WORKING PAPER · NO. 2018-51

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

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

JULY 2019

Becker Friedman Institute
FOR ECONOMICS AT UCHICAGO

5757 S. University Ave.
Chicago, IL 60637
Main: 773.702.5599
bfi.uchicago.edu

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

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

[1]University of Chicago
[2]Rhodium Group
[3]University of California, Berkeley
[4]Rutgers University
[5]London School of Economics
[6]Columbia University
[7]NBER
[8]Analysis Group

July 31, 2019

Abstract

We develop empirically-grounded estimates of willingness-to-pay to avoid excess mortality risks caused by climate change. Using 40 countries' subnational data, we estimate a mortality-temperature relationship that enables global extrapolation to countries without data and projection of its future evolution, accounting for adaptation benefits. Further, we develop a revealed preference approach to recover unobserved adaptation costs. We combine these components with 33 high-resolution climate simulations, which produces a right-skewed distribution of global WTP with a mean of \$38.1 per tCO_2 under a high emissions scenario. Projections generally indicate increased mortality in today's poor locations and higher adaptation expenditures in rich ones.

*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, and Tata Center for Development. We thank Trinetta Chong, Greg Dobbels, Radhika Goyal, Simon Greenhill, Dylan Hogan, Azhar Hussain, Theodor Kulczycki, Sebastien Annan Phan, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and we thank 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, the National Academies of Sciences, and the Econometric Society for comments, suggestions, and help with data.

1 Introduction

Understanding the likely global economic impacts of climate change is of tremendous practical value to both policymakers and researchers. Some observers claim that the climate change risk is existential, while others argue that it is a relatively small risk among the many faced by the planet. These differences in perspective are reflected in global climate policy, which is at once both lenient and wildly inconsistent. At the same time, the economics literature has struggled both to provide empirically founded estimates of the economic damages from climate change and to provide quantitative insight into what climate change will mean at the local level for diverse populations.

Decades of study have accumulated numerous insights and important findings regarding the economics of climate change, both theoretically and empirically, but a fundamental gulf persists between the two main types of analyses pursued. On the one hand, there are models able to capture the global nature of problem, such as stylized integrated assessment models (IAMs) (e.g. Nordhaus, 1992; Tol, 1997; Stern, 2006), whose great appeal is that they provide an answer to the question of what the global costs of climate change will be. However, the many necessary assumptions of IAMs weakens the authority of these answers. On the other hand, there has been an explosion of highly resolved empirical analyses whose credibility lie in their observation of real world data and careful econometric measurement (e.g. Schlenker and Roberts, 2009; Deschênes and Greenstone, 2007) but whose analyses tend to be limited in scope and rely on short-run variation in weather that may not fully account for adaptation to gradual climate change. At its core, this dichotomy persists because researchers must trade off between being complete in scale and scope or investing heavily in data collection and analysis. The result is that no study has delivered estimated effects of climate change that are as comprehensive as those of IAMs, while simultaneously being grounded in detailed econometric analyses using high-resolution globally representative micro-data.

This paper aims to resolve the tension between these two approaches in the context of mortality risk due to climate change. Specifically, it strives to provide the scope and global scale of IAMs, but transparently built upon highly resolved econometric foundations. In so doing, it aims to account for both the benefits and costs of adaptation. The net result is that the paper produces a global estimate of the full costs of climate change on mortality risk, which we label an "excess mortality partial social cost of carbon". Additionally, we develop separate estimates of the changes in mortality risk for 24,378 regions that together account for the entire world; this spatial resolution marks a substantial improvement upon existing IAMs, which represent (at most) 16 heterogeneous global regions (Tol, 1997).

Our analysis has three main steps. First, we estimate the mortality-temperature relationship around the world, both today and into the future. This is accomplished by using the most exhaustive dataset ever collected on annual, sub-national mortality statistics. These data cover the universe of deaths from 41 countries totaling 55% of the global population at a resolution similar to that of U.S. counties (2^{nd}-administrative level) for each year across multiple age categories (i.e. <5, 5-64, and >64). These data allow us to estimate the mortality-temperature relationship with substantially greater resolution and coverage of the human population than previous studies; the most comprehensive econometric analyses to date have been for a few countries within a single region or individual cities from several countries. We find that, in our sample, an additional 35°C day, relative to a day at 20°C, increases

the annual all-age mortality rate by 0.4 deaths per 100,000 on average, while an additional -5°C day increases this rate by 0.3. That is, we estimate a U-shaped relationship between annual mortality rates and daily temperatures in this multi-country dataset.

These data also reveal heterogeneity in the mortality-temperature response function within each age category, reflecting the benefits of adaptation. Specifically, we allow the effect of temperature to vary as a function of the long-run climate (e.g. Deschênes and Greenstone, 2011; Hsiang and Narita, 2012; Barreca et al., 2015; Auffhammer, 2018) and income per capita (e.g. Burgess et al., 2017; Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015). These variables were carefully chosen based on the intersection of prior evidence from the literature, economic theory, and variables that are included in standard projections of the global economy used to develop physical climate model projections (O'Neill et al., 2014). We find that there is substantial heterogeneity in the mortality-temperature relationship: moving from the poorest to richest tercile in our sample saves on average 1.1 deaths per 100,000 per day at 35°C. Similarly, moving from the coldest to hottest tercile of long-run average temperature saves on average 1.4 deaths per 100,000 at 35°C.

A critical feature of the analysis is that the explicit modeling of heterogeneity additionally allows us to predict the mortality-temperature relationship across the entire world today, at high spatial resolution: we estimate location-specific response functions for each of 24,378 impact regions across the globe.[1] This exercise is possible, because there are measures of income and climate for the parts of the world where micro-level mortality data do not exist; we can therefore "fill in" missing information regarding sensitivity to climate for the 45% of people for whom mortality data are unavailable. A striking finding is that globally, the effect of an additional 35°C day (relative to a location-specific minimum mortality temperature) for the vulnerable over 64 population is 10.1 deaths per 100,000, which is approximately 30% larger than the effect from the regions of the world where data are currently available. This result underlines that our understanding of the consequences of climate change is highly dependent on which regions of the world have data and suggests that available estimates may underestimate climate change impacts because they disproportionately rely on data from wealthy economies and temperate climates.

We then project the mortality-temperature relationship for each of these impact regions out to 2100, based on standard projections of the evolution of income, population, and climate. This is a substantial step forward from assuming that response functions are constant over time, as has been the norm in much of the literature to date (e.g. Deschênes and Greenstone, 2011). A key finding from this exercise is that the mortality consequences of an additional hot day decline substantially by the end of the century, due to both the protective effects of higher incomes and costly adaptations that individuals are predicted to undertake in response to warmer climates. Put plainly, climate adaptation and income growth have substantial benefits.

The second feature of our analysis is the development of a general revealed preference method capable of estimating the full adaptation costs that populations will incur to obtain the adaptation benefits that we project, even though adaptation costs cannot be directly observed. This is a critical step because a full accounting of the economic burden of warming must account for the opportunity costs of all resources used to achieve these reduced sensitivities to temperature through adaptive adjustments,

[1] In the U.S., these impact regions map roughly into a county.

in addition to the direct mortality impacts. Yet, previous research has been unable to empirically measure all of these costs directly, because the range of potential responses to warming—whether defensive investments (e.g. building cooling centers) or compensatory behaviors (e.g. exercising earlier in the morning)—is enormous, making enumeration of their costs extraordinarily challenging. Indeed, the previous literature has frequently noted that adaptation will involve costs, and occasionally produced partial costs estimates (e.g. Deschênes and Greenstone, 2011; Barreca et al., 2016), or estimates of total impacts net of costs (e.g. Schlenker, Roberts, and Lobell, 2013; Deryugina and Hsiang, 2017), but has made little progress in building a comprehensive empirical measure of the costs of adaptation.

The basis for the revealed preference approach comes from the fact that we can measure the benefits of adaptation in terms of reduced mortality sensitivities to temperature combined with the theoretical restriction that individuals will only make adaptation investments when their costs are less than or equal to their benefits. The examples of Seattle, WA and Houston, TX, which have similar income levels, institutions, and other factors, but have very different climates, provide intuition for our approach. On average in the historical record, Seattle has just 1 day per year where the average temperature exceeds 30°C, while Houston experiences just over 8 of these days annually.[2] Houston has adapted to this hotter climate; we estimate that a day at 35°C (relative to a day at 20°C) increases the annual all-age mortality rate by 1.3 *more* deaths per 100,000 in Seattle than it does in Houston. If these outcomes are the result of revealed preferences, then it must be the case that the costs required to achieve Houston-like temperature sensitivity exceed the benefits that Seattle would receive from adopting similar practices, which seems sensible since these extreme temperatures only occur $\frac{1}{8}$ as often. Further, the costs of adapting in Houston, compared to enduring Seattle's high temperature-sensitivity, must be less than or equal to the reduction in mortality that adaptation provides. Indeed, the difference in air conditioning penetration rates, which were 27% in Washington state and 100% in Texas as of 2000-2004 (Barreca et al., 2016), indicates that the observed differences in temperature sensitivities between these cities reflect cost-benefit decisions.

We leverage this intuition to build a formal theoretical framework that allows us to derive empirically implementable estimates of adaptation costs, even though they are not directly observable.[3] Thus, in our projections of the future, we are able to simultaneously estimate how populations will reduce direct mortality from the climate associated with their projected patterns of adaptation, while also tracking the costs incurred in order to achieve these adaptation benefits. Importantly, our approach allows for an arbitrarily large number of unknown adaptive adjustments, and it accounts for the possibility that some adaptations generate consumption value that is independent of their mortality benefits (e.g. the consumption value of air conditioning).

Together, these two features of the analysis allow us to develop measures of the full mortality-related costs of climate change for the entire world, reflecting both the direct mortality costs (accounting for adaptation) and all adaptation costs. We find that the mean estimate of the total mortality burden of climate change across 33 different climate models and accounting for statistical uncertainty is projected to be worth 85 death equivalents per 100,000 at the end of the century under a high emissions scenario,

[2]These values are calculated over the years 1968-2010.

[3]We note that Schlenker, Roberts, and Lobell (2013), Guo and Costello (2013) and Deryugina and Hsiang (2017) exploit similar arguments regarding the equality of marginal adaptation costs and marginal adaptation benefits, and we describe in detail in Section 2 how our approach relates to these earlier contributions.

with an interquartile range of [16, 121]. For reference, all cancers are responsible for approximately 125 deaths per 100,000 globally today (WHO, 2018). Further, failing to account for income and climate adaptation, as has been the norm in the literature, would overstate the mortality costs of climate change by a factor of about 2.6, on average.

Finally, we note that there is evidence of substantial heterogeneity in impacts around the globe. For example, at the end of the century we project an annual average increase of about 160 death equivalents per 100,000 population in Accra, Ghana, which amounts to an ∼19% increase over Ghana's average mortality rate today. In contrast, at the end of the century we project an annual decrease of about 230 per 100,000 in Oslo, Norway, which amounts to a ∼28% decrease in Norway's average mortality rate today (United Nations, 2017). These total impacts are derived both from physical deaths and from adaptation costs measured in death equivalents; on average, about 14% of the total burden of climate change is borne as adaptation costs, although the relative importance of each component differs greatly across the world.

The third step is the development of a transparent procedure to transform the econometrically-based results into a single number describing the full welfare cost of additional mortality risk imposed by the marginal emission of CO_2 today. Specifically, we estimate the full willingness-to-pay (WTP) to avoid the alteration of mortality risk associated with the release of an additional metric ton of CO_2, which is referred to as the excess mortality "partial" Social Cost of Carbon (a "full" SCC would encompass impacts across all affected sectors). In conducting this WTP calculation, we construct empirical global "damage functions," each of which describes the costs of excess mortality risk in a given year as a function of the overall level of global climate change (Nordhaus, 1992). We find that damages are nonlinear in the magnitude of warming, reflecting U-shaped relationships between mortality and temperature recovered in our econometric estimation. We combine these damage functions with a climate model at the scientific frontier to provide what are, to our knowledge, the first empirically-derived estimates of global WTP to avoid climate change for a single sector of the economy.

Our central values suggest that with a 3% discount rate, the excess mortality partial SCC is roughly $38.1 (in 2019 USD) with a high emissions scenario (i.e., Representative Concentration Pathway (RCP) 8.5, in which CO_2 emissions growth is sustained). In contrast, this value falls to $4.4 with a moderate emissions scenario (i.e., Representative Concentration Pathway (RCP) 4.5, in which CO_2 emissions are stable through 2050 and then decline), due to the nonlinearity of estimated damage functions. When accounting for climate model and statistical uncertainty, the respective interquartile ranges are [-$7.6, $87.4] for RCP8.5 and [-$26.9, $40.1] for RCP4.5.

The rest of this paper is organized as follows: Section 2 outlines a conceptual framework for the problem of projecting climate damages into the future, accounting for adaptation and its cost; Section 3 describes the data used in the estimation of impacts and in the climate change projected impacts; Section 4 details the econometric approach; Section 5 describes the results of the econometric analysis; Section 6 explains how we extrapolate mortality impacts across space and project them over time while computing adaptation costs and benefits; Section 7 presents global results from projections that use high-resolution global climate models; Section 8 details the calculation of a damage function based on these results and combines it with a simple climate model to compute a mortality partial SCC; Section 9 discusses limitations of the analysis; and Section 10 concludes.

2 Accounting for costs and benefits of adaptation to climate change empirically

Climate change is projected to have a wide variety of impacts on well-being, including altering the risk of mortality due to extreme temperatures. The ultimate effect on particular outcomes like mortality rates will be determined by the adaptations that are undertaken. Specifically, as the climate changes, individuals and societies will weigh the costs and benefits of undertaking actions that allow them to exploit new opportunities (e.g., converting land to new uses) and protect themselves against new risks (e.g., investments in air conditioning to mitigate mortality risks). The full cost of climate change will thus reflect both the realized *direct impacts* (e.g. changes in mortality rates), which depend on the benefits of these adaptations, and *the costs of these adaptations in terms of foregone consumption.* However, to date it has proven challenging to develop a theoretically founded and empirically credible approach to explicitly recover the full costs of climate change.[4]

This section develops an approach for empirically estimating populations' willingness to pay (WTP) to avoid the mortality risks from climate change that reflects *both* the costs and benefits of adaptation. Previous work has established methods to estimate the benefits of adaptation (e.g. Auffhammer, 2018) but not its costs. Thus, our key contribution is to devise an empirical strategy to construct adaptation cost estimates using a revealed preference approach, even when numerous individual margins of adaptation and their costs cannot be directly observed.

2.1 Definitions and intuition for the economics of adaptation

We define the *climate* of a location as the joint probability distribution over a vector of possible conditions that can be expected to occur at a given location over a specific interval of time. Let $\boldsymbol{C}$ be a vector of parameters describing the entire joint probability distribution over all relevant climatic variables.[5] For example, $\boldsymbol{C}$ might contain the mean and variance of daily average temperature and rainfall, among other parameters. Weather realizations are a random vector $\boldsymbol{c}$, describing the empirical distribution of the same parameters in $\boldsymbol{C}$ taken over some finite period of time, such as a year. Mortality risk is a function of both $\boldsymbol{c}$ and a vector of K endogenous economic variables $\boldsymbol{b} = \{b_1, ..., b_K\}$. The vector $\boldsymbol{b}$ captures all choice variables available to individuals, except consumption of a numeraire good x, including possible adaptive behaviors and investments that could interact with individuals' exposure to a warming climate, such as installation of air conditioning and time allocated to indoor activities. Mortality risk is then captured by the probability of death during a unit interval of time $f = f(\boldsymbol{b}, \boldsymbol{c})$.

Climate change will influence mortality risk through two pathways.[6] 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 $\boldsymbol{b}$ will capture all long-run adaptation to the climate (e.g. Mendelsohn, Nordhaus, and Shaw, 1994; Kelly, Kolstad, and Mitchell, 2005), since these adaptations

[4] See Deschênes and Greenstone (2011); Hsiang and Narita (2012); Schlenker, Roberts, and Lobell (2013); Lobell et al. (2014); Guo and Costello (2013); Deschênes, Greenstone, and Shapiro (2017); Deryugina and Hsiang (2017) for different discussions of this issue and some of the empirical challenges.

[5] See Hsiang (2016) and Deryugina and Hsiang (2017) for a more complete discussion of this approach to measuring climate with a finite vector.

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

are necessarily factor reallocations based on knowledge of the climate (Deryugina and Hsiang, 2017). Therefore, since the climate $\boldsymbol{C}$ determines both $\boldsymbol{c}$ and $\boldsymbol{b}$, for notational simplicity we rewrite the probability of death at initial climate $\boldsymbol{C}_1$ as:

$$\Pr(\text{death} \mid \boldsymbol{C}_1) = f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) \tag{1}$$

where we define $\boldsymbol{c}(\boldsymbol{C})$ to be a random vector $\boldsymbol{c}$ drawn from a distribution characterized by $\boldsymbol{C}$ (Hsiang, 2016). Equation 1 describes the structure through which climate and compensatory investments, captured in $\boldsymbol{b}$, enter the health production function (Grossman, 1972).

Many previous empirical estimates of the effects of climate assume no adaptation takes place (e.g. Deschênes and Greenstone, 2007; Houser et al., 2015). Specifically, these approaches calculate changes in an outcome variable imposed by changing the distribution of $\boldsymbol{c}$, assuming economic decisions embodied by $\boldsymbol{b}$ do not change. Following this approach, the change in mortality risk incurred due to a change in climate from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ would be calculated as:

$$\textit{mortality effects of climate change without adaptation} = f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_2)) - f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) \tag{2}$$

which ignores the fact that individuals will respond to the change in climate by altering $\boldsymbol{b}$. Thus, Equation 2 is equivalent to a partial derivative of mortality risk with respect to climate, as individuals are not allowed to make any adjustments in response to the changing climate.

In reality, optimizing populations will update their behaviors and technologies $\boldsymbol{b}$ in response to a changing climate as their beliefs about $\boldsymbol{C}$ evolve. To the extent that populations may attenuate mortality damages by adjusting $\boldsymbol{b}$ in response to climate change, these adaptations will generate benefits by countering the effect of climate on mortality. Thus, a more realistic estimate for the change in mortality due to a change in climate is:

$$\textit{mortality effects of climate change with adaptation} = f(\boldsymbol{b}(\boldsymbol{C}_2), \boldsymbol{c}(\boldsymbol{C}_2)) - f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) \tag{3}$$

Note that if the climate is changing such that the mortality risk from $\boldsymbol{C}_2$ is higher than $\boldsymbol{C}_1$ 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.[7] Stated another way, the change in Equation 3 must be weakly smaller than the change in Equation 2. Thus, Equation 3 is equivalent to a total derivative of mortality risk with respect to climate, as it reflects individuals' compensatory responses. The costs of these adaptations are the costs of adjustment to $\boldsymbol{b}$, and their benefits are the resulting weakly smaller changes in mortality.

Several analyses have estimated reduced-form versions of Equation 3, confirming that accounting for endogenous changes to technology, behavior, and investment may mitigate the direct effects of climate in a variety of contexts (e.g. Barreca et al., 2016).[8] 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

[7] If the change in climate decreases mortality risk when holding $\boldsymbol{b}$ fixed, then we would expect adaptation to lead to larger changes in mortality (i.e. smaller reductions) than indicated by Equation 2.

[8] 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); Burgess et al. (2017); Auffhammer (2018).

for $\boldsymbol{b}$ so that all comparable populations would be inoculated against the climate. But we do not observe this to be true: previous studies that demonstrate benefits of adaptation do so by demonstrating reduced sensitivity to marginal environmental changes ($\frac{\partial f}{\partial \boldsymbol{c}}$) in more adverse climates (e.g. Houston) and larger sensitivity in less adverse climates (e.g. Seattle) (Deschênes and Greenstone, 2011).[9] Thus, observed cross-sectional heterogeneity in climate sensitivity can be reconciled if individuals who are otherwise comparable face differential unobserved costs of adaptation. 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 adaptive reactions to these changes but also their cost. Thus, the total cost of changing mortality risks that result from a climate change $\boldsymbol{C}_1 \rightarrow \boldsymbol{C}_2$ is

full value of mortality risk due to climate change =

$$VSL \underbrace{[f(\boldsymbol{b}(\boldsymbol{C}_2), \boldsymbol{c}(\boldsymbol{C}_2)) - f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1))]}_{\text{observable change in mortality}} + \underbrace{A(\boldsymbol{b}(\boldsymbol{C}_2)) - A(\boldsymbol{b}(\boldsymbol{C}_1))}_{\text{adaptation costs}} \tag{4}$$

where VSL is the value of a statistical life. If the costs of adaptation $A(\boldsymbol{b})$ were omitted from this calculation, we might substantially underestimate the overall economic burden of warming.

A key objective of this paper is to measure the total costs of climate change's impact on mortality risk, consistent with Equation 4. The next subsection develops a model of optimizing behavior that leads to an expression for adaptation costs associated with the mortality risk from climate change. Importantly, this expression is composed solely of empirically identifiable elements that are in principle observable, despite the fact that the full range of K potential adaptations is unlikely to be observable.

2.2 Deriving an expression for adaptation costs from individual optimizing behavior

Several steps are necessary to derive an expression for total adaptation costs. First, we assume agents have rational expectations and can integrate over the distribution of realizations of $\boldsymbol{c}$ for each climate to compute the expected probability of death $\tilde{f}$ conditional on the climate and their actions: $\tilde{f}(\boldsymbol{b}(\boldsymbol{C}), \boldsymbol{C}) = \mathbb{E}_{\boldsymbol{c}}[f(\boldsymbol{b}(\boldsymbol{C}), \boldsymbol{c}(\boldsymbol{C})) \mid \boldsymbol{C}]$. Thus $1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})$ is the expected probability of survival, analogous to a health production function in other settings (Grossman, 1972); we assume that $\tilde{f}(\cdot)$ is continuous and differentiable. Second, we assume there exists some numeraire good x for which utility $u(x, \boldsymbol{b})$ is quasilinear, and that agents 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).[10] Third, we assume that markets clear for all technologies and investments represented by the choice vector $\boldsymbol{b}$, as well as for the numeraire good x, and that all choices $\boldsymbol{b}$ and x can be treated as continuous. Fourth, we assume that there is a competitive and frictionless rental

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

[10] For simplicity, here we assume that agents do not receive utility directly from the climate. However, it is straightforward to adapt the following derivation to allow $u(\cdot)$ to be a function of the climate. We discuss the minor implications of this addition below.

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 $\boldsymbol{b}$, . Because agents have accurate expectations over their current and future climates and access to a competitive rental market, markets will clear efficiently in each period as the climate warms. We assume markets clear quickly relative to warming, a reasonable assumption given the extremely slow pace of climate change (surface temperature rises just 0.0001°C per day on average under a high emissions scenario).

With this setup, agents jointly solve for all K margins of adaptation described in the choice vector $\boldsymbol{b}$ such that $\boldsymbol{b}^* = \arg\max u(x, \boldsymbol{b})[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, subject to a budget constraint $h(\boldsymbol{b}) + x = Y$, where $h(\boldsymbol{b})$ is the pecuniary cost of adaptation expended by households and Y is exogenous income. Thus, summing across $K+1$ first order conditions and substituting for the value of the statistical life (VSL), $\boldsymbol{b}^*$ satisfies:[11]

$$-VSL \sum_k \underbrace{\frac{\partial}{\partial b_k}\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}_{\substack{\text{marginal survival} \\ \text{benefit of } b_k}} = \underbrace{\sum_k \frac{\partial}{\partial b_k}\left[\underbrace{h(\boldsymbol{b}^*)}_{\substack{\text{pecuniary} \\ \text{costs of } \boldsymbol{b}}} - \underbrace{u(x^*, \boldsymbol{b}^*)\frac{1}{\partial u(x^*, \boldsymbol{b}^*)/\partial x}}_{\text{non-mortality benefits of } \boldsymbol{b}} \right]}_{\text{net marginal cost of } b_k = \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k}} = \sum_k \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k}. \quad (5)$$

Equation 5 governs expenditures on adaptation. Its lefthand side is the product of the sum of marginal changes to expected mortality risk due to any adjustments to $\boldsymbol{b}$ and the negative of the VSL, so it represents the total expected marginal benefit of adjusting $\boldsymbol{b}$ through its effect on mortality risk.

The righthand side of Equation 5 has two parts. The first term inside the brackets, $h(\boldsymbol{b}^*)$, represents all pecuniary expenditures required to achieve $\boldsymbol{b}^*$, such as spending on units of air conditioning. The second term in the brackets represents [minus] the dollar value of all non-mortality utility benefits and costs derived from $\boldsymbol{b}^*$, such as the utility of enjoying air conditioning or the disutility of exercising at midnight to avoid the daytime heat (expressed in dollars of WTP by dividing through by the marginal utility of consumption, $\frac{\partial u(\cdot)}{\partial x}$). The derivative of the sum of these two terms with respect to each element of $\boldsymbol{b}$ can thus be interpreted as the *net* marginal cost of each adaptive action, since all non-mortality marginal benefits and costs are differenced out.

We denote these net adaptation costs as $A(\boldsymbol{b})$, just as in the previous subsection, where we define $A(\boldsymbol{b})$ over the space $\boldsymbol{b}$ such that its derivative with respect to each b_k is equal to $\frac{\partial h(\boldsymbol{b}^*)}{\partial b_k} - \frac{\partial u(x^*, \boldsymbol{b}^*)/\partial b_k}{\partial u(x^*, \boldsymbol{b}^*)/\partial x}$.[12] For our purposes, it is neither important nor empirically feasible to separately identify the two components of net cost, and throughout our remaining analysis adaptation costs should be interpreted as net of any direct non-pecuniary benefits.[13] Equation 5 is therefore further simplified by substituting $A(\boldsymbol{b})$, showing that in equilibrium the total expected marginal benefits of adaptation (left-hand side of Equation 5), in terms of mortality risk, equal the total marginal net costs (far-right-hand side). Thus, agents will adjust every dimension of $\boldsymbol{b}^*$ in response to their climate $\boldsymbol{C}$ until their marginal costs equal their expected marginal benefits as adaptive strategies.

The marginal costs and benefits of adaptation described in Equation 5 are not directly observable.

[11] See Appendix A for derivation.

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

[13] In a variant of this model in which agents derive utility directly from the climate, the interpretation of net costs $A(\cdot)$ is modified to include an additional component representing changes in utility derived directly from the changing climate. However, this change of interpretation to include an additional term that is "netted out" in $A(\cdot)$ is the only implication of adding climate directly to the utility function.

In particular, there undoubtedly exists an enormous number of adaptive margins K. One strategy in the literature has been to construct partial cost estimates (A_k) that account for specific observable margins of adaptation (b_k) that can be easily priced. For example, Deschênes and Greenstone (2011) estimate the cost of increasing U.S. residential energy consumption used for cooling as an adaptation strategy to reduce heat-related mortality in homes. In principle, a large number of such studies that each measure the costs of specific margins of adaptation could be assembled to construct the total cost of all margins of adaptation $A(\boldsymbol{b}) = A_1(b_1) + ... + A_K(b_K)$. However, in practice it seems unlikely to be feasible to enumerate and identify the cost of all possible margins of adaptation, especially when numerous possible margins may not even be known. Thus, the "enumerative" approach to determining the total costs of adaptation seems infeasible.

It is possible to make the expression in Equation 5 of greater practical value if we eliminate $\boldsymbol{b}$ by noting that it is an implicit function of climate and income, $\boldsymbol{b}^* = \boldsymbol{b}^*(\boldsymbol{C}, Y)$.[14] This allows us to collapse all the information encoded in the large unobserved vector $\boldsymbol{b}$ into these observable terms. Thus, we take the total derivative of $\tilde{f}(\cdot)$ with respect to the climate $\boldsymbol{C}$, applying the Chain Rule to each element b_k, and obtain:

$$\frac{d\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{d\boldsymbol{C}} = \sum_k \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k} \frac{\partial b_k^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial \boldsymbol{C}} \tag{6}$$

which says that the total derivative, or the effect of climate on mortality once individuals have adapted to a change in climate, is the sum of two terms. The first term represents the expected impacts on mortality of all changes in adaptive investments induced by the change in climate; as discussed, this is of limited practical value because of data and estimation limitations.[15] The second term is the effect that the climate would have if individuals were prevented from adapting (i.e. the partial derivative).[16] For example, if 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, and the second term to be positive, reflecting the impacts of additional heat on fatalities without adjustment.

It is straightforward and incredibly useful to rearrange Equation 6 to yield an expression for the impossible-to-observe mortality benefits of adaptation. Specifically, we write:

$$\sum_k \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k} \frac{\partial b_k^*}{\partial \boldsymbol{C}} = \frac{d\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial \boldsymbol{C}} \tag{7}$$

expressing the unobservable term as the difference between the total and partial derivatives of the expected probability of death with respect to climate. This expression is powerful, because we have now written the mortality benefits of adaptation as a function of two terms that can in principle be estimated.

[14] Substituting $x^* = Y - h(\boldsymbol{b}^*)$ in Equation 5.

[15] 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.

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

We use this insight in combination with Equation 5 to develop an expression for the additional adaptation costs incurred as the climate changes gradually from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$. As the climate changes by an incremental $d\boldsymbol{C}$ (e.g. the warming that occurs during a single year) agents will respond by adjusting K dimensions of $\boldsymbol{b}^*$ incrementally until the sum of all K marginal costs continuously equal the sum of the K expected marginal benefits (Equation 5). We wish to integrate the sum of these marginal costs to compute total costs, but they cannot be directly observed. However, combining Equations 5 and 7, it is apparent that we can infer them by differencing the total and partial derivatives of mortality risk with respect to the climate, both of which are in principle observable. Using double substitution from both Equations 5 and 7 for the expression $\frac{\partial A(\boldsymbol{b}^*)}{\partial \boldsymbol{b}}$, the change in total adaptation costs along a continuous climate change trajectory from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ is:[17]

$$A(\boldsymbol{b}^*(\boldsymbol{C}_2, Y)) - A(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) = \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} \frac{\partial A(\boldsymbol{b}^*)}{\partial \boldsymbol{b}} \frac{d\boldsymbol{b}^*}{d\boldsymbol{C}} d\boldsymbol{C} = -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial \boldsymbol{C}} \right] d\boldsymbol{C}, \tag{8}$$

which is a line integral along the equilibrium locus $\boldsymbol{b}^*(\boldsymbol{C})$ from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$. The practical value of Equation 8 is that it outlines how we can use estimates of the total and partial derivatives to infer net adaptation costs, even though adaptation itself is not directly observable. Graphical intuition for this result is provided in Appendix A.1.

In the following sections, we develop an empirical panel model exploiting both short-run and long-run variation in which the total derivative $\frac{d\tilde{f}}{d\boldsymbol{C}}$ can be separated from the partial derivative $\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$. Using idiosyncratic year-to-year time-series variation within each location, we recover causal estimates of the partial derivative of mortality rates with respect to the entire daily temperature distribution, since these changes occur unexpectedly and thus individuals do not have the ability to adapt to them. Additionally, we can plausibly identify the total derivative by interacting long-run temperatures between locations with this variation, which captures the overall effect of a shift in the temperature distribution after populations have fully adapted.

The key innovation of our approach is that we use the difference between these two derivatives (Equation 8) to construct explicit estimates of all expenditures dedicated to unobservable adaptive adjustments. By observing how optimizing populations forgo benefits of marginal adaptation in their current climate, we can recover the structure of marginal costs they must face. Then, because we can observe the behavior of many different populations that all inhabit different climates at present, each of which is presumably adapted to their respective climates, we can compute the marginal benefits of marginal adaptations at each climate. This, in turn, allows us to infer marginal adaptation costs across these different climate regimes which, when integrated along the optimal adaptation pathway $\boldsymbol{b}^*(\boldsymbol{C})$ as the climate evolves from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$, recovers the structure of $A(\boldsymbol{b})$.

This revealed preference approach, outlined by Equation 8, relies on the equivalence of *marginal* adaptation benefits and *marginal* adaptation costs at all points along the equilibrium pathway $\boldsymbol{b}^*(\boldsymbol{C})$ (Equation 5). However, it does not imply that our estimates of *total* adaptation costs are equivalent to *total* adaptation benefits. Total adaptation benefits include all benefits of inframarginal adaptive investments, while we only measure adaptation costs on the margin. For example, the marginal

[17] See Appendix A for derivation. Note that Y is treated as exogenously determined, although it can be varied period-to-period in actual calculations.

mortality benefit of the most recently purchased unit of air conditioning must exactly equal its marginal cost, according to the first order condition (i.e. Equation 5), but the benefits of all inframarginal units of air conditioning can exceed their marginal cost (note that these inframarginal benefits are not used to compute changes in $A(\boldsymbol{b}^*)$). Thus, while adaptation surplus, the difference between total adaptation benefits and total adaptation costs under climate change, is zero for the marginal adaptive investment undertaken in each period, it may be nonzero for inframarginal investments (which may have been marginal in some prior period when the climate was different). The surplus generated from all adaptive investments taken by a population following the trajectory $\boldsymbol{C}_1 \rightarrow \boldsymbol{C}_2$ is therefore $\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} \left[\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C})}{\partial \boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}), \boldsymbol{C})}{\partial \boldsymbol{C}} \right] d\boldsymbol{C}$ (see Appendix A.3).[18]

The details of implementing adaptation cost estimates and this measure of surplus are discussed below (Section 6.3 and Appendix A.3), and we empirically quantify these values globally in Section 7. We now describe our data and the procedure through which we estimate $\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})$ for all populations on the planet.

3 Data

We believe that we have collected the most comprehensive data file ever compiled on mortality, historical climate data, and climate, population, and income projections. Section 3.1 describes the data necessary to estimate the relationship between mortality and temperature. 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 all of these datasets.

3.1 Data to estimate the mortality-temperature relationship

Mortality data. Our mortality data are collected independently from 41 countries.[19] Combined, this dataset covers mortality outcomes for 55% 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). Spatial coverage, resolution, and temporal coverage are shown in Figure 1A, and each dataset is summarized in Table 1 and detailed in Appendix B.1. We harmonize all records into a single multi-country 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.

Historical climate data. We perform analyses with two separate groups of historical data on precipitation and temperature. First, we use a reanalysis product, the Global Meteorological Forcing

[18] Note that the terms $\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_2)$ and its derivatives are not directly observable in the data, because $\boldsymbol{b}^*(\boldsymbol{C}_1)$ is an "off-equilibrium" level of adaptation for populations experiencing climate $\boldsymbol{C}_2$. However, we are able to compute this value by convolving a daily response function from a population at $\boldsymbol{C}_1$ with the daily temperature distribution experienced by a population at $\boldsymbol{C}_2$. This inequality will hold if the marginal adaptation investment in climate $\boldsymbol{C}_2$ has larger mortality benefits than do previous inframarginal investments undertaken in the prior climate $\boldsymbol{C}_1$. See Appendix A.3 for details.

[19] Our main analysis uses age-specific mortality rates from 40 of these countries. We 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 omission of India from our main regressions lowers our data coverage to 38% of the global population.

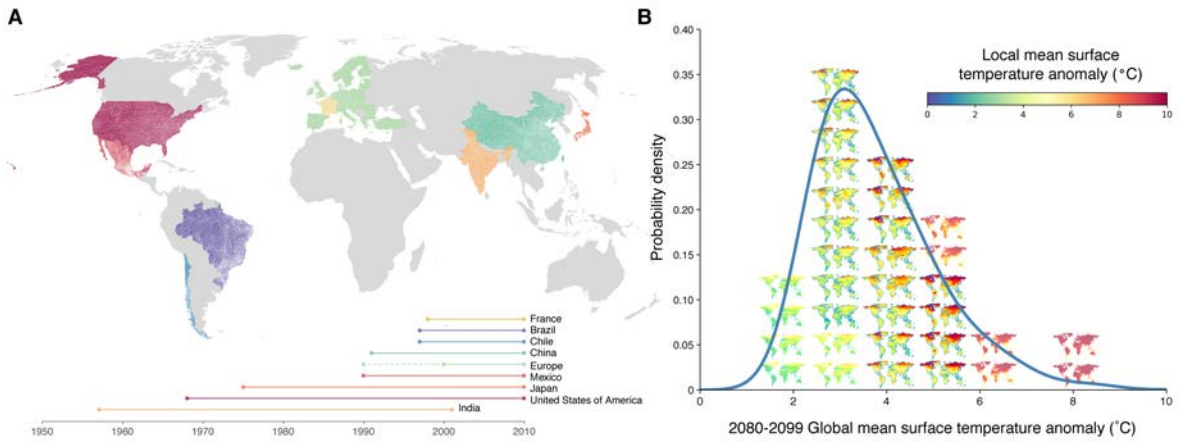

Figure 1: Mortality statistics and future climate projections used in generating empirically-based climate change mortality impact projections. Panel A shows the spatial distribution and resolution of mortality statistics from all countries used to generate regression estimates of the temperature-mortality 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). Panel B shows the 21 climate models (outlined maps) and 12 model surrogates (dimmed maps) 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.

Dataset (GMFD) (Sheffield, Goteti, and Wood, 2006), which relies on a climate 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 using a procedure detailed in Appendix B.2.4 that preserves nonlinearity in the mortality-temperature relationship.

Covariate data. Our 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). The construction of the income variable requires some estimation to downscale to ADM1 level; details on this procedure are provided in Appendix B.3.

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

Defining *impact regions* for projections. We partition the global land surface into a set of 24,378 regions onto which we generate location-specific projected damages of climate change. These regions (hereafter, *impact regions*) are constructed such that they are either identical to or represent a union

Mortality records

Country	N	Spatial scale[×]	Years	Age categories	Average annual mortality rate[*†] All-age	Average annual mortality rate[*†] >64 yr.	Global pop. share[◇]	Average covariate values[*□] GDP per capita[⊗]	Avg. daily temp.[⊘]	Annual avg. days $> 28°C$
Brazil	228762	ADM2	1997-2010	<5, 5-64, >64	525	4096	.028	11192	23.8	35.2
Chile	14238	ADM2	1997-2010	<5, 5-64, >64	554	4178	.002	14578	14.3	0
China	7488	ADM2	1991-2010	<5, 5-64, >64	635	7507	.193	4875	15.1	25.2
EU	13013	NUTS2[‡]	1990[▷]-2010	<5, 5-64, >64	1014	5243	.063	22941	11.2	1.6
France[⊕]	3744	ADM2	1998-2010	0-19, 20-64, >64	961	3576	.009	31432	11.9	.3
India[^]	12505	ADM2	1957-2001	All-age	724	–	.178	1355	25.8	131.4
Japan	5076	ADM1	1975-2010	<5, 5-64, >64	788	4135	.018	23241	14.3	8.3
Mexico	146835	ADM2	1990-2010	<5, 5-64,>64	561	4241	.017	16518	19.1	24.6
USA	401542	ADM2	1968-2010	<5, 5-64, >64	1011	5251	.045	30718	13	9.5
All Countries	**833203**	–	–	–	**780**	**4736**	**.554**	**20590**	**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.	Terrestrial Hydrology Group, Princeton
BEST	Rohde et al. (2013)	Interpolation	1°	temp.	Berkeley Earth
UDEL	Matsuura and Willmott (2007)	Interpolation	0.5°	precip.	University of Delaware

Table 1: Historical mortality & climate data

[*] 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.

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.

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) that provide daily temperature and precipitation through the year 2100.[20] [21] 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

[20] See Hsiang and Kopp (2018) for a description of climate model structure and output, as well as the RCP emissions scenarios.
[21] 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.

and Knutti, 2007; Rasmussen, Meinshausen, and Kopp, 2016).[22] 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 1B 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 mixed model ensemble (SMME). Gridded output from these projections are aggregated to impact regions; full details on the climate projection data are in Appendix B.2.

Population projections. Projections of national populations are derived from the Shared Socioeconomic Pathways (SSPs), which describe a set of plausible scenarios of socioeconomic development over the 21^{st} century in the absence of climate impacts or policy. In particular, 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). We use scenarios SSP2, SSP3, and SSP4, all of which are plausibly consistent with emissions trajectories RCP4.5 and RCP8.5 (Riahi et al., 2017). 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).

Income projections. As for population, national income per capita projections are derived from the SSP scenarios. We allocate 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.

4 Methods: Estimation of the mortality-temperature relationship

This section describes the methods we use to causally estimate temperature's impact on mortality using historical mortality records and climate data. It is divided into two parts. In the first subsection, we detail the estimating equation used to recover the average treatment effect of temperature on mortality rates across all administrative regions in our sample. In the second subsection, we describe a model of heterogeneous treatment effects that allows us to capture important differences in temperature sensitivity across distinct populations in our sample.

4.1 Estimating a pooled multi-country mortality-temperature response function

We begin by estimating a pooled multi-country 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

[22]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.

estimate the following equation on the pooled mortality sample from 40 countries,[23]

$$M_{ait} = g_a(\boldsymbol{T}_{it}) + q_{ac}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{9}$$

where a indicates age category with $a \in \{< 5, 5\text{-}64, > 64\}$, i denotes the second administrative level (ADM2, e.g. county),[24] 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 9 is the effect of temperature on mortality, represented by the response function $g_a(\cdot)$, which varies by age.[25] 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.[26] This process results in the annual, ADM2-level vector $\boldsymbol{T}_{it}$. We then choose $g_a(\cdot)$ to be a *linear* function of the *nonlinear* elements of $\boldsymbol{T}_{it}$. This construction allows us to estimate a linear regression model while preserving the nonlinear relationship between mortality and temperature that takes place at the grid-cell-by-day level (Hsiang, 2016). The nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ determine, through their linear combination in $g_a(\cdot)$, the functional form of the mortality-temperature response function.

In our main specification, $\boldsymbol{T}_{it}$ contains polynomials of daily average temperatures (up to fourth order), 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 when covariate interactions are introduced (see Section 4.2). 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 alternatively define $\boldsymbol{T}_{it}$ as a vector of binned daily average temperatures,[27] as a vector of restricted cubic splines of daily average temperatures,[28] and as a 2-part

[23] We omit India in our main analysis because mortality records do not record age.

[24] 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.

[25] There is substantial evidence (e.g. Deschênes and Moretti, 2009) that the mortality-temperature relationship varies across age groups. In some displays of results, we show an all-age response function $g(\cdot)$ where temperature effects are not age-specific, but all other terms remain as in Equation 9.

[26] 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:

$$\boldsymbol{T}_{it} = \left[\sum_{z \in i} w_{zi} \sum_{d \in t} h_1(T_{zd}), ..., \sum_{z \in i} w_{zi} \sum_{d \in t} h_P(T_{zd}) \right]$$

where transformations $h_p(\cdot)$ indicate each of P possibly nonlinear functions of gridded daily average temperature. 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).

[27] 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.

[28] In the restricted cubic spline specification, daily spline terms are summed across the year and knots are positioned

linear spline of daily average temperatures.[29] While the binned model is the most flexible functional form, this method is demanding of the data, a constraint that binds particularly in models that allow for heterogeneity in temperature sensitivity (see Section 4.2), and it results in discontinuous response surfaces. Results for these alternative functional form specifications are very similar to the fourth-order polynomial and are provided in Appendices D.1 and F.

The core appeal of Equation 9 is that the mortality-temperature response function 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 fit the multi-country pooled model in Equation 9 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.[30] Standard errors are clustered at the first administrative level (ADM1, e.g. state), instead of at the unit of treatment (ADM2, e.g. county), to reflect spatial correlation in temperatures. Robustness of this model to alternative fixed effects and error structures is shown in Section 5, and to alternative climate datasets in Appendix D.1.

4.2 Heterogeneity in the mortality-temperature response function based on climate and income

The average treatment effect identified through Equation 9 and its variants is likely to mask important differences in the sensitivity of mortality rates to changes in temperature across the diverse populations included in our sample. These differences in sensitivity reflect differential investments in adaptation – i.e. different levels of $\boldsymbol{b}^*$. To capture such heterogeneity, we develop a simple two-factor interaction model using average temperature (i.e. long-run climate) and average incomes to explain cross-sectional variation in the estimated mortality-temperature relationship. This approach provides estimates of adaptation and of income effects, as they are observed in the historical record.

Importantly, the two factors defining this interaction model come directly from the theoretical framework in Section 2. First, a higher average temperature incentivizes investment in adaptive behaviors, as the return to any given adaptive mechanism is higher the more frequently the population experiences days with life-threatening temperatures. In Section 2, this was represented by $\boldsymbol{b}^*$ being a function of climate $\boldsymbol{C}$, because the expected marginal mortality benefit of additional adaptation is dependent on the expected number of dangerously hot and cold days. In our empirical specification, we use a parsimonious parameterization of the climate, interacting our nonlinear temperature response

at the locations $\{-12, -7, 0, 10, 18, 23, 28, 33\}$ in °C.

[29] In the linear spline specification, heating degree days below 0°C and cooling degree days above 25°C are summed across the year.

[30] 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.

function with the location-specific long-run average temperature.[31] Second, higher incomes relax agents' budget constraints and hence facilitate adaptive behavior. In Section 2, this was captured by optimal adaptation $\boldsymbol{b}^*$ being an implicit function of income Y.

In addition to these theoretical arguments, there is a practical reason to restrict ourselves to these two covariates when estimating this interaction model. In order to predict responses around the world and inform projections of damages in the future, we require that all key covariates in the specification be available globally today, at high spatial resolution, and available for the standardized SSP-RCP scenarios that we utilize to generate future projections. Unlike other covariates that may be of interest, average incomes and climate can be extracted from the SSPs and the SMME climate simulations, respectively.

We capture heterogeneous patterns of temperature sensitivity via the interaction model:

$$M_{ait} = g_a(\boldsymbol{T}_{it} \mid TMEAN_s, \log(GDPpc)_s) + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{10}$$

where s refers to ADM1-level (e.g., state or province), $TMEAN$ is the sample-period average annual temperature, $GDPpc$ is the sample-period average of annual GDP per capita, and all other variables are defined as in Equation 9. We implement a form of $g_a(\cdot)$ that exploits linear interactions between each ADM1-level covariate and all nonlinear elements of the temperature vector $\boldsymbol{T}_{it}$. The model does not include uninteracted terms for $TMEAN$ and $GDPpc$ because they are collinear with α_{ai}. In contrast to the uninteracted models in Equation 9, we estimate Equation 10 without any regression weights since we are explicitly modeling heterogeneity in treatment effects rather than integrating over it (Solon, Haider, and Wooldridge, 2015). Note that while the interactions between each element of $\boldsymbol{T}_{it}$, long-run climate, and GDP per capita are linear, this specification allows for the same flexibility in the functional form of temperature as in Equation 9, it is just conditional on income and climate.[32]

The temperature interactions in Equation 10, which capture spatial differences in temperature sensitivity, rely on cross-sectional variation in climate and income across the administrative units in our mortality dataset, in combination with plausibly random year-to-year fluctuations in temperature. Thus, while the as-good-as random assignment of temperature allows us to interpret coefficients in Equation 9 as causal, the unavailability of random variation in long-run climate and income mean that stronger assumptions are required to interpret the coefficients associated with the temperature

[31] In Appendix D.5, we show robustness of this parsimonious characterization of the long-run climate to a more nuanced specification. There, we interact our nonlinear temperature variables with long-run average degree days below 20°C, and with long-run average degree days above 20°C to capture the fact that locations with similar annual average temperature may have very different distributions of degree days, due to heterogeneous seasonality. As is detailed in the Appendix, our results are robust to this alternative specification.

[32] To see how we implement Equation 10 in practice, note that in Equation 9, 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 10, 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 $\boldsymbol{T}_{it}$.

interactions in Equation 10 causally.[33]

5 Results: The mortality-temperature relationship

This section presents results describing temperature's impact on mortality, as observed in the historical record. Results are shown for a global average treatment effect first, followed by results from a model of heterogeneous treatment effects, in which we uncover substantial differences in temperature sensitivity across diverse populations around the world.

5.1 Pooled multi-country mortality-temperature response function

Pooling subnational mortality records across 40 countries, we first estimate Equation 9, showing results for the mortality-temperature response function for a randomly selected individual, obtained with a fourth-order polynomial in daily average temperature. Here, we estimate an all-age mortality response $g(\boldsymbol{T}_{it})$ in which an average treatment effect across all age categories is recovered.

Table 2 displays our main result, showing the marginal effects 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 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 column (3) is demanding of the data and we do not find statistical evidence for differential time trends in residual mortality rates across ADM1s within countries, once country-year-age fixed effects are accounted for.

The U-shaped response common in the prior literature is evident across all specifications. Examining column (2), we find that a day at 35°C (95°F) leads to an increase in the all-age mortality rate of approximately 0.4 extra deaths per 100,000, relative to a day at 20°C. A day at -5°C (23°F) similarly increases the all-age mortality rate, relative to a moderate day, by 0.3 deaths per 100,000.

Robustness to temperature functional form and climate data. Figure D1 demonstrates the robustness of our main specification to functional form assumptions (i.e. different nonlinear functions in $\boldsymbol{T}_{it}$) and to climate datasets. Across columns, distinct functional forms are shown. The binned functional form is an important benchmark, because it is the closest to being fully non-parametric; the similarity of the binned regression response functions with those from three other functional forms is reassuring. Across rows, different climate datasets are shown. The GMFD data (top) and BEST data (bottom) are drawn from independent sources, as described in Section 3; we find broad similarity in the response functions across all functional forms and both sets of data.

Age group heterogeneity. It is likely that age cohorts respond differently to temperature (Deschênes and Moretti, 2009). This heterogeneity is important to account for, as there exist large differ-

[33] We focus here on income and climate as the most important cross-sectional factors to consider based on the theoretical arguments and projection data limitations discussed above, as well as findings from across the empirical climate-economy literature demonstrating the importance of one or the other of these two factors for mediating climate impacts in cross-section (e.g. 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).

Table 2: Temperature-mortality response function estimated using pooled subnational data across 40 countries. This table shows coefficient estimates (standard errors) for a temperature-mortality response function estimated using pooled subnational data across 40 countries and 38% of the global population. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD temperature 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).

	All-age mortality rate (per 100,000)				
	(1)	(2)	(3)	(4)	(5)
35°C	0.356***	0.395**	0.183	0.610**	0.426***
	(0.130)	(0.169)	(0.118)	(0.260)	(0.163)
30°C	0.277***	0.280***	0.120*	0.310***	0.296***
	(0.070)	(0.081)	(0.066)	(0.100)	(0.077)
25°C	0.137***	0.131***	0.053	0.111***	0.138***
	(0.037)	(0.035)	(0.039)	(0.030)	(0.034)
20°C	–	–	–	–	–
	–	–	–	–	–
0° C	0.112	0.105	0.108	0.139**	0.065
	(0.122)	(0.125)	(0.098)	(0.066)	(0.120)
-5°C	0.293**	0.261*	0.194*	0.237**	0.201
	(0.149)	(0.154)	(0.104)	(0.095)	(0.145)
Adj R-squared	0.983	0.989	0.991	0.999	0.989
N	820697	820237	820237	820237	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

ences in age distributions across countries today and we expect considerable demographic transitions in the future. We test for heterogeneity across age groups by estimating Equation 9, which provides separate mortality-temperature response functions $g_a(\boldsymbol{T}_{it})$ for each of the three age categories. The regression results are shown in Table 3, with estimates from column (2) plotted in Figure D2.

Table 3 makes apparent that there is substantial heterogeneity across age groups within this multi-country sample. Again selecting column (2) as our preferred specification, people over the age of 64 experience approximately 4.5 extra deaths per 100,000 for a day at 35°C compared to a day at 20°C, and this is a substantially larger effect than that for the younger cohorts, which exhibit little response. This age group is also more severely affected by cold days than are the younger age cohorts; the estimates suggest that people over the age of 64 experience 3.4 extra deaths per 100,000 for a day at −5°C compared to a day at 20°C, while there is relatively little evidence for a mortality response to these cold days for other age categories. Overall, these results reveal that the elderly are disproportionately harmed by additional hot days and disproportionately benefit from reductions in cold days. These findings from 40 countries are consistent with prior evidence on age heterogeneity in the mortality-temperature relationship in the U.S. (Deschênes and Moretti, 2009). It is important

Table 3: 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	2.135***	-0.318	0.010	-0.056	-0.338
	(0.479)	(0.229)	(0.155)	(0.199)	(0.231)
30° C	1.268***	-0.187*	0.007	-0.029	-0.186*
	(0.218)	(0.097)	(0.066)	(0.087)	(0.099)
20°C	–	–	–	–	–
	–	–	–	–	–
0°C	-2.142***	-0.096	-0.089	-0.065	-0.129
	(0.318)	(0.114)	(0.106)	(0.045)	(0.115)
-5°C	-2.258***	-0.196*	-0.116	-0.030	-0.221*
	(0.389)	(0.117)	(0.105)	(0.076)	(0.121)
Panel B: 5 - 64 years of age					
35°C	4.533***	0.008	0.008	0.089	0.030
	(0.654)	(0.109)	(0.067)	(0.168)	(0.109)
30°C	2.538***	0.050	0.031	0.038	0.061
	(0.256)	(0.064)	(0.036)	(0.076)	(0.063)
20°C	–	–	–	–	–
	–	–	–	–	–
0°C	-4.140***	-0.125**	-0.093*	-0.009	-0.127**
	(0.299)	(0.063)	(0.049)	(0.040)	(0.059)
-5°C	-4.730***	-0.119	-0.090*	-0.004	-0.119*
	(0.377)	(0.079)	(0.050)	(0.056)	(0.072)
Panel C: >64 years of age					
35° C	-3.927**	4.542**	1.947	4.708**	4.714**
	(1.769)	(1.936)	(1.308)	(1.866)	(1.884)
30° C	-1.922**	2.598***	0.954	2.380***	2.679***
	(0.768)	(0.832)	(0.585)	(0.698)	(0.805)
20°C	–	–	–	–	–
	–	–	–	–	–
0°C	8.140***	2.009***	1.764***	1.230***	1.694**
	(0.744)	(0.725)	(0.507)	(0.372)	(0.710)
-5°C	10.305***	3.399***	2.506***	2.002***	2.891***
	(0.892)	(0.951)	(0.575)	(0.522)	(0.904)
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

to note, however, that the oldest age group (over 64 years) accounts for just 12% of the population in our historical sample, limiting its similarity to the average treatment effect in Table 2.

Alternative specifications. In both Tables 2 and 3, columns (4) and (5) provide results from alternative specifications. In column (4), we address the fact that some of our data sources are drawn from countries which may have less capacity for data collection than others in the sample. To account for the possibility that some data sources are systematically less precise, we re-estimate our model using Feasible Generalized Least Squares (FGLS) under the assumption of constant variance within each ADM1 unit.[34] In column (5), we address the finding that temperatures can exhibit a lagged effect on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). While lagged effects within and across months in the same calendar year are accounted for in the net annual mortality totals used as a dependent variable in all our specifications, in column (5) we define a 13-month exposure window to additionally account for January mortality that may result from up to a four-week lag of temperature exposure in the previous December.[35] Tables 2 and 3 show that the results for both of these alternative specifications are comparable in sign and magnitude to those from column (2).

5.2 Subnational heterogeneity in the mortality-temperature response function

The estimation of Equation 10 provides an opportunity to test for systematic heterogeneity in the mortality-temperature response function. Specifically, this model interacts the nonlinear temperature response with ADM1-level covariates of average climate and average income. Tabular results are reported in Table D2 for each of the three age groups of interest. As these terms are difficult to interpret, we visualize this heterogeneity 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 D2. 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 2 for the >64 age category (other age groups are shown in Appendix D.4). Each subpanel in Figure 2 shows a temperature response function for the >64 mortality rate, evaluated at a the average level of income and climate in each bin of the $\log(GDPpc) \times TMEAN$ joint distribution. Average incomes are increasing across bins vertically and average temperatures are increasing across bins horizontally.

The results in Figure 2 are consistent with the predictions from our theoretical framework in Section 2. Recall that we expect increased frequency of exposure to higher temperatures to incentivize

[34] To do this, we estimate the model in Equation 9 using population weights and our preferred specification (column (2)). Using the residuals from this regression, we calculate an ADM1-level weight, 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-age category 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.

[35] 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 $t-1$.

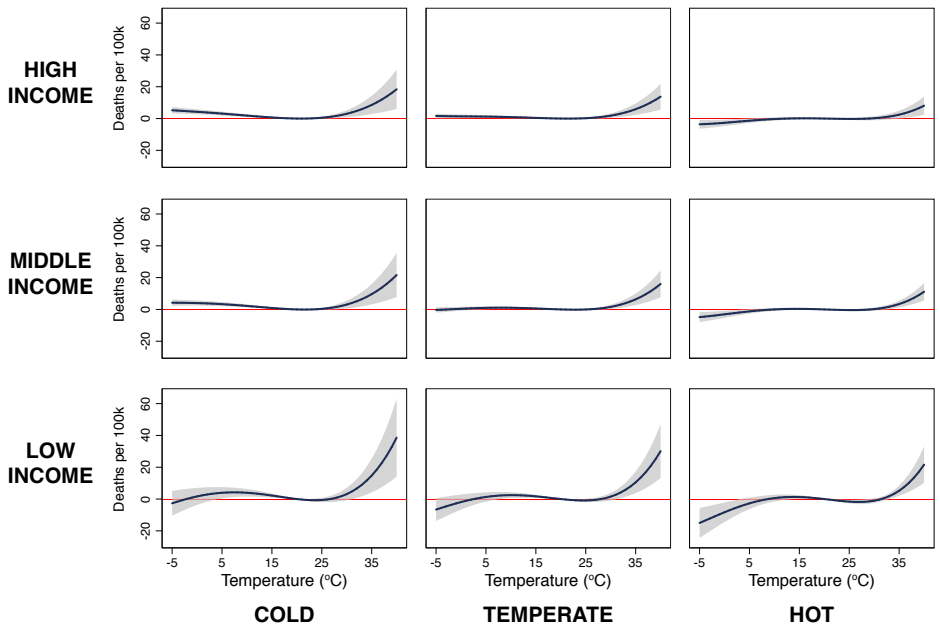

Figure 2: Heterogeneity in the mortality-temperature relationship (age >64 mortality rate). Each panel represents a predicted mortality-temperature response function for the >64 age group for a subset of the income-average temperature covariate space within our data sample. Response functions in the 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.

investment in adaptive behaviors or technologies, as the marginal mortality benefit of adaptation is higher in hotter locations. This would lead to lower temperature sensitivities to heat in places which are warmer. Indeed, in the >64 age category, moving from the coldest to the hottest tercile saves on average 8.6 (p-value=0.04) deaths per 100,000 at 35°C. Similarly, a loosening of the budget constraint, as proxied by increasing GDP per capita, should enable individuals to invest further in adaptation. Here too, the association is apparent as the impact of hot days on mortality declines with rising income; on average, moving from the poorest to the richest tercile in the sample saves approximately 6.0 (p-value=0.05) deaths per 100,000 at 35°C for the $>$ 64 age category.

Based on these results, it is apparent that a primary margin of adaptation to temperature relevant to climate change occurs at the extremes of the temperature distribution. Income is protective on the hottest days for the middle and older age category, and is additionally beneficial on the coldest days for the youngest age category. As evidenced in Figure 2 and Appendix D.4, the richest regions exhibit substantial flattening of the response curve. This is particularly the case on hot days, and is consistent with previous evidence from the U.S. (Barreca et al., 2016; Heutel, Miller, and Molitor, 2017).

6 Methods: Spatial and temporal projections of the mortality-temperature relationship out of sample

This section describes how we combine the empirical results from Section 5 with the theoretical framework from Section 2 to generate global projections of mortality risk under climate change. The first two subsections outline how we use the estimation of Equation 10 to spatially and temporally extrapolate mortality-temperature response functions. The third subsection details how we use those same response surface parameters to estimate adaptation costs as derived in Section 2. The final subsection describes our method for capturing statistical and climate uncertainty when using this response surface and adaptation cost estimates to construct future projections of the full mortality risk of climate change.

6.1 Spatial projection: Constructing a globally representative response

A key challenge for any complete analysis of the global mortality-temperature relationship is unavailability of data throughout much of the world. 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 analysis of the previous section, which is especially troubling because these populations have incomes and live in climates that may differ from the parts of the world where data are available.

Our solution to this data sparsity is to use the results from the estimation of Equation 10 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.

To see how this is done, 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 10. Second, because data on GDP per capita are not available at the impact region level, we use the procedure described in Section 3.2 to predict high-resolution impact region-level income using national income statistics in combination with night lights imagery from the NOAA's Defense Meteorological Satellite Program (DMSP). Third, we calculate average annual temperature (i.e. a measure of the long-run climate) for each impact region using the same temperature data that were assembled for the regressions in Equations 9 and 10.

Using these impact region-level data, we predict the shape of the response function for each age group a, impact region r, and year t, up to a constant: $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt})$. Note that the various fixed effects in Equation 10 are unknown and omitted, since they were nuisance parameters in the original regression. When presenting predicted impacts in the present, we evaluate the values of the covariates at the baseline period $t = 2015$. This results in a spatially heterogeneous and globally comprehensive set of predicted response functions, each of which is specific to an impact region. Once this step is complete, we have a unique set of age-specific predicted response functions for each location on Earth.

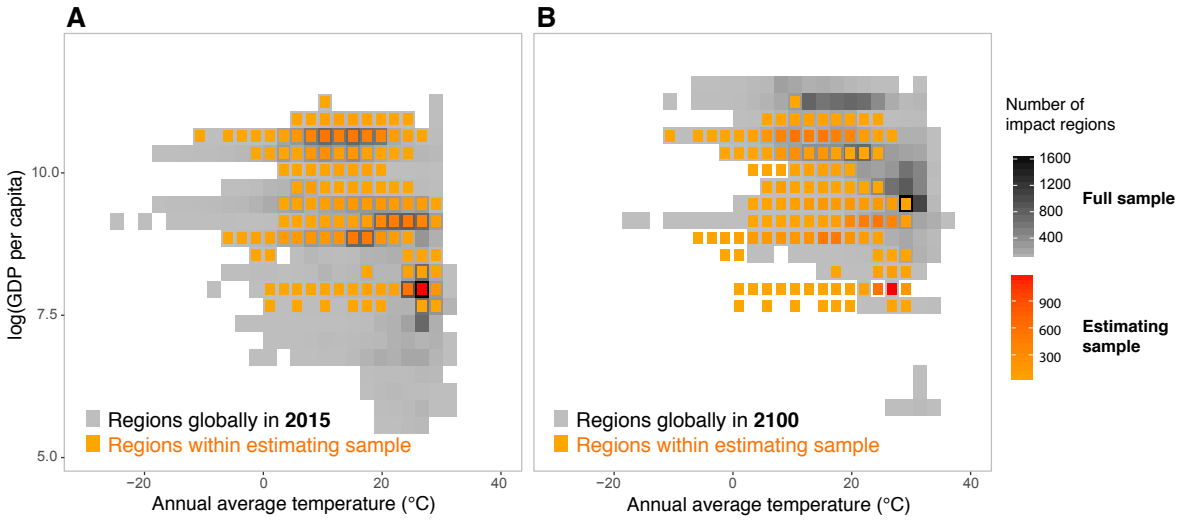

Figure 3: Joint coverage of income and long-run average temperature for estimating and full samples. Joint distribution of income and long-run average annual temperature in the estimating sample (red-orange), as compared to the global sample of impact regions (grey-black). Panel A shows in grey-black 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) 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.

The credibility of the predicted response functions in regions where mortality data are unavailable depends on the representativeness of the observed sample. Thus we evaluate the extent of common overlap between the two samples in Figure 3. Panel A 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 Equations 9 and 10. We note 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 scarcity of mortality data in poorer countries. At the end of the century, the overlap is generally better, although unsurprisingly the support of our historical data does not extend to the highest projected temperatures and incomes.

6.2 Temporal projection: Accounting for future adaptation benefits

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. To reflect this adaptation and account for benefits of income growth, we allow each impact region's mortality-temperature response function to evolve over time as incomes and climate change. Just as with the spatial extrapolation exercise described above, we model the evolution of response functions based on projected changes to average climate and GDP per capita, again using the estimation results from fitting Equation 10.

We allow the response function in region r and in year t to evolve over time as follows. First, a 13-year moving average of income per 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 $\log(GDPpc)_{rt}$.[36] The length of this time window is chosen based on a goodness-of-fit test across alternative window lengths (see Appendix E.1).[37] 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 $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt})$ are calculated for each region for each age group in each year with these updated values of $TMEAN_{rt}$ and $\log(GDPpc)_{rt}$.

The calculation of future mortality-temperature response functions is conceptually straightforward. However, as we are generating projections decades into the future, we must impose a set of reasonable constraints on this calculation in order to ensure plausible out-of-sample projections.

The following two constraints, guided by economic theory and by the physiological literature, ensure that future response functions are consistent with the fundamental characteristics of mortality-temperature responses that we observe in the historical record.[38] 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).[39] 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 10 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" and the response function would invert (Figure E1). This is guaranteed to occur mechanically if enough time elapses, because we have only allowed income and climate to interact with the response functions linearly.[40] However, such behavior, in which extreme temperatures are less damaging to mortality rates than more moderate temperatures, is inconsistent with a large body of epidemiological and econometric literature recovering U-shaped response functions for mortality-temperature relationships under a wide range of functional form assumptions and across diverse locations globally (Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011), as well as what we observe in our data. To avoid this unrealistic out-of-sample behavior, we impose weak monotonicity of mortality with respect to temperature; this assumption binds for the >64 age category at 35°C in 9% and 18% of impact regions in 2050 and 2100, respectively.[41] 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

[36] GDP per capita is provided in the SSP projections only at 5-year increments. The values are linearly interpolated between SSP 5-yearly updates of values. A uniform income growth rate is applied to all impact regions within each country.

[37] We can derive a duration over which updating occurs in the case of income due to substantial time series variation of income over the course of our observed sample. For temperature, the historical trends have so far been small, making derivation of a comparable duration infeasible.

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

[39] We identify for each impact region a "minimum mortality temperature" (MMT); this is the temperature at which mortality is minimized for each predicted impact region and age-specific response function in the baseline year 2015. We constrain this response function minimization problem to select a temperature within the range of 10°C to 30°C, following extensive evidence from epidemiology and ergonomics indicating that temperatures outside this range are unlikely to be physiologically optimal (see Appendix E.2 for details). 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).

[40] In Appendix D.5 we show that a nonlinear interaction between annual temperatures and long-run climate does not change the interaction effects recovered in Equation 10.

[41] 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.

fixed effects in Equation 10 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.[42]

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. We place no restrictions on the cross-sectional effect of income on the temperature sensitivity when estimating Equation 10, but we constrain the marginal effect of income on temperature sensitivity to be weakly negative in future projections. This assumption never binds for temperature sensitivity to hot days (>35°C).[43]

Under these two constraints, we estimate projected impacts separately for each impact region and age group for each year from 2015 to 2100 by applying projected changes in the climate to these spatially and temporally heterogeneous response functions. We compute the nonlinear transformations of daily average temperature that are used in the function $g_a(\boldsymbol{T}_{rt})$ 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. In Section 6.4 we discuss a simulation method we use to additionally capture uncertainty arising from econometric estimation of Equation 10, and in Section 8 we detail a method for extrapolation of damage estimates through 2300, which are used to construct values for the excess mortality partial Social Cost of Carbon (SCC).

6.3 Implementing empirical estimation of adaptation costs

The full cost of the mortality risk due to climate change is the sum of the observable change in mortality and adaptation costs (Equation 4). The latter cannot be observed directly; however, as derived in Section 2, we can recover an expression for adaptation costs that is empirically tractable. Theoretically, these 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, and integrating this difference (see Equation 8). Here, we describe a practical implementation for this calculation; details are provided in Appendix A.4.

The first key building block of this calculation is an empirically-derived expression for expected mortality risk. To construct this function, we take expectations of Equation 10, which allows us to specify empirically-estimated expected mortality risk for each age group a in region r for year t; we denote this function $\mathrm{E}[\hat{g}_{art}(\cdot)]$. Recall that $\hat{g}_{art}(\cdot)$ is a function of the summary climate parameter $TMEAN$ and log income per capita, as well as the estimated parameters associated with each, which we recover from estimation of Equation 10 (see Section 6.2). We assume that individuals use the recent climate past to make projections about current temperature realizations, so the expectation of $\hat{g}_{art}(\cdot)$ is computed over realizations of temperature $\boldsymbol{T}_{rt}$ for region r in year t from the prior 15 years,

[42]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.

[43]The assumption that rising income cannot increase the temperature sensitivity of mortality does not bind for hot days because our estimated marginal effects of income are negative for high temperatures (see Table D2). However, it does bind for the >64 age category under realized temperatures in 30% and 24% of impact regions days in 2050 and 2100, respectively.

with weights of historical observations linearly declining in time, analogously to the construction of $TMEAN$ (see Section 6.2 and Appendix E.1). We omit subscripts for clarity in remainder of this section, but the calculation presented below is conducted yearly for each age and region separately.

The construction of $\mathrm{E}[\hat{g}(\cdot)]$ allows us to calculate the difference between the total and partial derivatives of expected mortality risk with respect to the climate. First, the partial derivative is captured empirically through a change in events $\boldsymbol{T}$, where climate and income are "held fixed" in an initial period t_1, since the partial effect must exclude any compensatory behavior by populations. 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 are captured in our econometric framework through the $TMEAN$ interactions in $g(\cdot)$. Thus, the empirical estimate of the total derivative includes both the partial effect of changes to $\boldsymbol{T}$ and the effect of adaptive adjustments captured by the effect of $TMEAN$. Because income effects are not caused by changes in the climate, income Y_1 is held fixed.[44] Therefore, the *difference* between the total and partial derivatives of expected mortality risk with respect to the climate simplifies to the effect of $TMEAN$ on $\mathrm{E}[\hat{g}]$:[45]

$$\begin{aligned}\frac{d\hat{\tilde{f}}_t}{d\boldsymbol{C}} - \frac{\partial\hat{\tilde{f}}_t}{\partial\boldsymbol{C}} &= \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}}\bigg|_{\boldsymbol{C}_1,Y_1}\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} + \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C},Y_1}\frac{\partial TMEAN}{\partial \boldsymbol{C}} - \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}}\bigg|_{\boldsymbol{C}_1,Y_1}\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} \\ &= \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C},Y_1}\frac{\partial TMEAN}{\partial \boldsymbol{C}} \end{aligned} \tag{11}$$

where $\hat{\tilde{f}}$ denotes the empirical estimate (indicated by the "hat" notation) of expected mortality risk $\tilde{f}(\cdot)$, as in Section 2, and where $\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}}$ is the change in the all nonlinear elements of $\boldsymbol{T}$ that describe the daily temperature distribution, resulting from an incremental change in climate.

Recall that we aim to provide empirical estimates of the total costs of adaptation that are incurred as the climate evolves over some trajectory from an initial climate $\boldsymbol{C}_1$ in period t_1 to a final climate $\boldsymbol{C}_2$ in period t_2. Following Equation 8 in Section 2, the total adaptation costs due to such a change in climate can be estimated as the integral of the difference between the total and partial derivatives shown in Equation 11. In all high-resolution climate model projections that we use, climate changes in a region can be indexed by time. Thus, in each projection, we solve for adaptation costs as a region's climate evolves from $\boldsymbol{C}(t_1)$ to $\boldsymbol{C}(t_2)$, approximating the integral using Newton's method over discrete time-steps of one year:[46]

$$\Big[\widehat{A(\boldsymbol{b}^*(\boldsymbol{C}(t_2),Y))} - \widehat{A(\boldsymbol{b}^*(\boldsymbol{C}(t_1),Y))}\Big] \approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \left(\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C}_\tau,Y}\right)(TMEAN_\tau - TMEAN_{\tau-1}) \tag{12}$$

We underscore that while the partial derivative $\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}$ is evaluated at time-varying values of

[44] Note that throughout this paper, we treat income and demographic trajectories as exogenously given by the Shared Socioeconomic Pathways (SSPs) and thus these variables are not affected by warming (although they vary through time). In future work, these covariates, which determine the sensitivity of mortality rates to temperature exposure, could be endogenized into the climate change projection.

[45] See Appendix A.4 for derivation.

[46] Because the amount of climate change experienced in any single year is small (on average 0.04°C globally), our discretization is a reasonable approximation of the continuous integral in Equation 8.

the climate $\boldsymbol{C}$, income of a region is held constant in Equation 12, aside from its influence over the VSL. This is because the goal is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate, which is distinct from any expenditures due to rising incomes. These income effects were partialled out econometrically in the estimation of Equation 10.[47] Finally, in implementation 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 (see Section 8).

These adaptation cost estimates are calculated for each impact region, age group, and year, using $t_1 = 2015$ as the baseline year, for each of our 33 high-resolution climate model projections.

6.4 Accounting for uncertainty in projected mortality effects of climate change

An important feature of the analysis is to develop estimates of the mortality impacts of climate change that reflect the inherent uncertainty in these future projections. As discussed in Sections 6.2 and 6.3, we construct estimates of the mortality risk of climate change for each of 33 distinct climate projections in the SMME, capturing uncertainty in the climate system.[48] Additionally, there exists an important second source of uncertainty in our projected impacts that is independent of physical uncertainty, arising from the econometric estimates of response functions; i.e. uncertainty in the estimate of $\hat{g}_a(\cdot)$.

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, 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 10.[49] 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

[47] Note that throughout our main analysis, the parameterization of $\hat{g}(\cdot)$ is linear in the terms $\boldsymbol{T}$; thus, we can distribute the expectation operator to write $\mathrm{E}[\hat{g}\left(\boldsymbol{T}_t \mid TMEAN_t, \log(GDPpc)_t\right)] = \hat{g}\left(\mathrm{E}[\boldsymbol{T}_t] \mid TMEAN_t, \log(GDPpc)_t\right)$. With this simplification, and the specific functional form of $\hat{g}(\cdot) = \left(\boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_1 TMEAN_t + \boldsymbol{\gamma}_2 \log(GDPpc)_t\right)\boldsymbol{T}_t$ that we use to implement Equation 10, we compute Equation 12 as follows (as above, omitting age and impact region subscripts for clarity):

$$\begin{aligned}\left[\widehat{A(\boldsymbol{b}^*(\boldsymbol{C}(t_2), Y)) - A(\boldsymbol{b}^*(\boldsymbol{C}(t_1), Y))}\right] &\approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \left(\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C}_\tau, Y}\right)(TMEAN_\tau - TMEAN_{\tau-1}) \\ &\approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \; \hat{\boldsymbol{\gamma}}_1 \mathrm{E}[\boldsymbol{T}]_\tau \left(TMEAN_\tau - TMEAN_{\tau-1}\right)\end{aligned}$$

where the expectation $\mathrm{E}[\boldsymbol{T}]_\tau$ is computed over realizations of temperature from the prior 15 years, with weights of historical observations linearly declining in time, analogously to the construction of $TMEAN$ (see Section 6.2 and Appendix E.1).

[48] 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.

[49] Note that coefficients for all age groups are estimated jointly in Equation 10, such that across-age-group covariances are accounted for in this multivariate distribution.

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 full mortality risk (i.e. inclusive of adaptation benefits and costs) for each impact region for each year between 1981 and 2100. The resulting calculation is highly computationally intensive (requiring ∼94,000 hours of CPU time), but it incorporates uncertainty from climate and econometric sources. When reporting projected impacts in any given year, we report summary statistics (e.g. mean, median) of this entire distribution.

7 Results: The global mortality consequences of future climate change

Here, we use our empirical results from Section 5.2 to extrapolate responses to the parts of the world where historical mortality data are unavailable (using the approach described in Section 6.1), allowing us to create the first global average treatment effect of temperature on mortality. We then project these responses into the future to calculate the full mortality risk of climate change, accounting for both the benefits and costs of adaptation (using the approach outlined in Sections 6.2, 6.3, and 6.4).

7.1 Spatial extrapolation of temperature sensitivity

Figure 4 demonstrates our extrapolation of mortality-temperature response functions to the entire globe. In panel A, these predicted mortality-temperature responses are plotted for each impact region for baseline (2015) values of income and climate for the oldest age category and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Despite a shared overall shape, panel A shows substantial heterogeneity across regions in this temperature response. In panel B, we show an analogous figure for the youngest age category, demonstrating our model's ability to predict the substantial within-location differences in mortality risk across age groups. Geographic heterogeneity within our sample is shown for hot days in the maps in panels C and D. On each map, the colors indicate the marginal effect of a day at 35°C day, relative to a day at a location-specific minimum mortality temperature, for ages >64 (panel C) and for ages < 5 (panel D); the grey areas are locations where mortality data are unavailable.

Figure 4E–H shows analogous figures but now extrapolated to the entire globe, including locations where we lack subnational mortality data. We can fill in the estimated mortality effect of a 35°C day for regions without mortality data by using location-specific information on income and climate during 2015. 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;[50] this value is considerably larger than the estimated effect of 7.8 deaths per 100,000 within the sample of countries for which we obtain mortality data. At the current global temperature distribution, the application of the in-sample average treatment effect (as opposed to the global average treatment effect) would lead to an underestimate of global deaths due to hot days (>35°C) by ∼400,000 annually, for the oldest age group. This difference in sensitivity

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

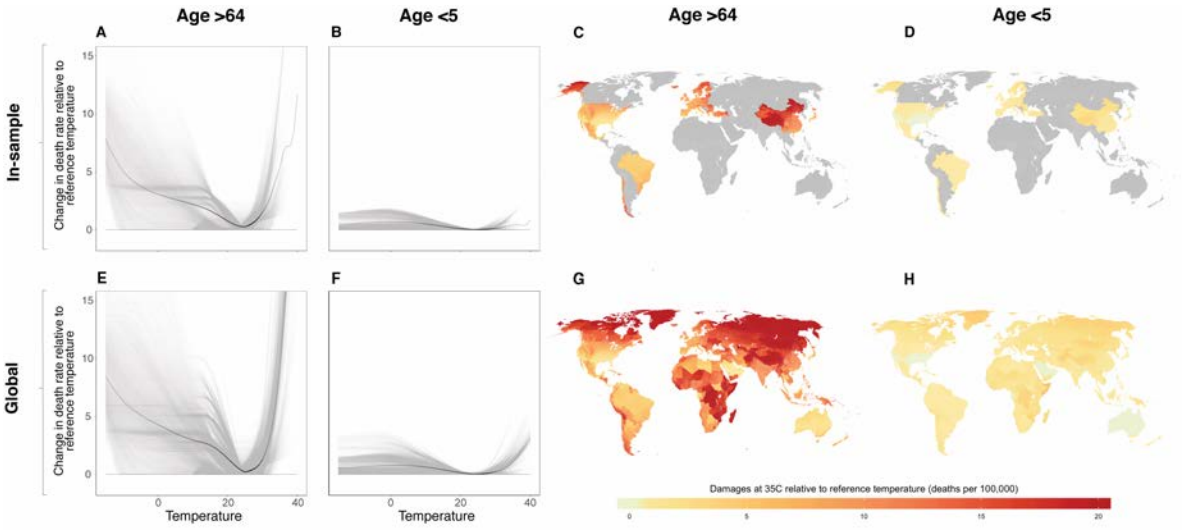

Figure 4: Using income and climate to predict current response functions globally. In panels A, B, E and F, grey lines are predicted response functions for impact regions, each representing a population of 276,000 on average. Solid black lines are the unweighted average of the grey lines, where the opacity 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 2001-2015.

is due to the fact that our sample represents wealthier populations where temperature responses are more muted. Overall, the bottom row in Figure 4 underscores that the places where data are available are not a random sample of all places on the planet, and that analyzing them alone can lead to biased estimates of climate change impacts.

7.2 Projection of future damages and adaptation

The previous subsection demonstrated that the model of heterogeneity outlined in Equation 10 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. We use 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 6.2. Uncertainty in these estimated response functions is accounted for through Monte Carlo simulation, as described in Section 6.4.

Projected adaptation benefits. Figure 5 provides an initial look into a key ingredient into projections of future damages and adaptation. In particular, we plot the spatial distribution of the *change* in the marginal damages of a 35°C day between 2050 and 2015 and between 2100 and 2015 for the >64 age category.[51] The maps reveal that in most regions of the world, there is a clear downward

[51] Specifically, these values are $\hat{g}_a(\boldsymbol{T}|TMEAN_{r,2050}, \log(GDPpc)_{r,2050}) - \hat{g}_a(\boldsymbol{T}|TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$ and $\hat{g}_a(\boldsymbol{T}|TMEAN_{r,2100}, \log(GDPpc)_{r,2100}) - \hat{g}_a(\boldsymbol{T}|TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$, all evaluated at daily temperature $T = 35$°C for age group a >64.

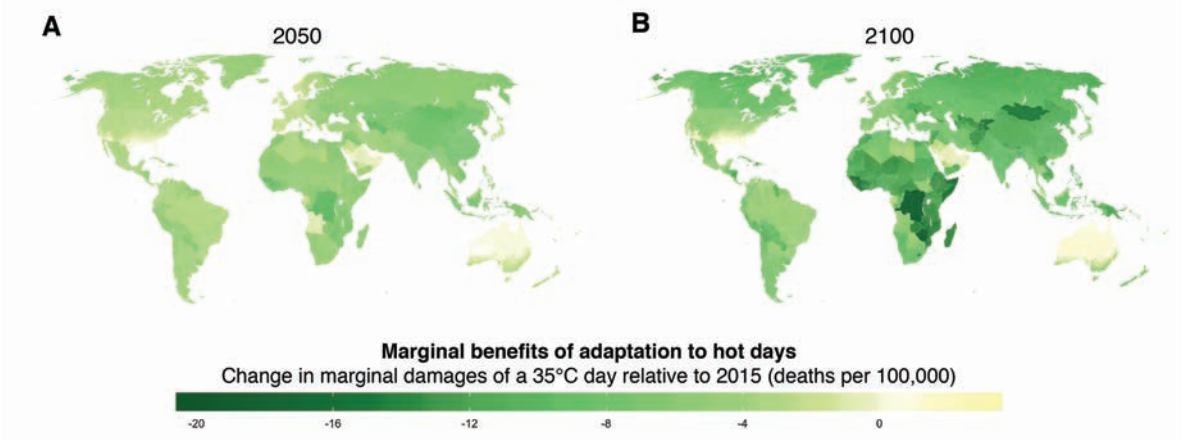

Figure 5: Spatial and temporal heterogeneity in temperature sensitivity. Panels A and B indicate the change in mortality sensitivity to hot days (35°C) for the oldest age category (>64) between 2015 and 2050 (A), and between 2015 and 2100 (B). Darker colors signify larger predicted adaptation to heat. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

trend in the sensitivity of mortality rates to high temperatures, as locations get both richer and hotter as the century unfolds. It is noteworthy that the average global increase in the mortality rate on a 35°C day (relative to a day at location-specific minimum mortality temperatures, as in Figure 4) declines by 7.7 per 100,000 between 2015 and 2100 such that it is just 2.4 per 100,000 in 2100. In 2015, similarly low sensitivities are only present in places like Houston, Texas, so our central estimates suggest that the average location in the world in 2100 exhibits current-day Houston-like levels of adaptation. Increasing incomes account for 77% of the decline in marginal damages for the >64 age category 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.

Defining four measures of expected climate change impacts. We now use our estimates of adaptation benefits, adaptation costs, and changes in climate exposure to develop measures of the expected costs of climate change induced mortality risk, reflecting Equations 2, 3, and 4 in Section 2. Here, we separate the protective role of income growth from that of adaptation to warming. As discussed, the central welfare metric is the full mortality-related costs of climate change, i.e. the sum of the increase in deaths and adaptation costs (Equation 4). The empirical estimation of each of these measures is displayed here in units of deaths per 100,000, although it is straightforward to monetize these measures using estimates of the value of a statistical life (VSL), as we do 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 does not adapt, such that their response function in each year t is a function of their 2015 level of income and average climate. In other words, mortality sensitivity to temperature is assumed not to change with future income or temperature. As discussed in Section 2, this is a benchmark model often employed in previous work. Specifically, the expected costs of 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):[52]

[52]Note that in all estimates of climate change impacts in (i)–(iv), population growth is accounted for as an exogenous

(i) *Mortality effects of climate change with neither adaptation nor income growth* (recall Equation 2):

$$\underbrace{\hat{g}(\boldsymbol{T}_t \mid TMEAN_{2015}, \log(GDPpc)_{2015})}_{\substack{\text{mortality risk with no adaptation} \\ \text{and climate change}}} - \underbrace{\hat{g}(\boldsymbol{T}_{2015} \mid TMEAN_{2015}, \log(GDPpc)_{2015})}_{\text{mortality risk without climate change}}$$

The second measure is the *mortality effects of climate change with benefits of income growth*, which allows the response function to change with future incomes, in both the warming and no-warming cases. This measure captures the change in mortality rates that would be expected from climate change if populations became richer, allowing them to spend more resources on adaptation, 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*:

$$\underbrace{\hat{g}(\boldsymbol{T}_t \mid TMEAN_{2015}, \log(GDPpc)_t)}_{\substack{\text{mortality risk with benefits of income growth} \\ \text{and climate change}}} - \underbrace{\hat{g}(\boldsymbol{T}_{2015} \mid TMEAN_{2015}, \log(GDPpc)_t)}_{\substack{\text{mortality risk with benefits of income growth,} \\ \text{without climate change}}}$$

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. 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* (recall Equation 3):

$$\underbrace{\hat{g}(\boldsymbol{T}_t \mid TMEAN_t, \log(GDPpc)_t)}_{\substack{\text{mortality risk with benefits of income growth} \\ \text{and adaptation under climate change}}} - \underbrace{\hat{g}(\boldsymbol{T}_{2015} \mid TMEAN_{2015}, \log(GDPpc)_t)}_{\substack{\text{mortality risk with benefits of income growth,} \\ \text{without climate change}}}$$

The final measure is the most complete, as it captures both the benefits and costs of adaptation. Recall that adaptation costs cannot be observed directly, but that we construct estimates using the revealed preference methodology detailed in Section 2. We call this measure the *full mortality risk of climate change*, and it captures the opportunity costs of investing in the adaptation benefits described by (iii):

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

$$\underbrace{\hat{g}(\boldsymbol{T}_t \mid TMEAN_t, \log(GDPpc)_t) - \hat{g}(\boldsymbol{T}_{2015} \mid TMEAN_{2015}, \log(GDPpc)_t)}_{\text{mortality effects of climate change with benefits of income growth and adaptation (iii)}} + \underbrace{\frac{1}{VSL}\left[A(TMEAN_t, GDPpc_t) - \widehat{A(TMEAN_{2015}, GDPpc_t)}\right]}_{\text{adaptation cost estimates}}$$

projection that does not depend on the climate.

In all of these measures, year $t = 2015$ is treated as the baseline year, meaning that climate change impacts are defined to be zero in this year. These four measures are all reported below in units of human lives. 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. Note that the use of the these units is why adaptation costs in expression (iv) are multiplied by $\frac{1}{VSL}$, as the definition of adaptation costs $A(\cdot)$ in Equation 8 is given in dollars.

Note that in all expressions of climate change impacts, the first term represents the predicted mortality rate under a future warming climate. The second term represents a counterfactual predicted mortality rate that would be realized under current temperatures. In expression (i), the counterfactual mortality rate is identical to that under current temperatures and incomes. In all other scenarios, the counterfactual mortality rate is one that would be realized under current temperatures, but in a population that benefits from rising incomes over the coming century (based on projections from the Shared Socioeconomic Pathways, or SSPs). These counterfactuals thus include 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. By subtracting off this counterfactual, our predicted mortality levels isolate the additional cost of climate change on a population experiencing economic growth.

The full mortality risk of climate change for 24,378 global regions. Figure 6 shows the spatial distribution of the full mortality risk of climate change (i.e. expression (iv)) in 2100 under the emissions scenario RCP8.5, expressed in death-equivalents per 100,000. All other measures of climate change impacts (i.e. expressions (i)–(iii)) 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 1981 and 2100 (equal to expression (iii)), separately for each age group. Following the approach outlined in Section 6 and Appendix A, we simultaneously compute associated measures of adaptation costs for each location and age at each point in time, and add them to expression (iii). The map displays the spatial distribution of our main results (expression (iv)), depicting the mean estimate across our ensemble of Monte Carlo simulations, accounting for both climate and statistical uncertainty and pooling across all age groups.[53] 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 6 makes clear that the costs of climate change induced mortality risks are distributed unevenly around the world. Despite the gains from adaptation shown in Figure 5, there are large losses in the global south. For example, in Accra, Ghana, we predict that climate change by end of century will cause damages equivalent to approximately 160 additional deaths per 100,000 annually under RCP8.5. In contrast, there are gains in many impact regions in the global north, including in Oslo, Norway, where we predict that the equivalent of approximately 230 lives per 100,000 are saved annually.

Aggregate global mortality consequences of climate change. To compute the global welfare implications of the results in Figure 6, we aggregate location-specific impacts across all regions for all

[53] 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.

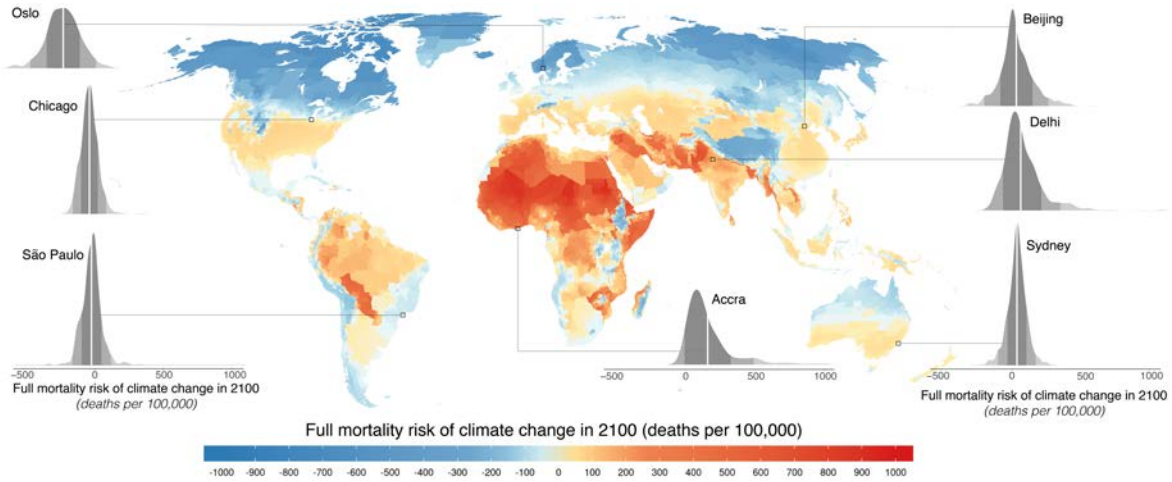

Figure 6: The mortality risk of future climate change. The map indicates the full mortality risk of climate change, measured in units of deaths per 100,000 population, in the year 2100. Estimates come from a model accounting for both the costs and the benefits of adaptation, and the map shows the climate 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.

four measures of climate change impacts (i.e. expressions (i)–(iv)). The time series in Figure 7 show predictions of global increases in the mortality rate (deaths per 100,000) due to climate change under emissions scenario RCP8.5. In panel A, each line shows an average 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 6.4. In panel B, the 10^{th}-90^{th} percentile range of the Monte Carlo simulation distribution is shown for the full mortality risk of climate change (expression (iv)) under RCP8.5, as well as the 25^{th}-75^{th} percentile range; the black line represents the same average value in both panels. While the full distribution of impacts is shown in Figure 7B for RCP8.5, the boxplots to the right summarize mortality impacts for both RCP8.5 and the moderate emissions scenario of RCP4.5 for comparison.

Figure 7A illustrates that if we failed to account for adaptation or income growth, we would estimate that the mortality cost of climate change would be 221 deaths per 100,000 by 2100, on average across simulation runs (orange line). This is an enormously large estimate; it 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 the effects of future income growth and adaptation to climate substantially reduce these projected impacts (Figure 7A). Higher incomes dramatically reduce the mortality effect, leading to an average of 102 deaths per 100,000 in 2100 (yellow line); climate adaptation reduces this further to 73 deaths per 100,000 (green line). Although substantially reduced from the *no adaptation* projection, these smaller counts of direct mortality remain economically meaningful—for comparison, the 2017 mortality rate from automobile accidents in the United States was 11.4 per 100,000.

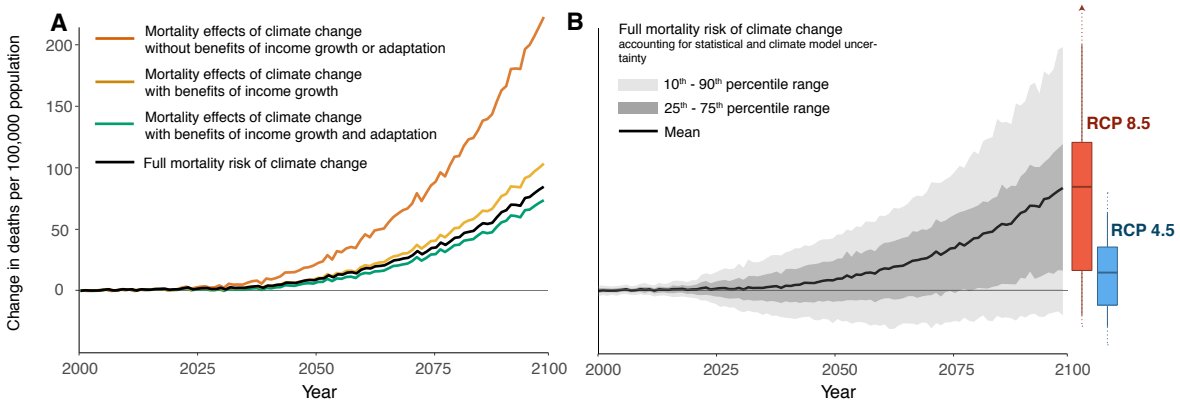

Figure 7: Time series of projected mortality risk of climate change. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In 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 7A also demonstrates that climate adaptation is projected to be costly. We compute that estimates of climate adaptation costs, calculated using the approach in Section 2, are valued at 12 death-equivalents per 100,000 in 2100. The net result is that the full mortality risk due to climate change (i.e., expression (iv)) is projected to equal to 85 deaths per 100,000 by the end of the century under RCP8.5. Had we accounted for the benefits of adaptation but failed to account for their costs, we would have underestimated the total aggregate impact of these changes, particularly in regions of the world where adaptation costs compose a substantial share of total damages.[54] Nonetheless, our estimate for the global average benefits of adaptation (29 deaths per 100,000) outweighs the costs of these adjustments (12 death-equivalents per 100,000), demonstrating that the adaptation surplus defined in Section 2 and detailed in Appendix A.3 is substantial.

The values in Figure 7A are mean values aggregated across results from 33 high-resolution climate models and all Monte Carlo simulation runs. However, as suggested in panel B of Figure 7, the full distribution of our estimated damages across climate models is right-skewed with a tail of potential mortality risk far higher than our central estimate. As evidence of this, the *median* value of the full mortality risk of climate change under RCP8.5 at end of century is 56 deaths per 100,000, as compared to a *mean* value of 85 and a 10^{th} to 90^{th} percentile range of [-21, 201]. We leave explicit pricing of the uncertainty arising from this distribution to future work, although we do propagate this uncertainty to compute a distribution of SCC estimates in Section 8.

Finally, Figure 7B and Appendix Figure F3 show the expected implications of mitigation for the

[54]We previously noted considerable heterogeneity across age-groups in our results. We will take this into account in our approach to valuing mortality damages monetarily in subsequent sections, and we display the underlying age group heterogeneity of these projections in Appendix F.

full mortality risk of climate change. The average estimate of the full mortality risk of climate change of 85 deaths per 100,000 by the end of the century under RCP8.5 falls to 14 under the emissions stabilization scenario of RCP4.5 (where emissions decline after 2050). For RCP4.5, the median end-of-century estimate is 9, and the 10^{th} to 90^{th} percentile range is [-45, 63].

Present day income and climate predict incidence of future climate change impacts. Whether the full mortality risk caused by climate change is realized through actual deaths (first term in expression (iv)), as opposed to costly compensatory investments (second term in expression (iv)), differs substantially across the globe. While some locations suffer large increases in mortality rates, others avoid excess mortality through expensive adaptation. Figures 8A-C demonstrate that current day income is strongly correlated with the composition of future damages. In panel A, the negative correlation indicates that today's poor locations tend to suffer large increases in mortality rates by end of century, while mortality rates tend to decline due to climate change in today's rich locations. However, there is large variance across impact regions within each income decile, implying that some poor regions are projected to experience mortality rate declines, and some wealthy regions mortality rate increases. In panel B, the positive correlation indicates 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. Panel C shows that the full mortality risk of climate change, the sum of both deaths and adaptation costs measured in death equivalents, is still borne disproportionately by regions that are poor today. On average, we find that in the poorest decile of today's income distribution, just 9% of the total burden of climate change induced mortality risk is borne as adaptation costs. In contrast, in the richest decile, on average approximately zero lives are lost in 2100 due to climate change, while adaptation costs are three times larger than they are in today's poorest regions. It is also apparent that poorer regions face higher uncertainty in the magnitude of their projected impacts (Figure F4).

Similar figures in panels D–F demonstrate that the hottest locations today suffer the largest increase in death rates, while the coldest pay the highest adaptation costs. The 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.

Climate change projection scenarios. The results in this section illustrate a single benchmark emissions and socioeconomic scenario (RCP8.5, SSP3). In Appendix F we report on the sensitivity of the 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 will depend greatly 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.

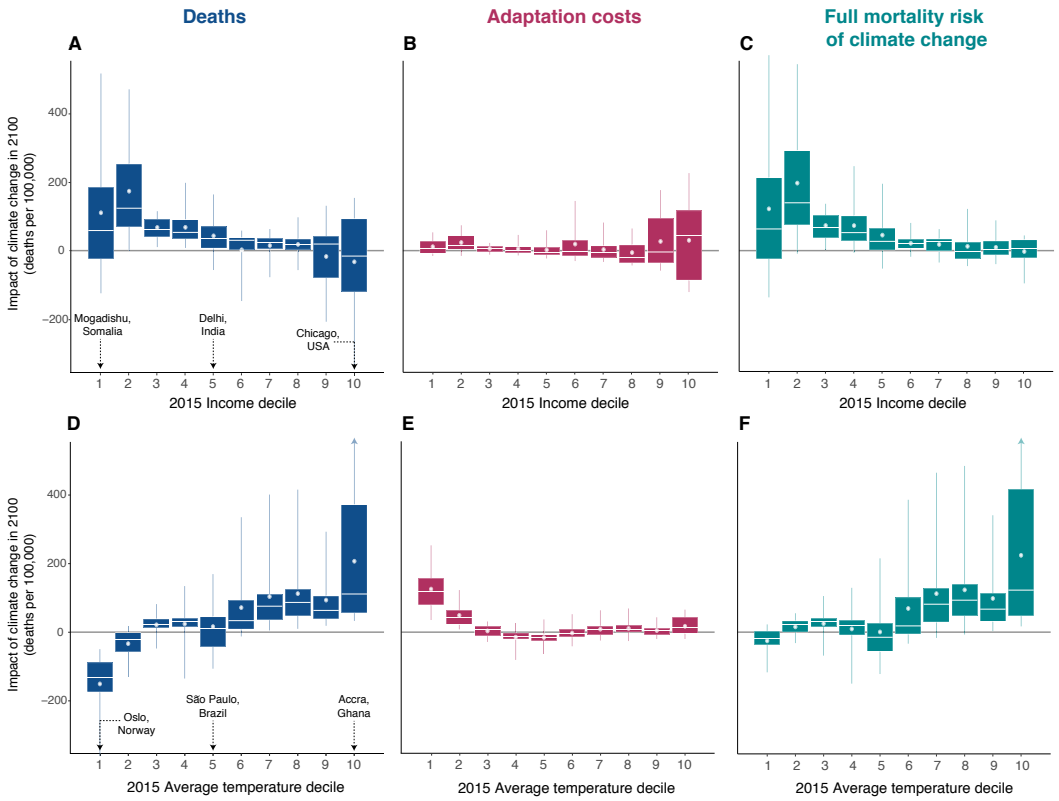

Figure 8: Climate change impacts and adaptation costs are correlated with present-day income and climate. Panels A and D show the change in annual mortality rates due to climate change in 2100 (RCP8.5, SSP3), accounting for the benefits of adaptation and income growth, against deciles of 2015 per capita income (A) and average annual temperature (D). Panels B and E show the annual adaptation costs incurred due to climate change in 2100, measured in death equivalents, from the same regions. Panels C and F show the full mortality risk due to climate change, which is the sum of deaths and adaptation costs measured in death equivalents. The income and average temperature deciles are calculated across 24,378 global impact regions and are population weighted using 2015 population values. All box plots show statistics of the distribution of estimated mean impacts across impact regions within a decile, where means are taken for each impact region across Monte Carlo simulations that account both for econometric and climate model uncertainty. Solid vertical lines in each box plot extend to the 5^{th} and 95^{th} percentiles of this distribution, boxes indicate the interquartile range, white horizontal lines indicate the median, and white circles indicate the mean.

8 A partial Social Cost of Carbon due to excess mortality risk

Time-series of projected costs due to excess mortality risk, as developed in the prior section, are alone insufficient to monetize the full social cost generated by the emission of a marginal ton of CO_2. This section outlines how we compute what we refer to as a "partial SCC" attributable to mortality risk. The mortality partial SCC represents the total global WTP (in net present value) of society to avert the excess mortality risk imposed by a marginal ton of emissions, including both the value of lost life as well as the cost of adaptations undertaken to protect populations. This calculation represents the component of the total Social Cost of Carbon that is mediated through excess mortality; however, it leaves out adverse impacts in other sectors of the economy, such as reduced labor productivity or changing food prices, which one would ideally include in the calculation of an efficient Pigouvian carbon tax. Hence, it is a mortality *partial* SCC.

There are three key steps to transforming projections of deaths due to climate change into a partial SCC attributable to excess mortality risk: monetizing damages, constructing a damage function, and computing marginal costs from a marginal CO_2 emission. This section describes these steps and reports results.

8.1 Monetizing damages

We follow the standard approach of using 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),[55] although the Appendix reports estimates based on the Ashenfelter and Greenstone (2004) estimate of $2.39 million (2019 USD).[56] We transform the VSL into a value per life-year lost using a method described in Appendix H.1. This approach 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.[57]

Assigning a total dollar value to global deaths also requires accounting for differences in income levels across the populations. In general, the VSL might vary with income because 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), in our primary estimate we 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.[58] However, we also show an alternative estimate where the VSL is adjusted only based on global average income and the lives of contemporaries are valued equally, regardless of their relative incomes. The former method 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 VSL and their own consumption (recall Equation 5). 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.

8.2 Constructing damage functions for excess mortality risk

The first step to computing a mortality partial SCC is combining econometric estimates with 33 high-resolution climate projections from the SMME into a damage function that describes the costs of excess mortality risk in a given year as a function of the overall level of climate change. Specifically, a damage function describes economic losses to an economy as a function of the change in *global mean surface*

[55] 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.

[56] 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.

[57] In our main specification, we use a simple life-years calculation that assigns each life-year lost the same economic value. 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 are slightly larger than those from the primary approach.

[58] 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.

temperature (ΔGMST),[59] a notion articulated at least as early as Nordhaus (1992). This function can be differentiated everywhere, allowing for marginal costs of a CO_2 impulse to be computed for any arbitrary global climate trajectory, as well as allowing us to disentangle the influence of economic uncertainty and climate science uncertainty.[60] Due to differences in the character of climate projections pre- and post-2100, there are some important differences in the approach for calculation of damage functions before and after 2100.

Computing damage functions through 2100. We generate estimates of the total value of mortality-related climate change damages (D_{irmt}) in each year (t) using many climate models (m), two emissions scenarios (r), and a resampling of parameter estimates that captures uncertainty in the estimation of mortality-temperature response functions (i). These multiple simulations lead to an empirically-derived distribution of potential economic outcomes that are conditional on the ΔGMST value for the year and climate model used to generate that projection. To construct damage functions, we then fit a conditional expectation function through these outcomes, following the approach in Hsiang et al. (2017). We interpret each projection of D_{irmt} as a potential realization of damages that result from the spatial distribution of warming in model m, conditional on the overall ΔGMST that is exhibited by that model under the emissions scenario r.[61]

An important difference between our approach and that in Hsiang et al. (2017) is that in our projections, the underlying population distribution and level of per capita income are both evolving over time and influencing the mortality sensitivity of warming. These changes over time necessitate year-specific damage functions. Thus, we condition on ΔGMST_{irmt} for each t to estimate a nonlinear damage function in each year:

$$D(\Delta GMST, t)_{irmt} = \alpha^t + \psi_1^t \Delta GMST_{rmt} + \psi_2^t \Delta GMST_{rmt}^2 + \varepsilon_{irmt} \tag{13}$$

using all 9,750 Monte Carlo simulation runs within a 5-year window of t, thereby allowing $D(\Delta GMST, t)_{irmt}$ to evolve flexibly over the century. We estimate Equation 13 as a quadratic, although all 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 this functional form choice in Appendix H.4.

Figure 9A illustrates the procedure for $t = 2097$, with D_{irmt} estimates from all Monte Carlo simulations shown as points[62] located along the horizontal axis based on their corresponding ΔGMST_{irmt}.

[59] *Global mean surface temperature* is defined as the global area-weighted average of surface air temperature over land and sea surface temperature over the oceans. Our climate change impacts are calculated relative to a baseline of 2001-2010. Therefore, we define changes in global mean surface temperature (ΔGMST) as relative to this same period.

[60] Appendix G provides a detailed description of why it is necessary to use a damage function to construct the mortality partial SCC, as opposed to a perturbation of each climate model with a pulse of CO_2.

[61] Note that the ΔGMST manifest in each model is an emergent summary parameter, resulting from the complex interaction of many physical elements of the model, including the *equilibrium climate sensitivity* of the model, a number that describes how much warming is associated with a specified change in greenhouse gas emissions. Differences in the spatial distributions of warming across models, and their mapping on populations around the world, remain an additional unresolved uncertain element of climate models that are idiosyncratic to each model.

[62] This scatterplot includes realizations under all RCP4.5 and RCP8.5 scenarios for all projections in our 33-member ensemble under our benchmark method of valuation – the EPA VSL, with an age-adjustment for expected life-years lost, and 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 D on GDP per capita and on demographics, we estimate separate damage functions for every SSP scenario. Results across different scenarios are shown in Appendix H.

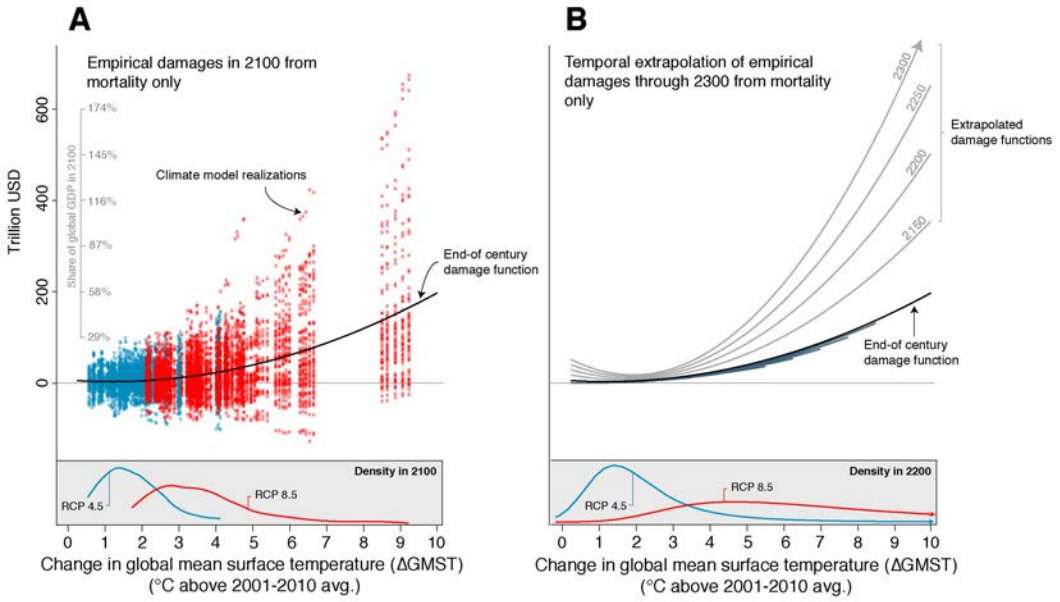

Figure 9: Empirically-derived mortality-only damage functions. Both panels show damage functions relating empirically-derived total global mortality damages to anomalies in global mean surface temperature (ΔGMST). In panel A, each point (red = RCP8.5, blue = RCP4.5) indicates the value of the full mortality 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. Additional damage functions are shown in blue for every 10 years pre-2100, each of which is estimated analogously to the end-of-century damage function, and in grey for every 50 years post-2100, each of which is extrapolated. 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 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 U.S. EPA VSL, an income elasticity of one applied to all impact regions, and a life-years adjustment.

The median end-of-century warming relative to 2001-2010 under RCP8.5 (red points) across our climate models is +3.7°C, while under RCP4.5 (blue points) it is +1.6°C. The black line is the quadratic damage function estimated for the year 2097, the latest year for which a full 5-year window of damage estimates can be constructed. Analogous curves are constructed for all years, 2015 to 2100.

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 13 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[63] 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 13, 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).[64] 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.

Panel B of Figure 9 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 and become steeper, as they did pre-2100. In Appendix H, we explore the importance of this extrapolation by using an alternative approach to estimating post-2100 damages, instead calculating partial SCC estimates using a damage function frozen at its 2100 shape for all years 2101-2300. With this alternative approach, our central estimate of the mortality partial SCC falls 19%, indicating that extrapolation of the damage function has a modest impact on our partial SCC estimates, due in part to the important role of discounting (Table H5).

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. The approach described above details the estimation 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.[65] As detailed above for the mean regression, extrapolation past the year 2100 is accomplished using a linear time interaction, estimated separately for each quantile. Central estimates of the mortality partial SCC reported below use the mean regression, while ranges incorporating damage uncertainty use the full set of time-varying quantile regressions.

[63] 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.

[64] The temporal trend over the entire 21^{st} century is convex, implying that our linearization is, if anything, conservative. The specific interaction model we estimate is: $D(\Delta GMST, t)_{irmt} = \alpha + \nu_1 \Delta GMST_{rmt} \times t + \nu_2 \Delta GMST^2_{rmt} \times t + \varepsilon_{irmt}$.

[65] We estimate a damage function for each of the following quantiles: 5, 25, 50, 75, and 95.

8.3 Computing marginal damages from a marginal carbon dioxide emissions pulse

The mortality partial Social Cost of Carbon at time t_0 is the marginal social cost from the change in mortality risk imposed by the emission of a marginal ton of CO_2 at t_0 holding everything else fixed, including the forecast trajectory of baseline greenhouse gas emissions. For a discount rate δ, this is:

$$Partial\ Social\ Cost\ of\ Carbon_{t_0} = \sum_{t_0}^{2300} e^{-\delta t} \frac{d\hat{D}(\Delta GMST, t)}{d\Delta GMST_t} \frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}} \tag{14}$$

where $\frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}}$ is the estimated increase in ΔGMST that occurs at each moment in time along the baseline climate trajectory as a result of a marginal unit of emissions at time t_0, which we approximate with an infinitesimally small pulse of CO_2 emissions occurring at time t_0. Its estimation requires a climate model capable of estimating the global temperature response in each year to a single pulse of CO_2 emissions. Because we are interested in computing this value for a large number of scenarios, including ones that reflect scientific uncertainty about the physical magnitude and timing of warming, referred to as *global climate sensitivity*, we use a "simple climate model" to estimate $\frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}}$. The values $\frac{d\hat{D}(\Delta GMST,t)}{d\Delta GMST_t}$ are the marginal damages at each moment in time 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.

Applying a simple climate model to the damage function. To calculate the change in ΔGMST_t due to a marginal pulse of CO_2 in 2020, we use the Finite Amplitude Impulse Response (FAIR) simple climate model that has been developed especially for this type of calculation (Millar et al., 2017).[66] We use FAIR to calculate ΔGMST_t trajectories for emissions scenarios RCP4.5 and RCP8.5, both with and without an exogenous impulse of 1Gt C (equivalent to 3.66Gt CO_2) in the year 2020, an approximation of an infinitesimal emission for which the model numerics are stable. In FAIR, this emissions impulse perturbs the trajectory of atmospheric CO_2 concentrations and ΔGMST for 2020-2300, with dynamics that are influenced by the baseline RCP scenario. In each scenario, the trajectory of ΔGMST_t in the "RCP + pulse" simulation is differenced from the baseline "RCP only" simulation to compute $\frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}}$, and the resulting damages are converted into USD per 1t CO_2.

Figure 10 shows the difference between the "RCP + pulse" and baseline RCP trajectories for emissions (panel A, showing the pulse), CO_2 concentrations (panel B), and ΔGMST (panel C) arising from the median values of FAIR's configuration parameter distributions.[67]

Finally, to compute the values in Equation 14, we calculate the stream of damages associated with both the baseline and perturbed trajectories, and discount their difference to compute the NPV of the excess mortality risk externality generated by this marginal emission, as illustrated in Figure

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

[67] Using the trajectories in Figure 10 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.

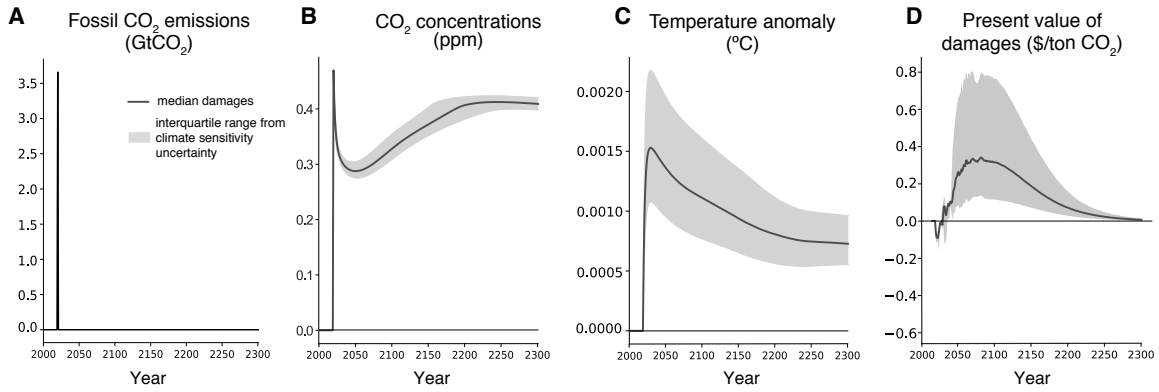

Figure 10: Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020. Panel A shows a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) in 2020 for emissions scenario RCP8.5. Panel B displays the effect of this pulse on atmospheric CO_2 concentrations, relative 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 using the FAIR simple climate model under median values for all climate sensitivity parameters. 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 function, using a 3% annual discount rate. The shaded area indicates the inter-quartile range of discounted marginal damages due to climate sensitivity uncertainty.

10D. The temporal pattern of the present value of mortality damages reflects the nonlinearity of the damage function (e.g. Figure 9), which itself depends on nonlinearities in location-specific mortality-temperature response functions, as well as the discount rate (3% is shown).

Uncertainty in global climate sensitivity. Throughout the analysis, we use the SMME to ensure our projections of climate change impacts fully reflect uncertainty in climate sensitivity. To ensure this same uncertainty regarding the climate system is carried through to our mortality partial SCC estimates, we generate simulations with the FAIR model to account for uncertainties in key parameter values that jointly determine the sensitivity of the climate to marginal emissions in the FAIR model. To do so, we resample from a joint distribution of four key FAIR parameters: 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 joint distribution is constrained such that the distribution of transient warming responses they produce matches the corresponding distributions from the IPCC Assessment Report 5 (AR5).[68] Because the SMME is also constructed to match the probabilistic distribution of equilibrium climate sensitivity in the IPCC AR5, this process ensures that the climate sensitivity uncertainty we capture within FAIR matches that from the SMME (see Figure G3).

We sample from this distribution and compute a mortality partial SCC for each vector of climate parameters drawn. The resulting distribution of mortality partial SCC values provides us with a sense of the overall uncertainty in the partial SCC that is attributable to unresolved physical uncertainty in global climate sensitivity. We report a range of SCC values accounting only for uncertainty in these climate sensitivity parameters (using the damage function estimated as the conditional expectation function through the full distribution of climate impact results, e.g. see Figure 9), in addition to a broader range of SCC values that additionally account for both climate sensitivity and damage function

[68] Details of this procedure are provided in Appendix G.

uncertainty. The latter use the damage function distribution estimated as quantile regressions through the full distribution of climate impact results in combination with all parameter combinations for FAIR.

8.4 Estimates of the partial Social Cost of Carbon due to excess mortality risk

Global aggregate estimates. Mortality partial SCC estimates are shown in panel A of Table 4, applying the three different annual discount rates used in prior estimates of the social cost of carbon (2.5%, 3%, and 5%) (Greenstone, Kopits, and Wolverton, 2013; National Academies of Sciences, Engineering, and Medicine, 2017) and values of the VSL that adjust for cross sectional variation in incomes among contemporaries and global income growth. In panel B, we use the alternative valuation approach described in Section 8.1, which assumes there is a globally homogeneous VSL in each year that evolves over time based on global income growth. The central estimates shown utilize the median values of FAIR's four key parameter distributions and the expected global damage function. Interquartile ranges (IQRs) are also shown in Table 4, reflecting uncertainty in climate sensitivity (uncertainty in the simple climate model FAIR), damage function uncertainty only (uncertainty arising from the empirical distribution of mortality-temperature response function parameters and from climate model uncertainty about the spatial distribution of warming), and both sources of uncertainty jointly. All values represent the global sum of each impact region's WTP today (2019 USD) to avoid changes to the trajectory of future mortality risk induced by warming from an additional metric ton of CO_2 emissions released in 2020, including both the costs and benefits of adaptation.

When following the Interagency Working Group on Social Cost of Greenhouse Gases (2016) preference for a discount rate of $\delta = 3\%$, the central estimate of the mortality partial SCC is \$4.4 per metric ton of CO_2 for the low to moderate emissions scenario (RCP 4.5) and \$38.1 per metric ton for the high emissions scenario (RCP 8.5). Utilizing the valuation that relies on a global average VSL in panel B, these central estimates increase to \$16.7 and \$64.5 per metric ton, respectively. In the primary specification ($\delta = 3\%$, Panel A) we estimate a mortality partial SCC IQR with full uncertainty of [-\$26.90, \$40.1] for the low to moderate emissions scenario (RCP 4.5) and an IQR of [-\$7.6, \$87.4] for the high emissions scenario (RCP 8.5). 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. 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 after accounting for climate sensitivity and damage function uncertainty, with a long right tail: the 95^{th} and 99^{th} percentiles of the partial SCC with $\delta = 3\%$ for RCP8.5 are \$387.2 and \$1,001.6, respectively.

Additional estimates, including different valuation assumptions and alternative socioeconomic scenarios, are presented in Appendix Tables H5-H4. It is noteworthy that alternative socioeconomic scenarios do not qualitatively alter the estimates, particularly for the high emissions scenario. In contrast, the choice of VSL plays a pivotal role in determining the SCC.

Table 4: Estimates of a partial Social Cost of Carbon for excess mortality risk incorporating the costs and benefits of adaptation

	Annual discount rate		
	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Globally varying valuation of mortality risk (2019 USD)			
RCP 4.5	5.0	4.4	2.2
Climate sensitivity uncertainty	[-6.5, 45.4]	[-3.1, 29.8]	[0.1, 8.5]
Damage function uncertainty	[-40.5, 52.2]	[-29.1, 38.4]	[-12.9, 17.0]
Full uncertainty	[-35.6, 58.0]	[-26.9, 40.1]	[-12.8, 15.2]
RCP 8.5	62.7	38.1	8.6
Climate sensitivity uncertainty	[25.4, 153.9]	[15.4, 94.0]	[3.5, 21.1]
Damage function uncertainty	[0.6, 103.7]	[-7.5, 71.6]	[-11.4, 26.3]
Full uncertainty	[1.0, 139.6]	[-7.6, 87.4]	[-11.2, 25.4]
Panel B: Globally uniform valuation of mortality risk (2019 USD)			
RCP 4.5	18.9	16.7	9.8
Climate sensitivity uncertainty	[-1.2, 82.4]	[2.5, 58.8]	[4.4, 22.9]
Damage function uncertainty	[-39.7, 98.3]	[-28.6, 73.5]	[-12.6, 34.4]
Full uncertainty	[-34.4, 104.8]	[-25.9, 76.0]	[-13.1, 33.2]
RCP 8.5	98.9	64.5	19.8
Climate sensitivity uncertainty	[42.7, 233.8]	[28.4, 150.6]	[9.7, 42.9]
Damage function uncertainty	[9.3, 156.1]	[-2.1, 112.5]	[-10.2, 45.9]
Full uncertainty	[6.6, 210.6]	[-3.6, 140.1]	[-10.4, 48.3]

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 G). 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, all regions have heterogeneous valuation, based on local income. In Panel B, all regions globally are given the population weighted global average VSL, after scaling using income. All estimates use a value of life years adjustment, valuing deaths by the expected number of life-years lost. Point 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. The uncertainty ranges are interquartile ranges [IQRs] showing the influence of climate sensitivity and damage function uncertainty (see Appendix G for details).

9 Limitations of this analysis

As the paper has detailed, the mortality risk partial SCC has many ingredients. We have tried to probe the robustness of the results to each of them, but there are four issues that merit highlighting when interpreting the results.

Migration responses First, the paper's estimates do not reflect the possibility of migration responses to climate change. If migration were free, it seems likely that the mortality risk partial SCC would be smaller, as people from regions with high damages (such as sub-Saharan Africa) may move to regions with low or even negative damages (such as Scandinavia). However, recent history in the U.S. and around the globe underscores that borders are meaningful and that there are substantial political costs to migration, in addition to migrants' moving costs, 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).

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. However, the limited availability of high-resolution humidity projections and the highly uncertain projections of humidity under climate change (Sherwood and Fu, 2014) mean that we leave exploration of this heterogeneity to future work.

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 the 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. The limitation of this approach is that in some location and year combinations, we must make out-of-sample predictions about how temperature sensitivity will diminish beyond that observed anywhere in the world today, as temperatures and incomes rise outside of the support in existing global cross-section.

A further limitation is that climate change is likely to lead to climate-biased technical change that lowers the *relative* costs of goods which reduce the health risks of high temperatures, and our estimates of the costs and benefits of adaptation do not account for this possibility. Thus, while the projections allow for continued trends in adaptation that have been observed historically, they may overstate the mortality risk of climate change if directed technological innovations lower the costs of adapting to high temperatures.

Uncertainty Throughout this paper, we have emphasized central estimates of the mortality risk partial SCC. However, our full set of estimates reveals a remarkable degree of uncertainty, stemming from three distinct sources: uncertain spatial distribution of warming, statistical uncertainty, and uncertainty regarding global climate sensitivity. While we are careful to account for and integrate these diverse sources of uncertainty in each step in the partial SCC construction, we hope that future work will enable a narrowing of these estimated ranges. However, at least some of this uncertainty is fundamentally unresolvable, and an important exercise for future research is to price this uncertainty, which would increase the estimated partial SCC when agents are risk averse (Weitzman, 2011; Traeger, 2014).

10 Conclusion

This paper has outlined a new method for empirically estimating the costs of climate change for a single sector of the economy and implemented it in the context of mortality risks associated with temperature change. There are several noteworthy methodological innovations and intermediate findings.

First, the relationship between mortality rates and temperature is highly nonlinear and varies with

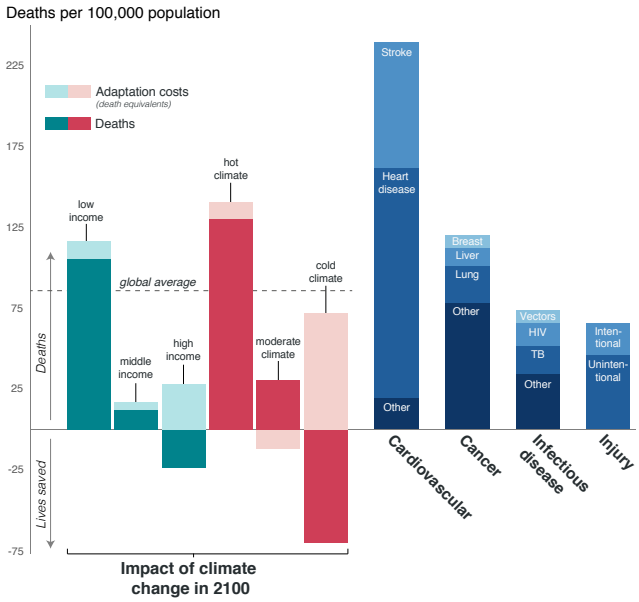

Figure 11: The impact of climate change in 2100 is comparable to contemporary leading causes of death. Impacts of climate change (teal and coral) are calculated for the year 2100 under RCP8.5 and SSP3 and include changes in death rates (solid colors) and changes in adaptation costs, measured in death equivalents (light shading). Income and average climate groups 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).

a location's income and climate. These findings were only possible due to the collection and analysis of highly resolved data covering roughly half of the world's population, which enabled us to estimate flexible empirical models relating mortality risk to temperature, climate, and income.

Second, the costs of climate change induced mortality risks are distributed unevenly across the 24,378 regions that we use to create local-level projections. 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 230 lives per 100,000 in Oslo, Norway. To put these numbers in perspective, our results suggest that the impact of climate change on mortality will be comparable globally to leading causes of death today, such as cancer and infections disease (Figure 11). Today's poor bear a disproportionately high share of the global mortality risks of climate change, as current incomes (as well as current average temperatures) are strongly correlated with future climate change impacts.

Third, the heterogeneous impacts appear to reflect investments that individuals and societies make based on the costs and benefits of responding to differences in climate. We outline and implement a revealed preference method to recover the costs of these adaptation investments, even though they cannot be directly observed. The results suggest that 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. Today's poor locations tend to bear a larger share of the burden in the form of direct mortality impacts, while today's rich face large increases in adaptation costs (Figure 11). Relatedly, estimated mortality impacts that do not account for adaptation overstate the mortality impacts of

climate change in 2100 by more than a factor of 2.6.

Fourth, there is substantial climate and statistical uncertainty around these estimates and we find that the distribution of projected losses is right skewed; for example the mean loss in 2100 is about 52% larger than the median loss. Although we do not account for risk aversion, it is evident that doing so would increase the valuations of these impacts.

The paper's ultimate goal is to develop the first empirically grounded partial SCC that reflects the consequences of climate change on mortality and investments in adaptation. Our central values suggest that with a 3% discount rate, the present value of excess mortality risk due to climate change imposed by a marginal metric ton of CO_2 emissions in 2020 is roughly \$4.4 (in 2019 USD) with a moderate emissions scenario (RCP4.5) and \$38.1 with a high one (RCP8.5). When accounting for climate model and statistical uncertainty, the respective interquartile ranges are [-\$26.9, \$40.1] for RCP4.5 and [-\$7.6, \$87.4] for RCP8.5; the positive skewness of these ranges reflects the risk of outcomes substantially more costly than our central estimate.

As a basis of comparison, the Obama Administrations's initial central estimate for the full SCC (including all sectors and using a discount rate of 3%) is \$50.6 (in 2019 USD) per metric ton of CO_2 emitted in 2020 (Interagency Working Group on Social Cost of Greenhouse Gases, 2016). This value was derived using simulations from the DICE, FUND, and PAGE integrated assessment models. The DICE and PAGE models do not document the origin of their damages in a manner that allows calculation of a mortality risk partial SCC that is comparable to our results (Rose, Diaz, and Blanford, 2017). However, Diaz (2014) computes comparable partial SCC values for the FUND model (3% discount rate, "business as usual" emissions scenario) and reports values for three comparable health impacts (diarrhea, vector born diseases, and cardiovascular/respiratory impacts) that total less than \$2.0 (2019 USD).[69] It is evident that our empirically grounded estimates of the costs of climate-induced mortality risks substantially exceed the available estimates from the models that underlie the existing estimates of the social costs of carbon. Indeed, with RCP8.5 (which is similar to the baseline assumptions underlying the Obama Administration's SCC), this paper's excess mortality *partial* SCC is 75% of the Obama Administration's *full* SCC.

The climate change challenge is considered existential by some and a relatively small risk by others, yet much of what is known about its overall impacts, particularly the SCC, comes from integrated assessment models that do not sit on a robust empirical foundation. In particular, many models currently used to compute the SCC are either not calibrated against data, have a calibration that is not documented, or are calibrated against empirical estimates that are not derived from modern empirical techniques and are unlikely to be globally representative. Advances in access to data and computing render these modeling choices unnecessary.

We believe that this paper has highlighted a key role for systematic empirical analysis in providing a clearer picture of how, why, and where costs of climate change are likely to emerge in the future. Looking ahead, this paper's general approach can be applied to other aspects of the global economy besides mortality risk, and we believe that doing so is a promising area for future research.

[69] These partial SCC values from FUND differ for numerous reasons, including that they were calibrated to different studies published between 1995-2004. See Diaz (2014) for details.

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. 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.

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.

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.

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

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.

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

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." .

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

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.

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.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2013. "Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation." *Review of Environmental Economics and Policy* 7 (1):23–46.

Grossman, Michael. 1972. "On the concept of health capital and the demand for health." *Journal of Political Economy* 80 (2):223–255.

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 and Robert E Kopp. 2018. "An Economist's Guide to Climate Change Science." *Journal of Economic Perspectives* 32 (4):3–32.

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 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.

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* .

Lobell, David B, Michael J Roberts, Wolfram Schlenker, Noah Braun, Bertis B Little, Roderick M Rejesus, and Graeme L Hammer. 2014. "Greater sensitivity to drought accompanies maize yield increase in the US Midwest." *Science* 344 (6183):516–519.

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.

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.

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.

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.

Rose, Steven K, Delavane B Diaz, and Geoffrey J Blanford. 2017. "Understanding the social cost of carbon: a model diagnostic and inter-comparison study." *Climate Change Economics* 8 (02):1750009.

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.

Schlenker, Wolfram, Michael J Roberts, and David B Lobell. 2013. "US maize adaptability." *Nature Climate Change* 3 (8):690–691.

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.

Traeger, Christian P. 2014. "Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity." *Economic Theory* 56 (3):627–664.

United Nations. 2017. "World Mortality 2017." Tech. rep., United Nations Economic and Social Affairs.

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* .

Weitzman, Martin L. 2011. "Fat-tailed uncertainty in the economics of catastrophic climate change." *Review of Environmental Economics and Policy* 5 (2):275–292.

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

ONLINE APPENDIX

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[70] 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 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_1 and C_2, 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

[70] In Appendix A.2 below, as in the main text, A are net costs since they are net any utility benefits of $\boldsymbol{b}$.

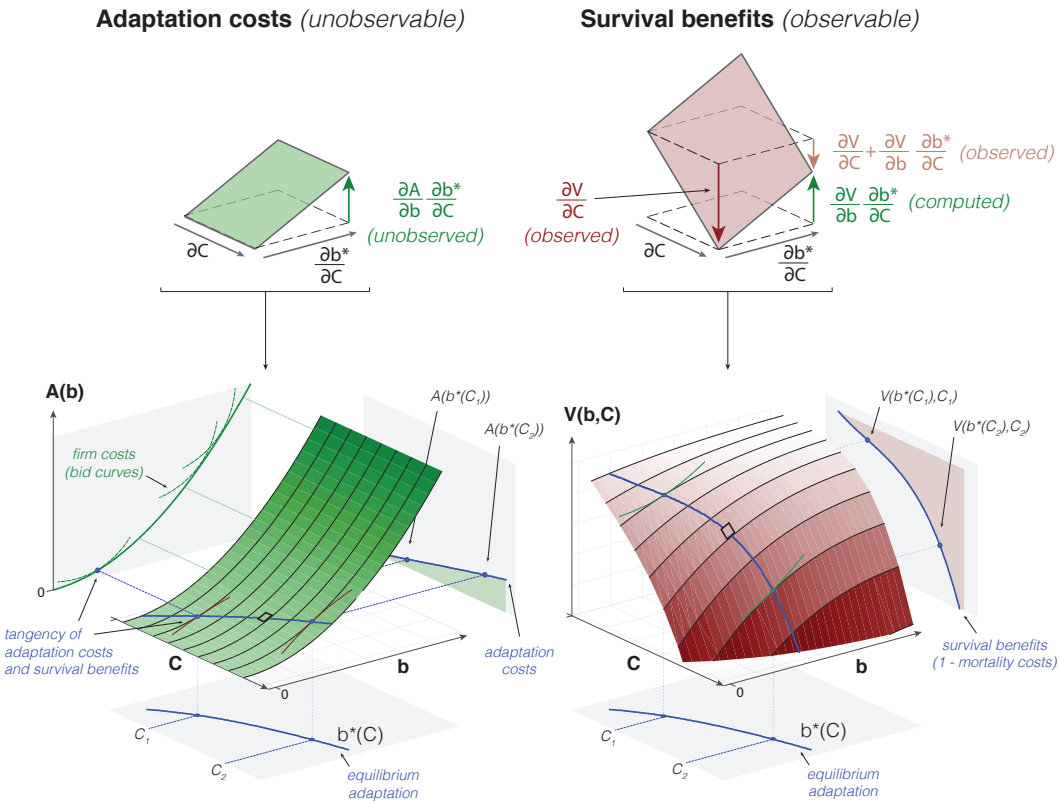

Figure A1: Use of revealed preference to recover WTP for an unobservable adaptation. Horizontal dimensions are climate C, representing temperature, and adaptation level b. Vertical dimensions are adaptation costs $A(b)$ in the left panel and expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b, C)]$ 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.

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 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_1 to C_2, individuals would continuously respond by adapting along $b^*(C)$ and traveling up the cost surface in the lower left panel, eventually incurring costs $A(b^*(C_2))$ rather than the initial costs $A(b^*(C_1))$ that they incurred prior to warming. We point out that the change in this total adaptation cost $A(b^*(C_2)) - A(b^*(C_1))$ 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_1 and C_2, 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 $A(b^*(C_2)) - A(b^*(C_1))$ for the non-marginal change in climate from C_1 to C_2. This intuition holds for an unknown number of margins of adaptation and a climate of arbitrary dimension, which we allow for below.

A.2 Mathematical details

The framework in Section 2 allows us to use observable relationships between mortality risk and changes in the climate to generate estimates of the unobserved cost of adaptation. Here, we provide details on the derivation of our final expression in Equation 8, which denotes the total adaptation costs incurred as a population experiences a change in climate from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$.

Recall that agents jointly choose K margins of adaptation described by the choice vector $\boldsymbol{b}$ such that $\boldsymbol{b}^* = \arg\max u(x, \boldsymbol{b})[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, subject to a budget constraint $h(\boldsymbol{b}) + x = Y$, where $h(\boldsymbol{b})$ is the pecuniary cost of adaptation expended by households and Y is exogenous income. $1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})$ represents expected survival probability, and we assume $\tilde{f}(\cdot)$ is continuous and differentiable. We assume that there exists a numeraire good x for which $u(\cdot)$ is quasilinear. The first order conditions of this optimization problem take the form:

$$[1-\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})]\frac{\partial u(x^*, \boldsymbol{b}^*)}{\partial x} = \lambda, \qquad [1-\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})]\frac{\partial u(x^*, \boldsymbol{b}^*)}{\partial b_k} - \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k} u(x^*, \boldsymbol{b}^*) = \lambda \frac{\partial h(\boldsymbol{b}^*)}{\partial b_k} \quad \forall k \in \{1, ..., K\}$$

where λ is the Lagrange multiplier. Note that because $\boldsymbol{b}$ is a vector, the second equality holds for K separate equations; there is a distinct first order condition for each b_k. Substitution of the first FOC

into the second equality gives the following expression, which holds for each b_k:

$$\underbrace{-\frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k} u(x^*, \boldsymbol{b}^*)}_{\substack{\text{marginal survival} \\ \text{benefit of } b_k \text{ (in utils)}}} = \underbrace{[1 - \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})] \left[\frac{\partial u(x^*, \boldsymbol{b}^*)}{\partial x} \frac{\partial h(\boldsymbol{b}^*)}{\partial b_k} - \frac{\partial u(x^*, \boldsymbol{b}^*)}{\partial b_k} \right]}_{\substack{\text{marginal adaptation cost of } b_k \text{, net of} \\ \text{marginal consumption value of } b_k \text{ (in utils)}}} \quad \forall k \in \{1, ..., K\}$$

The righthand side of the above expression represents the marginal cost of changing b_k, valued in utils, net of any change in utility realized from direct utility effects of adaptation. The lefthand side is simply the marginal mortality benefit realized from a change in b_k. We translate this expression into dollars by multiplying both sides of the expression by $\frac{1}{\lambda} = \frac{1}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\frac{\partial u(x,\boldsymbol{b})}{\partial x}}$, the inverse of the shadow value of the budget constraint, and substituting for the standard Value of a Statistical Life (VSL) definition: $VSL = \frac{u(x,\boldsymbol{b})}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\frac{\partial u(x,\boldsymbol{b})}{\partial x}}$.[71] Finally, summing across all choice variables b_k in the vector $\boldsymbol{b}$ gives the expression in Equation 5 from the main text:

$$-VSL \sum_k \underbrace{\frac{\partial}{\partial b_k} \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}_{\substack{\text{marginal survival} \\ \text{benefit of } b_k}} = \underbrace{\sum_k \frac{\partial}{\partial b_k} \left[\underbrace{h(\boldsymbol{b}^*)}_{\substack{\text{pecuniary} \\ \text{costs of } \boldsymbol{b}}} - \underbrace{u(x^*, \boldsymbol{b}^*) \frac{1}{\partial u(x^*, \boldsymbol{b}^*)/\partial x}}_{\text{non-mortality benefits of } \boldsymbol{b}} \right]}_{\text{net marginal cost of } b_k = \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k}} = \sum_k \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k} \quad \text{(A.1)}$$

where we define net adaptation costs $A(\boldsymbol{b})$ over the space $\boldsymbol{b}$ such that its derivative with respect to each b_k is equal to $\frac{\partial h(\boldsymbol{b}^*)}{\partial b_k} - \frac{\partial u(x^*, \boldsymbol{b}^*)/\partial b_k}{\partial u(x^*, \boldsymbol{b}^*)/\partial x}$. We use Equation A.1 (equivalent to Equation 5 in the main text), along with the definition of the total derivative in Equation 7, to derive our final expression of the total adaptation costs incurred over some path through the $\boldsymbol{C}$ subspace from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$, as follows

[71] Note that the VSL is defined as the willingness to pay for a marginal increase in the probability of survival. This willingness to pay is held constant here for clarity; however, in our empirical derivation, we allow the VSL to evolve with income, under an income elasticity of one.

(omitting the arguments of $\boldsymbol{b}^*$ for clarity after the first equality):[72]

$$A(\boldsymbol{b}^*(\boldsymbol{C}_2, Y)) - A(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) = \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} \frac{dA(\boldsymbol{b}^*(\boldsymbol{C}, Y))}{d\boldsymbol{C}} d\boldsymbol{C}$$

$$\textit{(Chain Rule)} \qquad = \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} \left[\sum_k \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k} \frac{\partial b_k^*}{\partial \boldsymbol{C}} + \underbrace{\frac{\partial A(\boldsymbol{b}^*)}{\partial \boldsymbol{C}}}_{=0} \right] d\boldsymbol{C}$$

$$\textit{(substitution from Equation A.1)} \qquad = -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL \sum_k \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k} \frac{\partial b_k^*}{\partial \boldsymbol{C}} d\boldsymbol{C}$$

$$\textit{(substitution from Equation 7)} \qquad = -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial \boldsymbol{C}} \right] d\boldsymbol{C} \tag{A.2}$$

The empirical implementation of Equation A.2 (equivalent to Equation 8) is detailed in Section 6.3.

A.3 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 $\boldsymbol{b}^*(\boldsymbol{C})$ (Equation 5) 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 7).

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 8). This surplus can be evaluated at any future climate $\boldsymbol{C}_2$. That is, adaptation surplus under a climate changing from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ can be written as:[73]

$$\textit{Adaptation surplus } (\boldsymbol{C}_1 \to \boldsymbol{C}_2) = \underbrace{-VSL[\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_2), \boldsymbol{C}_2) - \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_2)]}_{\text{total adaptation benefits}} - \underbrace{[A(\boldsymbol{b}^*(\boldsymbol{C}_2)) - A(\boldsymbol{b}^*(\boldsymbol{C}_1))]}_{\text{total adaptation costs}}$$

$$= -\int_{\boldsymbol{b}^*(\boldsymbol{C}_1)}^{\boldsymbol{b}^*(\boldsymbol{C}_2)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_2)}{d\boldsymbol{b}} d\boldsymbol{b} - \int_{\boldsymbol{b}^*(\boldsymbol{C}_1)}^{\boldsymbol{b}^*(\boldsymbol{C}_2)} \frac{\partial A(\boldsymbol{b})}{\partial \boldsymbol{b}} d\boldsymbol{b}^* \tag{A.3}$$

[72] Note that income Y is treated as exogenously determined, and therefore does not directly enter the expression for total adaptation costs. In actual calculations, its value can be varied period-to-period through changing the VSL (see Section 6.3 for details).

[73] Note that income only influences the calculation of surplus arising from climate-driven adaptation via changes in the VSL.

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.3 is the total benefits of adaptation, defined as [minus] the mortality effects of climate $\boldsymbol{C}_2$ with optimal adaptation (i.e. $\boldsymbol{b}^*(\boldsymbol{C}_2)$) minus the mortality effects of that same climate, but with adaptation fixed at its initial level (i.e. $\boldsymbol{b}^*(\boldsymbol{C}_1)$). The second term is the total costs of adaptation, defined as the adaptation costs under optimal adaptation in climate $\boldsymbol{C}_2$ minus adaptation costs under optimal adaptation in the initial climate $\boldsymbol{C}_1$. Adaptation benefits (the first term) can be computed by integrating $\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_2)}{d\boldsymbol{b}}$, the marginal mortality effect of adaptation evaluated at fixed climate $\boldsymbol{C}_2$. Note that this integration is not computed over the optimal pathway, as the climate is fixed at $\boldsymbol{C}_2$ and any $\boldsymbol{b} \neq \boldsymbol{b}^*(\boldsymbol{C}_2)$ 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.3 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 $\boldsymbol{C}$. 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 and Appendix A.2 above), total adaptation costs, the second term in Equation A.3, can be rewritten as an integral over $\boldsymbol{C}$ using a simple change of variables. However, rewriting total adaptation benefits, the first term in Equation A.3, as an integral over $\boldsymbol{C}$ 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.3), 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_1 to C_2 are represented by the the vertical difference 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 $b^*(C_2)$ when experiencing climate C_2, instead of holding adaptation fixed at its initial level $b^*(C_1)$. 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 $b^*(C_1)$ to $b^*(C_2)$. This integration along the green line represents the definition of total adaptation benefits written in Equation A.3. 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 $b^*(C_1)$ and integrating $V(b, C)$ over C from C_1 to C_2), and then subtracting off the value of the survival impacts of the optimal pathway from C_1 to C_2 (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.

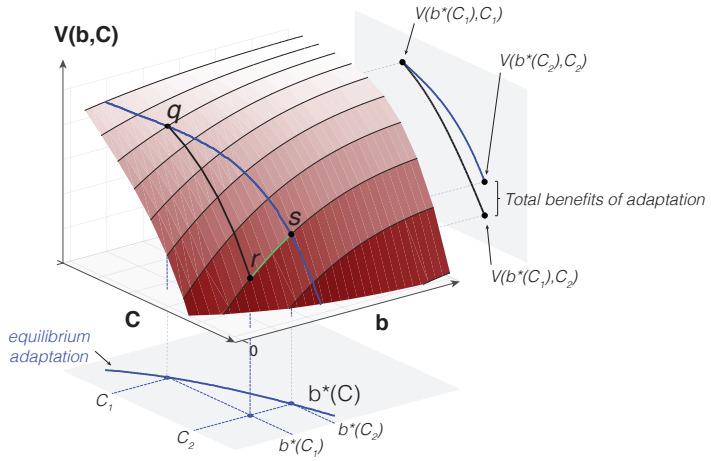

Figure A2: Recovering total benefits of adaptation using revealed preference. Horizontal dimensions are climate C, representing temperature, and adaptation level b. The vertical dimension is expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b, C)]$, in units of dollars of WTP. The equilibrium adaptation path $\{b^*(C), C\}$ is drawn in blue (line $q \to s$), and the off-equilibrium path $\{b^*(C_1), C\}$ is drawn in black (line $q \to r$). To derive the total benefits of adaptation under a change in climate from C_1 to C_2, we integrate the surface along the green line (line $r \to s$), evaluating changes in survival benefits at a fixed climate C_2, as adaptation evolves from $b^*(C_1)$ to $b^*(C_2)$. The magnitude of total adaptation benefits is shown on the $V \times C$ plane on the right panel.

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.3 writes total adaptation benefits using a path along the surface in the $\boldsymbol{b}$ dimension between the end points $\{\boldsymbol{b}^*(\boldsymbol{C}_2), \boldsymbol{C}_2\}$ and $\{\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_2\}$.[74] 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 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 $\{\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_1\}$ (analogous to point q in Figure A2) and traverses along the off-equilibrium path from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ with adaptation fixed at $\boldsymbol{b}^*(\boldsymbol{C}_1)$ (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 $\boldsymbol{C}_2$, to arrive at $\{\boldsymbol{b}^*(\boldsymbol{C}_2), \boldsymbol{C}_2\}$ (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

[74] 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.3. For parsimony, we use the latter notation here and in the subsequent expressions.

the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$ (analogous to the blue line between q and s in Figure A2):

$$\begin{aligned}\oint_L \nabla[VSL\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\cdot\partial\boldsymbol{b}\partial\boldsymbol{C} &= \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C})}{\partial\boldsymbol{C}}d\boldsymbol{C} + \int_{\boldsymbol{b}^*(C_1)}^{\boldsymbol{b}^*(C_2)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_2)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &\quad + \int_{\boldsymbol{C}_2}^{\boldsymbol{C}_1} VSL\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{d\boldsymbol{C}}d\boldsymbol{C} \\ &= 0 \end{aligned} \tag{A.4}$$

By rearranging Equation A.4 (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:

$$\begin{aligned}\textit{Total adaptation benefits} &= -\int_{\boldsymbol{b}^*(C_1)}^{\boldsymbol{b}^*(C_2)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_2)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &= -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C})}{\partial\boldsymbol{C}}\right]d\boldsymbol{C}\end{aligned} \tag{A.5}$$

Using Equation A.5 and a change of variables to rewrite the total costs of adaptation as an integral over $\boldsymbol{C}$, we can rewrite Equation A.3 as:

$$\begin{aligned}\textit{Adaptation surplus}\,(\boldsymbol{C}_1\to\boldsymbol{C}_2) &= -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\left[\underbrace{\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{d\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/ adaptation}}} - \underbrace{\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C})}{\partial\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/o adaptation}}}\right]d\boldsymbol{C} \\ &\quad - \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2}\frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}))}{\partial\boldsymbol{b}}\frac{\partial\boldsymbol{b}^*}{\partial\boldsymbol{C}}d\boldsymbol{C}\end{aligned} \tag{A.6}$$

While the total adaptation benefits term in Equation A.6 (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.6 in terms of objects that are measurable, using Equation 8 from the main text to substitute for the object $\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2}\frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}))}{\partial\boldsymbol{b}}\frac{\partial\boldsymbol{b}^*}{\partial\boldsymbol{C}}d\boldsymbol{C}$:

$$\begin{aligned}\textit{Adaptation surplus}\,(\boldsymbol{C}_1\to\boldsymbol{C}_2) &= -\int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C})}{\partial\boldsymbol{C}}\right]d\boldsymbol{C} \\ &\quad + \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{d\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{\partial\boldsymbol{C}}\right]d\boldsymbol{C} \\ &= \int_{\boldsymbol{C}_1}^{\boldsymbol{C}_2} VSL\left[\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C})}{\partial\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}),\boldsymbol{C})}{\partial\boldsymbol{C}}\right]d\boldsymbol{C}\end{aligned} \tag{A.7}$$

In Equation A.7, 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, $\boldsymbol{C}_1$. 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.7 will be identical when $\boldsymbol{C} = \boldsymbol{C}_1$, if they diverge at some point after $\boldsymbol{C}$ warms beyond $\boldsymbol{C}_1$, then surplus will be nonzero. Thus, a sufficient condition for positive surplus is:

$$\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C})}{\partial \boldsymbol{C}} > \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}), \boldsymbol{C})}{\partial \boldsymbol{C}} \quad \forall \boldsymbol{C} \in (\boldsymbol{C}_1, \boldsymbol{C}_2] \tag{A.8}$$

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.8 is weakly positive, and the total adaptation surplus over the climate trajectory $\boldsymbol{C}_1 \to \boldsymbol{C}_2$ is positive.

A.4 Implementation details for the empirical estimation of adaptation costs

In Section 6.3, we describe how we use econometric estimation of Equation 10 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 and Appendix A.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 8), and integrating this difference. To empirically construct an estimate of these costs, we begin by taking expectations of Equation 10 to specify our empirically estimated expected mortality risk for an age group a in region r for year t:

$$\hat{\tilde{f}}(.)_{art} \equiv \mathrm{E}[\hat{f}(.)_{art}] = \mathrm{E}[\underbrace{\hat{g}_a\left(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt}\right)}_{\hat{g}_{art}(\cdot)}] + ... \tag{A.9}$$

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 10. The expectation of $\hat{g}(\cdot)$ is computed over realizations of temperature for region r in year t from the prior 30 years, with weights of historical observations linearly declining in time, analogously to the construction of $TMEAN$ (see Section 6.2 and Appendix E.1). Below we omit subscripts for clarity, but the following calculation is conducted yearly for each age and region separately.

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 $\boldsymbol{C}_1$ ($TMEAN_1$) and income Y_1 ($\log(GDPpc)_1$). Climate and income are evaluated in an initial period t_1 and are "held fixed", since the partial effect must exclude any compensatory behavior by populations. The partial effect of the climate on expected mortality risk is then:

$$\frac{\partial \hat{\tilde{f}}}{\partial \boldsymbol{C}} = \frac{\partial \hat{\tilde{f}}}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} = \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_1, Y_1} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} \tag{A.10}$$

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 $TMEAN$. Because income effects are not caused by changes in the climate, income Y_1 is held fixed.[75] The total effect of the climate on expected mortality risk is then:

$$\begin{aligned} \frac{d \hat{\tilde{f}}}{d \boldsymbol{C}} = \frac{\partial \hat{\tilde{f}}}{\partial \boldsymbol{C}} + \frac{\partial \hat{\tilde{f}}}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{C}} &= \frac{\partial \hat{\tilde{f}}}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} + \frac{\partial \hat{\tilde{f}}}{\partial TMEAN} \frac{\partial TMEAN}{\partial \boldsymbol{C}} \\ &= \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_1, Y_1} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{C}} + \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN} \right|_{\boldsymbol{C}, Y_1} \frac{\partial TMEAN}{\partial \boldsymbol{C}} \end{aligned} \tag{A.11}$$

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 11 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.11 and A.10:

$$\frac{d \hat{\tilde{f}}_t}{d \boldsymbol{C}} - \frac{\partial \hat{\tilde{f}}_t}{\partial \boldsymbol{C}} = \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN} \right|_{\boldsymbol{C}, Y_1} \frac{\partial TMEAN}{\partial \boldsymbol{C}} \tag{A.12}$$

The righthand side of Equation A.12 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

[75] Note that throughout this paper, we treat income and demographic trajectories as exogenously given by the Shared Socioeconomic Pathways (SSPs) and thus these variables are not affected by warming (although they vary through time). In future work, these covariates, which determine the sensitivity of mortality rates to temperature exposure, could be endogenized into the climate change projection.

Equation A.12 into Equation 8 from the main text allows us to estimate non-marginal changes in adaptation costs incurred as the climate of each population changes. As discussed in the main text, we compute these non-marginal changes by integrating over time instead of over climates, as climate changes in a region for each of our 33 high-resolution climate model projections can be indexed by time. Thus, in each projection, we solve for adaptation costs as a region's climate evolves from $\boldsymbol{C}(t_1)$ to $\boldsymbol{C}(t_2)$:

$$\begin{aligned}\left[\widehat{A(\boldsymbol{b}^*(\boldsymbol{C}(t_2), Y)) - A}(\boldsymbol{b}^*(\boldsymbol{C}(t_1), Y))\right] &= -\int_{\boldsymbol{C}(t_1)}^{\boldsymbol{C}(t_2)} VSL_t \left(\frac{d\hat{\bar{f}}_t}{d\boldsymbol{C}} - \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}}\right) d\boldsymbol{C} \\ &= -\int_{t_1}^{t_2} VSL_t \left(\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C}_t, Y}\right) \frac{\partial TMEAN}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}(t)}{dt} dt \\ &\approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \left(\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C}_\tau, Y}\right) (TMEAN_\tau - TMEAN_{\tau-1})\end{aligned} \tag{A.13}$$

where the second equality results from substitution of Equation A.12 into Equation 8, as well as changing the variable of integration from $\boldsymbol{C}$ to t (indexing positions along the line integral by t). The third equality is a discretized approximation; we implement the computation of the integral using Newton's method over discrete time-steps of one year. 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. Note that Equation A.13 is identical to Equation 12 in the main text.

These adaptation cost estimates are calculated for each impact region, age group, and year, using $t_1 = 2015$ as the baseline year, for each of our 33 high-resolution climate model projections.

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.

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).[76] 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). 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).[77] 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) [78] and merged with the death counts to calculate mortality rates. Administrative boundary files

[76] http://datasus.saude.gov.br/sistemas-e-aplicativos/eventos-v/sim-sistema-de-informacoes-de-mortalidade.
[77] Data are available here: http://www.deis.cl/bases-de-datos-defunciones/.
[78] Data are available here: http://www.ine.cl/estadisticas/demograficas-y-vitales

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. 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.

B.1.4 European Union

The EU maintains a centralized statistical database known as EuroStat (Eurostat, 2013)[79] 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.[80] 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)[81] 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

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

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

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

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 Tables 2 and 3. 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 D2). However, these data are used to assess the external validity of our extrapolation methods, as discussed in Appendix D.6.

B.1.7 Japan

Japanese data on mortality and population at the prefecture-year[82] level were obtained from the National Institute of Population and Social Security Research[83] 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.[84] 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.[85] 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

[82] 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.

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

[84] 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.

[85] 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.

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)[86] 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.[87]

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. A full list of these administrative boundaries is available upon request.

[86] Partial data are freely available through the CDC Wonder database.
[87] Data are available here: https://www.census.gov/geo/maps-data/data/tiger-line.html.
[88] 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)[88]
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.[89] 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 $1° \times 1°$ 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 $0.5° \times 0.5°$ 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).[90] 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 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

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

[90] 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).

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 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 1B.

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 9 and 10. In this case, we begin with data on average temperatures for each day 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 2010 Gridded Population of the World dataset.[91] We then sum these daily polynomial terms T_{zd}^p over days in the year t. The vector

[91] Data are available here: https://sedac.ciesin.columbia.edu/data/collection/gpw-v4.

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 9 and 10, quadratic polynomials in precipitation are similarly calculated and weighted averages are taken over administrative units. In Appendix Figure D1, 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 10) 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**.[92] 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**.[93] 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

[92] Penn World Tables (PWT) database: https://www.rug.nl/ggdc/productivity/pwt/.
[93] Eurostat database: http://ec.europa.eu/eurostat/data/database.

subnational level income data for all countries outside the EU: Brazil, Chile, China, France,[94] 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 10. 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 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:

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

$$GDPpc_{sct} = \nu_{sct} \times GDPpc_{ct}^{PWT}$$

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 10 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 1A).

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.

[94] 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.

[95] EU subnational income data come from Eurostat (2013). For all other countries, subnational income data are obtained from Gennaioli et al. (2014).

Country	ISO code	Years in mortality sample	Years in income sample[95]
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.

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.[96]

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

[96] 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.

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.[97] 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 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 $\frac{GDPpc_{rc}}{\sum_{r \in c} GCPpc_{rc}}$ for impact region r in country c, which will allow us to predict income at the impact region level, given projections of national GDP per capita from the SSPs, $\sum_{r \in c} GDPpc_{rc} = GDPpc_c^{SSP}$. 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 $\frac{GDPpc_{ict}}{\sum_{i \in c} GDPpc_{ict}}$.[98] 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:

$$\frac{GDPpc_{ict}}{\sum_{i \in c} GDPpc_{ict}} = \alpha + \beta ZNL_{ict} + \epsilon_{ict} \tag{B.14}$$

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 estimated coefficients from Equation B.14 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 $\sum_{r \in c} GDPpc_{rc}$ comes from one of the SSP projected income scenarios. The result of this

[97] 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.

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

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).[99] 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. 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:

$$\widehat{pop}_{rtv} = \begin{cases} pop_{ctv}^{SSP}\left(\frac{pop_{r,2011}^{LandScan}}{\sum_{r\in c} pop_{r,2011}^{LandScan}}\right), & \text{if } r \in C \\ pop_{r,2011}^{LandScan}, & \text{if } r \notin C \end{cases} \tag{B.15}$$

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.

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

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.

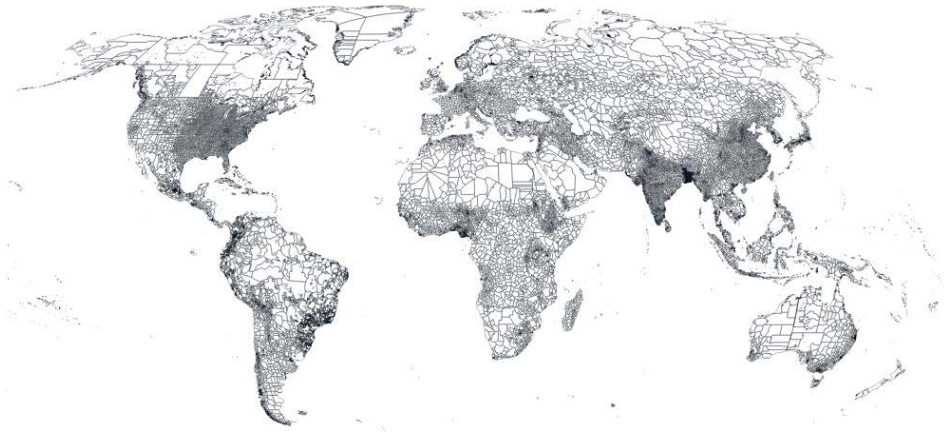

Figure C1: Map of the 24,378 "impact regions" for which location-specific projections are calculated. We use a clustering algorithm to form these regions from the full set of GADM administrative regions, such that they are roughly similar in total population, and so that they are approximately internally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

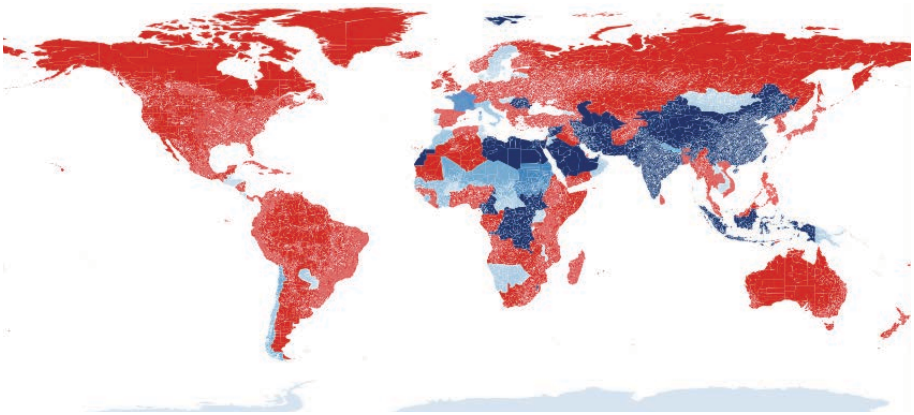

Figure C2: Categorization of countries based on the method used to construct impact regions out of GADM administrative regions. A country's target number of impact regions is $target_c$, as computed in the text. Countries in shades of blue have target values that can be approximated by one of the available 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),[100] 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.

[100] 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, and out-of-sample validation

This appendix provides additional illustrations of the main econometric regression results used and discussed throughout the main text (Figures D2, Figures D3 and D4, and Table D2), shows a set of robustness checks for those main results (Figures D1 and D5 and Table D3), and discusses an out-of-sample validation test designed to evaluate the accuracy with which our estimates predict mortality-temperature responses in locations that are not used for estimation (Figure D6).

D.1 Robustness of the pooled multi-country mortality-temperature response function

Figure D1 displays the results of estimating Equation 9 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 results for an all-age mortality response $g(\boldsymbol{T}_{it})$ in which an average treatment effect across all age categories is recovered (as in Table 2 in the main text). The four functional forms estimated are fourth-order polynomials, bins of daily average temperature, restricted cubic splines, and piecewise linear splines (details on these functional forms are in Section 4 of the main text). All regressions include $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects, and are population weighted. Robustness to alternative fixed effects specifications is shown in Table 2.

D.2 Age-specific pooled multi-country mortality-temperature response functions

Figure D2 displays the mortality-temperature responses for each of our three age categories (<5, 5-64, >64) estimated from Equation 9 and using the pooled 41-country sample. These curves correspond with our primary specification (column 2 in Table 3).

D.3 Heterogeneity in the mortality-temperature response function across countries

The administrative regions in our sample display substantial heterogeneity in incomes, climates, and demographics, among many other characteristics. To begin to examine this heterogeneity before estimating the two-factor model of heterogeneous mortality-temperature responses in Equation 10, here we investigate variation in mortality responses to temperature at the country level across our sample.[101] Table D1 displays differential mortality-temperature response functions for each of the 9 countries or regions (i.e. the EU, which is composed of 33 countries) in our data.

[101] For exposition purposes, we treat the EU here as a single "country". A dummy variable is used to estimate the EU only response, but each of the 33 countries in the EU sample have their own set of country-year-age fixed effects.

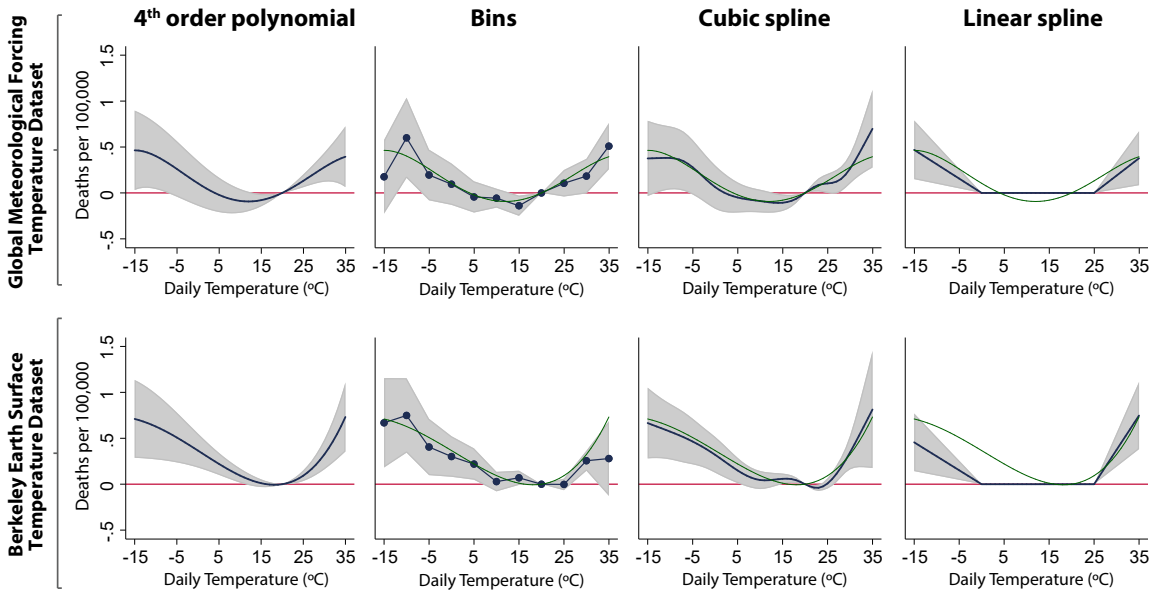

Figure D1: Robustness of the all-age mortality-temperature relationship to alternative functional forms and to different historical climate datasets. Row 1 shows the mortality-temperature response function as estimated using daily temperature and precipitation data from the Global Meteorological Forcing 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.

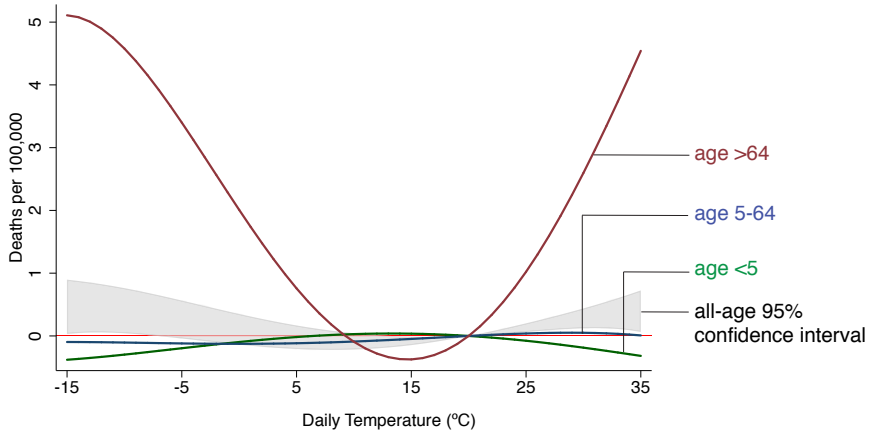

Figure D2: Mortality-temperature response function with demographic heterogeneity. Mortality-temperature response functions are estimated for populations <5 years of age (green), between 5 and 64 years of age (blue), and >64 years of age (red). Regression estimates shown are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects (Equation 9 in the main text).

Table D1: Heterogeneity by country in the mortality-temperature response function. 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. 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. Country-specific coefficients are generated by interacting all climate variables and fixed effects with country dummies. Point estimates are only shown for daily average temperatures that are actually experienced in each country over our sample period.

Temperature	BRA	CHL	CHN	FRA	JPN	MEX	USA	EUR
35°	0.684					0.212	0.547**	
	(0.452)					(0.286)	(0.244)	
30°	0.018	-1.009**	0.343	0	0.268*	-0.068	0.362***	0.863**
	(0.1)	(0.485)	(0.677)	(0.525)	(0.16)	(0.117)	(0.106)	(0.414)
25°	0.023	-0.152**	-0.191	0.308*	0.055	-0.023	0.165***	0.287**
	(0.08)	(0.077)	(0.237)	(0.172)	(0.059)	(0.068)	(0.039)	(0.144)
20°	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–
0°		-0.204	1.12***	-0.419*	0.272*	4.526***	0.017	1.426
		(0.284)	(0.421)	(0.249)	(0.154)	(0.991)	(0.119)	(1.063)
-5°		-1.058	0.817*	-0.478	0.451*	10.829***	0.161	1.968
		(0.732)	(0.456)	(0.377)	(0.24)	(2.801)	(0.146)	(1.435)
-10°		-3.34*	0.435	-1.341*	0.773*	21.986***	0.299*	2.248
		(1.736)	(0.541)	(0.782)	(0.437)	(6.384)	(0.175)	(1.575)
Adj R-squared				.989				
Observations				820237				
Adm2-Age FE				YES				
Cntry-Year-Age FE				YES				

Standard errors clustered at ADM1 level. Stacked regression is run with population weighting.
*** p<0.01, ** p<0.05, * p<0.1

D.4 Age-specific heterogeneity of the mortality-temperature response function by average income and average climate

The estimation of Equation 10 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 ($\log(GDPpc)$). Tabular results from this estimation are reported in Table D2 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).[102] 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 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 D2: 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	-0.928*	-0.102*	-0.236	-0.031*	-4.658*	-0.686**
	(0.535)	(0.053)	(0.162)	(0.018)	(2.388)	(0.332)
30°C	-0.264	-0.044	-0.017	-0.014	-0.080	-0.299**
	(0.277)	(0.027)	(0.068)	(0.009)	(0.919)	(0.141)
20°C	–	–	–	–	–	–
	–	–	–	–	–	–
0°C	-0.765	0.034	0.050	-0.029*	1.985	-0.718***
	(0.528)	(0.030)	(0.169)	(0.018)	(2.144)	(0.153)
-5°C	-0.873	0.035	0.206	-0.040**	6.019*	-0.888***
	(0.619)	(0.031)	(0.247)	(0.020)	(3.207)	(0.205)

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 2

[102] 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.

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 D2. 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 2 for the >64 age category. Here we replicate this figure for the other two age groups in our analysis.

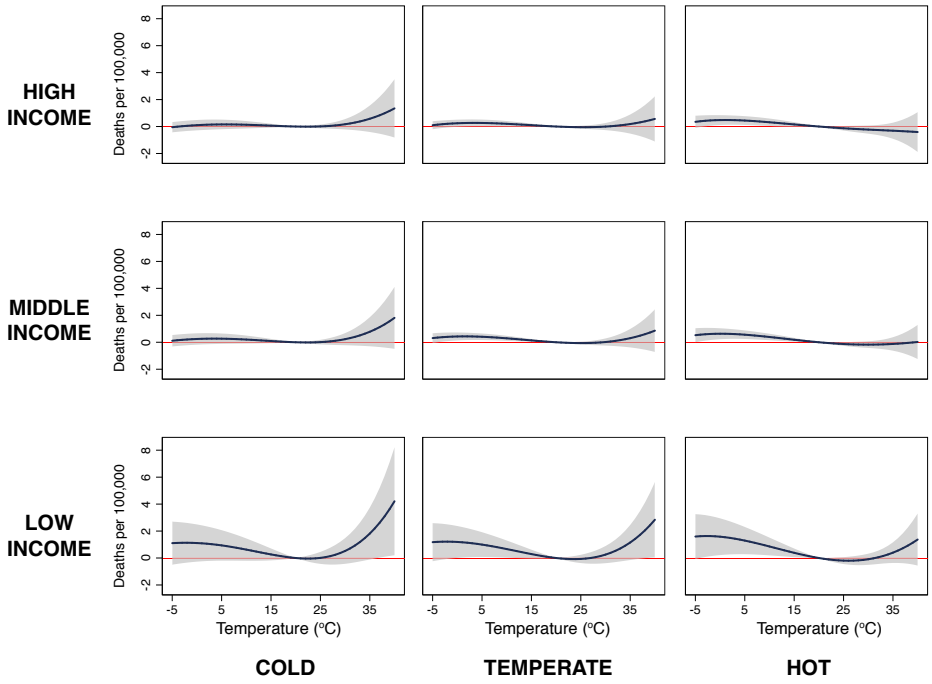

Figure D3: Heterogeneity in the mortality-temperature relationship (ages <5 mortality rate). Each panel represents a predicted response function for the ages <5 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left are 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.

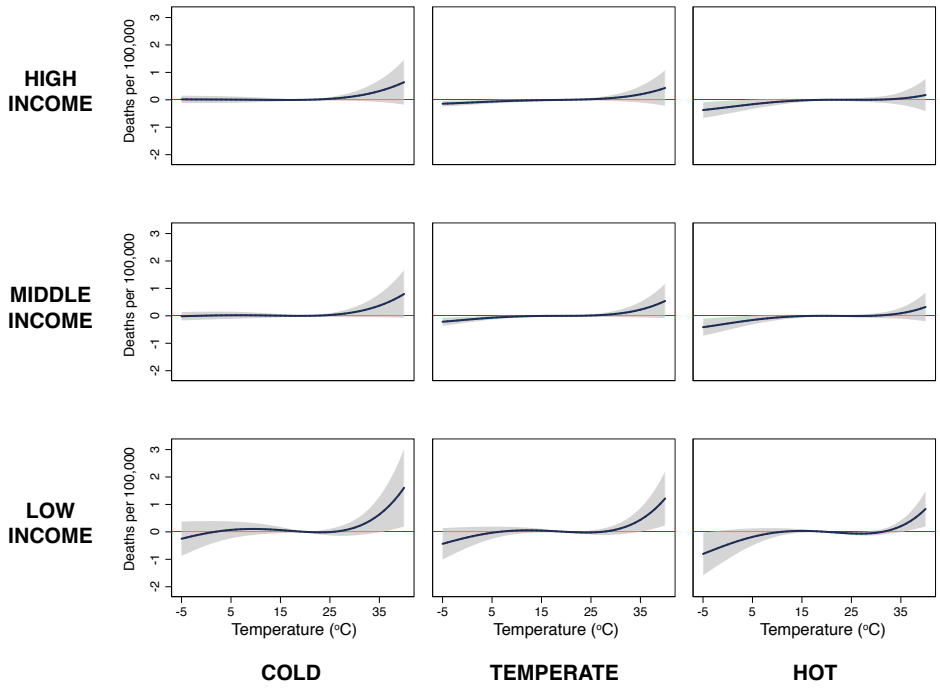

Figure D4: Heterogeneity in the mortality-temperature relationship (ages 5-64 mortality rate). Each panel represents a predicted response function for the ages 5-64 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left 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.

D.5 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 10 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 10 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 D2, 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 D2 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.

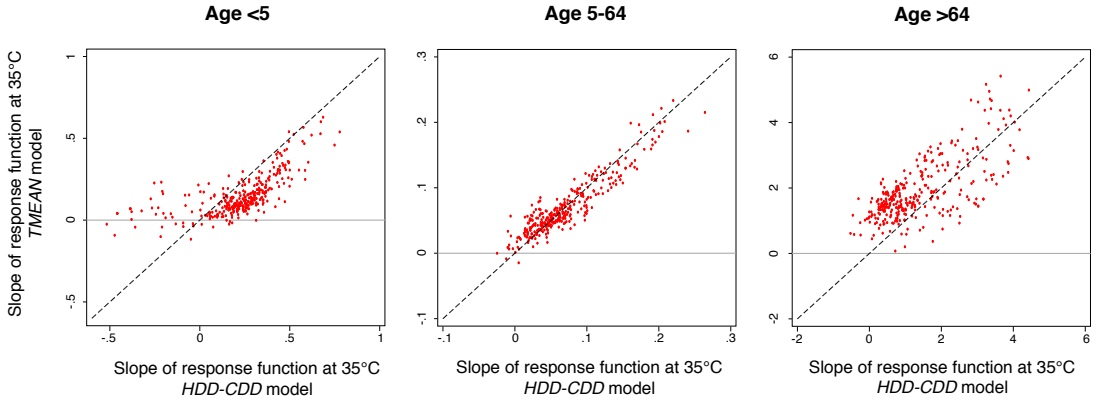

Figure D5: Predicted mortality-temperature response functions in-sample are similar under alternative characterizations of long-run average annual temperature. Each panel contains a scatter plot of the slope (i.e. derivative) of the predicted mortality-temperature response function, evaluated at 35°C, 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 10 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 10 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.6 Replication of Burgess et al. (2017) and out-of-sample model validation in India

Throughout our analysis, we use coefficients estimated from Equation 10 in the main text, in combination with local-level observations of $TMEAN$ and $\log(GDPpc)$, to generate predicted response functions in all regions of the world, including where mortality data are unavailable (see Section 6 for details). The accuracy of this extrapolation depends in part on the representativeness of the observed sample; 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. To evaluate the performance of our interaction model in this region of the global distribution, we use mortality data from India to conduct an out-of-sample validation exercise. India represents the poorest and hottest country for which we have

been able to obtain mortality records, and therefore provides an important check on our extrapolation performance.

To execute this validation test, we use data from Burgess et al. (2017), and begin by replicating the analysis conducted by those authors to confirm consistency with the existing literature. To allow for direct comparison with Burgess et al. (2017), in place of the fourth order polynomial we use as our main specification, here we estimate the mortality-temperature relationship using the Indian data with binned daily temperatures, where annual values are calculated as the number of days in region i in year t that have an average temperature within a bin range k. Bin edges in degrees Celsius are given by the following set: $K = \{-\infty, -15, -10, -5, 0, 5, 10, 15, 20, 25, 30, 35, +\infty\}$. We cluster our standard errors at the ADM1 level (in India, this is equivalent to the state level), as in all our specifications throughout the main text. Burgess et al. (2017) estimate a similar binned regression model using 2°C bins, with clustering at the ADM2 (i.e. district) level.

After constructing an all-age mortality-temperature response function for India alone, we compare this result to the predicted response functions generated from our main analysis. These impact region-level response functions are generated using estimation of Equation 10, which relies on mortality, climate, and income data from the sample of countries shown in Table 1, but excludes India.

In Figure D6, we show the result of this replication and out-of-sample validation exercise for India. Figure D6 compares our predicted responses for all impact regions in India (in grey) to the mortality-temperature response estimated using India's data alone (in red). We generate the all-age average response function for each impact region from our age-specific interaction model by taking a population-weighted average of the responses predicted for each age category, using age-specific population values for the year 2015. Our model performs remarkably 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. Moreover, our results are very similar to the findings in Burgess et al. (2017), with an approximately linear increase in deaths for temperatures above 20°C.

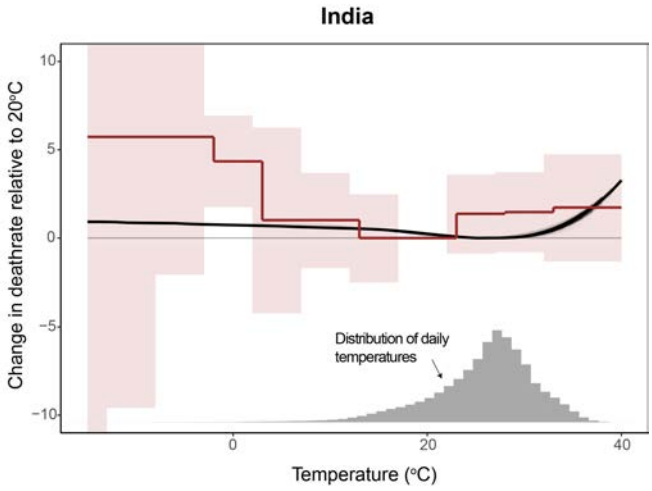

Figure D6: Out-of-sample validation of the mortality-temperature response function in India. Grey lines indicate predicted response functions for each impact region in India, predicted using coefficients from the interaction model in Equation 10 and using an estimation sample without Indian data. The solid black line is the unweighted average across all regions, while the red line is the estimated response function using only all-age mortality data from India. The red line is estimated using a nonparametric binned regression, as described in the text, to enable comparison with Burgess et al. (2017). The relative congruence between red and black 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 we have the least representation.

E Implementation details for 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, and 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.

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 10 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 7), we apply the estimated coefficients from Equation 10 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 6.2 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:

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

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:

$$M_{ait} - M_{ai,t-1} = \boldsymbol{\beta}_t \left[\boldsymbol{T}_{it} - \boldsymbol{T}_{i,t-1}\right] + q_a(\boldsymbol{R}_{it}) + \varepsilon_{it} \tag{E.16}$$

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 9). 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 10, 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 6, 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. Gas-

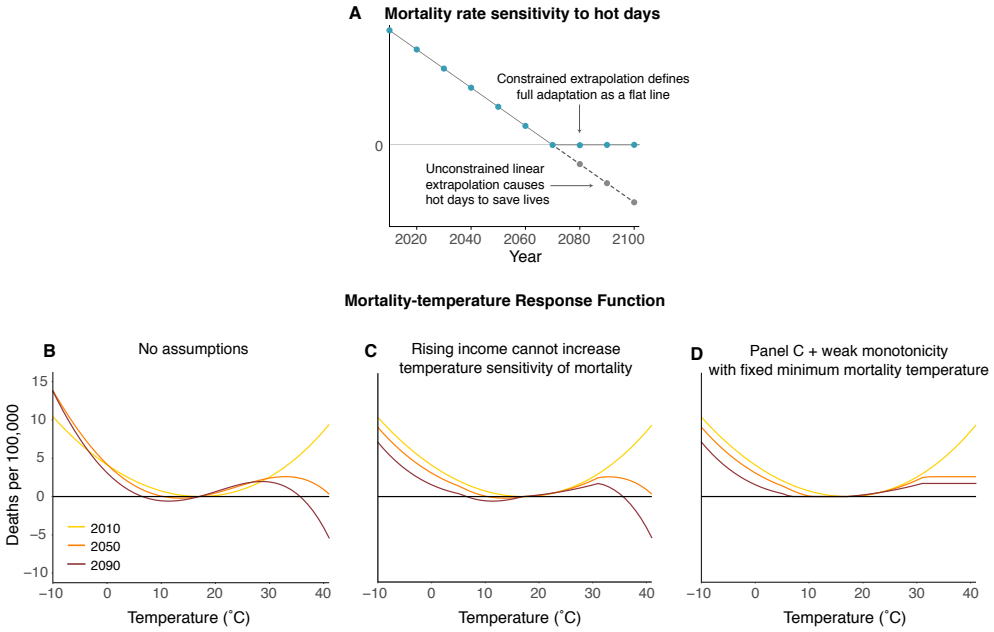

Figure E1: Two assumptions imposed in climate projections ensure that full adaptation is defined as a flat-line response function and that responses conform to basic physical and economic constraints. Panel A demonstrates heuristically the importance of imposing assumptions on the shape of response functions 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).

parrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011). As shown in Section 5, we also recover U-shaped relationships 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,[103] it is important to note that we allow these minimum mortality temperatures 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 10. This assumption is important because Equation 10 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.[104]

In imposing the weak monotonicity constraint, we fix the minimum mortality temperature at its baseline level in 2015 for each impact region. We do so because the use of spatial and temporal fixed effects in Equation 10 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.[105]

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 10, 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 10 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.

[103] E.g. Guo et al. (2014) demonstrate that mortality risk is smallest around the 75^{th} percentile of local temperatures in 12 different countries.

[104] See Appendix F.3 for results in which we explore a scenario with slower rates of adaptation. Under this alternative scenario, Assumption #1 binds much less frequently.

[105] 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.

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 7), as well as a robustness check and sensitivity analysis regarding the functional form of the mortality-temperature relationship and assumptions regarding the rate of adaptation, respectively.

F.1 Additional climate change projection results

Alternative measures of climate change impacts In Figure 6 of the main text, we show a map of impact region-level mean estimates of the full mortality risk of climate change, accounting both for adaptation and income benefits as well as adaptation costs. However, in Section 7 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; (iii) mortality effects of climate change with benefits of income growth and adaptation. Panels A-C in Figure F1 below show projected impacts for each of these three alternative measures; for comparison, panel D repeats the full mortality risk of climate change map from the main text.

Climate change projections by age group In the main text, Figure 7 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 F2, 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 8 and Appendix G).

Climate change projections by socioeconomic scenario Throughout Section 7 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, 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).

Correlation of uncertainty in climate change projections with present day income As shown in Figure 8 in the main text, there is a strong correlation between present day income and the composition of future damages between deaths and adaptation costs. In general, low income locations tend to suffer relatively large increases in mortality rates by end of century, while high income locations incur relatively large adaptation costs. In Figure F4 below, we examine whether low income locations also face higher *uncertainty* in projected impacts of climate change. We do so by demeaning each Monte Carlo simulation by an impact region-specific mean, and showing the spread in demeaned impacts for each decile of today's income distribution. As in Figure 8, panel A in Figure

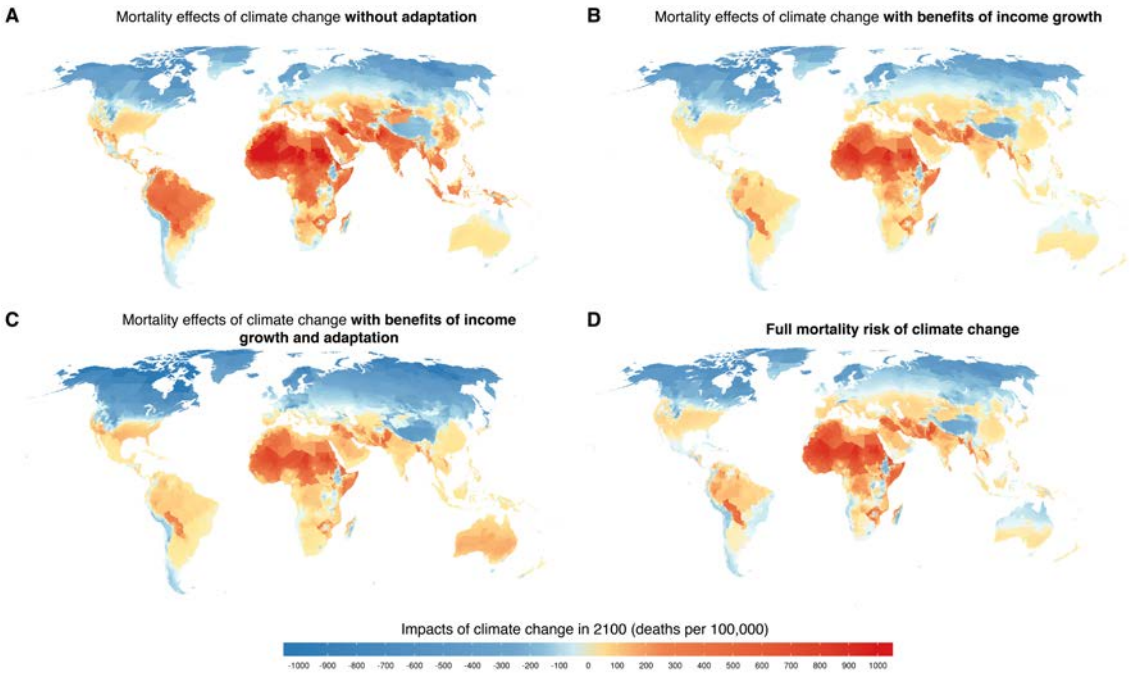

Figure F1: Mortality costs of climate change under alternative adaptation scenarios. All maps show predicted mortality effects of climate change and colors in each impact region represent the mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. 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.

F4 shows projected deaths, panel B shows adaptation costs, and panel C shows the full mortality risk of climate change (the sum of deaths and adaptation costs measured in death equivalents); each panel shows impacts plotted against deciles of today's income distribution. There is some evidence that lower income regions face higher uncertainty in the magnitude of their projected full mortality risk of climate change, arising from higher uncertainty in impacts on death rates.

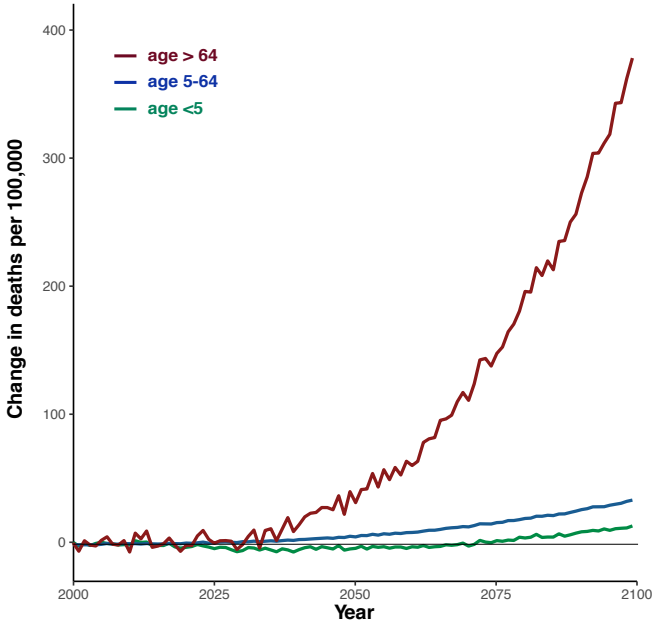

Figure F2: Heterogeneity in climate change impacts on mortality by age group. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. 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 7 in the main text represents the sum across these age-specific projections, where death rates are population weighted by age-specific population values.

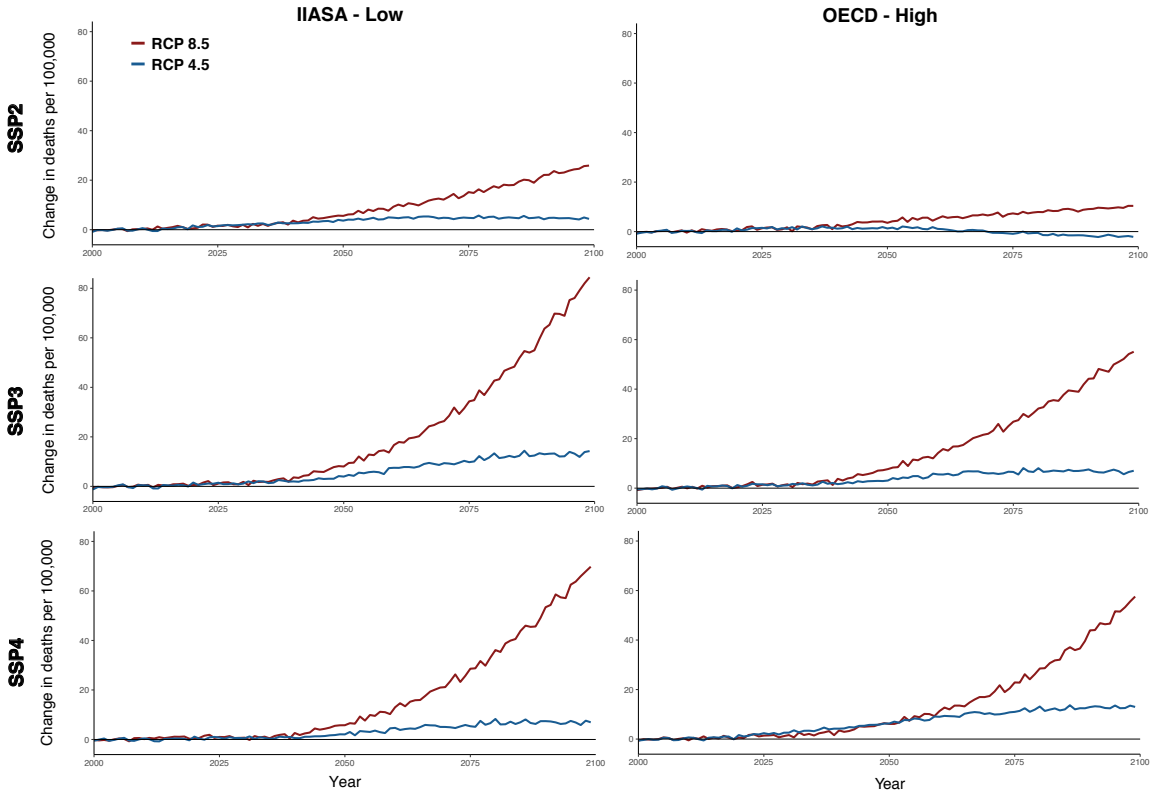

Figure F3: The full mortality risk of climate change under different scenarios of population growth, economic growth, and emissions. Rows denote different Shared Socioeconomic Pathway (SSP) scenarios, columns denote two separate modeling groups that produce data for each SSP, and each panel shows a time 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 lower 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.

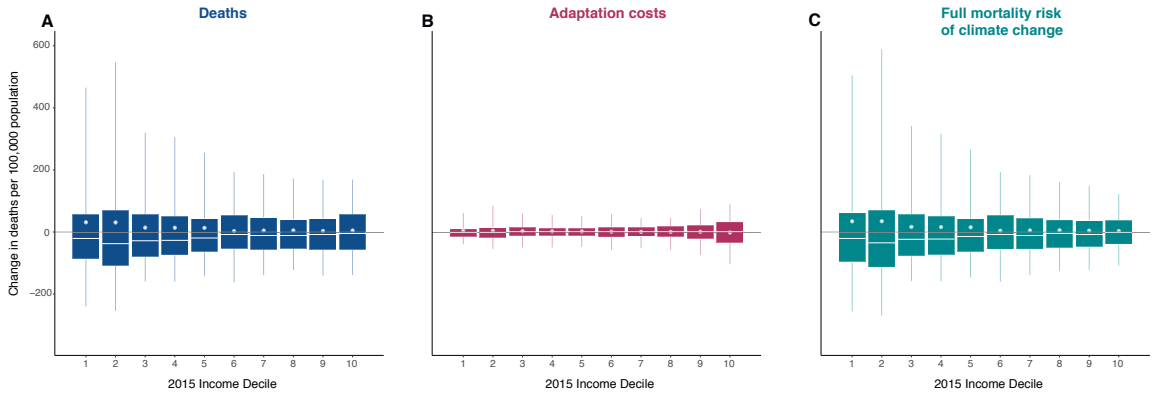

Figure F4: Uncertainty in climate change impacts and adaptation costs across present-day income groups. All three panels show impacts by present-day income decile for impact region estimates that have been demeaned across the full set of Monte Carlo simulations in order to demonstrate how uncertainty in future impacts may be correlated with contemporary incomes. This figure is analogous to Figure 8 in the main text, which shows distributions of impact region mean estimates. Panel A shows the range of annual mortality rates due to climate change in 2100, accounting for the benefits of adaptation, against deciles of 2015 income. Panel B shows the range of annual adaptation costs incurred due to climate change in 2100, measured in death equivalents. Panel C shows the range of full mortality risk due to climate change estimates, which are the sum of deaths and adaptation costs measured in death equivalents. Income deciles are calculated as in Figure 8. All box plots show moments of the distribution of impact region-specific demeaned impacts within an income decile. Solid vertical lines in each box plot extend to the 5^{th} and 95^{th} percentiles of this distribution, boxes indicate the interquartile range, white horizontal lines indicate the median, and white circles indicate the mean. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

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

As discussed in Section 4, we experiment with four distinct nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ in Equations 9 and 10 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 10 (see Section 4.2). 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 D1, we show that the binned and fourth order polynomial functional forms recover remarkably similar mortality-temperature response functions across our pooled multi-country sample. Below in Figure F5, we show that this similarity carries through to generate very similar climate change impact projections across the binned and polynomial functional forms. Both projections are constructed using estimation of the interaction model in Equation 10 in combination with high-resolution covariates $TMEAN$ and $\log(GDPpc)$ to generate impact region-specific response functions (see Section 6 for details).

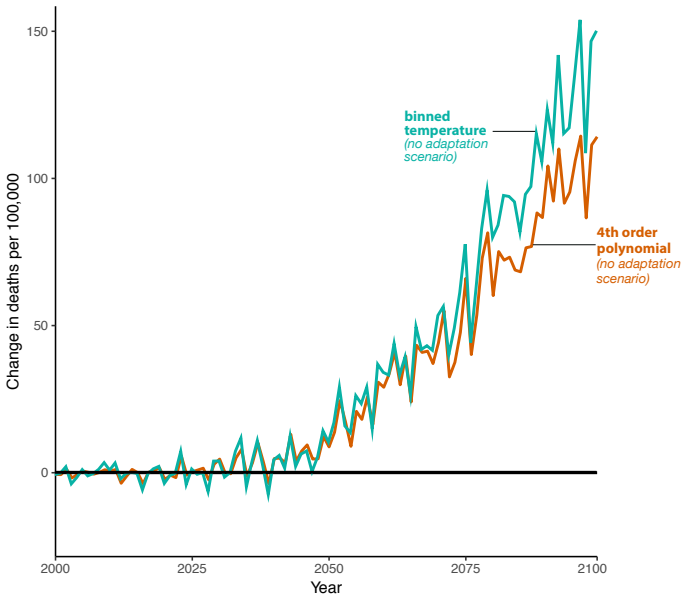

Figure F5: Robustness of impact projections to alternate functional forms of temperature. Each line represents the time series of changes to the mortality rate due to climate change under the socioeconomic scenario SSP3 and the emissions scenario RCP 8.5. Results shown are for a single climate model (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 10. 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 the rate of adaptation

In our main results, we use the estimated coefficients from Equation 10 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 10 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 10 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 a sensitivity analysis where the speed at which the mortality-temperature response function changes with time-varying covariates is deterministically reduced by half. This exercise is used to understand how climate change impact projections change if the evolution of the response function towards zero (see Figure E1) occurs more slowly.

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, 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 $TMEAN_{rt} = TMEAN_{r,t-1} + \frac{\Delta TMEAN_{rt}}{2}$. Note that both the primary specification and reduced growth analyses generate identical covariates (and hence, response functions) in 2015.

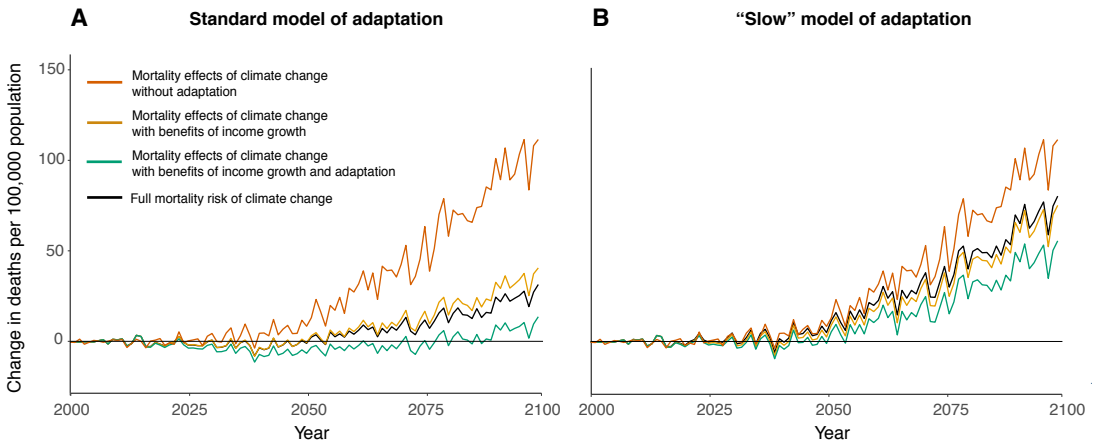

Figure F6: Impacts of climate change on mortality are qualitatively similar with a model of slower adaptation rates. Time series of projected mortality costs of climate change (black line), as compared to partial estimates from incomplete accounting of the costs and benefits of adaptation (other colors). 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 6. Panel B shows results for an alternative model of adaptation in which the rate of adaptation to both income growth and to warming climate is cut in half.

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",[106] in combination with our empirically-derived damage functions, to construct mortality partial SCC estimates. We detail this implementation below.

G.1 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:

$$\frac{dT_i}{dt} = \frac{q_i F - T_i}{d_i}, i \in \{0, 1\}, \tag{G.17}$$

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,

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

representing different uptake timescales:

$$\frac{dR_j}{dt} = a_j E - \frac{R_j}{\alpha_j \tau_j}, j \in \{0, 1, 2, 3\} \tag{G.18}$$

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.1.1, 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 10; as described below and in Section 8, this difference is used to construct mortality partial SCC estimates.

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

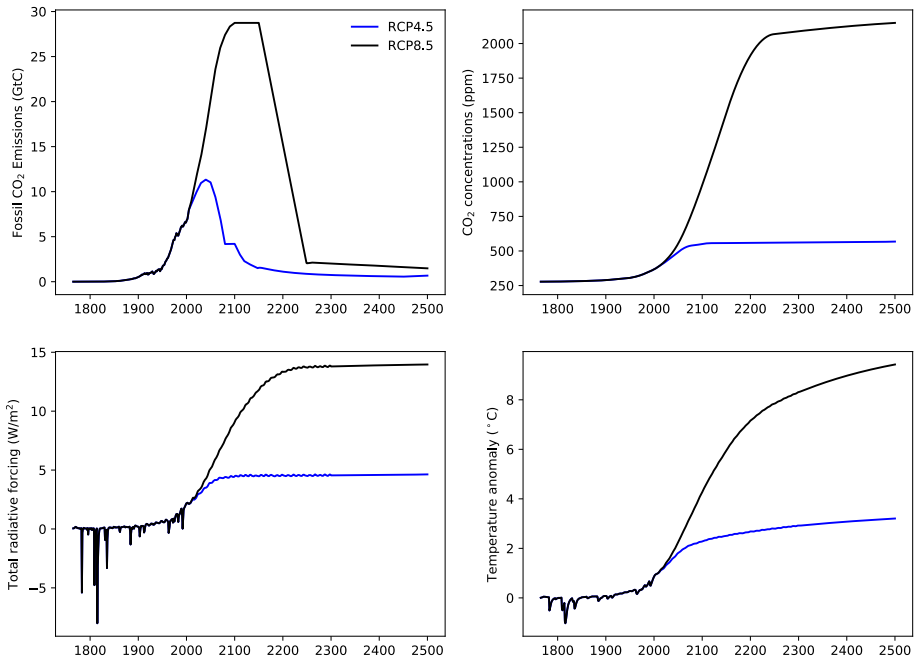

Figure G1: Behavior of key variables in the FAIR simple climate model under median climate parameters. Each panel shows the temporal trajectory of key variables in FAIR that are used in our calculation of the social cost of carbon. The trajectories shown arise under FAIR run with median climate parameter 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.1.1 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.

***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.

***τ_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 4 in the main text.

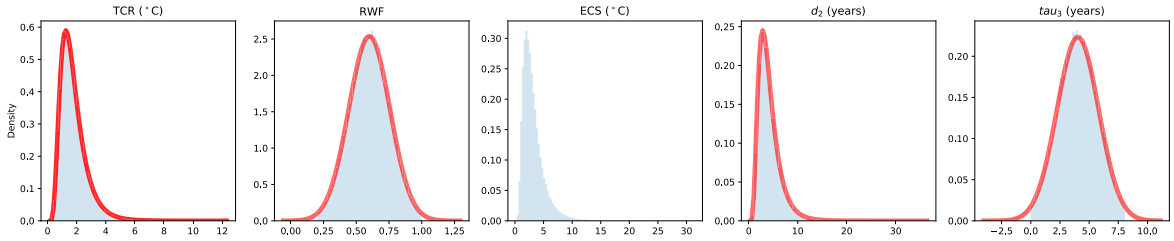

Figure G2: Distributions of key FAIR parameters for climate sensitivity uncertainty both before (red curve) and after (blue shading) applying constraints. Each panel indicates the distribution of a key parameter in the FAIR simple climate model, both before (in red) and after (in blue) the imposition 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).

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).

Finally, we assess the reasonableness of the "handoff" between the SMME models, on which the damage function is estimated (see Section 8.2), 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). 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 the handoff is reasonable (as would be expected based on the construction of the SMME).

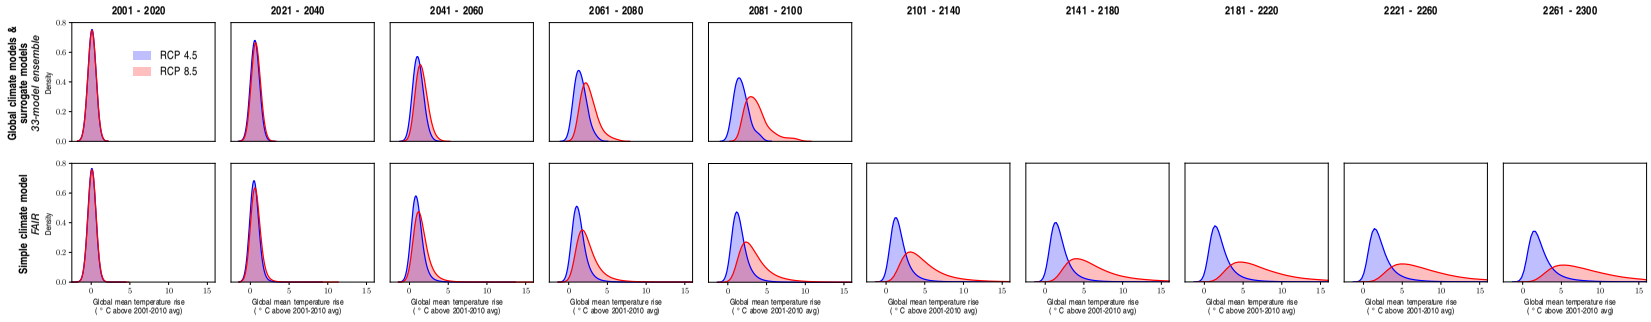

Figure G3: Distribution of changes in global mean surface temperature (ΔGMST) from an ensemble of global climate models and surrogate models (SMME) and from the simple climate model FAIR. Top row: Distribution of ΔGMST from 2001 to 2100, according to an ensemble of 33 GCMs and surrogate models that form 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.2 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 8. 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 Section 8 for a discussion of valuation of mortality-related costs of climate change). This figure contains the same information as Figure 9 in the main text, while additionally demonstrating substantial heterogeneity across distinct valuation scenarios (our primary valuation method uses a value of life-years adjustment with a globally-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) .

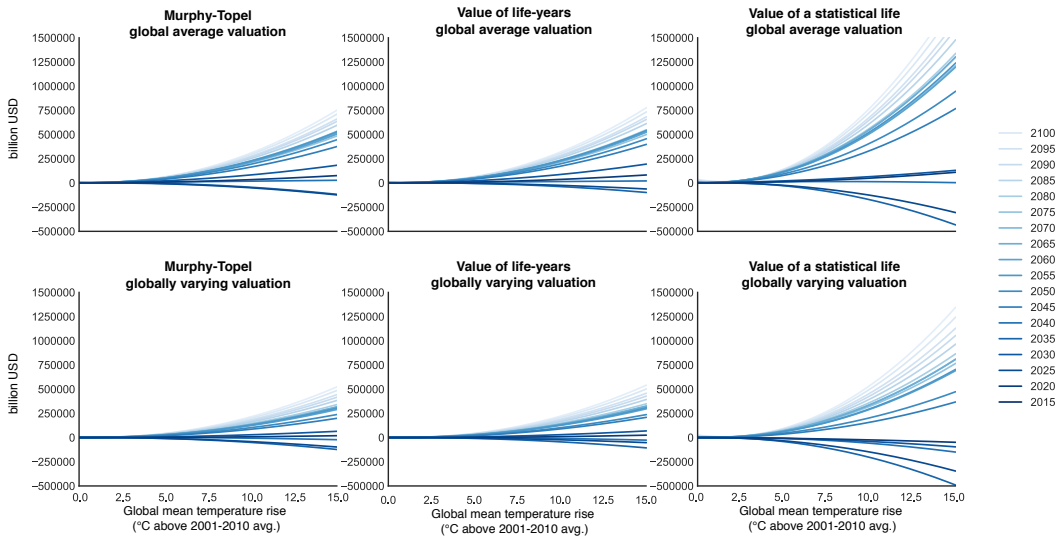

Figure G4: Temporal evolution of empirically derived damage functions (billion USD) as a function of global mean surface temperature anomaly. Each panel shows estimates of quadratic damage functions estimated independently for each 5-year period from 2015 to 2100 under various valuation assumptions regarding the valuation of lives lost or saved.

The coefficients on these quadratic damage functions are constructed for each year from 2015 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 2.5%, 3% and 5% discount rates, and summed through time to create a net present value, following Equation 14 in Section 8. This final value is the net present value of the full mortality risks caused of a marginal emission of CO_2. A more robust 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.

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 each of the following quantiles: 5, 25, 50, 75, and 95. As in the mean damage function estimation, extrapolation past the year 2100 is accomplished using a time interaction model (see Section 8). 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 SCCs associated with 25^{th} and 75^{th} percentile damage functions are reported as the damage function uncertainty IQR of the SCC.

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 8, in our primary specification we report mortality partial Social Cost of Carbon (SCC) calculations which value deaths according to the total value of expected life-years lost. We do so 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 table 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. Note that this procedure assumes that our estimated climate change driven deaths occur with uniform probability for all people within an age category.

The above methodology 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 and a valuation method that values deaths according to the total value of expected life-years lost (see

[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.

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

Appendix H.1). This appendix shows a range of mortality partial SCC estimates under alternative VSL values, without a life-years adjustment (so that all lives lost or saved are valued equally, regardless of age), and with a life-years adjustment that allows for age-specific values of remaining life, as derived by Murphy and Topel (2006). Table H2 provides mortality partial SCC estimates across these distinct valuation approaches under the "globally varying valuation" method shown in Panel A of the main text Table 4. For all estimates shown in Table H2, an income elasticity of one is used to adjust for cross sectional variation in incomes and for global income growth. Table H3 provides mortality partial SCC estimates across distinct valuation approaches under the "globally uniform valuation" method shown in Panel B of the main text Table 4. For all estimates shown in Table H3, a globally homogeneous VSL is used in each year that evolves over time based on global income growth.

Finally, Table H4 shows mortality partial SCC estimates under various socioeconomic projections (SSP3 is used throughout the main text).

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 adjust death valuation by expected life-years lost. Murphy-Topel life years adjusted estimates 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	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally varying valuation of mortality risk (2019 USD)					
Value of life years						
RCP 4.5	5.0	4.4	2.2	2.3	2.1	1.0
Climate sensitivity uncertainty	[-6.5, 45.4]	[-3.1, 29.8]	[0.1, 8.5]	[-3.0, 21.1]	[-1.4, 13.9]	[0.0, 3.9]
Damage function uncertainty	[-40.5, 52.2]	[-29.1, 38.4]	[-12.9, 17.0]	[-18.8, 24.3]	[-13.6, 17.9]	[-6.0, 7.9]
Full uncertainty	[-35.6, 58.0]	[-26.9, 40.1]	[-12.8, 15.2]	[-16.6, 27.0]	[-12.5, 18.7]	[-6.0, 7.1]
RCP 8.5	62.7	38.1	8.6	29.2	17.7	4.0
Climate sensitivity uncertainty	[25.4, 153.9]	[15.4, 94.0]	[3.5, 21.1]	[11.8, 71.7]	[7.1, 43.8]	[1.6, 9.8]
Damage function uncertainty	[0.6, 103.7]	[-7.5, 71.6]	[-11.4, 26.3]	[0.3, 48.3]	[-3.5, 33.4]	[-5.3, 12.3]
Full uncertainty	[1.0, 139.6]	[-7.6, 87.4]	[-11.2, 25.4]	[0.5, 65.0]	[-3.5, 40.7]	[-5.2, 11.8]
Value of statistical life						
RCP 4.5	-4.0	2.7	6.1	-1.9	1.3	2.8
Climate sensitivity uncertainty	[-26.4, 85.7]	[-12.3, 57.6]	[1.7, 17.8]	[-12.3, 39.9]	[-5.7, 26.8]	[0.8, 8.3]
Damage function uncertainty	[-35.7, 41.2]	[-30.1, 42.0]	[-16.9, 29.4]	[-16.6, 19.2]	[-14.0, 19.6]	[-7.9, 13.7]
Full uncertainty	[-28.8, 82.7]	[-21.3, 60.4]	[-13.2, 28.4]	[-13.4, 38.5]	[-9.9, 28.1]	[-6.2, 13.2]
RCP 8.5	148.8	90.0	20.7	69.3	41.9	9.7
Climate sensitivity uncertainty	[58.2, 371.4]	[35.8, 223.1]	[9.6, 46.6]	[27.1, 173.0]	[16.7, 103.9]	[4.5, 21.7]
Damage function uncertainty	[47.2, 218.8]	[21.5, 141.7]	[-6.4, 45.6]	[22.0, 101.9]	[10.0, 66.0]	[-3.0, 21.2]
Full uncertainty	[36.8, 336.7]	[20.3, 206.3]	[-5.5, 49.3]	[17.1, 156.8]	[9.5, 96.1]	[-2.5, 22.9]
Murphy-Topel life years adjusted						
RCP 4.5	8.8	7.6	4.2	4.1	3.5	2.0
Climate sensitivity uncertainty	[-3.3, 48.9]	[-0.5, 33.0]	[1.6, 10.7]	[-1.6, 22.8]	[-0.2, 15.3]	[0.8, 5.0]
Damage function uncertainty	[-35.4, 53.9]	[-25.9, 40.5]	[-11.7, 18.7]	[-16.5, 25.1]	[-12.1, 18.8]	[-5.4, 8.7]
Full uncertainty	[-32.3, 61.1]	[-24.8, 43.3]	[-11.5, 17.2]	[-15.0, 28.5]	[-11.5, 20.2]	[-5.4, 8.0]
RCP 8.5	63.7	39.6	10.2	29.7	18.4	4.7
Climate sensitivity uncertainty	[27.0, 152.4]	[17.0, 94.1]	[4.8, 22.6]	[12.6, 71.0]	[7.9, 43.8]	[2.2, 10.5]
Damage function uncertainty	[-5.5, 102.4]	[-9.6, 72.1]	[-10.4, 27.9]	[-2.6, 47.7]	[-4.5, 33.6]	[-4.9, 13.0]
Full uncertainty	[-4.0, 132.5]	[-8.5, 85.4]	[-10.2, 26.6]	[-1.9, 61.7]	[-4.0, 39.8]	[-4.8, 12.4]

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 adjust death valuation by expected life-years lost. Murphy-Topel life years adjusted estimates 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	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally uniform valuation of mortality risk (2019 USD)					
Value of life years						
RCP 4.5	18.9	16.7	9.8	8.8	7.8	4.6
Climate sensitivity uncertainty	[-1.2, 82.4]	[2.5, 58.8]	[4.4, 22.9]	[-0.6, 38.4]	[1.2, 27.4]	[2.0, 10.7]
Damage function uncertainty	[-39.7, 98.3]	[-28.6, 73.5]	[-12.6, 34.4]	[-18.5, 45.8]	[-13.3, 34.2]	[-5.9, 16.0]
Full uncertainty	[-34.4, 104.8]	[-25.9, 76.0]	[-13.1, 33.2]	[-16.0, 48.8]	[-12.0, 35.4]	[-6.1, 15.5]
RCP 8.5	98.9	64.5	19.8	46.1	30.0	9.2
Climate sensitivity uncertainty	[42.7, 233.8]	[28.4, 150.6]	[9.7, 42.9]	[19.9, 108.9]	[13.2, 70.1]	[4.5, 20.0]
Damage function uncertainty	[9.3, 156.1]	[-2.1, 112.5]	[-10.2, 45.9]	[4.3, 72.7]	[-1.0, 52.4]	[-4.7, 21.4]
Full uncertainty	[6.6, 210.6]	[-3.6, 140.1]	[-10.4, 48.3]	[3.1, 98.1]	[-1.7, 65.3]	[-4.9, 22.5]
Value of statistical life						
RCP 4.5	18.8	24.5	21.1	8.8	11.4	9.8
Climate sensitivity uncertainty	[-22.3, 161.8]	[-4.8, 117.8]	[9.7, 48.4]	[-10.4, 75.3]	[-2.2, 54.9]	[4.5, 22.6]
Damage function uncertainty	[-22.7, 98.8]	[-15.5, 87.6]	[-5.9, 53.8]	[-10.6, 46.0]	[-7.2, 40.8]	[-2.7, 25.0]
Full uncertainty	[-20.4, 163.3]	[-11.4, 122.6]	[-5.2, 58.0]	[-9.5, 76.0]	[-5.3, 57.1]	[-2.4, 27.0]
RCP 8.5	231.2	149.0	44.9	107.7	69.4	20.9
Climate sensitivity uncertainty	[95.5, 560.9]	[63.4, 355.1]	[22.3, 95.4]	[44.5, 261.2]	[29.5, 165.4]	[10.4, 44.4]
Damage function uncertainty	[105.2, 333.4]	[61.3, 222.7]	[8.7, 78.4]	[49.0, 155.3]	[28.5, 103.7]	[4.0, 36.5]
Full uncertainty	[68.1, 527.3]	[43.2, 338.6]	[8.3, 96.6]	[31.7, 245.6]	[20.1, 157.7]	[3.9, 45.0]
Murphy-Topel life years adjusted						
RCP 4.5	22.1	19.6	11.8	10.3	9.1	5.5
Climate sensitivity uncertainty	[1.8, 84.4]	[5.0, 61.1]	[6.0, 25.0]	[0.8, 39.3]	[2.3, 28.4]	[2.8, 11.7]
Damage function uncertainty	[-35.3, 94.2]	[-25.9, 71.7]	[-11.6, 34.8]	[-16.4, 43.9]	[-12.0, 33.4]	[-5.4, 16.2]
Full uncertainty	[-30.5, 101.6]	[-23.2, 74.9]	[-11.9, 33.9]	[-14.2, 47.3]	[-10.8, 34.9]	[-5.5, 15.8]
RCP 8.5	99.0	65.4	21.2	46.1	30.5	9.9
Climate sensitivity uncertainty	[43.9, 230.0]	[29.8, 149.2]	[10.9, 43.8]	[20.5, 107.1]	[13.9, 69.5]	[5.1, 20.4]
Damage function uncertainty	[7.9, 160.6]	[-1.9, 114.7]	[-8.9, 46.5]	[3.7, 74.8]	[-0.9, 53.4]	[-4.2, 21.7]
Full uncertainty	[5.3, 209.5]	[-3.3, 140.2]	[-8.9, 48.5]	[2.5, 97.6]	[-1.5, 65.3]	[-4.1, 22.6]

Table H4: 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, all regions have heterogeneous valuation, based on local income. In Panel B, all regions globally are given the global median VSL, after scaling using income. All estimates use a value of life years adjustment, valuing deaths by the expected number of life-years lost. 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		
	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Globally varying valuation (2019 U.S. Dollars)			
RCP 4.5			
SSP2	28.5	18.9	5.7
SSP3	5.0	4.4	2.2
SSP4	-19.3	-11.7	-2.9
RCP 8.5			
SSP2	51.4	31.7	6.7
SSP3	62.7	38.1	8.6
SSP4	42.0	27.2	7.5
Panel B: Global average valuation (2019 U.S. Dollars)			
RCP 4.5			
SSP2	44.8	32.6	13.6
SSP3	18.9	16.7	9.8
SSP4	-1.5	3.2	5.5
RCP 8.5			
SSP2	77.3	51.2	15.6
SSP3	98.9	64.5	19.8
SSP4	75.2	53.2	19.9

H.3 Alternative approach to estimating post-2100 damages

As discussed in Section 8, 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 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. Table H5 shows that this alternative approach to post-2100 damage estimation causes our central estimate of the SCC to fall by 19%.

Table H5: 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, 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. For the first column, damage functions continue to evolve over time in the years after 2100, according to the method described in Section 8. 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	\$16.9	\$16.9
Post-2100 damages	\$21.2	\$13.9
Total damages	\$38.1	\$30.7

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 8 for details). In Table H6, 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 H6: 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). Under globally varying valuation, all regions have heterogeneous valuation, based on local income. In global average valuation, all regions globally are given the global median VSL, after scaling using income. All estimates use a value of life years adjustment, valuing deaths by the expected number of life-years lost. 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		
	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Globally varying valuation (2019 USD)			
RCP 4.5	1.4	1.3	0.7
RCP 8.5	65.7	39.6	8.9
Global average valuation (2019 USD)			
RCP 4.5	24.0	18.3	8.8
RCP 8.5	96.9	63.5	19.7

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, 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.

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* .

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. *Chapter 12: Long-term Climate Change: Projections, Commitments and Irreversibility*. Intergovernmental Panel on Climate Change. 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.

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

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.

Global Administrative Areas. 2012. "GADM database of Global Administrative Areas, version 2.0." URL www.gadm.org.Accessed25December,2016.

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.

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

Hsiang, Solomon and Robert E Kopp. 2018. "An Economist's Guide to Climate Change Science." *Journal of Economic Perspectives* 32 (4):3–32.

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.

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.

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/.

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.

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.

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* (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* .

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.