4219.

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 (October):152–185. URL http://www.nber.org/papers/w13178.

Deschênes, Olivier and Enrico Moretti. 2009. "Extreme weather events, mortality, and migration." *The Review of Economics and Statistics* 91 (4):659–681.

Eurostat. 2013. "EuroStat Regional Yearbook." Available at http://ec.europa.eu/eurostat/data/database.

Gasparrini, Antonio, Yuming Guo, Masahiro Hashizume, Eric Lavigne, Antonella Zanobetti, Joel Schwartz, Aurelio Tobias, Shilu Tong, Joacim Rocklöv, Bertil Forsberg et al. 2015. "Mortality risk attributable to high and low ambient temperature: A multicountry observational study." *The Lancet* 386 (9991):369–375.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez De Silanes, and Andrei Shleifer. 2014. "Growth in regions." *Journal of Economic Growth* 19 (3):259–309. URL https://ideas.repec.org/a/kap/jecgro/v19y2014i3p259-309.html.

Geoffroy, Olivier, David Saint-Martin, Dirk JL Olivié, Aurore Voldoire, Gilles Bellon, and Sophie Tytéca. 2013. "Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments." *Journal of Climate* 26 (6):1841–1857.

Glaeser, Edward L, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. "Do institutions cause growth?" *Journal of economic Growth* 9 (3):271–303.

Global Administrative Areas. 2012. "GADM database of Global Administrative Areas, version 2.0." *Accessed 25 December, 2016* URL www.gadm.org.

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

Guo, Yuming, Antonio Gasparrini, Ben Armstrong, Shanshan Li, Benjawan Tawatsupa, Aurelio Tobias, Eric Lavigne, Micheline de Sousa Zanotti Stagliorio Coelho, Michela Leone, Xiaochuan Pan et al. 2014. "Global variation in the effects of ambient temperature on mortality: a systematic evaluation." *Epidemiology (Cambridge, Mass.)* 25 (6):781.

Hancock, Peter A, Jennifer M Ross, and James L Szalma. 2007. "A meta-analysis of performance response under thermal stressors." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 49 (5):851–877.

Heutel, Garth, Nolan H Miller, and David Molitor. 2017. "Adaptation and the Mortality Effects of Temperature Across US Climate Regions." *National Bureau of Economic Research* .

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

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer et al. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345):1362–1369.

Hsiang, Solomon and Robert E Kopp. 2018. "An Economist's Guide to Climate Change Science." *Journal of Economic Perspectives* 32 (4):3–32.

Hsiang, Solomon M, Kyle C Meng, and Mark A Cane. 2011. "Civil conflicts are associated with the global climate." *Nature* 476 (7361):438.

IIASA Energy Program. 2016. "SSP Database, Version 1.1 [Data set]." Tech. rep., National Bureau of Economic Research. URL https://tntcat.iiasa.ac.at/SspDb.Accessed25December,2016.

Isen, Adam, Maya Rossin-Slater, and Reed Walker. 2017. "Relationship between season of birth, temperature exposure, and later life wellbeing." *Proceedings of the National Academy of Sciences* 114 (51):13447–13452.

Joos, Fortunat, Raphael Roth, JS Fuglestvedt, GP Peters, IG Enting, W von Bloh, V Brovkin, EJ Burke, M Eby, NR Edwards et al. 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):2793–2825.

Kahn, Matthew E. 2005. "The death toll from natural disasters: the role of income, geography, and institutions." *Review of Economics and Statistics* 87 (2):271–284.

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

Kopp, Robert E and Bryan K Mignone. 2012. "The US government's social cost of carbon estimates after their first two years: Pathways for improvement." *Working paper* .

La Porta, Rafael and Andrei Shleifer. 2014. "Informality and development." *Journal of Economic Perspectives* 28 (3):109–26.

Mastrandrea, Michael D, Christopher B Field, Thomas F Stocker, Ottmar Edenhofer, Kristie L Ebi, David J Frame, Hermann Held, Elmar Kriegler, Katharine J Mach, Patrick R Matschoss et al. 2010. "Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties." Tech. rep., Intergovernmental Panel on Climate Change.

Matsuura, Kenji and Cort Willmott. 2007. "Terrestrial Air Temperature and Precipitation: 1900-2006 Gridded Monthly Time Series, Version 1.01." *University of Delaware.* URL http://climate.geog.udel.edu/climate.

Meinshausen, M., S. C. B. Raper, and T. M. L. Wigley. 2011. "Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – Part 1: Model description and calibration." *Atmospheric Chemistry and Physics* 11 (4):1417–1456. URL http://www.atmos-chem-phys.net/11/1417/2011/.

Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw. 1994. "The impact of global warming on agriculture: A Ricardian analysis." *The American Economic Review* :753–771.

Millar, Richard J, Zebedee R Nicholls, Pierre Friedlingstein, and Myles R Allen. 2017. "A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions." *Atmospheric Chemistry and Physics* 17 (11):7213–7228.

Millar, Richard J., Alexander Otto, Piers M. Forster, Jason A. Lowe, William J. Ingram, and Myles R. Allen. 2015. "Model structure in observational constraints on transient climate response." *Climatic Change* 131 (2):199–211.

Ministry of Health, Chile. 2015. "Department of Statistics and Information." Available at http://www.deis.cl/bases-de-datos-defunciones/.

Ministry of Health in Brazil. 2019. "Mortality Information System (SIM)." Available at http://datasus.saude.gov.br/sistemas-e-aplicativos/eventos-v/sim-sistema-de-informacoes-de-mortalidade.

Mitchell, Timothy D. 2003. "Pattern Scaling: An Examination of the Accuracy of the Technique for Describing Future Climates." *Climatic Change* 60 (3):217–242. URL http://link.springer.com/article/10.1023/A%3A1026035305597.

Moore, Frances C and David B Lobell. 2014. "Adaptation potential of European agriculture in response to climate change." *Nature Climate Change* 4 (7):610.

Murphy, Kevin M and Robert H Topel. 2006. "The value of health and longevity." *Journal of Political Economy* 114 (5):871–904.

National Academies of Sciences, Engineering, and Medicine. 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.* Washington, DC: The National Academies Press. URL https://www.nap.edu/catalog/24651/valuing-climate-damages-updating-estimation-of-the-social-cost-of.

National Center for Atmospheric Research Staff (Eds). 2015. “Global surface temperatures: BEST: Berkeley Earth Surface Temperatures.” *The Climate Data Guide* .

National Institute for the Study of Demography (INED). 2019. “Vital Registration System.” Available at https://www.ined.fr/en/.

O’Neill, Brian C, Elmar Kriegler, Keywan Riahi, Kristie L Ebi, Stephane Hallegatte, Timothy R Carter, Ritu Mathur, and Detlef P van Vuuren. 2014. “A new scenario framework for climate change research: The concept of shared socioeconomic pathways.” *Climatic Change* 122 (3):387–400.

Organization of Economic Cooperaton and Development. 2020. “OECD.Stat.” URL https://doi.org/10.1787/data-00285-en.

Rasmussen, D. J., Malte Meinshausen, and Robert E. Kopp. 2016. “Probability-weighted ensembles of U.S. county-level climate projections for climate risk analysis.” *Journal of Applied Meteorology and Climatology* 55 (10):2301–2322. URL http://journals.ametsoc.org/doi/abs/10.1175/JAMC-D-15-0302.1.

Riahi, Keywan, Shilpa Rao, Volker Krey, Cheolhung Cho, Vadim Chirkov, Guenther Fischer, Georg Kindermann, Nebojsa Nakicenovic, and Peter Rafaj. 2011. “RCP 8.5—A scenario of comparatively high greenhouse gas emissions.” *Climatic Change* 109 (1-2):33–57.

Ricke, Katharine L and Ken Caldeira. 2014. “Maximum warming occurs about one decade after a carbon dioxide emission.” *Environmental Research Letters* 9 (12):124002.

Rohde, Robert, Richard Muller, Robert Jacobsen, Saul Perlmutter, Arthur Rosenfeld, Jonathan Wurtele, J Curry, Charlotte Wickham, and S Mosher. 2013. “Berkeley Earth temperature averaging process.” *Geoinfor Geostat: An Overview* 1 (2):1–13.

Samir, KC and Wolfgang Lutz. 2014. “The human core of the shared socioeconomic pathways: Population scenarios by age, sex and level of education for all countries to 2100.” *Global Environmental Change* .

Seppanen, Olli, William J Fisk, and QH Lei. 2006. “Effect of temperature on task performance in office environment.” *Lawrence Berkeley National Laboratory* 1 (LBNL-60946).

Sheffield, Justin, Gopi Goteti, and Eric F Wood. 2006. “Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling.” *Journal of Climate* 19 (13):3088–3111.

Taylor, Karl E, Ronald J Stouffer, and Gerald A Meehl. 2012. “An overview of CMIP5 and the experiment design.” *Bulletin of the American Meteorological Society* 93 (4):485.

Tebaldi, Claudia and Reto Knutti. 2007. “The use of the multi-model ensemble in probabilistic climate projections.” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 365 (1857):2053–2075. URL http://rsta.royalsocietypublishing.org/content/365/1857/2053.

Thrasher, Bridget, Edwin P Maurer, C McKellar, and PB Duffy. 2012. “Technical note: Bias correcting climate model simulated daily temperature extremes with quantile mapping.” *Hydrology and Earth System Sciences* 16 (9):3309–3314.

U.S. Environmental Protection Agency. 2016. “Recommended Income Elasticity and Income Growth Estimates: Technical Memorandum.” Tech. rep., U.S. Environmental Protection Agency Office of Air and Radiation and Office of Policy.

Van Vuuren, Detlef P, Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque et al. 2011. “The representative concentration pathways: An overview.” *Climatic Change* 109 (1-2):5.

Viscusi, W Kip. 2015. “The role of publication selection bias in estimates of the value of a statistical life.” *American Journal of Health Economics* .

WHO. 2018. "Global Health Estimates 2016: Deaths by cause, age, sex, by country and by region, 2000–2016." Tech. rep., World Health Organization.

Wood, Andrew W, Lai R Leung, V Sridhar, and DP Lettenmaier. 2004. "Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs." *Climatic Change* 62 (1-3):189–216.

World Bank. 2020. "World Development Indicators." .

World Inequality Lab. 2020. "World Inequality Database." URL https://wid.world/data/.