Online Appendix for:

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

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

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

Appendix Table of Contents

A Using revealed preference to estimate adaptation costs

A.1 Graphical solution to inferring unobserved adaptation cost

In Sections 6.1 and A.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 A.1 depicts the adaptation problem 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 A.1, 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[1] 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 A.1, the red surface illustrates the expected benefits an individual would accrue for inhabiting some climate C and selecting adaptation b. The height of this surface is a total WTP for adaptation, conditional on the climate: it is equal to the VSL times the expected survival probability $1-\tilde{f}(b, C)$ at each position (b, C). For notational simplicity, we refer to this WTP surface as V. At low levels of adaptation, V declines rapidly with higher temperature C because survival probability declines quickly. At higher levels of adaptation, V declines more gradually with C because adaptation protects individuals against temperature. The solid black lines follow this WTP surface at fixed temperatures, showing how an individual in a given climate would benefit from additional adaptation (bid curves).

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

[1] In Appendix A.5 below, A are net costs since they are net any utility benefits or costs of $\boldsymbol{b}$.

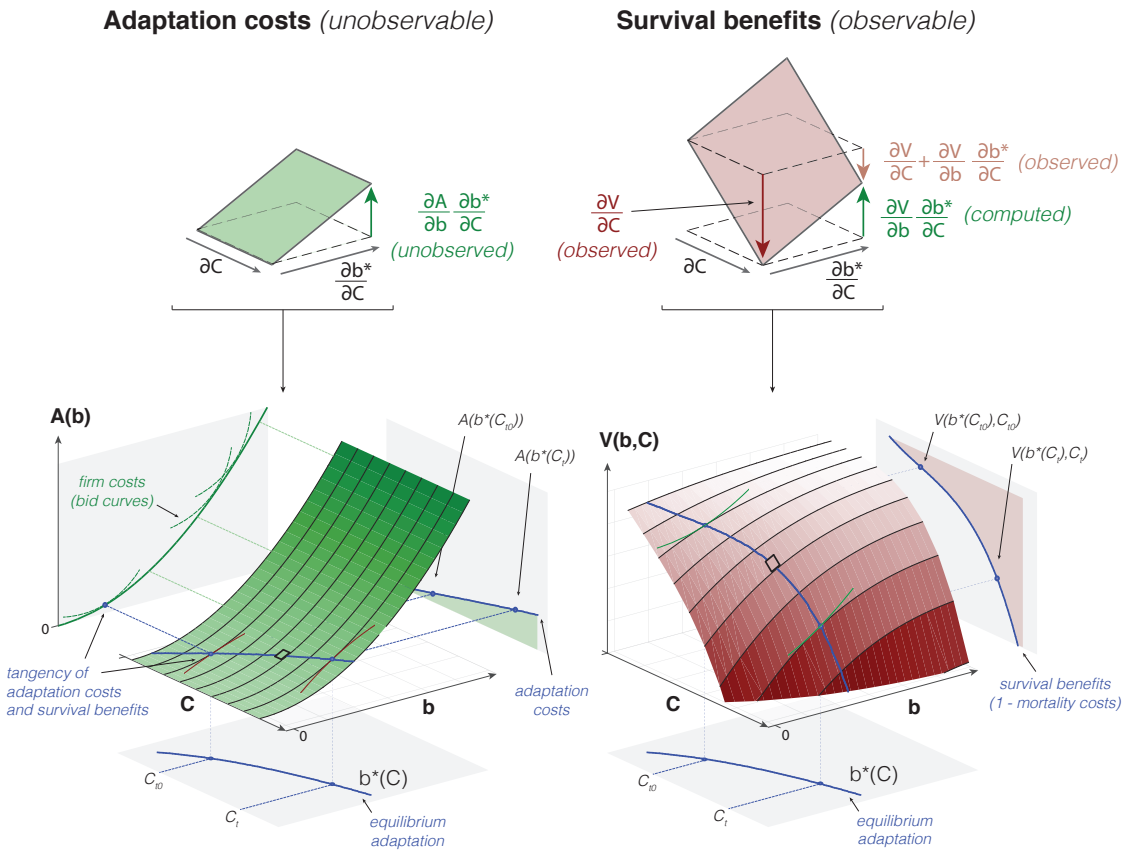

Figure A.1: Use of revealed preference to recover WTP for an unobservable adaptation. Horizontal dimensions are climate C, representing temperature, and adaptation level b. Vertical dimensions are adaptation costs $A(b)$ in the left panel and expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b,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

mortality costs rise with temperature. The sum of changes to these adaptation expenditures and the value of mortality costs is the full cost of changes to the climate.

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

To show this, at the top of Figure A.1 we draw tangency planes for both the costs and benefits surfaces

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

A.2 Mathematical details

The framework in main text Section 6.1 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 Equation 6, which denotes the total adaptation costs incurred as a population experiences a change in climate from time period t_0 to time period t.

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

Consider a single representative agent who derives utility in each time period t from consumption of a numeraire good x_t. This agent faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$, which depends both on the weather and on adaptive behaviors and investments captured by the composite good $\boldsymbol{b}_t$. As discussed in Section 6.1, changes in the climate $\boldsymbol{C}$ influence mortality risk through altering weather realizations $\boldsymbol{c}$ and through changing beliefs about the weather, hence changing adaptive behaviors $\boldsymbol{b}$.

In bringing this framework to our empirical analysis, we allow for 24,378 representative agents, one for each of the impact regions that together span the globe. We see this as a substantial improvement upon the existing estimates of global climate change damages that inform the SCC, even though there is heterogeneity

in preferences, climate, and income within these regions. For example, the DICE IAM assumes a single homogeneous global region (Nordhaus, 1992), the RICE IAM assumes 10 homogeneous regions (Nordhaus and Yang, 1996), the FUND IAM assumes 16 homogeneous regions (Tol, 1997), and the empirically-derived SCC estimates in Ricke et al. (2018) are country-level.

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

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

where $A(\boldsymbol{b}_t)$ represents expenditures for all adaptive investments, and Y is an income we take to be exogenous. Under these assumptions, the first order conditions of Equation A.1 define optimal adaptation as a function of income and the climate: $\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)$, which we sometimes denote below as $\boldsymbol{b}_t^*$ for simplicity.

We use this framework to derive an empirically tractable expression for the full mortality risk due to climate change, following Equation 3. We begin by rearranging the agent's first order conditions and using the conventional definition of the VSL (i.e., $VSL = \frac{u(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x}$ following, for example, Becker (2007) and Viscusi and Aldy (2003)[2]) to show that in any time period t,

$$\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}} = \frac{-u(x_t^*)}{\partial u/\partial x[1 - \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = -VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \tag{A.2}$$

That is, marginal adaptation costs (lefthand side) equal the value of marginal adaptation benefits (righthand side), when evaluated at the optimal level of adaptation $\boldsymbol{b}^*$ and consumption x^*. This expression enables us to use estimates of marginal adaptation benefits, which we obtain from the previous section's estimation results, to infer estimates of marginal adaptation costs.

To make the expression in Equation A.2 of greater practical value, we note that the total derivative of expected mortality risk with respect to a change in the climate is the sum of two terms:

$$\frac{d\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} = \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \tag{A.3}$$

The first term on the righthand side of Equation A.3 represents the expected impacts on mortality of all changes in adaptive investments induced by the change in climate; in practice, this term cannot be observed or estimated because of the countless elements of the $\boldsymbol{b}$ vector.[3] The second term is the direct effect that

[2]Note that this definition assumes the utility and marginal utility of consumption when dead is zero. However, this assumption is without loss of generality. If Equation A.1 included a term $v(x_t)$ indicating the (nonzero) utility of consumption when dead, the VSL would be $\frac{u(x)-v(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x + \tilde{f}(\boldsymbol{b},\boldsymbol{C})\partial v/\partial x}$. The result shown in Equation A.2 that marginal adaptation costs equal the value of marginal adaptation benefits, would be unchanged, and recovered adaptation costs would be identical to those derived here and throughout the main text.

[3]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.

the climate would have if individuals did not adapt (i.e., the partial derivative).[4] If, as is expected, climate change produces an increase in the frequency of heat events that threaten human health, it would be natural to expect the first term to be negative, as people make adjustments that save lives. In this case, we expect the second term to be positive, reflecting the impacts of heat on fatalities absent those adjustment.

Equation A.3 makes clear that we can express the unobservable mortality benefits of adaptation (i.e., $\frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}}$) as the difference between the total and partial derivatives of the expected probability of death with respect to climate. This has important practical value because both of these terms can be estimated, as we describe below.

The combination of this algebraic manipulation with Equation A.2 allows us to develop an expression for the *total* adaptation costs incurred as the climate changes gradually from t_0 to t that is entirely composed of elements which can be estimated:[5]

$$A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) = \int_{t_0}^{t} \frac{\partial A(\boldsymbol{b}_s^*)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds = -\int_{t_0}^{t} VSL_s \left[\frac{d\tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \tag{A.4}$$

Equation A.4 outlines how we can use estimates of the total and partial derivatives of mortality risk—with respect to the climate—to infer marginal adaptation costs, even though adaptation itself is not directly observable. In subsection A.4, we show how we use the empirical model described in the main text (Section 4) to separately identify the total derivative $\frac{d\tilde{f}}{d\boldsymbol{C}}$ and the partial derivative $\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$. We empirically quantify these values globally in Section 6.2.

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 $\boldsymbol{b}^*(Y, \boldsymbol{C})$ (Equation A.2) does not imply that our estimates of *total* adaptation costs are equivalent to *total* adaptation benefits for any given population at fixed climate $\boldsymbol{C}$. In general, we expect total adaptation benefits to exceed total adaptation costs, generating surplus from compensatory investments. Here, we define this surplus and illustrate why it is not zero. Empirically, we find that this surplus is substantial (see Section 5.3).

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 6). This surplus can be evaluated at any future climate $\boldsymbol{C}_t$. That is,

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

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

adaptation surplus under a climate changing from time period t_0 to t can be written as:[6]

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

$$= -\int_{\boldsymbol{b}^*(\boldsymbol{C}_{t_0})}^{\boldsymbol{b}^*(\boldsymbol{C}_t)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_t)}{d\boldsymbol{b}} d\boldsymbol{b} - \int_{\boldsymbol{b}^*(\boldsymbol{C}_{t_0})}^{\boldsymbol{b}^*(\boldsymbol{C}_t)} \frac{\partial A(\boldsymbol{b})}{\partial \boldsymbol{b}} d\boldsymbol{b}^* \tag{A.5}$$

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

The expression for adaptation surplus in Equation A.5 is represented as the difference between two integrals, each computed over the unobserved choice vector $\boldsymbol{b}$. To empirically identify adaptation surplus, we aim to rewrite this expression as a difference between integrals which are computed over the multi-dimensional climate $\boldsymbol{C}$, which changes over time. This is an important step, as changes in the climate $\boldsymbol{C}$ are empirically identifiable, while adjustments to $\boldsymbol{b}$ are unobserved by the econometrician. As shown below (as well as in Section 6.1 in the main text), total adaptation costs, the second term in Equation A.5, can be rewritten as an integral over time using a simple change of variables. However, rewriting total adaptation benefits, the first term in Equation A.5, as an integral over time (and hence, climate $\boldsymbol{C}$) requires multiple steps, which we outline below.

To see how we construct an empirically tractable expression for total adaptation benefits (first term in Equation A.5), we first consider a visual illustration. Figure A.2 shows the construction of total adaptation benefits using the same notation and format as the lower right panel of Figure A.1. As in Figure A.1, the red surface represents how expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b(C), C)]$ depend on both climate C and adaptation b, in the case where both climate and adaptation are univariate. The basic idea is that we want to quantify the vertical difference between points s and r (i.e., $s - r$), which can be computed empirically as the vertical difference $q - r$ minus the difference $q - s$. To see why, note that the total benefits of adaptation incurred under a climate change from C_{t_0} to C_t are represented by the vertical difference

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

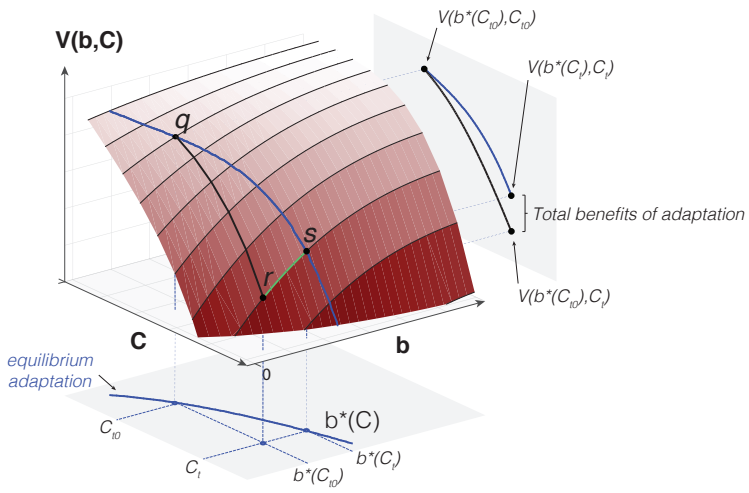

Figure A.2: Recovering total benefits of adaptation using revealed preference. Horizontal dimensions are climate C, representing temperature, and adaptation level b. The vertical dimension is expected survival benefits $V(b, C) = VSL[1 - \tilde{f}(b, C)]$, in units of dollars of WTP. The equilibrium adaptation path $\{b^*(C), C\}$ is drawn in blue (line $q \to s$), and the off-equilibrium path $\{b^*(C_{t_0}), C\}$ is drawn in black (line $q \to r$). To derive the total benefits of adaptation under a change in climate from C_{t_0} to C_t, we integrate the surface along the green line (line $r \to s$), evaluating changes in survival benefits at a fixed climate C_t, as adaptation evolves from $b^*(C_{t_0})$ to $b^*(C_t)$. The magnitude of total adaptation benefits is shown on the $V \times C$ plane on the right panel.

between points s and r (shown on the $V \times C$ plane on the right panel), because this height measures the total mortality benefits realized from optimally investing in adaptation $b^*(C_t)$ when experiencing climate C_t, instead of holding adaptation fixed at its initial level $b^*(C_{t_0})$. This difference can be computed in two ways. First, total benefits of adaptation can be computed by traversing along the off-equilibrium green line between points r and s; that is, by holding C fixed at C_2 and integrating $V(b, C)$ over b from $b^*(C_{t_0})$ to $b^*(C_t)$. This integration along the green line represents the definition of total adaptation benefits written in Equation A.5. However, this same vertical distance can alternatively be calculated by traversing along the off-equilibrium black line between points q and r (i.e., holding b fixed at $b^*(C_{t_0})$ and integrating $V(b, C)$ over C from C_{t_0} to C_t), and then subtracting off the value of the survival impacts of the optimal pathway from C_{t_0} to C_t (i.e., the height of the surface at point q minus point s). This integration over C (twice) is empirically identifiable, as changes in climate can, in principle, be observed.

Now, consider the construction of total adaptation benefits in an arbitrary multi-dimensional $\boldsymbol{b} \times \boldsymbol{C}$ space. We first note that the Gradient Theorem implies path independence of line integrals on smooth functions; thus, for a continuous and differentiable surface $VSL[1-\tilde{f}(\boldsymbol{b}, \boldsymbol{C})]$, the integral along any path on this surface depends only on the endpoints of that path. Equation A.5 writes total adaptation benefits using a path along the surface in the $\boldsymbol{b}$ dimension between the end points $\{\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t\}$ and $\{\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_t\}$.[7] 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

[7]Note that while Figure A.2 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.5. For parsimony, we use the latter notation here and in the subsequent expressions.

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

We define such a loop that begins at $\{\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_{t_0}\}$ (analogous to point q in Figure A.2) and traverses along the off-equilibrium path from $\boldsymbol{C}_{t_0}$ to $\boldsymbol{C}_t$ with adaptation fixed at $\boldsymbol{b}^*(\boldsymbol{C}_{t_0})$ (analogous to the black line between q and r in Figure A.2). In the second segment, it traverses in the $\boldsymbol{b}$ dimension, holding $\boldsymbol{C}$ fixed at $\boldsymbol{C}_t$, to arrive at $\{\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t\}$ (analogous to the green line in Figure A.2 and equal to the total benefits of adaptation). Finally, our path arrives back at its starting point by integrating along the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$ (analogous to the blue line between q and s in Figure A.2):

$$\begin{aligned}\oint_L \nabla[VSL\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\cdot\partial\boldsymbol{b}\partial\boldsymbol{C} &= \int_{t_0}^{t} VSL\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}),\boldsymbol{C}_s)}{\partial\boldsymbol{C}}\frac{d\boldsymbol{C}_s}{ds}ds + \int_{\boldsymbol{b}^*(C_{t_0})}^{\boldsymbol{b}^*(C_t)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_t)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &\quad + \int_{t}^{t_0} VSL\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s),\boldsymbol{C}_s)}{d\boldsymbol{C}}\frac{d\boldsymbol{C}_s}{ds}ds \\ &= 0 \end{aligned} \tag{A.6}$$

By rearranging Equation A.6 (including changing the direction of integration for the third segment), we can use this closed loop, which is composed of two computable segments and a third that is unobservable, to calculate the total benefits of adaptation:

$$\begin{aligned}\textit{Total adaptation benefits} &= -\int_{\boldsymbol{b}^*(C_{t_0})}^{\boldsymbol{b}^*(C_t)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_t)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &= -\int_{t_0}^{t} VSL\left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s),\boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}),\boldsymbol{C}_s)}{\partial\boldsymbol{C}}\right]\frac{d\boldsymbol{C}_s}{ds}ds \end{aligned} \tag{A.7}$$

Using Equation A.7 and a change of variables to rewrite the total costs of adaptation as an integral over $\boldsymbol{C}$, we can rewrite Equation A.5 as:

$$\begin{aligned}\textit{Adaptation surplus}\ (\boldsymbol{C}_{t_0}\to\boldsymbol{C}_t) = -\int_{t_0}^{t} VSL\left[\underbrace{\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s),\boldsymbol{C}_s)}{d\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/ adaptation}}} - \underbrace{\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}),\boldsymbol{C}_s)}{\partial\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/o adaptation}}}\right]\frac{d\boldsymbol{C}_s}{ds}ds \\ - \int_{t_0}^{t}\frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_s))}{\partial\boldsymbol{b}}\frac{\partial\boldsymbol{b}_s^*}{\partial\boldsymbol{C}}\frac{d\boldsymbol{C}_s}{ds}ds \end{aligned} \tag{A.8}$$

While the total adaptation benefits term in Equation A.8 (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.8 in terms of objects that are measurable, using Equation 6 from the main

text to substitute for the object $\int_{t_0}^{t} \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_s))}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds$:

$$\begin{aligned} \textit{Adaptation surplus}\,(\boldsymbol{C}_{t_0} \rightarrow \boldsymbol{C}_t) &= -\int_{t_0}^{t} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \\ &\quad + \int_{t_0}^{t} VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \\ &= \int_{t_0}^{t} VSL \left[\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_s}{ds} ds \end{aligned} \tag{A.9}$$

In Equation A.9, the first term inside the integral represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, but holding adaptation actions fixed at the levels that were optimal under the original climate, $\boldsymbol{C}_{t_0}$. In contrast, the second term represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, allowing adaptation actions $\boldsymbol{b}^*(\boldsymbol{C})$ to evolve optimally with the changing climate. Note that because the second term is a partial derivative, its integral is *not* the total change in the mortality rate. While the two partial derivatives in Equation A.9 will be identical when $\boldsymbol{C} = \boldsymbol{C}_{t_0}$, if they diverge at some point after $\boldsymbol{C}$ warms beyond $\boldsymbol{C}_{t_0}$, then surplus will be nonzero. Thus, a sufficient condition for positive surplus is:

$$\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_{t_0}), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} > \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_s), \boldsymbol{C}_s)}{\partial \boldsymbol{C}} \quad \forall s \in (t_0, t] \tag{A.10}$$

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.10 is weakly positive, and the total adaptation surplus over the climate trajectory $\boldsymbol{C}_{t_0} \rightarrow \boldsymbol{C}_t$ is positive.

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

In Section 6.1, we describe how we use econometric estimation of Equation 4 in combination with climate model projections to construct empirical estimates of changes in adaptation costs due to climate change. 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 A.4), and integrating this difference. To empirically construct an estimate of these costs, we begin by taking expectations of Equation 4 over weather realizations $\boldsymbol{T}$, to specify our empirically estimated *expected* mortality risk for an age group a in region r for year t:

$$\hat{\bar{f}}(.)_{art} \equiv \mathrm{E}[\hat{f}(.)_{art}] = \mathrm{E}[\underbrace{\hat{g}_a\left(\boldsymbol{T}_{rt}, TMEAN_{rt}, \log(GDPpc)_{rt}\right)}_{\hat{g}_{art}(\cdot)}] + ... \tag{A.11}$$

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 4. The expectation of $\hat{g}(\cdot)$ is computed over realizations of temperature for region r in year t from the prior 15 years, with weights of historical observations linearly declining in time. Below we omit subscripts for clarity, but the following calculation is conducted yearly for each age and region separately, for each of our 33 high-resolution climate models.

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

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

$$\frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}} = \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} = \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} \tag{A.12}$$

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 $\hat{\bar{f}}(.)$ with respect to the climate, we consider both the partial effect of changes to $\boldsymbol{T}$ and the effect of adaptive adjustments captured by the effect of $TMEAN$. The total effect of the climate on expected mortality risk is:

$$\begin{aligned} \frac{d \hat{\bar{f}}_t}{d \boldsymbol{C}} = \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} &= \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\bar{f}}_t}{\partial TMEAN} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \\ &= \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \left. \frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN} \right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \end{aligned} \tag{A.13}$$

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.

The *difference* between the total and partial derivatives of expected mortality risk with respect to the

climate is thus the difference between Equations A.13 and A.12:

$$\frac{d\hat{\tilde{f}}_t}{d\boldsymbol{C}} - \frac{\partial\hat{\tilde{f}}_t}{\partial\boldsymbol{C}} = \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \tag{A.14}$$

The righthand side of Equation A.14 is fully computable for years in our projection using a combination of empirically estimated parameters, $\hat{g}(\cdot)$, and climate projections, $\{\boldsymbol{T}, TMEAN\}$. Substituting Equation A.14 into Equation 6 from the main text allows us to estimate non-marginal changes in adaptation costs incurred as the climate of each population changes. In each projection, we solve for adaptation costs as a region's climate evolves from time period t_0 to t:

$$\begin{aligned}
\widehat{A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A}(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) &\approx -\int_{t_0}^{t} VSL_s \left[\frac{d\hat{\tilde{f}}_s}{d\boldsymbol{C}} - \frac{\partial\hat{\tilde{f}}_s}{\partial\boldsymbol{C}}\right] \frac{d\boldsymbol{C}_s}{ds} ds \\
&\approx -\sum_{\tau=t_0+1}^{t} VSL_\tau \left(\left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_\tau, Y_t}\right) (TMEAN_\tau - TMEAN_{\tau-1}) \\
&\approx -\sum_{\tau=t_0+1}^{t} VSL_\tau \, \hat{\boldsymbol{\gamma}}_{\mathbf{1}} \mathrm{E}[\boldsymbol{T}]_\tau \left(TMEAN_\tau - TMEAN_{\tau-1}\right),
\end{aligned} \tag{A.15}$$

where the second equality results from substitution of Equation A.14 into Equation A.4 and from employing a discretized approximation of the continuous integral (we use discrete time-steps of one year). As noted in the main text, recall that we hold income fixed at its endpoint value in the calculation of Equation A.15. This is because the goal of the calculation is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate only. Changes in adaptation expenditures due to rising incomes may change mortality risk under climate change, but these changes are voluntary and are not the consequence of the changing climate, and are therefore not included in our calculation of the total mortality-related costs of climate change. These income effects are accounted for econometrically in the estimation of Equation 4 through the interaction with income and they influence predicted temperature-mortality relationships in all of our calculations, but we do not track the cost of these effects and these costs are intentionally excluded from our calculation of climate-change-induced adaptation spending.

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

A.5 Alternative specification: Including adaptation in the utility function

Throughout the main text, we construct estimates of adaptation costs derived from a representative agent's problem in which utility is a function only of a consumption good x. In this simple model (see Equation 5), there is no direct utility benefit of adaptation behaviors or investments $\boldsymbol{b}$; instead, the actions represented by this composite good influence the agent's problem only through changing mortality risk. In an alternative specification shown here, we allow agents to derive utility both from consumption of x and also possibly

from the choice variables in $\boldsymbol{b}$ (for example, air conditioning might increase utility directly, regardless of its effect on mortality risk). We demonstrate that the implications of this alternative model are purely in the *interpretation* of our empirically derived adaptation cost estimates; the calculation described in Section 6.1 of the main text does not change.

As in Section 6.1 of the main text, we consider a single representative global agent who faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$ in each period t. We further assume there exists some numeraire good x_t for which utility $u(x_t, \boldsymbol{b}_t)$ is quasilinear. As above, this agent maximizes utility conditional on *expected* weather realizations, subject to an exogenous budget constraint and exogenously determined emissions. Letting $\tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t) = \mathbb{E}_{\boldsymbol{c}_t}[f(\boldsymbol{b}_t, \boldsymbol{c}(\boldsymbol{C}_t)) \mid \boldsymbol{C}_t]$ represent the expected probability of death, the agent solves:

$$\max_{\boldsymbol{b}_t, x_t} \; u(x_t, \boldsymbol{b}_t)\left[1 - \tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t)\right] \quad s.t. \;\; Y_t \geq x_t + A(\boldsymbol{b}_t), \tag{A.16}$$

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

Rearranging the agent's first order conditions and using the conventional definition of the VSL,[8] we can write:

$$\underbrace{\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}_t} - \frac{\partial u/\partial \boldsymbol{b}}{\partial u/\partial x}}_{\text{net marginal cost of } \boldsymbol{b}} = \frac{-u(x_t^*, \boldsymbol{b}_t^*)}{\partial u/\partial x[1 - \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = \underbrace{-VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}}}_{\substack{\text{marginal survival} \\ \text{benefit of } \boldsymbol{b}}} \tag{A.17}$$

This expression governs expenditures on adaptation. Its righthand side is the product of the negative of the VSL and the marginal change in expected mortality risk due to a change in adaptation, so it represents the expected marginal benefit (in dollar value) of adjusting $\boldsymbol{b}$ through its effect on mortality risk. This object is identical to its counterpart in Equation A.2 in subsection A.2. The lefthand side has two parts. The first term represents the marginal cost of all pecuniary expenditures incurred due to a marginal change in adaptation $\boldsymbol{b}$, such as spending on units of air conditioning. The second term represents [minus] the dollar value of all non-mortality marginal utility benefits or costs derived from a marginal change in $\boldsymbol{b}$, such as the utility of enjoying air conditioning or the disutility of exercising at midnight to avoid daytime heat (note that this object is expressed in dollars of WTP by dividing through by the marginal utility of consumption, $\partial u/\partial x$). Together, these two terms can be interpreted as the *net* marginal cost of all adaptive actions composing the composite $\boldsymbol{b}$, because non-mortality marginal benefits and costs are removed from the marginal pecuniary expenditures term $\partial A/\partial \boldsymbol{b}$.

Both terms composing net marginal costs in Equation A.17 are unobservable. In contrast, the marginal survival benefit can be rewritten as the product of the negative of the VSL and the difference between the total and partial derivatives of mortality risk with respect to the climate – i.e., $\frac{d\tilde{f}}{d\boldsymbol{C}} - \frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$ (see Equation A.3).

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

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

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

$$\begin{aligned} A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - A(\boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0})) - \frac{1}{\partial u/\partial x}\left[u(x^*(Y_t, \boldsymbol{C}_t), \boldsymbol{b}^*(Y_t, \boldsymbol{C}_t)) - u(x^*(Y_t, \boldsymbol{C}_{t_0}), \boldsymbol{b}^*(Y_t, \boldsymbol{C}_{t_0}))\right] \\ = \int_{t_0}^{t} \left[\frac{\partial A(\boldsymbol{b}_s^*)}{\partial \boldsymbol{b}} - \frac{\partial u(x_s^*, \boldsymbol{b}_s^*)/\partial \boldsymbol{b})}{\partial u(x_s^*, \boldsymbol{b}_s^*)/\partial x)}\right] \frac{d\boldsymbol{b}_s^*}{d\boldsymbol{C}} \frac{d\boldsymbol{C}_s}{ds} ds \\ = -\int_{t_0}^{t} VSL_s \left[\frac{d\tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_s^*, \boldsymbol{C}_s)}{\partial \boldsymbol{C}}\right] \frac{d\boldsymbol{C}_s}{ds} ds, \end{aligned} \tag{A.18}$$

where the last line relies on substitution from Equations A.17 and A.3. The righthand side of Equation A.18 can be approximated empirically as shown in Equation 7 in Section 6.1 in the main text. Thus, the only implication of this alternative model specification is that adaptation cost estimates should be interpreted as pecuniary expenditures net of direct utility benefits and costs.

Similarly, the mortality "partial" social cost of carbon shown in the main text, which relies on an estimate of adaptation costs, is unchanged under this alternative model specification. However, as in Equation A.18, the mortality partial SCC should be interpreted here as the marginal willingness to pay to avoid the alteration of mortality risk associated with a marginal increase in greenhouse gas emissions inclusive of the benefits and costs of adaptations undertaken to reduce mortality risk. Indeed, the omission of the direct utility benefits and costs of adaptation behaviors and technologies from the mortality partial SCC is intentional, because they are not a response to mortality-related risks. However, these utility effects are caused by climate change and should be included in a full, all-sector SCC.

B Data appendix

B.1 Mortality data

Our mortality data represent 41 countries. In some cases our data represent the universe of reported deaths in those countries, while in others (e.g., China), data are representative samples, as no vital statistics registry system exists. Combined, our dataset covers mortality outcomes for 55% of the global population. Data are drawn from multiple, often restricted, national and international sources, all mortality datasets contain information on deaths per 100,000 population from all causes at a monthly or annual frequency, and all except India contain age-specific mortality rates. Each of the countries' data are drawn from distinct databases, details of which are provided below. Figure B.1 displays the spatial coverage and resolution of all mortality records used, as well as their temporal coverage.

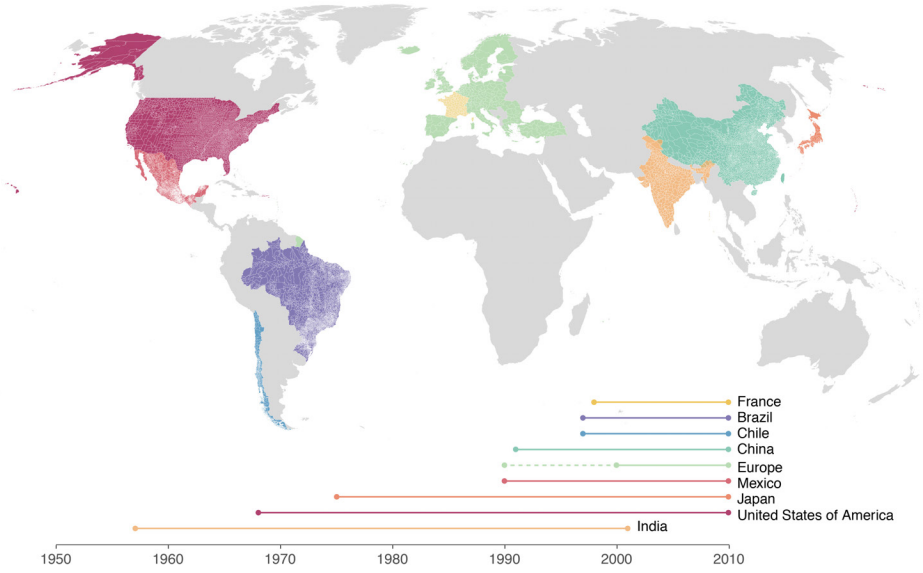

Figure B.1: Mortality statistics used to estimate the relationship between mortality, temperature, climate, and income. Figure shows the spatial distribution and resolution of mortality statistics from all countries used to generate regression estimates of the temperature-mortality relationship. Temporal coverage for each country is shown under the map (the dotted line for the European Union (EU) time series indicates that start dates vary for a small subset of countries).

B.1.1 Brazil

Brazilian mortality data at the ADM2-month level were obtained from the Mortality Information System (SIM) of the Ministry of Health in Brazil (Ministry of Health in Brazil, 2019).[9] 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).

[9]Data are available here: http://ghdx.healthdata.org/series/brazil-mortality-information-system-sim.

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).[10] 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)[11] and merged with the death counts to calculate mortality rates. Administrative boundary files were downloaded from GADM (Global Administrative Areas, 2012).

B.1.3 China

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

[10] Data are available at the first administrative unit (i.e., ADM1) here: https://repositoriodeis.minsal.cl/ContenidoSitioWeb2020/uploads/2019/06/Mortalidad_observada_ajustada_sexo_Region_Suicidios_1997_2017_publicar.xlsx. For access to the current version of the data at the second administrative unit (i.e., ADM2), contact the Ministry of health at: https://www.minsal.cl/.

[11] Data are available here: https://www.ine.cl/estadisticas/sociales/demografia-y-vitales.

boundaries.

B.1.4 European Union

The EU maintains a centralized statistical database known as EuroStat (Eurostat, 2013)[12] 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.[13] 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 B.1. 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)[14] for the years 1998-2010. Data from this source do not have a categorization of mortality for a <5 year old age group, as used in the main analysis. The youngest age group for which there are data is ages 0-19. In the main analysis, we assign the mortality rates in the French data for the 0-19 age group to the <5 age group when pooling across countries. As this introduces some measurement error, we perform a robustness check in which we alternatively assign the deaths in the 0-19 age group to our 5-64 age group; this leads to a minimal change in the multi-country pooled results shown in Table D.2. 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 D.3). However, these data are used to assess the external validity of our extrapolation methods, as discussed in Appendix D.8.

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

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

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

B.1.7 Japan

Japanese data on mortality and population at the prefecture-year[15] level were obtained from the National Institute of Population and Social Security Research[16] 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.[17] 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.[18] There were seven municipalities (less than 0.5% of total municipalities) that had inharmonious borders across data sets and years due to administrative splits or mergers; we assigned these municipalities into their respective unions before the splits or after the mergers.

B.1.9 United States

U.S. data on the universe of mortality and population at the county-year level were obtained from the Center for Disease Control (CDC) Compressed Mortality Files (CMF)[19] 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.[20]

[15] 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.

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

[17] Data are available here: https://en.www.inegi.org.mx/programas/mortalidad/.

[18] 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.

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

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

B.1.10 Aggregate data

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

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

Table B.1: 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)[21]
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 estimates of the mortality effects and full mortality risk of climate change 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 $0.25^\circ \times 0.25^\circ$ 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.[22] 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 $1^\circ \times 1^\circ$ 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 $0.5^\circ \times 0.5^\circ$ 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 ($\sim 1^\circ$ 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).[23] 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

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

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

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

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

B.2.3 SMME and model surrogates

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

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

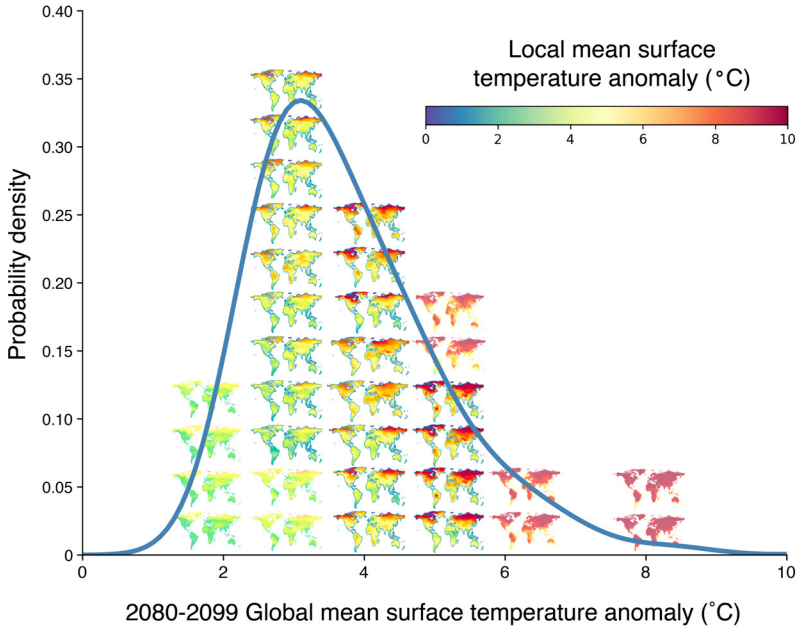

Figure B.2: Future climate projections from the surrogate model/mixed ensemble (SMME). Figure shows the 21 climate models (outlined maps) and 12 model surrogates (maps without outlines) that are weighted in climate change projections so that the weighted distribution of the 2080 to 2099 global mean surface temperature anomaly (ΔGMST) exhibited by the 33 total models matches the probability distribution of estimated ΔGMST responses (blue-gray line) under RCP8.5. For this construction, the anomaly is relative to values in 1986-2005.

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

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

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

To see how this calculation is operationalized, consider the fourth-order polynomial specification for temperature used in our main set of results for estimation of Equations D.21 and 4. In this case, we begin with data on average temperatures for each day d at each grid cell z, generating observations T_{zd}. These grid-level values must then be aggregated to the level of an administrative unit i in year t. To do this, we first raise grid-level temperature to the power p, computing $(T_{zd})^p$ for $p \in \{1, 2, 3, 4\}$. We then take a spatial average of these values over administrative unit i, weighting the average by grid-level population (and accounting for fractional grid cells that fall partially within administrative units). Population weights are time-invariant and calculated from the 2011 Landscan dataset (Bright et al., 2012). We then sum these daily polynomial terms T^p_{zd} over days in the year t. The vector of annual, administrative-level-by-year temperature variables we use for estimation is thus:

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

where w_{zi} is the share of i's population that falls into grid cell z, and where superscripts indicate polynomial powers. This nonlinear transformation performed prior to aggregation allows the aggregated measure of temperature to capture grid-by-day level exposure to very hot and very cold temperatures. In the econometric estimation of Equations D.21 and 4, quadratic polynomials in precipitation are similarly calculated and weighted averages are taken over administrative units. In Figure D.4, 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 4) 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**.[24] 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 B.1).
- **Eurostat (2013) subnational GDP**.[25] This dataset provides national and sub-national level income data for the European countries in our dataset. We use this dataset to obtain subnational income at the NUTS2 level of aggregation, which is the level at which we observe mortality records.
- **Gennaioli et al. (2014) subnational GDP**. This dataset provides national and sub-national income data for 1,503 administrative regions from 83 countries. We use this dataset to obtain subnational level income data for all countries outside the EU: Brazil, Chile, China, France,[26] 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 4. To do so, we use Eurostat (2013) and Gennaioli et al. (2014) to downscale the PWT national-level incomes. We prefer this approach to using the subnational data directly, as there are known inconsistencies in measurement of subnational GDP across countries. Thus, we make the assumption that the within-country distributions of GDP recorded in Eurostat (2013) and Gennaioli et al. (2014) are accurate, but the exact levels may not be. We rely on the PWT data as a consistent measure of GDP levels for all countries; thus, our subnational GDP estimates sum to national GDP from PWT for all countries in the sample. For administrative region s in country c in year t we calculate a weight, ν_{sct} that will apportion national income to subnational regions as follows:

[24] Penn World Tables (PWT) database: https://www.rug.nl/ggdc/productivity/pwt/.
[25] Eurostat database: http://ec.europa.eu/eurostat/data/database.
[26] As noted in Appendix B.1, we use higher resolution mortality data from France than that which is available through EuroStat. Therefore, we also rely on administrative income data from Gennaioli et al. (2014) instead of lower resolution income data from EuroStat.

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

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

where $GDPpc^{PWT}$ corresponds to per capita GDP drawn from the PWT dataset. Using these estimates of administrative-level GDP per capita, we construct the time-invariant income covariate $\log(GDPpc)_s$ used for estimation of Equation 4 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 B.1).

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 logarithmically. All subnational income data are in constant 2005 dollars PPP. A summary of the available years of data before interpolation is given in Table B.2.

Country	ISO code	Years in mortality sample	Years in income sample[27]
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 B.2: Temporal coverage of mortality records and years of available subnational income data.

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

B.3.2 Income projections and downscaling methodology

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

While there are many models within the SSP database, only the IIASA GDP model and OECD Env-Growth model provide GDP per capita projections for a wide range of countries. The IIASA GDP model describes incomes that are lower than the OECD Env-Growth model, so we produce results for both of these models to capture uncertainty within each socioeconomic scenario (we compute results for three socioeconomic scenarios: SSP2, SSP3, and SSP4). To construct annual estimates, we smoothly interpolate between the time series data in the SSP database, which are provided in 5-year increments. For each 5-year period, we calculate the average annual growth rate, and apply this growth rate to produce each year's estimate of GDP per capita.[28]

Throughout the main text, we show results relying on SSP3, although sensitivity of all main results to socioeconomic scenario are shown in the Appendix. While the methodology we develop to estimate future impacts of climate change on mortality, as well as a partial mortality-only SCC, can be applied to any available socioeconomic scenario, we emphasize SSP3 because its historic global growth rates in GDP per capita and population match observed global growth rates over the 2000-2018 period much more closely than either SSP2 or SSP4, as shown below in Table B.3.

Although the SSP scenarios provide national-level income projections, our high-resolution analysis requires estimates of location-specific GDP within country borders. To generate values of income for each of our 24,378 impact regions over time, we allocate national GDP per capita values from the SSPs across impact regions within a country through a downscaling procedure that relies on nightlights imagery from the NOAA Defense Meteorological Satellite Program (DMSP). This approach proceeds in two steps. First, we use available subnational income data from Gennaioli et al. (2014) in combination with higher-resolution income data from the U.S., China, Brazil, and India, to empirically estimate the relationship between GDP per capita and nightlight intensity.[29] Second, we use this estimated relationship to allocate national-level GDP data across impact regions within each country, based on relative intensity of night lights in the present. While this approach models heterogeneity in income levels across impact regions, each region grows in the future at the same rate as the national country projection from the SSPs. We detail these two steps below.

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

from Gennaioli et al. (2014).

[28] 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.

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

Table B.3: Comparison of SSP growth rates to observed data in the historical record This table shows global average growth rates in GDP per capita and in population from observational data (World Bank), as well as from each SSP scenario used in our analysis. Note that International Institute for Applied Systems Analysis (IIASA) GDP model (Samir and Lutz, 2014) only provides GDP per capita estimates after 2010. For both GDP per capita and population, and for each historical time period, SSP3 matches historical data more closely; we therefore show climate change projection results using this scenario throughout the main text.

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

economic output to nightlights intensity, we take an unconventional regression approach to recovering this relationship because our goal is to apportion national income within a country, as opposed to predict the level of income at any given location. In particular, we are interested in the ratio $\frac{GDPpc_{rc}}{\sum_{r \in c} w_{rc} GCPpc_{rc}}$ for impact region r in country c (where w_{rc} is a region-specific population weight), which will allow us to predict income at the impact region level, given projections of national GDP per capita from the SSPs, $\sum_{r \in c} w_{rc} GDPpc_{rc} = GDPpc_c^{SSP}$. Thus, we estimate a regression relating *relative* GDP per capita to *relative* nightlights intensity, where each administrative region's values are calculated as relative to the country mean. The dependent variable for administrative region i in country c and year t is thus $\frac{GDPpc_{ict}}{\sum_{i \in c} w_{ict} GDPpc_{ict}}$.[30] To construct a measure of location-specific relative nightlight intensity, we calculate a z-score of nightlights (ZNL) for each administrative region i within a country c using:

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

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

The regression we estimate is as follows:

$$\frac{GDPpc_{ict}}{\sum_{i \in c} w_{ict} GDPpc_{ict}} = \alpha + \beta ZNL_{ict} + \epsilon_{ict} \tag{B.19}$$

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

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

estimated coefficients from Equation B.19 to compute income at impact region level. To do so, we construct values $ZNL_{rct} = \frac{NL_{rct} - \overline{NL}_{ct}}{\sigma(NL_{ct})}$ for each impact region r using the average of stable nightlights from DMSP across the years 2008-2012. We then estimate $GDPpc_{rct}$ as follows:

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

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

B.3.3 Population projections and downscaling methodology

Future projections of national populations are derived from the International Institute for Applied Systems Analysis (IIASA) (Samir and Lutz, 2014) population projections as part of the Shared Socioeconomic Pathways (SSPs).[31] The IIASA SSP population projections provide estimates of population by age cohort, gender, and level of education for 193 countries from 2010 to 2100 in five-year increments. Each projection corresponds to one of the five SSPs, as defined in O'Neill et al. (2014). These populations are mapped to impact regions by country code using 3-digit country ISO-codes.

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

$$\widehat{pop}_{rtv} = \begin{cases} pop_{ctv}^{SSP} \left(\frac{pop_{r,2011}^{LandScan}}{\sum_{r \in c} pop_{r,2011}^{LandScan}} \right), & \text{if } r \in C \\ pop_{r,2011}^{LandScan}, & \text{if } r \notin C \end{cases} \tag{B.20}$$

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

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

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

Table B.4: Scale and scope of existing empirically-based estimates of the mortality effects of climate change. Table highlights papers quantifying the mortality effects of climate change that either are broad in spatial scope (e.g., covering multiple countries) or aim to account for at least one driver of adaptation (e.g., income). The first row highlights the present study, while the last row indicates the studies used to calibrate the mortality component of the FUND IAM, which is currently used to inform the U.S. government's social cost of carbon (SCC).

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

C Spatial units for projection: "Impact regions"

We create a set of custom boundaries that define the spatial units for which location-specific projected impacts 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 C.1, and we detail this agglomeration algorithm below.

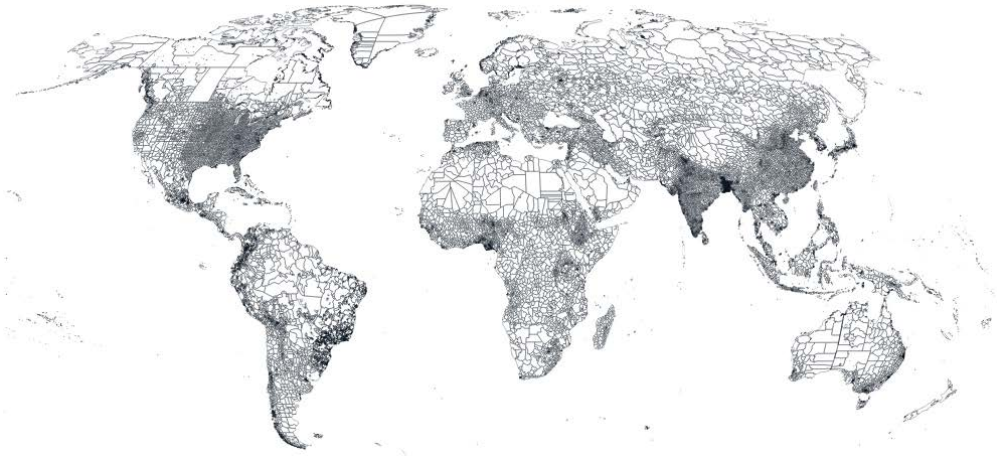

Figure C.1: Map of the 24,378 "impact regions" for which location-specific projections are calculated. We use a clustering algorithm to form these regions from the full set of GADM administrative regions, such that they are roughly similar in total population, and so that they are approximately internally 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 C.2. 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 C.2. 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 C.2. 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 C.2. 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 C.2, 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:

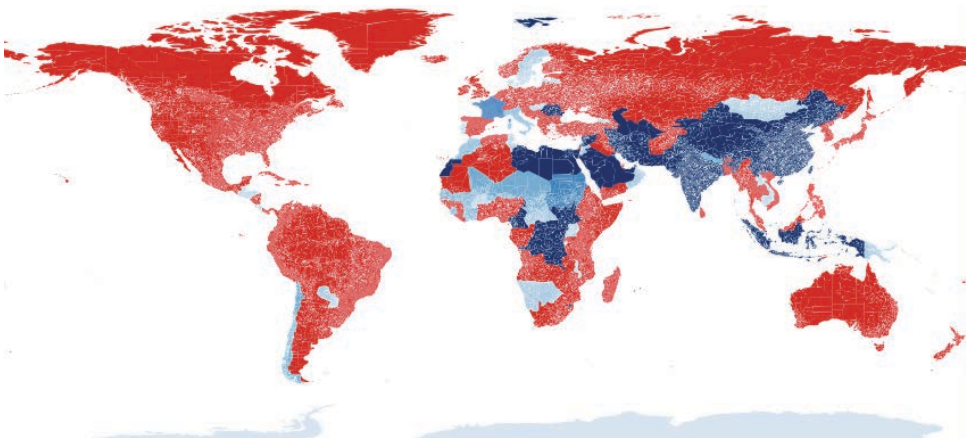

Figure C.2: Categorization of countries based on the method used to construct impact regions out of GADM administrative regions. A country's target number of impact regions is $target_c$, as computed in the text. Countries in shades of blue have target values that can be approximated by one of the available 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),[32] 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.

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

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

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

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

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

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

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

This appendix shows additional illustrations of and tabular results for the main econometric regressions used and discussed throughout the main text (Figures D.1 and D.2, and Table D.1), results obtained using a pooled version of the main model in which no heterogeneity in the mortality-temperature relationship is modeled (Figure D.3 and Table D.2), a set of robustness checks for the main empirical results (Figures D.4, D.6, D.7, and D.8 and Tables D.3 and D.4), and a set of out-of-sample validation tests designed to evaluate the accuracy with which our estimates predict mortality-temperature responses in locations and time periods that are not used for estimation (Figures D.9, D.10, and D.11 and Tables D.5 and D.6).

D.1 Age-specific heterogeneity in the mortality-temperature relationship by average income and average climate

The estimation of Equation 4 tests for systematic heterogeneity in the mortality-temperature response function by modeling interactions between the temperature variables ($\boldsymbol{T}$) and the ADM1-level covariates of average climate ($TMEAN$) and average income ($\log(GDPpc)$). To see how we implement Equation 4 in practice, note that in Equation D.21, 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 4, we allow $g_a(\boldsymbol{T}_{it})$ to change with climate and income by allowing each element of $\boldsymbol{\beta}_a$ to be a linear function of these two variables. We do not include a triple interaction between temperature, climate and income. Using this notation, our estimating equation is:

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

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

Tabular results from this estimation are reported in Table D.1 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).[33] 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

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

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

Table D.1: 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. $TMEAN$ is defined as the average annual temperature in °C over the sample period, while $\log(GDPpc)$ is the logarithm of the average annual GDP per capita over the sample period, measured in constant 2005 dollars PPP. Regression is a fourth-order polynomial in daily average temperature, estimated using GMFD weather data with a sample that was winsorized at the top 1% level. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects. Each temperature variable is interacted with each covariate.

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

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

As these terms are difficult to interpret, we visualize this heterogeneity in the main text in Figure 1 by dividing the sample into terciles of income and climate (i.e., the two interaction terms), creating nine discrete bins describing the $\log(GDPpc) \times TMEAN$ space. We plot the predicted response functions at the mean value of covariates within each of these nine bins, using the coefficients shown in Table D.1. 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 the main text in Figure 1 for the >64 age category. Figures D.1 and D.2 replicate this figure for the other two age groups in our analysis.

D.2 Results and robustness with pooled model

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

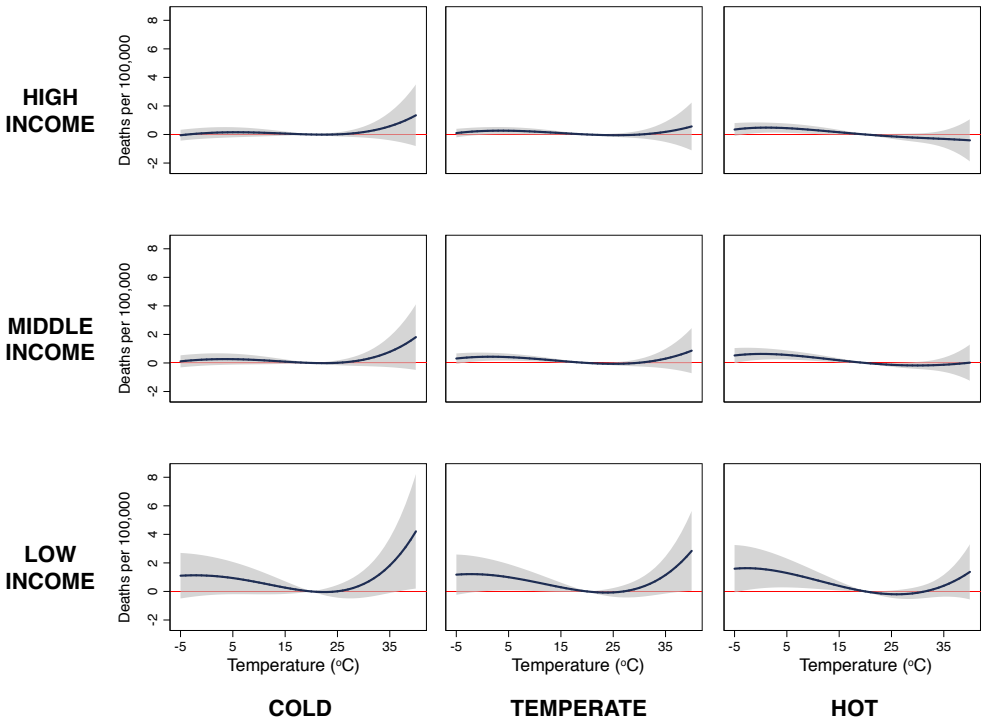

Figure D.1: Heterogeneity in the mortality-temperature relationship (ages <5 mortality rate). Each panel represents a predicted response function for the ages <5 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left 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.

demonstrate robustness to other functional forms of temperature and other climate datasets, and show results for alternate model specifications.

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

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

$$M_{ait} = g_a(\boldsymbol{T}_{it}) + q_{ac}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{D.21}$$

where a indicates age category with $a \in \{< 5, 5\text{-}64, > 64\}$, i denotes the second administrative level (ADM2, e.g., county),[34] c denotes country, and t indicates years. Thus, M_{ait} is the age-specific all-cause mortality

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

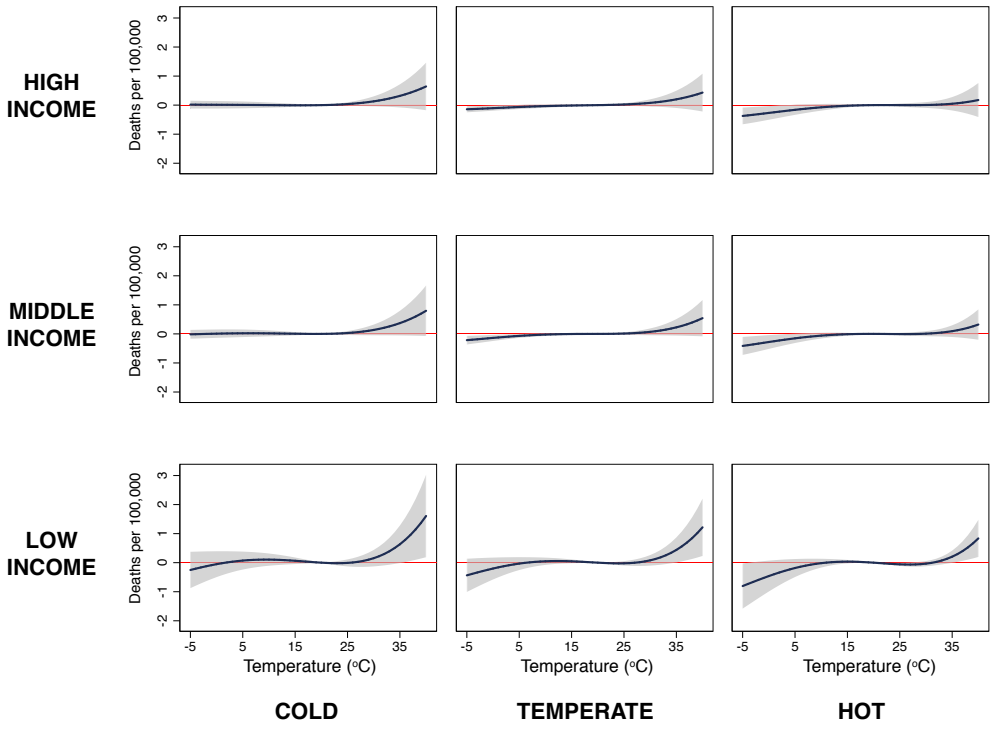

Figure D.2: Heterogeneity in the mortality-temperature relationship (ages 5-64 mortality rate). Each panel represents a predicted response function for the ages 5-64 mortality rate for a subset of the income-average temperature covariate space within our data sample. Response functions in the lower left 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.

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 D.21 is the effect of temperature on mortality, represented by the response function $g_a(\cdot)$, which varies by age. As in our our main specification, $\boldsymbol{T}_{it}$ contains polynomials of daily average temperatures (up to fourth order), summed across the year. These calculations are performed at the grid cell level before being aggregated up to the level of the administrative units in the data. Results for alternative functional form specifications are shown in Figure D.4 and the consequences of alternate functional forms for climate change projection results are shown in Appendix F. Analogous to temperature, we summarize daily grid-level precipitation in the annual ADM2-level vector $\boldsymbol{R}_{it}$. We construct $\boldsymbol{R}_{it}$ as a second-order polynomial of daily precipitation, summed across the year, and estimate an age- and country-specific linear function of this vector, represented by $q_{ac}(\cdot)$.

We fit the multi-country pooled model in Equation D.21 using weighted least squares, weighting by age-specific population because we are recovering treatment effects that average over multiple sources of

unmodeled heterogeneity (Solon, Haider, and Wooldridge, 2015).[35] 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 account for spatial as well as temporal correlation in error structure. Robustness of this model to alternative fixed effects and error structures is shown in Table D.2, and to alternative climate datasets in Figure D.4.

Age-specific pooled multi-country mortality-temperature response functions. As prior work has shown that age cohorts respond differently to temperature, and because we expect considerable demographic transitions in the future, we allow for heterogeneity across age groups in Equation D.21. Specifically, we allow for separate mortality-temperature response functions $g_a(\boldsymbol{T}_{it})$ for each of three age categories (<5, 5-64, >64). Figure D.3 displays the mortality-temperature responses for each age group, estimated from Equation D.21 and using the pooled 40-country sample and our preferred specification (column (2) in Table D.2). This reveals substantial heterogeneity across age groups within our multi-country sample: people over the age of 64 experience approximately 4.7 extra deaths per 100,000 for a day at 35°C (95°F) compared to a day at 20°C (68°F), a substantially larger effect than that for younger cohorts, which exhibit little response. This age group is also more severely affected by cold days; estimates indicate that people over the age of 64 experience 3.4 deaths per 100,000 for a day at −5°C (23°F) compared to a day at 20°C, while there is a relatively weak mortality response to these cold days for other age categories. Overall, these results demonstrate that the elderly are disproportionately harmed by additional hot days and disproportionately benefit from reductions in cold days, consistent with prior evidence from the U.S. (Deschênes and Moretti, 2009; Heutel, Miller, and Molitor, 2017). It is important to note, however, that the oldest age group (over 64 years) accounts for just 12% of the population in our historical sample.

Robustness to temperature functional form and climate data. Figure D.4 displays the results of estimating a version of Equation D.21 using a set of different functional forms of temperature (i.e., different formulations of the temperature vector $\boldsymbol{T}_{it}$) and using two different climate datasets to obtain those temperatures (see Appendix B.2 for details on these climate datasets). Here we show the mortality response $g_a(\boldsymbol{T}_{it})$ for the >64 age group. The four functional forms estimated are fourth-order polynomials, bins of daily average temperature, restricted cubic splines, and piecewise linear splines. The binned functional form is an important benchmark, as it is closest to being fully non-parametric; the similarity of the binned regression response functions with those from three other functional forms is reassuring. The GMFD climate data (top) and BEST climate data (bottom) are drawn from independent sources, as described in Appendix B.2, and lead to broadly similar response functions across all functional forms. All regressions include $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects, and are population weighted.

Alternative specifications. In Table D.2, marginal effects of temperature on age-specific mortality rates are shown for a range of alternative specifications. These estimates can be interpreted as the change in the number of deaths per 100,000 per year resulting from one additional day at each temperature, compared to the reference day of 20°C (68°F). Columns (1)-(3) increase the saturation of temporal controls in the model specification, ranging from country-year fixed effects in column (1) to country-year-age fixed effects in column (2), and adding age-specific state-level linear trends in column (3). Our preferred specification

[35] 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.

Table D.2: Temperature-mortality response function with demographic heterogeneity estimated using pooled subnational data. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the top 1% level. Point estimates indicate the effect of a single day at each daily average temperature value shown, relative to a day with an average temperature of 20°C (68°F).

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

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

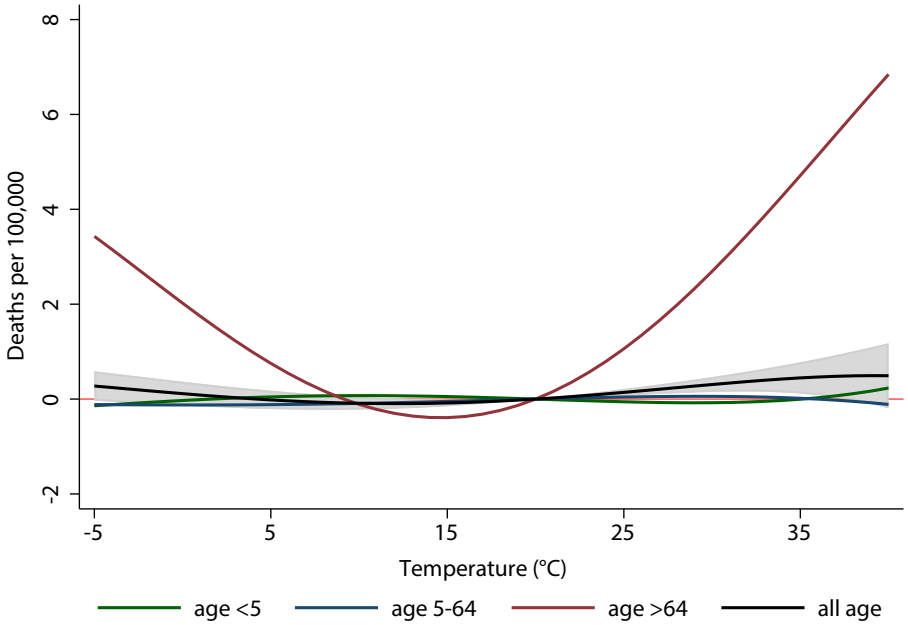

Figure D.3: Mortality-temperature response function with demographic heterogeneity. Mortality-temperature response functions are estimated for populations <5 years of age (green), between 5 and 64 years of age (blue), >64 years of age (red), and pooled across all ages (black, with associated 95% confidence intervals shaded in grey). Regression estimates shown are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level. All age-specific response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects (Equation D.21). Confidence intervals are shown only for the all-age response function; statistical significance for age-specific response functions can be seen in Table D.2.

is column (2), as column (1) does not account for differential temporal shocks to mortality rates by age group, while in column (3) we cannot reject the null of equal age-specific, ADM1-level trends. In column (4), we address the fact that some of our data are drawn from countries which may have less capacity for data collection than others in the sample. Because our mortality data are collected by institutions in different countries, it is possible that some sources are systematically less precise. To account for this, we re-estimate our model using Feasible Generalized Least Squares (FGLS) under the assumption of constant variance within each ADM1 unit.[36] In column (5), we address the possibility that temperatures can exhibit lagged effects on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). Lagged effects within and across months in the same calendar year are accounted for in the net annual mortality totals used in all specifications. However, it is possible that temperature exposure in December of each year affects mortality in January of the following year. To account for this, in column (5) we define a

[36]To do this, we estimate the model in Equation D.21 using population weights and our preferred specification (column (2)). Using the residuals from this regression, we calculate an ADM1-level weight that is equal to the average value of the squared residuals, where averages are taken across all ADM2-age-year level observations that fall within a given ADM1. We then inverse-weight the regression in a second stage, using this weight. All ADM2-age-year observations within a given ADM1 are assigned the same weight in the second stage, where ADM1 locations with lower residual variance are given higher weight. For some ADM2s, there are insufficient observations to identify age-specific variances; to ensure stability, we dropped the ADM2s with less than 5 observations per age group. This leads us to drop 246 (of >800,000) observations in this specification.

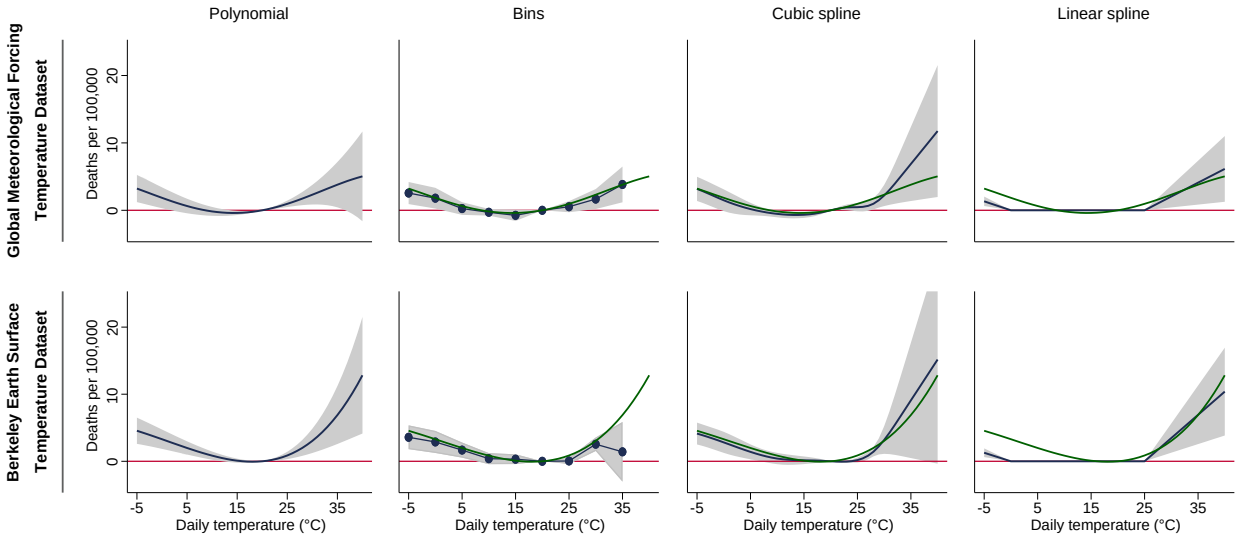

Figure D.4: Robustness of the mortality-temperature relationship to alternative functional forms and to different historical climate datasets (age >64). Row 1 shows the mortality-temperature response function as estimated using daily temperature and precipitation data from the Global Meteorological Forcing Dataset (GMFD). Row 2 shows the same response, using daily temperatures from Berkeley Earth Surface Temperature (BEST), and monthly precipitation from the University of Delaware. Each column displays a distinct functional form, with the fourth-order polynomial shown in column 1 overlaid in teal on each subsequent column. In the binned specification (column 2), annual values are calculated as the number of days that have an average temperature that falls within a fixed set of 5°C bins. The bin edges are positioned at the locations $\{-\infty, -15, -10, -5, 0, 5, 10, 15, 20, 25, 30, 35, +\infty\}$ in °C. In the restricted cubic spline specification (column 3), daily spline terms are summed across the year and knots are positioned at the locations $\{-12, -7, 0, 10, 18, 23, 28, 33\}$ in °C. In the linear spline specification (column 4), heating degree days below 0°C and cooling degree days above 25°C are summed across the year.

13-month exposure window to additionally account for temperatures previous December.[37] Table D.2 shows that the results for both of these alternative specifications are similar in sign and magnitude to those from column (2).

D.3 Additional results: Spatial extrapolation of temperature sensitivity

Figure D.5 reports on our extrapolation of mortality-temperature response functions to the entire globe for the <5 age group as well as the >64 age group shown in the main text in Figure 3. As in Figure 3, panels A and B show predicted mortality-temperature responses for each impact region for 2015 values of income and climate and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Geographic heterogeneity within our sample is shown for hot days in the maps in panels C and D, where colors indicate the marginal effect of a day at 35°C, relative to a day at a location-specific minimum mortality temperature. Grey areas are locations where mortality data are unavailable. Figure D.5E–H show analogous plots, but now extrapolated to the entire globe.

[37] 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 $t-1$.

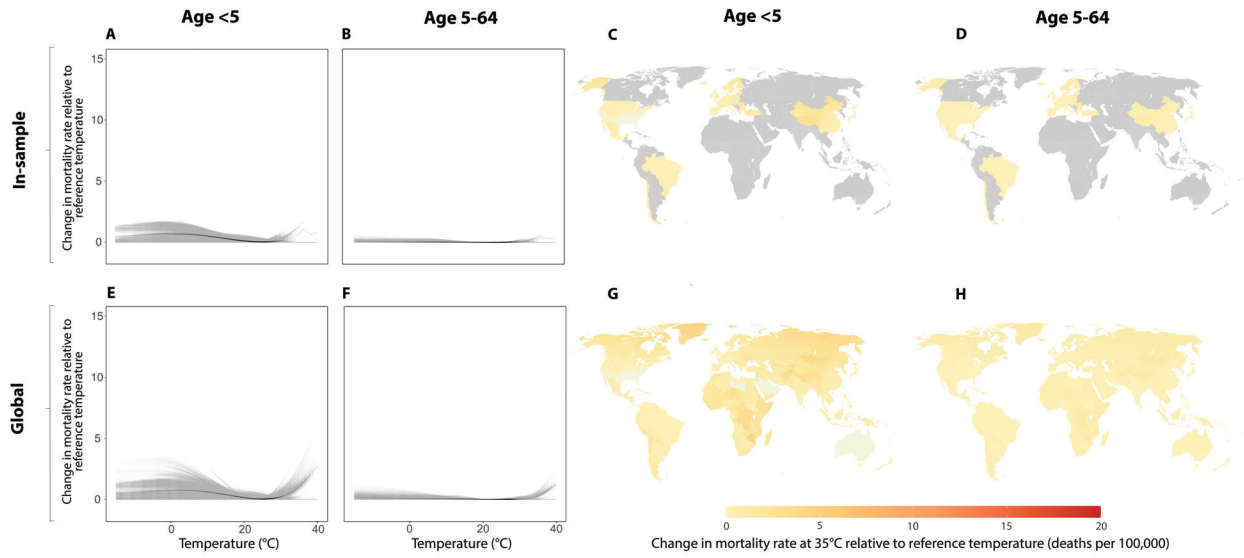

Figure D.5: Using income and climate to predict current response functions globally (ages <5 and 5-64). In panels A, B, E and F, grey lines are predicted response functions for impact regions, each representing a population of 276,000 on average. Solid black lines are the unweighted average of the grey lines, where the opacity indicates the density of realized temperatures (Hsiang, 2013). Panels C, D, G and H show each impact region's mortality sensitivity to a day at 35°C, relative to a location-specific minimum mortality temperature. The top row shows all impact regions in the sample of locations with historical mortality data (included in main regression tables), and the bottom row shows extrapolation to all impact regions globally. Column titles indicate corresponding age categories. Predictions shown are averages over the period 2001-2010 using the SSP3 socioeconomic scenario and climate model CCSM4 under the RCP8.5 emissions scenario. Main text Figure 3 shows the analogous figure for age >64.

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

Our primary results rely on a parsimonious representation of the climate: to capture adaptation to long-run climate, we interact our nonlinear temperature variables ($\boldsymbol{T}$) with the long run average annual temperature ($TMEAN$), conditioning on income ($\log(GDPpc)$). In this specification, $TMEAN$ acts as a summary statistic of the long-run average climate, and we find that the mortality sensitivity to high temperatures declines as $TMEAN$ rises. To test the robustness of this finding, here we use a richer characterization of the climate, replacing our climate interaction term $TMEAN$ in Equation 4 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 4 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 D.3. Consistent with our main results in Table D.1, 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 D.3: 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 D.3 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 D.6 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 D.1 and D.3, 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.

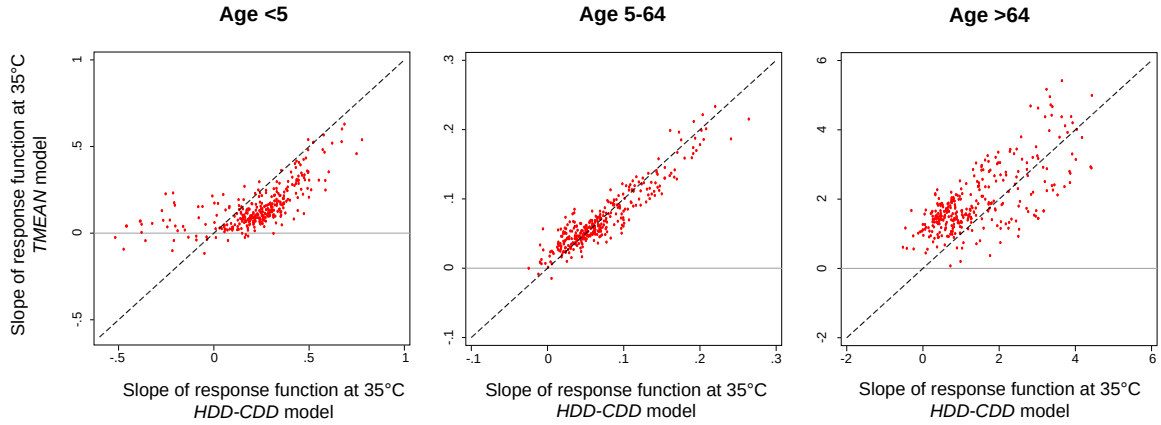

Figure D.6: Predicted mortality-temperature response functions in-sample are similar under alternative characterizations of long-run average annual temperature. Each panel contains a scatter plot of the slope (i.e., derivative) of the predicted mortality-temperature response function, evaluated at 35°C, 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 4 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 4 in which all nonlinear temperature variables are interacted with long-run annual average heating degree days (HDDs) below 20°C and cooling degree days (CDDs) above 20°C. Predictions shown are for all ADM1 regions within our estimating sample. Each column shows predictions for a different age category.

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

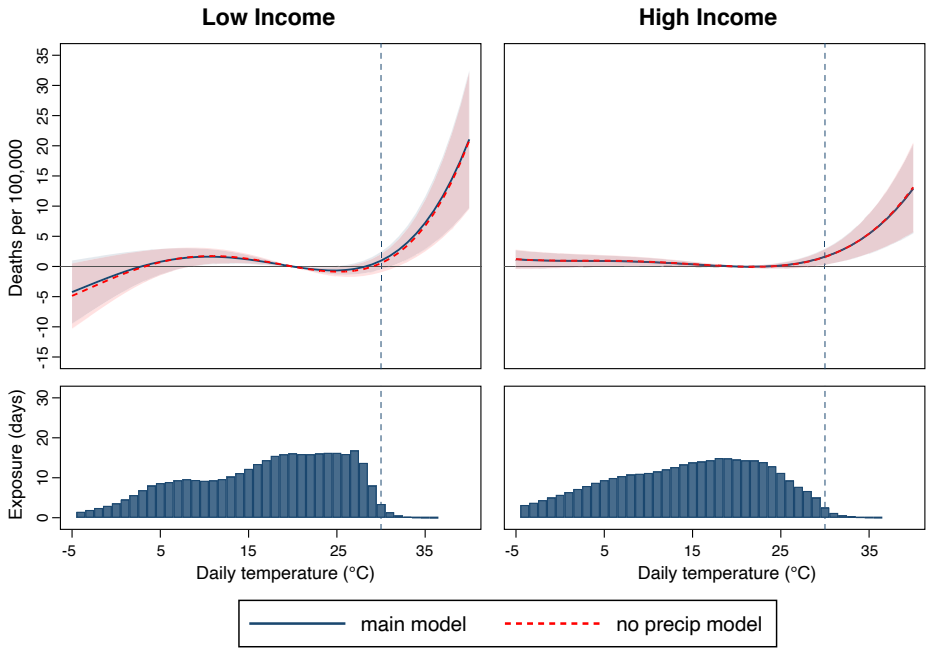

Figure D.7: Predicted mortality-temperature response functions are similar with versus without precipitation controls (age >64). Each panel shows the predicted mortality-temperature relationship resulting from the estimation of a version of Equation 4, and evaluated at the population-weighted mean value of the logarithm of GDP per capita and long run average temperature within the lowest income tercile of the estimation sample (left panel), and the highest income tercile of the estimation sample (right panel). The main regression model (Equation 4) is shown in the solid blue line, while a version of Equation 4 omitting the country-specific quadratic precipitation controls is shown in the dashed red line. Vertical dashed lines indicate the 99.5^{th} percentile of the daily temperature distribution in each income group. 95% confidence intervals for both regression models are shown in the shaded areas.

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

The analysis implemented in the main text relies on a two-factor model to explain heterogeneity in the mortality-temperature relationship. There are three primary reasons we use just income and long-run average climate to model heterogeneity. First, these covariates are conceptually intuitive determinants of adaptation; looser budget constraints enable more investment in adaptive technologies and behaviors, while increased exposure to a particular weather event leads to updated beliefs and corresponding adaptive investments. Second, both covariates have been shown to be important in explaining heterogeneity in climate impacts across a number of other contexts.[38] Finally, substantial research has been conducted to generate projected

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

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

However, a valid critique of this model is that other factors that likely explain heterogeneity in the mortality-temperature relationship are omitted from our main estimating equation (Equation 4). In this section, we examine five additional covariates that plausibly influence the mortality-temperature relationship, but for which projections into the future are not available (potentially due to the impossibility of that exercise). First, we show that the inclusion of these additional variables into the interaction model in Equation 4 has a negligible impact on predicted mortality-temperature relationships. Second, we show that including these covariates in estimation, but omitting them when generating predictions, as would be necessary when generating a climate change impact projection, leads to substantial bias.

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

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

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

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

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

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

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

$$M_{aict} = g_a(\boldsymbol{T}_{it}, TMEAN_s, \log(GDPpc)_s, COVAR_c) + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait}, \tag{D.22}$$

where $COVAR_c$ is one of the five alternative covariates mentioned above. All other variables are identical to our main model in Equation 4. Due to missing data for certain countries, merging these data results in sample sizes that are in most cases smaller than our original sample. Only the institutions regression has 100% of the observations as our original sample, while health (91.5%), education (81.7%), inequality (99.7%), and informality (72.8%) all have smaller sample sizes.

We conduct two tests using the results from estimating Equation D.22 for each alternative covariate. First, we assess whether each added interaction variable actually explains substantial heterogeneity that is not already captured by our two-factor model. These results are shown in Figure D.8, where we use the coefficients estimated in Equations 4 and D.22 to predict mortality-temperature response functions for our main model and for the model including the alternate covariate. We predict responses using average values of income and climate in the lowest income tercile of our data (left column) and in the highest income tercile of our data (right column). Since the sample sizes differ across variables due to data availability constraints, we re-estimate the main model using a comparable sample for each additional covariate, leading to some variation in the main model response function across rows. The daily distribution of temperatures in each income group is shown in the top row.

Figure D.8 shows that differences between the main model (blue solid lines) and the alternative models (red dashed lines) are small across most of the temperature support. Some larger differences do emerge, for example in the mortality-temperature response in poorer countries when including an interaction with education, but across all covariates, none of the differences from our main model are statistically significant. Moreover, the majority of these slightly larger differences occur in the coldest 0.5% of our sample (with the middle 99% of the temperature distribution indicated by vertical dashed lines). The implication of this test is that, while the alternative covariates do in some cases explain some heterogeneity in the temperature-mortality relationship, the difference in predicted mortality-temperature relationships between our main model and these alternatives is never statistically significant and is substantively small across most of the temperature support, reaching a maximum only at the coldest temperatures, which are rare in our sample and are likely to become increasingly rare globally as the Earth warms.

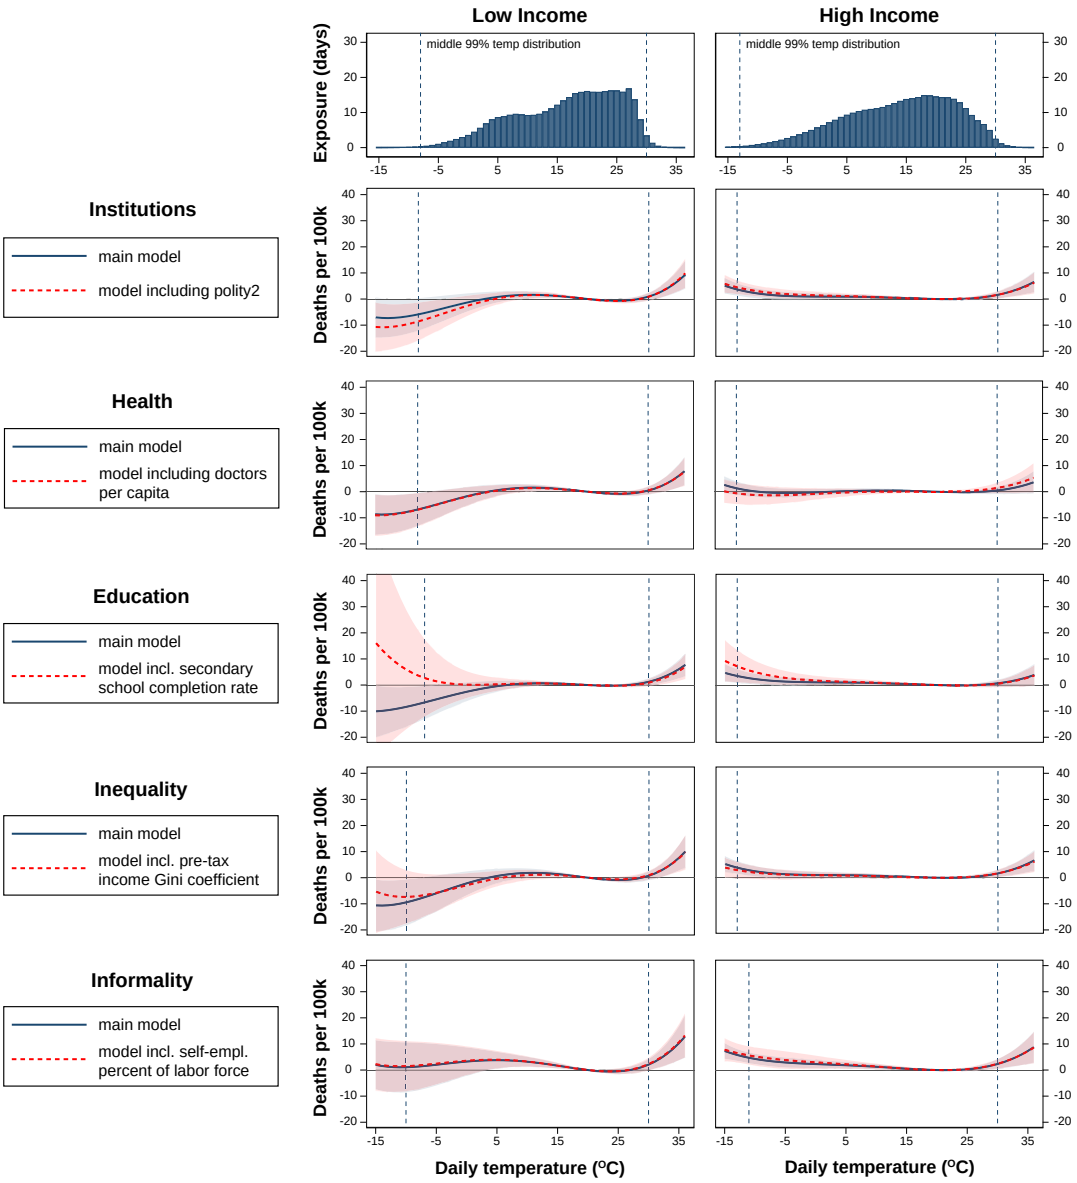

Figure D.8: Predicted mortality-temperature response functions are robust to inclusion of additional interaction terms (age >64). Comparison of response functions estimated using Equation 4 ("main model", blue line) and including additional covariates using Equation D.22 ("alternative model", red dashed line) for the >64 age group. Samples are adjusted according to data availability for each covariate. The left panel shows the predicted response evaluated at the population-weighted mean level of income, climate, and the corresponding covariate across all observations that fall into the lowest tercile of GDP per capita in our estimation sample, while the right shows the corresponding predicted response for observations in the highest tercile of GDP per capita in our estimation sample. The area between the vertical dashed lines is the middle 99% of the distribution of temperatures indicated by the histogram. As samples vary slightly (see Table D.4) across each comparison, the middle 99% varies slightly across figures, as does the main model line. Results for other age groups are similar.

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

In Table D.4 the RMSE values are shown for both cases. The "Main Model" column shows the RMSE from estimation of Equation 4 using the sample that varies slightly across rows based on data availability of each of the alternate covariates. The "Alternative Model" column shows the RMSE from estimation of Equation D.22 for each additional covariate, where predicted values are generated by setting the coefficient on the alternate covariate to zero. This second case mimics the situation in which we estimate a model with an interaction term for a variable that cannot be projected into the future. The key comparison is between the RMSE for the "Main Model" and the RMSE for the "Alternative Model", shown in the final "Difference" column. For each of these additional covariates, the alternative model exhibits a substantially worse model fit (i.e., the "Difference" column is negative), implying that the inclusion of a covariate in estimation for which there are no data in the future would lead to a model that performs strictly worse than the two-factor model used throughout the paper.

Table D.4: Investigating additional sources of heterogeneity in the mortality-temperature relationship. Each row corresponds to model performance metrics from the estimation of Equation D.22 with the inclusion of the named covariate. All covariates other than those in the main model (Equation 4) are observed at country level. Sample sizes differ across rows due to the availability of data for each of the covariates, with sample sizes ranging from 73% to 100% of our main estimating sample. The "RMSE Main Model" column shows the in-sample root mean squared error (RMSE) of our main model estimated using Equation 4. The "RMSE Alternative Model" column shows the in-sample RMSE of a model that is estimated using Equation D.22, but where predictions are generated omitting the impact of the additional covariate. This is done to mimic a situation in which estimation includes historical data on additional determinants of heterogeneity, but climate change projections must be made without projection data for that additional covariate. The differences, all negative, are shown in the difference column and indicate that the RMSEs of our main estimation equation are consistently lower than those for the alternative models.

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

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

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

Throughout our analysis, we use coefficients estimated from Equation 4 in the main text, in combination with local-level observations and projections of $TMEAN$ and $\log(GDPpc)$, to generate predicted response functions in all regions of the world, including where mortality data are unavailable, both in the present and into future (see Section 5.2 for details). In contrast, much prior work generates projected impacts of climate change using spatially and/or temporally homogeneous response functions (e.g., Hsiang et al., 2017; Deschênes and Greenstone, 2011). To assess the performance of our model in predicting mortality-temperature relationships out-of-sample, this section implements multiple custom cross-validation exercises designed to mimic the spatial and temporal extrapolation that is required when using available historical data to generate global climate change projections decades into the future.

We perform three cross-validation exercises, each of which provides multiple measures of the out-of-sample performance of Equation 4. In each case, we compare these performance metrics to the performance of a benchmark model that ignores adaptation, which has been the standard in much of the prior literature, and to a measure of in-sample model fit. We also show results for models that include either climate *or* income as drivers of heterogeneity in the mortality-temperature relationship, but not both. Because we are assessing the performance of our interaction model in predicting mortality sensitivity to temperature, as opposed to mortality rates overall, all measures of model fit are reported using residualized data, in which all fixed effects and controls are removed from identifying variation prior to estimation.[39] Results from all three tests are shown in Table D.5, and are discussed in the following subsections.

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

Sample	Observations	% of global population (2010)	% of global population (2100)	RMSE (income benefits model)	RMSE (climate adaptation model)	RMSE (full adaptation model)	RMSE (no adaptation model)	ΔRMSE: full adapt. versus no adapt.
A: In-sample model fit								
Full Sample In Sample	820,698	–	–	565.29	565.25	565.19	577.80	-12.62
B: 10-fold cross-validation at ADM1 level								
Full Sample Out of Sample	820,698	–	–	565.38	565.38	565.34	577.97	-12.62
C: 9-fold cross-validation across income × long-run temperature covariate space								
Full Sample Out of Sample	820,698	–	–	565.63	565.70	565.66	578.04	-12.37
Low Income - Cold	4156	6.5	0.0	879.72	875.84	877.00	903.69	-26.69
Low Income - Moderate	15,279	9.0	0.5	495.24	494.52	495.14	498.43	-3.29
Low Income - Hot	334,968	65.5	46.0	707.49	707.47	707.54	718.73	-11.19
Middle Income - Cold	3507	1.5	0.0	604.32	604.85	604.68	353.96	250.71
Middle Income - Moderate	15,108	1.0	0.0	256.90	256.99	256.96	264.73	-7.77
Middle Income - Hot	78,160	2.0	24.5	485.58	485.58	485.50	483.97	1.53
High Income - Cold	125,934	5.0	2.0	528.27	529.07	528.53	564.42	-35.89
High Income - Moderate	137,706	5.0	5.0	331.01	331.06	330.91	339.95	-9.04
High Income - Hot	105,880	4.0	22.0	401.76	401.68	401.81	415.12	-13.31
D: 2-fold cross-validation across time (post-2004 hold-out)								
Pre 2005 In Sample	607,979	–	–	565.79	565.83	565.72	577.85	-12.13
Post 2004 In Sample	212,719	–	–	563.50	563.38	563.03	578.38	-15.35
Post 2004 Out of Sample	212,719	–	–	564.11	563.69	564.00	578.36	-14.36

Table D.5: Evaluation of out-of-sample model performance This table presents results from three separate cross-validation exercises. Panel A shows in-sample model performance. For panels B, C, and D, a section of the data is omitted from the sample and the empirical model is re-estimated using the remaining observations. The mortality rates of the omitted observations are then predicted out-of-sample, and the root mean squared error (RMSE) is calculated. In the column titled "RMSE (full adaptation model)", this calculation is conducted using the estimation of Equation 4, which includes heterogeneity in the mortality-temperature response driven both by adaptation to climate and by income. In the column "RMSE (no adaptation model)", this calculation uses a regression equation that omits any model of heterogeneity or adaptation. Differences between the model with adaptation (i.e., Equation 4) and the model without adaptation are shown in the "ΔRMSE: full adapt. versus no adapt." column, with negative values indicating that our main estimating equation performs better than the benchmark model without adaptation. The columns "RMSE (income benefits model)" and "RMSE (climate adaptation model)" rely on versions of Equation 4 that include only income heterogeneity or only climate heterogeneity, respectively. Panel B reports results from k-fold cross-validation, panel C from a custom cross-validation that blocks data based on long-run income per capita and average temperatures, and panel D from a custom cross-validation that divides the data into pre-2005 and 2005-2010 samples. All results are shown using residualized data in which $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects were removed from all variables before cross-validation was conducted. See text for details.

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

First, we conduct a standard k-fold cross-validation analysis with 10 folds. Because of the panel structure of our data and because we use $ADM1$ level climate and income variables to determine mortality sensitivity to temperature (see Equation 4), we spatially block when defining these ten folds, ensuring that all $ADM2 \times year \times age$ observations that fall within the same $ADM1$ are removed from the sample in the same fold. This ensures that the strong serial and spatial correlation between observations within an $ADM1$ does not artificially inflate our measure of out-of-sample performance. These results are shown in Panel B of Table D.5, which reports the root mean squared error (RMSE) from our main interaction model (Equation 4), a benchmark model that does not account for any form of adaptation (the "no adaptation model"), and two alternative models that account for just one form of adaptation (the "income benefits model" and the "climate adaptation model"). These results reveal that our interaction model out-performs the no-adaptation benchmark model (RMSE decline of 12.62, or 2.2). Moreover, our interaction model obtains high out-of-sample performance when compared to the in-sample RMSE shown in Panel A (RMSE increase of 0.15, or 0.03%). The two partial adaptation models also out-perform the benchmark model, both in- and out-of-sample.

D.7.2 Cross-validation using blocking by covariate values

Second, we design a custom cross-validation analysis that systematically removes blocks of data based on long-run climate and income, the two covariates determining mortality sensitivity to temperature in Equation 4. This exercise is designed to mimic the spatial extrapolation we conduct to generate mortality-temperature relationships in locations without mortality data, based solely on their long-run climate and income (see Section 5.2 for details). To conduct this analysis, we split our sample into 9 "blocks", based on the tercile of the long-run climate and income distributions that each observation falls into.[40] For example, one block corresponds to observations in the lowest income tercile and the coldest climate tercile, while another block corresponds to observations in the middle income tercile and the hottest climate tercile. We then repeatedly remove each block of the data, re-estimate Equation 4 using the remaining 8 blocks of data, and generate out-of-sample predictions for the removed block. In Panel C we show out-of-sample RMSE results for each block, as well as for the full sample.

We draw three main conclusions from the results in Panel C of Table D.5. First, the interaction model overall performs well when predicting mortality rates in locations where income and climate fall outside the estimation sample ranges. The full-sample out-of-sample RMSE from this exercise is only slightly higher than the in-sample RMSE (RMSE increase of 0.47, or 0.08%), which relies on the full dataset for estimation. Additionally, this exercise produces only a slightly higher RMSE than the standard k-fold cross-validation shown in Panel A (RMSE increase of 0.32, or 0.06%). Note that this climate and income blocking exercise is much more challenging than k-fold cross-validation, as entire portions of the income and climate distributions are omitted from estimation. Second, the interaction model out-performs the no-adaptation benchmark model overall (RMSE decline of 12.37, or 2.1%), and in 7 out of 9 blocks of the data, encompassing 100%

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

of global population at 2100. This indicates that the model in Equation 4 provides more accurate mortality rate predictions than a model ignoring heterogeneity based on income and long-run climate. Notably, the interaction model performs well in the low income and hot climate block, conditions that will be experienced by 46% of the global population in 2100 under SSP3. We further demonstrate predictive power in hot and low income conditions in another out-of-sample exercise described in Section D.8 below. Finally, the partial adaptation models, which include only income or only climate as determinants of heterogeneity in the mortality-temperature relationship, also out-perform the benchmark model in this exercise, both overall and in most subsets of the data.

D.7.3 Cross-validation using temporal blocking

Our third out-of-sample analysis assesses our model's ability to predict mortality and mortality-temperature sensitivity into the future. To conduct this analysis, we split our sample into observations that fall before the year 2005, and those after (and including) 2005. We choose this cutoff because it ensures a roughly balanced set of countries in both samples (see Figure B.1). We then use the pre-2005 data to estimate the interaction model in Equation 4, and use these regression results to predict mortality rates in the 2005-2010 data. Results are shown in Panel D of Table D.5. These RMSE comparisons show that our interaction model performs well when predicting residualized mortality rates in years that were not used in estimation, showing little RMSE increase relative to an in-sample estimation using 2005-2010 data only (RMSE increase of 0.97, or 0.17%). Consistent with the prior two tests, the interaction model out-performs a model without adaptation in this out-of-sample test (RMSE decline of 14.36, or 2.5%), as do both partial adaptation models.

D.7.4 Visualizing out-of-sample performance

In addition to the tabular results shown in Table D.5, here we show visualizations of the predictive performance of the interaction model for the two cross-validation approaches based on spatial or temporal blocking. Specifically, we display the difference between out-of-sample predicted mortality-temperature response functions and estimates of those same response functions using a subset of the sample data. While the tabular results of out-of-sample performance in Table D.5 indicate overall measures of fit, these figures help visualize when and to what extent the model accurately extrapolates mortality-temperature sensitivity in different portions of the sample. This exercise differs from the RMSE exercise shown above in that the comparison is between two estimated response functions, neither of which is known with certainty. In contrast, the RMSE table reports error between the true mortality rate observations and predicted mortality rates.

First, Figure D.9 plots the differences in the >64 mortality-temperature relationship between an out-of-sample prediction and an "in-sample" estimation using a subset of the data, for four of the 9 blocks in the spatial blocking exercise described above. The solid red line indicates the difference between: *(i)* the mortality-temperature response function predicted for each block (e.g., high income and cold climate in the upper left) by the estimation of Equation 4 using the remaining 8 blocks of data, and *(ii)* the mortality-temperature response function estimated using a model without interactions with data from within that

block alone. 95% confidence intervals on the difference are shown in the shaded red area,[41] and vertical dashed lines indicate the middle 99% of the daily temperature distribution within each block.[42]

Figure D.9 demonstrates that the differences between in-sample and out-of-sample response functions are rarely statistically significant for the >64 age group (these results are similar for other age groups). Statistically significant differences only arise at the extreme cold end of the temperature distribution, often in blocks where those temperatures are rarely realized (e.g., cold temperatures in the hot climate blocks). Moreover, for the majority of observed temperatures, the magnitude of these differences are small relative to the overall mortality-temperature response function (see main text Figure 1).

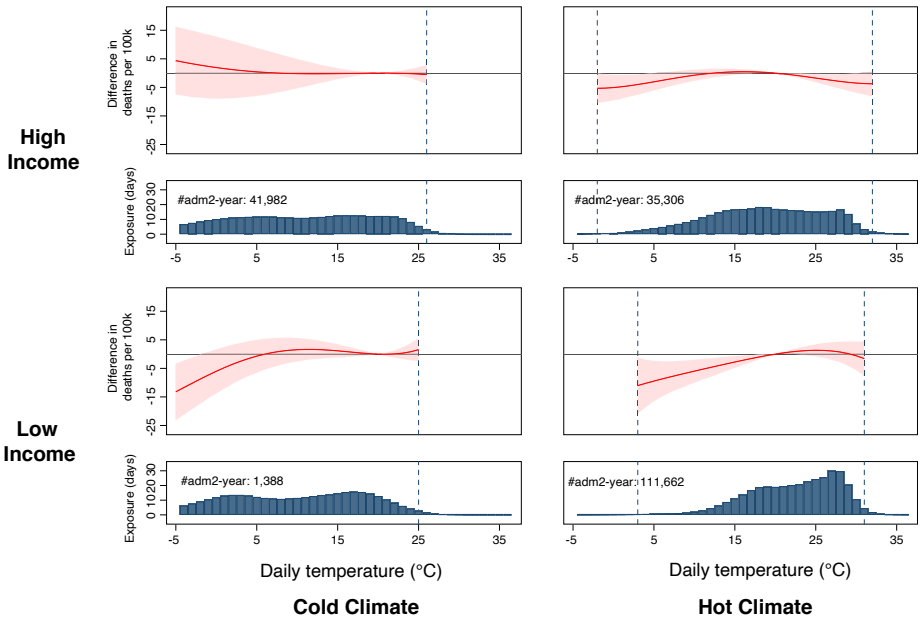

Figure D.9: Differences between spatial out-of-sample predicted mortality-temperature response functions and in-sample estimated response functions (age >64mortality rate). Each panel shows the difference between *(i)* an out-of-sample predicted mortality-temperature relationship resulting from the estimation of a version Equation 4 using all data *except* observations falling within the income and climate "block" indicated (e.g., high income and cold, low income and hot), and *(ii)* an in-sample estimated mortality-temperature relationship resulting from the estimation of a similar model without interactions using data from within the income and climate block alone. Positive values indicate the out-of-sample prediction overestimates mortality sensitivity to temperature relative to an in-sample estimation (and vice versa for negative values). Histograms show the distribution of daily temperature for all locations falling within the indicated block and vertical dashed lines indicate the middle 99% of the daily temperature distribution within each block. 95% confidence intervals are computed by running the out-of-sample and in-sample regressions in a stacked regression model saturated with block-level indicators. Results for other age groups are similar.

Analogous to Figure D.9, the temporal extrapolation performance of our interaction model is shown in Figure D.10, which plots the differences in the >64 mortality-temperature relationship between an out-of-

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

[42] Note that vertical dashed lines are omitted from Figure D.9 if the 0.5^{th} or 99.5^{th} percentile of the data do not fall between -5°C and 35°C.

sample prediction and an in-sample estimation for the lowest and highest income terciles of the full sample income distribution. The solid red line indicates the difference between: *(i)* the mortality-temperature response function predicted for each income tercile by the estimation of Equation 4 using pre-2005 data only, and *(ii)* the mortality-temperature response function estimated using a model without interactions with data from 2005-2010 only. The 95% confidence intervals on the difference are shown in the shaded red area,[43] and vertical dashed lines indicate the 99.5^{th} percentile of the daily temperature distribution within each income group. As above, this exercise differs from results shown in Table D.5 in that the comparison is between two estimated response functions, neither of which is known with certainty.

Figure D.10 demonstrates that, for the >64 age group, adaptation to heat is over-estimated in low income regions (left panel) and under-estimated in high income regions (right panel). That is, using data from before 2005 leads to predicted mortality sensitivity to temperature in 2005-2010 that is too *low* in low income regions and too *high* in high income regions, relative to an in-sample estimate. These differences are statistically significant for the low income group and of modest size (∼5 deaths per 100,000 per 30°C day). In contrast, differences between out-of-sample predicted response functions and in-sample estimated response functions are small and statistically insignificant for other age groups. To show the implications of both under- and over-estimating rates of adaptation, we show in Section F.4 sensitivity of our projected mortality impacts of climate change to assumptions about the rate of adaptation over the 21^{st} century.

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

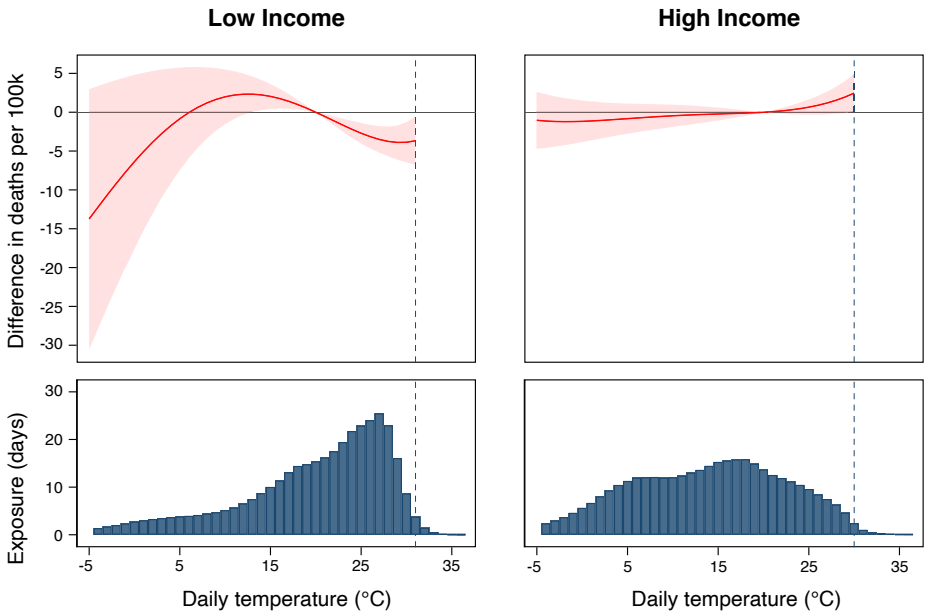

Figure D.10: Differences between temporal out-of-sample predicted mortality-temperature response functions and in-sample estimated response functions (age >64 mortality rate). Each panel shows the difference between *(i)* an out-of-sample predicted mortality-temperature relationship resulting from the estimation of Equation 4 using on all data before 2005, and *(ii)* an in-sample estimated mortality-temperature relationship resulting from the estimation of a similar model without interactions using data from 2005-2010 alone. Histograms show the distribution of daily temperature for all locations falling within the indicated block and vertical dashed lines indicate the 99^{th} percentile of the daily temperature distribution within each block. 95% confidence intervals are computed by running the out-of-sample and in-sample regressions in a stacked regression model saturated with block-level indicators. Results for other age groups show smaller and statistically insignificant differences between out-of-sample and in-sample response functions.

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

As discussed above in Appendix D.7, the accuracy of the spatial and temporal extrapolation of response functions conducted in the main text depends in part on the representativeness of the observed sample. In the ideal case, we would have data for countries that cover the full distribution of income and climate, but as shown in Figure 2 in the main text, our observed sample lacks coverage for the poorest and hottest regions of the global income-climate distribution. The results in Appendix D.7 provide some confidence that we are able to extrapolate our estimates to the poorer and hotter regions of the world within our sample. However, a reasonable concern is that countries with subnational, age-specific mortality rate data may be different than the countries without such data. These are often developing countries where mortality data are either inconsistently collected or not collected at all. To address this concern, here we expand upon the cross-validation experiments shown above to test how our model performs in a region that is both lower income and hotter than our estimation sample. India represents the poorest and hottest country for which we have been able to obtain mortality records, and therefore provides an important check on the extrapolation

performance of our interaction model.

To execute this validation test, we use data from Burgess et al. (2017) and compare mortality-temperature relationships estimated using these data to those predicted for India from our main estimating equation, Equation 4. The primary reason that we do not use data from India in our main estimation is that they do not contain age-specific mortality rates, and we show that age at death is a key source of heterogeneity (see Figure D.3).

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

$$M_{it} = f(T_{it}) + q(R_{it}) + \alpha_i + \lambda_r^1 t + \lambda_r^2 t^2 + \varepsilon_{it}. \tag{D.23}$$

The outcome M_{it} is the all-age mortality rate for district i in year t, which we estimate in levels rather than in logs, as in Burgess et al. (2017), to ensure direct comparability with our main specification. To generate comparable functional forms, we estimate Equation D.23 using a fourth-order polynomial, denoted by $f(T_{it})$, as we have used this as our main specification (Burgess et al. (2017) use a binned temperature specification). We control for precipitation, denoted by $q(R_{it})$, identically to Burgess et al. (2017) via a set of three dummy variables, each of which takes the value of 1 when total annual rainfall in district i and year t falls within each of three location-specific rainfall terciles. Due to the negligible effect of precipitation on our estimates of the mortality-temperature relationship (see Figure D.7), this choice makes little difference. Following Burgess et al. (2017), we also include a set of district fixed effects, α_i, and linear and quadratic trends for each "climate region" (of which the authors note there are four separate regions) r of India. Observations are weighted by district population and standard errors are clustered at the ADM2 level.

We then use the estimates from Equation 4 to predict the mortality-temperature relationship across India. To do so, we first predict mortality-temperature response functions for India for each of the three age groups in our main estimation, using the population-weighted average values of ADM1-level incomes and average temperatures across the country. Once we have the predicted age-specific national responses, we take the age-weighted average of these response functions to generate an all-age average mortality-temperature response function across India. We compute this average using age-specific population values from the year 2015, which are available in the Burgess et al. (2017). We cluster standard errors at the ADM1 level (in India, this is equivalent to the state level), as in all our specifications throughout the main text. For the purposes of assessing out-of-sample performance of our main model as it compares to alternative models estimated in the literature, we also predict an all-India response function using estimates for a version of Equation 4 that models only heterogeneity in long-run income and one that models only heterogeneity in long-run average temperature.

Figure D.11 shows the result of this replication and out-of-sample validation exercise for India. The figure compares our predicted responses in India (in blue) to the mortality-temperature response estimated using India's data alone (in red), following Burgess et al. (2017). Our model performs well, despite containing no information on Indian mortality rates: for the hotter end of the response function, where much of the low income world resides, our prediction is, if anything, conservative in extrapolating out-of-sample. Included in the figure are two dashed blue lines which show the predicted mortality-temperature relationship using

estimates from models with only one or the other of our two interaction terms. The model using estimates from Equation 4 replicates the country model more closely than both alternative models.[44]

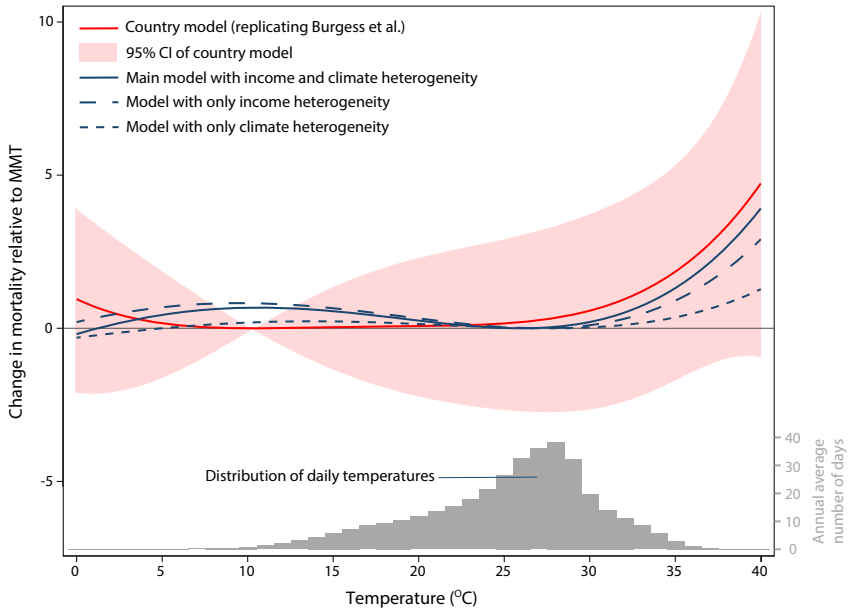

Figure D.11: Out-of-sample validation of the mortality-temperature response function in India. Dark blue lines indicate out-of-sample predicted response function using coefficients from the interaction model in Equation 4, as well as versions interacting with only income (long-dashed line) and only climate (short-dashed line). The red line is estimated following Burgess et al. (2017) using all-age mortality data for India, as described in the text. The relative congruence between red and solid dark blue lines shows that our interaction model generates reasonable predicted response functions in the poorest and hottest regions of the world, the subset of the covariate space for which the main estimating sample has the least representation. Note that all curves are centered on their respective minimum mortality temperatures, as we use these curves to compute predicted deaths below in Table D.6, and all predicted deaths reported in the paper use location- and model-specific minimum mortality temperatures.

In addition to the comparison of response functions in Figure D.11, Table D.6 reports quantitative differences in predicted mortality rates across the four models. This table summarizes the quantitative difference between the response functions shown in Figure D.11. Mortality rates are predicted in-sample for each model by taking the product of the historical distribution of temperature exposure with the response function.[45] Table D.6 shows that predicted mortality rates follow the same pattern as the response functions in Figure D.11, with all out-of-sample model predictions falling below Burgess et al. (2017) predicted mortality rates. These out-of-sample predictions under-estimate in-sample predicted mortality rates by 37% (full interaction model in Equation 4), 42% (interaction model with income only), and 78% (interaction model with long-

[44]Note that in Figure D.11, each response function is centered such that the predicted change in mortality is zero at the value of its minimum mortality temperature (MMT). This recentering is arbitrary, as the inclusion of fixed effects in Equation D.23 implies that the level of the response function is not recoverable; only slopes are causally identified. However, we choose this approach to match the quantitative exercise shown in Table D.6, which uses the MMT as the reference temperature in order to mimic the climate change projection approach we take throughout the main text.

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

run temperature only). However, all predicted values fall within the large confidence interval given by the Burgess et al. (2017) results. We conclude that our estimates compare relatively well to an in-sample model in India, although our out-of-sample values tend to be conservative.

Model	Predicted mortality	Lower 95% CI	Upper 95% CI	Difference	% difference
Burgess et al. (2017)	119.67	-867.85	1107.20	0.00	0.00
Full interaction (Eq. 4)	74.87	-8.26	158.00	-44.80	-37.44
Income-only interaction	69.22	-6.96	145.40	-50.45	-42.16
Climate-only interaction	26.15	-4.37	56.66	-93.53	-78.15

Table D.6: Evaluation of differences in predicted mortality rates when using an in-sample estimation for India versus an out-of-sample predicted response function. Each row refers to a different empirical model of the mortality-temperature relationship in India. The Burgess et al. (2017) model in the first row is estimated following Burgess et al. (2017) using all-age mortality data from India. In all three remaining rows, the mortality-temperature relationship is predicted from a form of Equation 4, relying on data from 40 other countries; no Indian data are used. Predicted mortality represents total deaths per 100,000 per year that are attributable to historical temperature variation. Differences indicate the difference between out-of-sample predicted mortality rates from models in rows 2-4, relative to the model in row 1. Out-of-sample predicted mortality in rows 2-4 are smaller than predicted mortality with the India data, but within the confidence interval.

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

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

E.1 Determining the temporal dynamics of income effects

We estimate the relationship between long-run average climate, average income, and mortality-temperature sensitivity via the estimation of Equation 4 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, we apply the estimated coefficients from Equation 4 over time, allowing impact region response functions to evolve as the climate warms and incomes grow. To do so, we must make an assumption regarding the rate at which the income and average climate covariates update. Here, we detail how we define this speed of adjustment in the case of income growth. While we can derive a duration over which updating occurs in the case of income due to substantial time series variation in incomes in our observed data, the historical trends for temperature have been small to date, making a similar derivation infeasible. Thus, for the case of updating based on long-run average climate, we use the standard definition of "climate" and assume a duration of 30 years.

In future projections, we estimate impact region response functions using time-varying measures of $\log(GDPpc)_{rt}$ (see Section 5.2 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:

$$M_{ait} - M_{ai,t-1} = \boldsymbol{\beta}_t \left[\boldsymbol{T}_{it} - \boldsymbol{T}_{i,t-1}\right] + q_a(\boldsymbol{R}_{it}) + \varepsilon_{it} \tag{E.24}$$

where M_{ait} is the mortality rate for region i in period t and age group $a > 64$, $\boldsymbol{T}_{it}$ is the vector of polynomials of daily average temperatures (up to the fourth order), $\boldsymbol{R}_{it}$ is the vector of cumulative monthly precipitation (up to the second order), as in the main text (see Equation D.21). 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 4, respectively. τ_p is a hyper-parameter which controls the rate of pooling of the data; if it is 0, inverse-variance weighting is used across individual year estimates.

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

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

As discussed in Section 5.2, we impose two assumptions when applying our econometrically-derived model of adaptation to generate projections of future climate change. These assumptions are guided by economic theory as well as the physiological literature and are used to ensure plausible out-of-sample projections over the 21^{st} century. Graphical intuition for these constraints is shown in Figure E.1.

Assumption #1: Weak monotonicity. A large body of epidemiological and econometric literature has recovered U-shaped relationships between mortality rates and daily temperatures, where both extreme cold and extreme heat increase the risk of death. These parabolic response functions have been recovered in studies using a wide range of functional form assumptions (e.g., binned daily temperatures, restricted cubic splines, or polynomials) and across diverse locations globally (e.g., Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011). As shown in Section 4, we also recover U-shaped relationships

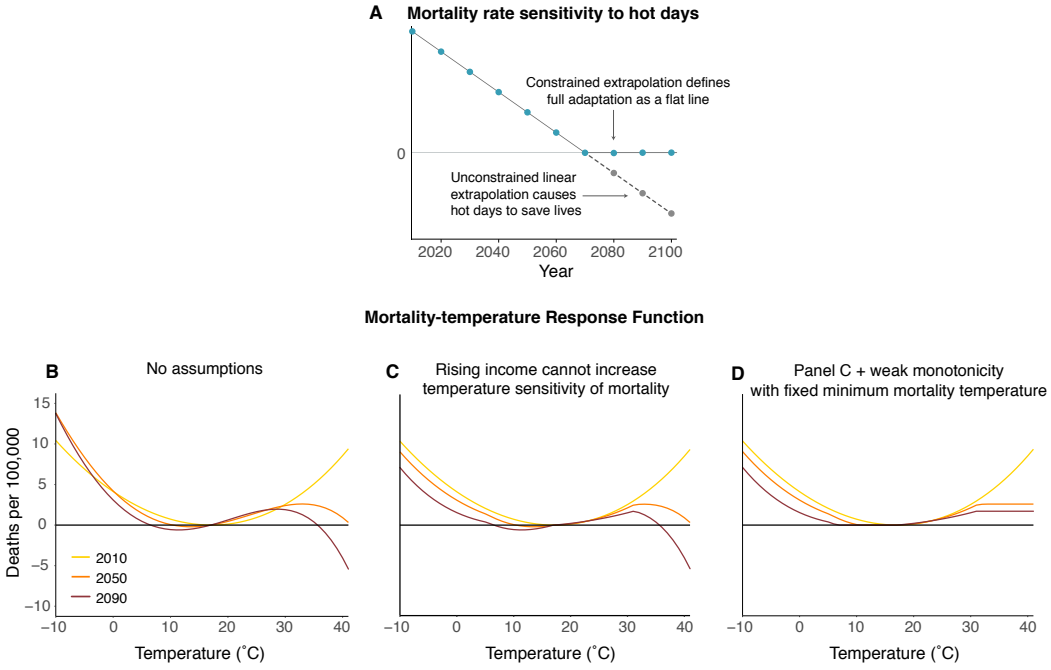

Figure E.1: Two assumptions imposed in climate projections ensure that full adaptation is defined as a flat-line response function and that responses conform to basic physical and economic constraints. Panel A demonstrates heuristically the importance of imposing assumptions on the shape of response functions under adaptation over the 21^{st} century. As shown, linearly declining mortality rate sensitivity to hot days occurs over the course of the century as populations adapt. However, linear extrapolation can lead to mortality benefits on hot days, as shown with the dashed line and grey dots. Our assumptions (shown in teal) ensure that full adaptation is realized when hot days impose zero additional mortality risk. Panels B through D represent an empirical example of how the imposition of these constraints can change the shape of the adapted response function, for the Chicago, Illinois impact region. Panel B has no assumptions, panel C imposes the assumption that income is weakly protective, and panel D imposes the assumption of weak monotonicity around a time-invariant minimum mortality temperature (MMT).

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

To implement this assumption, we first identify a range of physiologically optimal temperatures. Drawing on extensive research across epidemiology and medicine (e.g., Curriero et al., 2002; Guo et al., 2014), as well as ergonomics (e.g., Seppanen, Fisk, and Lei, 2006; Hancock, Ross, and Szalma, 2007), we let this range of possible minimum mortality risk cover the temperatures 10°C to 30°C. We then search, within this range, for the temperature at which the location-specific response function in each impact region r in the baseline years of 2001-2015 is minimized. Because distinct populations may differ substantially in the temperature at which mortality is minimized,[46] it is important to note that we allow these minimum mortality temperatures

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

(MMTs) to be spatially heterogeneous. With these optimal temperatures in hand, we impose the assumption that mortality rates must remain weakly increasing in daily temperatures to both the left and the right of this minimum. To operationalize this, we calculate impacts along an adjusted response function that is defined as the cumulative maximum to the right and left of the minimum mortality temperature along each region- and year-specific response function derived from our response surface estimated in Equation 4. Consistent with prior literature (Heutel, Miller, and Molitor, 2017; Curriero et al., 2002; Gasparrini et al., 2015), we find that these minimum mortality temperatures are highly correlated both with both long-run average temperature (positively) and with income (negatively).

This assumption is important because Equation 4 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. However, such behavior is inconsistent with a large body of epidemiological and econometric literature recovering U-shaped mortality-temperature relationships under many functional form assumptions and in diverse locations (Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011), as well as what we observe in our data. As a measure of its role in our results, the weak monotonicity assumption binds for the >64 age category at 35°C in 9% and 18% of impact regions in 2050 and 2100, respectively.[47] To avoid this unrealistic behavior, we impose weak monotonicity. An example of this assumption in practice is given in panel E of Figure E.1.[48]

In imposing the weak monotonicity constraint, we fix the MMT at its baseline level in 2015 for each impact region. We do so because the use of spatial and temporal fixed effects in Equation 4 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.[49]

Assumption #2: Rising income cannot increase the temperature sensitivity of mortality. Because increased income per capita strictly expands the choice set of individuals considering whether to make adaptive investments, it should not increase the effect of temperature on mortality rates. Consistent with this intuition, we find that income is protective against extreme heat for all age groups. However, for some age groups, the estimation of Equation 4 recovers statistically insignificant but positive effects of income on mortality sensitivity to extreme cold (Table D.1). Therefore, we constrain the marginal effect of income on temperature sensitivity to be weakly negative in future projections, although we place no restrictions on the cross-sectional effect of income when estimating Equation 4.[50] 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 E.1.

different countries.

[47] 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.

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

[49] 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.

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

A visual example of the influence of these constraints can be seen for one example impact region (Chicago, Illinois) in Figure E.1. 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.

E.3 Projected benefits of adaptation

Figure E.2 provides an examination of projected changes in the mortality-temperature relationship over time, which is a key ingredient for projections of the mortality effects of climate change. We plot the spatial distribution of the *change* in the mortality-temperature relationship evaluated at 35°C between 2015 and 2050 (panel A) and 2015 and 2100 (panel B) for the >64 age category. Specifically, these values are calculated as:

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

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

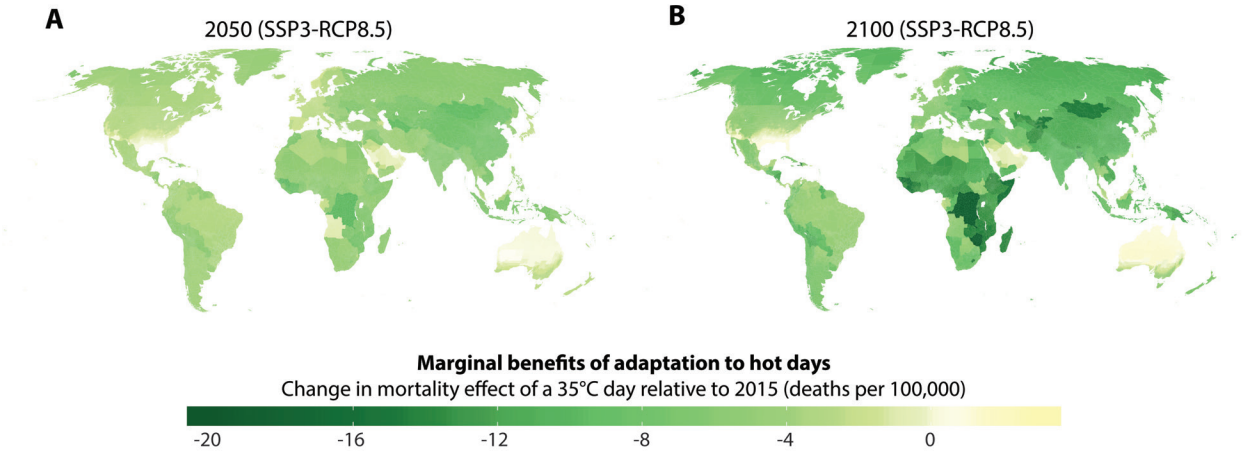

Figure E.2: Spatial and temporal heterogeneity in temperature sensitivity. Panels A and B indicate the change in mortality sensitivity to hot days (35°C relative to a location-specific minimum mortality temperature) for the oldest age category (>64) between 2015 and 2050 (A), and between 2015 and 2100 (B). Specifically, these values are $\hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2050}, \log(GDPpc)_{r,2050}) - \hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$ in panel A and $\hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2100}, \log(GDPpc)_{r,2100}) - \hat{g}_a(\boldsymbol{T}_{35}, TMEAN_{r,2015}, \log(GDPpc)_{r,2015})$ in panel B, where $\boldsymbol{T}_{35}$ is a fourth order polynomial for a daily temperature of 35°C and where the age group is $a > 64$. Darker colors signify larger predicted adaptation to heat. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

F Climate change projections: Additional results and robustness

This appendix provides additional illustrations of the main climate change projection results used and discussed throughout the main text (i.e., Section 5.3), as well as a robustness check and sensitivity analysis regarding the functional form of the mortality-temperature relationship, different assumptions about the behavior of the relationship outside of the historical sample values, and assumptions regarding the rate of adaptation.

F.1 Additional climate change projection results

Uncertainty in global estimates of climate change impacts Table 2 in the main text reports estimates of the mortality effects of climate change and the full mortality risk of climate change at global scale, including interquartile ranges (IQRs) that reflect uncertainty in both the climate and in the econometric estimation of the relationship between mortality and temperature. In Table F.1, we show how these two drivers of uncertainty each contribute to overall uncertainty for each of the estimates in main text Table 2. The first row shows the mean estimate of the mortality effects of climate change (based on Equations 2', 2a', 2b') and the full mortality risk of climate change (based on Equation 3') across the set of Monte Carlo simulations described in Section 5.2, which account for both climate model and econometric uncertainty. The *full uncertainty IQR* indicates the interquartile range (IQR) across all these Monte Carlo simulations. These means and IQRs are identical to the values shown in panel A of main text Table 2.

The row titled *full uncertainty* σ in Table F.1 indicates the standard deviation across the full set of Monte Carlo simulations. The row titled *econometric uncertainty* σ indicates the standard deviation across a set of Monte Carlo simulations in which a single climate model, CCSM4, was combined with 495 random draws from the empirical multivariate normal distribution characterized by the covariance between all of the parameters from the estimation of Equation 4. This standard deviation provides a measure of the importance of uncertainty in the estimated relationship between mortality and temperature, including uncertainty regarding how the mortality-temperature relationship is influenced by long-run climate and by income. The row titled *climate uncertainty* σ indicates the standard deviation across a set of Monte Carlo simulations in which all 33 climate models in the surrogate/model mixed ensemble (SMME) were combined with the point estimates from the estimation of Equation 4. This standard deviation provides a measure of the importance of climate model uncertainty. Because our estimated relationship between mortality and temperature is non-linear, the total variance in mortality effects of climate change will not equal the sum of the two component variances (Hsiang et al., 2017).

Results in Table F.1 show that for all measures of the mortality effects of climate change (columns 1-4), climate uncertainty is comparable to or larger than econometric uncertainty. When the full mortality risk of climate change is reported in death equivalents (column 6a), climate uncertainty continues to dominate econometric uncertainty. However, because uncertainty from different sources affects regions of the globe and age groups differentially, whether and how the full mortality risk of climate change is monetized influences the results of this decomposition. Column 6b of Table F.1 shows that econometric uncertainty is far more influential than climate uncertainty in driving overall uncertainty in the monetized value of the full mortality

risk of climate change (shown as % of GDP in 2100) when our preferred valuation approach is used (see Section 3 for details). This finding suggests that climate uncertainty is particularly influential in poor regions of the globe, where the value of life-years lost is relatively low (due to a VSL that scales with income), but where projected mortality rate impacts are relatively high. Indeed, we find that the climate uncertainty standard deviation of the full mortality risk of climate change when measured in deaths per 100,000 is 9 times larger in the poorest tercile of impact regions than it is in the wealthiest tercile, while the econometric uncertainty standard deviation is just 1.7 times larger.

Table F.1: Global estimates of the full mortality risk of climate change and associated uncertainty in 2100 (high emissions scenario, RCP8.5)

	No income growth or adaptation	**Benefits of income growth**	**Benefits of climate adaptation**	**Mortality effects of climate change**	**Costs of climate adaptation**	**Full mortality risk of climate change**	
	Eq. 2a' *deaths/100k* (1)	Eq. 2b' - Eq. 2a' *deaths/100k* (2)	Eq. 2' - Eq. 2b' *deaths/100k* (3)	Eq. 2' *deaths/100k* (4)	Eq. 7 *deaths/100k* (5)	Eq. 3' *deaths/100k* (6a)	*% of GDP* (6b)
Mean impacts	220.6	-116.5	-31.0	73.1	11.7	84.8	3.2
Full uncertainty IQR	[76.4, 258.8]	[-149.4, -39.2]	[-60.1, 3.8]	[5.6, 101.4]	[0.2, 19.4]	[17.4, 116.4]	[-5.4, 9.1]
Full uncertainty σ	239.8	145.0	83.4	114.7	19.8	118.1	9.6
Econometric uncertainty σ	59.9	40.9	23.4	48.6	11.9	47.2	5.4
Climate uncertainty σ	193.9	103.4	44.6	46.3	16.2	62.3	1.2

Table shows projections of the mortality effects of climate change and the full mortality risk of climate change across all age categories. Mean estimates are averages across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. The *full uncertainty IQR* indicates the interquartile range (IQR) across all Monte Carlo simulations, accounting for both sources of uncertainty. The *full uncertainty* σ indicates the standard deviation across these same simulations. The *econometric uncertainty* σ indicates the standard deviation across a set of Monte Carlo simulations in which a single climate model, CCSM4, was combined with 500 random draws from the empirical multivariate normal distribution characterized by the covariance between all of the parameters from the estimation of Equation 4. This standard deviation provides a measure of the importance of statistical uncertainty. Finally, the *climate uncertainty* σ indicates the standard deviation across a set of Monte Carlo simulations in which all 33 climate models in the surrogate/model mixed ensemble (SMME) were combined with the point estimate from the estimation of Equation 4. This standard deviation provides a measure of the importance of climate model uncertainty. See Tables 2 and F.2 for detailed descriptions of each column. All estimates shown rely on the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

Alternative measures of climate change impacts In Figure 4 of the main text, we show a map of impact region-level mean estimates of the mortality effects of climate change, following Equation 2'. However, in Sections 5 and 6.1 we also define three other measures of expected climate change impacts: (i) mortality effects of climate change without income growth or adaptation (Equation 2a'); (ii) mortality effects of climate change without adaptation (Equation 2b'); and (iii) the full mortality risk of climate change, accounting both for adaptation and income benefits as well as adaptation costs (Equation 3'). Panels A, B and D in Figure F.1 below show projected impacts for each of these alternative measures; for comparison, panel C repeats the mortality effects of climate change map from the main text.

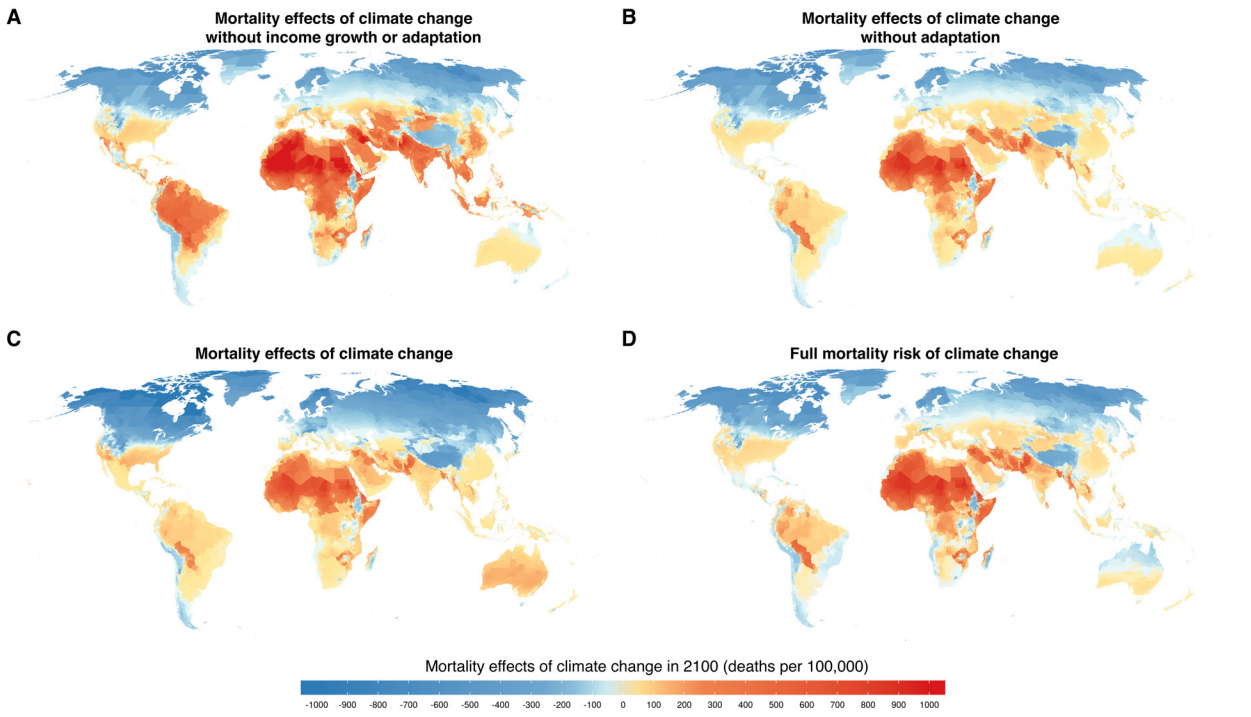

Figure F.1: The mortality effects of climate change under alternative adaptation scenarios and the full mortality risk of climate change. All maps show predicted mortality effects of climate change and colors in each impact region represent the mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Panel A shows estimates of the mortality effects of climate change without income growth or adaptation to climate change (Equation 2a'). Panel B shows estimates of the mortality effects of climate change without adaptation (Equation 2b'). Panel C shows estimates of the mortality effects of climate change (Equation 2'), which accounts for changes in the mortality sensitivity to temperature in response to both future income and future climate. Panel D shows the full mortality risk of climate change (Equation 3'), which is the sum of the mortality effects of climate change and inferred adaptation costs from Equation 7, measured in units of "death equivalents" (see main text for details). 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.

Section 5 also presents a time series of the mortality effects of climate change using measures in Equations 2', 2a', and 2b'. Figure F.2 adds estimates of the full mortality risk of climate change (Equation 3') to the same figure, showing the mean estimate over time, as well as the uncertainty surrounding this mean, as captured by Monte Carlo simulations.

Finally, Figure F.3 presents the same time series of aggregate global mortality consequences of climate change for measures (i) and (ii), as in Figure F.2, but adds shading to indicate the uncertainty surrounding the mortality effects of climate change with benefits of income growth. Just as in Table D.1, the uncertainty in the estimated relationship between income per capita and the temperature sensitivity of mortality is apparent.

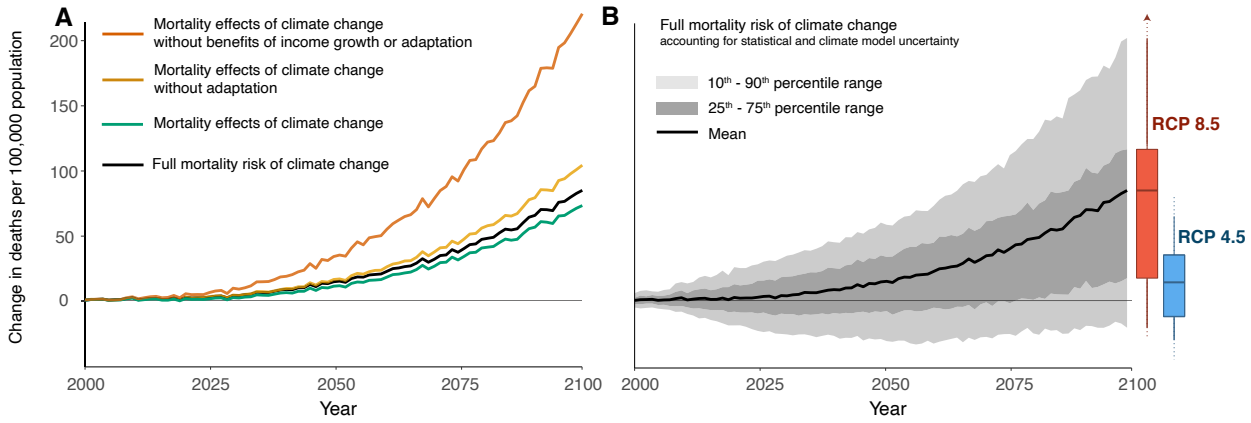

Figure F.2: Time series of projected full mortality risk of climate change. All lines show projected mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents one of three measures of the mortality effects of climate change. Orange (Equation 2a'): mortality effects of climate change without income growth or adaptation. Yellow (Equation 2b'): mortality effects of climate change without adaptation. Green (Equation 2'): mortality effects of climate change. Black (Equation 3'): full mortality risk of climate change, calculated as the sum of the mortality effects of climate change 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 (equivalent to the black line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of full mortality risk of climate change 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.

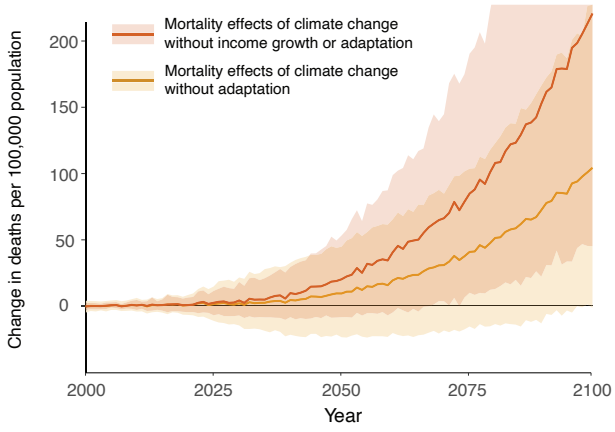

Figure F.3: Uncertainty in the mortality effects of climate change including benefits of income growth. Both solid lines show measures of the mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Shaded areas indicate the 10^{th}-90^{th} percentile range of the Monte Carlo simulations. The orange line and confidence interval shows the mortality effects of climate change without income growth or adaptation (2a'), while the yellow line and confidence interval shows the mortality effects without adaptation, but with income growth (Equation 2b'). Both projection estimates shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

Climate change projections by age group In the main text, Figure 5 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 F.4, we show each of these age-specific projections for SSP3 and RCP8.5 (for reference, main text Table 1 shows that the average mortality rate for the oldest age group is 4,736 deaths per 100,000 in our estimation sample). While all age groups have a mean estimate that is above zero by end-of-century, the oldest age group dominates our projections in terms of death rates. These large demographic differences are taken into account in our valuation steps (see Section 7 and Appendix G).

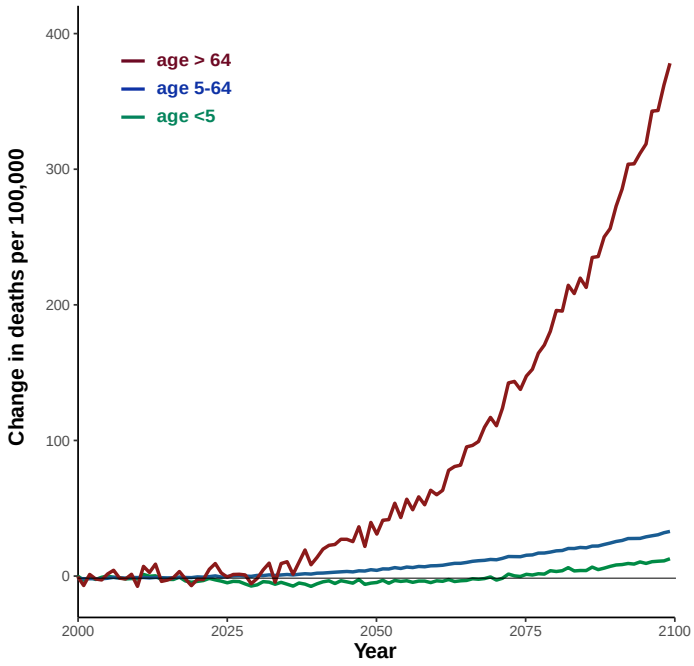

Figure F.4: Heterogeneity in the full mortality risk of climate change by age group. All lines show projections of the full mortality risk of climate change (Equation 3') across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate 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. Figure F.2 represents the sum across these age-specific projections, where death rates are population weighted by age-specific population values.

Climate change projections by socioeconomic scenario Throughout Section 5.3 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 full mortality risk 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).

Gains from mitigation spatially and in aggregate. In the main text, Figure 4 displays the mortality effects of climate change spatially under the socioeconomic scenario SSP3 for the entire globe. Figure F.6 shows a comparison between the mortality effects of climate change under RCP8.5 and RCP4.5, showing

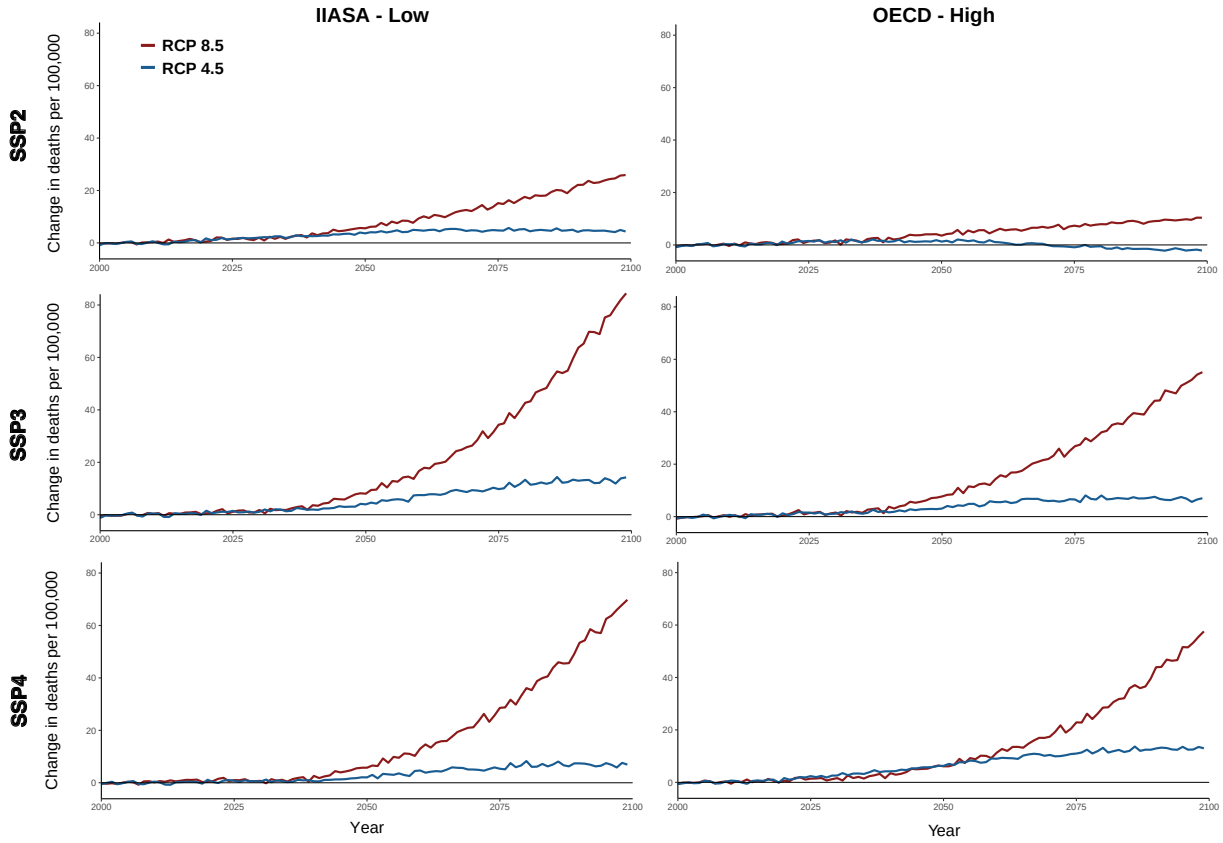

Figure F.5: The full mortality risk of climate change under different scenarios of population growth, economic growth, and emissions. Rows denote different Shared Socioeconomic Pathway (SSP) scenarios, columns denote two separate modeling groups that produce data for each SSP, and each panel shows a time series of the full mortality risk of climate change for RCP 4.5 and RCP 8.5. Both lines indicate the full mortality risk of climate change, accounting for both adaptation benefits and costs following Equation 3', and indicate the mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. RCP8.5 is a high-emissions scenario, while RCP4.5 is a scenario with aggressive emissions reductions. The OECD economic projections tends to exhibit slightly higher income growth than the IIASA economic projections. Throughout the main analysis, projection results relying on IIASA and OECD socioeconomic projections are both used and weighted equally.

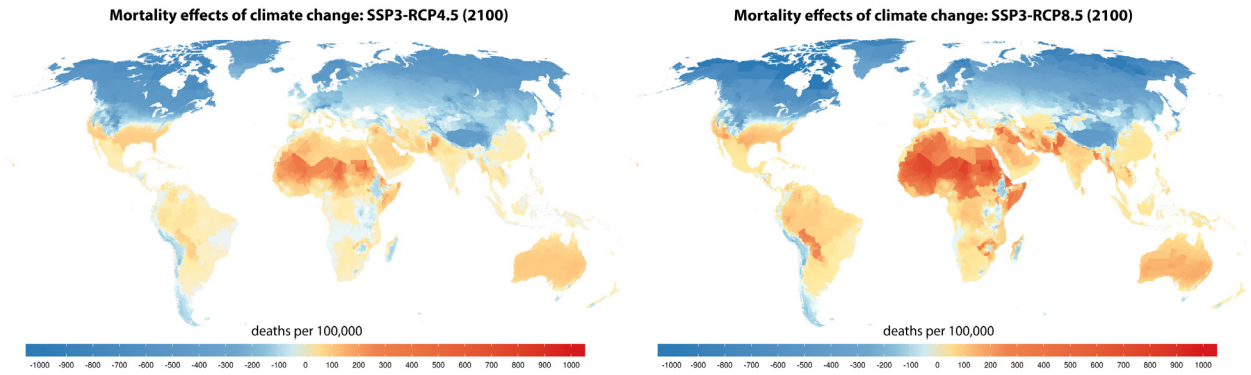

Figure F.6: The mortality effects of climate change under RCP4.5 and RCP8.5 for SSP3. These maps indicate the mortality effects of climate change (i.e., Equation 2'), measured in units of deaths per 100,000 population, in the year 2100. Estimates come from a model accounting for the benefits of adaptation and income growth, and the map shows the climate model weighted mean estimate across Monte Carlo simulations conducted on 33 climate models. All values shown refer to the SSP3 socioeconomic scenario.

the gains from mitigation. As expected, reducing emissions to the level of RCP4.5 is predicted to have substantial benefits in terms of reduced mortality effects of climate change. However, the spatial pattern of impacts remains, with clearly unequal distribution of impacts between places that are relatively poor today versus places that are relatively wealthy. Figure F.7 replicates the aggregate time series (Figure 5 in the main text) for RCP4.5 under the SSP3 scenario, and Table F.2 replicates the aggregate mortality effects and full mortality risk of climate estimates for 2100 (Table 2 in the main text) for RCP4.5 under SSP3.[51] The gains from reducing emissions are evident in both sets of aggregate results.

[51] In Table 2 in the main text and F.2 here, Europe includes the Aland Islands, Albania, Andorra, Austria, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Czech Republic, Denmark, Estonia, Faroe Islands, Finland, France, Germany, Gibraltar, Greece, Guernsey, Hungary, Iceland, Ireland, Isle of Man, Italy, Jersey, Kosovo, Latvia, Liechtenstein, Lithuania, Luxembourg, Macedonia, Malta, Moldova, Monaco, Montenegro, Netherlands, Norway, Poland, Portugal, Romania, Russia, San Marino, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Ukraine, United Kingdom, and Vatican City. Similarly, sub-Saharan Africa includes Angola, Benin, Botswana, Burkina Faso, Burundi, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Cote d'Ivoire, Democratic Republic of the Congo, Djibouti, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Gambia, Ghana, Guinea, Guinea-Bissau, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritania, Mauritius, Mayotte, Mozambique, Namibia, Niger, Nigeria, Reunion, Republic of Congo, Rwanda, Saint Helena, Sao Tome and Principe, Senegal, Seychelles, Sierra Leone, Somalia, South Africa, South Sudan, Sudan, Swaziland, Tanzania, Togo, Uganda, Western Sahara, Zambia, and Zimbabwe.

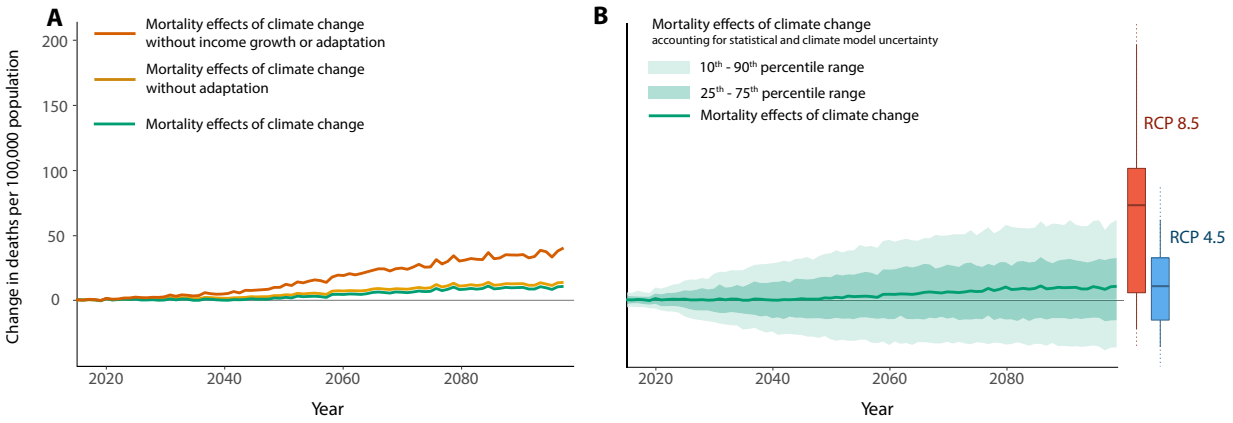

Figure F.7: Time series of projected mortality effects of climate change under RCP4.5 for SSP3. All lines show projected mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents one of three measures of the mortality effects of climate change. Orange (Equation 2a'): mortality effects of climate change without income growth or adaptation. Yellow (Equation 2b'): mortality effects of climate change without adaptation. Green (Equation 2'): mortality effects of climate change. Panel B shows the 10^{th}-90^{th} percentile range of the Monte Carlo simulations for the mortality effects of climate change (equivalent to the green line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of mortality effects of climate change in 2100 under both RCPs. All line estimates shown refer to the RCP4.5 emissions scenario and all line and boxplot estimates refer to the SSP3 socioeconomic scenario.

Table F.2: Global and regional estimates of the full mortality risk of climate change in 2100 (moderate emissions scenario, RCP4.5)

	No income growth or adaptation	Benefits of income growth	Benefits of climate adaptation	Mortality effects of climate change	Costs of climate adaptation	Full mortality risk of climate change	
	Eq. 2a'	Eq. 2b' - Eq. 2a'	Eq. 2' - Eq. 2b'	Eq. 2'	Eq. 7	Eq. 3'	
	deaths/100k	*deaths/100k*	*deaths/100k*	*deaths/100k*	*deaths/100k*	*deaths/100k*	*% of GDP*
	(1)	(2)	(3)	(4)	(5)	(6a)	(6b)
Panel A: Global estimates							
Mean impacts	40.3	-26.5	-3.0	10.7	3.5	14.2	0.6
IQR	[7.8, 57.9]	[-47.8, -2.6]	[-16.3, 8.1]	[-15.5, 32.5]	[-0.9, 7.5]	[-12.3, 35.2]	[-3.9, 4.6]
Panel B: Regional estimates							
China	14.8	-8.6	-12.5	-6.4	9.9	3.4	0.5
USA	-9.0	-1.0	-2.7	-12.7	12.7	-0.1	0.2
India	78.0	-69.6	3.6	12.0	-1.0	11.1	1.5
Pakistan	144.3	-44.3	-18.6	81.4	16.1	97.6	8.0
Bangladesh	77.2	-21.9	-13.3	42.0	7.1	49.1	4.5
Europe	-28.0	6.5	-42.0	-64.3	50.6	-12.8	-0.7
Sub-Saharan Africa	37.6	-16.1	-1.5	20.0	0.5	20.5	1.7

Table shows projections of the mortality effects of climate change and the full mortality risk of climate change across all age categories. Mean estimates are averages across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, brackets indicate the interquartile range (IQR). Columns 1-4 are computed using the three measures of the mortality effects of climate change detailed in Section 5, all in units of deaths per 100,000. Column 1 (Equation 2a'): mortality effects of climate change without benefits of income or adaptation to climate change. Column 2 (Equation 2b' - Equation 2a'): benefits of income growth. Column 3 (Equation 2' - Equation 2b'): benefits of adaptation to climate change. Column 4 (Equation 2', equal to the sum of columns 1-3): mortality effects of climate change. Column 5 shows the mortality-related costs of adaptation inferred using a revealed preference approach (Equation 7 divided by the VSL), measured in "death equivalents". Columns 6a-6b show the full mortality risk of climate change (Equation 3'), measured in deaths per 100,000 (column 6a) and represented as % of 2100 GDP (column 6b) using an age-adjusted value of the U.S. EPA VSL with an income elasticity of one applied to all impact regions. Column 6a is equivalent to the sum of columns 4 and 5. All estimates shown rely on the RCP4.5 emissions scenario and the SSP3 socioeconomic scenario.

The impact of climate change in 2100 under RCP4.5 compared to contemporary leading causes of death. Figure F.8 presents the same results as Figure 9 in the main text, but for RCP4.5. As can be seen, despite the overall decrease in the full mortality risk of climate change under SSP3 and RCP4.5 when compared to RCP8.5, much of the inequality in both the mortality effects of climate change and the adaptation costs that was evident in main text Figure 9 remains.

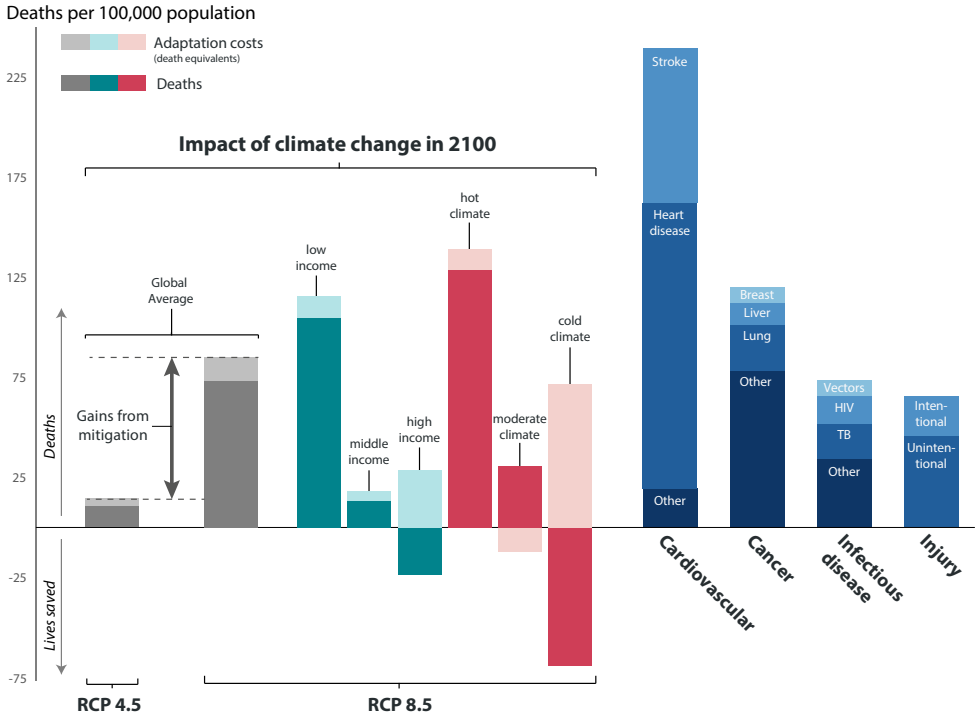

Figure F.8: The impact of climate change in 2100 under RCP4.5 compared to contemporary leading causes of death. Impacts of climate change (grey, teal, and coral) are calculated for the year 2100 for SSP3 and include changes in death rates (solid colors) and changes in adaptation costs, measured in death equivalents (light shading). Global averages for RCP 8.5 and RCP 4.5 are shown in grey, demonstrating the gains from mitigation. Income and average climate groups under RCP4.5 are separated by tercile of the 2015 global distribution across all 24,378 impact regions. Blue bars on the right indicate average mortality rates globally in 2018, with values from WHO (2018).

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

As discussed in Section D.2, we experiment with four distinct nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ in Equations D.21 and 4 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 4. 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 D.4, we show that the binned and fourth order polynomial functional forms recover similar mortality-temperature response functions across our pooled multi-country sample. Below in Figure F.9, we show that this similarity carries through to generate similar climate change impact projections across the binned and polynomial functional forms. Both projections are constructed using estimation of the interaction model in Equation 4 in combination with high-resolution covariates $TMEAN$ and $\log(GDPpc)$ to generate impact region-specific response functions (see Section 5.2 for details).

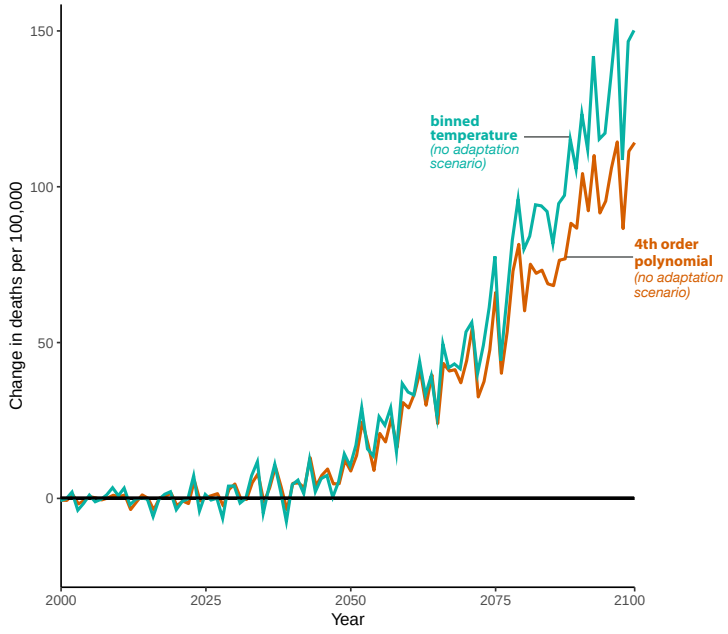

Figure F.9: Robustness of impact projections to alternate functional forms of temperature. Each line represents the time series of changes to the mortality rate due to climate change under the socioeconomic scenario SSP3 and the emissions scenario RCP 8.5. Results shown are for a single climate model (CCSM4). Lines shown refer to estimates of mortality effects of climate change without adaptation or benefits of income growth (Equation 2a'), 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 4. In green is the same object, but with binned daily temperatures used as a functional form in estimation. While the binned regression imposes far fewer restrictions on the regression than does the polynomial, the projected impacts under these two sets of parameterizations are strikingly similar.

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

The paper uses historical data to estimate the mortality-temperature response and uses the results to project the impacts of temperatures in the future. A key challenge, however, is that climate change will cause locations to experience temperatures that have not been observed in the historical record (e.g., see main text Figure 2), thus necessitating out of sample predictions.

Figure F.10 probes the sensitivity of the projected mortality effects of climate change to alternative assumptions about the relationship between mortality and temperature at temperatures that are not observed in available data sets. Specifically, for all temperatures above the maximum and below the minimum daily temperatures within our dataset, we alter the slope of the impact region-specific response functions in two ways. First for "constant out-of-sample extrapolation", we set the marginal effect of temperature fluctuations to equal the value at the maximum if above the maximum temperature, and vice versa for temperatures below the minimum (Figure F.10B). This implies that the response function is flat for all temperatures outside the observed range. For "linear out-of-sample" extrapolation, we set the marginal effect to be linearly increasing in the out-of-sample regions with a slope equal to the slope between the response function evaluated at the maximum (minimum) and the maximum minus 0.1°C (plus 0.1°C) (Figure F.10C). It is apparent that neither of these alternatives have a meaningful effect on the overall projected impacts; looking at projections from a single GCM, the projected mortality effects of climate change are 13.6 per 100,000 in 2100 under RCP8.5 in the paper's main specification (Panel A) and 12.6 per 100,000 and 13.4 per 100,000 in Panels B and C, respectively.

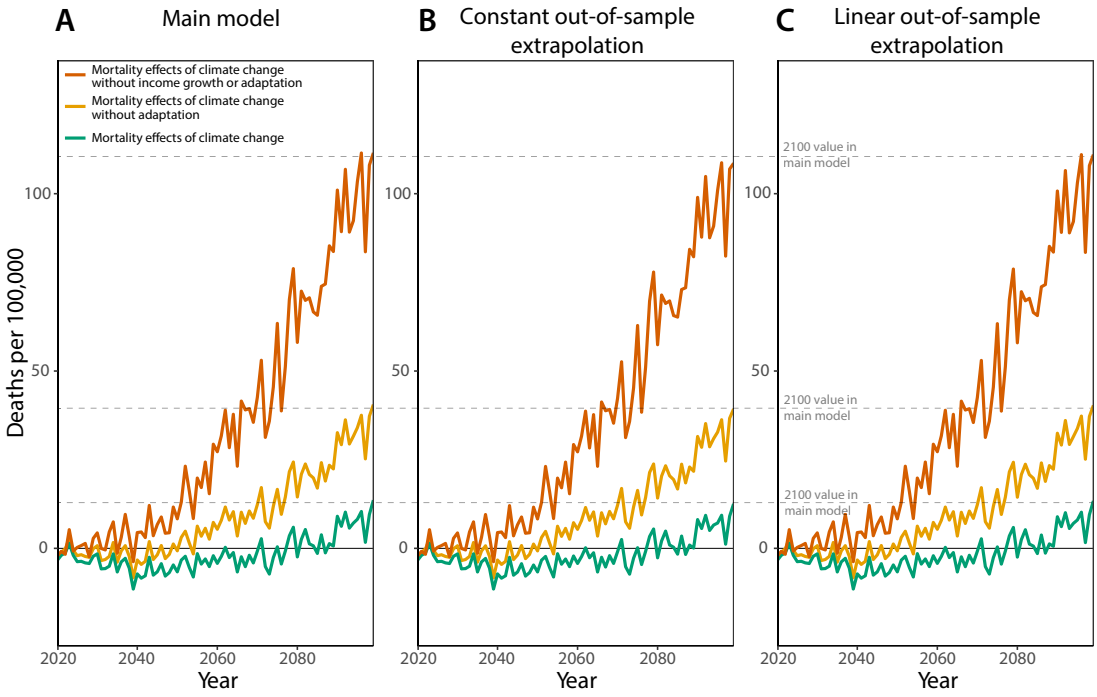

Figure F.10: Two alternative assumptions on out-of-sample extrapolation of response functions. All panels show time series projections of the mortality effects of climate change under RCP8.5 and SSP3 using the CCSM4 climate model. Panel A has no out-of-sample restrictions. Panel B imposes the restriction that the mortality-temperature response function is flat for all temperatures outside the observed range in the estimating sample. Panel C imposes the restriction that the mortality-temperature response function increases linearly for all temperature outside the observed range in the estimating sample. In this case, the linear slope is equal to the slope between the response function evaluated at the maximum (minimum) and the maximum minus 0.1°C (plus 0.1°C).

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

In our main results, we use the estimated coefficients from Equation 4 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 4 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 4 over time when computing future climate change impacts; in doing so, we must make an assumption regarding the rate at which mortality sensitivity to temperature declines with changing covariates. As discussed previously, our main specification relies on a 13-year Bartlett kernel for $\log(GDPpc)$ and a 30-year Bartlett kernel for $TMEAN$.

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

zero (see Figure E.1) is assumed to occur more slowly or more quickly.

In the main model, income grows for each impact region r according to $GDP_{rt} = \rho_{ct} GDP_{r,t-1}$, where c indicates the country that region r falls into, and ρ_{ct} is a country- and year-specific growth rate given exogenously by the SSP scenarios. The kernel-averaged climatic temperature for region r used in the main model is $TMEAN_{rt} = TMEAN_{r,t-1} + \Delta TMEAN_{rt}$. In this "slow adaptation" alternative approach (see Figure F.11B), we replace income growth with $GDP_{rt} = \left(\frac{\rho_{ct}-1}{2} + 1\right) GDP_{r,t-1}$ after the year 2015, and we reduce linear growth in temperature by replacing it with $TMEAN_{rt} = TMEAN_{r,t-1} + \frac{\Delta TMEAN_{rt}}{2}$. In the "fast adaptation" alternative approach (see Figure F.11C), we similarly replace income growth with $GDP_{rt} = 1.5\rho_{ct} GDP_{r,t-1}$ after the year 2015, and we reduce linear growth in temperature by replacing it with $TMEAN_{rt} = TMEAN_{r,t-1} + 1.5\Delta TMEAN_{rt}$. Note that both the primary specification and reduced growth analyses generate identical covariates (and hence, response functions) in 2015.

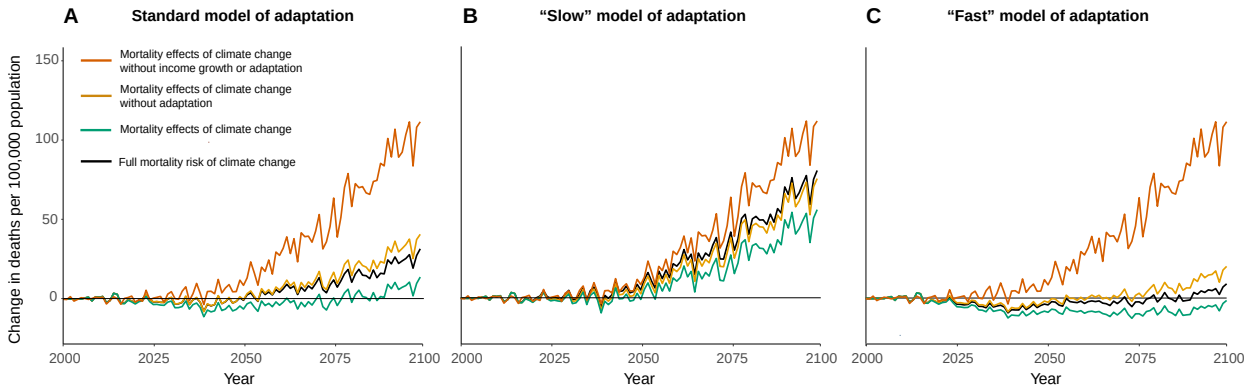

Figure F.11: Impacts of climate change on mortality under alternative assumptions about rates of adaptation. Time series of projected full mortality risk of climate change (black line), and three measures of the mortality effects of climate change (other colors). All lines show predicted mortality impacts of climate change across all age categories under the RCP8.5 emissions scenario, for the socioeconomic scenario SSP3, and using a single climate model (CCSM4). Panel A shows results for our standard model of adaptation, as described in Section 5.2. Panel B shows results for an alternative model of adaptation in which the rate of adaptation to both income growth and to a warming climate is cut in half. Panel C shows results for an alternative model in which the rate of adaptation to both income growth and to a warming climate is increased by 150%.

G Calculation of a mortality partial social cost of carbon

In principle, one could compute a mortality partial social cost of carbon (SCC) estimate by perturbing each global climate model (GCM) in the Surrogate Mixed-Model Ensemble (SMME) with a pulse of CO_2 and projecting mortality for each location in both the original and perturbed simulations. However, in practice, such a procedure is both prohibitively costly from a computational standpoint and would also prevent the calculation of an SCC for any climate trajectory that did not exactly coincide with one of the 33 models. Instead, we rely on a "simple climate model",[52] in combination with our empirically-derived damage functions, to construct mortality partial SCC estimates. We detail this implementation below.

G.1 Computing post-2100 damage functions

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

To estimate post 2100-damages, we develop a method to extrapolate changes in the damage function beyond 2100 using the observed evolution of damages near the end of the 21^{st} century. The year-specific damage functions estimated using Equation 9 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 damages rise over time is the net result of countervailing forces. On the one hand, later years are projected to have larger and older[53] 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 9, but interacting each term linearly with year t (we use 2085-2100 because the evolution of damages over time becomes roughly linear conditional on ΔGMST by this period). The temporal trend over the entire 21^{st} century is convex, implying that our linearization is, if anything,

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

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

conservative. The specific interaction model we estimate is:

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

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

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

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

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

$$\frac{dT_i}{dt} = \frac{q_i F - T_i}{d_i}, i \in \{0, 1\}, \tag{G.25}$$

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

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

$$\frac{dR_j}{dt} = a_j E - \frac{R_j}{\alpha_j \tau_j}, j \in \{0, 1, 2, 3\} \tag{G.26}$$

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

The two scenarios considered in this analysis, RCP4.5 and RCP8.5, represent two widely divergent emissions and climatic pathways, especially in years beyond 2050. Following the method used in previous estimates of the SCC, including in the National Academies of Sciences, Engineering, and Medicine (2017), we include projections starting in the current period (here defined as 2020) through the year 2300. Due to the long residence times of CO_2 in the atmosphere and the changes in global mean surface temperature associated with CO_2 emissions, SCC estimates can vary significantly depending on the definition of this window, especially when low discount rates are applied. To illustrate the large differences across RCP scenarios, Figure G.1 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 8; as described below and in Section 7, this difference is used to construct mortality partial SCC estimates.

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

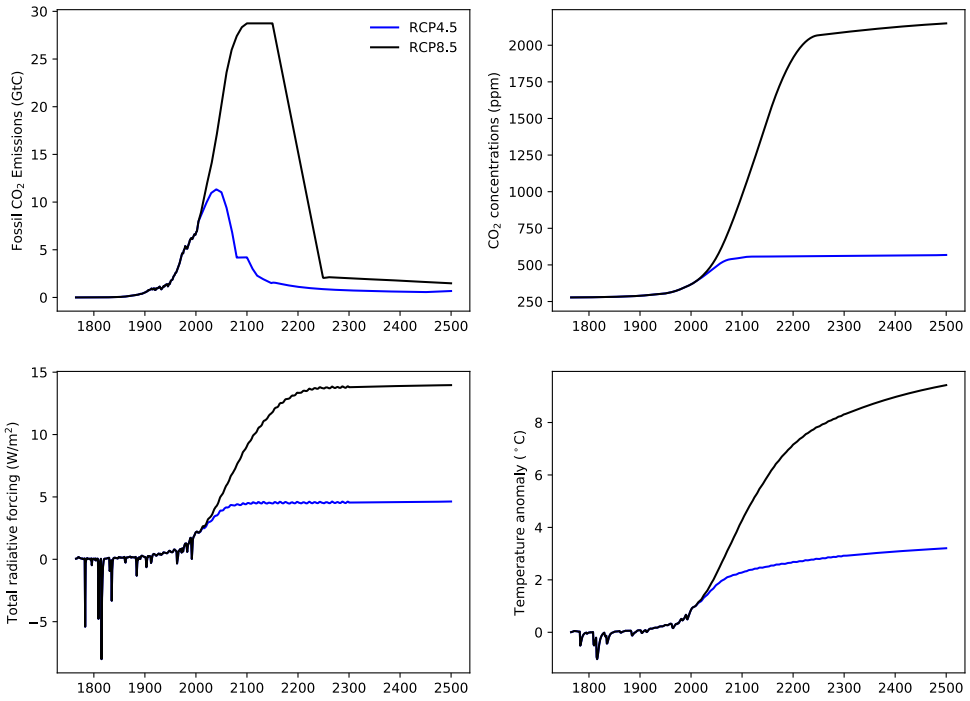

Figure G.1: Behavior of key variables in the FAIR simple climate model under median climate parameters. Each panel shows the temporal trajectory of key variables in FAIR that are used in our calculation of the social cost of carbon. The trajectories shown arise under FAIR run with median climate parameter values calculated from our uncertainty distributions for the equilibrium climate sensitivity, transient climate response, short thermal adjustment time, and time scale of rapid carbon uptake by the ocean mixed layer. The values are shown as anomalies from the year 1765, FAIR's reference state.

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

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

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

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

ranges of variables in the CMIP5 models as representing the "likely" range (central at least 66% probable range) due to structural uncertainty. Previous studies based on CMIP5 results (e.g., Joos et al. (2013); Ricke and Caldeira (2014)) and those using the CMIP5 5–95% range of TCR and ECS as 5-95% input ranges to their models (e.g., Millar et al. (2017)) thus show results that characterize only the central 66% of possibilities. Here we explicitly model the tails of the input and output distributions by generating TCR and ECS distributions with likely ranges as specified by the AR5 report. To preserve the expected correlation between TCR and ECS, rather than sampling ECS directly, we follow Millar et al. (2015) and instead sample the realized warming fraction (RWF, the ratio of TCR/ECS), which is nearly independent of TCR. We subsequently filter the parameter sets to ensure consistency with expectations regarding the initial pulse adjustment timescale (the time it takes the climate system to reach a warming peak following a pulse emission of CO_2).

Below we outline the sources used to construct the distributions of each parameter.

TCR: Collins, Knutti et al. (2013) conclude that "TCR is *likely* in the range 1°C to 2.5°C... is positive and extremely unlikely greater than 3°C" (p. 1112). In IPCC terminology (Mastrandrea et al., 2010), *likely* refers to a probability of at least 66%, *very likely* to a probability of at least 90%, and *extremely likely* to a probability of at least 95%. Thus we construct a log-normal distribution for TCR with the 17th to 83rd range of 1.0-2.5 °C.

RWF: As noted by the National Academies of Sciences, Engineering, and Medicine (2017), a RWF likely range of 0.45 to 0.75 is approximately consistent with the ECS likely range of 1.5 – 4.5°C (Collins, Knutti et al., 2013). We construct a normal distribution for RWF following this central 66% likelihood range, and sample this distribution, along with TCR, to construct the ECS distribution as TCR/RWF.

ECS: Collins, Knutti et al. (2013) conclude that "ECS is positive, *extremely unlikely* less than 1°C (high confidence), and *very unlikely* greater than 6°C (medium confidence)" (p. 1111) and *likely* between 1.5 and 4.5°C. To construct our sampling distribution, we randomly draw samples from the TCR and RWF distributions, and obtain ECS samples by calculating TCR/RWF. The constructed ECS samples follow a log-normal distribution with the 17^{th}-83^{rd} range of 1.60-4.65 °C.

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 G.1 and Figure G.2.

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 G.1).

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 G.1). Based on this post-filtering evaluation, we conclude that the resulting distribution is adequately consistent with our target constraints and the recommendations of the National Academies of Sciences, Engineering, and Medicine (2017). We apply the retained parameter sets to FAIR to produce climate projections that represent climactic uncertainties and are further used in calculating the SCC uncertainty, as described in the next section. The interquartile range of the final SCC values across the entire distribution of parameter sets are shown in Table 3 in the main text.

Finally, we assess the reasonableness of the "handoff" between the SMME models, on which the damage function is estimated, and FAIR, with which future damages due to a pulse of CO_2 are calculated using the difference in temperature between the pulse and control runs. A comparison of climate sensitivity uncertainty across these two climate projections is important, as the climate sensitivity uncertainty captured in the empirically-based projections of mortality damages derives from the SMME, while the uncertainty we proliferate through to the SCC relies on the simple climate model FAIR. Figure G.3 shows the distribution of GMST changes relative to 2001-2010 (ΔGMST) over time, according to the SMME (top row) and the simple climate model FAIR (bottom two rows). To ensure comparability, here and in damage function estimation we use smoothed values of the ΔGMST realizations from each SMME model, where smoothing is done using a 20-year centered moving average. SMME data are available until the year 2100; thus, the top two rows show a direct comparison between FAIR and the SMME models for these years, showing a strong amount of overlap in both RCP4.5 and RPC8.5 distributions of warming and indicating the handoff is reasonable (as

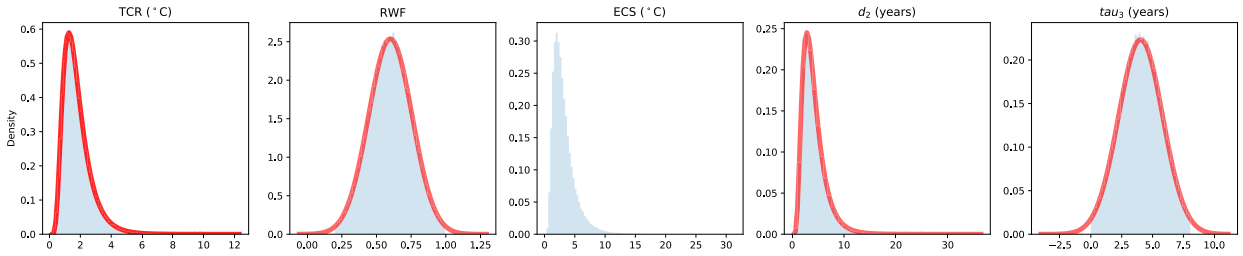

Figure G.2: Distributions of key FAIR parameters for climate sensitivity uncertainty both before (red curve) and after (blue shading) applying constraints. Each panel indicates the distribution of a key parameter in the FAIR simple climate model, both before (in red) and after (in blue) the imposition of constraints described in the text. Distributions shown are: **A** transient climate response (TCR); **B** realized warming fraction (RWF) used to define ECS (=TCR / RWF); **C** equilibrium climate sensitivity (ECS) shown only after applying constraints due to unrealistic values in the initial distribution occurring as RWF$\rightarrow$ 0; **D** short thermal adjustment time (d_2); **E** time scale of rapid carbon uptake by the ocean mixed layer (τ_3).

would be expected based on the construction of the SMME).

Parameter	*Distribution from literature*	*Pre-IPT distribution*	*Post-IPT distribution*	Distribution	*Source*
TCR (C)	[1.00, 2.50]	[1.00, 2.49]	[1.00, 2.50]	Lognormal	AR5
RWF	[0.45, 0.75]	[0.45, 0.75]	N/A	Normal	NAS (2017)
ECS (C)	[1.5, 4.5]	[1.60, 4.65]	[1.61, 4.61]	Lognormal	AR5
d_2 (years)	(1.6, 8.4)	(1.6, 8.4)	(1.6, 8.3)	Lognormal	Millar et al. (2017)
τ_3 (years)	Joos et al. (2013) point estimates	4.04 (1.07, 6.96)	4.04 (1.25, 6.79)	Normal	Joos et al. (2013)
Key metrics					
TCRE (C/TtC)	[0.8, 2.5]	N/A	[0.88, 2.34]	Normal	AR5
Time to T_{max} (years)	(6.6, 30.7)	(4.0, 100.0)*	(4.0, 100.0)*	N/A	Ricke and Caldeira (2014)
RCP 4.5 GMST					
2046 – 2065	1.4 [0.9, 2.0]	N/A	1.38 [0.73, 1.98] (0.51, 2.88)	Normal	AR5
2081 – 2100	1.8 [1.1, 2.6]	N/A	1.81 [0.93, 2.60] (0.65, 3.88)	Normal	AR5
2181 – 2200	2.3 [1.4, 3.1]	N/A	2.37 [1.13, 3.46] (0.78, 5.41)	Normal	AR5
2281 – 2300	2.5 [1.5, 3.5]	N/A	2.73 [1.24, 4.01] (0.85, 6.45)	Normal	AR5
RCP 8.5 GMST					
2046 – 2065	2.0 [1.4, 2.6]	N/A	2.05 [1.09, 2.90] (0.77, 4.20)	Normal	AR5
2081 – 2100	3.7 [2.6, 4.8]	N/A	3.71 [1.96, 5.31] (1.39, 7.73)	Normal	AR5
2181 – 2200	6.5 [3.3, 9.8]	N/A	7.34 [3.82, 10.60] (2.69, 15.35)	Normal	AR5
2281 – 2300	7.8 [3.0, 12.6]	N/A	8.86 [4.48, 12.84] (3.11, 18.84)	Normal	AR5

Table G.1: Comparisons of the distributions of key FAIR parameter values. This table compares the distributions of key FAIR parameter values that pass the initial pulse-adjustment timescale (IPT) constraint against the relevant distributions from the literature (included in the IPT constraint is filtering of τ_3 and RWF as specified in the text). Distributions shown are: transient climate response (TCR); realized warming fraction (RWF); equilibrium climate sensitivity (ECS); short thermal adjustment time (d_2); time scale of rapid carbon uptake by the ocean mixed layer (τ_3); transient climate response to emissions (TCRE); and the change in global mean surface temperature (GMST) from the reference period 1986-2005 at various points in the projections. Note that RWF is only used to create our ECS distribution, and so the post-IPT distribution of RWF is not reported. Distributions reported are determined by the reference values from the literature, so that different parameters have different descriptions of their associated distributions: 5 to 95% ranges are given in ( ), 17 to 83% ranges (*likely* ranges for AR5) are given in [ ], and means are given without ( ) or [ ].

* We only consider the first 100 years post-pulse to be consistent with the length of the simulations in Ricke and Caldeira (2014).

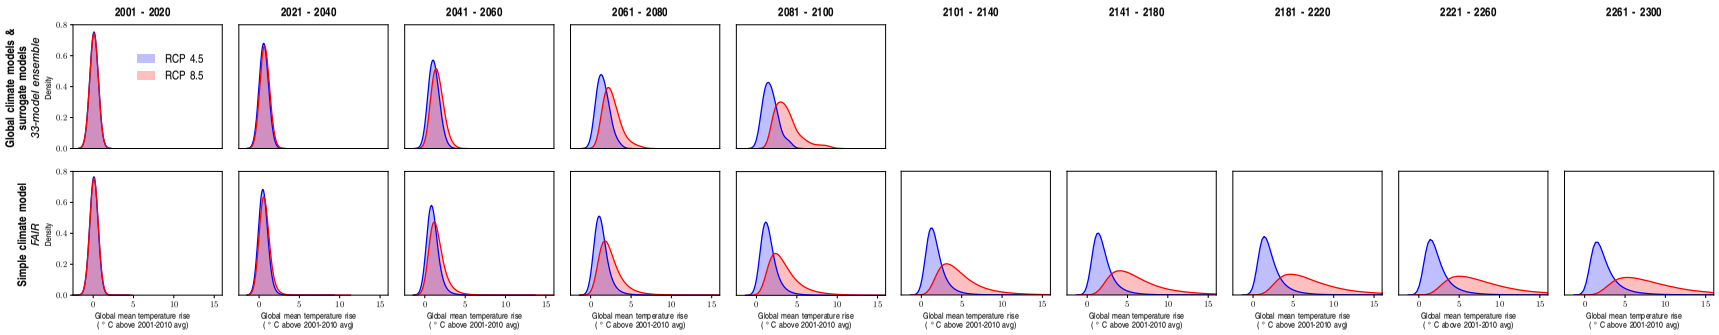

Figure G.3: Distribution of changes in global mean surface temperature (ΔGMST) from an ensemble of global climate models and surrogate models (SMME) and from the simple climate model FAIR. Top row: Distribution of ΔGMST from 2001 to 2100, according to an ensemble of 33 GCMs and surrogate models that form the SMME. Second row: Distribution of ΔGMST from 2001 to 2300, according to 96,306 of simulation runs of the simple climate model FAIR.

G.3 Converting temperature scenarios to mortality partial SCC

We convert the temperature scenarios developed in the climate module into estimates of mortality-related damages using the global damage functions described in Section 7. These damage functions characterize valued mortality damages as a function of ΔGMST (changes in GMST relative to 2001-2010). Figure G.4 shows these functions in 5-year time steps for each combination of valuation assumptions using the US EPA VSL (see Sections 3 and H.1 for discussion of valuation of the full mortality risk of climate change). This figure contains the same information as Figure 7 in the main text, while additionally demonstrating substantial heterogeneity across distinct valuation scenarios (our primary valuation method uses an age-varying VSL in which impact region-specific VSLs are constructed using an income elasticity of one; this valuation is shown in the bottom row and second column of Figure G.4) .

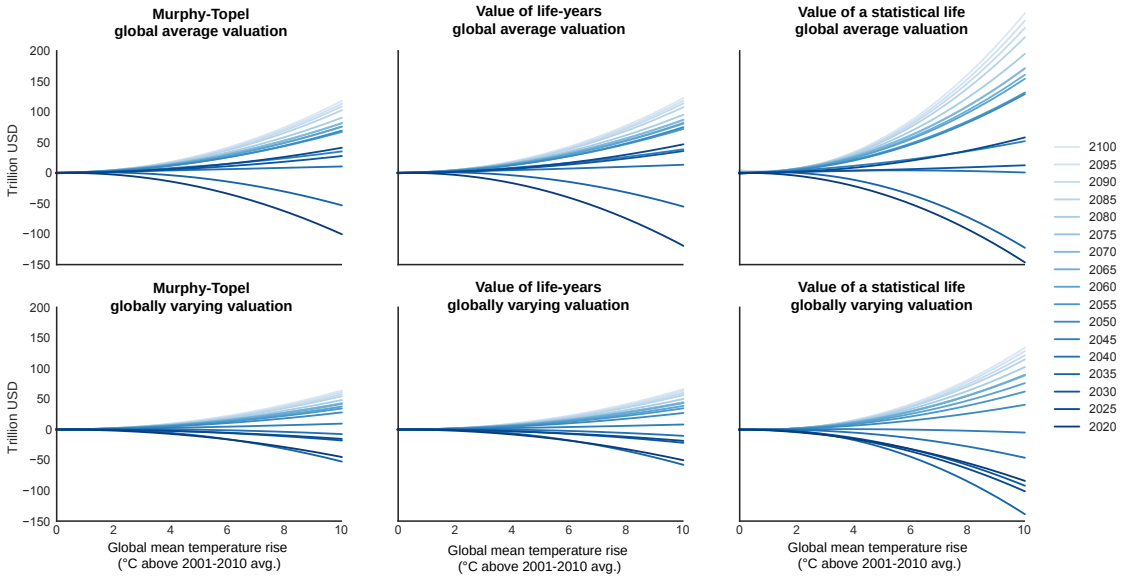

Figure G.4: Temporal evolution of empirically derived damage functions (trillion USD) as a function of global mean surface temperature anomaly. Each panel shows estimates of quadratic damage functions estimated independently for each 5-year period from 2015 to 2100 under various valuation assumptions regarding the valuation of lives lost or saved.

The coefficients on these quadratic damage functions are constructed for each year from 2020 to 2300, as described in the main text. We then generate annual estimates of temperature-related mortality damages by applying the ΔGMST values from both the control FAIR scenarios (RCP4.5 and RCP8.5), as well as pulse scenarios, to the empirically derived damage functions. After computing mortality damages associated with each scenario, we subtract each pulse scenario from the corresponding control scenario and divide by the pulse amount to estimate the marginal effect of the pulse. This time series is then discounted using 2.0%, 2.5%, 3% and 5% discount rates, and summed through time to create a net present value, following Equation 10 in Section 7. This final value is the net present value of the full mortality risks caused of a marginal emission of CO_2. An alternative estimate would make use of Ramsey-like discounting, accounting for the relationship between consumption growth and the discount rate, but we leave this for future study.

Figure G.5 replicates the SCC calculation graphically shown in Figure 8 for RCP 4.5.

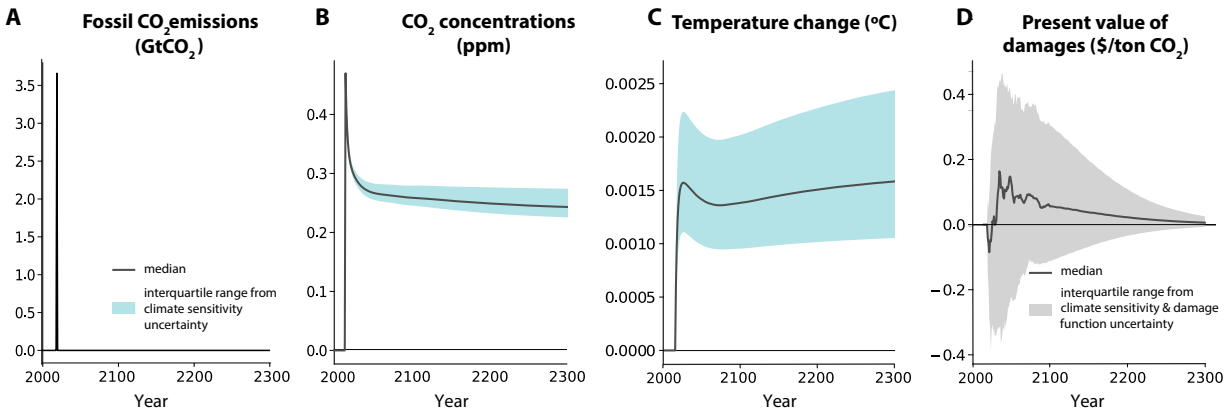

Figure G.5: Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020 under RCP4.5. Panel A shows a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) in 2020 for emissions scenario RCP 4.5. Panel B displays the effect of this pulse on atmospheric CO_2 concentrations, relative to the baseline. In panel C, the impact of the pulse of CO_2 on temperature is shown where the levels are anomalies in global mean surface temperature (GMST) in Celsius. In panels A-C, shaded areas indicate the inter-quartile range due to climate sensitivity uncertainty, while solid lines are median estimates. Panel D shows the change in discounted damages over time due to a 1 Gt pulse of CO_2 in 2020, as estimated by our empirically-derived damage functions, using a 2% annual discount rate and the age-varying U.S. EPA VSL with an income elasticity of one applied to all impact regions. The shaded area indicates the inter-quartile range due to climate sensitivity and damage function uncertainty, while the solid line is the median estimate.

In the main text, we report uncertainty in the mortality partial SCC in three ways: accounting for climate sensitivity uncertainty only, damage function uncertainty only, and full uncertainty (both climate and economic). Here we briefly describe how these values are generated.

Mortality partial SCC estimates accounting for both climate sensitivity and damage function uncertainty: Using our Monte Carlo projections of damages, for each year from 2015 to 2100 we pool all Monte Carlo results for the associated 5-year window. We then run quantile regressions to fit quantile-specific damage functions for 19 quantiles (i.e., every 5^{th} percentile from the 5^{th} to 95^{th}). As in the mean damage function estimation, extrapolation past the year 2100 is accomplished using a time interaction model (see Section 7). In this extrapolation, we allow each quantile of the Monte Carlo distribution to evolve over time heterogeneously, based on the observed changes over time that we estimate at the end of the 21^{st} century.

We run each quantile-specific damage function through each of the 96,306 sets of FAIR parameters and up-weight runs in order to reflect probability mass in the damage function uncertainty space. This process reflects a joint sampling from the full space of damage function uncertainty and climate sensitivity uncertainty. The relevant SCC interquartile range (IQR) is resolved from the resulting distribution of mortality partial SCCs.

Mortality partial SCC estimates accounting for climate sensitivity uncertainty only: To isolate uncertainty in the mortality partial SCC that derives from climate sensitivity uncertainty, we run the central estimate of our damage function through each of the 96,306 sets of FAIR parameters. The corresponding SCC IQR is resolved from the resulting distribution of mortality partial SCCs.

Mortality partial SCC estimates accounting for damage function uncertainty only: To isolate uncertainty in the mortality partial SCC that derives from uncertainty in the damage function, we run the

set of quantile-year damage functions through FAIR with each climate parameter fixed at its median value (as is done in the central mortality partial SCC estimates). The corresponding SCC IQR is resolved from the resulting distribution of mortality partial SCCs.

H Sensitivity of the mortality partial social cost of carbon

The mortality partial social cost of carbon (SCC) estimates shown in the main text depend upon a set of valuation and functional form assumptions and are reported for a particular socioeconomic scenario (SSP3). In this appendix, we detail our valuation approach and provide a wide range of additional mortality partial SCC estimates under alternative valuation approaches, alternative functional forms and extrapolation approaches for the damage function, and under multiple different socioeconomic scenarios. Throughout, we show multiple discount rates and emissions trajectories.

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

As described in Section 7 of the main text, panel A of main text Table 3 utilizes a valuation approach that adjusts the VSL by the total value of expected life-years lost. We provide this metric in order to accommodate the large heterogeneity in mortality-temperature relationships that we uncover across age groups. To adjust VSL values accordingly (see Table H.1 for a set of commonly used VSLs), we first calculate the value of lost life-years by dividing the U.S. EPA VSL by the remaining life expectancy of the median-aged American. This recovers an implied value per life-year. We then apply an income elasticity of one[55] to convert this life-year valuation into a per life-year VSL for each impact region in each year. To calculate life-years lost for a given temperature-induced change in the mortality rate, we use the SSP projected population values, which are provided in 5-year age bins, to compute the implied conditional life expectancy for people in each age bin. We take the population-weighted average of remaining life expectancy across all the 5-year age bins in our broader age categories of <5, 5-64, and >64. This allows us to calculate total expected life-years lost, which we multiply by the impact-region specific VSL per life-year to calculate total damages.

Table H.1: 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

[55] As noted in the main text, the EPA recommends VSL income elasticities of 0.7 and 1.1 (U.S. Environmental Protection Agency, 2016), while a review by Viscusi (2015) estimates an income-elasticity of the VSL of 1.1.

This procedure assumes that our estimated climate change driven deaths occur with uniform probability for all people within an age category. Without historical data containing information on age-specific mortality rates at higher resolution than our three age categories, or information on other chronic health conditions that may lower the life expectancy of individuals in each age group, we cannot empirically parameterize a more detailed life expectancy calculation. However, it is plausible that older individuals within an age category and those with chronic conditions are more likely to die due to extreme temperatures, which would imply that our mortality risk values, when computed using a value of life-years lost approach, are overstated. While we do not have the data sufficient to test this hypothesis, prior evidence from pollution-related mortality in the United States suggests this bias may be substantial (Deryugina et al., 2019).

The above methodology also values each life-year lost identically. In an alternative set of calculations (see results in Appendix H.2), we adjust the life-year values based on the age-specific value of remaining life derived by Murphy and Topel (2006). Murphy and Topel (2006) provide estimates of the value of remaining life for each age group. The authors do not estimate the level of the VSL, but instead provide age-specific values *relative* to a given population-wide VSL. We use these relative values of remaining life by age to adjust the U.S. EPA VSL, such that life-years lost are heterogeneously valued for each impact region in each year, by age. The resulting SCC calculations are shown in Tables H.2 and H.3.

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

In the main text, mortality partial SCC values are shown using a combination of the US EPA VSL, an income elasticity of one, and valuation methods that value deaths using an age-varying value of a statistical life calculation (see Appendix H.1). This appendix shows a range of mortality partial SCC estimates under alternative VSL values, an age-invariant VSL, alternative assumptions about the role of income in valuation, with a life-years adjustment to the VSL that allows for age-specific values of remaining life, as derived by Murphy and Topel (2006), and under alternative socioeconomic scenarios.

Table H.2 provides mortality partial SCC estimates across these distinct valuation approaches under the method shown in the main text Table 3, in which an income elasticity of one is used to adjust VSLs across the globe and over time. Table H.3 provides mortality partial SCC estimates across distinct valuation approaches under a globally uniform valuation method in which a globally homogeneous VSL is used in each year, which evolves over time based on global income growth. Under this alternative, the lives of contemporaries are valued equally, regardless of their relative incomes. The method shown in the main text is most consistent with the revealed preference approach we use to estimate costs of adaptation, given that we observe how individuals make private tradeoffs between their own mortality risk and their own consumption (Equation A.2). 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.

Comparing mortality partial SCC values across panels A and B in Table H.2 shows the influence of age-adjusting the valuation of mortality rates. Throughout the main text, we show that older ages are more susceptible to dying in response to temperature extremes (e.g., see Table D.2). Therefore, age-adjusting the

valuation of mortality rates, which down-weights the valuation of the oldest age group, has an ambiguous influence on the SCC, as this group is more vulnerable both to heat and to cold. Table H.2 shows that under RCP4.5, the age-adjusted SCC is higher than the age-invariant SCC, while under RCP8.5, age-adjustments result in a lower SCC. This difference between RCPs is due to the fact that under RCP4.5, there are many high-income and temperate regions of the globe in which mortality rates in the oldest age group decline on net due to climate change (see Figure F.6). These avoided deaths lower the SCC, counteracting increased net mortality increases in other regions of the globe. When these avoided deaths are valued at a lower rate, the SCC increases. This effect does not dominate under a high emissions scenario (RCP8.5), because increases in heat-related mortality exceed avoided cold-related mortality once temperatures rise to sufficiently high levels.

Table H.2: Globally varying valuation: Estimates of the mortality partial social cost of carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions have heterogeneous valuation, based on local income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix H.1. The first row of every valuation shows our estimated mortality partial SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] showing the influence of climate sensitivity and damage function uncertainty (see Appendix G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally varying valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	28.5	17.1	7.9	2.9	13.3	7.9	3.7	1.3
Climate sensitivity uncertainty	[12.8, 70.1]	[8.3, 39.3]	[4.4, 15.8]	[2.0, 4.3]	[6.0, 32.7]	[3.9, 18.3]	[2.1, 7.4]	[0.9, 2.0]]
Damage function uncertainty	[-38.9, 75.2]	[-23.8, 43.2]	[-13.2, 24.8]	[-6.8, 10.9]	[-18.1, 35.0]	[-11.1, 20.1]	[-6.1, 11.5]	[-3.2, 5.4]
Full uncertainty	[-35.6, 88.5]	[-24.7, 53.6]	[-15.2, 26.3]	[-8.5, 11.5]	[-16.6, 41.2]	[-11.5, 25.0]	[-7.1, 12.2]	[-3.9, 5.3]
RCP 8.5	66.4	36.6	14.2	3.7	30.9	17.0	6.6	1.6
Climate sensitivity uncertainty	[33.3, 140.4]	[18.8, 76.6]	[7.7, 28.3]	[2.4, 6.2]	[15.5, 65.4]	[8.7, 35.7]	[3.6, 13.2]	[1.1, 2.9]
Damage function uncertainty	[-5.7, 116.6]	[-8.0, 67.0]	[-7.0, 35.3]	[-7.1, 13.8]	[-2.7, 54.3]	[-3.7, 31.2]	[-3.3, 16.4]	[-3.1, 6.4]
Full uncertainty	[-2.8, 126.5]	[-7.8, 73.0]	[-11.4, 32.9]	[-8.9, 13.0]	[-1.3, 58.9]	[-3.6, 34.0]	[-5.3, 15.3]	[-4.1, 6.1]
Panel B: Value of statistical life								
RCP 4.5	24.6	14.9	6.7	1.7	11.4	7.0	3.1	0.8
Climate sensitivity uncertainty	[0.9, 99.0]	[2.4, 52.9]	[2.5, 18.3]	[1.0, 2.1]	[0.4, 46.1]	[1.1, 24.6]	[1.2, 8.5]	[0.5, 1.0]
Damage function uncertainty	[-17.6, 69.4]	[-16.1, 44.8]	[-11.6, 24.2]	[-7.4, 12.3]	[-8.2, 32.3]	[-7.5, 20.9]	[-5.4, 11.3]	[-3.4, 5.7]
Full uncertainty	[-25.5, 102.9]	[-21.2, 63.5]	[-15.7, 32.1]	[-11.8, 14.7]	[-11.9, 47.9]	[-9.9, 29.6]	[-7.3, 15.0]	[-5.5, 6.9]
RCP 8.5	123.9	65.1	22.1	3.5	57.7	30.3	10.3	1.6
Climate sensitivity uncertainty	[56.2, 280.0]	[30.0, 147.0]	[10.8, 48.3]	[2.2, 5.6]	[26.2, 130.4]	[14.0, 68.5]	[5.0, 22.5]	[1.0, 2.6]
Damage function uncertainty	[43.9, 182.1]	[20.6, 100.4]	[-2.2, 40.0]	[-7.8, 13.9]	[20.4, 84.8]	[9.6, 46.6]	[-1.0, 18.6]	[-3.6, 6.5]
Full uncertainty	[13.7, 253.6]	[3.0, 139.0]	[-5.6, 53.4]	[-9.3, 16.0]	[6.4, 118.1]	[1.4, 64.7]	[-2.6, 24.9]	[-4.3, 7.5]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	28.9	17.5	8.3	3.1	13.4	8.1	3.9	1.4
Climate sensitivity uncertainty	[13.3, 69.9]	[8.8, 39.6]	[4.7, 16.3]	[2.1, 4.8]	[6.2, 32.5]	[4.1, 18.4]	[2.2, 7.6]	[1.0, 2.2]
Damage function uncertainty	[-27.3, 90.0]	[-24.6, 52.1]	[-11.5, 26.8]	[-5.8, 12.2]	[-12.7, 41.9]	[-11.4, 24.3]	[-5.4, 12.5]	[-2.7, 5.7]
Full uncertainty	[-36.5, 93.9]	[-25.3, 56.6]	[-15.6, 27.7]	[-8.6, 12.2]	[-17.0, 43.8]	[-11.8, 26.4]	[-7.3, 12.9]	[-4.0, 5.7]
RCP 8.5	65.5	36.3	14.3	4.0	30.5	16.9	6.7	1.9
Climate sensitivity uncertainty	[33.0, 237.7]	[18.8, 75.5]	[7.9, 28.3]	[2.6, 6.6]	[15.4, 64.1]	[8.7, 35.2]	[3.7, 13.2]	[1.2, 3.1]
Damage function uncertainty	[-4.2, 116.3]	[-10.1, 70.9]	[-7.5, 34.4]	[-6.4, 13.7]	[-2.0, 54.2]	[-4.7, 33.0]	[-3.5, 16.0]	[-3.0, 6.4]
Full uncertainty	[-2.4, 120.5]	[-8.0, 70.9]	[-11.6, 33.0]	[-8.8, 13.6]	[-1.1, 56.1]	[-3.7, 33.0]	[-5.4, 15.4]	[-4.1, 6.3]

Table H.3: Globally uniform valuation: Estimates of the mortality partial social cost of carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions are given the global median VSL, after scaling using income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix H.1. The first row of every valuation shows our estimated mortality partial SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] showing the influence of climate sensitivity and damage function uncertainty (see Appendix G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally uniform valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	58.1	37.5	19.9	9.0	27.0	17.5	9.3	4.2
Climate sensitivity uncertainty	[27.8, 136.3]	[19.4, 82.2]	[11.4, 38.7]	[5.8, 15.1]	[13.0, 63.5]	[9.0, 38.3]	[5.3, 18.0]	[2.7, 7.1]
Damage function uncertainty	[-20.8, 143.7]	[-15.5, 83.3]	[-7.3, 48.5]	[-3.9, 22.0]	[-9.7, 66.9]	[-7.2, 38.8]	[-3.4, 22.6]	[-1.8, 10.2]
Full uncertainty	[-19.1, 166.3]	[-13.3, 101.7]	[-8.4, 50.0]	[-5.0, 21.7]	[-8.9, 77.5]	[-6.2, 47.4]	[-3.9, 23.3]	[-2.3, 10.1]
RCP 8.5	123.8	72.3	32.0	11.5	57.7	33.6	14.9	5.3
Climate sensitivity uncertainty	[15.1, 228.1]	[15.3, 134.1]	[12.3, 60.2]	[7.2, 21.4]	[29.4, 120.5]	[17.6, 69.4]	[8.2, 29.2]	[3.2, 9.4]
Damage function uncertainty	[40.9, 210.2]	[7.6, 127.2]	[1.3, 59.2]	[-2.8, 25.2]	[19.1, 97.9]	[3.6, 59.2]	[0.6, 27.6]	[-1.3, 11.8]
Full uncertainty	[14.7, 236.2]	[4.6, 141.1]	[-2.9, 64.6]	[-4.6, 24.9]	[6.9, 110.0]	[2.2, 65.7]	[-1.4, 30.1]	[-2.1, 11.6]
Panel B: Value of statistical life								
RCP 4.5	68.2	46.2	25.9	11.9	31.8	21.5	12.1	5.5
Climate sensitivity uncertainty	[15.1, 228.1]	[15.3, 134.1]	[12.3, 60.2]	[7.2, 21.4]	[7.0, 106.2]	[7.1, 62.4]	[5.7, 28.0]	[3.4, 10]
Damage function uncertainty	[36.4, 167.3]	[17.6, 105.6]	[7.5, 56.1]	[-1.7, 28.0]	[17.0, 77.9]	[8.2, 49.2]	[3.5, 26.1]	[-0.8, 13.1]
Full uncertainty	[11.1, 246.6]	[2.8, 148.2]	[-4.1, 71.0]	[-4.2, 30.2]	[5.2, 114.9]	[1.3, 69.0]	[-1.9, 33.1]	[-2.0, 14.1]
RCP 8.5	260.1	143.9	57.5	17.6	121.1	67.0	26.8	8.2
Climate sensitivity uncertainty	[121.3, 578.5]	[68.8, 317.6]	[29.2, 121.7]	[10.1, 32.9]	[56.5, 269.4]	[32.0, 147.9]	[13.6, 56.7]	[4.7, 15.3]
Damage function uncertainty	[132.8, 349.7]	[73.7, 203.1]	[26.6, 92.2]	[2.2, 32.2]	[61.8, 162.9]	[34.3, 94.6]	[12.4, 42.9]	[1.0, 15.0]
Full uncertainty	[72.1, 508.2]	[39.0, 287.0]	[11.9, 117.4]	[-2.0, 37.9]	[33.6, 236.7]	[18.2, 133.7]	[5.5, 54.7]	[-1.0, 17.6]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	55.5	35.8	19.0	8.6	25.9	16.7	8.8	4.0
Climate sensitivity uncertainty	[26.3, 131.7]	[18.4, 79.1]	[10.9, 37.0]	[5.5, 14.4]	[12.2, 61.3]	[8.6, 36.8]	[5.1, 17.2]	[2.6, 6.7]
Damage function uncertainty	[-23.8, 133.2]	[-12.9, 90.7]	[-4.0, 46.6]	[-4.4, 19.7]	[-11.1, 62.0]	[-6.0, 42.2]	[-1.9, 21.7]	[-2.0, 9.2]
Full uncertainty	[-19.1, 164.3]	[-14.2, 99.9]	[-9.0, 48.8]	[-5.7, 21.3]	[-8.9, 76.5]	[-6.6, 46.5]	[-4.2, 22.7]	[-2.7, 9.9]
RCP 8.5	121.0	70.1	30.7	10.9	56.3	32.6	14.3	5.1
Climate sensitivity uncertainty	[61.5, 253.2]	[36.6, 144.7]	[17.0, 60.2]	[6.6, 19.1]	[28.6, 117.9]	[17.1, 67.4]	[7.9, 28.0]	[3.1, 8.9]
Damage function uncertainty	[20.8, 189.2]	[10.6, 122.5]	[0.5, 58.0]	[-3.1, 22.6]	[9.7, 88.1]	[4.9, 57.0]	[0.2, 27.0]	[-1.5, 10.5]
Full uncertainty	[13.1, 224.6]	[3.7, 134.5]	[-2.7, 61.7]	[-5.2, 24.0]	[6.1, 104.6]	[1.7, 62.7]	[-1.3, 28.7]	[-2.4, 11.2]

Table H.4 shows mortality partial SCC estimates under various socioeconomic projections (SSP3 is used throughout the main text; see Appendix B.3.2 for a discussion of this choice). We note that under SSP4 and a moderate emissions scenario (RCP4.5), the central estimate of the partial SCC is negative under all discount rates shown. While SSP4 shows global average increases in the full mortality risk of climate change by 2100 under both emissions scenarios (see Figure F.5), the negative SCC is driven by different income and demographic changes projected under SSP4 relative to the other SSPs, both of which influence the valuation of lives lost. In particular, SSP4 projects that today's wealthy and relatively cold locations will experience dramatically higher future incomes, with much older populations, when compared to SSP2 or SSP3. This increase in income and rapid aging of the population leads to many lives saves in cold regions of the world as the climate warms, and each life is valued highly due to income growth raising the VSL (recall that we use an income elasticity of one for the VSL throughout the text). In contrast, SSP4 projects very low income growth in today's hot and poor locations, such that lives lost due to warming in these regions receive little value in this scenario. Note that with sufficiently high emissions (RCP8.5), heat-related deaths outweigh cold-related lives saved even in today's wealthy and relatively cold regions of the world, such that the partial SCC for SSP4 is no longer negative.

Table H.4: Estimates of the mortality partial social cost of carbon (SCC) under various socioeconomic projections. In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD. In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Each row shows, for a different SSP scenario, our estimated SCC using the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function.

	Annual discount rate			
	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 USD)				
RCP 4.5				
SSP2	45.6	25.7	10.4	2.9
SSP3	28.5	17.1	7.9	2.9
SSP4	-23.5	-14.5	-7.5	-3.7
RCP 8.5				
SSP2	63.6	33.3	11.0	1.2
SSP3	66.4	36.6	14.2	3.7
SSP4	41.6	22.5	7.9	1.2
Panel B: Globally varying value of a statistical life (2019 USD)				
RCP 4.5				
SSP2	4.1	2.0	-0.9	-3.3
SSP3	24.6	14.9	6.7	1.7
SSP4	-97.7	-64.3	-36.1	-18.5
RCP 8.5				
SSP2	90.5	43.9	10.7	-2.5
SSP3	123.9	65.1	22.1	3.5
SSP4	53.3	23.1	1.2	-6.4

Finally, Table H.5 shows mortality partial SCC estimates under both SSP2 (repeating values in Table H.4) and a "hybrid" SSP designed to approximate a scenario in which climate change impacts on economic growth are endogenized. Throughout our main analysis, we treat income as exogenously given by the Shared Socioeconomic Pathways (SSPs). However, a growing literature indicates that the level and/or growth rate of income is influenced by temperature (e.g., Burke, Hsiang, and Miguel, 2015; Kalkuhl and Wenz, 2020). Following this literature and allowing income to respond to emissions could influence our mortality partial SCC estimates both by changing location-specific VSLs, and by changing income-driven adaptation in location-specific mortality-temperature relationships. While a full treatment of this topic is beyond the scope of this analysis, here we create a hybrid SSP that is constructed to approximate the impact of endogenous economic growth on the mortality partial SCC. Our analysis involves two steps.

First, we choose two scenarios from the three SSPs included in the main analysis (SSP2, SSP3, and SSP4) for which differences in income across SSPs for each quintile of the global income distribution approximately match the impacts of climate change from Burke, Hsiang, and Miguel (2015). To see this visually, Panel A of Figure H.1 shows the estimated impacts of climate change on GDP per capita from Burke, Hsiang, and Miguel (2015), where impacts are shown for each quintile of the 2010 country-level income distribution. The level of these curves indicate the difference between incomes under climate change following RCP8.5 versus without climate change. Panel B of Figure H.1 shows the difference between incomes under SSP2 versus under a hybrid SSP in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010. As can be seen by comparing across panels, the difference between SSP2 and our hybrid scenario closely approximates the estimated GDP per capita climate change impacts in Burke, Hsiang, and Miguel (2015).

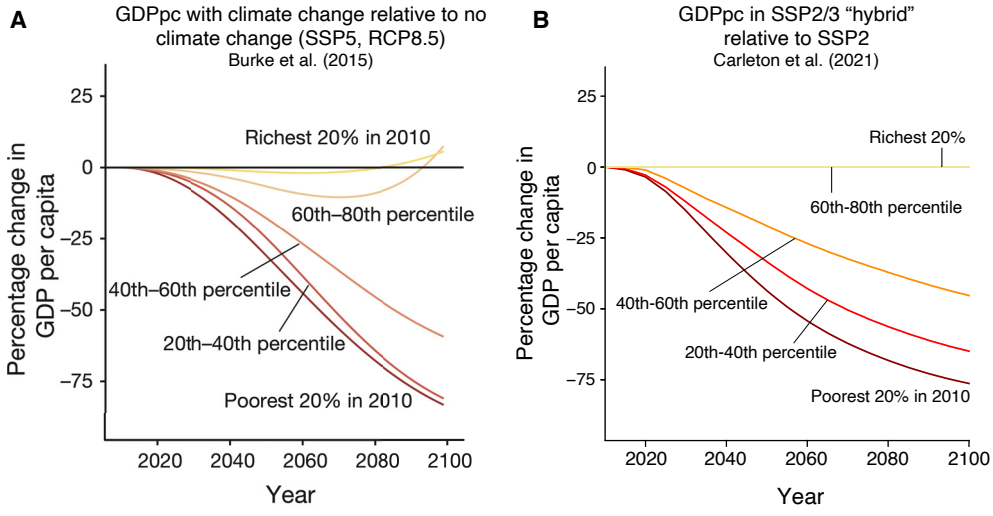

Figure H.1: Constructing a hybrid Shared Socioeconomic Scenario (SSP) to approximate an income trajectory that is endogenous to climate change, following Burke, Hsiang, and Miguel (2015). Panel A is reproduced from Burke, Hsiang, and Miguel (2015) and shows mean impacts of climate change by 2010 income quantile for the authors' benchmark empirical model. The right panel shows the mean difference in income between SSP2 and a hybrid socioeconomic scenario in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010.

Second, we compute the mortality partial SCC under SSP2 as well as under our hybrid scenario, and compare SCC estimates. The difference in SCCs across these two scenarios approximates the effect of endogenizing income growth to climate change (as estimated by Burke, Hsiang, and Miguel (2015)) on the mortality partial SCC. Table H.5 shows this comparison. Despite the extraordinary income differences shown in Figure H.1, under our central valuation approach (δ=2%) and RCP8.5 emissions, the SCC rises by just 6% when using the hybrid socioeconomic scenario, relative to SSP2. Note that we do not report SCCs in Table H.5 under RCP4.5, as the hybrid scenario was calibrated to match estimates from Burke, Hsiang, and Miguel (2015) for RCP8.5.

Table H.5: Estimates of the mortality partial social cost of carbon under a hybrid socioeconomic scenario designed to approximate an endogenous growth trajectory. Partial mortality SCC estimates are shown for a reference socioeconomic scenario (SSP2, "Reference"), as well as a hybrid scenario ("Hybrid") in which SSP2 projected income is replaced by SSP3 projected income only for the poorest 60% of countries in 2010. An income elasticity of one is used to scale the U.S. EPA VSL value and all estimates correspond to RCP8.5 emissions. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year. See text for details on the hybrid socioeconomic scenario.

	Annual discount rate			
	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$
SSP2 (Reference)	63.6	33.3	11.0	1.2
Hybrid SSP	65.6	35.3	12.7	2.5

H.3 Alternative approach to estimating post-2100 damages

As discussed in Section 7, we rely on an extrapolation of estimated damage functions to capture mortality impacts of climate change after the year 2100, due to data limitations. In this appendix, we explore the importance of this extrapolation by using an alternative approach to estimating post-2100 damage functions. Here, we calculate mortality partial SCC estimates using a set of damage functions in which the estimated 2100 damage function is applied to all years from 2100-2300. Effectively, this freezes the damage function at its 2100 level for all later years. Values shown are for SSP3, RCP8.5, with a discount rate of 2% and an age-varying VSL. Table H.6 shows that this alternative approach to post-2100 damage estimation causes our central estimate of the SCC to fall by 21%.

Table H.6: The influence of damage function extrapolation in years after 2100 on estimates of the mortality partial social cost of carbon (SCC). In this table, an income elasticity of one is used to scale the U.S. EPA VSL value, and all SCC values are for the year 2020 under RCP8.5 emissions, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. The VSL is age-varying, so that these values are directly comparable to panel A in Table 3 in the main text. For the first column, damage functions continue to evolve over time in the years after 2100, according to the method described in Section 7. In the second column, the damage function estimated for the year 2100 is used for all years after 2100. All mortality partial SCC estimates use the median values for the four key input parameters of the simple climate model FAIR and a conditional mean estimate of the damage function.

	Extrapolating post-2100 damage function	*Holding post-2100 damage function fixed*
Pre-2100 damages	\$12.8	\$12.8
Post-2100 damages	\$23.8	\$16.0
Total damages	\$36.6	\$28.8

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

Throughout the main text, we report mortality partial SCC estimates that rely on a quadratic damage function estimated through all damage projections from all Monte Carlo simulation runs (see Section 7 for details). In Table H.7, 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 H.7: Estimates of the mortality partial social cost of carbon (SCC) using a cubic polynomial damage function In this table, an income elasticity of one is used to scale the U.S. EPA VSL value. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. Damage functions are estimated as a cubic polynomial, instead of a quadratic (as in the main text). In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Estimates rely on the median values of the four key input parameters into the simple climate model FAIR and a conditional mean estimate of the damage function.

	Annual discount rate			
	$\delta = 1.5\%$	$\delta = 2\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying value of a statistical life (2019 USD)				
RCP 4.5	15.0	9.4	4.9	2.4
RCP 8.5	85.5	44.5	16.1	4.0
Panel B: Globally varying value of a statistical life (2019 USD)				
RCP 4.5	31.1	18.7	9.1	3.8
RCP 8.5	135.3	68.4	21.9	2.8

Appendix references

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

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

Auffhammer, Maximilian, Solomon M Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using weather data and climate model output in economic analyses of climate change." *Review of Environmental Economics and Policy* 7 (2):181–198.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro. 2016. "Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century." *Journal of Political Economy* 124 (1):105–159. URL http://dx.doi.org/10.1086/684582.

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

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

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

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

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

Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li. 2013. "Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy." *Proceedings of the National Academy of Sciences* 110 (32):12936–12941.

Collins, Matthew, Reto Knutti et al. 2013. *Long-term Climate Change: Projections, Commitments and Irreversibility*, chap. 12. Intergovernmental Panel on Climate Change, 1029–1136. URL http://www.ipcc.ch/report/ar5/wg1/.

Curriero, Frank C, Karlyn S Heiner, Jonathan M Samet, Scott L Zeger, Lisa Strug, and Jonathan A Patz. 2002. "Temperature and mortality in 11 cities of the eastern United States." *American Journal of Epidemiology* 155 (1):80–87.

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

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

Deryugina, Tatyana, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif. 2019. "The mortality and medical costs of air pollution: Evidence from changes in wind direction." *American Economic Review* 109 (12):4178–4219.

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

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, 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.

Eurostat. 2013. "EuroStat Regional Yearbook." Available at http://ec.europa.eu/eurostat/data/database.

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

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

Geoffroy, Olivier, David Saint-Martin, Dirk JL Olivié, Aurore Voldoire, Gilles Bellon, and Sophie Tytéca. 2013. "Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments." *Journal of Climate* 26 (6):1841–1857.

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

Global Administrative Areas. 2012. "GADM database of Global Administrative Areas, version 2.0." *Accessed 25 December, 2016* URL www.gadm.org.

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

Guo, Yuming, Antonio Gasparrini, Ben Armstrong, Shanshan Li, Benjawan Tawatsupa, Aurelio Tobias, Eric Lavigne, Micheline de Sousa Zanotti Stagliorio Coelho, Michela Leone, Xiaochuan Pan et al. 2014. "Global

variation in the effects of ambient temperature on mortality: a systematic evaluation." *Epidemiology (Cambridge, Mass.)* 25 (6):781.

Hancock, Peter A, Jennifer M Ross, and James L Szalma. 2007. "A meta-analysis of performance response under thermal stressors." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 49 (5):851–877.

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

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

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

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

Joos, Fortunat, Raphael Roth, JS Fuglestvedt, GP Peters, IG Enting, W von Bloh, V Brovkin, EJ Burke, M Eby, NR Edwards et al. 2013. "Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: A multi-model analysis." *Atmospheric Chemistry and Physics* 13 (5):2793–2825.

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

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

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

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

Mastrandrea, Michael D, Christopher B Field, Thomas F Stocker, Ottmar Edenhofer, Kristie L Ebi, David J Frame, Hermann Held, Elmar Kriegler, Katharine J Mach, Patrick R Matschoss et al. 2010. "Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties." Tech. rep., Intergovernmental Panel on Climate Change.

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

Meinshausen, M., S. C. B. Raper, and T. M. L. Wigley. 2011. "Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – Part 1: Model description and calibration." *Atmospheric Chemistry and Physics* 11 (4):1417–1456. URL http://www.atmos-chem-phys.net/11/1417/2011/.

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

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

Millar, Richard J., Alexander Otto, Piers M. Forster, Jason A. Lowe, William J. Ingram, and Myles R. Allen. 2015. "Model structure in observational constraints on transient climate response." *Climatic Change* 131 (2):199–211.

Ministry of Health, Chile. 2015. "Department of Statistics and Information." Available at http://www.deis.cl/bases-de-datos-defunciones/.

Ministry of Health in Brazil. 2019. "Mortality Information System (SIM)." Available at http://datasus.saude.gov.br/sistemas-e-aplicativos/eventos-v/sim-sistema-de-informacoes-de-mortalidade.

Mitchell, Timothy D. 2003. "Pattern Scaling: An Examination of the Accuracy of the Technique for Describing Future Climates." *Climatic Change* 60 (3):217–242. URL http://link.springer.com/article/10.1023/A%3A1026035305597.

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

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

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

National Center for Atmospheric Research Staff (Eds). 2015. "Global surface temperatures: BEST: Berkeley Earth Surface Temperatures." *The Climate Data Guide* .

National Institute for the Study of Demography (INED). 2019. "Vital Registration System." Available at https://www.ined.fr/en/.

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

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

O'Neill, Brian C, Elmar Kriegler, Keywan Riahi, Kristie L Ebi, Stephane Hallegatte, Timothy R Carter, Ritu Mathur, and Detlef P van Vuuren. 2014. "A new scenario framework for climate change research: The concept of shared socioeconomic pathways." *Climatic Change* 122 (3):387–400.

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

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

Riahi, Keywan, Shilpa Rao, Volker Krey, Cheolhung Cho, Vadim Chirkov, Guenther Fischer, Georg Kindermann, Nebojsa Nakicenovic, and Peter Rafaj. 2011. "RCP 8.5—A scenario of comparatively high greenhouse gas emissions." *Climatic Change* 109 (1-2):33–57.

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

Ricke, Katharine L and Ken Caldeira. 2014. "Maximum warming occurs about one decade after a carbon dioxide emission." *Environmental Research Letters* 9 (12):124002.

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

Samir, KC and Wolfgang Lutz. 2014. "The human core of the shared socioeconomic pathways: Population scenarios by age, sex and level of education for all countries to 2100." *Global Environmental Change* .

Seppanen, Olli, William J Fisk, and QH Lei. 2006. "Effect of temperature on task performance in office environment." *Lawrence Berkeley National Laboratory* 1 (LBNL-60946).

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

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

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.

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

U.S. Environmental Protection Agency. 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* .

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

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

Wood, Andrew W, Lai R Leung, V Sridhar, and DP Lettenmaier. 2004. "Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs." *Climatic Change* 62 (1-3):189–216.

World Bank. 2020. "World Development Indicators." .

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